



A University of Sussex PhD thesis

This thesis is protected by copyright which belongs to the author

It cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Please visit Figshare for more information and further details

<https://sussex.figshare.com/Theses>

FARE-KE

Framework for Affordable, Reliable Kinesthetic Evaluation

By

Jose Luis Berna Moya



School of Engineering and Informatics

UNIVERSITY OF SUSSEX

A dissertation submitted to the University of Sussex
in accordance with the requirements of the degree of
DOCTOR OF PHILOSOPHY in the School of Engineering
and Informatics.

Jun 2024

Declaration

I hereby declare that this thesis has not been and will not be, submitted in whole or in part to another University for the award of any other degree.

Signature: Jose Luis Berna Moya

Acknowledgements

Reflecting on my doctoral journey, I am profoundly grateful for the opportunities it has presented and for the remarkable individuals I have met along the way. This venture has reshaped not just my academic career but has also profoundly impacted my personal growth.

I extend my deepest gratitude to Professor Sriram Subramanian and Professor Marianna Obrist, who provided the initial opportunity to join their lab as a Research Technician. This position not only introduced me to the fascinating world of human-computer interaction but also ignited my passion for research, setting the stage for my PhD journey. The vibrant atmosphere and unwavering support from my colleagues at the Interact Lab and the SCHI Lab were central to my motivation. Their passion and dedication during my time as a technician helped shape my path toward pursuing a PhD.

I am especially thankful to my supervisor, Diego Martínez, whose passion, dedication, and enthusiasm for his career profoundly influenced my decision to embark on this PhD. His approach, filled with curiosity and a sense of play, has been a continuous source of inspiration and guidance even at our 5am meetings.

I would also like to thank Yanan Li, Elizabeth Rendon-Morales, Gianluca Memoli, and the many others who dedicated their time and provided support after the Interact Lab transitioned to new ventures.

Most importantly, I wish to thank my family, who embarked on this journey with me. The news of my son Sean's impending arrival coincided with the receipt of my PhD acceptance letter. Three years later, we welcomed my daughter Aurora. Throughout this academic and personal journey, my wife Catherine, and my children, Sean and Aurora, have been my pillars of support. Their love, silliness and encouragement, especially during challenging times such as the COVID pandemic, have been my stronghold. Their laughter, love, and unwavering support have made all the difference, making this challenging journey not only possible but also a joyous adventure.

Abstract

Over the past decade, there has been a marked increase in accessible and user-friendly design tools and hardware platforms for haptic device development. This democratisation has enabled individuals from diverse expertise and academic backgrounds to participate in haptic design and prototyping. However, despite significant progress in design capabilities, the field of haptic device characterisation has not experienced comparable advancement in accessible methodologies and tools.

This thesis chronicles the development of the Framework for Affordable, Reliable Kinesthetic Evaluation (FARE-KE), which emerged organically from practical challenges encountered while creating and characterising diverse haptic controllers. The research journey began with developing novel haptic prototypes, each presenting unique measurement challenges that revealed limitations in existing characterisation approaches. These challenges were particularly pronounced during COVID-19 restrictions, which necessitated innovative solutions with limited resources and equipment access.

Through iterative development across multiple projects, methodologies for measurement, calibration, and analysis evolved into a coherent framework. FARE-KE provides a standardised approach to haptic device characterisation using affordable, readily available components. Key technical contributions include: (1) validated low-cost load cell configurations capable of measuring forces from 0.5N to 80N with accuracy better than 2%; (2) a custom data acquisition system supporting sampling rates up to 10kSPS across multiple channels; (3) calibration methods that reduce crosstalk in multi-axis measurements by a factor of 10; and (4) standardised analysis procedures for consistent performance evaluation.

To demonstrate the framework's development and application, this thesis presents three case studies of haptic controllers that contributed to its evolution: a pneumatically-actuated single-axis controller (VR Recoil), which revealed that lower levels of linear kinesthetic feedback can enhance immersion without compromising performance; a study of self-touch haptic sensations (ProxyTouch), which confirmed somatosensory attenuation effects in pinch interactions; and a flywheel-based multi-axis haptic device (HapticWhirl), which demonstrated the ability to generate five distinct haptic modalities from a single actuator system. Each project not only exemplifies specific kinesthetic feedback modalities but also documents the characterisation challenges that shaped FARE-KE's development.

The FARE-KE framework represents a significant contribution to the haptic research community by lowering barriers to rigorous device characterisation. By providing accessible methods that complement existing approaches without requiring specialised expertise or expensive equipment, the framework supports the ongoing democratisation of haptic technology development. The standardised methodology enables more consistent reporting of device performance metrics, facilitating meaningful comparisons between different haptic systems and accelerating innovation in the field.

Parts of the work presented in this thesis have been previously published in the following papers:

Berna Moya JL, van Oosterhout A, Marshall MT, Martinez Plasencia D. "HapticWhirl, a Flywheel-Gimbal Handheld Haptic Controller for Exploring Multimodal Haptic Feedback". *Sensors (Basel)*. 2024 Jan 31;24(3):935. Doi: 10.3390/s24030935. PMID: 38339652; PMCID: PMC10857638.

Berna-Moya, Jose Luis, and Diego Martinez-Plasencia. "Exploring the effects of replicating shape, weight and recoil effects on VR shooting controllers." *Human-Computer Interaction-INTERACT 2019: 17th IFIP TC 13 International Conference, Paphos, Cyprus, September 2-6, 2019, Proceedings, Part I 17*. Springer International Publishing, 2019.

Contents

Acknowledgements.....	1
Abstract	2
Contents	a
List of Figures.....	a
List of Tables.....	18
Acronyms	22
1 Introduction.....	24
1.1 Ethos & Main Contribution.....	28
1.2 Research Questions	29
2 Situating our Work	32
2.1 Understanding Haptic Perception: Tactile and Kinesthetic Feedback.....	33
2.1.1 Tactile Feedback	34
2.1.2 Kinesthetic Feedback.....	36
2.2 Haptic technologies	38
2.2.1 Tactile Technologies	38
2.2.2 Kinesthetic Feedback Technologies	42
2.3 Haptic Controllers Characterisation.....	51
2.3.1 Tactile Feedback Characterisation	52
2.3.2 Kinesthetic Characterisation	61
2.3.3 Haptic controller metrics.....	84
2.3.4 Summary/Learning from Previous Literature	93
3 FARE-KE Framework for Affordable, Reliable Kinesthetic Evaluation.....	98
3.1 FARE-KE Framework Overview	101
3.2 Initial Considerations	104
3.3 Testbed Designs	105
3.4 Load Cell Count and Type.....	106
3.5 Considering Range and Capacity.....	107
3.6 DAQ Selection	108
3.7 Data Logging	112
3.8 Calibration Process	113
3.9 Device Characterisation Process	117
3.10 Conclusion	122
4 FARE-KE Data Acquisition Board	124
4.1 Proposed design	124
4.2 Reading Data from the DAQ - Microcontroller.....	130
4.2.1 Microcontroller Selection Criteria.....	131
4.2.2 Library Implementation.....	132
4.3 Performance Metrics	132

5	FARE-KE Sensor/Hardware Validation	140
5.1	Direct Single-Axis Force.....	141
5.1.1	Sensor Validation.....	142
5.1.2	Specific Setups for Single Axis Force.....	155
5.2	Single-Axis Impact Force	156
5.2.1	Sensor validation	156
5.2.2	Specific Setups for Single Axis Impact.....	157
5.3	In-line Forces	158
5.3.1	Specific Setups for In-Line force.....	159
5.4	Rotational Joint Torque.....	161
5.4.1	Specific Setups for Joint Torque.....	161
5.5	Force Torque	165
5.5.1	Sensor Design & Operation	166
5.5.2	3-Axis Load Cells.....	166
5.5.3	6-axis Force-Torque Plate.....	170
5.5.4	Specific Setups for Torque measurements.....	177
5.6	Closing Section.....	178
6	VR Recoil	180
6.1	Introduction.....	181
6.2	Related Work.....	182
6.3	Recoil VR	185
6.3.1	Haptic Controller Design	185
6.4	FARE-KE & VR RECOIL.....	188
6.4.1	Kinesthetic feedback Modelling.....	189
6.5	User Studies.....	191
6.5.1	User Study 1: replicating shape and weight.....	191
6.5.2	Experimental setup.....	191
6.5.3	User Study 2: Linear Kinesthetic feedback vs Passive.....	194
6.5.4	User Study 3: Assessing the Impact of Kinesthetic Feedback Levels	196
6.6	Conclusions.....	199
7	Proxy Touch.....	201
7.1	Introduction.....	202
7.2	Related Work.....	202
7.2.1	Technologies for Kinesthetic feedback in VR.....	202
7.2.2	Techniques exploiting visual dominance in VR.....	203
7.2.3	Somatosensory Attenuation	204
7.3	FARE-KE & Force Measuring Setup.....	207
7.4	Proxy-Touch SA Evaluation.....	209
7.4.1	Experiment Design and Testing Procedure.....	209
7.4.2	Participants.....	210
7.4.3	Results	210

7.5	Conclusions.....	212
7.5.1	Future work.....	212
8	HapticWhirl.....	214
8.1	Introduction.....	214
8.2	Related Work.....	216
8.2.1	Torque Gyroscope Feedback Devices.....	216
8.2.2	Haptic Interfaces for VR.....	217
8.2.3	Multimodal Kinaesthetic Feedback	219
8.3	Controller Design.....	220
8.3.1	Hardware Implementation.....	220
8.3.2	Working Principle.....	221
8.3.3	Controller Haptic Features.....	223
8.4	FARE-KE & HapticWhirl.....	225
8.5	Experimental Validation	226
8.5.1	Torque Feedback.....	226
8.5.2	Impact Simulation	227
8.5.3	Vibration Feedback	228
8.5.4	Whirl Effect (Rotational Forces).....	229
8.6	Applications.....	230
8.7	Conclusions.....	231
9	Conclusion	233
9.1	FARE-KE Framework Validation and Summary	Error! Bookmark not defined.
9.2	Research Contributions	Error! Bookmark not defined.
9.3	Limitations and Considerations	Error! Bookmark not defined.
9.4	Future Work.....	Error! Bookmark not defined.
9.5	Concluding Remarks	Error! Bookmark not defined.
10	Bibliography	237
11	Appendix – Commercial Load Cells Comparison Table	264
12	Appendix – VR Recoil Modelling	265
13	Appendix – HapticWhirl Modelling.....	268
14	Appendix – Previous Amplifier Designs Used (alternatives to HX711).....	276
15	Appendix – FARE-KE Metrics Table Comparison	279
16	Appendix – FARE-KE Proposed Table Of Metrics.....	280
17	Appendix – FARE-KE DAQ MCP3654 PCB design.....	281
18	Appendix – FARE-KE Calibrating Single Axis Load Cell.....	282
19	Appendix – FARE-KE 3-Axis Load Cell Assembly Guide	286
20	Appendix – FARE-KE 3-Axis Single Point Calibration	288
21	Appendix – FARE-KE 3-Axis Multipoint Calibration Method.....	292
22	Appendix – FARE-KE 6-Axis Load Cell Assembly.....	298
23	Appendix – FARE-KE 6-Axis Load Cell Multipoint Calibration	299

24 Appendix – FARE-KE Analysis Implementations	310
25 Appendix – FARE-KE ADC Libraries	311

List of Figures

Figure 1. Flowchart outlining the structure of the literature review in this chapter, categorising haptic perception, haptic technology, haptic characteristics, and evaluation metrics.....	32
Figure 2. The multimodal sensory system, illustrating how touch, once considered a single sense, is now divided into multiple distinct modalities, including mechanoreception, proprioception, and kinesthesia, among others [7].....	33
Figure 3. Diagram illustrating the receptive fields of Merkel cells, Ruffini endings, Meissner's corpuscles, and Pacinian corpuscles, highlighting their density, function, frequency response, receptive field area, and adaptation properties. Adapted from [11].....	34
Figure 4. Illustration of proprioceptors involved in sensing force and body movement, including muscle spindles and Golgi tendon organs. Adapted from [14].....	35
Figure 5. Summary of tactile technologies: (a) Eccentric rotating mass motor, (b) linear resonant actuator, (c) solenoid actuator, and (d) piezoelectric actuator.	38
Figure 6. Various modalities of handheld kinesthetic haptic devices. Direct actuation: (a) Torc [32], (b) Claw [33]. Indirect actuation, where force is applied as torque through a handle or grip: (c) Thor's hammer [34], (d) Itorqu [35] and (e) AirRacket [36].....	43
Figure 7. Examples of handheld devices that deliver direct actuation: (a) CapstanCrunch [37], (b) NormalTouch [38], (c) Bstick [39], (d)Haptivec [40], (e) Pupop [41] and (f) pneu-multi-tools [42].....	43
Figure 8. Examples of Dynamic Passive Haptic Feedback (DPHF): (a) Shifty[46], (b) Drag:on [47], (c) Transcalibur[48], (d) ElastOscillation[49] and (e) ElaStick[50].	45
Figure 9. Examples of aerodynamic haptic devices: (a) WindCage[51], (b) Leviopole[52], (c) Wind-blaster[53] , (d) Aero-plane[54] and (e) JetController[55].....	46
Figure 10. Examples of flywheel torque haptic devices: (a) Golf-club type controller[57], (b) Itorqu[37], (c) GyroVR[58] and (d) TorqueScreen[59].	47

Figure 11. Examples of exoglove devices categorized by function. Active force exogloves (red): (a) CyberGrasp [69], (b) [25]. Braking-based exogloves (green): (c) Wolverine [70], (d) DextrES [71]. Alternative methods (blue), such as pneumatic actuation, offer kinesthetic feedback with limited control: (e) ExoTen-Glove (TSA) [72], (f) ω -jet [73].....49

Figure 12. Haptuator mk-II's magnitude-frequency response under different mechanical loadings: (a) freely suspended, (b) sandbag placement, and (c) stone base attachment.....54

Figure 13. Touch-Box device designed to measure normal forces applied to its top panel, enabling precise analysis of vibrotactile feedback delivery.55

Figure 14. Haptic ring and its characterization [82]. (a) Vibrotactile device worn on the proximal phalanx of the index finger. (b) 3D-printed housing with an opening for the actuator's moving platform, ensuring direct skin contact and secured with Velcro straps. (c) Experimental setup measuring actuator displacement under load using an accelerometer. (d) Displacement amplitude of the loaded actuator as a function of driving frequency and input voltage, peaking around 160-190 Hz.56

Figure 15. Experimental device featuring an accelerometer mounted on a stylus (left). User holding the prototype during the user study (right).....56

Figure 16. Vitaki vibrotactile actuator response: standard operation (left) compared to overdrive and braking techniques (right).57

Figure 17. Comparison of four transformations (SA321, SoC321, Mag321, DFT321) applied to a three-dimensional texture vibration signal. Left: one-dimensional signals in the time domain. Right: smoothed energy spectral density (ESD) of transformed signals (colored) and the original 3D signal (black). The DFT321 method achieves the best spectral ($M_{sm} - sm''sm = 1.00$) and temporal match ($M_{tm} - tm''tm = 0.64$), outperforming the others.59

Figure 18. Laser Doppler vibrometry (LDV) applications for haptic measurements: (a) CHASM device for vibrotactile feedback, (b) testbed for characterizing CHASM vibrations, (c) LDV setup measuring palm vibrations induced by haptic actuators [82] and (d) LDV analysis of mid-air haptic stimuli for remote tactile feedback [81].60

Figure 19. Force-sensing and haptic feedback device. (a) Standard force-sensing resistor (FSR). (b) The "Touch-and-Fold" prototype, a compact fingertip-mounted haptic device.

(c) Internal components of "Touch-and-Fold," including the force sensor, Linear Resonant Actuator (LRA), and mechanical transmission system.....	61
Figure 20. HapticLinks prototype for measuring displacement and applied force: (a) layer hinge design, (b) prototype mounted on a lathe with a force gauge, and (c) torque–angle curves illustrating mechanical response characteristics.....	64
Figure 21. CapstanCrunch prototype for direct actuation haptic feedback. (Left) Assembled device. (Right) Characterization setup for performance measurements.....	65
Figure 22. Generic force sensor design: (a) cross-section of a load cell with flexure, strain gauge placement, and mounting point, (b) detailed strain gauge construction with a resistive path on polymer film, and (c) Wheatstone bridge circuit configuration for stressed and unstressed strain gauges.....	65
Figure 23. Various types of load cells: (a) beam load cell, typically used for precision weighing; (b) bottom load cell, designed for compression force measurements; (c) S-type load cell for both tension and compression; (d) 3-axis load cell for multi-directional force measurement; and (e) force-and-torque load cell measuring both rotational and linear loads.....	66
Figure 24. Evaluation of the HairTouch prototype: (a) the prototype in use, (b) testbed setup for device characterization, and (c) plot illustrating the force versus hair-pressed distance relationship.	67
Figure 25. ImpactVest haptic prototype: (a) worn by a user, (b) detail of device mechanics, (c) characterization setup, and (d) measured impact force over time, demonstrating performance during testing.....	68
Figure 26. PropelWalker haptic device: (a) prototype worn, (b) characterization setup for the single ducted fan, and (c) graph illustrating the relationship between the fan's PWM cycle and output force.....	68
Figure 27. CLAW device: (a) in use, illustrating how it is held, (b) detail highlighting the integration of a beam load cell for real-time control, and (c) force versus angle measurements.....	69

Figure 28. OsciHead device: (a) Upward force application along Z-axis with vertically oriented actuators; (b) Rotational torque generation through asymmetric actuation along positive and negative Z-axes; (c) Characterization setup using two beam load cells to measure force output; (d) Force-extension distance linear regression showing relationship between elastic band extension and measured force.....70

Figure 29. CMG Haptic Backpack: (a) User wearing the scissored-pair CMG prototype in backpack configuration; (b) Schematic diagram of torque measurement setup showing metal plate anchored by ball joint and attached to load cells; (c) Testing setup with labeled components including load cells, controller, and data logging laptop; (d) Graphs showing measured torques at various flywheel velocities, with sagittal torques outpacing transverse ones due to the device’s gimbal motions.71

Figure 30. Thor’s Hammer device: (a) Device held by a user; (b) Device mechanics with labeled components; (c) Characterization setup with load cells and mounting frame; (d) Scatter plot comparing intended versus measured force on the X-Z plane; circles represent target force vectors, with red vectors indicating deviation from expected values.72

Figure 31. Pull-Ups Kinesthetic Feedback System: (a) User performing pull-ups assisted by a suspension kit with integrated pneumatic actuators; (b) Close-up of a load cell in line with a pneumatic muscle; (c) Characterization of the device by loading weights; (d) Graphs showing displacement and actuation time across different loads, used for system calibration and feedback control.73

Figure 32. Components of the Wrist Soft Exosuit: (a) Glove with attached 3D-printed flexible structure and Inertial Measurement Unit (IMU); (b) User wearing the complete exosuit; (c) Side view showing the wrist configuration; (d) Actuation mechanism with remote motor and pulley system; (e) Close-up of the force sensor used to detect interaction forces; (f) Forearm strap components including Bowden cable, electronic cable, and sensor acquisition board. 74

Figure 33. Shape Memory Alloys actuator characterization: (a) SMA actuator designed using the ANISMA toolkit applied to skin; (b) Testing setup with SMAs anchored to load cell for force measurement; (c) Graph showing time-dependent pull force of SMAs at various actuation currents, demonstrating response characteristics and actuation speed.....74

Figure 34. DextrES Haptic Glove characterization: (a) Glove prototype utilizing electrostatic brakes for kinesthetic feedback and piezoelectric elements for tactile sensation; (b) Pulling tester with force sensor for characterizing actuator performance, with graphs showing dynamic friction forces recorded at varying voltages; (c) Bar chart showing friction force correlation with applied voltage, demonstrating the proportional relationship between voltage and braking force.75

Figure 35. Aero-plane Haptic Device characterization: (a) Aero-plane controller held at its centre-of-mass, with air jets producing a downward force; (b) Measurement setup with the propeller actuator mounted on an aluminum frame and force sensor attached at the bottom; (c) Close-up of the 3D printed air duct and force sensor configuration; (d) Step response graph showing force measurements with raw and filtered data, demonstrating the device's dynamic force feedback capabilities.76

Figure 36. SCALE Force Interaction Framework: (a) 3-axis load cell integrated with peripheral circuits for force measurement; (b) Schematic representation of touch interaction, illustrating the dynamic nature of human contact and the intersection points used to determine touch location; (c) Augmented workbench setup showcasing real-time object status tracking through touch interactions.....77

Figure 37. Virtual Super-Leaping Haptic Feedback System: (a) User engaged in a VR environment with the haptic device designed to simulate extreme jumping; (b) Testbed setup for characterization with labeled 6-axis force sensor and sound level meter; (c) Close-up of the measurement fixture showing coordinate axes and reference measurements; (d) Graph depicting the device's force response over time for different input pulse widths, illustrating the system's step response and latency characteristics.78

Figure 38. Pneumatic haptic system: (a) Main components including accumulator, valve, and user coupling device that functions without an external pressure source by adjusting stiffness through valve operation; (b) Characterization setup with labeled components including motor, actuator, force sensor, valve, linear bearing, lead screw, pressure sensors, and capacitor, designed to measure pressure changes and output force.79

Figure 39. Wearable haptic devices and characterization methods: (a) Column showing PneuSleeve device worn on forearm with labeled components, its actuation mechanism with

lead screw and linear part, and the complete prototype; (b) Column displaying Bellowband wrist-worn device, its force measurement setup, and internal component view with embedded sensor; (c) Column presenting Tasbi wrist-worn research prototype, its force characterization setup using ATI Nano 17 sensor, and the device components.....80

Figure 40. iTorqU haptic controller evaluation: (a) iTorqU controller operated by a user; (b) Characterization setup with the controller mounted on an ATI Mini40 sensor, clamped to a work surface; (c) Response comparison graph highlighting the correlation between predicted and measured torque values along the Z-axis over a specific actuation period.81

Figure 41. Haptic golf club feedback device: (a) Proposed device in use, showcasing the direction of generated torque and the measurement point on a golf club; (b) Characterization setup with the device mounted on a six-axis force sensor with labeled axis of measurement; (c) Graph presenting a comparison between expected theoretical values and measured torque values across different motor speeds (100-140 rpm).81

Figure 42. Mantis haptic device: (a) User interaction with a multi-armed Mantis haptic device in VR environment; (b) Detailed view of the Mantis components including labeled force sensor, controller, and brushless motors; (c) Close-up of the custom-built acrylic force sensor.82

Figure 43. Multi-axis force sensor assembly: (a) Dimensions of the Alps Alpine HSFPAR003A force sensor (2.0mm × 1.6mm); (b) PCB with four HSFPAR003A load cells mounted, necessary for multi-axis force measurement; (c) Exploded schematic view detailing the sensor components including jaw attachment, sensing plate, load cell arrays, wire clamp, and strain relief bracket; (d) Assembled prototype showing the complete sensing device.....83

Figure 44. Table from Wireality [147] comparing key features (e.g., actuator type and feedback mechanism) across various haptic devices. While it omits detailed dynamic performance metrics (such as those revealed by step-response testing) it provides a straightforward overview of each device’s primary attributes.85

Figure 45. Ramp analysis plot showing the relationship between actuator input and output force, highlighting hysteresis patterns (blue and red paths) with annotated characteristics including output force resolution, static friction breakaway point, and actuation steps in 10% increments.87

Figure 46. Aero-plane characterization graphs: (a) Output force and noise level vs. pulse width driving the propeller as it increases from 1.2ms to 2.0ms, showing minimum, maximum, steady-state forces and decibel readings; (b) Raw sensor measurement and filtered sensor measurement of the jet propeller’s step response over time, with filtered curve removing high-frequency noise and showing 0.6265s response time.89

Figure 47. Step response analysis plot showing key performance metrics: input signal (blue) and output force (red) over time, with labeled measurements including rise time, settling time, fall time, peak force, maximum force, and steady-state error (magnified in inset). The plot illustrates standard haptic device dynamic behavior characteristics with 10% and 90-105% threshold markers.90

Figure 48. FARE-KE framework standardization flowchart illustrating the decision tree for haptic device characterization, from sampling frequency selection through sensor type determination, measurement approach, and analysis methods, with references to relevant appendices. 100

Figure 49 Examples of testbed configurations for haptic device characterization: (a-c) Technical diagrams showing proposed measurement setups, and (d-f) corresponding device implementations for PacaPa rotational joint [45], FacePush head-mounted system [163], and OsciHead oscillating mechanism [119]..... 105

Figure 50. Load cell designs for haptic device characterization: (a) Beam load cell; (b) Button load cell; (c) S-type load cell; (d) 3-axis load cell; (e) 6-axis force/torque load cell. 106

Figure 51. Multi-axis force sensing configurations: (a) Assembly of single-axis load cells arranged to form a 3-axis load cell setup; (b) 3-axis load cell complete with housing; (c) 6-axis force-torque plate constructed using four 3-axis load cells mounted between acrylic plates..... 107

Figure 52. Haptipedia data visualization showing the relationship between workspace size and maximum force/torque capabilities of haptic devices: (left) translational force peak vs. translational workspace size; (right) rotational torque peak vs. rotational workspace size, with both graphs using logarithmic scales to display the wide range of device capabilities. 108

Figure 53. Load cell amplification and data acquisition options: (a) HX711 breakout board; (b) MCP3564 breakout board; (c) INA125 instrument amplifier; (d) FARE-KE proposed custom DAQ board featuring multiple channels for multi-axis force measurement.	109
Figure 54. Example wiring circuit diagram for an INA125 instrument amplifier connection to a load cell, showing power supply, reference voltage configuration, and signal conditioning components required for proper load cell interfacing.....	110
Figure 55. DAQ board prototype implementation: (left) Complete PCB layout, (right) Close-up of a single channel highlighting the gain amplification resistor used to control signal sensitivity.	111
Figure 56. TyCommander software interface for managing serial communication with the Teensy 3.2 microcontroller used in the FARE-KE data acquisition system.	112
Figure 57. Telemetry Viewer interface showing real-time multi-axis force and torque measurements with color-coded channels, displaying both raw data trends and processed values for comprehensive haptic device characterization.	113
Figure 58. Complete haptic system diagram showing the integration of components: Teensy microcontroller for data acquisition, custom multi-channel DAQ board, Arduino controller for actuator management, and 6-axis force-torque plate for measurement. The diagram illustrates signal paths and connections for comprehensive characterization, with the Arduino controlling compressed air flow to the pneumatic actuator while the force-torque plate measures the resulting mechanical response. This setup enables both step response testing and ramp response analyses through various control methods including Boolean states, PWM signals, and serial communication, ensuring precise manipulation of the haptic device's actuation parameters AirRacket [38].	118
Figure 59. The FARE-KE custom data acquisition (DAQ) board with 3 MCP3564 24-bit ADCs.....	124
Figure 60. Pi Pinout connection diagram for the FARE-KE DAQ board, showing the mapping between MCP3564 ADC channels, device ports (0A-2D), scale indices (0-11), and 3-axis load cell connections (X,Y,Z axes), with color-coded channel groupings to facilitate proper wiring and software configuration.....	125

Figure 61. Simplified schematic of the FARE-KE DAQ system showing four main functional blocks: sensing, amplification/filtering, digitalisation, and voltage regulation with SPI connection to microcontroller.	126
Figure 62. Power system configuration for the FARE-KE DAQ board featuring separate LDO regulators for digital (3.3V _{Digital}) and analog (3.3V _{Analog}) supplies.	127
Figure 63. Basic circuit design between load cell and MCP3564 ADC showing basic component for digitalisation of the signal.	128
Figure 64. Detailed view of the signal chain for one channel of the FARE-KE DAQ system with load cell in Wheatstone bridge configuration, differential amplification using precision op-amps with programmable gain resistors, and direct connection to the MCP356x 24-bit ADC with separate analog and digital grounds.....	129
Figure 65. Visual comparison of ADC signal noise for different oversampling ratio (OSR) settings of the MCP3564 DAQ board, showing eight colour-coded noise bands with corresponding SNR values (24.34 to 25.68 dB), demonstrating how higher OSR values reduce noise bandwidth but maintain similar peak-to-peak amplitude.	133
Figure 66. Comparison of MCP3564 and HX711 ADC performance: (a) characterization testbed with beam load cell connected to both data acquisition systems and solenoid actuator, and (b) response curves showing lower latency (8.93 ms) for MCP3564 compared to HX711 (20.69 ms raw, 43.38 ms filtered) during impact force measurement.	135
Figure 67. Left: Experimental setup. Right: Graph showing unfiltered and filtered ADC data from MCP3564 and HX711 during the test, illustrating the sharp transient peak captured by MCP no filter and the smoothed responses due to filtering similar to that of the HX711.	136
Figure 68 Voltage readings over time for the HX711 (blue) and MCP3564 (green) ADC boards.....	138
Figure 69. Examples of haptic devices suitable for single-axis load cell characterization: (a) tactile feedback glove with bend sensors embedded in tubes on fingers for directional forces, and (b) walking-based haptic feedback device (PropelWalker) showing direction of drag force applied to user's leg during locomotion.....	141

Figure 70. Digital load cell calibration data showing linear relationship between known weight (in gram-force) and ADC digital values for three different load cells (blue, orange, green), with each demonstrating consistent response across both compression and tension forces between -3000 gf and +4000 gf. 142

Figure 71. Comparison of three calibration methods for beam load cells showing linear relationships between applied weight (-3000 to 4000 gf) and digital readings: (a) Zero Offset approach adjusting for unloaded sensor values, (b) Midpoint Zero Adjustment using average of positive and negative values, and (c) Ignoring Zero Readings method excluding zero points from regression, all yielding highly linear responses ($R_{\text{fit}} \geq 0.999$) for TAL220, Degraw, and CZL635 sensors. 145

Figure 72. Calibration model comparison for TAL220, Degraw, and CZL635 load cells showing nearly identical performance between single scaling factor, linear and polynomial regression approaches across the full measurement range (-3000 to +4000 g). 147

Figure 73. Noise profiles for three beam load cells under unloaded static conditions, showing TAL220 with highest noise amplitude (6 gf peak-to-peak), compared to Degraw and CZL635 (4 gf and 3 gf respectively), measured over a 200 μs time window. 148

Figure 74. Noise distribution histograms for three load cells under no-load conditions, showing normal distributions with different offset ranges: CZL635 (-12 to -4 gf), Degraw (-20 to -10 gf), and TAL220 (-25 to -7.5 gf), indicating consistent noise characteristics with CZL635 exhibiting the narrowest distribution. 149

Figure 75. Frequency spectrum analysis of three load cells showing similar noise characteristics with dominant peaks at approximately 10 kHz and smaller harmonic components at 20-40 kHz, with TAL220 exhibiting stronger high-frequency components compared to Degraw and CZL635. Peaks in their frequency spectrum at 0 removed. 150

Figure 76. Temperature sensitivity analysis of three beam load cells, showing their drift response (12 grams) to ambient temperature fluctuations (19.7-23.3 $^{\circ}\text{C}$) over time, with CZL635 showing greatest stability (blue line, 2 gf), Degraw exhibiting moderate temperature-related drift (green line), and TAL220 displaying highest temperature dependency (purple line). 150

Figure 77. Repeatability test for three load cells over 8 loading/unloading cycles (160 minutes), showing consistent force responses of approximately -2300 gf when loaded (red background) and return to zero during unloaded periods (green background), demonstrating stable performance and minimal hysteresis for all sensors. 152

Figure 78. Average values for three load cells across eight loading/unloading cycles, showing: (top row) consistent loading responses between 2280-2325 gf with Degraw exhibiting smallest variation around 2305 gf; (bottom row) unloading (zero) return characteristics with notable offset patterns, where CZL635 shows positive drift (2-12 gf), Degraw shows negative bias (-10 to 0 gf), and TAL220 maintains positive values (2-10 gf). 153

Figure 79. Simplified proposed measurement setup for characterising single-axis force: (a) top and side view schematics showing the testing arrangement with load cell placement, and (b) photograph of the implemented system with labelled components including aluminium extrusion structure, DAQ board and analysis tools. 155

Figure 80 Characterisation approaches for wearable haptic devices: (a) ANISMA actuator prototype mounted on skin for direct force testing, (b) test fixture for precise alignment on optical table, and (c) schematic diagram showing vertical beam load cell configuration for push/pull force measurement. 155

Figure 81. Characterisation setup for the FacePush haptic feedback device: (a) schematic diagrams of force application paths for direct head pressure and side motor-pull mechanisms with measurement points indicated by arrows, and (b) photograph of the actual device being worn by a user with VR headset and attached motorised haptic actuators. 156

Figure 82. ImpactVest characterisation setup: (a) photograph showing the experimental apparatus with labelled impactor and force sensor components for measuring direct impact forces, and (b) schematic diagram illustrating the arrangement of the impactor mechanism relative to the load cell for accurate impact measurement. 157

Figure 83. Air-based haptic device characterisation setups: (a) BLASNEL prototype schematic showing measurement approach for air-based tactile feedback, (b) user wearing the BLASNEL prototype, (c) Ω -jet prototype measurement configuration, and (d) photograph of the implemented Ω -jet actuator with backside view of nozzle arrangement. 158

Figure 84. Pull-Ups pneumatic VR system: (a) wearable haptic device with cables exerting upward forces to simulate climbing sensations in VR, (b) close-up of pulley mechanism with tension cables, and (c) schematic diagram of S-type load cell measurement setup for characterising the system’s in-line tension forces..... 159

Figure 85. Pulling tester mechanisms for characterizing tension forces in haptic devices: (a) Instron 3344L industrial tester, (b) manual pulling tester with force gauge, (c) linear actuator equipped with servo motor, and (d) proposed implementation with S-load cell for characterizing similar actuators as Dextres..... 160

Figure 86. Examples of various haptic devices using passive kinesthetic feedback mechanisms: (a) Wireality [147] hand tracking system with string-based resistance for physical object simulation in VR, (b) DextrES [71] electrostatic braking glove for variable finger resistance, and (c) wearable haptic device with cable-based force 160

Figure 87. Integration of force sensors in haptic interface designs: (a) ProxyTouch characterization testbed showing wrist rest, button-type force sensor, motor-controlled lever, and thumb strap for measuring finger forces during interaction, and (b) TORC [34] hand-held controller with embedded miniature force sensors for measuring grip pressure in VR applications..... 161

Figure 88. Force sensor setup diagrams for kinesthetic feedback controllers: (a,b) schematics showing proposed characterization method for measuring forces perpendicular to finger movement in pinching devices, and (c,d) photographs of the actual Claw and CapstanCrunch haptic controllers with their mechanical pinch-grip mechanisms that can be tested using this arrangement..... 162

Figure 89. Haptic links characterization setups: (a) alternative setup with linear actuator top view, (b) alternative setup with linear actuator side view, and (c) lathe-based testing apparatus components including spindle, haptic link, force gauge, and slide..... 162

Figure 90. Pacapa haptic device characterization: (a) device at 0fl position, (b) device rotated to 90fl position showing joint movement, and (c) schematic of testbed setup with beam load cell aligned to measure force and torque at distance L from rotation axis 163

Figure 91. OsciHead haptic feedback device: (a) linear vertical actuation movement along Z-axis, (b) rotational movement around vertical axis in clockwise direction, and (c) schematic

of proposed testbed design to measure torque around the rotation axis with distance t between measurement points.....	164
Figure 92. Alternative OsciHead characterisation testbed: schematic diagram showing measurement approach with force vectors indicated at distance L from rotation axis, enabling calculation of torque produced by the headset-mounted actuators.	165
Figure 93. Multi-axis force measurement systems: (a) schematic of standard 6-axis force plate application for robotic arm base force and torque measurement, and (b) FARE-KE custom force plate design showing arrangement of four 3-axis load cells with labeled dimensions (w , L) and coordinate system directions.	165
Figure 94. Commercial and custom multi-axis force sensors: (a) Forsentek 6-axis force/torque sensor with compact cylindrical design, (b) industrial multi-axis load cell with square mounting pattern, (c) 3-axis load cell using stacked single-axis design with aligned beam cells, and (d) custom-built load cell with integrated electronics in green housing.	166
Figure 95. Heatmap table comparison for the three calibration methods presented, showing the error as a percentage of the load applied to the axis.	169
Figure 96. Heatmap table comparison for the three calibration methods presented, showing the error in grams when loading the sensor after calibration.	170
Figure 97. 6-axis force/torque plate design used to measure force and torque on handheld devices. The four three-axis load cells are labeled A, B, C and D with dimensions L and W used to compute the output of the plate.....	170
Figure 98 Force plates applications: (a) commercial force plate being used for gait analysis and biomechanical assessment of foot impact forces during walking, and (b) schematic illustration of a robotic arm system with integrated force plate at its base to measure operational forces and torques during movement.	171
Figure 99. Error distribution analysis for load cell calibration: (a) histogram comparing raw and calibrated readings showing convergence towards zero error after calibration, and (b) box plot demonstrating significantly reduced interquartile range and median values closer to zero for calibrated measurements.	174

Figure 100. Mean Absolute Error (MAE) by axis for raw versus calibrated measurements, showing substantial error reduction across all force (F_x , F_y , F_z) and torque (M_x , M_y , M_z) components, with most significant improvements observed in the F_z axis..... 175

Figure 101. Crosstalk effect visualization: (a) error magnitude on F_z axis and (b) error magnitude on M_x axis when other axes are actuated sequentially, demonstrating significant reduction in undesired cross-axis interference after applying the calibration matrix. 176

Figure 102. Gyroscope-based haptic controllers with force/torque measurement: (a) iTorqu handheld device with gimbal-mounted flywheel, and (b) experimental characterisation setup using ATI Mini45 force/torque sensor with control moment gyroscopes..... 177

Figure 103. Arrangement for handheld devices testbed mounting: (a) side view of haptic device mounted at the base of the handle on force/torque plate, (b) Ideally mounting will consider user handling (c) Forces at a distance can be calculated based on rotation and length arms variables..... 177

Figure 104. In the left image, you can see the measuring setup used to capture the controller’s response. A button load cell (outlined in blue) is employed to record the actuation force exerted by the controller. To reduce friction and enable the controller to move freely in a horizontal direction, it is mounted on a bearing carriage (highlighted in yellow). In the right image, you can see the enclosure (blue) that houses the pneumatic piston (red). When the trigger is pressed, the piston which on the home position is retracted moves the weight (white) back and forth within the enclosure colliding with the end cap of the enclosure.
..... **Error! Bookmark not defined.**

Figure 105. left simplified model of the haptic prototype to estimate force impact vs Pressure. Right impulse response from our three calibrated settings. 190

Figure 106. image of the layout of the virtual environment with all the bullseye shown as well as the platform where the participant w..... 191

Figure 107. From left to right Distribution of Hit Distance results on metres to the centre of the target. Time per shot on seconds, % of hit and miss per condition. Right boxplot shows result from Q1, Q2 and Q3 from a 0-7 Likert scale..... 193

Figure 108. From left to right Distribution of Hit Distance results for experiment one. Time per shot for the study 1, % of hit and miss per condition. Right boxplot shows result from questions 1, 2 and 3	195
Figure 109. Linear model regression for the 4 conditions and r-square value, shows the models obtained and the fitting model for each. The X axis represents the normalized distance from 0 to 10 while the Y axis represents the number of shots taken.....	198
Figure 110. (a) Hand gesture during knob interaction in virtual space. (b) Real hand gesture in real space. (c) Retargeting representation from virtual to real space the user perceives an object pin.....	201
Figure 111. Depiction of the Force Matching Paradigm. (a) Presented Force, characterized by constant forces applied to the subject's left index finger using a motor actuator. (b) Matching Force produced by the subject's right index finger pressing a sensor directly above the left index finger (0 cm), which controls the force output on the opposing finger. (c) Matching Force generated by the subject's right index finger pressing a sensor positioned 25 cm away from their left index finger, controlling the force output on the opposite finger. (d) Attenuation index calculated as the difference between matched force and presented force.	205
Figure 112. device used for the Force Matching Paradigm.....	207
Figure 113. (a) Presented Force applied by a lever (1N, 1.5N, 2N, 2.5N, 3N, 3.5N) on participants fixed thumb. (b) Direct Matching Force applied by participants' index finger (skin contact). (c) Indirect Matching Force applied by the index finger with a physical proxy in between (5cm). In both b and c, the setup was either visible (directly in front) or not visible (from top view) by participants. (d) Plot for comparison of Presented Vs Matching Forces. (e) Mean attenuation Matched Force - Presented Force.	211
Figure 114. The physical components of the HapticWhirl controller. Blue color is used to highlight the components that from part of the Z axis gimbal, green for the Y axis and yellow for all the components of the braking mechanism.	220
Figure 115. The controller's inputs and outputs. Left components represent the gimbal control sensors and actuators, on right flywheel components and communication.	221

Figure 116. Operational Modes of the HapticWhirl Controller. (a) Torque Mode: Actuation of gimbals alters the flywheel’s orientation, generating torque. (b) Vibration Mode: Rapid pitch-axis oscillation of the flywheel produces tactile vibrations. (c) Impact Mode: Sudden halting of the flywheel simulates an impact. (d) Inertial Mode: Delayed flywheel response creates a sensation of drag. (e) Whirl Mode: Continuous yaw-axis actuation induces a stirring sensation..... 223

Figure 117. (a) Measurement of the magnitude of HapticWhirl’s output on a single axis. (b) A breakdown of forces for each axis - X (red), Y (green), and Z (blue)..... 227

Figure 118. Output torque resulting from suddenly stopping the flywheel while it is rotating at full speed. 227

Figure 119. (a) Output from back-and-forth rotation of the flywheels gimbal (pitch) with the flywheel at rest. (b) Output the same rotation, but with the flywheel rotating at full speed. In both scenarios, the gimbal was actuated at a frequency of 25Hz..... 228

Figure 120. Measured outputs of the HapticWhirl when creating a whirl effect. The top row shows the rotation of the flywheel around its center over the handle. The middle row displays the outputs and position when rotated with alignment to the handle’s X-axis. The bottom row illustrates the rotation of the flywheel around the handle’s Z-axis. In all instances, the active rotation axis was rotated at approximately 450 RPM..... 229

Figure 121. Left 3D model of a momentum wheel, illustrating the angular velocity. Right 3D model of a steered momentum wheel, showcasing the rotation vectors..... 268

Figure 122. Illustration of the HapticWhirl kinematic model showing the A, B, C, and D coordinate frames with axes labeled in red, green, and blue, representing the different components of the system including the disk and gimbals..... 269

Figure 123. Rotation about the Y-axis illustrating the rotation matrix “(R’-C “rightarrow A” “). 270

Figure 124. Left Rotation from frame A to frame D through intermediate frame C. Right Rotation matrix for Z-axis transformation from frame C to D 274

Figure 125. Left final 3 axis load cell, Right layout and overlap of the single load cell bars. 286

Figure 126. Heatmap table comparison for the three calibration methods presented, showing the error on grams when loading the sensor after calibrated.	295
Figure 127. Heatmap table comparison for the three calibration methods presented, showing the error as a percentage of the load applied to the axis.	296
Figure 128. Error Distribution for Raw and Calibrated Readings histogram and Error Distribution (Box Plot) for Raw and Calibrated Readings.....	306
Figure 129. MAE by Axis for Raw vs. Calibrated bar chart.....	307
Figure 130. As axis are actuated, they have an undesired effect on other axis called crosstalk. This plot shows the crosstalk effect on the Fz axis above and the Mz axis below when another axis are actuated. On orange are the raw values when compensated with the calibration matrix K^{cal}	307
Figure 131. Scatter plots for Actual vs. Estimated Fx, Fy, Fz, Mx, My, and Mz, with readings on Newton and Newton/meter respectively.	307

List of Tables

Table 1. Aggregated Sensory Thresholds and Body Capabilities for Kinesthesia: This table is an adapted summary of various tables from Samur et al., focusing on the kinesthetic perception of limb motion and position, sensory thresholds of the human hand, and human body capabilities related to kinesthetic perception. Data sourced from various studies as detailed in Performance metric for haptic [1].....	37
Table 2. Expanded and adapted table from [18] [5], [19], [20], [21]. These values are generalized; actual parameters may differ significantly depending on specific models, manufacturers, and application environments.....	39
Table 3. Summary of accelerometers used on previous work to measure vibrations of vibrotactile actuators. Cost might vary depending on the package of the chip. *1 prices based on https://www.lcsc.com/ as 01/23.....	52
Table 4. Comprehensive haptic device characterization metrics table showing three analysis categories (Physical Properties, Ramp Analysis, and Step Response Analysis) with associated variables, units of measurement, and relevance to device performance.	92
Table 5. Metrics and variables for physical characterization of haptic devices.	117
Table 6. Variables measured using the ramp response analysis, including seven key metrics for haptic device characterization. Each metric is presented with its appropriate measurement unit and relevance to device performance evaluation.	119
Table 7. Variables measured using the step response analysis, presenting five time-domain performance metrics for haptic device evaluation. Each variable is shown with its appropriate unit of measurement and functional relevance to device performance assessment.	121
Table 8. The MCP3564 can internally provide an OSR from 32 to 98304. A higher OSR increases the number of samples taken and averaged for each output data point, filtering out noise but also reducing the data rate. To evaluate the OSR's impact on DAQ performance, digital raw readings were collected from the ADC connected to a beam load cell at various OSR settings down to 220 SPS.....	133

Table 9. Sampling rate comparison at different oversampling ratio (OSR) settings for the FARE-KE DAQ board, showing exponential decrease from 11.59 kSPS to 0.22 kSPS for single-channel operation and from 2.8975 kSPS to 0.055 kSPS when using four channels simultaneously.	134
Table 10. Lag delay comparison between MCP3564 and HX711 ADCs, showing MCP3564 filtered data with only 8.9 ms delay compared to 20.64 ms for unfiltered HX711 and 42.50 ms for filtered HX711, all measured relative to unfiltered MCP3564 signal used as reference baseline.....	135
Table 11. Comparison of impulse and force measurements between MCP3564 and HX711 ADCs, showing MCP3564 capturing higher values (0.39 Ns impulse, 1.56 N peak force) compared to HX711 (0.345-0.335 Ns impulse, 1.38-1.34 N peak force), with filtering having minimal effect on MCP3564 measurements but slightly reducing HX711 values.....	137
Table 12. Excitation voltage stability comparison between HX711 and MCP3564 ADCs, showing slightly better performance from MCP3564 with lower standard deviation (0.0133 V vs 0.0140 V) and smaller voltage range (0.2533 V vs 0.3119 V), despite different mean voltage levels (3.2840 V vs 4.2913 V).	137
Table 13. Projected sampling rate comparison for the FARE-KE DAQ board across different configurations, showing current rates with internal 5 MHz clock for single channel (11.59 to 0.22 kSPS) and four channels (2.8975 to 0.055 kSPS), alongside theoretical rates using 20 MHz external clock which would yield four-channel throughput from 38.4 to 0.6 kSPS...	139
Table 14. Performance characteristics of three beam load cells tested in the FARE-KE framework, comparing TAL220 (10 kg capacity), Degraw (5 kg), and CZL635 (5 kg) models with their respective sensitivities (1787.65, 2517.44, and 3033.73 counts/gf), operational ranges, and resolution (0.56, 0.4, and 0.33 mgf).....	143
Table 15. Zero-load offset values for three beam load cells (TAL220, Degraw, CZL635) in different orientations, showing significant differences between upward and downward mounting positions that must be accounted for during calibration.	144
Table 16. Average zero-load offset values for three beam load cells, showing TAL220 with -914462.5 counts (-10.90% of full scale), Degraw with -659200 counts (-7.86%), and CZL635	

with 689936.5 counts (8.22%), highlighting the importance of offset compensation during calibration..... 145

Table 17. Comparison of three calibration methods for each beam load cell, showing consistent slopes (1787.65, 2517.44, and 3033.73 counts/gf) across all methods with excellent linearity (R-squared ≥ 0.9999) but varying intercept values, demonstrating that calibration method primarily affects offset compensation while maintaining similar maximum errors (11.45-15.63 gf). 146

Table 18. Calibration equation comparison for three load cells using single scaling factor, linear regression, and polynomial regression approaches, showing minimal variation in primary conversion coefficients across methods..... 147

Table 19. Calibration method comparison for three load cells showing RMS errors and Rffl values, with polynomial regression yielding lowest errors (6.98-9.00 gf) followed by linear regression (7.37-9.17 gf) and single scaling factor (8.61-10.27 gf), though all methods achieve excellent linearity (Rffl ≥ 0.9999). 148

Table 20. Noise characteristics comparison for three load cells showing CZL635 with lowest standard deviation (1.193 gf) and peak-to-peak noise (9.784 gf), Degraw with moderate values (1.226 gf SD, 10.152 gf peak-to-peak), and TAL220 with highest variability (2.163 gf SD, 17.105 gf peak-to-peak), all maintaining good SNR (45-48 dB). 149

Table 21. Temperature-drift correlation analysis for three load cells using multiple statistical measures, showing TAL220 with strongest temperature dependence (correlation coefficients 0.93-0.97), Degraw with strong negative correlation (-0.78 to -0.93), and CZL635 with least temperature sensitivity (coefficients 0.07-0.54).result indicates that the TAL220 has a strong linear relationship between temperature and readings..... 151

Table 22. Repeatability data comparison for three load cells, showing standard deviation during repeated loading and unloading cycles with Degraw offering best loading consistency (4.06 gf), TAL220 showing best unloading consistency (2.36 gf), and CZL635 exhibiting highest variability during loading (10.26 gf). 153

Table 23. Maximum absolute deviation comparison for three load cells during repetitive testing, showing Degraw with best consistency under load and TAL220 with most stable unloaded readings. 154

Table 24. Comparison of four load cell options for haptic device characterization, including Honeywell FSS015WNSB (1.5 kg miniature compression, €70), Phidget button load cells in 50 kg and 200 kg capacities (both €45, 25 mm diameter), and Micro Load Cell DYLY-106 S-type (10 kg tension/compression, €50), each offering different form factors and measurement ranges for various testing scenarios.	159
Table 25. Validation performance comparison across different modelling techniques for load prediction, quantifying models' accuracy using Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Maximum Error in grams (g).....	168
Table 26. Mean Absolute Percentage Error (MAPE) and R-squared metrics before and after using the calibration matrix, demonstrating improved measurement accuracy with calibrated values.	174
Table 27. Crosstalk error measurements comparing raw and calibrated values after applying the calibration matrix to loading values along the Fy axis, demonstrating reduced interference across all force and torque channels.....	175
Table 28. Statistical comparison of measurement errors before and after calibration showing substantial improvements in standard deviation, mean values and maximum errors across all six axes, confirming the effectiveness of the calibration matrix.	176
Table 29. Readings from the 3axis load cell with a 150g load aligned on each of the axis.	290
Table 30 Digital reading for the loading of each of the 3 axis to compute calibration matrix	293
Table 31. Validation Performance Comparison across different modelling techniques for load prediction. The table quantifies the models' accuracy using Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Maximum Error, all measured in grams (g).	295
Table 32. shows an example of the calibration data collected for the force and torque measurements.	304
Table 33 Mean Absolute Percentage Error (MAPE) and R-squared metrics before and after using the calibration matrix.	306

Table 34. Crosstalk Error From Raw Readings and calibrated Values after applying the calibration matrix to the loading values along the Fy axis. 308

Table 35. Standard deviation, mean values and max error for the error on the raw readings, and the errors after applying the calibration matrix to the data. 309

Acronyms

Acronym	Definition
ADC	Analog-to-Digital Converter
AR	Augmented Reality
CAD	Computer-Aided Design
CI	Confidence Interval
CMRR	Common-Mode Rejection Ratio
CMG	Control Moment Gyroscope
CSV	Comma-Separated Values
DAQ	Data Acquisition System
DC	Direct Current
DOF	Degrees of Freedom
DPHF	Dynamic Passive Haptic Feedback
EA	Error Angle
EDR	Effective Data Rate, used in context of ADC data handling
EMI	Electromagnetic Interference
EMS	Electrical Muscle Stimulation
ENOB	Effective Number of Bits
ERM	Eccentric Rotating Mass
ES	Electrostatic
ESC	Electronic Speed Controller
ESD	Energy Spectral Density
F/T	Force/Torque
FA	Fast-Adaptive
FARE-KE	Framework for Affordable, Reliable Kinesthetic Evaluation
FMP	Force Matching Paradigm
FPS	Frames Per Second
FSR	Force Sensing Resistor
GUI	Graphical User Interface
HD	Hit Distance
HCI	Human-Computer Interaction
HMD	Head-Mounted Display
IC	Integrated Circuit
IMU	Inertial Measurement Unit
IQR	Interquartile Range, a measure of statistical dispersion
KP	Knowledge Of Performance
KR	Knowledge Of Results
LDO	Low-Dropout Regulator
LDV	Laser Doppler Vibrometry
LFF	Linear Force Feedback
LRA	Linear Resonant Actuator
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error
MCU	Microcontroller Unit
MEMS	Microelectromechanical Systems
MR	Mixed Reality
MXL	Miniature Extra Light

OSR	Oversampling Ratio
PA	Passive Actuation
PCB	Printed Circuit Board
PD	Proportional-Derivative
PGM	Pneumatic Gel Muscles
PLA	Polylactic Acid
PSRR	Power Supply Rejection Ratio
PWM	Pulse Width Modulation
RA	Rapid-Adaptive
RMA	Rotating Motor Actuator
RMIS	Robot-assisted Minimally Invasive Surgery
RMS	Root Mean Square
RMSE	Root Mean Square Error
RPM	Revolutions Per Minute
SA	Somatosensory Attenuation
SA	Slow-Adaptive
SD	Standard Deviation
SEM	Standard Error of the Mean
SMA	Shape Memory Alloy
SNR	Signal-to-Noise Ratio
SP-CMG	Scissored-Pair Control Moment Gyroscope
SPI	Serial Peripheral Interface
SPS	Samples Per Second
THD	Total Harmonic Distortion
TPS	Time Per Shoot
TSA	Twisted String Actuation
TSE	Traditional Sandwich Element
USB	Universal Serial Bus
VE	Virtual Environment
VR	Virtual Reality

1 Introduction

The rapid evolution of hardware and software technologies has transformed haptic device design and prototyping, incorporating advanced digital fabrication methods. These innovations have outpaced evaluation methodologies, which remain crucial for ensuring consistent assessment across different devices. Foundational texts such as "Performance Metrics for Haptic Interfaces" [1] have tackled these issues, but updated methods that reflect current advances in device design and fabrication are needed.

This thesis chronicles how my research journey, which began with developing novel haptic prototypes, ultimately led to the creation of a comprehensive characterisation framework. Each prototype presented unique measurement challenges that revealed limitations in existing approaches to haptic device evaluation. These experiences, particularly during COVID-19 restrictions when laboratory access was limited, necessitated creative solutions with accessible components. The methodologies developed across these projects gradually lead to FARE-KE - a framework that systematises the insights gained from these practical challenges into a cohesive approach for haptic device characterisation.

Broader Impact of VR Research and Accessibility

Virtual reality technology has profound implications beyond the technical research community, affecting multiple sectors of society. The transformative potential of VR in education and training is particularly significant, as immersive environments enable experiential learning previously impossible in traditional settings. Medical students can practice surgical procedures without risk to patients, emergency responders can train for disaster scenarios safely, and technical skills can be developed through repeated practice with immediate feedback—all facilitated by effective haptic interfaces that provide realistic tactile responses.

In psychology and cognitive science, VR offers unprecedented opportunities to study human behavior and perception in controlled yet realistic environments. Research on spatial cognition, social interaction, and therapeutic interventions for conditions like PTSD and phobias has advanced significantly through VR applications. However, these societal benefits depend critically on accurate, reliable haptic feedback that mimics real-world interactions.

Making VR research more accessible through frameworks like FARE-KE therefore serves a broader societal purpose. When haptic device evaluation becomes more standardized and accessible, researchers from diverse fields—not just engineering specialists—can contribute to advancing VR technology. This democratization accelerates innovation and ensures that VR applications are developed by multidisciplinary teams with varied perspectives, ultimately creating more inclusive and effective virtual experiences that can benefit education, healthcare, psychological wellbeing, and many other domains.

Standardization in Haptic Device Evaluation

The standardisation of evaluation processes is critical for maintaining reliable assessments across various haptic device models. Traditional evaluation methodologies connected to mechanical and engineering frameworks have not kept pace with changes brought by digital fabrication technologies. These methods often require costly and specialised equipment, making them less accessible and adaptable to modern rapid prototyping practices. This creates a significant gap in the evaluation landscape, highlighting the need for innovative methodologies that align with contemporary haptic device design while facilitating easier access for diverse designers and researchers.

This thesis is dedicated to developing hardware and software solutions for characterising kinesthetic haptic devices, with a focus on handheld models. The insights gained from creating diverse haptic prototypes inform a framework whose components are released as open source to foster further development and community contributions. A comprehensive review of both kinesthetic and vibrotactile characterisation provides a solid foundation for future enhancements. The framework includes an array of sensors, electronic design tools, data collection techniques, and processing methods chosen for their accessibility and cost-effectiveness, allowing researchers to easily acquire components, program development boards, connect controllers, and follow straightforward instructions to obtain detailed physical performance metrics.

Foundational works such as "Performance Metrics for Haptic Interfaces" [1] and "Performance Measures for Haptic Interfaces" [2] have received limited citation counts (84 and 299 respectively) compared to the number of devices presented since their publication. The methodologies detailed within these texts primarily cater to an engineering audience, necessitating specialised knowledge for implementation. This requirement reflects the

inherently complex nature of haptic device design, as described by Kern et al. in "Engineering Haptic Devices" [3]:

"Haptic systems are mechatronic systems, incorporating powerful actuators, sophisticated kinematic structures, specialised sensors, and demanding control structures as well as complex software. The development of these parts is normally focus of specialised areas of specialists, i.e. mechanical engineers, robotic specialists, sensor and instrumentation professionals, control and automation engineers and software developers. A haptic system engineer should be at least able to understand the basic tasks and procedures of all of these professions, in addition to the required basic knowledge about psychophysics and neurobiology outlined in the last chapters".

This complexity creates challenges for junior researchers and PhD students who may lack the extensive knowledge typical of an experienced haptic system engineer. The implementation and evaluation methods detailed in academic literature often demand significant effort and expertise, creating barriers to entry without the support of a skilled, interdisciplinary team. Thus, there is a need for more accessible evaluation methods that can be integrated into the haptic design process by researchers at various levels of expertise.

Advancements in technology, including digital fabrication techniques, 3D modelling, printing, and electronic development platforms, have significantly lowered barriers to entry in electronics expertise and manufacturing. These developments have broadened haptic controller design, incorporating diverse approaches such as pneumatic, fluidic, chemical, EMS (Electrical Muscle Stimulation), and turbine-based systems into handheld kinesthetic controllers. Moreover, the growing accessibility of affordable virtual reality headsets has further democratised VR technology, leading to a focus on untethered handheld devices. This evolution highlights the need for evaluation methods that accommodate the expanding variety of haptic device designs and technologies.

Another recurring issue in haptic device characterisation is the lack of practical guidelines that can be implemented without extensive mechanical engineering knowledge. This deficiency often limits the adoption of characterisation processes, presenting significant challenges for designers and researchers new to this field. The persistent concern regarding the absence of standardisation and detailed reporting protocols has been documented in numerous scholarly articles. For example:

“Haptipedia [4](CHI – 2019) Quote–Hopes for improved metrics and standardization – Several of our interviewees noted the lack of standards in the field and hoped that the Haptipedia taxonomy can provide metrics and protocols for device characterisation. P7DevD:“In all the papers they just tell you ... I built it, and it works ... what does that mean?”

The solution, however, was not straightforward. P7DevD: “I don’t really know what to do... there’s not a specific standard in haptics.”

“Getting Your Hands Dirty Outside the Lab: A Practical Primer for Conducting Wearable Vibrotactile Haptics Research [5] (TRANSACTIONS ON HAPTICS – 2019) Quote– While this article identifies a number of areas for future improvement, we argue that the most urgent issue in the haptic community at this moment is the lack of generalizability and comparability of studies’ findings. We therefore invite researchers to pay particular attention to developing:

- Standardised assessment methods.
- Standardised coupling methods and reporting, such as measurement procedures (e.g., load cell placement).

Contributions to these two aspects would allow the generation of more cohesive, consistent, repeatable, and comparable results that would strengthen the haptics community, bringing it closer to that of its audio and visual counterparts.”

Implementation and Characterisation of Vibrotactile Interfaces [6] (2018)“the characterisation of haptic prototypes—together with their technical documentation—allows reproducible implementations and enables other users and designers to carry on research and development, rather than resulting in one-of-a-kind devices.”

To tackle these challenges, this thesis introduces FARE-KE, the “Framework for Affordable, Reliable Kinesthetic Evaluation,” which emerged organically from my experiences developing and characterising several novel haptic controllers. This framework offers a standardised and accessible method for measuring and modelling the operation of kinesthetic feedback controllers. By employing affordable hardware and providing clear, step-by-step instructions,

FARE-KE assists researchers of varying backgrounds and skill levels in effectively evaluating their haptic devices.

1.1 Ethos & Main Contribution

The framework guiding this thesis is based on open-source technology principles. This approach aims to build a collaborative community among researchers and practitioners in haptic technology, promoting innovation, collaborative problem-solving, and democratising access to haptic developments. By providing comprehensive information freely available to the public in the form of a detailed "Tutorial," haptic technology becomes more accessible, supporting a circular economy model that empowers individuals with limited resources.

Publications advocating an open-source ethos highlight several common themes: firstly, haptic technology becomes more inclusive, facilitating entry for new practitioners and nurturing community-oriented values; secondly, open-source haptic designs encourage iterative improvements, leading to better interfaces and collectively advancing haptic design; thirdly, these open designs offer an educational platform supporting academic programmes, embedding the characterisation of haptic devices within the design process.

The primary contribution of this thesis is the development of FARE-KE, a framework aimed at promoting the democratisation of haptic research within the open-source community. The framework offers detailed guidance on selecting appropriate hardware and software, along with methods for measuring and modelling the operation of kinesthetic feedback controllers. It establishes clear analytical procedures, identifies comparable features across various haptic devices, and outlines properties to ensure controller performance is assessed using standardised metrics and methods.

Moreover, this thesis includes an extensive literature review covering the most widely used haptic modalities and practical methods employed by researchers to characterise haptic devices, supported by insightful case studies. The review primarily emphasises previous works providing visual or clearly descriptive information about researchers' measurement setups and methodologies. This analysis serves as a valuable resource, offering guidance on sensor applications and integration within the framework outlined in this study.

Additionally, this research presents three novel haptic controllers developed to investigate various scenarios and interaction properties within a virtual environment. Each controller is

designed with specific objectives, which will be explored in greater depth through subsequent research questions. By combining the Kinesthetic Feedback Modelling Framework, case studies, and the literature review, this thesis contributes to the field of haptic research by providing a standardised and accessible method for evaluating kinesthetic feedback controllers, clearly organising the core research questions across five chapters.

1.2 Research Questions

RQ1: How can a standardised and accessible framework be developed to make haptic device characterisation more widely adopted by the research community? This research question is addressed in the thesis through the development and presentation of the FARE-KE framework in Chapters 3 and 4. The framework provides comprehensive guidelines for the physical evaluation of haptic devices, making the characterisation process more accessible and adoptable by researchers with varying levels of expertise.

RQ2: What are the key components and methodologies required to establish a foundation for standardised measurements, enabling the comparison of diverse haptic devices across various modalities? The thesis successfully addresses this question by proposing the essential elements and techniques needed for a standardised approach to characterizing and comparing haptic devices. Chapters 3, 4, and 5 present the key components of the FARE-KE framework, including guidelines for hardware and software selection, data acquisition, performance analysis, and sensor validation.

RQ3: How can the versatility of the proposed framework be demonstrated in characterizing a wide range of kinesthetic feedback modalities? The versatility of the FARE-KE framework is demonstrated through case studies in Chapters 6, 7, and 8. Each case study applies the framework to different kinesthetic feedback modalities, validating its adaptability and effectiveness in evaluating diverse kinesthetic haptic devices.

RQ4: What is the impact of linear kinesthetic feedback on user performance and experience in virtual reality shooting scenarios? The VR Recoil case study in Chapter 6 addresses this research question by examining the effects of linear kinesthetic feedback on user performance, immersion, and learning in VR shooting experiences. The study provides insights into how linear kinesthetic feedback influences user experience in virtual environments.

RQ5: Can kinesthetic feedback be effectively recreated using the user’s own body, without relying on external proxies? The Proxy Touch case study in Chapter 7 explores the feasibility of using the user’s own body to generate kinesthetic feedback in VR interactions. The study examines the somatosensory attenuation (SA) effect in pinch gestures, demonstrating the potential for recreating kinesthetic feedback without external proxies.

RQ6: How can continuous rotational flywheels be utilised to simulate fluid motion haptic effects, and what are the key considerations for characterizing such devices? The HapticWhirl case study in Chapter 8 addresses this research question by investigating the capabilities of continuous rotational flywheels in recreating fluid motion haptic effects. The study discusses the challenges and techniques for characterizing such devices using the FARE-KE framework, providing insights into key considerations for evaluating continuous rotational flywheel-based haptic devices.

Research Question	Objective	Chapter
RQ1: How can a standardised and accessible framework be developed to make haptic device characterisation more widely adopted by the research community?	<ol style="list-style-type: none"> 1. Develop a standardized, accessible framework for haptic device characterisation 2. Provide open-source hardware and software components with step-by-step instructions for various expertise levels 3. Encourage collaboration and foster a sense of community within the haptic research field 	<p>Chapter 3: FARE-KE Framework Overview Chapter 4: FARE-KE Data Acquisition System</p>
RQ2: What are the key components and methodologies required to establish a foundation for standardised measurements, enabling the comparison of diverse haptic devices across various modalities?	<ol style="list-style-type: none"> 1. Identify and define comparable features and properties of haptic devices across various modalities 2. Establish consistent metrics and procedures for evaluating and comparing haptic devices 3. Integrate findings from a comprehensive literature review to inform framework development 	<p>Chapter 3: FARE-KE Framework Overview Section 3.4: Load Cell Count and Type Section 3.5: Considering Range and Capacity Section 3.8: Calibration Process Section 3.9: Device Characterisation Process</p>
RQ3: How can the versatility of the proposed framework be demonstrated in characterizing a wide range of kinesthetic feedback modalities?	<ol style="list-style-type: none"> 1. Investigate and evaluate different techniques for kinesthetic feedback 2. Determine the most suitable and versatile approach for various kinesthetic feedback modalities 3. Validate the versatility and effectiveness of the approach through case studies and practical implementation 	<p>Section 3.2: Initial Considerations Section 3.3: Testbed Designs Chapter 6: VR Recoil Case Study Chapter 7: Proxy Touch Case Study Chapter 8: HapticWhirl Case Study</p>
RQ4: What is the impact of linear kinesthetic feedback on user performance and experience in virtual reality shooting scenarios?	<ol style="list-style-type: none"> 1. Assess the effect of linear recoil on aim task performance, immersion, and learning in VR shooting experiences 2. Adapt and apply the developed framework to accommodate different haptic device modalities 3. Validate the effectiveness of the approach through user studies and practical implementation 	<p>Chapter 6: VR Recoil Case Study Section 6.4: FARE-KE & VR RECOIL Section 6.5: User Studies</p>
RQ5: Can kinesthetic feedback be effectively recreated using the user's own body, without relying on external proxies?	<ol style="list-style-type: none"> 1. Investigate the feasibility of leveraging the user's own body to generate kinesthetic feedback in VR interactions 2. Assess the potential of utilising mutual finger contact during pinch gestures as a means of providing haptic feedback 3. Validate the effectiveness of the approach through user studies and practical implementation 	<p>Chapter 7: Proxy Touch Case Study Section 7.3: FARE-KE & Force Measuring Setup Section 7.4: Proxy-Touch SA Evaluation</p>
RQ6: How can continuous rotational flywheels be utilised to simulate fluid motion haptic effects, and what are the key considerations for characterizing such devices?	<ol style="list-style-type: none"> 1. Explore the capabilities of continuous rotational flywheels in recreating fluid motion haptic effects 2. Identify the challenges and techniques for characterizing such devices using the FARE-KE framework 3. Validate the effectiveness of the approach through practical implementation and performance analysis 	<p>Chapter 8: HapticWhirl Case Study Section 8.4: FARE-KE & HapticWhirl Section 8.5: Experimental Validation</p>

2 Situating our Work

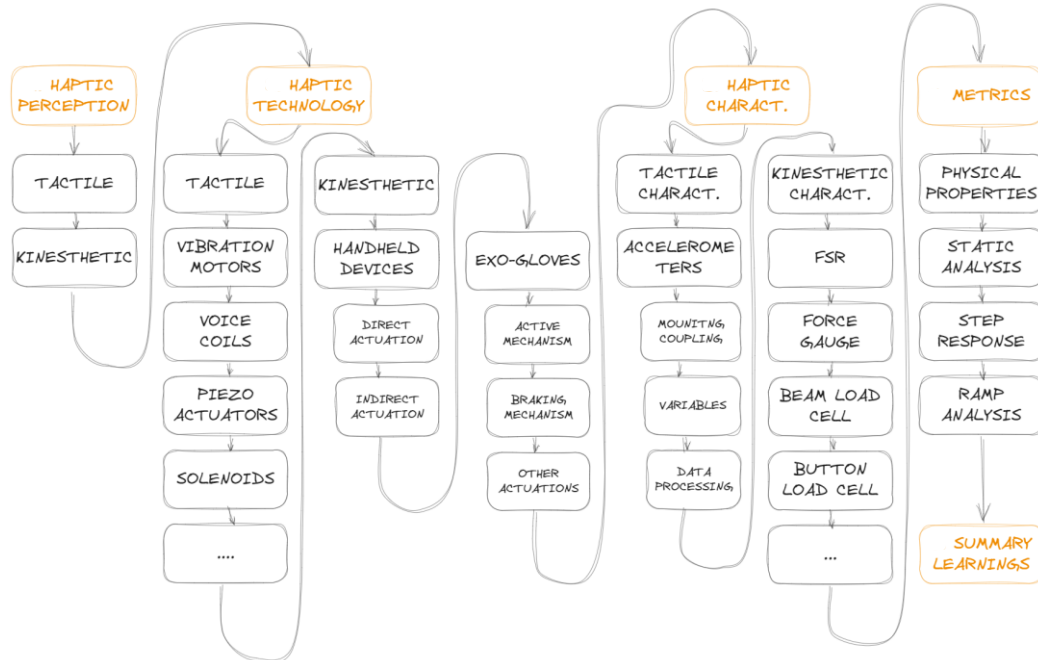


Figure 1. Flowchart outlining the structure of the literature review in this chapter, categorising haptic perception, haptic technology, haptic characteristics, and evaluation metrics.

The creation of a comprehensive framework for evaluating and characterising haptic devices requires a thorough understanding across various fields. Haptic design involves specialisations such as mechanical engineering, robotics, sensor technology, control and automation, software development, and psychophysics. This chapter explores key areas relevant to establishing the groundwork for the FARE-KE framework.

The analysis explores haptic perception, focusing on tactile and kinesthetic senses, and examining human dynamics, capabilities, and limitations essential for haptic device design. The discussion then covers haptic technologies, from eccentric rotating motors (ERM) to innovative mid-air technologies, and reviews kinesthetic actuators documented in existing research.

The third section examines haptic controller characterisation methods, setups, and sensors previously employed in research, addressing both tactile and kinesthetic techniques. It critically assesses elements of the characterisation process to form the FARE-KE framework

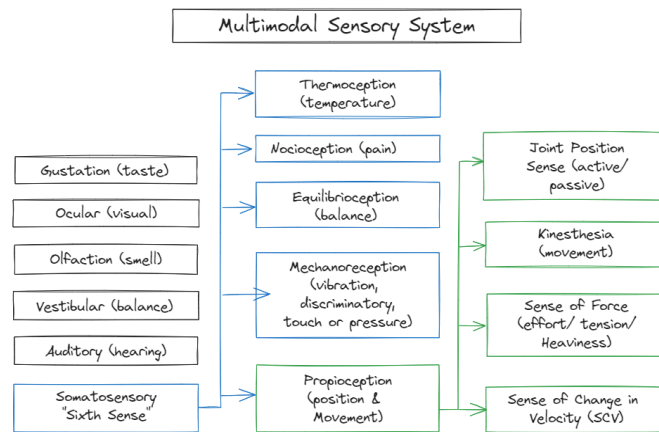


Figure 2. The multimodal sensory system, illustrating how touch, once considered a single sense, is now divided into multiple distinct modalities, including mechanoreception, proprioception, and kinesthesia, among others [7]

core, providing guidelines for force range and sampling rate requirements. The concluding section summarises insights to ensure clarity and practical applicability.

2.1 Understanding Haptic Perception: Tactile and Kinesthetic Feedback



Haptic perception, known as the sense of touch, is a sophisticated sensory system enabling environmental interaction through skin contact. This "somatic sense," often called the sixth sense [7], detects physical contact and responds to temperature, pressure, texture, and vibration variations. It relies on proprioceptors and mechanoreceptors distributed throughout the body, which transmit information to the central nervous system and cerebral cortex for processing [8]. Understanding this system's neurophysiological basis and perceptual limitations is crucial for developing effective haptic interfaces that deliver meaningful tactile experiences for daily activities like object manipulation and tool use.

This literature review specifically examines mechanoreception (tactile feedback) and proprioception (kinesthetic feedback), two senses with distinct functions and sensory receptors:

- **Tactile Feedback:** This involves detecting sensations like pressure, vibrations, and textures through the skin, enabling individuals to perceive and explore objects.

- Kinesthetic Feedback:** This sense uses information from proprioceptors to inform the brain about the body's position, size, and weight relative to its surroundings, assisting in navigation and spatial awareness.

Understanding these senses is vital for developing effective haptic interfaces that offer users rich and immersive tactile and kinesthetic experiences.

2.1.1 Tactile Feedback

Tactile feedback (Figure 3), mediated through the skin (the body's most extensive sensory organ), enables perception of pressure, temperature, pain, and other sensations essential for environmental interaction. Various mechanoreceptors within the skin respond to distinct mechanical stimuli. The primary mechanoreceptors Merkel cells, Ruffini endings, Meissner's corpuscles, and Pacinian corpuscles, each facilitate specific aspects of tactile perception [9][10].

These mechanoreceptors have distinct receptive fields that transmit sensory information. Pacinian and Ruffini corpuscles with their large receptive fields convey information about




Figure 3. Diagram illustrating the receptive fields of Merkel cells, Ruffini endings, Meissner's corpuscles, and Pacinian corpuscles, highlighting their density, function, frequency response, receptive field area, and adaptation properties. Adapted from [11].

general stimulus outlines, while Meissner’s corpuscles and Merkel disks with smaller receptive fields identify precise spatial boundaries of smaller stimuli [11].

Additionally, these mechanoreceptors show varying adaptation rates to continuous stimuli. Pacinian and Meissner’s corpuscles, categorised as fast-adaptive receptors (FA or RA), exhibit reduced sensitivity to prolonged stimulation and cease transmitting information if the stimulus remains constant on the skin [10]. In contrast, Merkel and Ruffini receptors, identified as slow-adaptive receptors (SA), maintain their responsiveness to sustained pressure on the skin [10].

Mechanoreceptor distribution and density vary across the body, significantly influencing tactile sensitivity. Peripheral regions such as arms and legs have higher receptor density compared to the trunk and proximal areas, resulting in heightened sensitivity thresholds in these peripheral areas [12].

The detection of a vibrotactile stimulus is explained by the four-channel theory, which posits that each mechanoreceptor type corresponds to a psychophysical channel sensitive to distinct input frequencies. Understanding these sensitivity ranges is crucial, as detection thresholds for complex vibrotactile stimuli are determined by the spectral component with the highest energy, adjusted for human psychophysical sensitivity [13]. Most findings, except for texture perception, derive from passive touch scenarios where the hand remains stationary. In contrast, active touch, involving hand and finger movement, presents different dynamics and can yield varied results.

Understanding the mechanisms of tactile feedback is vital for the development of the FARE-KE framework. Although the framework primarily focuses on kinesthetic devices, designs

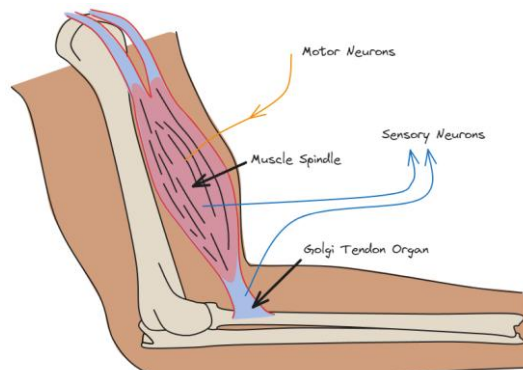


Figure 4. Illustration of proprioceptors involved in sensing force and body movement, including muscle spindles and Golgi tendon organs. Adapted from [14].

often incorporate both tactile and kinesthetic modalities, or actuators may inadvertently introduce unwanted vibrations during early prototyping, affecting user experience. Aligning FARE-KE with human tactile perception ensures that the evaluation process comprehensively measures and quantifies all facets of the haptic experience as well as ensuring researchers are aware of any potential limitations of the framework.

2.1.2 Kinesthetic Feedback

Our perception of position, movement, and muscle and joint tension is facilitated by the kinesthetic sense, or proprioception (Figure 4). This sensory system allows us to determine the spatial location of our body parts without visual input. Proprioceptors, sensory receptors that detect internal body stimuli, particularly from muscles and joints movement, play a crucial role in this process.

Proprioception involves three primary main types of proprioceptors: muscle spindles, Golgi tendon organs, and joint receptors [14]. Muscle spindles, embedded within muscle fibres, detect changes in muscle length and their rate of change, providing critical information about muscle stretch. This detection aids in regulating muscle contractions and reflex responses. Golgi tendon organs, located near muscle-tendon junctions, monitor muscle tension and relay data on muscle force, helping to regulate contractions and prevent damage from excessive force. Joint receptors, found in the joint capsule, detect changes in joint position and movement, providing data on angles, velocity, and acceleration, which is essential for coordinating joint movements and maintaining stability [16].

Slow-adapting Ruffini endings, which detect skin stretching, also contribute to proprioception. Though not as dominant as muscle spindles, Golgi tendon organs, and joint receptors, Ruffini endings help perceive body position and movement by providing information on the length and direction of skin stretch.

Developing a kinesthetic haptic controller requires understanding the human body's capabilities and operational range to accurately assess metrics associated with these phenomena. Table 1 lists parameters related to human kinesthetic perception, including sensory thresholds that define our capabilities and limitations. Considering these parameters is essential for the characterisation framework, ensuring operation remains within safe ranges. Exceeding these limits could increase injury risk from the controller, highlighting the importance of well-defined safety protocols in design and application.

Category	Parameter	Type	Value	Notes	
Kinesthesia	Position threshold	Absolute	0.6–1.1°	Elbow & shoulder position precision	
		Differential	1.7–2.7°	Smallest detectable change in position	
			2.0/0.8°	Wrist & elbow/shoulder minimum movement detection	
	Movement threshold		9 %	Finger & elbow minimum position change	
		Absolute	1°	Minimum detected movement in a specified velocity range	
		Differential	0.2–0.7°	Wrist, elbow & shoulder smallest detectable change in movement	
Information transfer		8 %	Elbow smallest detectable change in movement		
		1.92 bits	Amount of information transmittable about joint angle positions		
Sensing bandwidth		20–30 Hz	Frequency range of proprioceptive and kinesthetic sensor detection		
Sensory Thresholds	Force	Absolute	0.06 N	Smallest detectable force change	
		Differential	7 %	Detectable force change over a range of 2.5–10 N	
			7 %	Elbow muscle force change over a range of 25–400 N	
			15–27 %	Detectable force change for forces less than 0.5 N	
	Pressure	Absolute	0.019 g	Women's smallest detectable pressure change	
			0.055 g	Men's smallest detectable pressure change	
	Friction	Accurately scaled	0.43–2.79	Friction coefficient range accurately perceived	
	Shear	Accurately scaled	0.15–0.70 N	Range of shear force accurately perceived	
	Vibration	Intensity (Differential)	25 %	Detectable change in vibration intensity over a range of 10–20 dB sensation level	
	Spatial resolution		1.2 mm	Smallest detectable change in grating orientation	
			0.6 mm	Smallest detectable change in grating	
	Temporal resolution		5 ms	Minimum time interval for detecting two successive mechanical pulses	
	Body Capabilities	Force	Maximum	50/60 N	Maximum exertable force for the finger/wrist
				100 N	Maximum exertable force for the elbow & shoulder
Control			11–15 %	Control precision of force exertion	
			1 %	Detectable force change over a range of 10 to 20 N	
Bandwidth			1 %	Detectable force change over a range of 20 to 50 N for the wrist, elbow & shoulder	
			10 %	Detectable force change over a range of 5 to 18 N	
			2–3 Hz	Frequency range for detecting changes in force	
			5 Hz	As above	
			2–6 Hz	As above	
Motion		Maximum speed		17.6 rad/s	Maximum achievable speed of movement
		Bandwidth		<2 Hz	Frequency range for detecting changes in motion during active touch for sensing
				2–4 Hz voluntary movements
				2–7 Hz periodic tracking
			4–8 Hz skilled actions like hand writing, typing, tapping, playing musical instruments	
			10 Hz reflexive actions	
	<10 Hz	Frequency range for detecting changes in motion			
	8–12 Hz	Frequency range for detecting changes in motion during finger tremor			

Table 1. Aggregated Sensory Thresholds and Body Capabilities for Kinesthesia: This table is an adapted summary of various tables from Samur et al., focusing on the kinesthetic perception of limb motion and position, sensory thresholds of the human hand, and human body capabilities related to kinesthetic perception. Data sourced from various studies as detailed in Performance metric for haptic [1]




Figure 5. Summary of tactile technologies: (a) Eccentric rotating mass motor, (b) linear resonant actuator, (c) solenoid actuator, and (d) piezoelectric actuator.

2.2 Haptic technologies

Haptic technologies simulate touch and force sensations, significantly enhancing user experience by increasing intuitiveness and immersion across various applications. This section examines a range of technologies designed to generate tactile and kinesthetic feedback, highlighting their unique characteristics and specific applications. The section begins with an exploration of the most prevalent tactile technologies (Figure 5), followed by an overview of kinesthetic haptic actuators.

2.2.1 Tactile Technologies



Tactile feedback technologies stimulate skin mechanoreceptors using vibrotactile, thermal, and electro-tactile displays that evoke sensations of texture, temperature, and pressure. This subsection explores tactile feedback methods, from widely used vibrotactile approaches to novel technologies like chemical haptics. Table 2 evaluates different vibrotactile actuators that generate feedback through amplitude and frequency variations applied directly to the skin, operating across ranges up to several kilohertz.

Research indicates that human skin's sensitivity to tactile stimuli varies, covering a frequency spectrum from 0.4 to 1000 Hz [15], with frequencies beyond this range generally falling outside human tactile perception [16]. Optimal perceptual sensitivity is found within the 150 to 300 Hz range [17]. The amplitude of perceived vibrations varies significantly across different body part, emphasising the need for precise actuator performance tuning for specific application areas.

Actuator Type	Affordability and Availability	Mechanical Simplicity	Electrical Simplicity	Customizability	Expressiveness	Frequency Range	Force Output	Displacement Output	Voltage Required	Power Consumption
Eccentric Rotating Mass (ERM)	H	H	H	L	L	100-500 Hz	L	L	L	L
Linear Resonant Actuator (LRA)	M	H	M	M	H	~200 Hz	L	L	L	L
Voice Coil Actuator	L	H	L	H	H	10Hz- 2KHz	L	L	M	L
Solenoid Actuator	M	L	M	H	M	<100 Hz	L	L	L	M
Piezoelectric Actuator	L	M	L	M	H	Few Hz to several kHz	H	L	H	L
Shape Memory Alloy (SMA) Actuator	M	L	M	M	H	<10Hz	M	M	M	M
Pneumatic/ Hydraulic Actuator	L	M	H	H	L	<100Hz / <10Hz	H	H	H	H
Electrostatic Actuator	M	L	M	H	H	1-10 kHz	L	L	H	L

Table 2. Expanded and adapted table from [18] [5], [19], [20], [21]. These values are generalized; actual parameters may differ significantly depending on specific models, manufacturers, and application environments.

2.2.1.1 ERM Vibration Motors

The Eccentric Rotating Mass (ERM) motor is a widely used actuator for haptic feedback. It functions by rotating an imbalanced mass on a shaft, generating multi-axis vibrations primarily in the plane of rotation. The simplicity, ease of control, and durability of ERMs make them a cost-effective and reliable choice for various applications. However, ERMs have inherent limitations, such as the direct correlation between vibration strength and rotational speed [19], which restricts the complexity of haptic patterns they can generate. Additionally, actuation latency (the time required for the mass to accelerate or decelerate) further limits their effectiveness in delivering precise haptic feedback.

To mitigate some of these drawbacks, modern integrated circuit (IC) drivers [22] have been developed, incorporating features such as overdrive to accelerate motor startup and active braking mechanisms to decrease stopping time. While these advancements reduce some limitations, ERMs still struggle to produce complex haptic feedback patterns. Nevertheless, their affordability and reliability sustain their widespread adoption across various industries.

2.2.1.2 Vibrotactile linear actuators

Vibrotactile linear actuators offer a flexible haptic technology, distinguished by their ability to produce diverse tactile sensations. They generate vibrations along a single axis, allowing enhanced control over stimuli. Within this category, linear resonant actuators (LRAs) and voice coils (VCs) are the predominant designs. Both function on principles similar to loudspeakers, where electric current passing through a coil within a magnetic field causes movement along the field's axis, converting electrical signals into physical vibrations. In LRAs, this mechanism includes a mass-spring system which amplifies vibrations but limits the frequency range due to the resonant frequency of the mass. VCs offer a wider frequency range, enabling more expressive stimuli, albeit with reduced amplitude. Linear actuators feature quicker response times and broader frequency capabilities compared to ERMs, with initial research into this technology for haptic applications dating to the 1930s [23].

While ERM actuators operate with a simple DC voltage, linear actuators require a more complex electrical "audio" signal. This "audio" input enables LRA to generate effects that can simulate "virtual textures" on surfaces [24]. Modern haptic IC controllers have been developed to manage these actuators efficiently, incorporating pre-programmed haptic sensations to streamline their operation.

2.2.1.3 Solenoid Actuators

Solenoids are electromagnetic actuators commonly used in haptic technology to create vibrations and simulate kinesthetic feedback. These devices function by passing an electric current through a coil of wire, generating a magnetic field that causes a ferromagnetic core to retract or extend rapidly. This motion produces a strong force, enabling solenoids to generate high-amplitude, low-frequency vibrations, functioning similarly to linear resonant actuators (LRAs). Due to their significant actuation force, reliability, and simple design, solenoids are widely used in applications requiring robust kinesthetic feedback.

Despite their advantages, solenoids have inherent limitations due to the mass of the internal core moving within the coil. These include high response latency, output variability based on device orientation, and a restricted frequency range. These factors can limit the versatility of solenoids in nuanced haptic applications requiring rapid response and a broad frequency spectrum.

2.2.1.4 Piezoelectric Actuators

Piezoelectric actuators leverage the piezoelectric effect, where solid materials alter their shape upon application of electrical voltage. These actuators typically comprise a metal disk topped with a ceramic layer that responds to an electric field by expanding, causing linear displacement perpendicular to the disk. While the displacement produced is minimal, piezoelectric actuators can generate significant forces at high frequencies. This feature makes them particularly effective in creating various tactile sensations, such as simulating textures, by manipulating the frequency and amplitude of the applied electrical signal.

Due to their low power consumption and high reliability, piezoelectric actuators are well-suited for portable and wearable haptic devices. However, they operate at high voltages, often exceeding 100V, which requires specialised circuitry for voltage amplification and precise control. Furthermore, the limited displacement of these actuators may restrict the diversity of haptic sensations they can produce, posing challenges in applications requiring a broader range of tactile feedback.

2.2.1.5 Shape Memory Alloy

Shape memory alloys (SMAs) offer a novel approach in haptic technology. These materials deform in response to temperature or electrical current changes, generating force, pressure, or movement for haptic feedback. SMAs are advantageous due to their compact size, low power consumption, and ability to produce consistent motion. However, they have limitations, including slow actuation and a restricted range of motion that can limit their use in dynamic applications. Additionally, precise control of temperature or current is necessary to ensure accurate feedback, adding complexity to their integration into haptic systems.

2.2.1.6 Mid-air technologies

Mid-air technologies use phased array ultrasound transducers to generate skin vibrations via acoustic pressure without direct physical contact. This technology adjusts transducer phase shift and integrates hand-tracking technology, enabling accurate control over the vibration focal point. Such control allows simulation of various haptic patterns that mimic the sensation of touching shapes or edges across an interface. The intensity and frequency of stimuli produced by this technology depend on the ultrasound speaker wavelength [25]. For instance, a 40 kHz ultrasound speaker, with a wavelength of approximately 8.5 cm, can produce an output force of 16 mN, creating tangible haptic feedback in mid-air.

2.2.1.7 Air actuators



Air actuators, using fans or pneumatic systems, stimulate skin mechanoreceptors to create tactile haptic feedback. This feedback is experienced as pressure, discernible through skin hair movement. Devices like the Wind Tactor [26] use directed airflow to provide tactile stimulation to specific body areas such as the neck. Similarly, other systems use controlled air vortices to generate tactile sensations on the face [27]. While effective, these technologies pose implementation challenges, and their performance is easily influenced by environmental conditions.

Other designs such as Aerial delivers vortices capable of generating forces up to $9.8 \mu\text{N}$ and reach users up to 1.2 meters away. However, their travel speed introduces a latency of a few seconds before the user perceives the effect. This design demonstrated an 84% accuracy in delivering targeted tactile feedback to a user at a distance, showcasing the potential for remote haptic interaction using air actuator systems.

2.2.1.8 Chemical

Chemical haptics is an emerging field within haptic technology that enables sensory experiences beyond conventional vibrotactile or mechanical methods. This approach stimulates nociceptors to simulate temperature variations, using substances like peppermint oil for cooling and low doses of capsaicin for warmth. For instance, the Douleur haptic device [28] combines capsaicin with a micro-needling process to induce pain and improve dermal absorption. Additionally, a system by Lu et al. [29] utilises up to four chemical agents to enhance virtual reality realism. Brooks et al. [30] further investigate chemical haptics' ability to generate temperature illusions by activating trigeminal nerve receptors on the face. The potential of chemical haptics to create precise and varied tactile sensations highlights its promise as an innovative area within haptic technology.

2.2.2 Kinesthetic Feedback Technologies

As we explore kinesthetic feedback technologies, it is important to acknowledge the wide range of actuators used in this field. Unlike the previous section, which examined individual actuators, taking a broader perspective. It highlights how actuators function within kinesthetic systems to generate feedback across different modalities. Instead of classifying




Figure 6. Various modalities of handheld kinesthetic haptic devices. Direct actuation: (a) Torc [32], (b) Claw [33]. Indirect actuation, where force is applied as torque through a handle or grip: (c) Thor's hammer [34], (d) Itorqu [35] and (e) AirRacket [36].

actuators by their properties, the focus here is on their roles in delivering kinesthetic feedback. Various kinesthetic devices are discussed, each illustrating distinct feedback methods. This overview underscores the multitude different ways in which stimuli are conveyed to users, forming a foundation for later discussions on measurement techniques.

This examination of kinesthetic devices follows the classification system proposed by Adilkhanov et al. [31], with particular attention to wearability. Two primary categories of haptic devices are analysed: handheld kinesthetic devices and exogloves, selected for their relevance in portable and handheld applications.

2.2.2.1 Hand-held kinesthetic devices

Hand-held kinesthetic devices are designed to be held and manipulated by the user, as shown in Figure 6, enabling interaction with haptic content through the device itself. These devices fall into two main subtypes based on the nature of their actuation: direct and indirect.

2.2.2.1.1 Hand-held Devices - Direct Actuation

The first category examined comprises hand-held devices that exert force directly upon the user's skin, such as on the hand or fingers (Figure 7). An exemplary model of direct actuation is the Claw [32], which provides both kinesthetic and tactile feedback to the user's index finger. This device employs an actuated lever to support pinching motions between the thumb




Figure 7. Examples of handheld devices that deliver direct actuation: (a) CapstanCrunch [37], (b) NormalTouch [38], (c) Bstick [39], (d) Haptivec [40], (e) Pupop [41] and (f) pneu-multi-tools [42].

and index finger whilst delivering feedback responsive to trigger actions. Additionally, Claw integrates a voice coil actuator at the lever's extremity in contact with the index fingertip, providing tactile feedback that enhances interaction with virtual surfaces.

A similar approach is presented in CapstanCrunch [33], which uses a lever mechanism to simulate pinching actions and material stiffness. Unlike Claw, CapstanCrunch incorporates a capstan system between the motor and lever, amplifying the force from a lower-power motor and allowing for adjustable resistance. This design demonstrates the range of methods for achieving similar kinesthetic feedback in hand-held devices that use different actuators. In both cases, the devices can be characterised by the force exerted on the fingertip or by their torque generation capabilities.

Benko et al. [34] introduced an alternative hybrid approach that integrates tactile and kinesthetic feedback to represent virtual shapes. They developed two devices: NormalTouch, which employs a tiltable platform to simulate the normal surface of a virtual object for kinesthetic feedback, and TextureTouch, which features a matrix of actuated pins to recreate the texture of a virtual object. In NormalTouch, the platform adjusts its angle of contact with the fingertip, suggesting that its characterisation could be simplified to a single-axis force when the platform moves in one direction.

Bstick [35] is a kinesthetic device that engages both fingers and the hand, featuring five linear actuators capable of exerting forces up to 22 N on the contact plates. Designed with a cylindrical form, Bstick employs a multipoint actuation strategy, positioning four micro linear actuators to correspond with each finger grasping the handle. This configuration enables movement forward or backward, effectively altering the device's diameter. PaCaPa [36] is another device designed to apply forces directly to the user's hand or fingers while being held. It uses a servo motor to pivot a "wing" through a hinge mechanism, adjusting palm pressure across a 0-90° hinge angle range, thereby providing a uniaxial force.

Haptivec [37] uses a distinct method to impart impulse forces onto the user's hand with compact linear actuators, such as small solenoids, that exert pressure at the user's palm. This approach offers directional haptic pressure vectors to the user's fingers and hand during interaction with virtual environments. The device's round hand-held controller contains solenoid actuators arranged in eight principal directions, providing directional pressure feedback in an x-y coordinate plane. For devices like Bstick, characterisation clearly involves

measuring the linear force exerted by each actuator. However, with designs such as PaCaPa, the process of measuring this linear force and identifying the specific point of measurement is less straightforward.

Inflatable props offer a novel alternative to traditional mechanical kinesthetic devices, utilising shape-changing properties to simulate interactions with virtual objects. Popop [38] design employs a range of basic inflatable shapes that expand under controlled pressure when users interact with virtual objects. By combining these basic forms, more complex structures become possible. An advancement in interactivity is demonstrated by the Pneu-multi-tools[39], which incorporates folding hinges integrated with several basic shapes. This design enables the creation of rudimentary mechanisms that mimic real-world tools, such as pliers, offering passive stand-ins for these implements. While such designs can replicate stiffness sensations based on system pressure, quantifying this stiffness remains challenging. The perceived stiffness varies with the hand area contacting the device, and differences in individual users' grips significantly affect the force they experience. This variability presents a considerable challenge, common across the other designs discussed, in characterizing users' subjective force experiences during device interaction.

2.2.2.1.2 Indirect Actuation

Indirect actuation techniques use technologies like air turbines, moving components, or energy-storing elements such as elastics to provide kinesthetic feedback. These methods position the actuator at a distance from the user's hand, typically connected by a handle. The haptic feedback manifests as torque experienced upon the device's actuation.

Dynamic Passive Haptic Feedback (DPHF) is a design approach that leverages an object's natural dynamics to generate counterforce against user movements. Several designs incorporate weight-shifting or shape-altering mechanisms to modify controller attributes




Figure 8. Examples of Dynamic Passive Haptic Feedback (DPHF): (a) Shifty[46], (b) Drag:On [47], (c) Transcalibur[48], (d) ElastOscillation[49] and (e) ElaStick[50].

(Figure 8). For instance, Shifty [40] adjusts the controller’s centre of mass to modify its moment of inertia, altering the perceived weight and resistance to movement, enhancing the sensation of mass distribution. Meanwhile, Drag:on [41] utilises foldable fans to generate passive haptic feedback via air resistance, creating either resistive force or rotational effects based on fan positioning. Similarly, Transcalibur [42] builds upon Shifty’s approach by integrating a pivot mechanism that enables both weights to rotate 90 degrees, altering mass distribution. While these techniques offer the advantage of low energy consumption, their effectiveness is restricted to moments of motion. This poses a challenge for characterisation, as assessing these controllers requires active analysis of their performance during user interaction.

An alternative approach to passive haptic feedback mechanisms involves designs that incorporate elastic components for actuation. One such design is ElastOscillation [43], which features a mass suspended between two elastic actuators. A DC motor shifts this mass across a plane, modifying the tension in the elastic elements and, consequently, influencing the mass’s movement. This setup can replicate the sensation of swirling motion within a container. Similarly, ElaStick [44] adjusts the firmness of an elastic component affixed to a VR controller to produce haptic feedback. This device employs four motors to either tighten the elastic body into various configurations or release it abruptly, providing the user with the tactile sensation of interacting with virtual objects of varying physical properties.

Aerodynamic haptic devices, as illustrated in Figure 9, utilise components such as propellers, turbines, or fans to generate force feedback. Thor’s Hammer [45], for example, is a handheld virtual reality controller that houses six propellers within a cube, capable of producing up to 4N of three-dimensional force. WindCage [46] builds upon this concept, optimising it for serious gaming by extending the design into a pole equipped with two actuators, each generating up to 1.8N of force. Leviopole [47] follows a similar design approach, incorporating a pole with actuators at each end and eight propellers arranged in a quadcopter configuration.




Figure 9. Examples of aerodynamic haptic devices: (a) WindCage[51], (b) Leviopole[52], (c) Wind-blaster[53], (d) Aero-plane[54] and (e) JetController[55].

This setup allows for effects such as simulating super-leaping [48], where the device generates an upward pull when the user jumps in a virtual environment, modifying their perception of jump height. WindBlaster [49] takes a different approach by employing wrist-mounted turbines instead of propellers, providing adjustable orientation and rotational speed. This flexibility enables the device to produce unidirectional force or rotational feedback with up to 1.5N of force.

AeroPlane [50] similarly uses dual high-velocity turbines on a handheld device, enabling uni-axial force or rotational feedback through variable turbine speeds, capable of delivering up to 7N of force per turbine. JetController [51] utilises compressed air, incorporating multiple nozzles on a handheld controller to provide up to 6N of linear kinesthetic feedback. By controlling air release via electro-valves across the nozzles, the device can generate three-dimensional forces by combining the output from multiple nozzles.

These untethered designs allow force generation without physical tethering. However, they present challenges by factors such as vibrations, noise, and wind produced during actuation [45]. Typically, these devices provide linear kinesthetic feedback along a single axis, though some designs support multi-axial or more complex feedback by integrating multiple actuators or dynamically adjusting their orientation.

The final mechanism under discussion employs the flywheel approach, as shown in Figure 10. This system is particularly effective for untethered haptic controllers, offering the advantage of generating feedback without physical attachments. The flywheel operates by dynamically adjusting its spinning mass, producing a gyroscopic force perpendicular to both its tilt and rotational axis. This requires a sophisticated gimbal system to precisely control the flywheel's orientation relative to its actuation shaft, to ensure accurate and responsive kinesthetic feedback.




Figure 10. Examples of flywheel torque haptic devices: (a) Golf-club type controller[57], (b) Itorqu[37], (c) GyroVR[58] and (d) TorqueScreen[59].

Early flywheel-based prototypes appeared in the 2000s [52], [53], [54], [55], contributing to research on directional cues for navigating virtual three-dimensional environments [52]. Later designs focused on enhancing virtual reality (VR) interactions, such as the Golf Club Device [56] and handheld systems like iTorqu [57], [58]. However, these devices encountered washout time issues, where the flywheel needed to return to its original angular state before generating new feedback. This reset process produced an opposing torque, which could disrupt the intended haptic experience. To mitigate this, researchers introduced scissored-pair flywheels, where two counter-rotating flywheels cancel out unwanted reaction torques [59], [60], [61], [62]. Flywheel-based devices generate momentum stimuli across three axes, requiring force/torque (F/T) sensors capable of simultaneously measuring force and angular momentum along all three dimensions for accurate characterisation.

This review outlines the diverse techniques employed by kinesthetic haptic devices to deliver force and torque. These approaches range from directly applying force to the user's hand through lever-based mechanisms, as seen in devices like Claw and Capstan Crunch, to the dynamic reorientation of flywheels in systems such as the Golf Club Device and iTorqu. Despite their differences, these mechanisms can be categorised into broader groups for a more structured understanding.

2.2.2.2 Exogloves & Exoskeleton

As we conclude the exploration of kinesthetic haptic devices, we turn to Exohaptic Devices, which include both exogloves and exoskeletons (Figure 11). These devices provide kinesthetic feedback while allowing users to maintain hand mobility for interaction with their environment. Exohaptic devices utilise different actuation methods, categorised here into active mechanisms, locking mechanisms, and alternative actuation techniques.

Exogloves and exoskeletons with active mechanisms use motors or other force-generating components to provide kinesthetic feedback, often modulating resistive forces to simulate properties like stiffness. Locking mechanisms, in contrast, restrict movement by limiting the user's range of motion. Additionally, the literature presents alternative actuation methods, such as particle jamming, Bowden cables, pneumatic systems, twisted string actuation (TSA), and pneumatic gel muscles (PGM). Each of these methods exhibits unique actuation properties that influence the overall haptic experience and will be discussed separately to highlight their distinctive characteristics. Exoskeleton systems, including both exogloves and

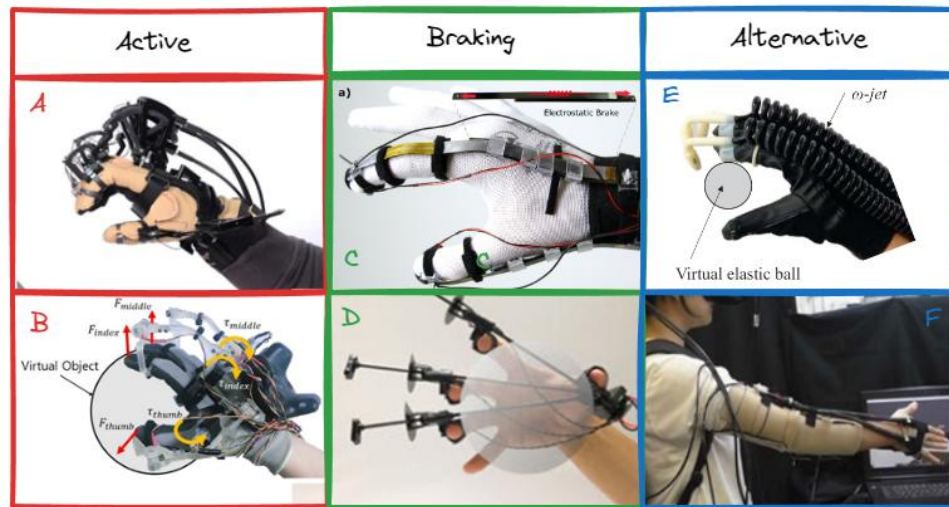


Figure 11. Examples of exoglove devices categorized by function. Active force exogloves (red): (a) CyberGrasp [69], (b) [25]. Braking-based exogloves (green): (c) Wolverine [70], (d) DextrES [71]. Alternative methods (blue), such as pneumatic actuation, offer kinesthetic feedback with limited control: (e) ExoTenGlove (TSA) [72], (f) ω -jet [73]

full-body exoskeletons, offer the advantage of hands-free operation. However, as body-grounded haptic devices, they introduce force feedback not only at the intended stimulation point but also at the body attachment points, which may affect user perception and comfort.

Commercially available glove-based haptic exoskeletons, such as Dexmo [63] and CyberGrasp [64], provide active and resistive kinesthetic feedback, enhancing user interaction within virtual environments. These devices are used across various fields, including rehabilitation, medical training, and telemanipulation. The CyberGrasp glove employs a Bowden cable system that functions as a series of tendon actuators along the back of the hand, this system simulates different textures and forces while tracking the user's hand movements and finger positions in real time for accurate virtual interaction feedback. In contrast, Dexmo is a mechanical device that utilises link structures to deliver kinesthetic feedback. Its initial design [63] featured a braking mechanism that provided resistive force only during virtual object interactions. Later commercial versions [64] have expanded its capabilities, aligning them more closely with those of CyberGrasp.

An alternative mechanical approach to haptic actuation is hydraulic actuation, as demonstrated in the Hand of Hope [65] project. This wearable haptic glove was specifically developed for hand movement rehabilitation, employing a linkage-based system powered by linear actuators to facilitate controlled finger movements. The design integrates a slim, flexible artificial muscle, a compact hydraulic module for actuation, and a micro pressure sensor that

accurately measures internal hydraulic pressure without interference or distortion. By modulating hydraulic pressure, the system provides both active and passive feedback, adapting to the user’s specific rehabilitation needs. Exogloves that deliver both active (force-generating) and resistive (force-restricting) kinesthetic feedback require a characterisation process that accounts for both modalities. When evaluating such actuators with a single-axis force sensor, it is crucial that the sensor can measure both compressive (push) and tensile (pull) forces to ensure a comprehensive and accurate assessment of the device’s performance.

Incorporating active kinesthetic feedback into exogloves presents significant challenges, particularly due to the need for multiple actuators, five per hand when actuating each finger individually. These actuators must be lightweight enough to be worn comfortably while generating sufficient force for effective feedback. Striking this balance is crucial, as the actuators must replicate the natural capabilities of the human hand, which can exert forces of up to 60N [1]. Ensuring accuracy in both design and functionality is essential for delivering realistic and effective haptic feedback.

A more straightforward approach in exogloves focuses solely on delivering resistive force. This concept is exemplified by the ExoTen-Glove [66], which employs a Twisted String Actuation (TSA) system to restrict movement in the user’s fingers. The system integrates a DC motor and encoder to regulate the movement of twisted strings, while a force sensor monitors user input, providing responsive feedback in a closed-loop system. The TSA system allows for variable braking force control by adjusting string tension, enhancing force output from otherwise low-power DC motors. Wolverine [67] utilises a latching mechanism that restricts finger movement along a sliding rod. This mechanism increases braking force as the user’s hand closes, using a mechanical latch to enable the use of lightweight motors. Another lightweight latching design achieving high braking forces is Dextres [68], a haptic glove that employs electrostatic braking for feedback during grasping activities in virtual reality environments. The mechanism functions by engaging or releasing jointless flexible metallic links attached to the fingers, leveraging electrostatic braking to constrain movement.

Alternative designs in exogloves and exoskeletons use elastic actuators to reduce reliance on electromechanical components. One such design, ElasticVR [69], features an elastic link between the user’s palm and forearm. This link can be actuated by a motor to change the length of the elastic actuator or to release it suddenly, generating resistive forces of up to

14N. Other passive approaches, such as Elastic-arm [70] and Elasticcon [71], provide egocentric haptic feedback by docking the hand to the shoulder with an elastic link used for complementing other interaction techniques. The elasticity of these components allows for a diverse range of feedback. The force-displacement characteristics of the elastic link can be assessed by measuring the force required to stretch or compress the link to a specified displacement. This measurement determines the kinesthetic feedback that the device provides at any given position.

Overall, the exoglove devices discussed primarily use single-axis actuators, providing kinesthetic feedback in a single direction. Thorough dynamic characterisation is essential, for devices employing Dynamic Passive Haptic Feedback (DPHF), as these generate feedback only in response to user movement. By contrast, systems such as flywheels and actuators with more complex force outputs require momentum measurements across multiple axes, requiring force/torque (F/T) sensors. Despite varying complexity in haptic device designs, in many instances the characterisation process can be simplified to measuring a single actuator along one axis.

2.3 Haptic Controllers Characterisation

After reviewing an extensive array of actuators and kinesthetic approaches, this thesis now examines how these devices are characterized, along with an evaluation of suggested methods for their characterisation. The core aim is to develop a robust foundation for a haptic characterisation framework that encompasses both tactile and kinesthetic technologies. This is essential because these two modalities often converge; for instance, certain actuators might inadvertently produce vibrations while providing kinesthetic feedback. On this section both modalities will be cover, beginning with a detailed review of literature related to vibrotactile actuators, including aspects such as coupling methods, sensor technologies, and data processing techniques. The examination will then shift to kinesthetic devices, exploring their implementation and testing throughout the research process, including the selection of transducers, amplifiers, analog-to-digital converters (ADCs), signal conditioning, data analysis, and evaluation methods.

This detailed review aims to support the future enhancement of frameworks for more reliable, accurate, and effective characterisation of haptic controllers. It is worth noting that the setup

Accelerometer	Measuring Range	Frequency Range (Khz)	Data Type	Cost 1	Sensitivity	Dimensions (Mm) - Package
LIS344ALH	±6 G	1.8	Analog	\$14	Vdd/15	4x4x1.5 Lga-16l
352C23 ICP	±1000 G Pk	10	Analog	\$200	5 Mv/G	4.1x8.6x2.8
LSM6DS3	±16 G	6.7	I2c 16bit	\$9.43	0.488 Mg/Lsb	2.5x3x0.86 Lga-14l
ADXL335	±3 G	1.6 (X, Y) 0.55 (Z)	Analog	\$ 9.6	300 Mv/G	4x4x1.5 Lfesp-16
ADXL345	±16g	3.2hz	Spi, I ² C 13bit	\$1.9	256lsb/G	3x5x1.5 Lga-14
ADXL356	±40g	2.4hz	Analog	\$24	80 Mv/G	6x5.6x2.05 Lcc-14

Table 3. Summary of accelerometers used on previous work to measure vibrations of vibrotactile actuators. Cost might vary depending on the package of the chip. *1 prices based on <https://www.lesc.com/> as 01/23

procedures for characterizing haptic controllers are often insufficiently detailed in the literature, with only a small portion of studies providing extensive overviews of the process. This review will focus on papers that include visual representations and images of the characterisation setup, offering key insights into sensor placement and implementation strategies that are not readily apparent from textual descriptions.

2.3.1 Tactile Feedback Characterisation



As with kinesthetic feedback, there is no fixed standard process in the literature for defining vibrotactile characterisation. Reliable modelling requires measuring amplitude, frequency, and acceleration of mechanical vibrations. This review identified two main techniques for measuring vibrations: accelerometers and non-contact vibrometers. Due to cost-effectiveness, accessibility, and common use in haptic research, accelerometers are the most frequently used method for measuring vibrations (Table 3). Accelerometers enable direct measurement of mechanical vibrations by attaching the sensor to the vibrating surface. Non-contact vibrometers, such as laser Doppler vibrometers, measure vibrations without physical contact by detecting changes in the object’s surface position or displacement. However, these devices are specialised, costly, and less accessible. The subsequent sections will detail the characterisation process based on research papers, examining accelerometer selection and characteristics, followed by coupling, data acquisition, and processing strategies.

2.3.1.1 Accelerometers

Commercial accelerometers, or Inertial Motion Units (IMUs), employ microelectromechanical systems (MEMS) technology, integrating electrical and mechanical components on a single chip to detect acceleration. The prevalent method in these devices involves measuring capacitance changes caused by an internal mass movement, which converts into an electrical signal. When characterising haptic devices and selecting an accelerometer, several sensor properties must be considered:

- **Measuring Range:** The accelerometer's measuring range is critical for characterising devices that produce large magnitude vibrations, as it determines the maximum acceleration the sensor can capture.
- **Frequency Range:** This parameter determines the vibration frequencies an accelerometer can detect. For devices generating vibrations across a broad spectrum, choosing a sensor that covers these frequencies is essential. As noted in Section 2.2.1, human vibration perception is limited to frequencies below 1 kHz, so a sensor should at least cover this range.
- **Data Output:** Commercial accelerometers provide either analog or digital output. Digital accelerometers have resolution limited by the Analog-to-Digital Converter (ADC), with data update rates dependent on the digital communication protocol and ADC capabilities. These aspects matter when integrating with microcontrollers or data acquisition systems (DAQs).
- **Sensitivity:** Sensitivity is the minimum detectable change in acceleration, enabling capture of haptic stimuli nuances, particularly in devices with small actuators. In analog accelerometers, sensitivity is typically measured as output voltage per unit of acceleration (mV/g), affecting measurement precision.
- **Dimensions:** The accelerometer's physical size and properties are vital, especially for space-constrained applications. The sensor's dimensions and weight can influence the haptic device's performance; heavier accelerometers might mechanically dampen the actuator. Selecting an appropriately sized and weighted accelerometer ensures accurate, reliable measurements without compromising actuator functionality.

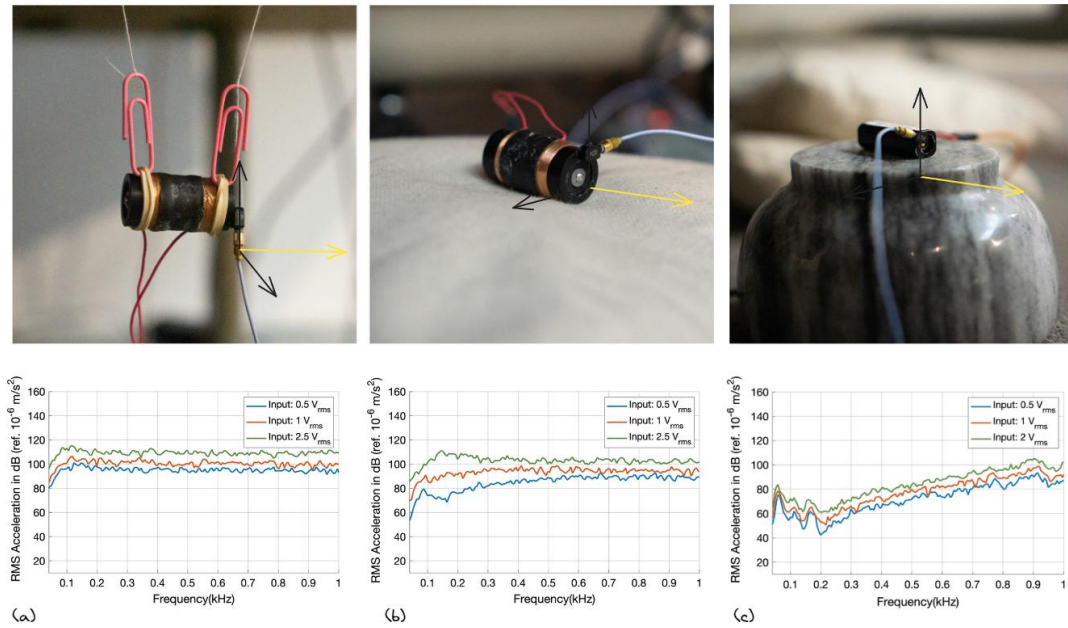


Figure 12. Haptuator mk-II's magnitude-frequency response under different mechanical loadings: (a) freely suspended, (b) sandbag placement, and (c) stone base attachment.

2.3.1.1.1 Characterisation Using Accelerometers: Mounting and Coupling

The initial step in measuring vibrations within a haptic interface involves mounting or coupling the accelerometer to the haptic device. Bukkapatnam et al. [72] investigated various methods for characterising voice coil actuators using accelerometers, aiming to develop an affordable and portable vibrotactile toolkit for digital musical instruments. The same principle of operation also applies to the study of vibrotactile interfaces as explored by Fortin et al. [73]. The authors conducted various tests, including magnitude-frequency response and total harmonic distortion, under different conditions such as freely suspended, placed on a sandbag, or attached to a rigid surface (Figure 12). These tests assessed how different loading conditions influence actuator performance.

A lightweight accelerometer (0.2 g) was attached directly to the actuator, aligned with its main motion axis. Data was captured using a PCB Piezotronics 352C23 ICP accelerometer and processed through a National Instruments NI USB-4431 DAQ system, with signals analysed in MATLAB. The analysis revealed that loading conditions altered the actuator's magnitude-frequency characteristics and introduced damping effects. The total harmonic distortion (THD) levels varied significantly under different loading conditions, sometimes

exceeding 100%, indicating increased non-linear behaviour. This presents a challenge for vibrotactile interfaces designed to deliver purely sinusoidal stimuli, particularly in psychophysical experiments requiring precision. Such non-linearities can distort intended stimuli, affecting experimental result reliability and user experience quality in scenarios requiring high-fidelity sensory feedback.




Figure 13. Touch-Box device designed to measure normal forces applied to its top panel, enabling precise analysis of vibrotactile feedback delivery.

The study offers a detailed evaluation of measuring vibrations in haptic actuators, covering loading condition effects, and coupling and equalisation methods on voice coil actuators' frequency response. While actuator orientation during testing is not directly examined, it is recognised as a factor influencing both actuator performance and haptic perception. Although the research uses sophisticated hardware for data capture, the availability of cost-effective alternatives is also noted.

The coupling of the vibrator affects not only its performance but also the perception of vibration magnitude as shown by Papetti et al. [6]. Touch-Box (Figure 13) investigates the dynamics between user-applied force and effects on perceived vibrotactile feedback, offering insights for hapticians studying human sensitivity to vibrations under varying pressure levels.

The Touch-Box setup includes a CZL635 load cell for measuring force, a Tactile Labs Haptuator Mark II voice coil actuator to generate tactile feedback, and a Wilcoxon 736T piezoelectric accelerometer for vibration measurement. An Arduino UNO controls the setup, with an RME Fireface 800 audio interface capturing vibratory signals. Users interacted with the vibration plate; by analysing the relationship between vibration amplitude and finger-applied force, researchers found that perceived vibration intensity increases with applied force. This finding highlights the influence of actuator coupling in haptic device characterisation, significantly impacting the user's tactile experience.

Gaudeni et al. [74] demonstrated how to account for coupling effects during actuator characterisation by adding a 100g weight to simulate the force experienced when a haptic interface is used. This method established the relationship between input haptic signal and




Figure 14. Haptic ring and its characterization [82]. (a) Vibrotactile device worn on the proximal phalanx of the index finger. (b) 3D-printed housing with an opening for the actuator's moving platform, ensuring direct skin contact and secured with Velcro straps. (c) Experimental setup measuring actuator displacement under load using an accelerometer. (d) Displacement amplitude of the loaded actuator as a function of driving frequency and input voltage, peaking around 160-190 Hz.

output response for a haptic ring prototype. An Inertial Measurement Unit (IMU) sensor (LSM6DS3 from STMicroelectronics) was positioned under the weight and attached to the actuator's moving platform (Figure 14). The accelerometer's output was captured using a Teensy 3.2 Development Board and analysed to assess average peak acceleration induced by different input signals.

In the study, the authors presented displacement amplitude in millimetres rather than traditional units such as "g" or acceleration. They actuated the voice coil with sinusoidal signals and calculated displacement amplitude using the formula where acceleration amplitude is divided by the square of the product of 2π and the frequency (f) of the sinusoidal signal. This method emphasises the need to consider load and coupling variations when evaluating an actuator's displacement. Such factors can notably affect the frequency at which peak displacement occurs, underlining their relevance in accurate measurement of actuator capabilities. Accelerometers typically report in G-forces or acceleration, directly indicating vibration intensity across multiple axes. Alternative measurements, such as displacement in millimetres, are sometimes used, but decibels might offer a more suitable representation due to the logarithmic nature of vibration intensity perception, especially when comparing relative intensities.




Figure 15. Experimental device featuring an accelerometer mounted on a stylus (left). User holding the prototype during the user study (right).

Alternatively, the characterisation process can incorporate the user as part of the evaluation, where a study with multiple participants interacting with the device captures live performance data as shown by Gourishetti et al. [75]. Their work investigated the use of low-cost DC motors for creating high-fidelity vibrations on rigid tools such as a stylus.

To evaluate the rotating motor actuator (RMA) performance, a comparative study used a reference actuator (Haptuator Mark II) and an RMA as shown in Figure 15. Participants adjusted a slider to match the vibration magnitude of the test actuator to that of the reference actuator. During practice sessions, the experimenter verified that participants could perceive the vibration magnitude. The acceleration generated by the actuators was measured using a STMicroelectronics LIS344ALH MEMS-based accelerometer rigidly attached to a stylus. For accurate measurements, acceleration data was digitally high-pass filtered at 10 Hz using a second-order Butterworth filter to remove gravity and low-frequency user-induced motion influences. A National Instruments DAQ (USB-9925) captured the data, generated command signals to power the prototype through two current amplifiers, and measured accelerometer data.

Expanding upon practical evaluation methods for vibrotactile devices, "Vitaki: A Vibrotactile Prototyping Toolkit for Virtual Reality and Video Games" [76] presents an authoring tool designed for characterising Eccentric Rotating Mass (ERM) actuators. This toolkit enables exploration of various feedback techniques, including overdrive and active braking (Figure 16), to enhance vibrotactile feedback effectiveness.

The methodology employed by the researchers involves using a Samsung L760 disk-type actuator paired with an ADLX335 analog accelerometer to measure performance. Measurements captured via an oscilloscope provide detailed data on the actuator's




Figure 16. Vitaki vibrotactile actuator response: standard operation (left) compared to overdrive and braking techniques (right).

performance metrics like start and stop times, frequency, and amplitude for each pulse. This data is essential for optimising the actuator's parameters, particularly when implementing voltage adjustments for effective overshoot during startup and efficient active braking. This approach identified the most efficient voltage for overshoot startup and active braking techniques used to drive the ERM, demonstrating significant reductions in startup time and braking to less than 20ms, compared to 100ms without active braking.

The author also explains how to calculate the theoretical output force of an ERM actuator, which can be compared with accelerometer data to gauge actuator efficiency. This is possible when the actuator properties are available; otherwise, determining the eccentric weight mass requires destructive, irreversible disassembly of the actuator.

$$F_{vibration} = m \cdot r \cdot \omega^2$$

where :

m - the mass of the eccentric weight

r - mass offset distance

ω - speed of the motor ($rads^{-1}$), $\omega = 2\pi f$

The simplicity of Vitaki's approach (directly reading the analog output from the accelerometer using a digital oscilloscope) demonstrates an accessible method for measuring vibrations. This technique requires minimal technical skills, making it valuable for quickly assessing actuator behaviour and adjusting parameters for optimal performance in real-time.

The method used to process multi-axial data significantly impacts the accuracy of its representation. Currently, there is no established standard for processing such data. Landin et al. [77] investigated various techniques for converting 3-axis accelerometer readings into a single time-varying signal, aiming to preserve the original vibrations' relevant characteristics.

The authors examined several methods, including single-axis projection (SA321), sum of components (SoC321), vector magnitude (Mag321), and principal component analysis (PCA321). SA321 uses only one axis of the 3D signal, discarding potentially valuable information. SoC321 adds the three components together, which can cause destructive interference due to correlations between axes. Mag321 calculates the 3D acceleration vector's




Figure 17. Comparison of four transformations (SA321, SoC321, Mag321, DFT321) applied to a three-dimensional texture vibration signal. Left: one-dimensional signals in the time domain. Right: smoothed energy spectral density (ESD) of transformed signals (colored) and the original 3D signal (black). The DFT321 method achieves the best spectral ($M_{sm}_{sm} = 1.00$) and temporal match ($M_{tm}_{tm} = 0.64$), outperforming the others.

magnitude, but introduces non-linear redistribution of spectral content. PCA321 projects the 3D signal onto its principal variation axis, but may not capture all relevant information.

To address these limitations, Landin et al. proposed using a Discrete Fourier Transform (DFT321) as an alternative approach. The DFT321 method provides an orthogonal basis for component summation, effectively preserving the signal’s spectral integrity while minimising destructive interference. The authors compared the DFT321 method with other techniques using two metrics: spectral match (M_{sm}) and temporal match (M_{tm}).

The results showed that the DFT321 method provided a better fit to the input signal compared to other approaches. For stationary texture vibrations, it achieved a perfect spectral match and a high temporal match, outperforming other transformations (Figure 17). This highlights the DFT321 approach’s effectiveness in reducing 3D vibrations’ dimensionality while preserving their perceptually relevant characteristics.

While accelerometers have proven versatile and effective for characterising vibrotactile feedback in haptic interfaces, they have several limitations. The requirement for direct contact with the vibrating surface can be impractical and may alter the actuator’s output. External factors such as temperature fluctuations and electromagnetic interference can affect measurement accuracy. For high-frequency or subtle vibrations, non-contact methods like Laser Doppler Vibrometry offer an alternative approach, which will be examined in the following section.

2.3.1.2 Laser Vibrometers

A Laser Doppler Vibrometer (LDV) is a non-contact measurement technique used to measure the vibration, velocity, or displacement of a target surface. The process involves directing a




Figure 18. Laser Doppler vibrometry (LDV) applications for haptic measurements: (a) CHASM device for vibrotactile feedback, (b) testbed for characterizing CHASM vibrations, (c) LDV setup measuring palm vibrations induced by haptic actuators [82] and (d) LDV analysis of mid-air haptic stimuli for remote tactile feedback [81].

laser beam onto the target surface and analysing the frequency shift in the reflected beam caused by the vibrating target. From this shift, vibration velocity or displacement can be calculated. LDV has been used to measure haptic interface vibrations, as it is accurate, non-interfering, and capable of measuring displacements as small as 100 microns. Several studies have reported using LDV to validate haptic interface performance and provide more precise vibration assessment. For example, Preechayasomboon et al. [78] used LDV to characterise the Chasm (Figure 18, a & b), demonstrating LDV application in this context. The authors employed a Polytec PSV-500 LDV to measure surface displacement and compared the results with readings from a built-in encoder. LDVs are particularly suited for situations requiring a contactless approach, as they avoid the dampening effect caused by attaching an accelerometer. This non-intrusive method enables precise vibration measurement on surfaces, such as assessing mid-air haptics stimuli [79] (Figure 18 d) or in cases where contact sensor placement would distort readings [80] (Figure 18 d).

Currently, only a limited number of LDV manufacturers exist, with products primarily aimed at precision industries. This introduces a key limitation: high initial cost. Although companies like Polytec offer rental options, the expense remains substantial compared to accelerometers. Additional considerations include the need for both sensor and target to remain static during measurements, as readings are sensitive to the distance between them. Unstable mounting can introduce significant noise. For example, when measuring a participant’s hand, movement must be constrained, as shown in Figure 18 c, which was part of the user study conducted by Dandu et al. [80].

Haptic feedback often combines vibrotactile and kinesthetic elements. This interaction plays a critical role in the overall perception and effectiveness of the feedback. For instance, the Touchbox [6] study showed that force exerted during a vibrotactile stimulus directly influences perceived intensity. In other cases, vibrations are unwanted byproducts of high-

speed actuators, and characterisation helps identify side effects and thresholds. As the focus shifts from vibrotactile to kinesthetic feedback characterisation, understanding measurement approaches will enhance comprehension of the complex dynamics between force and vibration in haptic interactions.

2.3.2 Kinesthetic Characterisation



This section describes the process for characterizing kinesthetic controllers. This characterisation depends on the specific device design, as these controllers feature various physical attributes and operational modalities highlighted in the previous literature review. For example, a haptic exoskeleton for the hand differs in design parameters such as size and torque generation from a handheld controller, requiring different sensors and measurement approaches. The section begins with a detailed review of previous characterisation setups organised by sensor type, providing an overview of methods to measure and quantify kinesthetic feedback. The discussion then examines the variables commonly used to define a kinesthetic controller’s properties, identifying a subset of variables descriptive of the main device characteristics based on previous literature.

2.3.2.1 Characterisation Setups

Characterisation setups from past studies are discussed based on the sensor’s complexity and the number of measurement axes. This systematic approach helps improve understanding of various experimental configurations. The section concludes by synthesizing common features among the setups, highlighting key insights and trends from the literature.




Figure 19. Force-sensing and haptic feedback device. (a) Standard force-sensing resistor (FSR). (b) The "Touch-and-Fold" prototype, a compact fingertip-mounted haptic device. (c) Internal components of "Touch-and-Fold," including the force sensor, Linear Resonant Actuator (LRA), and mechanical transmission system.

2.3.2.1.1 FSR Force Sensing Resistors

Force Sensing Resistors (FSRs) provide an affordable method for measuring single-axis force. These variable resistors comprise two conductive layers separated by an air gap, with resistance changing proportionally to applied pressure. Their compact size (as thin as 0.1 mm) and adaptability suit haptic characterisation setups requiring limited space and non-intrusive placement between object surfaces and user contact points, as demonstrated by Touch&Fold [81] (Figure 19).

Touch&Fold is a nail-mounted foldable haptic device for mixed reality (MR) environments, using the sensor for closed-loop control. The design provides tactile feedback by pressing against the user's fingertip when a virtual object is touched. The control loop incorporates a Force Sensing Resistor (FSR 400, Interlink Electronics) and a photo interrupter (SG-105F, Kodenshi) as limit switches. Here, the force sensor serves as a closed-loop feedback signal, enabling precise tuning of finger pad pressure and measuring normal force on the sensor's surface.

FSRs offer several advantages for force measurement in haptic feedback systems, including cost-effectiveness, compact size, and ease of use. Their high sensitivity enables accurate measurement of small resistance changes. However, FSRs exhibit limitations that may affect accurate force measurement, which must be considered. Some of these limitations include:

- **Linearity:** FSRs exhibit non-linear resistance output, affecting calibration ease and increasing output scaling uncertainty. However, using a trans-impedance amplifier can improve the sensor's output linearity [82]. This can be further enhanced through polynomial equation calibration [83], [84].
- **Drift stability:** Output changes over time when under strain lead to stability issues. Researchers have examined this phenomenon [85], [86] and some have attempted to model the creep effect as a solution. However, it remains unclear if this approach can be generalised to all signals [86].
- **Hysteresis:** Output variability during load application and removal can reduce force measurement accuracy. Complex mathematical models incorporating moving integrals can counteract hysteresis effects [84], [87].
- **Homogeneity:** The method of applying force on the sensor surface significantly affects output readings, making sensor surface characterisation challenging. To

address this, researchers have used durable coating shells [83] to cover the sensing area, applying uniform actuation and improving accuracy.

- **Repeatability:** This refers to response uniformity when tested under identical conditions multiple times. Resistive pressure sensors often present repeatability challenges [88], [89].

The literature review reveals that FlexiForce A201/A401 [90] and Interlink Electronics FSR 40X [91] are the two most commonly used force sensors in previous literature. Numerous studies have aimed at improving force measurement accuracy for both sensors and addressing the limitations noted above [82], [83], [88], [92], [93]. Hollinger et al. presented a comprehensive comparison of drift and hysteresis values for these sensors [94], with response differences attributed to their manufacturing processes, providing a starting point for FSR selection. The FlexiForce uses a Traditional Sandwich Element (TSE) design, while Interlink sensors employ Non-Aligned Electrode Elements to avoid lumped capacitance [95], which can cause force measurement inaccuracies.

Paredes-Madrid et al. [95] proposed a technique for improving FSR force sensing reliability by driving the FSR with an oscillating frequency matching the sensor's capacitance and resistance. This approach significantly improves force measurement accuracy and reduces hysteresis and drift by factors of 3 and 9, respectively. However, this technique applies only to FlexiForce designs, as it relies on opposite signs of capacitance and resistance to correct drift characteristics.

Despite FSR limitations, their compact size and affordability have made them popular for force sensing in the haptics community. These sensors have measured force on soft actuators [96], [97], [98] and sensed interaction between actuators and user skin [81], [97], [99], [100], [101], [102]. However, the literature shows limited references to the characterisation process or compensation for sensor limitations, suggesting these limitations are either overlooked or measurements are taken as guidance.

The techniques presented by Paredes-Madrid et al. [95] to improve FSR performance require advance electronic knowledge, potentially limiting adoption within the haptic community. This may explain why some researchers choose more reliable sensors and techniques, such as load cells or force gauges, when precise force measurements are critical to their research.




Figure 20. HapticLinks prototype for measuring displacement and applied force: (a) layer hinge design, (b) prototype mounted on a lathe with a force gauge, and (c) torque–angle curves illustrating mechanical response characteristics.

2.3.2.1.2 Force Gauges

Force gauges are instruments designed to measure and quantify force or weight. Their design typically relies on strain gauge technology, which converts sensor deformation upon force application into a quantifiable electrical signal. Simple, durable, and portable, these devices provide real-time force measurement in an easy-to-use standalone solution. Lower-cost models may lack the ability to store data or high sampling rates, potentially making them unsuitable. In contrast, high-end force gauges, frequently used in research, feature RS232 communication and computer interfaces.

A practical application was demonstrated in Haptic Links [103], where researchers presented a haptic device capable of rendering variable stiffness between two virtual reality controllers coupled using a mechanical link (Figure 20). They used a force gauge to compare three prototypes: Chain, Layer-Hinge, and Ratchet-Hinge. By securing one end of the haptic device to the lathe spindle and the other to the force gauge fixed to the carriage, researchers could precisely move the carriage while taking sensor readings, effectively evaluating each design and allowing for comparable results between the various Haptic Links designs. Though the specific force gauge model is not mentioned in the paper, visual inspection suggests it was similar to a Salter Brecknell ES-22 ElectroSamson [104].

The size of force gauges can be a disadvantage when characterizing haptic devices. However, in quasi-static or static characterisation setups where the haptic device is mounted on a table or measuring jig, sensor size may not be significant, as it can be incorporated into the setup design. This is evident in the characterisation of CapstanCrunch [33], a palm-grounded haptic controller using a friction-based capstan mechanism known as the Twisted String Actuator (TSA). Controlled by a small internal motor, the TSA provides haptic feedback for




Figure 21. CapstanCrunch prototype for direct actuation haptic feedback. (Left) Assembled device. (Right) Characterization setup for performance measurements.

virtual reality environments. For CapstanCrunch, force measurements were acquired using a force gauge (Figure 21 - right), recording the applied force at 10 kHz sampling rate, while a computer-controlled geared stepper motor moved the finger rest towards the thumb position. The authors also used a high-speed camera to measure the haptic device's actuation latency.

Beyond reporting force dynamic range, the authors used the force gauge to calculate the device's stiffness capabilities. Stiffness quantifies the resistance force encountered during interaction and is characterised by the relationship between applied force and resulting actuator displacement. For haptic controllers designed to simulate object properties, stiffness measurement is essential as it enables meaningful comparison between similar systems. For instance, CLAW [32], can deliver a stiffness of 5.73 N/mm, while CapstanCrunch delivers a maximum stiffness of 5.88 N/mm. Load cells are also less affected by temperature changes compared to other sensors such as FSR and can be designed to be very small, making them suitable for use in small-scale haptic devices.

2.3.2.1.3 Load Cell

A load cell is a sensor that measures force or weight acting upon it, calculating these measurements from the deformation of its solid body (flexure, Figure 22 - left). Typically,




Figure 22. Generic force sensor design: (a) cross-section of a load cell with flexure, strain gauge placement, and mounting point, (b) detailed strain gauge construction with a resistive path on polymer film, and (c) Wheatstone bridge circuit configuration for stressed and unstressed strain gauges.

load cells utilise strain gauges mounted to a flexure made of aluminium or steel. When a force is applied at the measuring end, it deforms the strain gauges bonded to the flexure's body. This deformation alters the electrical resistance of the strain gauges' resistive path ((Figure 22 – middle).

These changes in resistance are then translated into an electrical signal with the aid of a Wheatstone bridge circuit (Figure 22 – right). The bridge operates on the principle that when the ratio of resistances is equal, the output voltage remains at zero. However, when deformation alters the resistance of the strain gauges, the bridge becomes unbalanced, creating a voltage difference at the output. This change in output voltage is directly proportional to the resistance variation, allowing precise measurement of force or deformation.

The resulting signal, typically in the millivolt range, is then amplified and digitised to calculate the magnitude of the applied force. The Wheatstone bridge circuit enables the indirect measurement of current changes by precisely monitoring resistance changes. This approach is especially useful when dealing with small resistance changes, which produce correspondingly small current changes that would be difficult to measure directly.

Load cells are available in several configurations (Figure 23), including beam type or cantilever designs, widely used for accurate measurement of both small and large objects. Button/cylinder types suit limited space applications, measuring compression forces and inline force measurements along a single axis. The S-beam variant is engineered for inline tension or compression measurements, making it appropriate for characterisations where force must be measured in direct alignment with the actuator. Force/Torque (F/T) sensors integrate a more sophisticated arrangement of strain gauges, enabling this specialised load cell to measure both forces and torques in three dimensions. These sensors are particularly suitable for handheld kinesthetic haptic devices where user-experienced stimuli can be directly measured



Figure 23. Various types of load cells: (a) beam load cell, typically used for precision weighing; (b) bottom load cell, designed for compression force measurements; (c) S-type load cell for both tension and compression; (d) 3-axis load cell for multi-directional force measurement; and (e) force-and-torque load cell measuring both rotational and linear loads.




Figure 24. Evaluation of the HairTouch prototype: (a) the prototype in use, (b) testbed setup for device characterization, and (c) plot illustrating the force versus hair-pressed distance relationship.

at the handle; however, their complex mechanical structure significantly increases both cost and implementation complexity.

2.3.2.1.4 Beam Load Cell

Beam load cells feature a simple rectangular body design that measures force applied at one end relative to the grounded opposite end. Their manufacturing simplicity makes them the most affordable load cell option, with hobbyist versions available for under £10, including ADC HX711 amplifier shields compatible with microcontroller boards. Arduino libraries provide calibration and data capture solutions, making beam load cells accessible and versatile for force measurement applications. This section examines previous literature on various characterisation modalities based on haptic device force output complexity.

One Axis Linear Force

HairTouch [105], is a haptic handheld device that uses two brushes/pins to interact with user's thumb (Figure 24). By adjusting the length and bending direction of brush hairs mounted on the pins, the controller modulates hair elasticity and tip direction, hence rendering various stiffness and roughness levels at each pin. To measure the resistive force of the hairs on the brush and quantify the haptic feedback, the authors mounted a beam load cell (TAL220) to an extruded aluminium frame, opposite to a single mounted brush pin which pushes linearly onto the load cell. The signal output from the beam load cell was amplified with an HX711 amplifier, and a microcontroller was used to read the digital value. The resistive force and hair-pressed distance relationship was measured at different hair lengths by repeatedly moving the actuator with the pin to various distances from the force sensor. In this characterisation setup, the actuator moves in a single axis orthogonal to the measuring surface from the beam load cell, simplifying the device characterisation process and providing consistent measurements for the various brush hair lengths.




Figure 25. ImpactVest haptic prototype: (a) worn by a user, (b) detail of device mechanics, (c) characterization setup, and (d) measured impact force over time, demonstrating performance during testing.

One Axis Impact, elastic actuators

The TAL220 load cell has also been used to measure impact forces on a single axis. ImpactVest [106] (Figure 25) demonstrates a haptic vest that renders impacts across a user’s chest using elastic actuators to generate the impact. The vest features nine impactors in a 3×3 matrix, generating interactions using stored energy from elastic actuators. To characterise the actuator, the authors mounted it onto an aluminium frame and captured the impact generated by releasing the elastic band, causing a rubber ball to collide with the beam cell. The readings were collected for approximately three seconds, measuring the impact force. Although HX711 amplifiers are commonly used in haptic characterisation, their maximum sampling rate is limited to 80 SPS (samples per second), with a default configuration of only 10 SPS. This limitation makes them potentially unsuitable for applications requiring higher sampling rates, such as impact force measurements. Several other haptic characterisation setups for elastic actuators have been presented in studies[69], [107], [108], [109], providing additional visual details and methodological descriptions.

Force Sensing for Wind Actuators

In addition to elastic interfaces, wind actuators such as propellers or turbines have been used as approaches for generating untethered pulling or pushing forces in haptic applications.




Figure 26. PropelWalker haptic device: (a) prototype worn, (b) characterization setup for the single ducted fan, and (c) graph illustrating the relationship between the fan’s PWM cycle and output force

PropelWalker [110], Figure 26, exemplifies this approach with a calf-worn haptic device employing ducted fans to simulate buoyancy and resistance forces associated with different fluids and materials in virtual environments. Although the design features four ducted fans (two per calf), the authors' characterisation setup isolated a single ducted fan to measure generated forces per actuator. The setup resembles a weight scale, with the ducted fan mounted atop and applying even force to the load cell positioned on the base frame platform. Using a CZL601 load cell capable of measuring forces up to 40 kg, the system connected to an HX711 amplifier with data transmitted to an Arduino UNO microcontroller and displayed on an LCD1602 screen. The authors controlled the haptic device's force generation using Pulse Width Modulation (PWM) through an experimental C# program that sent the microcontroller input signals with 0% to 100% duty cycle in 5% increments to the electronic speed controller (ESC) which controlled the turbine, a process similar to a ramp response analysis. The force output was stabilized for 5 seconds, and 21 measurements were recorded from the sensor as the PWM was incremented and decremented. This setup was also employed to measure the acceleration and deceleration time, or step response, of the device.

Sensing Integration for Close Loop Control

Impedance haptic interfaces require closed-loop control, where the force applied by the user is constantly monitored. This principle of operation was demonstrated by CLAW [32] Figure 27, which utilises a beam load cell as part of the device structure. By incorporating a load cell into the index finger exoskeleton actuation mechanism, CLAW measures the force exerted by the user's index finger, enabling haptic feedback for various virtual reality (VR) manipulation tasks, such as grasping, touching and triggering. As an admittance display, the controller senses force as input and renders position as output. It employs a 333 Hz PD control loop for force estimation using a beam load cell with a 0-5 kg range. The amplifier, an HX711




Figure 27. CLAW device: (a) in use, illustrating how it is held, (b) detail highlighting the integration of a beam load cell for real-time control, and (c) force versus angle measurements.

ADC communicating with a Teensy 3.2 microcontroller, provides a 24-bit signal with high resolution (theoretical resolution step of 0.000023 N). To address the 80 SPS limitation of the HX711, the authors implemented a slope extrapolator to calculate force value changes at a higher sampling rate (333Hz). Their characterisation reported outputs in both force (N) and stiffness (N/mm), with the prototype generating up to 30 N of grasping force and stiffness up to 10 N/degree (5.73 N/mm).

Characterizing Oscillation with Single Axis Sensors

Other ways to use beam load cells include characterizing oscillation kinesthetic interfaces, a common feature of devices that use elastic actuators. OsciHead [111] is an example of such a system. Other ways to use beam load cells include characterising oscillation kinesthetic interfaces, a common feature of devices using elastic actuators. OsciHead [111] exemplifies such a system (Figure 28) featuring elastic actuators mounted on either side of a head-mounted display (HMD). This design simulates both inertia and impact feedback by adjusting elastic band tension and release. The authors employed a characterisation setup to investigate relationships between extension distance for energy storage/impact. The characterisation setup utilised two TAL220 beam load cells mounted in parallel on an extruded aluminium frame. The researchers positioned these load cells at the same distance from each other as in the actual oscillator design, ensuring the actuation forces aligned properly with the sensors. Using HX711 ADC amplifiers, they measured the forces generated by the OsciHead oscillator proxy and analyzed the stored force magnitude. This analysis produced two distinct linear regression lines - one representing inertia and another for impact. As the direction of the force is the same on both load cells the readings from both sensors can be combined to calculate the total force that a user would experience during actual use of the device.




Figure 28. OsciHead device: (a) Upward force application along Z-axis with vertically oriented actuators; (b) Rotational torque generation through asymmetric actuation along positive and negative Z-axes; (c) Characterization setup using two beam load cells to measure force output; (d) Force-extension distance linear regression showing relationship between elastic band extension and measured force.

Measuring Force on 3 Axis

Measuring the force generated by multi-axis kinesthetic haptic devices presents greater challenges, particularly for handheld devices where actuators positioned at a distance from the hand cause users to perceive torque rather than direct force. However, when a device employs multiple actuators aligned along perpendicular axes, each providing forces normal to their alignment axis, multiple load cells can effectively measure each axis in a quasi-static setup. Thor's Hammer [45] demonstrates this approach as a handheld haptic device using propeller propulsion to generate ungrounded, 3-DOF kinesthetic feedback. The device incorporates six motors with propellers generating strong air thrusts, capable of applying up to 4 N of force in arbitrary directions. For evaluation, researchers constructed an aluminium extrusion measurement frame equipped with six calibrated CZL635 load cells, each measuring up to 5 kg of force. The prototype was suspended within this frame using steel wires that aligned each motor to a corresponding load cell Figure 30. These load cells connected to HX711 amplifiers, with data collected via an Arduino Uno sampling at 80 SPS and transmitting to a PC. Since Thor's Hammer generates arbitrary forces in any direction, the authors utilised their characterisation setup to assess output accuracy when combining various force vectors by converting force readings from multiple axes into force vectors with direction and magnitude for analysis.

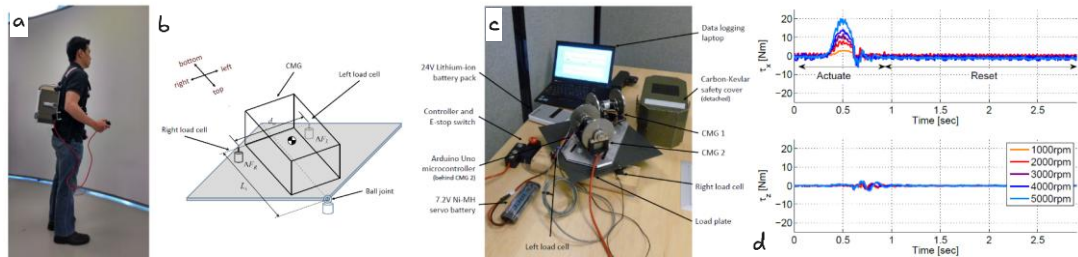


Figure 29. CMG Haptic Backpack: (a) User wearing the scissored-pair CMG prototype in backpack configuration; (b) Schematic diagram of torque measurement setup showing metal plate anchored by ball joint and attached to load cells; (c) Testing setup with labeled components including load cells, controller, and data logging laptop; (d) Graphs showing measured torques at various flywheel velocities, with sagittal torques outpacing transverse ones due to the device's gimbal motions.

Although the characterisation approach in this study effectively analyses actuator performance and interaction, it does not directly measure the torque felt by users, as it fails to account for frame rigidity or force transfer to the handle. More accurate torque measurements would require a force/torque sensor mounted on the handle, but these sensors typically cost $\sim 100\times$ more than the presented setup and may be inaccessible to many researchers. Additionally, given Thor's Hammer's high-speed propellers, integrating an

accelerometer into the characterisation setup would be beneficial, as device-generated vibrations can significantly affect the perceived haptic experience. The authors acknowledge these limitations and indicate that future studies will aim to incorporate such sensors to improve characterisation accuracy and reliability.

Measuring Torque

Characterising controllers that generate forces in multiple directions by combining actuator forces presents challenges when using single beam load cells, especially with complex actuators like flywheels that produce intricate force/torque outputs. Flywheels and gyroscopes represent some of the most complex force actuators in kinesthetic haptic devices. While force/torque sensors capable of measuring forces and torques along three axes (providing 6 metrics over time) are the most common characterisation method, beam load cells can be suitable in specific scenarios. Chiu et al. [59] demonstrated this on the characterisation of a wearable scissored-pair control moment gyroscope (SP-CMG) for human balance assistance. Their device applies torque to the human trunk's sagittal plane via a backpack to aid postural balance restoration. The SP-CMG's design isolates output torque to a single axis, eliminating off-axis torque effects typical in flywheel operation, thus allowing measurement with a single-axis beam load cell (Figure 29). To measure the generated torque, the authors mounted the device on a plate with a ball joint at one end and two load cells on opposite sides of the other end. This configuration allowed the device to pivot over the ball joint while the load cells measured the output torque. Using this setup, the researchers assessed the device's performance by measuring output torque at various flywheel RPMs.




Figure 30. Thor's Hammer device: (a) Device held by a user; (b) Device mechanics with labeled components; (c) Characterization setup with load cells and mounting frame; (d) Scatter plot comparing intended versus measured force on the X-Z plane; circles represent target force vectors, with red vectors indicating deviation from expected values.




Figure 31. Pull-Ups Kinesthetic Feedback System: (a) User performing pull-ups assisted by a suspension kit with integrated pneumatic actuators; (b) Close-up of a load cell in line with a pneumatic muscle; (c) Characterization of the device by loading weights; (d) Graphs showing displacement and actuation time across different loads, used for system calibration and feedback control.

2.3.2.1.5 Button & S-Type Load Cells

Beam load cells are widely used, but their cantilever design may not be suitable for certain characterisations, such as when measuring in-line forces (e.g., along a wire). In such cases, alternative designs like button or S-type load cells may be more appropriate, as they are specifically designed to fit in line with tension or compression forces while offering similar advantages. However, these alternative designs typically involve more complex manufacturing processes, resulting in higher costs compared to beam load cells.

Ye et al. presented Pull-Ups [112], a ceiling suspension kit using pneumatic artificial muscle groups to enable users to move up to 15 cm by pulling them while they hold the mechanism (Figure 31). In their study, the researchers employed two load sensors (ZLBM-102) to measure real-time user-exerted force. These load cells were installed between the pneumatic artificial muscle group's up-hook and ceiling hook to measure force in each hand and total force. The authors used this data for closed-loop operation and initial calibration, mapping the relationship between force and displacement generated by the pneumatic muscle under various loads. While the authors did not detail the analog-to-digital converter (ADC), they included an image showing a microcontroller and an amplifier with two HX711 modules.

S-Type Load Cells

Button and S-type load cells offer a compact form factor, making them well-suited for integration into haptic designs. For example, the Futek LSB200 has dimensions of 17.5x19 mm. Chiaradia et al. [113] presented a cable-driven soft wrist exosuit designed for flexion assistance and utilised the Futek to characterize its operation (Figure 32). The force sensor is installed along the cable actuator, measuring cable tension to estimate the weight of the load lifted by the actuator.

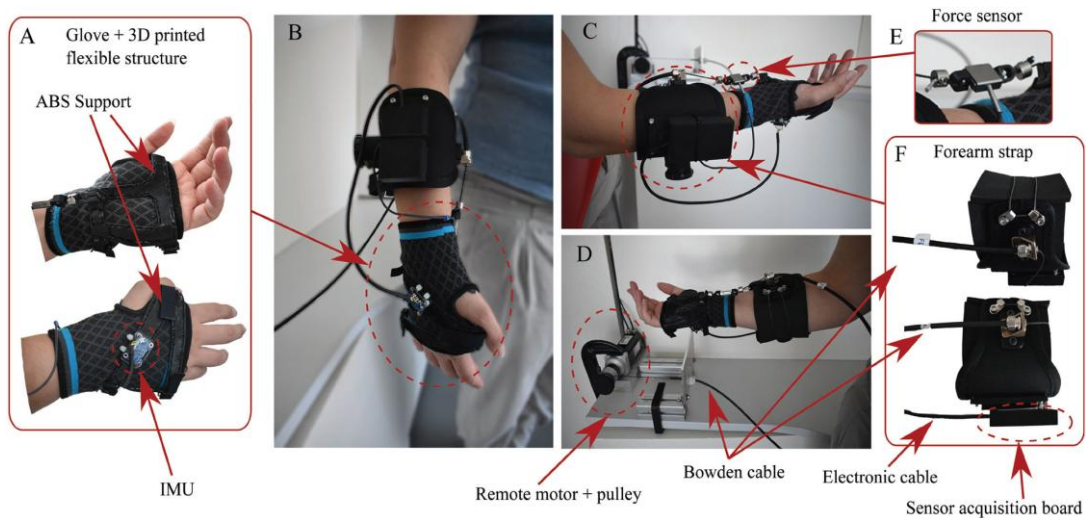


Figure 32. Components of the Wrist Soft Exosuit: (a) Glove with attached 3D-printed flexible structure and Inertial Measurement Unit (IMU); (b) User wearing the complete exosuit; (c) Side view showing the wrist configuration; (d) Actuation mechanism with remote motor and pulley system; (e) Close-up of the force sensor used to detect interaction forces; (f) Forearm strap components including Bowden cable, electronic cable, and sensor acquisition board.

To conduct a preliminary evaluation of their proposed glove design, the authors compared the stretching of a conventional fabric glove with their reinforced design, analysing glove stretching against cable tension (measured in Newtons). The authors used the sensor to quantify torque around the wrist joint midpoint between the forearm, where the actuator is mounted, and the glove/hand, where the cable is anchored. The exosuit is controlled by an admittance controller, which requires real-time force sensing to calculate output torque. Forces measured by the load cell are transformed into joint torque estimates via the force Jacobian matrix, providing the necessary feedback for the admittance control strategy. Signals from the motor encoder angle, IMU flexion angle and speed, as well as load cell force, are low-pass filtered at a frequency of 10 Hz to eliminate high-frequency noise or vibrations that could compromise force measurement accuracy.




Figure 33. Shape Memory Alloys actuator characterization: (a) SMA actuator designed using the ANISMA toolkit applied to skin; (b) Testing setup with SMAs anchored to load cell for force measurement; (c) Graph showing time-dependent pull force of SMAs at various actuation currents, demonstrating response characteristics and actuation speed.

Messerschmidt et al. presented an alternative use case for a Futek LSB250, employing the load cell to characterise a wearable Shape Memory Alloys (SMAs) pull mechanism designed in their prototyping toolkit ANISMA [114]. The authors used the sensor to record force measurements with varying actuation current values to determine SMA actuation speed for their proposed designs. These tests were conducted using 30 coils, starting from the maximum recommended elongation distance (Figure 33), with the pulling force digitised via an NI DAQ USB-6218 interface. The SMAs were securely connected to the load cell on one end and a solid mount on the other. Using this setup, the authors conducted over 40 test measurements with different SMA sizes to compare their properties.

Both previous publications demonstrate use cases for S-type load cells with active actuators. However, passive actuators may require a different approach where a mechanism consistently increases pulling force. These setups, known as pull tests, are computer-controlled and include a load cell and grips to securely hold the tested component. The load cell measures the force applied as the component is pulled or compressed. Dextres [68] exemplifies this approach, an exoglove utilising an electrostatic clutch that generates holding forces up to 20 N on each finger by modulating electrostatic attraction between flexible elastic metal strips. The authors employed an Instron 3344L pull tester with a 50 N S-type load cell (Instron 2519) to measure the braking force of the electrostatic (ES) brake (Figure 34). As a passive brake, the actuator allowed the pull tester to create displacement at constant speed along the brake axis, measuring braking force over a 10 mm displacement range. This method enabled force measurement at various voltages (a variable controlling friction force) facilitating precise characterisation applicable to haptic devices with passive actuators.

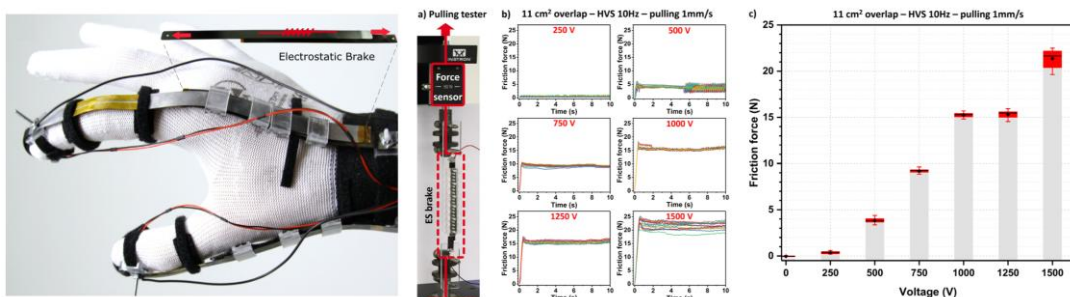


Figure 34. DextrES Haptic Glove characterization: (a) Glove prototype utilizing electrostatic brakes for kinesthetic feedback and piezoelectric elements for tactile sensation; (b) Pulling tester with force sensor for characterizing actuator performance, with graphs showing dynamic friction forces recorded at varying voltages; (c) Bar chart showing friction force correlation with applied voltage, demonstrating the proportional relationship between voltage and braking force.

Button and S-shaped load cells offer advantages when measuring forces aligned with actuators, as shown in [112], [113]. These sensors can be extremely compact, making them ideal for integration. However, they typically cost more than beam cells. For single-axis characterisation, as with ANISMA or Dextres, a beam load cell with comparable accuracy can serve as a cost-effective alternative without compromising characterisation quality.

All sensor types presented so far measure force along only a single axis, limiting their ability to accurately capture forces applied at angles or in multiple directions simultaneously. When the force perceived by users contains multiple axial components, more complex designs such as 3-axis load cells or force/torque (F/T) sensors are recommended to provide a more complete picture of the haptic experience.

2.3.2.1.6 3-Axis Load Cell Sensors

3-axis load cells can measure forces on x, y, and z components simultaneously. This sensor type has been utilised to characterise haptic actuators like Aero-plane [50], a handheld device using wind turbines to simulate centre of mass shifting in two degrees-of-freedom (Figure 35). By modulating the speed of two jet propellers, the device generates dynamic force feedback, creating the illusion of a weighted object moving on a virtual plane.

To characterise the turbine actuator, the authors mounted it on an aluminium frame similar to the previously discussed Thor’s hammer cage. A VARIENSE-FSE103 force sensor with a 250 Hz sampling rate was attached to the cage bottom and connected to the actuator via a steel wire to capture propeller-exerted force. To accurately recreate the upward force generated by the haptic controller, the actuator was characterised using an upward force mounted to mimic the real use case.




Figure 35. Aero-plane Haptic Device characterization: (a) Aero-plane controller held at its centre-of-mass, with air jets producing a downward force; (b) Measurement setup with the propeller actuator mounted on an aluminum frame and force sensor attached at the bottom; (c) Close-up of the 3D printed air duct and force sensor configuration; (d) Step response graph showing force measurements with raw and filtered data, demonstrating the device's dynamic force feedback capabilities.




Figure 36. SCALE Force Interaction Framework: (a) 3-axis load cell integrated with peripheral circuits for force measurement; (b) Schematic representation of touch interaction, illustrating the dynamic nature of human contact and the intersection points used to determine touch location; (c) Augmented workbench setup showcasing real-time object status tracking through touch interactions.

The motor speed was controlled using an Electronic Speed Controller (ESC) with a minimum speed rating of 1.25 ms pulse and a maximum pulse width of 2 ms, covering a range from 0-100% speed (ramp analysis). To understand the relationship between motor speed and resulting force, the authors collected data at 28 actuation speed steps. Despite using a 3-axis load cell, they characterised operation using readings from a single axis. This study demonstrates how characterisation can be used to conduct step response analysis for evaluating actuator performance metrics, including rise time, settling time, and maximum/minimum settling force. These metrics were calculated using MATLAB's `stepinfo` function with the captured data. To reduce noise and improve data quality, a low-pass filter with a half-power frequency of 5 Hz (Butterworth zero phase filtering) was applied prior to processing.

Although the use of 3-axis load cells in characterisation applications is limited, alternative force sensing approaches have been presented that could be applied to future characterisation frameworks. One notable example is the "SCALE: Enhancing Force-based Interaction by Processing Load Data from Load Sensitive Modules" [115] (Figure 36). In this publication, the authors present a framework incorporating load-sensitive modules (equipped with a 3 axis load cell) for force-based interaction, which can recognise a wide range of touch interactions based on force direction.

In the SCALE study, a 3-axis load cell (FNZ100N from Forsentek.inc) with a load capacity of 10 kg and three ADC HX711 operating at 80 Hz, one per axis, was employed. This setup enabled the measurement of multi-point touch and real-time computation of force directions. The authors provided details on various force-measuring implementations, including 2D touch detection, 3D touch detection and touch point classification. The load cell used in SCALE is




Figure 37. Virtual Super-Leaping Haptic Feedback System: (a) User engaged in a VR environment with the haptic device designed to simulate extreme jumping; (b) Testbed setup for characterization with labeled 6-axis force sensor and sound level meter; (c) Close-up of the measurement fixture showing coordinate axes and reference measurements; (d) Graph depicting the device's force response over time for different input pulse widths, illustrating the system's step response and latency characteristics.

made of a single aluminium body, but its design emulates a combination of three beam load cells oriented along different directions to measure forces in three distinct axes. This feature will be further discussed in Section 5.5.2, where this design serves as a principle for the proposed 3-axis design manufactured using rapid prototyping tools. However, these load cells are not suitable for measuring torques as the design is not optimised for such operation. When torque measurement is required, researchers use force/torque sensors, which are introduced in the next section.

2.3.2.1.7 F/T Load Cells

Torque sensors play a key role in characterising haptic devices that generate output torques with complex dynamics or require measurements at a distance from the actuator for more realistic assessment. These industrial sensors typically form part of a software suite that includes calibration, compensation for crosstalk, and a dedicated data acquisition system, ensuring reliable and accurate data capture.

Single axis

The most straightforward approach to using an F/T sensor is to characterise an actuator on a single axis. Sasaki et al. [48] demonstrate this approach in their characterisation of the "Virtual Super-Leaping" system (Figure 37), an interaction technique combining visual and haptic feedback to simulate extreme jumping. The haptic device consists of a pole with propeller units generating kinesthetic feedback. To characterise it, the device was mounted on an extruded aluminium frame, positioned 400mm above the ground to minimise ground effect. A 6-axis force sensor (Leptirino FFS055YA501U6) was attached to the frame's centre to measure the device's output force. Sensor readings were monitored through USB




Figure 38. Pneumatic haptic system: (a) Main components including accumulator, valve, and user coupling device that functions without an external pressure source by adjusting stiffness through valve operation; (b) Characterization setup with labeled components including motor, actuator, force sensor, valve, linear bearing, lead screw, pressure sensors, and capacitor, designed to measure pressure changes and output force.

communication, while system noise was measured using a sound level meter (FUSO SD-7000A) placed at ground level. The authors reported the device's capability to deliver 11.7N of force and 86.7dB of audio noise with an input of 1500ms (50% propeller speed) and 27.8N of force and 95.3dB of audio noise with an input of 2.0ms (100% speed).

Although a beam or button load cell could potentially measure the force output of a device like Virtual Super-Leaping, using a Force/Torque (F/T) sensor offers advantages beyond simple force measurement. The F/T sensor enables identification of unbalanced actuation between both sides of the device or undesired torques produced by manufacturing imperfections, although the authors did not specify this in their publication.

Measuring of Pneumatic System Applied Force

F/T sensors often have a small footprint, making them ideal for integration into designs or testbeds, as demonstrated in the characterisation of "Cutting the Cord" [116] (Figure 38), a pneumatic device employing the user's motion to pump fluid in a closed pneumatic circuit. The authors used the F/T sensor to measure single-axis force and evaluated performance by inflating a fabric pneumatic actuator coupled with the sensor mounted on a linear bearing system. An ATI Nano 25 F/T sensor, mounted behind the stationary end plate, measured the forces applied by the device. The characterisation setup ran on a desktop computer with control and sensors managed through a National Instruments (NI) DAQ RIO. Additionally, the authors measured system pressure by inserting a T-coupling with a pressure sensor. This data determined the relationship between parallel plate displacement and actuator volume, as well as measuring the device's refill time.



Figure 39. Wearable haptic devices and characterization methods: (a) Column showing PneuSleeve device worn on forearm with labeled components, its actuation mechanism with lead screw and linear part, and the complete prototype; (b) Column displaying Bellowband wrist-worn device, its force measurement setup, and internal component view with embedded sensor; (c) Column presenting Tasbi wrist-worn research prototype, its force characterization setup using ATI Nano 17 sensor, and the device components.

Compression haptic interfaces for the wrist or forearm, such as PneuSleeve [117], Bellowband [118], and Tasbi [119], (Figure 39) were characterised using a similar approach. These devices provide haptic feedback through mechanical compression of the wrist or forearm. In all these cases, researchers used an F/T sensor to measure single-axis force across a radial solid and analyse the relationship between the actuation mechanism and the output compression force. These setups commonly employed the popular ATI Nano 17 F/T sensor from ATI Industrial Automation, which has a diameter of 17mm, making it well-suited for integration into small haptic devices.

F/T Measuring of Gyroscope devices

F/T sensors excel in addressing the complexities of multi-axis torque/force actuators, such as flywheels, which produce intricate output torque. Flywheel haptic devices have been prevalent in the haptic community as they provide torque without grounding elements. Multiple prototypes have explored this effect over the years [55], [59], [120], [121], [122].




Figure 41. Haptic golf club feedback device: (a) Proposed device in use, showcasing the direction of generated torque and the measurement point on a golf club; (b) Characterization setup with the device mounted on a six-axis force sensor with labeled axis of measurement; (c) Graph presenting a comparison between expected theoretical values and measured torque values across different motor speeds (100-140 rpm).

Recent publications have characterised these controllers using F/T sensors. For example, Nakamura et al. [56] developed a golf club-type device with a variable-angle flywheel, offering ungrounded kinesthetic feedback to modify the user’s posture during a golf swing (Figure 41). The authors used a 6-axis force sensor (WEF-6 A 200-4 RCD, WACOH-TECH Inc., Japan) to assess the haptic device’s performance at a 240 Hz frequency under five different driving conditions. Comparisons between theoretical output and real-time measurements highlighted discrepancies, prompting the authors to provide hypotheses explaining these differences. This demonstrates the value of F/T sensors in identifying areas for design improvement in haptic controllers. Gyroscope haptic controllers featuring multiple rotating gimbals and flywheel actuators produce intuitive output torques. In such cases, accurate characterisation becomes essential for verifying operation. For example, iTorqU 2.1 [57] presented a haptic controller with a single flywheel actuated by two gimbals, enabling the flywheel to pivot over two axes simultaneously. An ATI Mini40 F/T sensor assessed the device’s performance at the handle position as a ground point (Figure 40). This configuration measured the F_x , F_y , F_z , M_x , M_y , and M_z that the user would perceive. The authors’ comprehensive analysis and mathematical modelling allowed comparison with sensor readings and verification of the haptic modelling and performance.




Figure 40. iTorqU haptic controller evaluation: (a) iTorqU controller operated by a user; (b) Characterization setup with the controller mounted on an ATI Mini40 sensor, clamped to a work surface; (c) Response comparison graph highlighting the correlation between predicted and measured torque values along the Z-axis over a specific actuation period.

F/T sensors offer distinct advantages, including industrial-grade quality, high accuracy, and multi-axis force and torque measurements. However, their cost can restrict accessibility for labs with limited funding or independent researchers. F/T sensor design remains an active research area, with numerous open-source publications aiming to provide low-cost alternatives to this technology.

2.3.2.1.8 Open Source & Rapid Prototyping F/T sensors

Extensive research has focused on designing alternative force sensors using various technologies, including traditional strain gauge designs [86], [123], [124], piezoelectric sensors [125], [126], [127], capacitance sensors [128], [129], [130], and optical sensors [131], [132], [133]. This analysis concentrates on alternative and open-source designs applicable to kinesthetic haptic device characterisation.

Solutions employing rapid fabrication techniques like 3D printing and laser cutting have also emerged. Kesner et al. [134] provided design guidelines for 3D-printed force sensors, highlighting challenges such as high hysteresis observed with soft thermoplastics. Hendrik et al. [133] demonstrated low-cost, multi-axis force sensors produced below \$40 by using optical sensors instead of strain gauges to address plastic flexure deficiencies. The inherent elasticity in 3D printed components makes strain gauge technologies less effective in combination with these materials, as these require linear deformation and suit only small changes. Consequently, 3D printed force sensor designs typically employ optical sensors [135], [136], [137] to measure deformation, better compensating for high hysteresis and slow response time.




Figure 42. Mantis haptic device: (a) User interaction with a multi-armed Mantis haptic device in VR environment; (b) Detailed view of the Mantis components including labeled force sensor, controller, and brushless motors; (c) Close-up of the custom-built acrylic force sensor.

A custom-built F/T sensor was designed and integrated into Mantis design [138] (Figure 42), a haptic device requiring accurate force and torque measurements for its admittance control system (close loop). This sensor was manufactured using common makerspace tools such as laser cutters, with overall cost below £50, excluding amplifiers. Inspired by commercial load cells, the design features a 3-axis force sensor constructed from laser-cut acrylic. Each sensor arm incorporates four strain gauges arranged in a Wheatstone bridge configuration connected to a differential amplifier, enabling determination of forces in XYZ directions. The prototype used 8mm laser-cut acrylic with rapid-set epoxy to mount strain gauges. Temperature drift was mitigated through a 3D printed shroud regulating airflow. The signal was amplified and digitalised using a MCP6N16 instrumentation amplifier's and a 16-bit differential ADC channel on an ARM Cortex M4 microcontroller. Once calibrated, the sensor measures a maximum range of $\pm 3\text{kg}$ at a 10kHz sampling rate.

The approach taken by Mantis [138] offers an accessible entry-level method for designing custom in-built force sensors, as acrylic can be easily shaped using laser cutters. However, acrylic sheets have a tensile strength of 10-12 ksi, significantly lower than aluminium's 45 ksi. While acrylic outperforms 3D printed sensors, it retains some limitations common to plastic materials. This approach requires precise strain gauge mounting techniques, as minor misalignments in orientation or coupling between gauge and flexure surface can significantly impact sensor accuracy, linearity and overall performance, factors requiring careful consideration during implementation.

An alternative approach uses the Alps Alpine HSFPAR Force Sensors Family which can measure forces up to 8N with a compact $2 \times 1.6\text{mm}$ form factor, providing cost-effective small-




Figure 43. Multi-axis force sensor assembly: (a) Dimensions of the Alps Alpine HSFPAR003A force sensor ($2.0\text{mm} \times 1.6\text{mm}$); (b) PCB with four HSFPAR003A load cells mounted, necessary for multi-axis force measurement; (c) Exploded schematic view detailing the sensor components including jaw attachment, sensing plate, load cell arrays, wire clamp, and strain relief bracket; (d) Assembled prototype showing the complete sensing device.

scale sensing solutions. Chua et al. [139] utilised these sensors in their 3-degree-of-freedom force sensor for Robot-assisted Minimally Invasive Surgery (RMIS), available as open-source and adaptable for applications beyond RMIS. Their design incorporates eight miniature button load cells mounted on a sandwich PCB, combining commercial load cells with PCB fabrication to create a multi-axis F/T sensor (Figure 43). Each sensor has a range of ± 5 N axially and ± 3 N laterally, with average root mean square errors below 0.15 N in all directions. Its compact design (9.5×8.5×23.8mm, 3.33 g) suits integration in RMIS, robotic joints or haptic devices.

The force sensor employs two standard 2-layer FR4 PCBs, each containing four HSFPAR003A button load cells, with a high-stiffness 303 stainless steel sensing plate between them. Each load cell measures compression forces up to 8 N using a piezoresistive full Wheatstone bridge, ensuring good temperature stability. Bridge outputs are amplified using AD623 instrumentation amplifiers with a gain of 21.

Having reviewed previous work on actuator types, their hardware characterisation, and commonly employed sensors, section 2.3.3 will shift focus to the data, what to measure. Previous literature reveals characterisation often concentrates on measuring actuator output during operation, providing insight into range, capabilities and performance. The upcoming section analyses this data, exploring variables that most effectively describe operational aspects of haptic controllers.

2.3.3 Haptic controller metrics



This section of the literature review examines variables describing kinesthetic haptic device properties and predominant analysis approaches. We have surveyed publications on haptic devices and works specifically addressing characterisation and metrics for haptic technologies [3] [4] [1], [2], [5], [21], [140], [141], [142]. Initially, this section provides an overview of physical properties and features measurable without force sensors. While this data enables basic device comparisons and provides informative insights, it offers limited context regarding actual haptic controller performance (Figure 44).

To address this limitation, two common analytical approaches applicable to kinesthetic devices are explored. First, quasi-static assessment focuses on test bench performance measurements, analysing the haptic controller’s response as a ramp signal incrementally increases controller output. This approach reveals critical actuator performance parameters, including maximum and minimum force capabilities. Similarly, step response analysis investigates the haptic controller’s behaviour under sudden input changes, providing essential insights into dynamic behaviour and responsiveness.

	Wireality	VR Haptic Drones [26]	CLAW [10]	Normal Touch [5]	ElasticVR[51]	Wolverine [8]	DextrES [25]	PuPoP [50]	TORC [33]	Haptic Revolver [53]	Thor's Hammer [24]	Haptic Link [46]
Sensing Points	whole hand	whole hand	thumb + index	one finger	whole hand	thumb, index + middle fingers	thumb + index	whole hand	thumb, index + middle fingers	one finger	arm	arm
Actuator Type	ratchet gear with solenoid pawl + strings	drones	servo motor + voice coil actuator	servo motor	rotation motor + elastic band	one-way brake	electrostatic brake + piezoelectric actuator	pneumatic	vibrotactile motor + voice coil actuator	DC motor + wheel	motor + propeller	linear actuator + ball-and-socket
Haptic feedback	collision	collision	collision + grip + texture	collision + compliance	impact + compliance	grip	grip + compliance	grip + compliance	texture + compliance	texture	tug	bimanual grip
Max. Force per Finger	183N	n/a	30N	n/a	14N	106N	20N	n/a	n/a	3.35N	4N	80N
Peak Power Consumption	2.19W	20.40W	5W	n/a	n/a	2.89W	<0.12W	24W	n/a	2.50W	204.70W	7.20W
Weight on the Hand	11g	0g	420g	150g	150g	55g	16g	n/a	n/a	237g	692g	>651g
Cost	\$35	\$30	>\$150	n/a	n/a	\$40	n/a	n/a	<\$90	n/a	n/a	n/a

Figure 44. Table from Wireality [147] comparing key features (e.g., actuator type and feedback mechanism) across various haptic devices. While it omits detailed dynamic performance metrics (such as those revealed by step-response testing) it provides a straightforward overview of each device’s primary attributes.

2.3.3.1 Physical Properties

The most readily accessible properties describing the features of a kinesthetic haptic controller are based on its fundamental operation and physical properties. The list below includes variables commonly cited in previous research:

- **Workspace:** This refers to the physical space within which the device operates, enabling interactions with virtual or remote environments. The workspace is reported in meters as a measure of the dimensions of the workspace ($X * Y * Z$). Fully untethered handheld devices benefit from not being constrain to a workspace.
- **Degrees of Freedom (DOFs):** The DOFs indicate the device’s haptic output directional capabilities. This affects the range of motion and flexibility, allowing complex interactions across multiple planes.

- **Electrical properties:** Power consumption and voltage are crucial features, often detailed alongside power supply types when the device is battery-powered (e.g., 12.6V LiPo 3S battery).
- **Dimensions:** Expressed as (X * Y * Z) in meters. An outside bounding box can be used to determine the total dimensions.
- **Weight:** Measured in grams, the weight affects portability and ease of use. Lighter devices are preferable for wearable applications to minimise user fatigue and allow longer interaction periods.
- **Operational noise:** The noise level generated during operation can affect the device's perception. Noise levels are typically measured using a sound level meter, especially important for devices like propeller-actuated controllers known for higher noise outputs, it is typically reported as an absolute level in decibels. [45], [50]
- **Vibration:** Critical for haptic devices, especially those using high-speed actuators like turbines or flywheels. System vibrations can be measured using an accelerometer mounted on the end effector. While many kinesthetic devices are early prototypes or conceptual, refining the design and manufacturing processes can mitigate vibrations and noise.

These properties provide the foundation for describing physical attributes of haptic controllers. The following section explores advanced analysis techniques requiring force sensors and test beds, including ramp analysis and step response. These approaches build upon previously presented characterisations, offering deeper insights into haptic device performance and behaviour.

2.3.3.2 Ramp Analysis

Ramp analysis is a widely implemented technique to evaluate haptic devices. In this approach, the haptic controller is mounted on a frame or testbed and connected to a force sensor. A ramp input signal, which can be voltage, current, pressure, or time pulses (ESC), is gradually increased from minimum to maximum actuation (Figure 45 while the corresponding output force is measured. For optimal results, tests should be conducted both incrementally and decrementally.




Figure 45. Ramp analysis plot showing the relationship between actuator input and output force, highlighting hysteresis patterns (blue and red paths) with annotated characteristics including output force resolution, static friction breakaway point, and actuation steps in 10% increments.

From this analysis, several variables and metrics can be derived that describe the device performance characteristics. The input/output trajectory, commonly referred to as the calibration curve, delineates the relationship between input signal and output force (e.g., kPa/N, V/N, A/N). This relationship defines metrics including:

- **Maximum Force/Torque:** Incremental ramp testing identifies the relationship between force and actuation percentage, enabling model creation for calculating haptic controller output. This approach provides a safe method for testing maximum force output, as potential damage or overheating can be detected before maximum actuation power.
- **Minimum Force/Torque (static friction breakaway):** Most mechanical and electromechanical devices exhibit a minimum threshold before movement occurs, appearing as an output force control with a zero offset. In the case of Aero-plane [50] the motors are controlled using an ESC with a typical PWM range of 1ms to 2ms. However, in their control system, they use a range of 1.25ms to 2ms, as the rotor response was significantly slower when operating below 1.25ms.
- **Hysteresis:** This parameter quantifies the difference between device response to increasing versus decreasing input. Measured by loading and unloading the device under identical conditions, hysteresis represents the maximum deviation between increasing and decreasing output curves. The units correspond to the haptic device's output values, typically force units (Newtons) when output is expressed in force. The formula for calculating hysteresis error is:

$$h_{\max} = \max(v_{\text{dec}} - v_{\text{inc}})$$

Here, v_{dec} represents haptic output during the decreasing phase of the input signal, while v_{inc} denotes the output during the increasing phase. The term h_{\max} captures the largest difference between phases, providing a measure of the hysteresis error.

- **Sensitivity:** This parameter represents output force or torque change per unit input change. For Aero-plane [50], this denotes output Newtons per 0.1ms increment. During ramp response testing, sensitivity is calculated by dividing output force change by input signal change. Common units include N/V, representing force change (Newtons) per voltage change (volts), varying with the control system employed.
- **Output Force Resolution:** The smallest detectable change in output force or torque. Resolution accuracy depends on both D/A resolution and force resolution, determined by incrementing input in small steps while recording output force.
- **Dynamic Range:** The ratio between maximum and minimum producible values, often expressed in decibels (dB), a logarithmic scale for large ratio comparisons. The calculation formula is:

$$\text{Dynamic Range (dB)} = 20 \times \log_{10} \left(\frac{\text{Maximum Force}}{\text{Minimum Force}} \right)$$

- **Stiffness:** Defined as the resistance to deformation, stiffness is quantified by the ratio of the change in force to the change in displacement ($\Delta F/\Delta x$). It is determined by applying a range of input forces to the device and measuring the corresponding displacements, either in linear units (e.g., millimetres) or angular units (e.g., degrees) using a goniometer. Stiffness is a critical parameter for haptic feedback because it enables the simulation of contact with solid virtual objects by accurately reproducing the sensation of a rigid surface. This property not only underpins the realistic feel of virtual materials but also contributes to maintaining a stable grip on remote objects. Stiffness is typically reported in units of N/mm or N/deg, and from a psychophysical perspective, objects with a stiffness above 25 N/mm [143] are perceived as rigid.

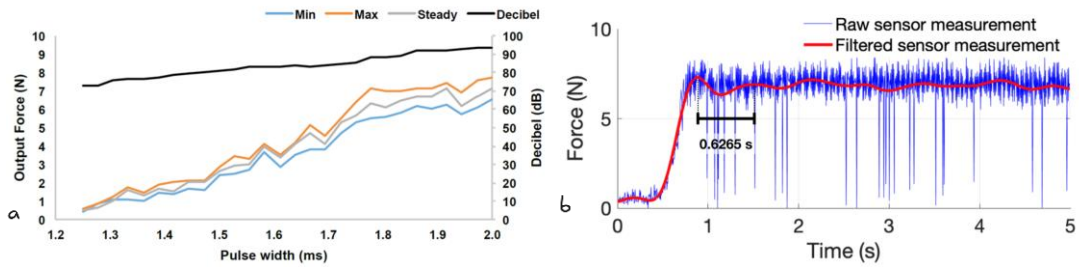


Figure 46. Aero-plane characterization graphs: (a) Output force and noise level vs. pulse width driving the propeller as it increases from 1.2ms to 2.0ms, showing minimum, maximum, steady-state forces and decibel readings; (b) Raw sensor measurement and filtered sensor measurement of the jet propeller's step response over time, with filtered curve removing high-frequency noise and showing 0.6265s response time.

In the Aero-plane paper referenced throughout this section, ramp analysis characterised static performance metrics, identifying potential issues in device performance and stability while evaluating the controller's response to constant inputs and actuator output linearity. This approach also established maximum peak force capabilities before conducting step/impulse analysis, procedures that impose greater stress on the controller with potential for permanent damage.

Following static analysis completion, step response analysis provides deeper insights into haptic controller operation and revealed additional properties. The Aero-plane authors employed this approach to develop a transfer function model of the actuator, a mathematical representation describing the dynamic input-output relationship. Step response analysis provides the controller's dynamic behaviour and stability when rendering sudden output changes. The resulting transfer function enabled designers to predict actuator responses to arbitrary input signals, an essential capability for developing stable, responsive control systems for haptic devices.

2.3.3.3 Step Response Analysis

Step response analysis identifies several key performance metrics, including Peak Force, Max Continuous Force, Rise Time, Settling Time and Output Error. These variables summarise the controller's capabilities, allowing evaluation of speed, stability, responsiveness and accuracy. This analysis assesses as well as device performance its adaptation to rapidly changing input conditions.

The step response analysis examines system behaviour over time when subjected to a sudden input change, typically characterised by transition from one constant state to another, often

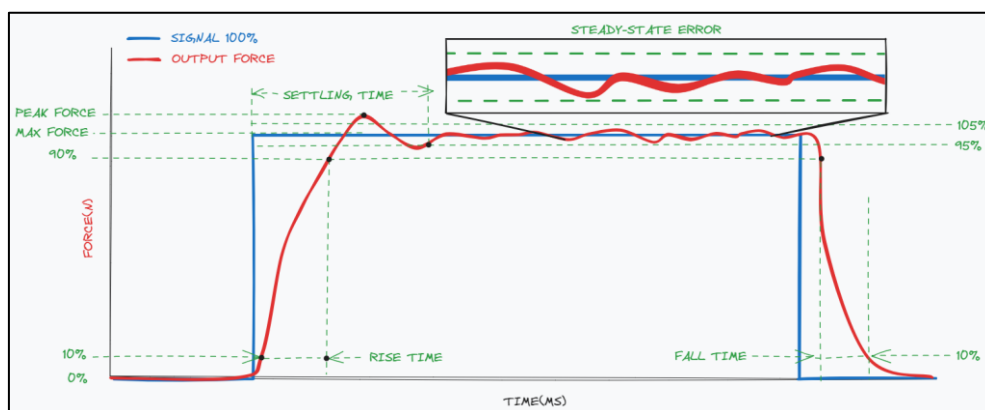


Figure 47. Step response analysis plot showing key performance metrics: input signal (blue) and output force (red) over time, with labeled measurements including rise time, settling time, fall time, peak force, maximum force, and steady-state error (magnified in inset). The plot illustrates standard haptic device dynamic behavior characteristics with 10% and 90-105% threshold markers.

the maximum operational signal. For devices using ESC to control actuators, this input change equates to a time pulse from 0 to 2ms, the maximum actuation pulse for an ESC. The primary aspect that step response analysis reveals is system stability. Specifically, it assesses the maximum stable actuation input where the device achieves and maintains its target position (e.g., moving from point A to point B) without overshooting or oscillating into instability. This behaviour is essential for kinesthetic haptic devices, as it establishes the precise relationship between input signal (voltage, current or pneumatic pressure) and output force or torque required for accurate performance. This analytical approach, standard in industrial settings, appears in various works discussed previously, such as Thor’s Hammer and JetPropeller, providing valuable insights into dynamic response characteristics and control loop effectiveness of haptic feedback systems.

Step response offers a straightforward yet effective method for evaluating haptic controller performance, identifying the following metrics:

- **Peak Force:** The maximum force output in response to maximum step input signal. This quantifies the device’s absolute force capacity and stability thresholds. Ideally, a mechanical device delivers peak force in minimal time without overshooting (which can lead to instability).
- **Max Continuous Force:** Once peak force is reached and the device settles after acceleration/rise time. This represents sustainable output force over prolonged periods with constant input signal, without overheating or internal damage.

- **Rise Time:** The time required for an actuator to reach desired output within a certain range, typically 10% to 90% of continuous force output. This indicates response speed to input signal, essentially the actuator’s response latency and time.
- **Settling Time:** An indication of response speed after input change and stabilisation. Shorter settling time indicates faster, more stable device response, sometimes termed acceleration or deceleration time.
- **Output Error:** Also known as steady-state error, this measures the difference between desired and actual output values when a haptic device reaches steady-state. Many actuators exhibit fluctuation error, particularly in handheld devices where actuator orientation directly affects performance. This appears in Aero-plane’s step response plot (Figure 46), showing average actuator force fluctuation around 0.8837 N.

Once collected, step response data can be analysed using software such as MATLAB, which offers built-in tools for system identification. MATLAB automatically identifies parameters like damping ratio and natural frequency by fitting mathematical models to response data, helping derive transfer functions that describe dynamic behaviour. The software calculates performance metrics including rise time, settling time, overshoot and steady-state error, characterising system responsiveness and stability. Ideally a haptic device will exhibit fast rise and settling times, minimal overshoot and negligible steady-state demonstrating stability and responsiveness to input signals.

2.3.3.4 Selected variables

Previous works by Samur et al. [1] and Kern et al. [140] offer comprehensive metrics describing kinesthetic haptic device operation and properties (full table in Appendix 15). These metrics typically encompass three main analyses: frequency response, step response and ramp response. However, in the kinesthetic device publications reviewed for this work, the more commonly reported analyses include physical device properties followed by step and ramp analysis [32], [45], [51], [103]. This highlights the emphasis in current research on analyses that have gained wider adoption within the community.

Based on the literature, the most common analyses performed by researchers were identified

Analysis	Variable	Metrics	Relevance
Physical Properties	Workspace	XYZ mm	Determines the physical range in which the device can operate
	Degrees of Freedom	2x rot., 1x transl.	Identifies the independent directions in which the device can move or rotate
	Electrical Properties	Voltage, Power	Defines the electrical requirements for the device to function properly
	Device Dimensions	XYZ mm	Specifies the physical size of the device, important for compatibility with use cases
	Weight	grams	Indicates the heaviness of the device which can impact user comfort and device portability
	Operational Noise	decibels	Provides an idea of the acoustic disturbance caused by the device during operation
	Vibrations	m/s ²	Measures the mechanical vibrations generated by the device, which can affect user experience
	Ramp Analysis	Max Force/Torque	Newtons
Min Force/Torque		Newtons	Specifies the smallest detectable force or torque that the device can exert
Hysteresis		Newtons	Describes the lag in response exhibited by the device when subjected to changing inputs
Sensitivity		Newtons/Volt	Measures the change in force output per unit change in input signal
Output Force Resolution		Newtons	Identifies the smallest detectable change in force output by the device
Dynamic Range		dB	Captures the range of force output the device can provide, from the minimum to maximum force
Stiffness		Newtons/mm	Quantifies the resistance of the device to deformation in response to an applied force
Step Response Analysis		Peak Force (overshoot)	Newtons
	Continuous Force	Newtons	Measures the force that the device can maintain continuously over a prolonged period
	Rise time	Seconds	Notes the time it takes for the device to reach the desired output from an initial state
	Settling Time	Seconds	Measures the time the device takes to stabilize after a change in input
	Output Error	%	Provides an estimation of the error between the desired and actual output of the device

Table 4. Comprehensive haptic device characterization metrics table showing three analysis categories (Physical Properties, Ramp Analysis, and Step Response Analysis) with associated variables, units of measurement, and relevance to device performance.

and synthesised into a more accessible metric set. This approach maintains the robustness of previous methodologies while reducing data capture complexity.

The proposed table comprises three primary components, each addressing different aspects of controller performance. 'Physical Properties' evaluates tangible and electrical device characteristics. 'Ramp Analysis' assesses controller performance under incrementally

increasing input conditions. 'Step Response Analysis' examines controller response to sudden input changes. These methods are summarised in Table 4. This framework aims for efficient categorisation through a comprehensive yet simplified evaluation approach. The streamlined measurement process for ramp and step response analyses enhances practicality and applicability. Section 3.5 will expand on this approach, outlining a minimalist yet effective setup for conducting these analyses with reproducibility and affordability at its core, providing clear guidelines for more consistent and accessible haptic device characterisation.

2.3.4 Summary/Learning from Previous Literature

5. SUMMARY
LEARNINGS

To conclude the literature review, this section presents key findings from previous work on haptic devices and their characterisation setups. Although numerous approaches generate kinesthetic feedback (including propellers, turbines, elastic elements, pneumatics and hydraulics) most employ electric motors as core actuators. The following Chapter 3, which focuses on the characterisation framework, will build upon the points presented in this summary of previous literature.

Characterisation Setup Build Materials

When it comes to the design of the testbed, previous work has used specialised beds or mechanisms like lathes or pulling testers, but the majority use custom-built testbeds designed and built using aluminium extrusions. These profiles provide numerous advantages, including cost-effectiveness, ready availability, and a high degree of customisation. They facilitate the assembly of various accessories and joints, allowing for versatile configurations tailored to specific experimental needs. Additionally, aluminium extrusions can be purchased in specific lengths, reducing assembly complexity and ultimately saving time.

Torque Vs Force Vs Stiffness

The decision to measure torque or force is particularly relevant for kinesthetic handheld controllers like Thor's Hammer[45] or Aero-plane [50], which generate force but deliver torque as the actuator sits at a distance from the user's hand on a handle. While torque measurement provides a more realistic assessment of user experience, it requires industrial-grade equipment

such as F/T sensors, resources not always accessible to researchers. Consequently, many opt for single load cells or multiple load cell combinations to measure actuator force, a more accessible technique.

The choice between torque or force measurement depends on the intended measurement objective. By employing a force/torque (F/T) sensor, researchers can evaluate both the user's perceived torques and the device's capabilities. However, the studies reviewed primarily focus on measuring force to assess device capabilities. Although measuring force alone cannot provide a comprehensive understanding of the user's haptic experience, it offers a foundation for describing device properties. This approach has a significant limitation: reporting torque measurements at the handle position inherently accounts for the distance from actuator to hand, allowing direct comparability between devices, while reporting actuator force only allow for actuator comparison. Even when actuator-to-hand distance is included (which is not always reported), this approach does not facilitate straightforward performance comparison between devices.

Drawbacks Of Measuring a Simplified Actuator

When characterising haptic controllers, if a controller employs multiple actuators, characterisation should include an equivalent number of actuators. In "Propel Walker" [110], the authors presents a system where a single ducted fan generates the output force, yet the actual device incorporates two actuators. Turbines and propellers do not generate twice the thrust by doubling actuator numbers. Therefore, the final output perceived by users remains undetermined. Accounting for actuator number and arrangement when designing force measurement setups is essential to accurately represent device capabilities.

What Sensor do I use?

Sensor selection for haptic device characterisation depends on the controller's output force magnitude, torque requirements, and number of actuated axes. While research labs with specialised equipment often use F/T sensors or industrial S-type sensors, beam load cells offer viable, low-cost alternatives with sufficient accuracy for many applications. Load cells provide advantages in accuracy and load range handling, though they require additional electronic components for amplification and data logging. Force gauges, built on load cell technology, offer ease of use and portability but have limitations in size and electronic integration for

closed-loop systems. Chapters 4 and 5 will detail these considerations further. Among reviewed papers, single-axis characterisation emerges as the most common approach for actuator assessment. The sensor should capture data in a setup that recreates the device's expected operation. For handheld devices delivering torque, an F/T sensor provides the ideal setup, while for devices providing direct skin actuation perceived along a single axis, beam load cells may be suitable.

Role of Automation During Characterisation Processes

The automation of actuator control during analyses such as ramp or step response is rarely explicitly stated in research papers. Automating haptic controller actuation during characterisation offers several advantages over manual methods. It ensures consistency in the measurement data, enables testing of the controller's lifespan [35], and reduces researchers' workload. The framework presented in Section 3.9 considers actuation and measurement as part of the same process, so analyses can be performed automatically to gather multiple samples for the same test to guarantee operational consistency.

Amplifier Sampling Rate & Resolution

A very popular choice for a load cell amplifier is the HX711, with 24-bit resolution. It is theoretically capable of providing a resolution of up to $5.85 \mu N$ when using the full resolution range over 10kg. However, the load cell's actual resolution is significantly lower than 24 bits due to its operation, noise, and other errors [144]. Furthermore, the amplifier is limited to a maximum sampling rate of 80 SPS. The ideal sampling rate remains an open question in the literature, with earlier research recommending a refresh rate of 320 Hz for tactile sensation and 30 Hz for proprioception [145], while more recent studies suggest rates of 1000 Hz [146], [147], [148]. However, Scandola et al. (2011) challenged the "1 kHz rule" [149], [150], finding no effect of refresh rates on their experiment, and concluding that force perception can be studied using low-frequency rates of 250Hz. Based on these studies, a higher refresh rate than 80 Hz is recommended to accurately capture kinesthetic feedback. A refresh rate of at least 250 Hz would be more suitable than the 80 Hz provided by the HX711 amplifier, based on findings from Scandola et al [150].

What Range Should the Characterisation Testbed Be Capable to Measure?

When developing a new haptic prototype, determining the magnitude of the target output torque or force range is one of the first steps. Ensuring the characterisation setup has bandwidth matching the controller's range enhances force resolution and measurement accuracy.

In order to provide some context to the ranges that can be expected, a survey of Haptipedia [4], including over 100 haptic devices, indicates that more than 80% are rated for peak force or torque under 100 N or 4 Nm. Additionally, Samur et al. [1] (p.49) present human sensory thresholds suggesting maximum force expected at the finger/wrist is approximately 50/60 N and 100 N at the elbow and shoulder. These are peak forces, and lower values are recommended to avoid injuries and fatigue. Previous literature can provide insights into torque or force ranges typically generated by similar devices, helping establish appropriate initial thresholds before a prototype is ready and guide characterisation setup design.

Signal Filtering.

Signal filtering removes unwanted signals or noise from data. For sensors used on kinesthetic feedback devices like load cells, low-pass filters eliminate high-frequency noise caused by electrical interference, vibration, or other sources. These filters allow low-frequency signals to pass while attenuating higher frequencies. When capturing movement or force readings, data is often filtered using a low-pass filter with cut-off frequency between 2 and 10 Hz [50], [53], [113], [151], [152], [153], [154]. This frequency range is not arbitrary but based on the fact that most human movements are performed at low frequency. Previous research suggests skilled actions such as typing or playing instruments are performed at 4-8 Hz and rarely above 10 Hz [155], [156]. Similarly, human body capabilities limit force frequency range to 2-6 Hz [155], [156], [157]. Therefore, for most kinesthetic haptic devices characterisation, filtering data with a 10 Hz cutoff frequency maintains signal integrity and remove undesired electrical noise or vibrations.

Previous Force Sensors Used.

A review of load cell types in previous research guided the initial selection process during framework development, later expanded through broader literature review. The analysis revealed that many publications omit specific sensor details. Of the 276 papers on kinesthetic haptic devices examined, only 60 provided specific sensor information. This is cover in

Appendix – Commercial Load Cells Comparison Table, categorised by type, publication reference and cost. Cost comparison demonstrates that single beam load cells average approximately £10, compared to entry prices of £3500 for F/T sensors (excluding amplifier costs).

The literature review provides valuable insights into haptic device design, actuation modalities and characterisation approaches. These findings informed the development of the proposed characterisation framework, FARE-KE, which are presented in the following chapter.

3 FARE-KE Framework for Affordable, Reliable Kinesthetic Evaluation

4. FARE-KE

HARDWARE

DATA
ACQUISITION

METHODOLOGY

The FARE-KE (Framework for Affordable, Reliable Kinesthetic Evaluation) streamlines the characterisation process for kinesthetic haptic devices through an integrated hardware and software toolkit. It establishes an initial standardisation framework while offering ready-to-implement, cost-effective, and replicable solutions for researchers during prototyping stages.

Developed through extensive research and hands-on experience, FARE-KE addresses key challenges in haptics performance metrics identified by Samur et al. [1]:

- **Equipment Limitations:** Many researchers lack access to advanced characterisation equipment, limiting their ability to document device performance comprehensively.
- **Complex Systems:** The characterisation process presents entry barriers for those without mechanical engineering backgrounds, necessitating a more straightforward analytical approach.
- **Variable Measurements:** Multiple characterisation techniques and measurement variables create inconsistency across the field, requiring standardised metrics to represent kinesthetic controller performance effectively (Appendix – Measurement Variables).

FARE-KE offers a comprehensive solution spanning literature review to hardware design and validation:

- **Literature Integration:** The framework examines both kinesthetic and vibrotactile haptic devices, presenting techniques from vibrotactile characterisation applicable to kinesthetic evaluation. This cross-domain approach enhances researchers' understanding of the characterisation process.

- **Support Resources:** Technical guidelines, hardware specifications, software documentation, and tutorials guide users through data sampling, hardware assembly, and characterisation procedures.
- **Data Acquisition System:** Research-informed sampling approaches ensure appropriate rates and bandwidth for haptic evaluation. The framework includes a 12-channel DAQ board capable of sampling three analog signals at approximately 10 kSPS or twelve signals at around 2.9 kSPS.
- **Hardware Innovation:** The framework emphasises cost-effective, replicable solutions through:
 - Validation of affordable beam load cell sensors
 - Design of custom 3-axis load cells using single-axis components
 - Development of 6-axis sensors measuring force and torque on x, y, and z axes using rapid prototyping technologies

This educational approach enables iterative design and refinement through direct performance feedback, particularly valuable during initial prototyping stages. The framework's comprehensive scope spans three chapters, addressing methodology, data acquisition system design, and hardware validation respectively.

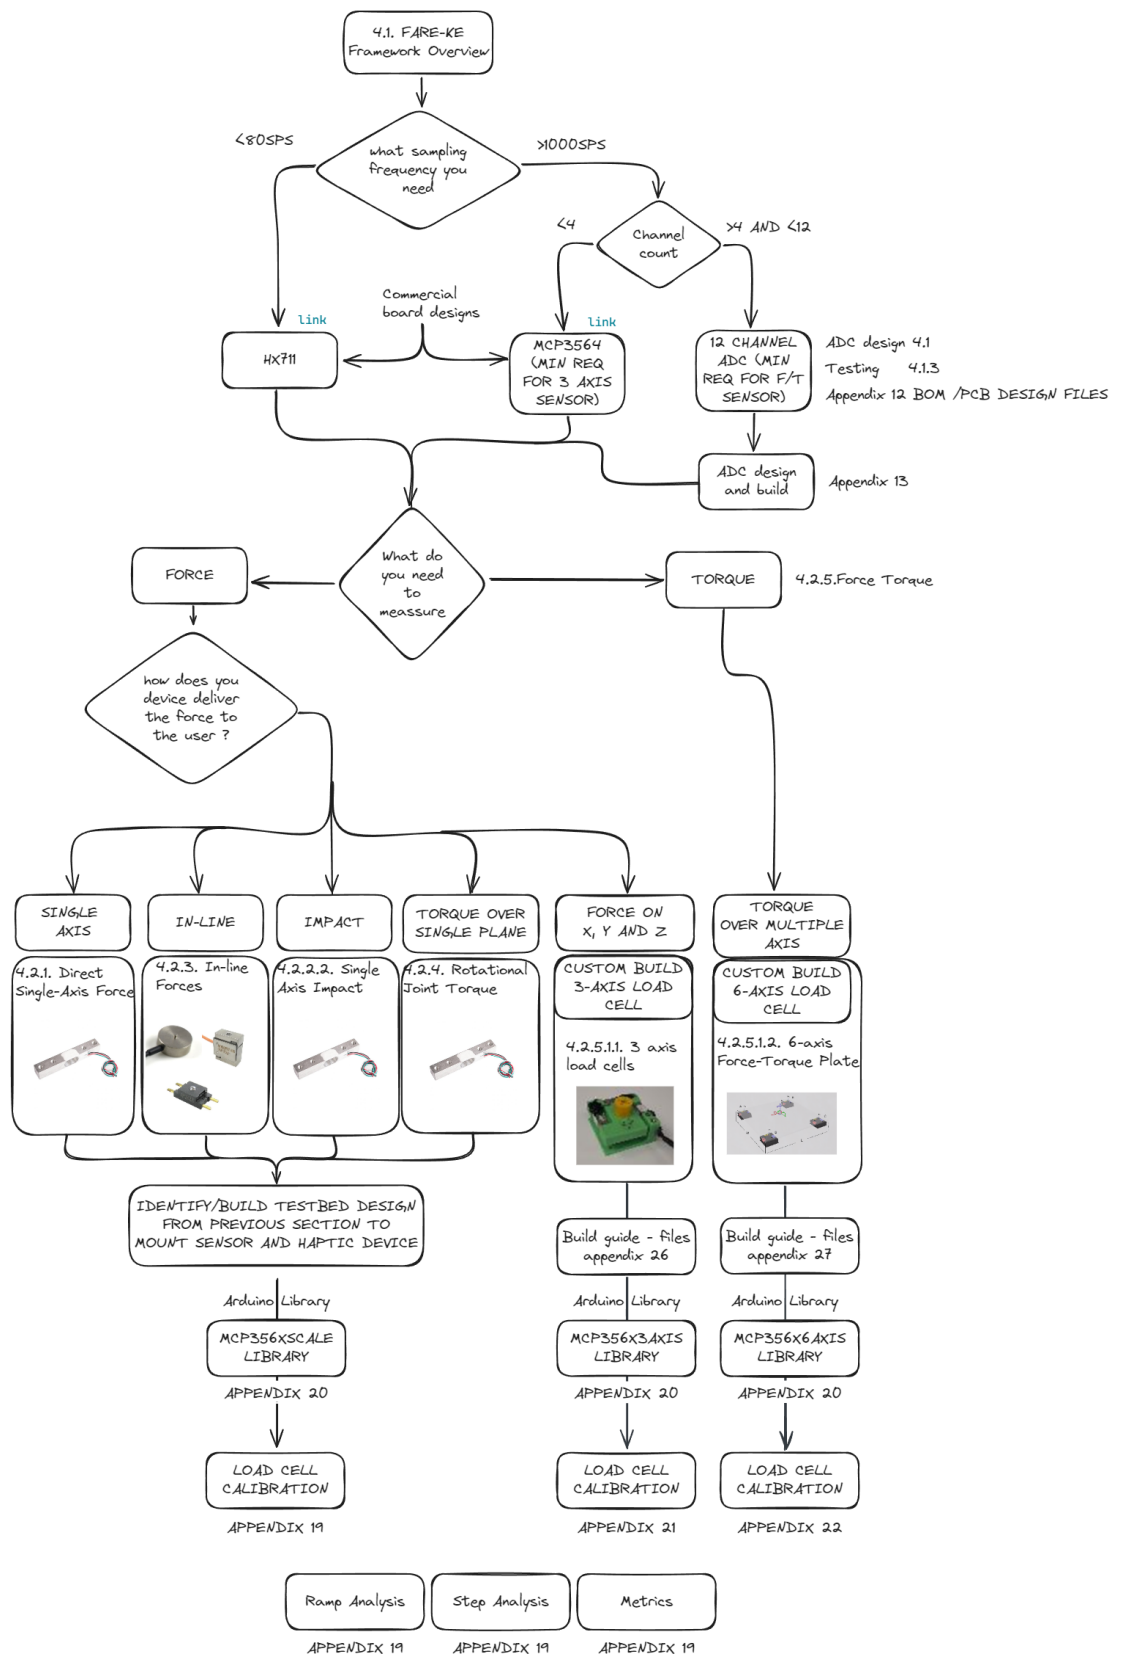


Figure 48. FARE-KE framework standardization flowchart illustrating the decision tree for haptic device characterization, from sampling frequency selection through sensor type determination, measurement approach, and analysis methods, with references to relevant appendices.

3.1 FARE-KE Framework Overview

FARE-KE establishes a structured approach to haptic device characterisation as illustrated in Figure 48. The framework addresses the identified gap between rapid prototyping technologies and traditional evaluation methods, providing consistent tools and procedures for performance assessment. The framework guides researchers through a logical progression from device classification to comprehensive performance analysis:

1. Initial Considerations and Classification (Section 3.2)

The characterisation process begins with identifying the haptic device’s actuation modality from five categories:

- Direct single-axis force
- Single-axis impact force
- In-line forces
- Rotational joint torque
- Multi-axis force/torque

This classification influences subsequent decisions on testbed design, sensor selection, and evaluation methodologies.

2. Testbed Design and Construction (Section 3.3)

Effective testbed design ensures measurements reflect actual device operation. The framework outlines principles for testbed construction using accessible materials like aluminium extrusions, prioritising:

- Structural rigidity to minimise unwanted vibrations
- Relevance of alignment between actuators and sensors
- Accommodation for various mounting configurations
- Scalability to different device dimensions

High-level design principles are supplemented by modality-specific implementation details in Chapter 5, where Sections 5.1.2, 5.2.2, and 5.3.1 provide detailed testbed configurations tailored to different kinesthetic feedback categories.

3. Load Cell Selection and Configuration (Sections 3.4-3.5)

Load cell selection depends on multiple factors:

- Device Complexity (Section 3.4): The framework correlates haptic modality with required sensor configurations, from single-axis beam load cells to complex multi-axis arrangements.
- Force Range and Capacity (Section 3.5): Guidelines for selecting appropriate force ranges based on expected device output and human sensory thresholds, ensuring optimal resolution while providing adequate safety margins.

4. Data Acquisition System (Section 3.6)

Three scalable approaches to data acquisition accommodate different expertise levels:

- Commercial Amplifier Breakout Boards: Simple, accessible solutions for basic applications
- Instrumentation Amplifiers with Microcontroller ADCs: Intermediate solutions with enhanced flexibility
- Custom-Designed Data Acquisition System: Advanced solution providing optimal performance across multiple channels

Section 3.6 introduces these approaches, while Chapter 4 provides comprehensive details on the custom DAQ system's design, implementation, and performance characteristics. Appendix 17 offers additional technical specifications, whilst Appendix 14 presents alternative amplification options for specific applications.

5. Signal Processing and Communication (Section 3.7)

Guidelines for effective signal processing and data logging using:

- Integration with appropriate microcontrollers
- Communication protocols for real-time data transfer
- Firmware implementation guidelines

These concepts are expanded in Section 4.2, which details specific implementation approaches for the custom DAQ system. The framework includes comprehensive software libraries, documented in Appendix 25 and available through GitHub repositories, covering ADC configuration, calibration, crosstalk compensation, and data logging across different load cell configurations.

6. Calibration Procedures (Section 3.8)

Calibration approaches for different load cell configurations:

- • Single-axis load cells: Simple scaling or regression-based approaches
- • Multi-axis configurations: Matrix-based methods addressing crosstalk effects
- • Force/torque plates: Complex calibration addressing interactions across 12 channels

Detailed implementation instructions for these calibration procedures appear in Appendices 18, 20, 21, and 23, complemented by software tools that automate calibration calculations. Section 5.5 provides additional context on calibration approaches for multi-axis designs.

7. Characterisation Analysis (Section 3.9)

Standardised analysis techniques for consistent device evaluation:

- • Physical property assessment: Dimensions, weight, degrees of freedom
- • Ramp response analysis: Maximum/minimum force, hysteresis, sensitivity
- • Step response analysis: Rise time, settling time, continuous force capabilities

supported by implementation code in Appendix 24

Comprehensive Validation Approach

FARE-KE extends beyond theoretical methodology to provide thorough validation of its components. Chapter 5 presents detailed validation across different kinesthetic modalities:

- Direct single-axis force (Section 5.1): Validation of affordable beam load cells
- Single-axis impact (Section 5.2): Assessment of rapid transient force measurement
- In-line forces (Section 5.3): Validation of tension/compression measurement
- Rotational joint torque (Section 5.4): Verification of torque measurement techniques
- Multi-axis force/torque (Section 5.5): Validation of custom-built multi-axis configurations

By following this structured approach, researchers can systematically implement effective device characterisation while maintaining consistent evaluation standards. The framework's modular structure enables adaptation to different expertise levels and research contexts.

3.2 Initial Considerations

To standardise the characterisation process for a range of devices, FARE-KE categorises devices by their actuation modality. The first step is to identify the force feedback modality of the haptic device, determining the testbed design, load cell type, and key experimental parameters. By considering the specific kinesthetic feedback delivery method, researchers can ensure load cell suitability and replicate the haptic controller's normal operating conditions, including orientation and degrees of freedom. FARE-KE summarises the actuation methods into five groups:

- **Direct Single-Axis Force:** Devices that apply force directly to the user's body, such as those providing continuous or instantaneous push or pull sensations, can be characterised using single-axis force sensors like beam load cells. These sensors are ideal for measuring linear forces with a well-defined vector component.
- **Single-Axis Impact Force:** Haptic devices that generate sudden bursts of energy, such as those using elastic bands or pneumatic actuation, require load cells with high sampling rates to accurately capture rapid force changes. Single-axis beam load cells are suitable, though they need amplifiers with high sampling rates and fast response times for measuring impact forces.
- **Inline Force:** Exoskeletons and devices that transmit force through cables or linkages can benefit from inline force sensors, such as S-type or button load cells. These load cells can measure both tension and compression forces and can be easily integrated into the device's mechanical structure.
- **Rotational Joint Torque:** Devices that generate torque through motors, gears, or other rotational mechanisms can be characterised using beam load cells when torque operates along a single plane by measuring force at a known distance from the rotation axis. The choice between a dedicated torque sensor and a force sensor depends on the specific device design and desired accuracy.
- **Force/Torque:** Haptic devices generating forces and torques in multiple axes require multi-axis force/torque sensors. Six-axis load cells can measure forces and torques in three orthogonal axes, providing comprehensive device output characterisation. The FARE-KE framework offers a cost-effective solution by utilising an array of single-axis load cells to create a custom 6-axis force/torque sensor.

A more detailed analysis of these characterisation groups will be presented in Chapter 5 , FARE-KE Sensor/Hardware Validation, introducing examples demonstrating approaches for different haptic modalities.

3.3 Testbed Designs

When designing a testbed for haptic device characterisation, understanding how the force is delivered to the user is crucial, guiding the alignment, placement, and load cell design to integrate with the actuator. For straightforward scenarios involving direct single-axis forces or impacts, a single-axis load cell typically suffices. The cell should be mounted in a position or orientation that accurately reflects the actuator’s operation. The testbed must emulate the device’s operational context (how it will be held or mounted by the user) and accommodate the same quantity and configuration of actuators as the final prototype to ensure representative data capture.

Extruded aluminium is frequently employed in testbed construction due to its versatility, standardisation, and wide availability of attachments and coupling systems. This allows for a flexible setup that can be easily adjusted to meet specific requirements. The 40x40 mm aluminium extrusions are sturdy enough to prevent bending or deformation under the weight or forces applied by most handheld kinesthetic haptic devices.

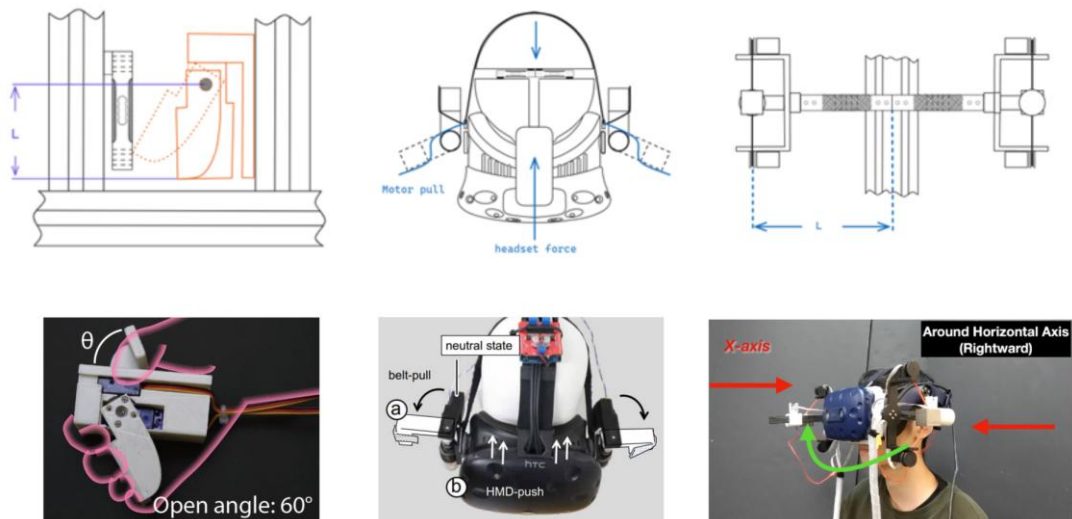


Figure 49 Examples of testbed configurations for haptic device characterization: (a-c) Technical diagrams showing proposed measurement setups, and (d-f) corresponding device implementations for PacaPa rotational joint [45], FacePush head-mounted system [163], and OscilHead oscillating mechanism [119].

Figure 49 shows various examples of alternative testbed design for measuring actuation representative of haptic device force generation. The Pacapa [36] testbed setup showcases torque measurement via the lever arm of the device actuator. In the FacePush [158] configuration, motors pulling straps generate push forces on a mannequin head equipped with single beam load cells on each side, directly measuring design capabilities. Oscihed [111] setup demonstrates rotating torque measurement from oscillating masses using single beam load cells, taking into account the actuator length from the anchor point.

3.4 Load Cell Count and Type

The selection and number of load cells are determined by the haptic device’s modality. For devices exerting single-axis forces, load cells measuring linear forces (Figure 50 a) provide a suitable solution. This type is most prevalent due to its mechanical simplicity and availability of low-cost models under \$10.

Multi axis force outputs necessitate multi-axis load cells. Three-axis load cells (Figure 50 c) capture multidimensional force vectors, while 6-axis load cells (Figure 50 e) are typically employed for multi-axial torque. Although 3-axis load cells are uncommon in haptic device characterisation literature, 6-axis load cells are preferred for handheld devices where the actuator is distant from the user’s grip point, hence capture the final use perceive actuation.

Given the high cost and limited availability of commercial multi-axis load cells, the proposed framework introduces sensor designs leveraging low-cost beam load cells. Figure 50 d illustrates a conventional 3-axis load cell design, mechanically consisting of an equivalent to three single beam load cells aligned along orthogonal axes. Utilising 3D printing technology, we designed a comparable design (Figure 51 b) where three single load cells are mounted in a specific configuration to mimic the industrial design. Chapter 5 provides detailed validation and calibration overview, with Appendix 18: FARE-KE 3-Axis Single Point Calibration providing a guide for 3-axis load cell calibration.



Figure 50. Load cell designs for haptic device characterization: (a) Beam load cell; (b) Button load cell; (c) S-type load cell; (d) 3-axis load cell; (e) 6-axis force/torque load cell.




Figure 51. Multi-axis force sensing configurations: (a) Assembly of single-axis load cells arranged to form a 3-axis load cell setup; (b) 3-axis load cell complete with housing; (c) 6-axis force-torque plate constructed using four 3-axis load cells mounted between acrylic plates.

Although 3-axis load cells are not widely used in kinesthetic haptic device characterisation, they provide a foundational approach for 6-axis force-torque sensor design. By assembling four 3-axis load cells between two solid plates, researchers can create a 6-axis force-torque measurement plate. This arrangement allows summation of data from twelve channels (three per load cell) to calculate force and torque relative to the plate’s centre. The configuration enables development of affordable 6-axis measurement plates, facilitating precise monitoring of haptic device force output and torque generation. Section 5.5 will examine the design and validation of this 6-axis load cell, with comprehensive validation, calibration, and assembly instructions provided in Appendix 23: FARE-KE 6-Axis Load Cell Multipoint Calibration.

3.5 Considering Range and Capacity

Before selecting a load cell, it is crucial to determine the required range or capacity. Industrial load cells are typically classified in newtons, while low-cost variants are often rated in kilograms. At this stage, researchers can estimate the expected output by calculating the actuator end effector mass or actuator speed, using this value as an initial reference point.

The literature review provides a contextual foundation for setting realistic force or torque targets, most haptic devices operate within peak forces under 100 N or torques of 4 Nm (Figure 52). However for sensors, it is crucial to select a device with a range close to the maximum expected output to maintain high resolution and measurement accuracy.

For impact force measurements, selecting a load cell with high resonant frequency is critical to capture rapid force variations and accommodate the expected force range. Manufacturer-specified overload values provide essential protection against unanticipated high-force events.

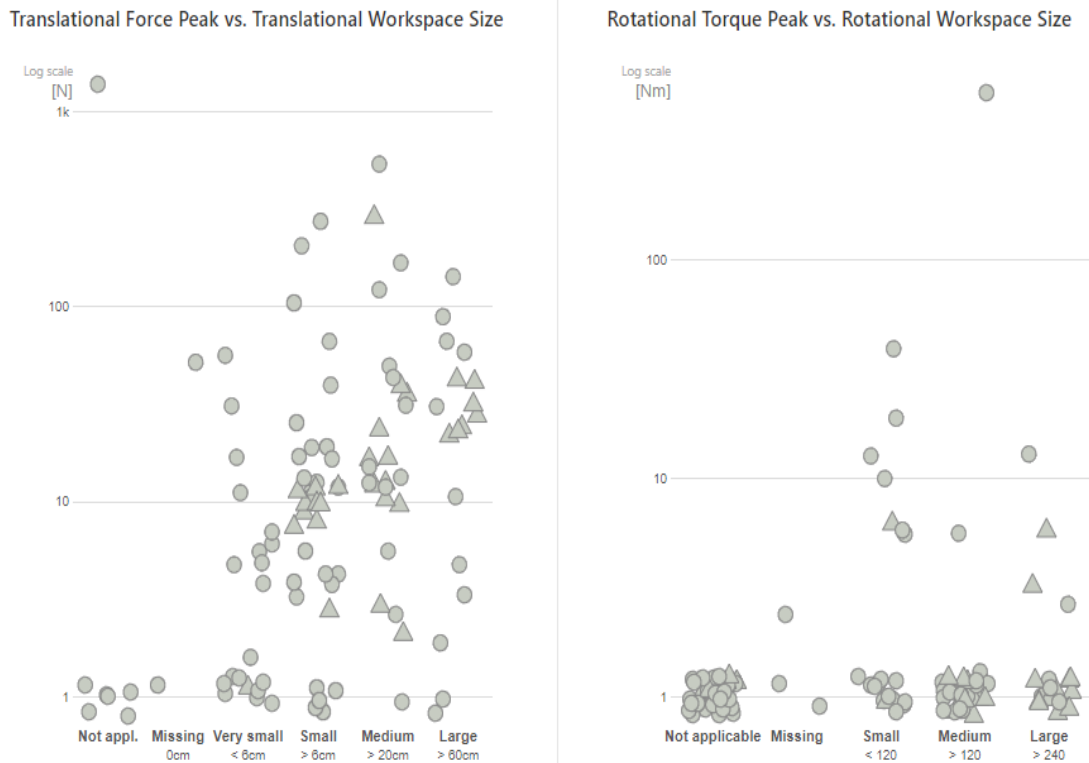


Figure 52. Haptipedia data visualization showing the relationship between workspace size and maximum force/torque capabilities of haptic devices: (left) translational force peak vs. translational workspace size; (right) rotational torque peak vs. rotational workspace size, with both graphs using logarithmic scales to display the wide range of device capabilities.

By employing four 3-axis load cells (each constructed from three independent single-axis Beam load cells, the force) torque plate integrates 12 individual sensing channels. For any given force direction, only the channels aligned with that axis contribute to the measurement. In each 3-axis unit, the capacity on each axis is limited by the rating of the individual beam load cell used. However, when arranged in a force-torque plate, the load is distributed among the four load cells, effectively increasing the overall load capacity for that direction and enhancing the system's resistance to bending moments.

3.6 DAQ Selection

To measure the forces exerted on the force sensor, the resultant signals must be amplified and digitalised. Load cell technology relies on a Wheatstone bridge, which outputs small voltage changes in response to strain gauge deformation. An amplifier is necessary to upscale these voltages to a measurable level within the ADC range.




Figure 53. Load cell amplification and data acquisition options: (a) HX711 breakout board; (b) MCP3564 breakout board; (c) INA125 instrument amplifier; (d) FARE-KE proposed custom DAQ board featuring multiple channels for multi-axis force measurement.

This section explores three practical signal amplification and conversion approaches, arranged by increasing implementation complexity:

1. **Commercially Available Amplifier Breakout Boards:** The simplest method involves selecting a readily available amplifier. Two amplifier breakout boards were evaluated during system development:
 - **HX711:** A commonly used amplifier in haptic device characterisation (Figure 53a). These amplifiers are widely accessible through online stores. The extensive availability of online examples and Arduino libraries makes it an entry-level choice for load cell amplification. However, its maximum sampling rate is limited to 80 SPS.
 - **MCP3564 Breakout:** Available on [Tindie](#) (Figure 53b), this 24-bit sigma-delta ADC supports up to four differential channels for measuring load cell output changes. The ADC can sample the load cell at rates up to 12.8 kSPS with its internal oscillator, and up to 153.6 kSPS with an external clock.
2. **Using Instrument Amplifiers:** These differential amplifiers (Figure 52c) are optimised to amplify low-level signals while rejecting noise and interference. Used in combination with the microcontroller’s integrated Analog-to-Digital Converter (ADC), the amplified analog signal is converted to a digital format for processing. This setup (Figure 54) offers two primary advantages:
 - **Gain Control:** The amplification system’s gain can be fine-tuned using a variable resistor. This adjustability enables calibrating signal amplification to utilise the ADC’s full range. Practically, one can connect a load cell, apply the maximum expected load, and adjust the circuit’s

gain to span the entire ADC range. This approach optimally matches the load cell's measurement range to the ADC, achieving the highest possible resolution for the specific load cell.

- **Sampling Flexibility:** As an analog component, the instrument amplifier does not impose sampling rate limitations. The load cell's sampling rate is determined by the microcontroller's onboard ADC, potentially offering faster rates compared to digital amplifiers like the HX711.

While using instrumentation amplifiers with a microcontroller's ADC offers benefits, this approach has some limitations:

- **Limited ADC Channels:** The number of load cells that can be read simultaneously is limited by the number of ADC channels the microcontroller has. This can be restrictive when multiple load cells monitoring is required.
- **Cost per Channel:** Each instrumentation amplifier typically costs between £5 and £10. When your system necessitates multiple channels, the overall cost increases, which will increase the overall cost.
- **Resolution Constraints:** Typical microcontroller ADCs, such as those on Arduino boards, provide up to 12-bit resolution. This contrasts with specialised boards like the HX711 or MCP3564, which offer 24-bit

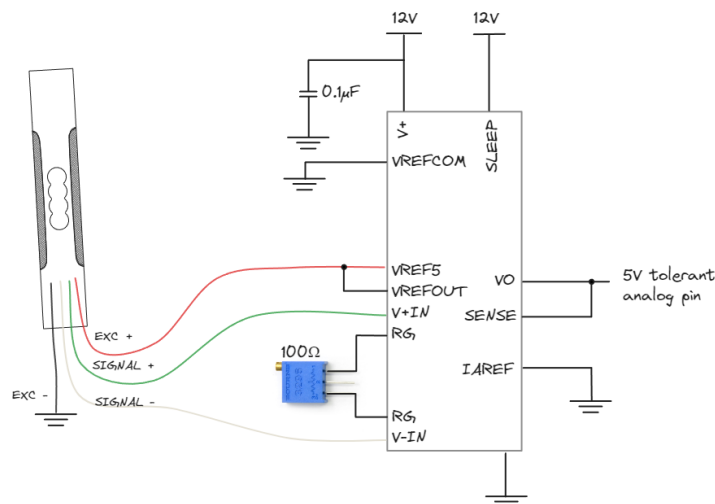


Figure 54. Example wiring circuit diagram for an INA125 instrument amplifier connection to a load cell, showing power supply, reference voltage configuration, and signal conditioning components required for proper load cell interfacing.

resolution. Higher resolution captures more data. Lower resolution may yield less accurate readings, potentially compromising applications demanding high-fidelity measurements, although lower bits or resolution also suffer from significant higher signal to noise ratio. For a summary of previously tested instrumentation amplifiers and a detailed comparison of their characteristics, refer to Appendix 14.

3. **Custom-Built Data Acquisition System:** The framework introduces a custom-built Data Acquisition System (DAQ) (Figure 55) designed using the MCP3564 ADC chip. This chip has the capability to sample up to four differential channels (equating to four load cells) with technical capabilities of over 150 kSPS. Since the proposed force/torque plate design requires readings from up to 12 channels, a custom board that implements three MCP3564 was designed.

Chapter 5 will cover performance testing and validation of this board in detail. The validation process includes tests for linearity, noise, crosstalk, and temperature stability, ensuring that the DAQ meets the requirements for accurate and reliable data acquisition in haptic device characterisation.

Although the MCP3564 includes built-in amplification, the board design incorporates an extra amplifying stage using op-amps configured similar to an instrument amplifier, with a single resistor controlling amplification gain. The current design uses fixed-value resistors to set gain levels, with future versions aiming to incorporate trimmer potentiometers for adjustable gain to fully utilise the amplifier's range. In the proposed design, a 2 k Ω resistor produces a gain

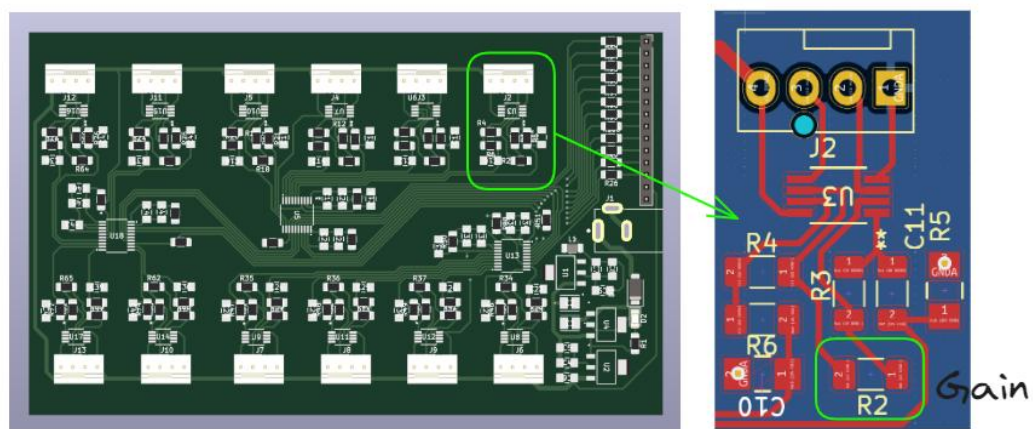


Figure 55. DAQ board prototype implementation: (left) Complete PCB layout, (right) Close-up of a single channel highlighting the gain amplification resistor used to control signal sensitivity.

of $\times 101$, with gain being inversely proportional to the resistor value. The PCB accepts 5 to 12 volts via a barrel jack and includes voltage regulation and protection circuitry to ensure stable operation across this range.

3.7 Data Logging

Force and torque values computed on the microcontroller must be transmitted to a computer for logging and analysis via USB serial port. During FARE-KE framework development, various data logging tools were evaluated, with TyCommander and Telemetry Viewer outperforming other solutions.

TyTools comprises standalone utilities designed for communication with Teensy microcontrollers. TyCommander offers real-time data monitoring and message printing via serial communication, enabling users to transmit and receive data through the serial port and save the received data as CSV files. TyCommander demonstrated reliable performance, establishing it as a dependable option for the FARE-KE framework. Other serial communication software, such as PuTTY or Arduino Serial Monitor, failed to capture data reliably at high baud rates.

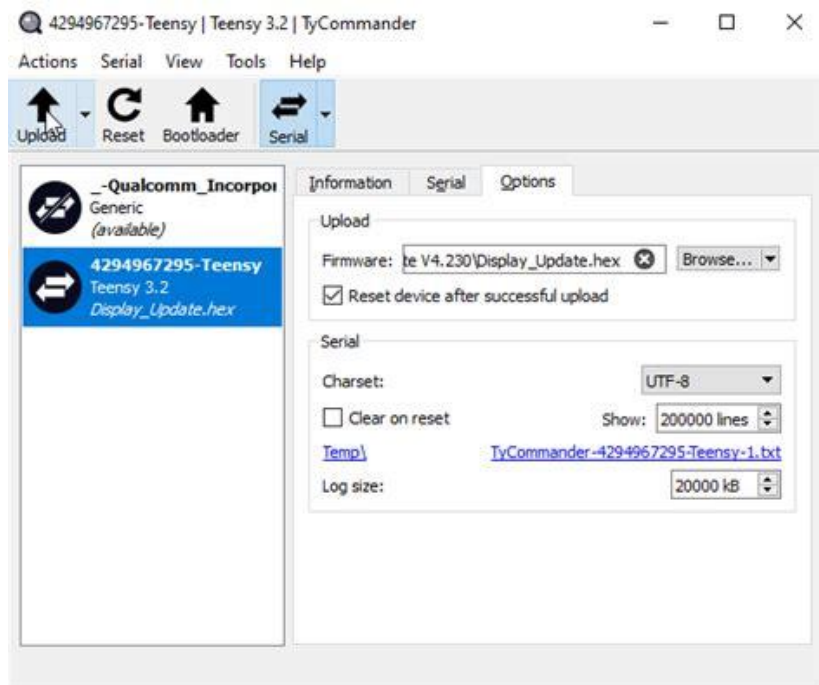


Figure 56. TyCommander software interface for managing serial communication with the Teensy 3.2 microcontroller used in the FARE-KE data acquisition system.




Figure 57. Telemetry Viewer interface showing real-time multi-axis force and torque measurements with color-coded channels, displaying both raw data trends and processed values for comprehensive haptic device characterization.

Telemetry Viewer extends functionality with advanced data visualisation and logging capabilities. This tool allows users to customise telemetry panels to monitor data streams from diverse sources, including multiple serial ports and Wi-Fi. It enables real-time data visualisation, customisation of the data view layout, and data export to CSV format for further analysis. By providing compatibility with these reliable tools, the FARE-KE framework streamlines the data logging process, making it more accessible to researchers and practitioners.

3.8 Calibration Process

Accurate calibration establishes a precise relationship between digital readings from the DAQ and forces being measured, ensuring collected data represents the physical phenomena under investigation. We have introduced three main systems of load cells that need calibration:

Single-Axis Load Cell Calibration

Single-axis load cells are well-documented sensors. A step-by-step calibration guide can be found in Appendix – FARE-KE Calibrating Single Axis Load Cell. Load cells exhibit high linearity, which simplifies the calibration process. Tests on three low-cost beam load cells confirmed their linear response with R^2 greater than 0.98. Due to their high linearity, calibration typically establishes a single scaling factor by applying a known weight and

recording the corresponding digital reading. This creates a mapping between digital values and applied force in grams or gram-force. Mathematically, the scaling factor can be calculated as follows:

$$\text{Scaling Factor} = \frac{\text{Known Weight}}{\text{Digital Reading}}$$

Once obtained, this scaling factor converts digital readings from the ADC into gram-force values. To convert gram-force (gf) to newtons (N), multiply the gram-force values by gravitational acceleration (approximately 9.81 m/s²), then divide by 1000 to convert grams to kilograms.

3-axis Load Cell Calibration

The calibration process for 3-axis load cells resembles single-axis calibration but includes orienting the assembled load cell along all three axes and suspending weights on each. Industrial 3-axis load cells constructed from a single body can achieve high accuracy using a single-point calibration matrix process. This process employs an inverse matrix to correct for crosstalk and manufacturing inaccuracies. By applying a known load to one axis and measuring outputs, calibration coefficients (including sensitivity and crosstalk) are determined and organised into a matrix. The inverse of this matrix transforms raw data into accurate force values, isolating crosstalk influence and enhancing precision for each axis. A detailed guide is available in Appendix – FARE-KE 3-Axis Single Point Calibration.

The calibration matrix K encapsulates the load cell's response to applied forces, accounting for both direct effects on each axis and indirect effects (crosstalk). To determine the actual forces F from observed raw readings R , the inverse of the calibration matrix K^{-1} is applied as follows:

$$F = K^{-1} \cdot R$$

For custom-built 3-axis load cells incorporating 3D-printed components and single-axis load cell arrays, thorough performance and linearity evaluation is essential. Assembly quality and alignment precision variations can affect linearity and introduce crosstalk, compromising measurement accuracy. For those reasons to assess performance and account for potential nonlinearities and crosstalk, a multipoint calibration method is employed on the proposed

design. This method advances beyond single-point calibration by considering multiple loads, more effectively addressing nonlinear behaviour and inter-axis crosstalk.

The multipoint calibration process involves applying a series of known weights along each axis and recording the corresponding digital readings. The process is outlined in Appendix – FARE-KE 3-Axis Multipoint Calibration Method. Least-squares regression optimises the calibration matrix by minimising discrepancies between observed outputs and the predicted model, incorporating all inter-axis interactions. This method's strength lies in its comprehensive accounting for crosstalk, resulting in a refined calibration matrix that yields precise force measurements from the multi-axis load cell.

The objective is to identify the calibration matrix K that minimises squared differences, ensuring load cell outputs R closely match theoretical predictions from F , after adjusting for crosstalk and interactions among the axes:

$$\min_K \| R - KF \|^2$$

6-axis Force Torque Plate Calibration

The calibration of a 6-axis load cell follows a similar approach to single axis or 3-axis load cell calibration. Known weights are suspended along each axis, and for torque measurements, a lever arm of a specific length is used with the suspended weights. The calibration procedure for 6-axis load cells is provided in Appendix – FARE-KE 6-Axis Load Cell Multipoint Calibration.

Unlike previous calibrations where digital readings from load cells were mapped to a known weight, to obtain force and torque readings from the 6-axis force-torque plate, calibrated readings (in Newtons) from individual 3-axis load cells are utilised. The force-torque plate comprises four 3-axis load cells, yielding a total of 12 individual force readings. During the 6-axis F/T plate calibration, these 12 readings are combined to calculate the forces and torques along the X, Y, and Z axes using the following equations:

$$\begin{bmatrix} F_x \\ F_y \\ F_z \\ \tau_x \\ \tau_y \\ \tau_z \end{bmatrix} = \begin{bmatrix} F_{x,A} + F_{x,B} + F_{x,C} + F_{x,D} \\ F_{y,A} + F_{y,B} + F_{y,C} + F_{y,D} \\ F_{z,A} + F_{z,B} + F_{z,C} + F_{z,D} \\ L/2(F_{z,D} + F_{z,C} - F_{z,B} - F_{z,A}) \\ L/2(F_{z,A} + F_{z,D} - F_{z,C} - F_{z,B}) \\ L/2(F_{x,A} - F_{x,B} - F_{x,C} + F_{x,D}) + W/2(F_{y,A} + F_{y,B} - F_{y,C} - F_{y,D}) \end{bmatrix}$$

Where:

- $F_{i,j}$ represents the force reading along the i -axis from the 3-axis load cell j (A, B, C, or D).
- L and W are the length and width of the force-torque plate, respectively.

These readings have been compensated for crosstalk within each 3-axis load cell by its own calibration matrix. However, due to imperfections in assembly, alignment, and construction of the force-torque plate, the combined readings may still exhibit crosstalk and inaccuracies. To compensate for these errors, a multipoint calibration approach using the least-squares method is employed.

The multipoint calibration process for the 6-axis force-torque plate, involves suspending known forces and torques and recording the uncalibrated readings from the computation of the above equations. The least-squares method finds the calibration matrix K that minimises the sum of squared differences between the observed force readings R and the model's predictions based on the applied forces and torques F , similarly to the 3 axis calibration:

$$\min_K \| R - KF \|^2$$

This approach accounts for crosstalk and interactions between individual 3-axis load cells within the force-torque plate, resulting in a refined calibration matrix that provides accurate force and torque measurements. The calibration matrix transforms the forces and torques computed from individual 3-axis load cells into the true forces and torques acting on the force-torque plate.

3.9 Device Characterisation Process

Variable	Metrics	Relevance
Workspace	XYZ mm	Determines the physical range in which the device can operate
Degrees of Freedom	2x rot., 1x transl.	Identifies the independent directions in which the device can move or rotate
Electrical Properties	Voltage, Power	Defines the electrical requirements for the device to function properly
Device Dimensions	XYZ mm	Specifies the physical size of the device, important for compatibility with use cases
Weight	grams	Indicates the heaviness of the device which can impact user comfort and device portability
Operational Noise	decibels	Provides an idea of the acoustic disturbance caused by the device during operation
Vibrations	m/s ²	Measures the mechanical vibrations generated by the device, which can affect user experience

Table 5. Metrics and variables for physical characterization of haptic devices.

The characterisation process within the FARE-KE framework is based on the following elements: physical properties, ramp analysis, and step response analysis, which were introduced in Chapter 3.

The first phase focuses on physical attributes of the device. Data acquisition for variables such as degrees of freedom, electrical properties, dimensions, and weight typically requires no specialised equipment. However, measuring operational noise and vibrations necessitates specific instruments. In the absence of specialised acoustic equipment, smartphone applications can serve as an interim solution for assessing noise, though with reduced precision. Given the proximity of users to haptic devices, it is recommended to measure noise levels at 50 cm from the device during activation. Vibrations can be measured using accelerometers to evaluate mechanical stability and unintended motions.

Ramp and Step Response Analyses.

The next phase involves conducting ramp and step response analyses. These provide insights into the device’s performance under varying input conditions and its response to sudden changes. For both analyses, the haptic device or its actuators are attached to a testbed or measuring rig integrated with appropriate load cells. Measurements are recorded during activation to capture the dynamic response.

To standardise and automate the process, a microcontroller handles data acquisition from load cells and actuator control. Various methods, including pulse-width modulation, digital signals, SPI, or serial communication, accomplish this. The microcontroller, interfacing with




Figure 58. Complete haptic system diagram showing the integration of components: Teensy microcontroller for data acquisition, custom multi-channel DAQ board, Arduino controller for actuator management, and 6-axis force-torque plate for measurement. The diagram illustrates signal paths and connections for comprehensive characterization, with the Arduino controlling compressed air flow to the pneumatic actuator while the force-torque plate measures the resulting mechanical response. This setup enables both step response testing and ramp response analyses through various control methods including Boolean states, PWM signals, and serial communication, ensuring precise manipulation of the haptic device's actuation parameters AirRacket [38].

the Analog-to-Digital Converter (ADC), also issues control commands to the haptic device, coordinating position, speed, or direction adjustments as necessary.

Figure 58 illustrates a haptic device with a dedicated controller modulating the actuator's input. The actuator is positioned on a 6-axis force-torque plate, while the data acquisition microcontroller manages operation by transmitting control commands. These commands activate or deactivate the actuator for step response testing or adjust input parameters for ramp response analyses.

Performing data capture and device control on the same microcontroller synchronises actuation signals with the measurement timeline and force-torque readings. This enables accurate assessment of additional properties such as latency during data analysis and provides deeper insights into the haptic device's performance characteristics.

Data Analysis

The referenced section provides a comprehensive guide on data processing techniques, tools, and metrics for evaluating device performance.

The FARE-KE framework includes Python-based data processing tools and techniques that integrate with the data acquisition system. These tools include libraries for data importing, synchronisation, filtering, signal processing, and custom functions for calculating standardised

metrics and generating plots. The Python scripts streamline the data analysis process and facilitate result interpretation.

Ramp Analysis

Variable	Metrics	Relevance
Max Force/Torque	Newtons	Determines the maximum force or torque that the device can exert
Min Force/Torque	Newtons	Specifies the smallest detectable force or torque that the device can exert
Hysteresis	Newtons	Describes the lag in response exhibited by the device when subjected to changing inputs
Sensitivity	Newtons/Volt	Measures the change in force output per unit change in input signal
Output Force Resolution	Newtons	Identifies the smallest detectable change in force output by the device
Dynamic Range	dB	Captures the range of force output the device can provide, from the minimum to maximum force
Stiffness	Newtons/mm	Quantifies the resistance of the device to deformation in response to an applied force

Table 6. Variables measured using the ramp response analysis, including seven key metrics for haptic device characterization. Each metric is presented with its appropriate measurement unit and relevance to device performance evaluation.

Ramp analysis determines static characteristics like maximum and minimum force/torque and dynamic properties such as hysteresis and stiffness, offering a complete profile of the device’s operational capabilities. This analysis predicts how the device will perform in real-world applications, where inputs vary in magnitude and speed. To perform the ramp analysis using the FARE-KE framework, follow these steps:

1. **Set up the test environment:**
 - Mount the haptic device securely on the testbed or measuring rig and couple it with the load cell(s).
2. **Configure the data acquisition system:**
 - Connect the Teensy microcontroller to the haptic device and load cell(s).
 - Load the `"rampResponseTelemetry.cpp"` script onto the Teensy.
 - Adjust the script parameters according to your specific haptic device requirements, including:
 - Serial communication or PWM values
 - Minimum and maximum input signal values
 - The number of ramp cycles
3. **Run the ramp analysis:**
 - Run the `"rampResponseTelemetry.cpp"` script on the Teensy.
 - The microcontroller will generate a gradually increasing and decreasing input signal based on the defined parameters.

- Simultaneously, it will acquire force/torque data from the load cell(s) at a sampling rate higher than 250 Hz to ensure accurate capture of kinesthetic feedback.

4. **Data collection and storage:**

- The collected data will be sent via serial communication to a connected computer.
- The data format includes timestamp, force values (ForceX, ForceY, ForceZ), torque values (TorqueX, TorqueY, TorqueZ), and the input signal value (Actuation).
- Save the received data as a CSV file for further analysis.

5. **Data analysis:**

- Use the provided [rampResponseAnalysis.py](#) Python script to process and analyse the collected data.
- The script will calculate the following metrics:
 - Maximum force/torque
 - Minimum force/torque
 - Hysteresis
 - Sensitivity
 - Output force resolution
 - Dynamic range
- The script accounts for actuator latency by applying a 25% threshold when calculating metrics such as hysteresis, sensitivity, and output force resolution.

This procedure facilitates comprehensive ramp analysis through the FARE-KE framework. The scripts streamline the analysis process, enabling researchers to focus on result interpretation and understanding of device performance characteristics under varying load conditions. The `rampResponseAnalysis.py` script establishes a foundation for analysis that accommodates customisation based on specific requirements or additional metrics. The subsequent section addresses step response analysis, another essential method for characterising dynamic performance of haptic devices.

Step Response Analysis

Variable	Metrics	Relevance
Peak Force (overshoot)	Newtons	Records the maximum force output by the device in response to a sudden change in input
Continuous Force	Newtons	Measures the force that the device can maintain continuously over a prolonged period
Rise time	Seconds	Notes the time it takes for the device to reach the desired output from an initial state
Settling Time	Seconds	Measures the time the device takes to stabilise after a change in input
Output Error	%	Provides an estimation of the error between the desired and actual output of the device

Table 7. Variables measured using the step response analysis, presenting five time-domain performance metrics for haptic device evaluation. Each variable is shown with its appropriate unit of measurement and functional relevance to device performance assessment.

Step response analysis characterises the dynamic performance of haptic devices. This analysis focuses on the device’s reaction to a sudden change in input, providing insights into response time and behaviour under peak load conditions. Key performance indicators such as rise time, overshoot, and settling time are documented during this analysis. To perform the step response analysis using the FARE-KE framework, follow these steps:

1. **Set up the test environment:**
 - Ensure that the haptic device is securely mounted on the testbed or measuring rig and actuator is coupled to the load cell.
2. **Configure the data acquisition system:**
 - Connect the Teensy microcontroller to the haptic device and load cell(s).
 - Load the `”stepResponseTelemetry.cpp”` script onto the Teensy.
 - Adjust the script parameters according to your specific haptic device, including:
 - Serial communication or PWM values
 - Input signal value for the step response (typically the maximum safe operating value)
 - The duration of the step response test
3. **Run the step response analysis:**
 - Run the `”stepResponseTelemetry.cpp”` script on the Teensy by sending a 0 via serial command to the microcontroller.
 - The microcontroller will apply a sudden input signal change to the haptic device and maintain it for the specified duration.
 - Simultaneously, it will acquire force/torque data from the load cell(s) at a sampling rate higher than 250 Hz to ensure accurate capture of the device’s dynamic response.
4. **Data collection and visualisation:**
 - Data will be sent via serial communication to a connected computer.

- The data format includes timestamp, force values (ForceX, ForceY, ForceZ), torque values (TorqueX, TorqueY, TorqueZ), and the input signal value (Actuation).
- TelemetryViewer will visualise the collected data in real-time.
- Save the received data as a CSV file for further analysis.

5. **Data analysis:**

- Use the provided "[stepResponseAnalysis.py](#)" Python script to process and analyse the collected data.
- The script will calculate the following metrics:
 - Peak force (overshoot)
 - Continuous force
 - Rise time
 - Settling time
 - Output error
- The script will identify the step segments based on the changes in the input signal value and calculate the metrics for each step response.
- It will also provide visualisations of the step response segments for a more intuitive understanding of the device's behaviour.

6. **Interpretation of results:**

- Interpret the results obtained from the "[stepResponseAnalysis.py](#)" script to understand your haptic device's dynamic characteristics.

This procedure facilitates comprehensive step response analysis within the FARE-KE framework. The scripts automate data acquisition and analysis processes, enabling focused interpretation of results and deeper understanding of device dynamic performance characteristics. The `stepResponseAnalysis.py` script establishes a foundation for analysis that accommodates customisation based on specific requirements or additional metrics. Its modular structure facilitates extension and adaptation to diverse research applications.

3.10 Conclusion

The FARE-KE framework's standardised characterisation process, coupled with its data analysis tools and techniques, promotes transparency and collaboration within the haptic research community. This structured approach to device characterisation and performance evaluation enables researchers to compare different devices and replicate experiments more effectively.

The scripts, code, and analysis techniques in the framework have been tested on a limited number of setups. They serve as a baseline from which researchers can adjust and modify to

meet the specific needs of their devices and experimental designs. The framework provides a template for a standardised characterisation process, adaptable to the distinct requirements of each haptic device. The subsequent section explores the technical details of the data acquisition system and its design, followed by an in-depth discussion of the load cells and testbeds used in the characterisation process.

4 FARE-KE Data Acquisition Board




Figure 59. The FARE-KE custom data acquisition (DAQ) board with 3 MCP3564 24-bit ADCs.

Taking into consideration insights from the literature review, a custom DAQ system was designed to capture data from load cell sensors. The system accommodates up to 12 load cells read simultaneously at up to 10ksps. Although primarily intended for load cell data capture, the DAQ (an instrument amplifier combined with an analog-to-digital converter (ADC)) can capture other types of analog signals irrespective of sensor type. This section examines the DAQ design, highlighting its role in enabling high-fidelity haptic data acquisition aligned with the characterisation framework's demands.

4.1 Proposed design

Based on hands-on experience and testing, the DAQ system design in Figure 59,60 was developed. Constraints and recommendations from the literature review informed this design as essential criteria. This methodology ensured the DAQ system met required ranges and frequencies for haptic applications and facilitated integration into existing workflows. The main design constraints are:

Refresh Rate: A minimum rate of 250Hz is recommended for teleoperation and admittance devices [150]. For haptic controllers delivering impact forces, higher rates are preferable. The

design targets a minimum data capture rate of 1kHz to facilitate digital oversampling techniques, enhancing resolution and data quality.

Resolution: Higher ADC resolutions enable more precise data capture. While commercial boards like the HX711 offer 24-bit resolution, this design aims to deliver similar performance by incorporating high-resolution ADCs. The goal maintains medium to high resolution, utilising at least 16 out of 24 bits effectively, as noise often affects the lowest bits.

Gain Flexibility: Adjusting gain to match expected force outputs optimises ADC utilisation across its full resolution. The design includes adjustable gain settings to maximise effective resolution and minimise quantisation errors.

Accuracy and Consistency: Measurement accuracy depends on stability, noise rejection, and thermal offsets. These factors were considered to maintain the highest number of resolution bits and ensure consistency across varying ambient temperatures.

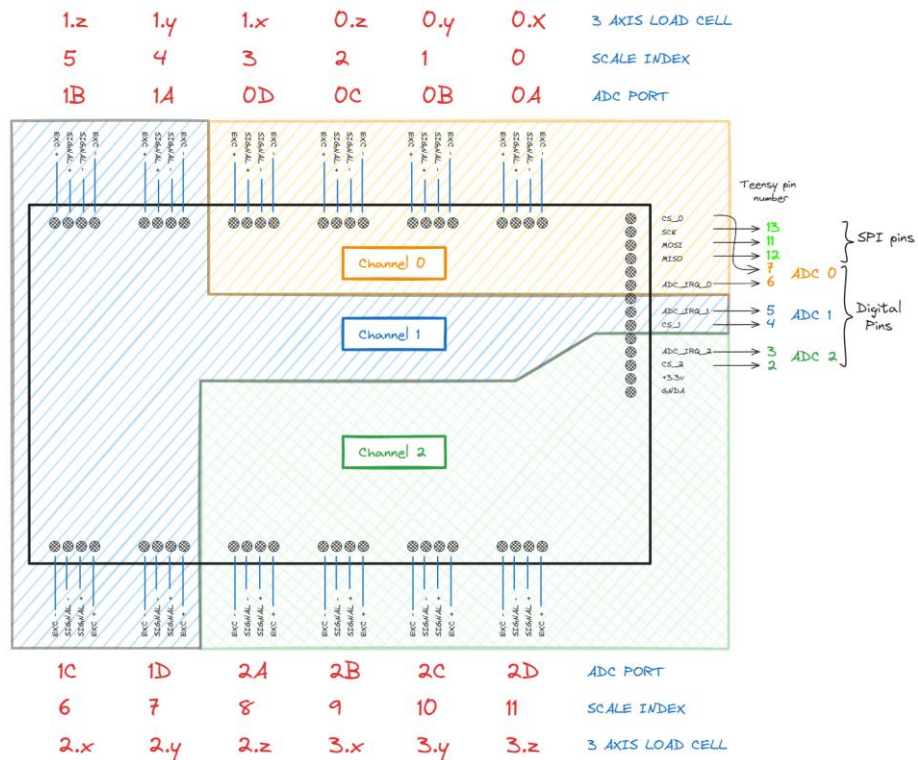


Figure 60. Pi Pinout connection diagram for the FARE-KE DAQ board, showing the mapping between MCP3564 ADC channels, device ports (0A-2D), scale indices (0-11), and 3-axis load cell connections (X,Y,Z axes), with color-coded channel groupings to facilitate proper wiring and software configuration.

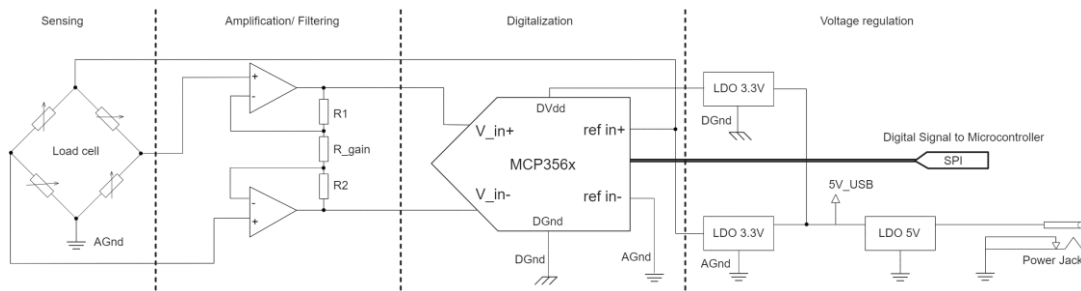


Figure 61. Simplified schematic of the FARE-KE DAQ system showing four main functional blocks: sensing, amplification/filtering digitalisation, and voltage regulation with SPI connection to microcontroller.

USB Power: The design operates on standard USB power, eliminating additional power supply components. It consumes less than 500mA, compatible with the maximum power output of a standard USB 2.0 port.

Cost: The objective develops an affordable DAQ system priced below \$150, enhancing accessibility. This cost excludes load cells, which vary based on specific requirements.

Throughout this research, numerous instrumentation amplifiers were tested, starting on commercial options like the HX711. However, limited SPS lead to exploring analog options. Initial experiments utilised instrumentation amplifiers including the AD623, INA129, and AD8237. While effective for a limited number of channels, microcontroller built-in ADCs typically offer limited resolution (up to 16 bits) and ADC channels. Additionally, the cost of individual instrumentation amplifiers can significantly increase the overall board cost compared to alternative approaches.

The following sections detail the principal components of the proposed Data Acquisition (DAQ) system (Figure 61), examining their roles and selection considerations. These components include voltage regulation, amplification, and analog-to-digital conversion. The sensing component will be expanded in consequent sections covering various types based on testbed requirements.

Voltage Regulation

Load cells operate using an excitation voltage that enables force measurement through changes in electrical resistance within strain gauges. If the excitation voltage contains noise, it directly affects signal readings. Therefore, the quality and stability of power supplies, particularly the excitation voltage provided to load cells, are critical. Despite noise rejection

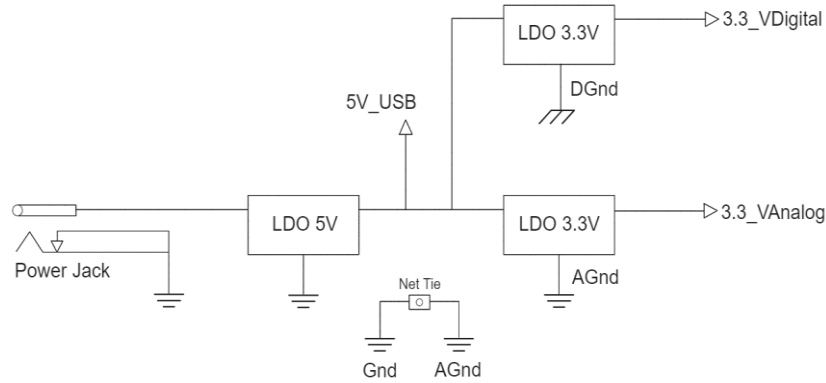


Figure 62. Power system configuration for the FARE-KE DAQ board featuring separate LDO regulators for digital (3.3V_VDigital) and analog (3.3V_VAnalog) supplies.

capabilities of amplifiers or ADCs, noisy excitation voltage invariably distorts measured output signals.

To mitigate this, the design incorporates multiple linear low-dropout regulators (LDOs) (Figure 62) and low-pass filters to provide stable, low-noise voltage. The initial power regulation stage utilises the MIC39100-5.0WS, an LDO delivering consistent 5V output from a 5-12VDC input. This design protects against input fluctuations and allows use of various unregulated voltage supplies up to 12V.

The 5V output from the MIC39100-5.0WS distributes to two separate LDOs, supplying distinct power sources for digital and analog components on the same PCB. The MCP1825S-330, known for low output voltage noise, generates a 3.3V output dedicated to the analog domain, powering components such as bridge excitation and amplification stages. Conversely, digital components receive power from the MCP1754TH-3302E, selected for higher output current capabilities and low dropout voltage, ensuring efficient power management. This supply voltage division serves several purposes:

- **Noise Isolation:** Digital circuit noise, particularly from rapid switching operations like clocks, remains isolated from sensitive analog signals, such as those from load cells.
- **Voltage Stability:** Analog components, critical for high-precision measurements, receive steady voltage, ensuring measurement integrity.
- **Interference Minimisation:** Separate power domains reduce electromagnetic interference (EMI) risk from digital circuits affecting analog signals.

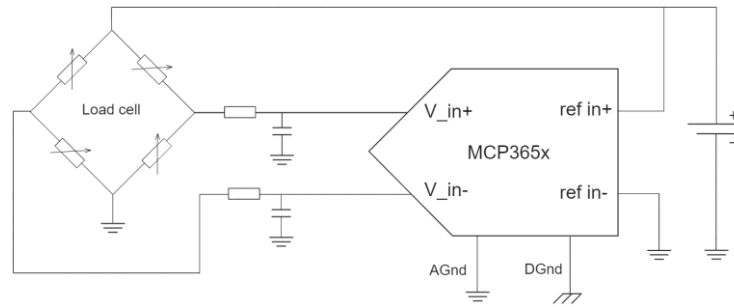


Figure 63. Basic circuit design between load cell and MCP3564 ADC showing basic component for digitalisation of the signal.

Analog to Digital Conversion

The proposed design incorporates 3 MCP3564 4-channel 24-bit analog-to-digital converter (ADC). This ADC sequentially reads up to four differential channels operating with an internal oscillator providing $\sim 5\text{MHz}$ clock, enabling sampling rates up to $\sim 12\text{kSPS}$ on a single channel. This ADC can be used without other components to amplify and digitise the signal (Figure 63).

Sampling rates can increase to 153.6 kSPS with an external 20MHz clock. The MCP3564 offers flexible oversampling ratios (OSRs) ranging from 32 to 98304, enhancing signal-to-noise ratio (SNR) and effective number of bits (ENOB) which reduces advance filtering requirements. Data communication occurs via SPI interface, handling digital readings post-filtering and conversion, alleviating microcontroller processing load. Additional advantages include resilience to low temperature drift with theoretical specifications below $1\ \mu\text{V}/^\circ\text{C}$ for offset and gain, minimising measurement errors from environmental fluctuations. This ADC also features programmable gain from $\times 0.33$ to $\times 64$, simplifying the circuit by potentially eliminating external instrumentation amplifiers and allowing system adjustment to load cell range and measuring range, optimising system sensitivity.

Amplification

While the MCP3564's integrated programmable gain amplifier offers initial signal amplification capabilities, incorporating a discrete op-amp-based amplification stage (Figure 64) provides advantages for precision load cell measurements. This additional stage offers more refined gain control than the fixed steps available in the MCP3564's internal amplifier,

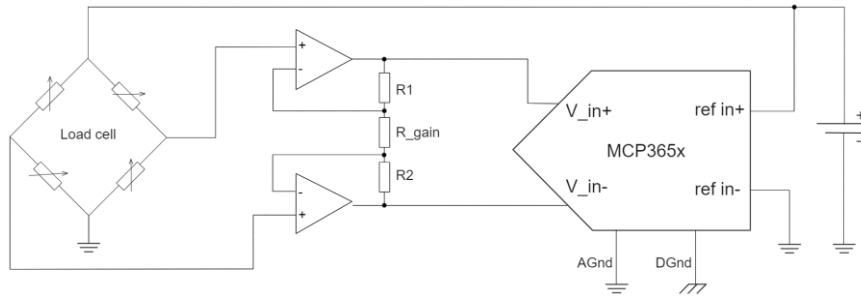


Figure 64. Detailed view of the signal chain for one channel of the FARE-KE DAQ system with load cell in Wheatstone bridge configuration, differential amplification using precision op-amps with programmable gain resistors, and direct connection to the MCP356x 24-bit ADC with separate analog and digital grounds.

allowing for optimal signal conditioning tailored to specific load cell characteristics. The following section details this amplification approach and its benefits for high-precision force measurement. The proposed design utilises a buffer instrumentation amplifier configuration (Figure 64 **Error! Reference source not found.**) using MCP6V97 op-amps. The amplification stage optimises the necessary gain while minimising noise. Increasing the value of the gain resistor in an instrumentation amplifier will decrease the overall gain because the gain is inversely proportional to the value of the gain resistor. This relationship is described by the formula:

$$Gain = 1 + \frac{2R}{R_{gain}}$$

where R is the value of the resistors in the input buffer and R_{gain} is the resistor connected between the output of the two input buffer amplifiers. The MCP6V97 features low-offset voltage that make it well-suited for use with load cell sensors:

High Input Impedance: The MCP6V97's high input impedance minimises current draw from the strain gauge-based load cell sensor, preventing self-heating effects that could impact the accuracy of the readings. High input impedance functions similarly to a measurement device that does not affect the circuit being measured.

Low Output Impedance: Allows the amplifier to effectively drive the ADC input, ensuring that the amplified signal is delivered without significant loss or distortion.

High Common-Mode Rejection Ratio (CMRR): This property helps the amplifier ignore noise that appears on both inputs (for example, noise due to temperature fluctuations), resulting

in more accurate readings. CMRR represents the ability to filter unwanted common signals whilst amplifying differential signals.

Adjustable Gain: While the current design utilises a fixed gain, the MCP6V97 supports adjustable gain settings to accommodate different sensor types and measurement ranges. This adaptability allows for optimisation with various load cell sensitivities.

Precision Specifications: With a maximum offset of only 25 μV and a noise level as low as 0.24 μV , the MCP6V97 provides clean signal amplification. Its power supply rejection ratio (PSRR) of 110 dB means that even fluctuations in the power supply are ignored, ensuring that the minute voltage changes from load cells are amplified.

Incorporating a programmable-gain amplifier improves the ADC's functionality. This setup accommodates various load cells with different sensitivities, ensuring that the ADC is not saturated and can utilise its entire range. Unlike the internal amplifier in the MCP3564, which has a fixed programmable gain, adjusting the gain resistor on the MCP6V97 enables fine-tuning of the gain.

4.2 Reading Data from the DAQ - Microcontroller

The DAQ operates as a two-part system: a data acquisition component covered in the previous section and the digital processing stage handled by the microcontroller. These parts communicate via SPI. After data is amplified and digitised by the MCP3564, it is transmitted to the microcontroller for processing. The microcontroller performs key operations:

Data Synchronisation with ADCs: The microcontroller governs SPI, ensuring synchronised reading across multiple ADCs at consistent intervals. The microcontroller selects each ADC using chip select pins after an interrupt signals new values are ready.

Filtering: Acquired data is typically filtered using a software 10Hz low-pass Butterworth filter. For fast transient signals like impacts, this filter should be removed to capture the complete representation when measuring impact feedback.

Data Processing: After the ADC converts the load cell's signal to a digital value, the microcontroller applies calibration and compensation, computing the applied force magnitude.

Real-time Synchronisation with Haptic Device: For characterisation tests, it synchronises ADC readings with haptic actuator control, facilitating automated characterisation.

4.2.1 Microcontroller Selection Criteria

The requirements for processing load cell data at high speeds with multiple channels influenced the microcontroller selection. We focused on 32-bit MCUs capable of handling SPI communication with the MCP3564 at full speed while computing forces or torques in real-time. This architecture provides several critical advantages:

- **High-Speed SPI Communication:** The MCP3564 operates with SPI at 20 MHz. Older 8-bit microcontrollers like the Atmega328P support a maximum SPI speed of only 8 MHz, creating potential communication bottlenecks with the proposed ADCs design.
- **Real-time Processing Capability:** Computing forces and torques from multiple load cells, particularly when implementing matrix operations for cross-talk compensation, benefits from faster processing capabilities that 32-bit architecture provides.
- **Efficient Data Transmission:** Modern 32-bit boards deliver significantly faster serial communication between the microcontroller and computer, essential when streaming data from multiple channels simultaneously.

We selected the Teensy 4.0 for its optimal balance of performance and affordability. With a 600 MHz clock speed (compared to 40-80 MHz on typical 32-bit boards like the ESP32), it provides ample processing power at approximately £25 (versus higher-end options like the Arduino Portenta at £100). Its native USB speeds of up to 480 Mbps enable efficient data transfer, and it offers sufficient memory for complex calculations and data buffering. The Teensy 4.0's processing capabilities make it particularly suitable for the matrix manipulations required to convert raw data from multiple load cell channels into real-time force/torque metrics. When selecting alternative boards, verifying supported USB and SPI speeds remains essential to guarantee adequate performance. Suitable alternatives include the ESP32, Arduino Due, or Adafruit Feather M0. The main PCB design incorporates only the amplifying and digitisation stages, with a pinout output to connect to the MCU. This design choice future-proofs the system, as development boards frequently reach end-of-life and cease production.

4.2.2 Library Implementation

To simplify interaction with the DAQ system, an Arduino library with three abstraction levels was developed:

- **MCP356x Scale:** The lowest abstraction level, delivering basic load cell functionalities including weight measurement in grams-force and taring.
- **MCP356x 3-axis:** An intermediate level that merges channel readings to output three-dimensional force calculations with calibration and cross-torque compensation techniques.
- **MCP356x 6-axis:** The highest abstraction level, processing data from all 12 channels to calculate forces and torques across x, y, and z axes.

The library incorporates functions for data filtering, oversampling, and signal processing techniques. For 3-axis load cell configurations, three sequential channels are combined to represent a single 3-axis load cell reading. These associations are predefined in the firmware, available in a GitHub repository along with comprehensive documentation and example sketches.

4.3 Performance Metrics

Benchmarking the DAQ's capabilities was part of the design process. This section outlines metrics significant to the performance and configuration profiles of the DAQ, considering the impact of the MCP3564's configuration settings. The Oversampling Rate (OSR) mitigates noise while influencing resolution and sample acquisition rate. This examination of the system's operational dynamics offers insights into its performance capabilities. During testing, the board was powered via the jack port with 5V, resulting in an average current consumption of 26mA when no load cells were connected. Each connected load cell adds approximately 4.2mA to the circuit's load, totalling around 76.4mA when all 12 channels are scanned simultaneously. This power consumption makes the design suitable for USB power supply, as it consumes less than the average 500mA provided by a standard computer USB port.

Oversampling Rate

Increasing the OSR improves signal quality, as evidenced by the reduction in standard deviation, from 2.99 gf with 32 samples averaged to 1.41 gf with 4096 samples. However, an OSR of 4096 limits the capture rate to 220 samples per second. To maintain a 1 kSPS sampling

OSR Setting	Sampling kSPS	Mean ADC Reading	Standard Deviation	SNR (dB)	SD (gf)	95% CI (gf)
32	11.59	-81082.92	4921.15	24.34	2.99	±5.98
64	5.79	-65060.29	4247.45	23.70	2.59	±5.17
128	2.90	-53700.42	3714.84	23.20	2.26	±4.53
256	1.45	-46071.12	3211.98	23.13	1.96	±3.91
512	0.72	-44473.23	2864.04	23.82	1.74	±3.49
1024	0.54	-44384.22	2707.46	24.29	1.65	±3.30
2048	0.36	-44451.76	2489.31	25.04	1.52	±3.03
4096	0.22	-44584.01	2318.21	25.68	1.41	±2.82

Table 8. The MCP3564 can internally provide an OSR from 32 to 98304. A higher OSR increases the number of samples taken and averaged for each output data point, filtering out noise but also reducing the data rate. To evaluate the OSR's impact on DAQ performance, digital raw readings were collected from the ADC connected to a beam load cell at various OSR settings down to 220 SPS.

rate, a maximum OSR of 256 is recommended. Figure 65 illustrates the OSR effects, displaying signals using raw digital readings. The signals are layered corresponding to OSR levels, from lowest to highest. As load cells have an inherent offset from zero, the data was adjusted around the 0-axis based on the mean value of the dataset. The signal with the lowest OSR (32) appears longest on the graph, while the highest OSR (4096) appears shortest. Visual inspection shows higher OSR values yield more consistent readings close to the expected zero value.

Signal-to-Noise Ratio (SNR) analysis was performed using the OSR test data, comparing signal quality relative to noise level. The SNR results align with the standard deviation findings, showing improvement from 24.34 dB at an OSR of 32 to 25.68 dB at an OSR of 4096. However, this improvement is marginal considering the load cell's calibrated reading




Figure 65. Visual comparison of ADC signal noise for different oversampling ratio (OSR) settings of the MCP3564 DAQ board, showing eight colour-coded noise bands with corresponding SNR values (24.34 to 25.68 dB), demonstrating how higher OSR values reduce noise bandwidth but maintain similar peak-to-peak amplitude.

range of ± 5000 gf. While signal quality improves with increased oversampling, the practical significance might be limited, especially for applications where the force range spans the full load cell capacity (± 5000 gf). This underscores the importance of researchers evaluating trade-offs to determine appropriate force range, load cell, and oversampling rate for their specific application requirements.

Sampling data for multiple channel

OSR Setting	Single Channel (kSPS)	Four Channels (kSPS)
32	11.59	2.8975
64	5.79	1.4475
128	2.90	0.725
256	1.45	0.3625
512	0.72	0.18
1024	0.54	0.135
2048	0.36	0.09
4096	0.22	0.055

Table 9. Sampling rate comparison at different oversampling ratio (OSR) settings for the FARE-KE DAQ board, showing exponential decrease from 11.59 kSPS to 0.22 kSPS for single-channel operation and from 2.8975 kSPS to 0.055 kSPS when using four channels simultaneously.

As sampling rate decreases with OSR increase, the number of channels sampled also affects the sampling capabilities proportionally. The MCP3564's ability to capture a single signal at up to 11.59 kSPS using the internal clock was validated. When configured in scan mode to read multiple channels, the ADC cycles through the specified channels, sampling each in turn. The effective data rate (EDR) divides by the number of channels read. To explore multichannel reading limitations, sample numbers were recorded over approximately 100 seconds while sampling all differential channels. To maintain a 1 kSPS rate while scanning four differential channels, the maximum OSR setting should be 64. For slower transient forces where a 250 SPS rate per channel suffices, an OSR of up to 256 can be used. However, as the board includes multiple ADCs, it can scan data from three individual channels at 11.56 kSPS with an OSR of 256 readings.

Force Reading

To illustrate performance differences between a low sampling rate and a high-speed ADC, the HX711 and MCP3564 were compared using identical beam load cells (model TAL220, 10kg). Each load cell connected to a different ADC, with readings taken simultaneously via interrupt triggers. Both load cells were mechanically coupled to a stepper motor with a lever, applying

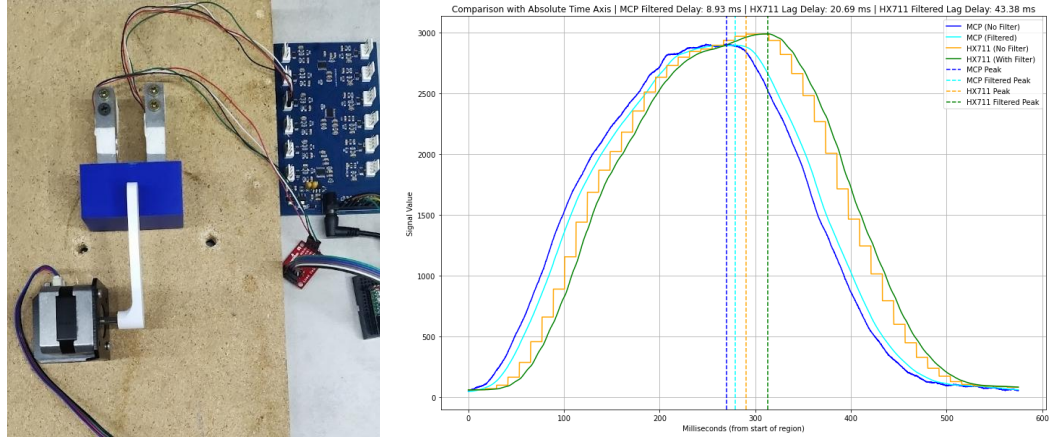


Figure 66. Comparison of MCP3564 and HX711 ADC performance: (a) characterization testbed with beam load cell connected to both data acquisition systems and solenoid actuator, and (b) response curves showing lower latency (8.93 ms) for MCP3564 compared to HX711 (20.69 ms raw, 43.38 ms filtered) during impact force measurement.

force evenly during actuation. A total of 20 actuations were performed to average the readings from both load cells

The average force measurement lasted approximately 600 ms. During this time, the MCP3564 captured an average of ~3,300 distinct data points, with each reading taking about 90 microseconds. In contrast, the HX711 required 13 milliseconds to process each data point, as shown in Figure 66, where the raw data exhibits a stepwise pattern. On average, the MCP3564 took over 120 samples per reading of the HX711. Both load cells' data underwent a 10Hz low-pass Butterworth filter to reduce noise. However, filtering introduced two disadvantages: it smoothed out data, potentially removing significant signal variations, and added latency between raw and filtered data. In visual comparison, raw data from the MCP3564 led the HX711 by ~20ms. When considering filtered data, the time lag between signals from both ADCs extended to approximately 33.6ms.

Fast response in data capture is required for accurately quantifying system latency, since the human body can discern force impulses that are separated by as little as 5 ms [1] the system should ideally account for that as baseline. While the sampling rate is not so critical in force

Signal Type	Average Lag Delay (ms)
MCP_no_filter	-
mcp_filtered_data	8.9
HX_no_filter	20.64
HX_with_filter	42.50

Table 10. Lag delay comparison between MCP3564 and HX711 ADCs, showing MCP3564 filtered data with only 8.9 ms delay compared to 20.64 ms for unfiltered HX711 and 42.50 ms for filtered HX711, all measured relative to unfiltered MCP3564 signal used as reference baseline.

measurement, it becomes even more critical when dealing with fast transient impact stimuli, where the force duration may be only a few milliseconds.

Impact Reading

Sampling rate is critical for kinesthetic devices delivering fast transient stimuli, such as actuators with elastic mechanisms. If the stimulus exceeds the ADC's capture rate, the recorded data may deviate from the original event. For this test, the testbed was modified to include a stepper motor raising a M12 threaded bar 160mm above the load cells before releasing it to impact the coupling mechanism.

In these tests, the same filter as in previous experiments was used to illustrate filtering implications on fast transient signals. The readings and filtered data are shown in Figure 67. Here, the MCP no filter signal (in blue) captures a short, sharp transient peak at the moment the M12 metal rod impacts, as all system parts are rigid. The peak, lasting about 1 millisecond, depicts the initial collision, followed by signal flattening for approximately 50 milliseconds until the load cell stabilises, then returns to zero at around 210 milliseconds. This plateau and subsequent decay represent the energy dissipation from the impact.

The MCP filtered data and the HX711 present a different view of the impact. While they capture the general trend (rise and fall), the exact magnitude and timing of the peak are lost. Filters smooth out sharp transients to reduce noise, which resulted in the most important part of the signal being lost. This demonstrates that for fast transient events, filtering distorts




Figure 67. Left: Experimental setup. Right: Graph showing unfiltered and filtered ADC data from MCP3564 and HX711 during the test, illustrating the sharp transient peak captured by MCP_no_filter and the smoothed responses due to filtering similar to that of the HX711.

Signal Type	Average Total Impulse (Ns)	Average Force (in Newtons)
MCP no filter	~0.39	~1.56
mcp filtered data	~0.39	~1.56
HX no filter	~0.345	~1.38
HX with filter	~0.335	~1.34

Table 11. Comparison of impulse and force measurements between MCP3564 and HX711 ADCs, showing MCP3564 capturing higher values (0.39 Ns impulse, 1.56 N peak force) compared to HX711 (0.345-0.335 Ns impulse, 1.38-1.34 N peak force), with filtering having minimal effect on MCP3564 measurements but slightly reducing HX711 values.

data by removing critical information. The MCP no filter offers the most accurate depiction of the impact event, capturing the rapid transient peak and subsequent mechanical oscillations.

To contextualise this, the area under the curve for the four signals was analysed. The MCP no filter and MCP filtered data show nearly identical Average Total Impulse and Average Force values. The MCP signal retains the integral properties of the original event, despite peak smoothing. The HX711 (both filtered and unfiltered) displays a reduction in both impulse and average force compared to the MCP signals.

While the MCP signals provide closely aligned readings, the HX711 demonstrates a deviation and poor representation in magnitude and timing. The decision to use filtered versus unfiltered data should be guided by specific application needs and the importance of capturing fast transient peaks versus achieving a smoother signal.

Excitation Voltage Noise

When capturing data from load cells, maintaining consistent and stable excitation voltage directly influences noise and quality of the output amplified signal. To assess the MCP3564 circuit performance, the excitation voltage from the board design was compared with that generated by a HX711 prototyping board. For this comparison, data from both ADC chips' AVCC pins were captured using a digital oscilloscope.

Parameter	HX711 ADC (V)	MCP3564 ADC (V)
Mean voltage	4.2913	3.2840
Standard deviation	0.0140	0.0133
Range (Max - Min)	0.3119	0.2533

Table 12. Excitation voltage stability comparison between HX711 and MCP3564 ADCs, showing slightly better performance from MCP3564 with lower standard deviation (0.0133 V vs 0.0140 V) and smaller voltage range (0.2533 V vs 0.3119 V), despite different mean voltage levels (3.2840 V vs 4.2913 V).




Figure 68 Voltage readings over time for the HX711 (blue) and MCP3564 (green) ADC boards.

The first difference is that the HX711 operates at 5V, whereas the MCP3564 uses 3.3V. Typically, higher excitation voltages ranging from 10 to 15V are preferable for load cells. Higher voltages reduce current through the strain gauges, minimising self-heating and diminishing noise impact on the output signal.

Our tests revealed that both ADC circuits exhibit similar interference levels in the frequency domain, with both designs demonstrating comparable performance. However, the MCP3564 showed a slightly smaller standard deviation, likely due to its operation within a lower voltage range. Both PCBs employ similar LDO technology and hardware filtering techniques, which points to an opportunity for further improvements in future PCB iterations to decrease noise on the DC excitation voltage. Based on the excitation voltage noise and the MCP3564 operating range, this noise could potentially affect the Effective Number of Bits (ENOB) of the force signal, reducing its resolution by approximately 7 bits.

Conclusion

The Data Acquisition System (DAQ) consumes less than 200 mA while sampling 12 channels, making it suitable for USB power. With no external oscillator clock, the ADC captures data at 11.59 kSPS on a single channel per ADC with an oversampling ratio (OSR) of 32. When multiple ADCs read 12 channels simultaneously, this configuration limits the sampling rate and OSR to 64, achieving ~ 1.45 kSPS. Future development plans include incorporating an

OSR	Single Channel (Internal Clock at ~5Mhz)	Four Channels (Internal Clock at ~5Mhz)	Expected rate for Four Channels using External Clock at 20Mhz
32	11.59	2.8975	38.4
64	5.79	1.4475	19.2
128	2.90	0.725	9.6
256	1.45	0.3625	4.8
512	0.72	0.18	2.4
1024	0.54	0.135	1.8
2048	0.36	0.09	1.2
4096	0.22	0.055	0.6

Table 13. Projected sampling rate comparison for the FARE-KE DAQ board across different configurations, showing current rates with internal 5 MHz clock for single channel (11.59 to 0.22 kSPS) and four channels (2.8975 to 0.055 kSPS), alongside theoretical rates using 20 MHz external clock which would yield four-channel throughput from 38.4 to 0.6 kSPS.

external oscillator to increase ADC speed and sampling rate, potentially allowing 2048 OSR across 12 channels while maintaining sampling above 1 kSPS.

Increasing OSR reduces signal noise but lowers sampling rate. This trade-off becomes critical when assessing fast transient responses, where lower OSR better captures rapid signal dynamics. The impact of noise relates to the force magnitude being measured. A load cell with ± 5000 grams-force (gf) range showed noise standard deviation (SD) of 2.9 gf. For devices without sudden forces, high OSR of 2048 while scanning 12 channels may provide sufficient speed to capture necessary details, achieving performance comparable to the HX711 ADC. This research provides guidelines to help researchers choose optimal settings for specific applications, balancing noise reduction and sampling speed based on their experimental requirements.

Although the excitation voltage circuit performs similarly to the HX711, future PCB iterations will explore increasing excitation voltage or adding filtering layers to minimise noise on analog excitation voltage, improving resolution bit count.

Further design revisions will consider integrating a precision trimmer potentiometer instead of fixed-value resistors to fine-tune amplifier circuitry gain. This adjustment would match amplification to the load cell's range with expected applied force, optimising the ADC's 24-bit resolution use and reducing noise impact on lower bits. Enhancing the ADC's signal-to-noise ratio remains a focus on the future development of the DAQ system.

5 FARE-KE Sensor/Hardware Validation

This chapter examines the various load cells employed in the characterisation process, focusing on their selection, validation, calibration, and application. Load cell selection directly correlates with the type of kinesthetic haptic feedback a device produces, as established in the literature review. The following sections categorise load cells based on the specific force types they measure, progressing from simple to complex measurements.

The simplest form is **direct single-axis force**, where devices apply force directly to the user's body with a well-defined vector component. These forces are measured using beam load cells, which offer versatility and cost-effectiveness for numerous applications. This category includes both continuous and momentary force applications.

Single-axis impact represents the second category, involving devices that generate sudden energy bursts through quick release mechanisms (e.g., elastic bands or pneumatic actuation). These rapid force changes require sensors with high sampling capabilities to accurately capture the force profile.

In-line forces, commonly found in exoskeletons employing wire and capstan mechanisms, require precise measurement using S-type and button load cells. These sensors can measure both tensile and compressive forces, offering mounting fixtures at both ends for inline configurations.

Moving beyond single-axis forces, **rotational joint torque** occurs in devices incorporating rotational components, typically identified by joints, levers, or mechanical transmission methods. Systems actuated across a single Cartesian plane can be effectively characterised using strategically placed single beam load cells.

Finally, for devices exerting **multi-axis forces or complex torques**, decomposing these to characterise individual actuators becomes essential. Force-torque sensors simultaneously measuring linear forces and rotational torques across three axes provide the most comprehensive measurement approach. This includes devices like high-speed flywheels generating both linear forces and torques. Since multi-axis force-torque sensors are not widely

available at budget prices, this chapter presents an alternative approach to building a 6-axis F/T sensor.

Multiple factors contribute to device characterisation beyond sensor measurements, including ergonomics (device handling, hand plasticity, and skin mechanics) and other elements shaping haptic perception. While acknowledging these factors' importance for complete characterisation, the current approach focuses on creating a universally accessible methodology that aligns with existing practices while addressing known challenges. The validation methodology progresses systematically through each category, beginning with the simplest single-axis measurements and culminating in complex multi-axis force-torque sensing. For each category, the chapter presents sensor validation procedures, practical setup considerations, and application examples from current haptic research.

5.1 Direct Single-Axis Force

In direct single-axis setups, the haptic device applies a direct push or pull force on the load cell, mimicking the kinesthetic feedback experienced by the user (Figure 69). PropelWalker [110], a wearable device with ducted fans attached to the user's calf, exemplifies this by simulating push and pull forces during walking in a VR environment. The ω -jet design utilises pneumatic nozzles on an exoglove to generate forces on the fingers.

Single-axis force characterisation represents the most direct kinesthetic modality. The beam load cell is fixed on one end, while the actuator is attached to apply force to the other end. Prior to examining the characterisation process in detail, this section validates three low-cost load cells that could characterise such devices. Budget beam load cells appear frequently in previous literature, but limited information exists regarding their accuracy or performance. Therefore, the initial step in the framework assesses the appropriateness of budget beam load cells for the intended application.




Figure 69. Examples of haptic devices suitable for single-axis load cell characterization: (a) tactile feedback glove with bend sensors embedded in tubes on fingers for directional forces, and (b) walking-based haptic feedback device (PropelWalker) showing direction of drag force applied to user's leg during locomotion.

5.1.1 Sensor Validation

To evaluate the performance of low-cost load cells and their suitability for measuring haptic device properties, three commonly used load cells identified from prior research were analysed: Degraw 5kg, TAL220, and CZL635. A comprehensive validation of these low-cost beam load cells was conducted under the following conditions:

1. **Linearity and Deviation:** This property ensures the load cell's output is directly proportional to the applied force. Non-linearity introduces additional complexity to calibration and could lead to errors, particularly in precision applications.
2. **Repeatability:** Ensuring the load cell consistently provides the same reading over time is essential for reliability. If readings under identical conditions vary, the measurements are unreliable, and calibration must be performed regularly.
3. **Drift:** Over extended durations, the load cell's output may change even if the applied load remains constant. Drift is particularly important in applications where loads are maintained for prolonged periods, as factors such as environmental temperature or current flowing along the strain gauge can significantly impact results.

Testing Methodology:

The same DAQ was used with the 3 load cell during testing, eliminating amplifier-induced variability. Calibrated laboratory weights were employed for testing and calibration, with beam load cells mounted horizontally and loaded by suspending weights from one end. When used in this orientation (e.g., in a weighing scale), the beam load cell measures force exerted due to gravity. This force can be converted to newtons by first dividing the mass in grams by

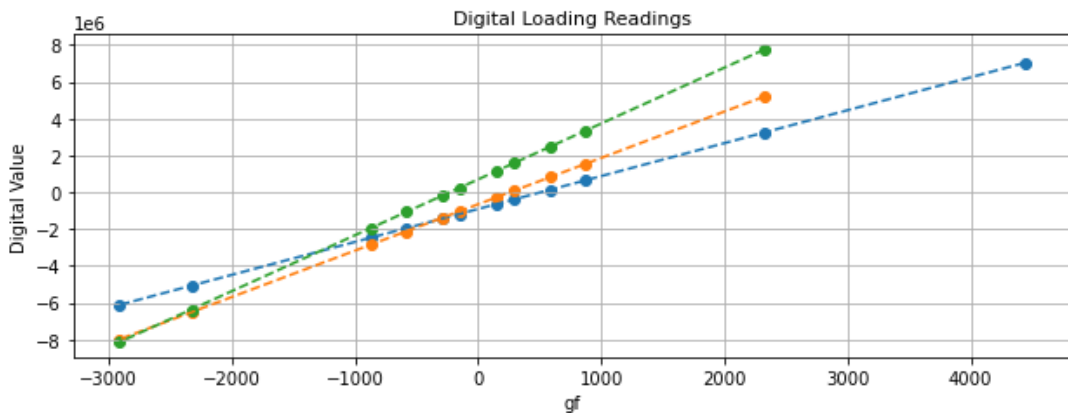


Figure 70. Digital load cell calibration data showing linear relationship between known weight (in gram-force) and ADC digital values for three different load cells (blue, orange, green), with each demonstrating consistent response across both compression and tension forces between -3000 gf and +4000 gf.

1000 to obtain kilograms and then multiplying by the gravitational acceleration constant (9.81 m/s^2), or it can be presented directly as gram-force. To account for power-up temperature fluctuations, the beam load cells were connected to the DAQ and stabilised for over 4 hours before data collection. System temperature was monitored throughout testing to assess potential effects on measurements.

To calibrate and determine linearity and potential deviation, readings were recorded during a stepwise weight loading procedure. The process began with a 100g weight and incrementally increased to 150g, 600g, 900g, 2300g, and 4500g. Since beam load cells frequently measure both compressive and tensile forces, data collection included loading in both directions by rotating the load cell body 180 degrees.

Slope (Sensitivity) & Intercept Analysis

Load	Capacity (kg)	Slope	Intercept	Operation Range (gf)	Sensitivity (mgf)
TAL220	10	1787.65	1306.72	+/- 4692.5	0.56
Degraw	5	2517.44	2180.97	+/- 3332.2	0.4
CZL635	5	3033.73	3678.78	+/- 2765.1	0.33

Table 14. Performance characteristics of three beam load cells tested in the FARE-KE framework, comparing TAL220 (10 kg capacity), Degraw (5 kg), and CZL635 (5 kg) models with their respective sensitivities (1787.65, 2517.44, and 3033.73 counts/gf), operational ranges, and resolution (0.56, 0.4, and 0.33 mgf).

Linear regression analysis of the three beam load cells revealed a significant linear response, with slope indicating the rate of change in digital reading per unit force change. The slope, determined by the load cell's mechanical properties, is steeper for load cells designed for smaller forces (Degraw or CZL635, rated for 5 kg) compared to those for higher capacities (TAL220, rated for 10 kg). The intercept represents digital output at zero force, corresponding to zero-load offset. The slope also determines sensitivity, defined as the minimum detectable weight increment based on digital resolution.

Noise can corrupt the least significant data bits, as demonstrated in the oversampling rate testing, limiting practical resolution. The operational range, defined as maximum and minimum measurable weights without ADC saturation, is determined by the slope and fixed gain of $\times 101$ (2 k Ω resistors) in the PCB design. Future designs incorporating a precision trimmer resistor will allow gain adjustment to match load cell capacity, enhancing system adaptability to various load cell specifications. In the current setup, load cells operate at approximately 50% of their rated capacity (e.g., CZL635, rated for 5 kg, limited to ± 2765 g) to ensure measurements remain within the linear response range and prevent ADC saturation.

Zero offset Compensation

Before calibrating the load cell, the zero offset must be considered (the digital reading when no load is applied). As shown in Figure 70, none of the three tested load cells intersect the digital 0 value. This common side effect in budget load cells arises from several factors:

- **Manufacturing Irregularities:** Visual inspection of low-cost beam load cells reveals limited quality in mechanical machining, particularly in the middle hole of the beam cell body.
- **Strain Gauge Resistance Tolerance:** The output reading assumes all strain gauges forming the Wheatstone bridge have identical resistance and tolerance. Low-cost strain gauges may have lower quality than industrial-grade products, with higher tolerance percentages.
- **Strain Gauge Misalignment:** alignment of all the strain gauges ensures that when the load cell deforms it affects all the strain gauges equally when loaded. Misalignment contributes to zero offset issues.

This offset is intrinsic to the load cell orientation, and changing the up direction affects the readings (Table 15). To determine the best approach for addressing zero offset during calibration, all three beam load cells were tested to characterise their zero-offset properties.

Weight	TAL220	Degraw	CZL635
0 (up)	-897071	-622022	664264
0 (down)	-931854	-696378	715609

Table 15. Zero-load offset values for three beam load cells (TAL220, Degraw, CZL635) in different orientations, showing significant differences between upward and downward mounting positions that must be accounted for during calibration.

Given zero offset values for two orientations are Z_{up} and Z_{down} , the mean zero offset \bar{Z} is computed as:

$$\bar{Z} = \frac{Z_{up} + Z_{down}}{2}$$

This mean zero offset is used to calculate the percentage offset relative to the ADC's total range R_{ADC} , determined by its resolution. Assuming a 24-bit ADC, $R_{ADC} = 2^{24}$. The percentage offset P_{offset} is given by:

$$P_{offset} = \left(\frac{\bar{Z}}{R_{ADC}} \right) \times 100$$

This formula provides the relative percentage of zero offset compared to the total ADC range (Table 16). The observed percentage offsets for TAL220, Degraw, and CZL635 highlight the inherent variability and impact of manufacturing differences on budget load cell precision.

Load Cell	Average Reading	Percentage Offset from Max
TAL220	-914462.5	-10.90%
Degraw	-659200	-7.86%
CZL635	689936.5	8.22%

Table 16. Average zero-load offset values for three beam load cells, showing TAL220 with -914462.5 counts (-10.90% of full scale), Degraw with -659200 counts (-7.86%), and CZL635 with 689936.5 counts (8.22%), highlighting the importance of offset compensation during calibration.

Load Cell Calibration Approach

The data obtained during loading was used to determine the best method for processing the zero offset point when calibrating the load cells, aiming to obtain a better fit model. The zero offset data was processed using three techniques:

- **Zero Offset:** Initially zeroing or "taring" the load cell in each direction set the initial state to zero as a reference point. Subsequent measurements were taken as deviations from this zero state in both positive and negative directions.
- **Midpoint Zero Adjustment:** Using the mean value between the two readings for the zero value.




Figure 71. Comparison of three calibration methods for beam load cells showing linear relationships between applied weight (-3000 to 4000 gf) and digital readings: (a) Zero Offset approach adjusting for unloaded sensor values, (b) Midpoint Zero Adjustment using average of positive and negative values, and (c) Ignoring Zero Readings method excluding zero points from regression, all yielding highly linear responses ($R^2 > 0.999$) for TAL220, Degraw, and CZL635 sensors.

- **Ignoring Zero Readings:** Performing calibration using values in both directions while omitting both zero value readings.

The results from models generated with linear regression were compared to determine which approach yielded the most accurate and reliable calibration method. Data was processed based on these three conditions, and linear regression analysis conducted for each dataset. By comparing performance metrics, the R^2 value was computed to determine which model best fit the data. The three load cells exhibited highly linear behaviour across the tested weight range, with R^2 values above 0.999% for all load cells and data methods (Table 17).

Calibration Method	Load Cell	Slope (Digital Value/gram)	Intercept	R-Squared	Max gf of Error
tal220	Zero Offset	1787.65	1306.72	0.999977	13.66
tal220	Ignoring Zero Readings	1787.64	-912936.97	0.999981	13.54
tal220	Midpoint Zero	1787.65	-913155.78	0.999981	13.54
Degraw	Zero Offset	2517.44	2180.97	0.999948	15.63
Degraw	Ignoring Zero Readings	2517.50	-656607.88	0.999969	11.94
Degraw	Midpoint Zero	2517.44	-657019.03	0.999969	11.78
CZL635	Zero Offset	3033.73	3678.78	0.999966	11.67
CZL635	Ignoring Zero Readings	3033.82	694308.79	0.999973	11.45
CZL635	Midpoint Zero	3033.73	693615.28	0.999973	11.45

Table 17. Comparison of three calibration methods for each beam load cell, showing consistent slopes (1787.65, 2517.44, and 3033.73 counts/gf) across all methods with excellent linearity (R-squared > 0.9999) but varying intercept values, demonstrating that calibration method primarily affects offset compensation while maintaining similar maximum errors (11.45-15.63 gf).

Among the three data capture methods tested, no significant difference was observed. However, the "mid-point zero" and "ignoring zero readings" methods yielded almost identical R^2 metrics for all load cells. Based on this data, any method could be adopted without significant impact on measurement accuracy. Despite small differences, ignoring zero values during calibration proved most suitable, offering the best fit model and simplest methodology. This approach was employed for all further load cell calibrations.

Load Cell Calibration Approach

To minimise cumulative errors in data collection and calibration (particularly important as these load cells would later form a 3-axis load cell), three calibration techniques were explored:

1. **Single Scaling Factor Calibration:** A rudimentary approach using a single scaling factor based on the mean ratio between actual weight and load cell reading.
2. **Linear Regression Calibration:** Using linear regression models to determine the slope and intercept that best fit the observed data.

3. Polynomial Regression Calibration: Fitting a second-degree polynomial to capture nuanced or non-linear behaviours in the load cell.

The calibration data consisted of digital readings corresponding to known weights, used without additional processing or offset adjustments. A Python script computed the models and evaluated their fit to the calibration data, resulting in three calibration models for each load cell (Table 18).

Load Cell	Approach	Model
TAL220	Single Scaling Factor	$y = 0.0005583x$
	Linear Regression	$y = 0.0005594x + 510.8$
	Polynomial Regression	$y = -2.923e-13x^2 + 0.0005596x + 513.9$
Degraw	Single Scaling Factor	$y = 0.0003959x$
	Linear Regression	$y = 0.0003972x + 261$
	Polynomial Regression	$y = -1.143e-13x^2 + 0.0003968x + 261.8$
CLZ635	Single Scaling Factor	$y = 0.0003286x$
	Linear Regression	$y = 0.0003296x + -228.6$
	Polynomial Regression	$y = -5.538e-14x^2 + 0.0003295x + -227.8$

Table 18. Calibration equation comparison for three load cells using single scaling factor, linear regression, and polynomial regression approaches, showing minimal variation in primary conversion coefficients across methods.

The three load cells show minimal deviation for a linear behavior (Figure 72). The Polynomial Regression method demonstrated highest accuracy across the three tested beam load cells, evidenced by consistently lower root mean square error (RMSE) values and R^2 values close to above 0.99 (Table 19). Considering the marginal improvement offered by Polynomial Regression compared to other methods, this additional precision makes it the preferable choice, particularly given its computational feasibility on current 32-bit microcontrollers for real-time measurement and conversion.




Figure 72. Calibration model comparison for TAL220, Degraw, and CZL635 load cells showing nearly identical performance between single scaling factor, linear and polynomial regression approaches across the full measurement range (-3000 to +4000 g).

Calibration Method	RMS TAL220 (gf)	RMS Degraw (gf)	RMS CLZ635 (gf)	R ² TAL220	R ² Degraw	R ² CLZ635
Single Scaling Factor	8.98	10.27	8.61	0.99997	0.99994	0.99995
Linear Regression	8.27	9.17	7.37	0.99998	0.99995	0.99997
Polynomial Regression	6.98	9.00	7.27	0.99998	0.99995	0.99997

Table 19. Calibration method comparison for three load cells showing RMS errors and R² values, with polynomial regression yielding lowest errors (6.98-9.00 gf) followed by linear regression (7.37-9.17 gf) and single scaling factor (8.61-10.27 gf), though all methods achieve excellent linearity (R² > 0.9999).

Analysis of Load Cell Noise

Analysing noise effects in load cell readings helps understand the accuracy and reliability of measurements obtained during characterisation. This section examines noise present in readings from three load cells: TAL220, Degraw, and CZL635.

The load cell readings were captured under no-load conditions and converted from digital format to grams-force using a polynomial regression model derived from previous calibration analysis. Figure 73 illustrates the noise profiles of the three load cells under static conditions.

The amplitude of noise differs among the load cells, expected due to their varying sensitivity. Upon visual inspection, certain patterns appear predominant across the three designs, suggesting potential resonances or interference. To investigate further, additional noise analysis was performed. The noise distribution reveals that most readings for the three load cells exhibit a roughly Gaussian distribution (Figure 74), implying that noise values cluster




Figure 73. Noise profiles for three beam load cells under unloaded static conditions, showing TAL220 with highest noise amplitude (± 6 gf peak-to-peak), compared to Degraw and CZL635 (± 4 gf and ± 3 gf respectively), measured over a 200 μ s time window.




Figure 74. Noise distribution histograms for three load cells under no-load conditions, showing normal distributions with different offset ranges: CZL635 (-12 to -4 gf), Degraw (-20 to -10 gf), and TAL220 (-25 to -7.5 gf), indicating consistent noise characteristics with CZL635 exhibiting the narrowest distribution.

around the mean value. Gaussian distribution of noise does not inherently indicate periodicity or repetitive patterns; rather, it suggests that the noise, while varying in amplitude, is random without specific directional bias or repeating cycles.

Table 20 summarises key statistical parameters of noise for each load cell. The CZL635 load cell, with higher sensitivity, exhibits lower peak-to-peak noise (9.784 gf) and a standard deviation of 1.193 gf, indicating typical variability of just over 1 gram-force in individual readings. However, measurements range from -5.99 gf to 3.794 gf, providing a peak-to-peak noise range of 9.784 gf.

Load Cell	mean	SD (gf)	min	max	peak-to-peak	SNR (dB)
CZL635	-1.062	1.193	-5.99	3.794	9.784	45.3050
Degraw	1.009	1.226	-4.084	6.068	10.152	47.0741
TAL220	0.831	2.163	-7.815	9.29	17.105	47.7898

Table 20. Noise characteristics comparison for three load cells showing CZL635 with lowest standard deviation (1.193 gf) and peak-to-peak noise (9.784 gf), Degraw with moderate values (1.226 gf SD, 10.152 gf peak-to-peak), and TAL220 with highest variability (2.163 gf SD, 17.105 gf peak-to-peak), all maintaining good SNR (45-48 dB).

To identify the nature of noise present in the signal, a Fourier transform analysis was performed on the digital readings from the three load cells. As shown in Figure 75, all three load cells exhibit a pronounced peak in their frequency spectrum at around 5 MHz, corresponding to the frequency of the internal clock of the ADC (MCP3564). The presence of this peak suggests that the internal oscillator of the ADC could be a significant noise source in the system. Future work will focus on improving grounding, filtering, and utilising an external oscillator, which could increase the conversion rate by approximately four times and enhance signal quality by mitigating such noise sources.




Figure 75. Frequency spectrum analysis of three load cells showing similar noise characteristics with dominant peaks at approximately 10 kHz and smaller harmonic components at 20-40 kHz, with TAL220 exhibiting stronger high-frequency components compared to Degraw and CZL635. Peaks in their frequency spectrum at 0 removed.

Creep /Drift

Understanding creep or drift in load cell measurements is essential for assessing long-term stability and reliability of sensors used in device characterisation. This section examines measurement variability over a 24-hour period and explores how temperature fluctuations impact readings from three load cells: CZL635, Degraw, and TAL220 (Figure 76).

Experimental Setup

The three load cells were loaded with the same weight (2300g), and data was logged at four readings per second over a 24-hour period. Ambient temperature was monitored to assess correlation between temperature changes and load cell readings. Load cell values were converted to gram-force using a polynomial regression model specific to each load cell type.

The data demonstrates that behaviour of the three load cells is highly influenced by temperature fluctuations, with each load cell showing distinct drift characteristics. The




Figure 76. Temperature sensitivity analysis of three beam load cells, showing their drift response (± 12 grams) to ambient temperature fluctuations (19.7-23.3°C) over time, with CZL635 showing greatest stability (blue line, ± 2 gf), Degraw exhibiting moderate temperature-related drift (green line), and TAL220 displaying highest temperature dependency (purple line).

TAL220 load cell exhibited a maximum error of about 10 grams-force during the test period, showing greater deviation from the expected value compared to other load cells. However, when temperature stabilised to levels similar to those at the start of the experiment, the TAL220 readings most closely matched the initial values, suggesting a more linear temperature response and reduced creep.

Correlation Analysis

To determine the relationship between temperature and load cell readings, various correlation analyses were conducted, including Pearson Correlation for linear relationships, Spearman Correlation for monotonic relationships, and Kendall Tau Correlation for rank correlation. The findings, detailed in Table 21, reveal that the TAL220 load cell demonstrates a strong positive linear relationship with temperature, evidenced by high correlation coefficients in all three tests. Conversely, the Degraw load cell shows a strong negative linear and monotonic relationship, indicating a decrease in readings as temperature rises. The CZL635 load cell exhibits a moderate positive linear relationship with temperature but weaker monotonic relationships, suggesting the influence of non-linearities or other external factors.

Load Cell	Pearson Correlation	Spearman Correlation	Kendall Tau Correlation
635	0.5405	0.1159	0.0731
Degraw	-0.7826	-0.9268	-0.7895
TAL220	0.9302	0.9712	0.8804

Table 21. Temperature-drift correlation analysis for three load cells using multiple statistical measures, showing TAL220 with strongest temperature dependence (correlation coefficients 0.93-0.97), Degraw with strong negative correlation (-0.78 to -0.93), and CZL635 with least temperature sensitivity (coefficients 0.07-0.54).result indicates that the TAL220 has a strong linear relationship between temperature and readings.

Conclusion

All three load cells demonstrate sensitivity to temperature variations, though their responses differ. The 10gf deviation over 24 hours observed in the TAL220 corresponds to a 0.435% deviation from the original readings. The Degraw and TAL220 load cells show predictable linear and monotonic relationships with temperature, whereas the CZL635 load cell, despite being less affected by temperature, appears more susceptible to non-linear factors, likely due to its smaller range compared to the others. Understanding these correlations is essential for ensuring measurement reliability and accuracy, particularly in environments subject to temperature variations. Temperature compensation techniques or conducting calibration and

characterisation within a controlled environment can effectively minimise the impact of temperature drift on load cell readings.

Repetition Load Test

Evaluating precision and consistency of measurements across multiple actuations is critical for assessing the uniformity of response from kinesthetic haptic devices. This section examines the performance of three load cells (CZL635, Degraw, and TAL220) during repeated loading and unloading cycles, focusing on their capacity to deliver consistent and reliable force measurements.

Experimental Setup

Three load cells were subjected to a constant load of 2300 grams-force (gf), maintained for 10 minutes, then unloaded for another 10 minutes. The loading and unloading processes were conducted simultaneously for all three cells, excluding the initial and final 2.5% of readings during both phases to eliminate transient fluctuations and ensure stabilisation. Temperature remained stable throughout the test.

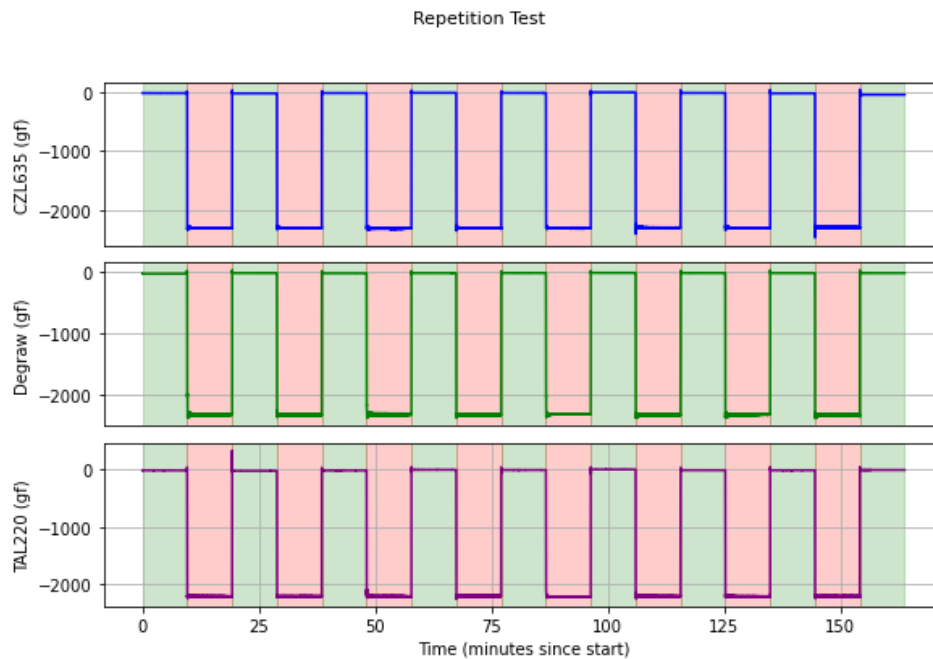


Figure 77. Repeatability test for three load cells over 8 loading/unloading cycles (160 minutes), showing consistent force responses of approximately -2300 gf when loaded (red background) and return to zero during unloaded periods (green background), demonstrating stable performance and minimal hysteresis for all sensors.

Results and Observations

Load Cell	Standard Deviation (Loading)	Standard Deviation (Unloading)
CLZ635	10.26 gf	3.93 gf
Degraw	4.06 gf	5.21 gf
TAL220	6.56 gf	2.36 gf

Table 22. Repeatability data comparison for three load cells, showing standard deviation during repeated loading and unloading cycles with Degraw offering best loading consistency (4.06 gf), TAL220 showing best unloading consistency (2.36 gf), and CZL635 exhibiting highest variability during loading (10.26 gf).

Figure 77 illustrates the loading and unloading cycles for the three load cells over 8 cycles of 20 minutes each. The digital readings were converted to gram-force using the calibration polynomial regression model specific to each load cell type. The CZL635 exhibits the highest standard deviation during the loading phase, suggesting potential challenges in scenarios demanding high precision. However, this load cell demonstrates significant improvement in consistency over the unloading phase. The Degraw load cell maintains relatively low standard deviation values for both loading and unloading phases, indicating balanced performance. The TAL220 demonstrates moderate variability during the loading phase but provides more consistent performance over the unload events.

The CZL635's plot (Figure 78) displays a notable peak in the last event, likely an outlier, contributing to its higher standard deviation during the loading phase. The Degraw load cell

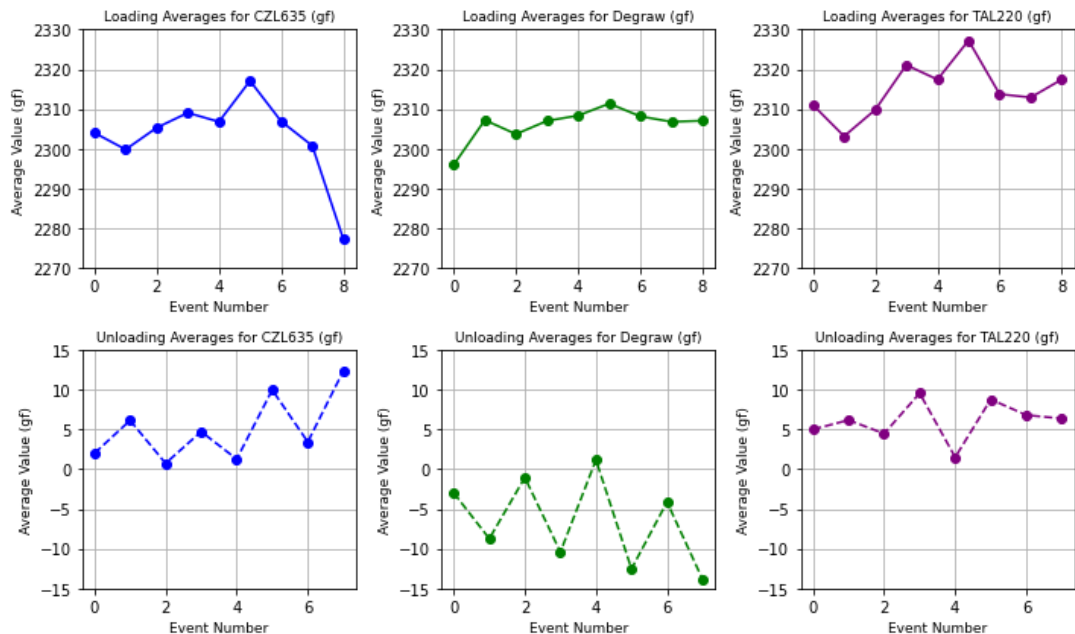


Figure 78. Average values for three load cells across eight loading/unloading cycles, showing: (top row) consistent loading responses between 2280-2325 gf with Degraw exhibiting smallest variation around 2305 gf; (bottom row) unloading (zero) return characteristics with notable offset patterns, where CZL635 shows positive drift (2-12 gf), Degraw shows negative bias (-10 to 0 gf), and TAL220 maintains positive values (2-10 gf).

demonstrates stable performance across all events, with readings consistently around 2310 gf, indicating reliability for applications requiring consistent measurements. Conversely, the TAL220 plot shows greater variability, with several readings drifting higher than expected, suggesting occasional deviations from anticipated values.

The CZL635 and Degraw load cells exhibit minor fluctuations and highest deviations during unloading at 12.35 gf and 13.86 gf, respectively (Table 23). In contrast, the TAL220 demonstrates greater stability with the lowest unloading deviation at 9.55 gf, suggesting consistent performance post-loading.

	CZL635	Degraw	TAL220
Max Absolute Deviation (Unloading)	12.35	13.86	9.55
Max Absolute Deviation (Loading)	22.72	11.37	27.24

Table 23. Maximum absolute deviation comparison for three load cells during repetitive testing, showing Degraw with best consistency under load and TAL220 with most stable unloaded readings.

Conclusion

The polynomial regression calibration methodology showed only marginal accuracy improvements, with R^2 values nearly identical to those of the linear regression model. The analog-to-digital converters (ADCs) utilised in this study feature sensitivity in the milligrams-force range; however, noise remains a significant challenge. Future work will focus on revising the PCB to include an external clock, thereby increasing data sampling speeds, which allows for higher Oversampling Rates (OSR) and could potentially reduce signal noise.

Ambient conditions, such as temperature fluctuations, affect load cell readings and induce measurement drifts, as notably exhibited by the TAL220. Maintaining thermal stability or implementing compensation techniques in characterisation setups is recommended. Additionally, preconditioning the system by allowing it to warm up for several hours enhances stability, and operating in a controlled environment yields more consistent results.

In conclusion, any of the analysed load cells can achieve accurate and consistent force measurements within a 1% error range and would be suitable. The study underscores the importance of minimising noise and environmental influences to ensure reliable force measurements over time and the requirements for regular calibrations to prevent external factors from altering sensor calibration.

5.1.2 Specific Setups for Single Axis Force

The specific setup for each haptic device varies to physically accommodate mounting the device and coupling it to the load cell in a manner representative of how the user perceives the force. For prototypes like PropelWalker, which exert a linear force directly onto the user's legs, the measuring setup should align the beam load cell with the actuator's force vector (Figure 79). This can be achieved using aluminium extrusion builds, a versatile and easy-to-implement approach commonly seen in previous literature for accommodating both the load cell and the actuator. When designing the test, characterising a prototype with a matching number of actuators to those present on the haptic device ensures the reported capabilities match those of the presented device.




Figure 79. Simplified proposed measurement setup for characterising single-axis force: (a) top and side view schematics showing the testing arrangement with load cell placement, and (b) photograph of the implemented system with labelled components including aluminium extrusion structure, DAQ board and analysis tools.

Mounting the actuator close to the load cell improves characterisation accuracy. For example, in PropelWalker, the frame's mass contributes to its inertia, affecting the latency between turbine (mounted on top) actuation and force transmission to the load cell at the base of the frame. For devices generating push/pull forces, such as ANISMA, S-type load cells are suitable but costlier. Similar characterisation can be performed using appropriately ranged beam load cells mounted vertically, providing an affordable alternative for active kinesthetic actuators which do not require a pulling test mechanism.




Figure 80 Characterisation approaches for wearable haptic devices: (a) ANISMA actuator prototype mounted on skin for direct force testing, (b) test fixture for precise alignment on optical table, and (c) schematic diagram showing vertical beam load cell configuration for push/pull force measurement.




Figure 81. Characterisation setup for the FacePush haptic feedback device: (a) schematic diagrams of force application paths for direct head pressure and side motor-pull mechanisms with measurement points indicated by arrows, and (b) photograph of the actual device being worn by a user with VR headset and attached motorised haptic actuators.

Characterising devices like FacePush, (Figure 81) which generates pressure on the face by pulling VR headset straps, requires a custom approach. A potential solution to measure the force the user will perceive is to modify a polystyrene head to include internal beam load cells, allowing for measurement of the force on each side independently when one or both actuators are activated. This ensures the collected data reflects the stimuli delivered to the user, providing an accurate basis for comparison with future designs.

5.2 Single-Axis Impact Force

Elastic or pneumatic actuators generate a sudden release of energy, often used to simulate an impact. When the actuator is directly coupled to the participant, and the force vector applies directly to the user's body, this can be characterised as a single-axis impulse. However, to accurately capture such fast stimuli, the measuring setup must capture data at a higher sample rate than when characterising static forces. Failure to do so may result in data that does not accurately represent the event's dynamics, which is critical for understanding the haptic system's behaviour and ensuring its safe and effective operation.

5.2.1 Sensor validation

Extensive testing and comparison of various load cells' performance were conducted in the previous section. The findings demonstrate the suitability of the selected load cells for measuring forces, provided that the data acquisition system can achieve high sampling rates.

The key considerations for impact force measurement include:

- **High sampling rate:** Capturing the rapid changes in force during an impact event requires a high sampling rate, ideally in the range of a few kHz.
- **Load cell rigidity:** The load cell's rigidity affects the measured force profile, with stiffer sensors resulting in sharper peaks and faster rise and fall times compared to more compliant materials like human skin.

- **Impulse consistency:** Despite differences in force profiles, the impulse (integral of force over time) should remain consistent for a given impact event, regardless of the sensor's rigidity.

When selecting a load cell for impact force measurement, consider the expected magnitude of the forces involved. Due to the rigidity of the load cell, it will not deform as much as softer materials like human skin. Consequently, the measured peak force will be significantly higher than what a user would experience so the load cell should be rated to match these forces to prevent damage.

5.2.2 Specific Setups for Single Axis Impact

When evaluating impact kinesthetic haptic actuators, the testing setup involves inherent limitations. The force-time profiles measured by rigid sensors differ significantly from those experienced by users due to skin plasticity. Rigid materials like aluminium or steel beam load cells produce sharp peaks with rapid rise and fall times, while materials with elasticity such as human skin generate smoother, elongated force-time curves.

Despite these profile differences, the impulse (calculated as the integral of force over time) remains consistent when mass and initial velocity are constant. This physical property represents the change in momentum of the impacting object. When characterising impacts for users, this creates an important consideration: while the measured impulse values remain consistent across materials, the perceived intensity varies between rigid sensors and human skin. Researchers must account for this discrepancy when translating characterisation results to expected user experiences.

For impact-delivering devices, secure mounting of all components minimises vibration noise and ensures accurate readings. Testbed design for devices delivering impacts directly to users follows methodology similar to that used for direct single-axis force characterisation.




Figure 82. ImpactVest characterisation setup: (a) photograph showing the experimental apparatus with labelled impactor and force sensor components for measuring direct impact forces, and (b) schematic diagram illustrating the arrangement of the impactor mechanism relative to the load cell for accurate impact measurement.




Figure 83. Air-based haptic device characterisation setups: (a) BLASNEL prototype schematic showing measurement approach for air-based tactile feedback, (b) user wearing the BLASNEL prototype, (c) Ω -jet prototype measurement configuration, and (d) photograph of the implemented Ω -jet actuator with backside view of nozzle arrangement..

Actuators delivering force directly to users can be measured by connecting the actuator to a load cell, as demonstrated in the ImpactVest study which uses pneumatic air for impact simulation.

For air applied directly to users (Figure 83) , a beam load cell combined with a surface panel of sufficient diameter can capture the air puff when positioned at typical user distance. This approach aligns with characterisation methods used for Aireal, a device employing loudspeakers to generate air vortices reaching approximately 1 meter from the user. When air-releasing nozzles couple to the user, mounting the nozzle to the load cell will provide measurements of the push force from the actuator, as seen in the Omega-jet case.

5.3 In-line Forces

While beam-type load cells offer a versatile and cost-effective solution for many applications, certain use-cases require alternative sensors for more reliable readings. In scenarios where forces are suspended or require alignment as with Pullups [112], S-type and button-type load cells prove more appropriate. These sensors measure both tensile and compressive forces and feature mounting fixtures at both ends for inline configurations. Their smaller size facilitates integration into space-limited devices.

Button and S-type load cells cost more due to complex manufacturing, with limited budget-friendly options available. Table 23 lists some of the models tested, alternative options with smaller form factors like the SBT630B Miniature S-type will be explored on future work.

Cost	Max capacity	Cost	Type	Dimensions (mm)
Honeywell FSS015WNSB	1.5kg	~70€	Compression Load Cell	9.14x3.18x5.59
Phidget Button Load Cell ID: 3136 0	50kg	~45€	Compression Load Cell	9mm height 25mm \varnothing
Phidget Button Load Cell ID: 3137 0	200kg	~45€	Compression Load Cell	9mm height 25mm \varnothing
Micro Load Cell DYLY-106 of the S-type	10kg	~50€	S-Type load Cell	25x30x12mm

Table 24. Comparison of four load cell options for haptic device characterization, including Honeywell FSS015WNSB (1.5 kg miniature compression, €70), Phidget button load cells in 50 kg and 200 kg capacities (both €45, 25 mm diameter), and Micro Load Cell DYLY-106 S-type (10 kg tension/compression, €50), each offering different form factors and measurement ranges for various testing scenarios.

5.3.1 Specific Setups for In-Line force

This section explores setups for inline force characterisation, focusing on S-type and bottom load cells designed for precise measurements of forces applied directly or in-line with their surfaces. Button load cells, with their compact circular design, suit single-axis force measurements particularly well. Their structural characteristics allow more direct and localised force measurement compared to beam load cells.

S-type load cells and some button cell designs provide anchoring mechanisms on both ends, enabling integration in kinesthetic devices to measure in-line forces. Pull-Ups [112] employed this approach. This active actuator using pneumatic muscles to pull participants upward was characterised similarly to the proposed use case. Loading weights suspended from the system included an in-line force sensor measuring real-time load while monitoring actuator contraction distance. Sensor integration within the kinesthetic implementation also enabled real-time force monitoring and closed-loop control during user studies.




Figure 84. Pull-Ups pneumatic VR system: (a) wearable haptic device with cables exerting upward forces to simulate climbing sensations in VR, (b) close-up of pulley mechanism with tension cables, and (c) schematic diagram of S-type load cell measurement setup for characterising the system's in-line tension forces.




Figure 86. Examples of various haptic devices using passive kinesthetic feedback mechanisms: (a) Wireality [147] hand tracking system with string-based resistance for physical object simulation in VR, (b) DextrES [71] electrostatic braking glove for variable finger resistance, and (c) wearable haptic device with cable-based force

Passive actuator instances like those in the exoglove design by Chiaradia et al. [113], Wireality [159] or DextrES [68] use braking mechanisms to restrict user movements (Figure 86). Characterising these devices requires an actuation mechanism applying load to determine dynamic response and capabilities. Haptic-Links researchers demonstrate a fundamental example, utilising a lathe to connect one end of the kinesthetic device to the spindle and the other to a carriage producing precise linear displacement. More advanced methodologies include dedicated push-pull testing apparatus or linear actuation systems, as demonstrated by Dexmo.




Figure 85. Pulling tester mechanisms for characterizing tension forces in haptic devices: (a) Instron 3344L industrial tester, (b) manual pulling tester with force gauge, (c) linear actuator equipped with servo motor, and (d) proposed implementation with S-load cell for characterizing similar actuators as DextrES.

Push-pull testers use precise mechanical linear actuators enabling linear displacement while measuring applied force (Figure 85). Mechanical linear actuators from CNC equipment, such as the RATMMOTOR ZBX80, offer an alternative for low-force applications. Beam or S-type load cells can be integrated with these type of actuators to enable dynamic characterisation and accurate force measurement during both pushing and pulling movements on brake-force kinesthetic devices.

Button load cell sensors offer the advantage of compact size, exemplified by the Honeywell FSG15N1A miniature button load cell measuring just 9.14×3.18×5.59mm. TORC showcased




Figure 87. Integration of force sensors in haptic interface designs: (a) ProxyTouch characterization testbed showing wrist rest, button-type force sensor, motor-controlled lever, and thumb strap for measuring finger forces during interaction, and (b) TORC [34] hand-held controller with embedded miniature force sensors for measuring grip pressure in VR applications.

this sensor’s integration in a closed-loop pinching kinaesthetic interface. Similar small form factor sensors were employed in ProxyTouch (Chapter 7), placed behind the user’s fingertip to measure force during pinching actions (Figure 87 a).

5.4 Rotational Joint Torque

For some kinesthetic devices, force generates over a pivot point or at a distance from the user as with Pacapa or Haptic Links (devices that operate with a pivot). Haptic Links latches its joint, offering a force-locking mechanism to which VR controllers mount. Meanwhile, Pacapa consists of an active mechanism operating a joint with a servo pivoting one side while the user holds the device to deliver active kinesthetic torques. In facePush, elastic weights mounted on both sides of a VR headset generate forces when released, creating torque with the pivoting point at the centre between the two actuators. For such devices, torque characterisation provides a more representative property allowing direct comparison between designs.

To characterise devices that generate torque, the setups described in this section use single-axis beam load cells based on the principle that torque is the product of force applied at a distance from a pivot point. Using beam or S-type load cells, torque is directly computed by multiplying the force by the distance from the pivot point to the force measurement point.

$$\text{Torque}(\tau) = \text{Force} \times \text{Distance}$$

5.4.1 Specific Setups for Joint Torque

In cases where measurement involves force in a single axis and the actuator includes mechanical linkages or rotating joints, as in devices like Claw or CapstanCrunch, determining




Figure 88. Force sensor setup diagrams for kinesthetic feedback controllers: (a,b) schematics showing proposed characterization method for measuring forces perpendicular to finger movement in pinching devices, and (c,d) photographs of the actual Claw and CapstanCrunch haptic controllers with their mechanical pinch-grip mechanisms that can be tested using this arrangement.

the characterisation approach may be less straightforward. Even though the mechanism might seem to exert torque, it is designed to reproduce a force to the user's fingertip. Employing beam load cells aligned with the actuation lever point can accurately characterise the controller's properties, as shown by the authors. Additionally, reporting torque capabilities offers deeper insights into the device's capabilities, as force capabilities depend on the actuator's torque and the arm's distance to the user's fingers.

Studies on devices such as Claw and CapstanCrunch explored additional properties beyond force, including the stiffness capabilities of the actuating mechanism. This helps determine the device's achievable stiffness for simulating virtual objects, measured by the displacement of the lever under a specific force, reported in N/mm or $N/^\circ$.

This section discusses three prototypes demonstrating this method: Haptic Links, which utilises a latching mechanism to generate braking torque; Pacapa, an active mechanical system that directly exerts torque on the user; and OsciHead, which produces rotational torque for a user wearing a VR headset.




Figure 89. Haptic links characterization setups: (a) alternative setup with linear actuator top view, (b) alternative setup with linear actuator side view, and (c) lathe-based testing apparatus components including spindle, haptic link, force gauge, and slide.

Braking mechanism

Haptic Links presents a detailed characterisation methodology, including visual aids (Figure 85). The characterisation method resembles a pull test systems, where the device functions as a passive mechanism. Here, Haptic Links is mounted to a lathe bed, with one end attached to the lathe chuck and the other connected to a sensor coupled with the lathe's saddle component. The lathe saddle moves, exerting a pull on the levers. By measuring the distance from the clamping point to the hinge mechanism, torque can be calculated. If a lathe is unavailable, a linear ball screw actuator can provide the required linear pull.

Active Mechanism

Devices like Pacapa (Figure 90), which feature a rotational joint, use a servo motor as an actuator to exert force on the user's hand by opening the side. A similar methodological approach used for single-axis forces can create the testbed, where a beam load cell aligns with the actuator. This setup allows the actuator to open and exert pressure on the load cell's measurement end. This technique directly measures the force the device produces, and multiplying this force by the arm length (l) calculates the torque. This illustrates that measuring torque and force are complementary; in this case, the author could report both metrics, enabling easier data comparison.

Inertia Mechanism

Devices such as OsciHead (Figure 91) use actuators on both sides of a VR headset to deliver twisting effects when released in opposite directions. When actuators release simultaneously, the proposed testbed for FacePush could be used. Due to opposite actuation directions, it will generate torque around the axis at the centre between the two forces generated on each side. A potential approach to characterise such a system would mount the actuators on beam load




Figure 90. Pacapa haptic device characterization: (a) device at 0° position, (b) device rotated to 90° position showing joint movement, and (c) schematic of testbed setup with beam load cell aligned to measure force and torque at distance L from rotation axis




Figure 91. OscillHead haptic feedback device: (a) linear vertical actuation movement along Z-axis, (b) rotational movement around vertical axis in clockwise direction, and (c) schematic of proposed testbed design to measure torque around the rotation axis with distance t between measurement points.

cells fixed at the same distance from the centre as they are on a headset. This method approximates the stimulus delivered to the user. The individual torques created by each actuator are measured as τ_1 and τ_2 :

$$\tau_1 = F_1 \times L_1$$

$$\tau_2 = F_2 \times L_2$$

In this case, F is the applied force and L is the distance from the rotation axis to the point where force is applied. Assuming both L_1 and L_2 are equal (i.e., the actuators are equidistant from the rotation axis), and if the forces act in a plane perpendicular to the rotation axis, then the total torque τ_{total} would be the sum of the individual torques, since one actuator pulls forward and the other pulls backward, making their rotation effects additive:

$$\tau_{\text{total}} = \tau_1 + \tau_2$$

However, this simplified approach measures force directly at the actuators and does not account for how stimulus delivery might change when the actuators are mounted in the actual headset. Following the guiding principle that the characterisation process should closely mirror the final device or prototype, an alternative is to integrate the beam load cell into a modified polystyrene mannequin head, as shown in Figure 92. This setup provides an assessment of the kinesthetic experience closer to what the user will experience. In scenarios such as, beam load cells better capture the experience only when both actuators operate simultaneously in the same direction (e.g., applying force). For accurately measuring torque, more advanced sensors are needed. These approaches will be explored in the next section on haptic devices and user cases.




Figure 92. Alternative OsciHead characterisation testbed: schematic diagram showing measurement approach with force vectors indicated at distance L from rotation axis, enabling calculation of torque produced by the headset-mounted actuators.

5.5 Force Torque

Handheld devices often feature actuators mounted at a distance from the user's hand, causing the output to be perceived as torque rather than linear force. As shown in the literature review, these devices are characterised using single-axis beam load cells coupled to the actuator, which report the actuator's capabilities but not necessarily the user's experience. Considering the distance from the actuators and reporting the output as torque allows for more comparable results.

While using single or multiple axis load cells remains common, more accurate multi-axis force-torque sensors, such as those used to characterise devices like iTorqU, provide more precise reporting from the user's point of contact. This section introduces a custom-built 6-axis force and torque sensor, constructed using affordable beam load cells and 3D-printed components. The design principle and operation are detailed, followed by validation, and concludes with use cases.




Figure 93. Multi-axis force measurement systems: (a) schematic of standard 6-axis force plate application for robotic arm base force and torque measurement, and (b) FARE-KE custom force plate design showing arrangement of four 3-axis load cells with labeled dimensions (w , L) and coordinate system directions.



Figure 94. Commercial and custom multi-axis force sensors: (a) Forsentek 6-axis force/torque sensor with compact cylindrical design, (b) industrial multi-axis load cell with square mounting pattern, (c) 3-axis load cell using stacked single-axis design with aligned beam cells, and (d) custom-built load cell with integrated electronics in green housing.

5.5.1 Sensor Design & Operation

The 6-axis F/T sensor design draws inspiration from force plate configurations, utilising four 3-axis load cells (Figure 93). Force plates are used at the base of robotic arms to calculate forces and torques during operation. Although commercial 3-axis load cells cost less than 6-axis versions, they remain prohibitively expensive. To address this, the research developed a more accessible method to construct 3-axis load cells using single-axis beam load cells. This section outlines the design, calibration, and performance evaluation of both the 3-axis load cells and the final 6-axis F/T sensor.

This approach results in a larger footprint compared to commercial sensors such as the ATI Nano17, making it unsuitable for integration on wearables and handheld devices. However, it provides a cost-effective solution for quasi-static evaluation of kinesthetic haptic devices where the device is mounted on a measuring rig.

5.5.2 3-Axis Load Cells

The 6-axis F/T sensor computes force and torque across various axes by combining the force readings from four 3-axis load cells. This section focuses on the design and considerations for constructing these 3-axis load cells. The final design uses single beam load cells, with 3D printing facilitating component alignment and coupling, aiming to provide a practical and affordable multi-axis force measurement solution based on established load cell technologies.

Design Overview and Assembly

As shown in Figure 94, 3-axis load cell designs range from intricate mechanical configurations with complex strain gauge placements (a and b) to simpler designs that align three beam load cells along the X, Y, and Z axes (c). The custom-built load cell (Figure 94 d) is based on commercial designs such as the one shown in (d). It utilises widely available TAL220 beam

load cells aligned along the x, y, and z directions, connected end-to-end using PLA 3D-printed parts for the central body and alignment elbows. Aluminium epoxy glue and M3 screws secure the load cells, with overall dimensions of 80x80x35 mm. Accurate alignment of the load cell centres is crucial for measurement accuracy and is achieved by precisely aligning the centres on the same plane using the 3D-printed parts. If available, using stiffer materials like aluminium for the alignment parts could further improve sensor performance. Detailed assembly instructions are provided in Appendix X - 3 Axis Load Cell Assembly.

The arrangement of the load cells allows for aligned force measurements on each axis but introduces crosstalk challenges, where forces applied along one axis affect readings on other axes. To address this, filtering techniques and a crosstalk calibration matrix, as methods used by U-Xuan Tan et al [160], were implemented in the firmware capturing the sensor data. The subsequent sub-section explains calibration methods employed to correct these inaccuracies, with a comprehensive step-by-step approach detailed in Appendix 20 – FARE-KE 3-Axis Single Point Calibration.

Calibration Methods Overview

The configuration in the custom design means each beam measures forces relative to its own grounding, leading to crosstalk interference. Various calibration methods were explored to mitigate this issue and enhance sensor accuracy.

Single-Point Calibration Method: This approach calculates the inverse of the calibration matrix, obtained by applying a single known load and using the corresponding sensor readings to calculate calibration coefficients. The calibration matrix (K) is constructed with main sensitivity coefficients along the diagonal and crosstalk coefficients as off-diagonal elements.

The inverse of this matrix (K^{-1}) transforms raw sensor readings into actual force values.

Multipoint Calibration Method: To overcome potential limitations of single-point calibration, this method involves multiple load cases to consider non-linearity. Two specific approaches were implemented:

Least-Squares Regression: This method finds the calibration matrix (K) by minimising the sum of squared differences between observed sensor outputs (R) and model predictions. It

inherently accounts for crosstalk by treating calibration as a system of linear equations that includes interactions between axes.

Polynomial Regression: This method extends the calibration model to include non-linear relationships by fitting polynomial equations to the data points. To account for crosstalk, polynomial terms involving multiple sensor readings were incorporated into the calibration model. This approach provides a more comprehensive calibration model for the 3-axis load cell, considering both non-linear behaviour and crosstalk effects.

To implement these calibration methods, data is collected under known loads applied along each axis of the load cell. The corresponding sensor readings are recorded and used to calculate the calibration coefficients. Detailed mathematical formulations, numerical examples, and code implementations for each calibration method are provided in Appendices 20 and 21.

To evaluate the effectiveness of the calibration methods, a separate validation dataset, not used during the calibration process, assesses the accuracy and robustness of each method in predicting forces based on sensor readings.

Validation and Performance Evaluation

A comprehensive validation process was conducted to evaluate the custom-built 3-axis load cell and assess the effectiveness of the implemented calibration methods. One of the assembled load cells was selected for this validation study. A separate validation dataset, consisting of sensor readings corresponding to known applied forces, was used to ensure an unbiased evaluation.

Validation Results

The validation phase highlighted the performance of the Polynomial Regression method in terms of both Mean Absolute Error (MAE) and Root Mean Square Error (RMSE).

Method	MAE (g)	RMSE (g)	Max Error (g)
Single-Point	8.91	15.24	70.43
Least Squares	8.70	14.72	67.29
Polynomial Regression	7.54	11.28	44.93

Table 25. Validation performance comparison across different modelling techniques for load prediction, quantifying models' accuracy using Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Maximum Error in grams (g).

Error Percentages (%)

Load(g)			Least Squares Error			Polynomial Error			Single Point Error		
x	y	z	x	y	z	x	y	z	x	y	z
-60	0	0	10.17	1.67	5.33	12.00	12.67	2.17	10.17	2.00	5.50
-285	0	0	2.81	0.21	3.12	3.44	2.14	0.81	2.74	0.42	3.30
-465	0	0	1.78	0.13	3.03	2.32	1.14	1.23	1.72	0.34	3.20
-1055	0	0	1.18	0.12	3.11	1.73	0.31	1.78	1.12	0.35	3.27
-1925	0	0	0.97	0.17	3.50	1.66	0.05	2.33	0.90	0.38	3.66
0	60	0	9.67	1.33	1.83	9.33	7.50	2.17	9.83	1.33	1.83
0	285	0	3.30	0.67	0.25	1.75	0.28	0.70	3.33	0.74	0.35
0	465	0	2.67	0.32	0.41	0.99	0.49	1.10	2.71	0.41	0.49
0	1055	0	2.16	0.51	0.95	0.40	0.19	1.02	2.22	0.59	1.03
0	1925	0	2.13	0.81	1.89	0.33	1.17	0.24	2.19	0.88	1.97
0	0	-60	8.00	0.50	5.17	8.00	9.83	0.17	7.67	0.50	5.00
0	0	-285	1.89	0.56	1.61	0.88	0.77	1.47	1.58	0.70	1.58
0	0	-465	1.35	0.43	1.03	0.22	0.04	1.14	1.05	0.58	0.97
0	0	-1055	0.81	0.70	0.76	0.45	1.14	0.56	0.51	0.85	0.69
0	0	-1925	0.48	0.84	0.66	0.82	1.57	0.54	0.18	0.99	0.60

Figure 95. Heatmap table comparison for the three calibration methods presented, showing the error as a percentage of the load applied to the axis.

As illustrated in the heatmap Figure 95, error rates are notably lower for the Least Squares and Polynomial Regression methods compared to the Single-Point method, with the Polynomial Regression method demonstrating the lowest error percentages overall.

The validation data was processed using each calibration method to estimate the forces and evaluate the performance of each technique. These estimated forces were then compared to the actual force measurements to calculate the absolute errors. Figure 96 presents the absolute errors for each method across all validating measurements, visualising performance across a range of load conditions. As shown on Table 25, even the highest-performance calibration method can have a MAE of up to 7.5 grams. This particularly affects measurements under the 100-gram range, where this error value could represent up to a 7.5% error. However, as higher forces are measured, this noise becomes less relevant and tends to remain below 2% error over the loading reading.

The Polynomial Regression method proved most effective, offering improved accuracy compared to the other methods. However, it is important to acknowledge that even this method has limitations, particularly when measuring low forces. Future work will focus on enhancing ADC noise immunity and increasing the sampling rate to acquire more data for better noise filtering and compensation. Additionally, regular validation and performance monitoring of the calibrated sensor are essential to ensure consistent and accurate force

Error Values (grams)

Load(g)			Least Squares Error			Polynomial Error			Single Point Error		
x	y	z	x	y	z	x	y	z	x	y	z
-60.0	0.0	0.0	6.1	1.0	3.2	7.2	7.6	1.3	6.1	1.2	3.3
-285.0	0.0	0.0	8.0	0.6	8.9	9.8	6.1	2.3	7.8	1.2	9.4
-465.0	0.0	0.0	8.3	0.6	14.1	10.8	5.3	5.7	8.0	1.6	14.9
-1055.0	0.0	0.0	12.5	1.3	32.8	18.3	3.3	18.8	11.8	3.7	34.5
-1925.0	0.0	0.0	18.7	3.2	67.3	31.9	1.0	44.9	17.3	7.4	70.4
0.0	60.0	0.0	5.8	0.8	1.1	5.6	4.5	1.3	5.9	0.8	1.1
0.0	285.0	0.0	9.4	1.9	0.7	5.0	0.8	2.0	9.5	2.1	1.0
0.0	465.0	0.0	12.4	1.5	1.9	4.6	2.3	5.1	12.6	1.9	2.3
0.0	1055.0	0.0	22.8	5.4	10.0	4.2	2.0	10.8	23.4	6.2	10.9
0.0	1925.0	0.0	41.0	15.5	36.3	6.4	22.5	4.7	42.2	16.9	37.9
0.0	0.0	-60.0	4.8	0.3	3.1	4.8	5.9	0.1	4.6	0.3	3.0
0.0	0.0	-285.0	5.4	1.6	4.6	2.5	2.2	4.2	4.5	2.0	4.5
0.0	0.0	-465.0	6.3	2.0	4.8	1.0	0.2	5.3	4.9	2.7	4.5
0.0	0.0	-1055.0	8.5	7.4	8.0	4.7	12.0	5.9	5.4	9.0	7.3
0.0	0.0	-1925.0	9.2	16.2	12.8	15.7	30.3	10.4	3.4	19.0	11.6

Figure 96. Heatmap table comparison for the three calibration methods presented, showing the error in grams when loading the sensor after calibration.

measurements over time. Establishing a systematic approach to recalibration and performance evaluation will maintain the sensor’s precision and reliability across various applications.

5.5.3 6-axis Force-Torque Plate

Building upon the design and validation of the custom 3-axis load cell, this section presents a novel 6-axis force-torque plate that leverages multiple 3-axis load cells to measure both forces and torques with high precision. This design is particularly well-suited for characterising the output of handheld kinesthetic haptic controllers, offering a cost-effective solution for capturing complex force and moment data.




Figure 97. 6-axis force/torque plate design used to measure force and torque on handheld devices. The four three-axis load cells are labeled A, B, C and D with dimensions L and W used to compute

The proposed force-torque plate integrates four calibrated 3-axis load cells, arranged near the corners of two parallel acrylic plates, each measuring 12mm in thickness: a base plate and a top plate. The bottom plate serves as a grounding base fastened to a table or other flat surface. The upper plate functions as the mounting surface for the device under measurement. Each load cell is securely bolted between the plates, resulting in a total of 12 force channels for the entire system. These 12 force readings calculate the force (F_x, F_y, F_z) and moment (M_x, M_y, M_z) acting on the plate.

An advantage of this multi-sensor force plate design is the ability to achieve higher load capacities by summing the rated loads of the four individual sensors. Additionally, by distributing the load cells towards the edges of the plates, the force plate enhances its moment capacity and stiffness compared to single-sensor configurations. To ensure accurate measurements, the force plate should be calibrated using a multipoint calibration approach, which involves collecting data at multiple force and torque levels across the sensor's measurement range.

The assembly and manufacturing processes for the 6-axis force-torque plate are streamlined to ensure ease of access and replication. Using components that are easily obtained and basic manufacturing techniques, the force plate can be reproduced by researchers equipped with standard digital fabrication tools. The force plate's modular design allows for significant customisation, with users able to adjust load capacities or dimensions by selecting appropriate load cells and modifying the size of the acrylic plates.

Design Overview and Assembly

In contrast to industry-standard force plates (Figure 98), which typically measure forces in the range of hundreds of newtons, this force plate operates within a range aligned with the capabilities of kinesthetic haptic devices ($\pm 80\text{N}$).




Figure 98 Force plates applications: (a) commercial force plate being used for gait analysis and biomechanical assessment of foot impact forces during walking, and (b) schematic illustration of a robotic arm system with integrated force plate at its base to measure operational forces and torques during movement.

The force plate is constructed using four 3-axis load cells, strategically arranged as shown in Figure 97. These beam load cells, having undergone rigorous testing as described in Section 5.5.2, have demonstrated reliable performance for forces up to 20N. The four 3-axis load cells are positioned between two acrylic plates, each measuring 12mm in thickness. The overall dimensions of the assembled force plate measure 200x200x65mm.

During assembly, meticulous attention is given to the alignment of the load cells, as precise alignment affects the mathematical calibration processes and significantly influences crosstalk levels. The 3-axis load cells are carefully aligned and mechanically coupled to the plates using M4 screws. Before final assembly, each individual 3-axis load cell undergoes calibration, essential for accurate force and torque measurements.

The 6-axis force plate calculates forces and torques by utilising data obtained from each 3-axis load cell after applying compensation for crosstalk. A total of 12 readings compute the forces and torques (F/T). The mathematical model supporting this approach is represented by the following equation:

$$\begin{bmatrix} F_x \\ F_y \\ F_z \\ M_x \\ M_y \\ M_z \end{bmatrix} = \begin{bmatrix} F_{x_A} + F_{x_B} + F_{x_C} + F_{x_D} \\ F_{y_A} + F_{y_B} + F_{y_C} + F_{y_D} \\ F_{z_A} + F_{z_B} + F_{z_C} + F_{z_D} \\ (W/2)(-F_{z_A} - F_{z_B} + F_{z_C} + F_{z_D}) \\ (L/2)(-F_{z_A} + F_{z_B} + F_{z_C} - F_{z_D}) \\ (L/2)(F_{y_A} - F_{y_B} - F_{y_C} + F_{y_D}) + (W/2)(F_{x_A} + F_{x_B} - F_{x_C} - F_{x_D}) \end{bmatrix}$$

In this equations, $M_x, M_y,$ and M_z denote the torques around the X, Y, and Z axes, respectively, while $F_x, F_y,$ and F_z represent the forces along those axes. The subscripts A, B, C and D refer to the individual 3-axis load cells positioned at each corner of the plate. It is important to note that these equations assume measurements are taken at the origin of the top plate, which is the centre and equidistant from all the load cells. Since the force and torque calculations rely on readings from the previously calibrated 3-axis load cells, the calibration process at this stage focuses on compensating for crosstalk and nonlinearities added when combining the force from the 3-axis load cells.

Calibration Methods Overview

The 6-axis load cell requires further calibration to ensure accurate force and torque measurements. The primary challenge is addressing nonlinearities and crosstalk effects that arise from the arrangement of the individual load cells. For the 6-axis sensor, a multipoint calibration approach is employed. This technique involves collecting calibration data at multiple force/torque levels across the sensor's measurement range. The calibration process involves the following steps:

- **1. Data Collection:** Known forces and torques are applied at multiple calibration points, with the corresponding raw sensor readings recorded. The collected data is organised into matrices: \mathbf{F} (applied forces/torques) and \mathbf{R} (raw sensor readings).
- **2. Calibration Matrix Computation:** The least squares method computes the calibration matrix \mathbf{K} , which minimises the difference between the product of the raw sensor readings \mathbf{R} and the calibration matrix \mathbf{K} , and the known applied forces and torques \mathbf{F} . This relationship can be expressed mathematically as:

$$\min_{\mathbf{K}_i} \sum_{i=1}^n \|\mathbf{F}_i - \mathbf{R}_i \cdot \mathbf{K}_i\|^2$$

where n is the total number of calibration points. The resulting calibration matrix \mathbf{K} converts the raw sensor readings into calibrated measurements of force or newtons.

Validation and Performance Evaluation

Using an assembled F/T plate, the calibration proceeded as described in Appendix 24 – FARE-KE 6-Axis Load Cell Multipoint Calibration, from which the calibration matrix was obtained.

$$\mathbf{K}_{\text{cal}} = \begin{bmatrix} 0.9996 & 0.0102 & 0.0126 & -0.0004 & -0.0217 & 0.0051 \\ -0.0303 & 0.9946 & 0.0150 & 0.0258 & 0.0010 & 0.0049 \\ -0.0091 & -0.0159 & 0.9946 & 0.0011 & -0.0094 & -0.0000 \\ -0.1138 & -0.1117 & -0.1168 & 1.1472 & 0.0196 & -0.0137 \\ -0.4016 & -0.0569 & 0.5885 & 0.0120 & 1.1564 & 0.0466 \\ -0.0173 & 0.0068 & 0.1061 & 0.0170 & -0.0043 & 1.0059 \end{bmatrix}$$

To assess performance various tests were conducted comparing the readings corrected with the matrix to the original readings obtained during calibration. First, calibration effectiveness was evaluated using Mean Absolute Percentage Error (MAPE) and R-squared metrics. The results indicated that the readings obtained during calibration already achieved an R^2 value of 0.98. However, after applying the calibration matrix, the sensor accuracy further improved to an R^2 value of 0.999.

Metric	Raw Readings	Calibration
MAPE (%)	2.814737	1.849263
R-squared	0.983804	0.999691

Table 26. Mean Absolute Percentage Error (MAPE) and R-squared metrics before and after using the calibration matrix, demonstrating improved measurement accuracy with calibrated values.

Next, histograms and box plots illustrated the distribution of errors between the raw and corrected data (Figure 99). The histograms show a notable convergence of the calibrated readings towards zero error, signifying a substantial reduction in error magnitude relative to the raw data. The box plots confirmed these results, showing that the interquartile range (IQR) of the calibrated data was significantly narrower and the median closer to zero, suggesting enhanced accuracy.




Figure 99. Error distribution analysis for load cell calibration: (a) histogram comparing raw and calibrated readings showing convergence towards zero error after calibration, and (b) box plot demonstrating significantly reduced interquartile range and median values closer to zero for calibrated measurements.

The MAE plot (Figure 100) showed in more detail the magnitude of improvement in each axis in Newtons for the force axes and Newton-meters for the torque axes. Overall, calibrated readings offered lower error readings compared to the raw readings.




Figure 100. Mean Absolute Error (MAE) by axis for raw versus calibrated measurements, showing substantial error reduction across all force (Fx, Fy, Fz) and torque (Mx, My, Mz) components, with most significant improvements observed in the Fz axis.

Additional tests were conducted to evaluate the calibration matrix’s influence on crosstalk. To better characterize the crosstalk behaviour in the 6-axis load cell design, crosstalk errors were plotted and compared to the readings after applying the compensation matrix across all calibration points. Figure xx illustrates the calibration data points for Fz and Mz, serving as examples. The calibration sequence, depicted in the graph, began with applying a force along Fy, followed by Mx, Fx, My, Mz, and finally Fz. This plot highlights the crosstalk on the Fz and Mz axes and demonstrates how crosstalk intensity escalates with increased load on the principal axis. The compensated readings, shown in orange, indicate a substantial reduction in crosstalk across these axes, with similar outcomes observed for the other axes.

Raw Force/Torque Crosstalk						Calibrated Force/Torque Crosstalk (N)					
Fx(N)	Fy(N)	Fz(N)	Mx(Nm)	My(Nm)	Mz(Nm)	Fx(N)	Fy(N)	Fz(N)	Mx(Nm)	My(Nm)	Mz(Nm)
0.026	0.049	0.013	0.001	0.002	0.002	0.026	0.049	0.013	0.000	0.003	0.002
0.006	0.000	0.000	0.01	0.002	0.002	0.023	0.000	0.011	0.003	0.003	0.001
0.015	0.000	0.018	0.029	0.001	0.002	0.025	0.000	0.008	0.004	0.002	0.006
0.042	0.000	0.031	0.045	0.002	0.011	0.015	0.000	0.006	0.001	0.003	0.000
0.055	0.000	0.048	0.059	0.000	0.002	0.022	0.000	0.001	0.004	0.001	0.013
0.107	0.000	0.074	0.093	0.000	0.022	0.000	0.000	0.005	0.006	0.000	0.001
0.180	0.000	0.129	0.156	0.000	0.037	0.004	0.000	0.011	0.006	0.001	0.001
0.301	0.000	0.22	0.258	0.001	0.062	0.01	0.000	0.021	0.004	0.000	0.001
0.453	0.000	0.283	0.362	0.001	0.089	0.017	0.000	0.004	0.007	0.000	0.003
0.594	0.000	0.363	0.483	0.003	0.116	0.002	0.000	0.015	0.003	0.001	0.000

Table 27. Crosstalk error measurements comparing raw and calibrated values after applying the calibration matrix to loading values along the Fy axis, demonstrating reduced interference across all force and torque channels.




Figure 101. Crosstalk effect visualization: (a) error magnitude on Fz axis and (b) error magnitude on Mx axis when other axes are actuated sequentially, demonstrating significant reduction in undesired cross-axis interference after applying the calibration matrix.

After calibration, there was a clear reduction in crosstalk error across all axes. The calibration proved particularly effective in mitigating crosstalk on the Mx axis, reducing errors to near zero. On other axes, the crosstalk magnitudes were significantly lowered. Overall, the calibration matrix substantially reduced crosstalk, particularly along the torque axes, where the standard deviation decreased by a factor of 10 for all metrics used to assess crosstalk errors.

	Fx	Fy	Fz	Mx	My	Mz
Std Dev raw	0.1139	0.0821	0.3907	0.0846	0.0979	0.0528
Std Dev Cal	0.0224	0.0377	0.0343	0.0154	0.0085	0.0043
Mean Raw	0.0683	0.0541	0.2465	0.0332	0.0538	0.0369
Mean Cal	0.0172	0.0231	0.0310	0.0072	0.0049	0.0045
Max Raw	0.5940	0.4120	2.2360	0.4830	0.4100	0.2580
Max Cal	0.0861	0.1700	0.1287	0.0783	0.0517	0.0158

Table 28. Statistical comparison of measurement errors before and after calibration showing substantial improvements in standard deviation, mean values and maximum errors across all six axes, confirming the effectiveness of the calibration matrix.

The statistical analysis, demonstrating decreases in standard deviation, mean errors, and maximum error values post-calibration, confirms the calibration matrix's role in enhancing sensor performance and the capabilities of the proposed 6-axis design. The calibration and subsequent error evaluations were conducted with forces and torques applied separately to individual axes. This method ensures precise calibration but does not consider scenarios where

multiple axes are actuated simultaneously. While no adverse effects on calibration accuracy have been observed under these conditions, future work should compare these results with those obtained using industrial-grade sensors under multi-axis loading conditions to provide deeper insights into the robustness of the calibration matrix.

5.5.4 Specific Setups for Torque measurements




Figure 102. Gyroscope-based haptic controllers with force/torque measurement: (a) iTorqu handheld device with gimbal-mounted flywheel, and (b) experimental characterisation setup using ATI Mini45 force/torque sensor with control moment gyroscopes.

The primary target for this design is handheld haptic controller. The proposed 6-axis F/T sensor aims to be used as in previous research with devices such as iTorqu or similar handheld devices [57], [62] that utilise industrial F/T sensors by firmly attaching the handle to the measuring plate (Figure 102).




Figure 103. Arrangement for handheld devices testbed mounting: (a) side view of haptic device mounted at the base of the handle on force/torque plate, (b) Ideally mounting will consider user handling (c) Forces at a distance can be calculated based on rotation and length arms variables.

As depicted in the optimal mounting should replicate the typical distance at which a user would hold the device. This setup supports quasi-static characterisation where the controller maintains an upward orientation as a baseline. Various devices can be evaluated with this approach, as illustrated in Figure 103.

5.6 Closing Section

The development of the FARE-KE framework, as described in Chapters 3, 4, and 5, originated from focused research on designing various kinesthetic haptic prototypes requiring thorough validation and quantitative analysis. This need intensified during the COVID-19 pandemic, as traditional user studies became impossible during lockdown periods. The FARE-KE framework emerged from identified gaps and inconsistencies in existing methods when characterisations had to be conducted with limited resources and without laboratory access.

Early Experiences and Project Insights

Before the FARE-KE framework, the research was primarily project-based, investigating methodologies to deliver kinesthetic haptic feedback. Each project explored kinesthetic methods of increasing complexity, from linear actuation to multi-axis torque actuation. These projects (VR Recoil, Proxy Touch, and HapticWhirl) are detailed in subsequent chapters and presented challenges that underscored the need for a unified framework in haptic device characterisation.

Throughout this development, most characterisation tools and approaches were one-off designs, leading to the integration of diverse insights into a coherent methodology addressing the wide-ranging aspects of haptic device functionality. By addressing these challenges systematically, the FARE-KE framework was developed as a tool for haptic device evaluation, focusing on consistency and reproducibility of research outputs.

Chapter 6 (VR Recoil) examines the development and evaluation of a gun-shaped controller for virtual reality environments. This chapter determines the impact of linear force feedback on aiming tasks in virtual environments, focusing on how different controller configurations influenced user perceptions and interactions. Before conducting user studies, a systematic evaluation method was required to measure and compare the various kinesthetic levels delivered and their effects across different setups. This project initiated the development of methodologies later incorporated into the FARE-KE framework.

Chapter 7 discusses "Proxy Touch", an initial study determining the extent to which the user's own body can serve as a haptic medium through visual retargeting techniques. This project explored psychological and physiological aspects of haptic perception to enhance the perception of virtual objects. The challenges in measuring such interactions precisely

highlighted shortcomings of previous evaluation methods. While earlier instances used sensors like FSR, further research revealed potential accuracy limitations despite their low profile, emphasising the need for precise and reliable characterisation tools.

Chapter 8 describes "HapticWhirl", focusing on the development and testing of a haptic controller employing gyroscopic effects to provide diverse tactile responses. This prototype presented the greatest characterisation challenge, as the actuator was placed at a distance from the user's hand and generated complex output torque. The difficulties in consistently characterising these outputs underscored the necessity for a robust, scalable evaluation framework. Insights from the HapticWhirl project were crucial in defining the evaluation criteria used in the FARE-KE framework, ensuring it could address the complexities of devices producing complex force and torque outputs.

6 VR Recoil

"VR Recoil" was presented at INTERACT 2019 under the title "Exploring the Effects of Replicating Shape, Weight, and Recoil Effects on VR Shooting Controllers". Virtual Reality (VR) controllers that mimic the shape, weight, and recoil of real firearms are increasingly utilised in first-person shooter (FPS) games. The influence of these physical attributes on user performance and experience, however, remains insufficiently understood. This thesis explores how tactile augmentation and linear force feedback (LFF) influence shooting accuracy, subjective appreciation, and skill acquisition in VR FPS games.

A prototype VR controller replicating characteristics of commercial gun controllers was developed, and a series of user studies conducted to evaluate the impact of shape, weight, and varying LFF levels on participants' performance and subjective assessments. Results indicated that while adding gun-like shape and weight did not significantly alter immediate shooting accuracy, it enhanced subjective experience compared to standard VR controllers. LFF introduction improved immersion and engagement at low levels, with no additional benefits observed at higher levels.

Higher LFF levels accelerated participants' learning curves, allowing them to achieve maximum aiming proficiency more quickly. These findings suggest VR controller manufacturers could optimise user satisfaction while reducing production costs and technical complexity by implementing lower LFF levels. Nonetheless, higher LFF levels might benefit training scenarios where rapid skill development is prioritised.

This research highlights the necessity for standardised reporting and characterisation of LFF in future studies to enhance reproducibility and enable comparisons across various VR controllers and findings. The insights advance understanding of haptic feedback in VR controllers and guide development of future devices for both gaming and professional training purposes

6.1 Introduction

The rapidly advancing field of virtual reality (VR) emphasizes immersion and realism to enhance user experiences. A key factor in enriching presence within VR systems is the accurate simulation of natural forces during virtual environment interactions. This section examines the impact of recoil (the backward movement experienced when firing a projectile or exerting force) on user experience in VR. Precise replication of recoil enhances immersion, offering authentic experiences and narrowing the gap between virtual and real-world interactions.

As demand for more immersive VR experiences grows, commercial solutions have emerged to replicate recoil with increasing realism. VR controllers are transitioning from basic vibrotactile feedback to advanced kinesthetic feedback mechanisms. Controllers specifically designed for first-person shooter games, such as Strike VR [160], MAG P90 [161], Delta Six [162], and Haptec [163], emulate the shape, weight, and recoil of real firearms to increase realism and presence. Nevertheless, the effects of these realistic controllers on user performance remain inadequately documented, with sparse understanding of their effectiveness compared to passive controllers like the Wii/NES Zapper [164]. Furthermore, both commercial and research prototypes frequently provide insufficient details about their haptic characteristics, complicating systematic evaluation and comparison.

This research examines how controller properties influence general VR users' performance (those without specialised precision aiming skills) in virtual environments (VE). The study primarily focuses on assessing the effects of shape, weight, and linear force feedback (LFF) on VR controller performance, initially investigating a controller with only shape and weight before evaluating the addition of LFF to this configuration.

This research comprises three distinct studies, each examining specific elements of VR controller design. The initial study assesses the impact of a gun-shaped controller with realistic weight on shooting performance compared to a standard VR controller. The second study evaluates the incorporation of linear kinesthetic feedback into the gun-shaped controller, analysing how various LFF intensities affect performance and preferences. The third study examines LFF influence on learning curves in shooting tasks, specifically investigating how different LFF levels affect maximum aiming distance achievement.

The outcomes indicate that higher LFF levels enable users to achieve their peak distance skills with fewer errors, highlighting LFF’s potential to enhance immersion and reduce errors in progressively challenging tasks. Interestingly, while the results show minimal impact on immediate precision aiming performance (despite potential fatigue from weight and recoil) they also reveal improvements in users’ subjective experience of the interaction.

6.2 Related Work

Controllers that realistically mimic the shape and weight of real-world objects enhance learning times, presence, and engagement. Studies on augmented cues in mobile game-based learning and presence in video games support this observation [165], [166], [167]. Integrating kinesthetic feedback improves hand-eye coordination [168], boosts performance, and potentially reduces skill acquisition time [169]. However, downsides exist; excessive recoil can impair aiming accuracy [170], [171], and lead to fatigue and injuries [172].

Tethered vs. Untethered Kinesthetic feedback Devices.

Tethered kinesthetic feedback devices, such as the 3DSystems Phantom [173] and HAPTICS’ Virtuouse [174], offer high accuracy and precision. These devices serve primarily in VR training environments for tasks requiring precise linear kinesthetic feedback, including needle insertion [175], surgical training [176], and teleoperation [177]. Their high cost and substantial maintenance requirements make them more suitable for research or industrial applications rather than consumer-grade VR. Additionally, physical tethering often limits operational range, challenging the development of untethered VR headsets.

Untethered interfaces typically compromise accuracy or haptic fidelity to enhance portability. Exoglove systems such as the Dexmo [63] exert active forces on users’ fingers, with studies showing participants using an exoglove in VR archery achieving higher accuracy compared to those without such assistance. Alternative approaches include asymmetric vibration, utilised by devices like Waves [178] and Traxion [179], which, despite minimal force magnitude, inadequately replicate recoil effects. Electric Muscle Stimulation (EMS) offers robust kinesthetic feedback, exemplified by simulating impacts like Punches [180]; however, it currently lacks capability to provide controlled, precise linear kinesthetic feedback vectors.

Task-specific Untethered Controllers

Controllers such as NormalTouch [34] and Haptic Revolver [181] employ mechanical actuators to manipulate physical shapes, providing spatially-registered haptic feedback directly to the user's fingers. This setup aims for a more accurate and realistic representation of virtual objects. User studies with these controllers have demonstrated improvements in haptic feedback accuracy and fidelity over traditional visual-only and vibrotactile feedback systems.

Additionally, CLAW [32] controller enhances immersion by enabling users to grasp virtual objects, interact with surfaces, and trigger actions using kinesthetic feedback and actuated movement specifically designed for the index finger. The integration of controller arm movement and vibrotactile rumble simulates a realistic recoil effect, a feature not typically available in commercial controllers. However, the impact of this recoil simulation on user performance remains unaddressed.

Shooting Precision and Feedback in FPS VR Games

Companies such as Strike VR [160] and Haptec [163], have commercialised advanced controllers equipped with linear kinesthetic feedback (LFF) aimed at multiplayer and arena-style games. While these controllers aim to increase user immersion, their actual impact on user performance remains undetermined.

Shooting precision in FPS VR games involves two main components: visual aiming and proprioception, which includes gun-holding and posture. Research indicates that providing feedback on performance (KP) or outcomes (KR) can significantly enhance user performance [182], [183], [184], [185]. This is particularly relevant in FPS VR games, where the view camera (i.e., the user's head) is not directly aligned with the controller (i.e., the hands) [234].

Linear Kinesthetic feedback (LFF) in 2D and 3D Interfaces

Pointing tasks within 2D graphical user interfaces (GUIs) have been a significant area of research since the development of Fitts's law [186], [187]. . Studies have particularly explored the role of linear kinesthetic feedback (LFF) in enhancing 2D interfaces. Akamatsu et al. compared auditory, tactile, and visual feedback, discovering that haptic feedback facilitated faster motor responses than either audio or visual feedback [188]. Further research by Akamatsu, involving a multimodal mouse equipped with LFF and drag force control, found

that kinesthetic feedback could reduce both stop time and target selection time after the cursor ceased movement [189]. Additionally, Cockburn et al. demonstrated that tactile feedback could decrease average target acquisition time in 2D interfaces [190]. However, these studies primarily provided feedback before user selection (feedforward) and focused exclusively on 2D environments.

Expanding into 3D environments, Bhagat et al. investigated user performance using a wireless infrared controller modified to provide LFF [191]. Their findings indicated that 3D selection tasks were completed faster with LFF compared to visual or auditory cues alone [169]. Subsequent studies with similar setups reported only modest improvements in performance attributed to haptic technology [192]. The diversity in hardware implementations across these studies has presented challenges in reproducing results consistently. While there is extensive research on VR applications in military training, focusing mainly on skill transfer to experienced shooters, the impact on general performance remains less explored [191], [193].

Previous research on haptic feedback applications for VR aiming has yielded mixed results, often influenced by the differing levels of feedback implemented across studies [169], [194]. A significant challenge arises from commercial VR gun controllers that do not disclose their feedback intensity levels, coupled with a lack of standardization in reporting these metrics within the academic literature [182], [183], [184], [185].

Measuring Recoil Effect in Training

Controllers designed to replicate recoil typically utilise electric linear mechanisms with a moving mass. Although these systems produce weaker force magnitude compared to actual firearm recoil, they still generate noticeable impact force along the controller barrel. To determine appropriate levels of kinesthetic feedback for experiments, existing methods for measuring linear kinesthetic feedback (LFF) were reviewed. Two primary methods prevail in ballistic measurements: force readings via load cells and displacement tracking using accelerometers. The former was selected due to more extensive documentation available on the design process of measuring rigs in the literature [195], [196], [197]. Measuring force enables comparisons with other prototypes and existing controllers, with prior research indicating that peak force is a more representative parameter for quantifying the recoil experienced by users [195]. The setup is designed to capture both impulse and peak force values.

While injuries from game controller vibrations remain uncommon, they have been documented [198] thus, characterizing and limiting LFF to safe levels is. Individual physiological differences influence these effects [172]; however, there are established guidelines to mitigate risk. For example, Spine et al. suggest a maximum recoil limit of 13.33 N/s to prevent injuries typically associated with real firearms [171]. Additionally, the UK's Health & Safety Executive recommends a daily maximum vibration exposure limit of 2.5 m/s². These guidelines informed the experimental attachment design.

This literature review highlights both advancements and obstacles in haptic VR controller development. It underscores the need for further research to elucidate effects of various controller properties (shape, weight, and linear kinesthetic feedback) on user performance and immersion. By synthesising findings from past studies and undertaking new investigations, this chapter provides insights into the design and application of haptic controllers in first-person shooter VR games and other contexts.

6.3 Recoil VR

Three user studies assessed the effects on user performance within a virtual reality (VR) shooting game, where performance was measured by the highest score achieved during an aiming exercise at bullseyes placed at various distances. The first study evaluated shooting performance using a gun-shaped controller with realistic weight compared to a standard VR controller. The second study explored the benefits of integrating linear kinesthetic feedback (LFF) into the gun-shaped controller. The third study focused on the impact of LFF on the learning curve in the VR shooting experience.

6.3.1 Haptic Controller Design

A linear kinesthetic feedback attachment was developed to emulate the weight distribution, shape, and LFF characteristics of existing VR recoil controllers. This section outlines the design and operation of the attachment, along with the test bench setup used to characterise the impulse and peak force, which relate to the perceived intensity of recoil. Additionally, a hand dynamometer was designed to measure handgrip strength, aiding in the assessment of user fatigue, and contextualising results with objective parameters. Design constraints for the haptic device included:

- **Compatibility:** The attachment integrated with the original HTC Vive controller, leveraging its built-in tracking and buttons. This compatibility ensured that participants in the user study were familiar with the controller operation, requiring minimal instruction.
- **Form Factor:** The device replicated the external shape of a pistol while maintaining minimal enclosure to avoid occluding the IR receivers on the controller, thus preserving optimal tracking performance.
- **Actuation Mechanism:** The mechanism delivered high impulse force and demonstrated durability to withstand numerous trials, as required in the user study.
- **Safety Measures:** The design ensured participant safety, with maximum impulse force remaining below recommended levels and safe mechanisms in place during VR headset use.
- **Visual and Physical Alignment:** The visual representation in VR aligned with the physical shape of the device, serving as a passive haptic prop.
- **LFF Variability:** The attachment reproduced various levels of LFF to meet different experimental conditions.
- **Weight Considerations:** The total weight of the device was comparable to commercially available VR pistol controllers.
- **Control Integration:** A microcontroller was integrated for communication with Unity to control the pressure and duration of the impulse.

6.3.1.1 Haptic Attachment Design

The design used the original HTC Vive controller as a foundation due to its ergonomic resemblance to a pistol grip, with the top ring serving as an attachment point. This setup leveraged the controller's built-in tracking and buttons, focusing efforts on developing the linear kinesthetic feedback (LFF) attachment.

Commercial high power haptic recoil devices typically employ electromagnetic linear actuators that allow for precise force adjustments and nuanced haptic effects. However, these systems present technical design and control challenges, which would divert from the primary research

focus on LFF effects. Consequently, alternative solutions were explored that could offer rapid response and substantial impulse force without the complexities of electromagnetic systems.

A pneumatic piston with a weighted end was selected as the most effective substitute. The pneumatic system provided several advantages: simpler design, easier implementation, quick response times, and powerful impulse forces that effectively simulated the characteristics of electromagnetic actuators. The system's reliability, low maintenance requirements, and flexibility to adjust parameters ensured precise control over force output.

The attachment incorporated a pneumatic cylinder system that actuated a weight attached to the piston end, thereby increasing kinetic energy. For user safety and to prevent entrapment, this mechanism was housed within a cylindrical aluminium tube. A metal cap at the enclosure end aided in converting kinetic energy into a sharp, momentary impulse when the piston head collided with it. The completed device weighed 950g, comprising 470g for the controller and 480g for the recoil attachment, comparable to the MAG P90 VR Pistol controller's total weight of 500g plus the controller weight.

An SMC Double Action Roundline Cylinder, model CD85N20-125-B, served as the pneumatic piston, coupled to a 5/2 electrovalve (VUVG-BK10-B52) and powered by a 24V 0.2A power supply. Modifications to the pneumatic components and connectors minimised airflow restrictions. A Bambi air compressor, models 150/500, supplied up to 8 bars of pressure. Airflow from the compressor to the electrovalve used 4mm diameter pipes, while 2mm diameter pipes connected the electrovalve to the piston.

Control was achieved via an Arduino Nano microcontroller paired with a TIP120 transistor circuit, which operated the electrovalve. This configuration permitted precise regulation of pressure and impulse duration, allowing adjustment of various LFF levels during the user study. Communication with Unity was facilitated through a COM port, operating at a baud rate of 2000000.

Principle of operation

The actuator maintained the pneumatic piston in a retracted position. Upon activation, the piston extended under specified pressure, propelling the attached weight forward at a velocity dictated by this pressure. The motion caused the weight to collide with the enclosure end, generating an impact force directed towards the user holding the controller. A relief valve

restricted the piston’s return to its starting position, ensuring it moved back at reduced speed to minimise inertia.

The recoil haptic attachment featured the piston rigidly mounted at one end, with an end-of-travel cap fixed at the opposite end, providing minimal elasticity. This configuration resulted in a brief collision time, producing significant deceleration and substantial impact force over a short duration. To determine the average impact force when the moving mass attached to the piston collided with the end cap, the impulse-momentum equation was rearranged, estimating a collision duration (Δt) of 1ms.

$$F_{impact} = \frac{M \times \Delta v}{\Delta t}$$

This model estimated impact force to range from 739.35N at 2 bars to 1478.70N at 8 bars. However, calculating average impact force is a simplified approach that neglects factors such as friction, airflow resistance, air compressibility, damping, elasticity, and time-varying forces. These factors affect system dynamics and user-experienced force, necessitating verification and testing of the actuator’s actual response under controlled conditions using sensors like load cells to obtain accurate force measurements during operation.

6.4 FARE-KE & VR RECOIL

The characterization of the VR Recoil controller presented several technical challenges that later informed key elements of the FARE-KE framework. When designing this haptic device, we needed to make decisions regarding measurement approach, sensor selection, and data acquisition methodology.

To validate the previously presented average force calculations and characterize the device, we constructed a measuring setup using V-slot linear aluminum extrusions. This design choice was influenced by the ready availability of materials and the previous literature review done on the topic. V-shape wheels that allowed the controller to move freely in a horizontal direction while maintaining alignment with the force sensor were used.

For sensor selection, we evaluated the measurement requirements for characterizing pneumatic impact forces and selected a Phidget button load cell (0/200kg). This choice aligned with previous literature on ballistic impact measurement [248] and principles described

by Spine et al. [171], which proved effective for measuring impulse forces on a single axis. The experience with this sensor type later informed our understanding of appropriate sensor selection based on force modality. In hindsight a beam load cell could have potentially suffice to perform this characterization.

Data acquisition was accomplished using a INA125 instrumentation amplifier, with signals logged via a Picoscope C9000. This setup represented the available equipment at the time and provided high-resolution data capture capabilities that ensured accurate measurement of the impact force between the moving mass and the enclosure’s end cap. The load cell was calibrated using a series of known weights, establishing a linear relationship between applied force and sensor output—a procedure that would be further refined in subsequent projects.

The measuring testbench design enabled us to validate the calculated average force values and ensure that the pneumatically generated impact force was within a safe and desirable range for users. The obtained data allowed us to compare curves with existing data [199] and assess LFF effects against related guidelines [171]. An additional advantage of this measuring system was its independence from user pose or grip, ensuring accurate evaluation of haptic actuator performance regardless of how users held the device.

The insights gained from this characterization process informed the subsequent user studies detailed in the following sections and contributed to our understanding of effective measurement approaches for kinesthetic haptic devices. The challenges encountered and solutions developed during this project would later be systematised in the FARE-KE framework, reflecting a progression from specific implementation to generalised methodology.

6.4.1 Kinesthetic feedback Modelling

To set the kinesthetic feedback levels for the user study, the measuring setup described previously was employed. The initial phase of the setup involved defining the actuation limits for operation and assessing the accuracy of the model predictions. While the model projected output forces of 739.35N at 2 bars and 1478.70N at 8 bars, actual measurements indicated substantially lower forces. At the 2-bar setting, the peak force recorded was 337.5N, lasting approximately 1 ms. When utilising the maximum pressure available from the compressor (8 bars), the peak force observed was \sim 500N, also lasting around 1 ms.

These observed forces were significantly lower than the model's estimates. This discrepancy can largely be attributed to the limitations of the 2mm tubing used in the electrovalve circuit and the piston's capacity to evacuate air from the opposite end, factors which were not fully accounted for in the initial modelling. This underestimation highlights the importance of practical testing in conjunction with theoretical models to accurately gauge system performance under operational conditions.

The operation of the haptic prototype did not exhibit the quasi-linear response with pressure variations as initially suggested by the simplified model (Figure 105 Left). To refine the understanding and adjust for practical user studies, we employed the measuring setup to identify an appropriate mid-range stimulus point between the minimum and maximum operational limits of the actuator.

Consequently, a middle linear kinesthetic feedback (LFF) setting at 3 bars pressure was introduced, resulting in a peak force of 412N, which lies at the midpoint between the peak forces achieved at the other two settings. The collected data was used to compute the impulse for each condition; 2 bars generated a 0.028 Ns impulse, 3 bars a 0.035 Ns impulse, and 8 bars a 0.042 Ns impulse. For clarity in discussion and analysis, the settings corresponding to 0.8, 0.3, and 0.2 bars are referred to as HIGH, MED, and LOW, respectively

It is critical to note that the short duration of the measured responses as well as the high force values is attributable to the inelastic impact dynamics of the setup. A real user handling the device will introduce a more elastic response due to human factors, although the final impulse (i.e., the total force over time) will remain consistent across different operational conditions.




Figure 104. left simplified model of the haptic prototype to estimate force impact vs Pressure. Right impulse response from our three calibrated settings.




Figure 105. image of the layout of the virtual environment with all the bullseye shown as well as the platform where the participant w.

To monitor fatigue and mitigate potential adverse effects during the experiments, each participant's hand grip strength was assessed. A precision digital hand dynamometer was designed using CAD software and produced with a MakerBot 3D printer with PLA filament. For accurate force measurement, the dynamometer was equipped with a Phidget 0/50 Kg load cell.

The same INA125 integrated circuit (IC) amplifier used in the test bench was employed to amplify the signal from the load cell, ensuring consistency across the measurement systems. The same calibration procedure was followed for the dynamometer as had been established for the testbed, which allowed for reliable and comparable force measurement across different experimental setups.

6.5 User Studies

6.5.1 User Study 1: replicating shape and weight.

Previous research indicates that tactile augmentation, which involves replicating an object's shape and weight, can enhance immersion in virtual environments. However, the added weight of the device used for tactile augmentation might negatively impact user experience and performance. This study evaluated the side effects of the LFF attachment in passive mode (without mechanical actuation) on user performance, compared to a standard controller.

6.5.2 Experimental setup

In this user study, participants wore a VR headset and performed a shooting task within a virtual environment. This environment consisted of an open field where users were positioned

on an elevated platform, enabling them to engage targets at various distances and angles (up to a maximum of 45 degrees). The virtual platform provided cues to assist participants in maintaining consistent positioning throughout the experiment.

Participants were presented with a 3x3 grid of targets, each appearing sequentially at four different distances (25, 35, 45, and 55 meters from the starting point). The central target aligned with the user's eye level, measured at the beginning of the experiment, while surrounding targets were positioned 2.5 meters from the central target, with a 0.5-meter gap between adjacent targets. Each target appeared twice, resulting in 72 trials per condition (4 distances x 9 targets x 2 repetitions). A 5-second countdown preceded the task, and an in-game questionnaire was administered after each set of 72 trials for each condition. The virtual environment was created using Unity 2017, and participants used an HTC Vive headset.

6.5.2.1 Method

This initial user study assessed performance differences between two conditions: a controller equipped with the LFF attachment in passive mode (without pneumatic activation) and a standard HTC controller. A within-subjects design was adopted for the experiment, with the order of conditions counterbalanced among participants. During the study, participants aimed and fired at sequentially appearing targets, randomly positioned at various distances. Each target was displayed for 3 seconds, followed by a 1.5-second pause before the next target appeared.

Successful hits were acknowledged with audio feedback, and a visual sphere indicating the ray collision point was displayed. This setup used ray-casting techniques to simulate the shooting mechanism, eliminating the influence of gravity on the bullet's trajectory.

Dependent variable indicators of performance in FPS games [200] were automatically recorded; Time per shot (TPS) was measured as the time elapsed from the moment a target appeared until the participant pressed the trigger. A Boolean value, 'Hit', registered whether the target was struck, while hit distance (HD) quantified the distance in meters between the hit point and the target's centre. Following each task, participants rated their experience using a 7-point Likert scale on three questions designed to assess their enjoyment of the activity:

- Enjoyment of the experience (Q1)
- Perceived arm/hand fatigue (Q2)

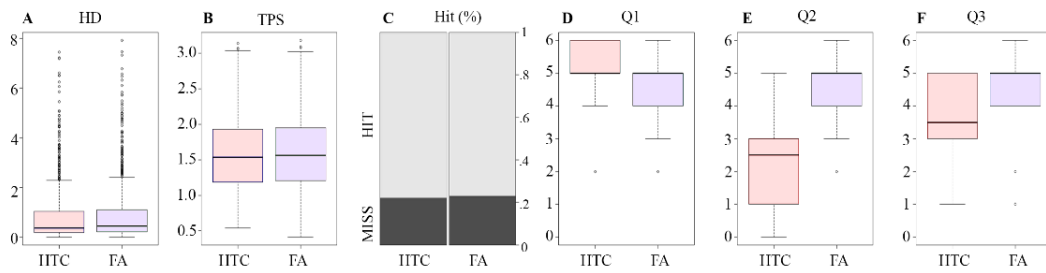


Figure 106. From left to right Distribution of Hit Distance results on metres to the centre of the target. Time per shot on seconds, % of hit and miss per condition. Right boxplot shows result from Q1, Q2 and Q3 from a 0-7 Likert scale.

- Perceived effect of the controller on aiming (Q3)
- Preferred condition (Q4).

Each testing session commenced with participants completing a background questionnaire to collect demographic data and assess prior experience with similar tasks. This was followed by a concise introduction to the experimental task. Participants were informed about task objectives, shown how to hold the controller using a two-handed grip (e.g., weaver stance), and instructed on using the sights for accurate target acquisition. They were told to begin shooting whenever ready and informed they would have only one attempt per target.

A preliminary training task featured 36 targets at the nearest distance of 25 meters to familiarise participants with the shooting mechanics. Following this, participants proceeded to the main part of the study, which involved two experimental conditions: using the passive LFF attachment or a standard HTC controller. To minimise fatigue, a mandatory 5-minute break without the VR headset was enforced between these tasks.

6.5.2.2 Results

The study aimed to determine if tactile augmentation could improve immersion in virtual environments without negatively impacting performance. Various performance indicators were analysed, including accuracy (Hit Distance), time per shot (TPS), hit rate, and participant responses to a post-experiment questionnaire assessing enjoyment, perceived fatigue, and the perceived effect of the controller on aiming.

The statistical analysis of the accuracy factor (Hit Distance, HD) indicated a similar trend across both conditions as depicted in Figure 107 A. The F-test yielded no significant difference between the two conditions ($F(1,13)=1.161, p > 0.1$). Similarly, the time per shot (TPS) results, illustrated in Figure 107 B, also showed no significant variation

($F(1,13) = 0.016, p > 0.1$). Participant expertise did not significantly influence TPS or hit rate, as demonstrated in Figure 107 C. Repeating the analysis with the group segmented by experience confirmed the consistency of these results.

Regarding subjective measures, the analysis of enjoyment (Question 1, Fig. 6.D) and the effect on aim (Question 3, Fig. 6.F) showed no significant difference between the controller conditions. However, the perceived fatigue (Question 2, Fig. 6.E) revealed a significant disparity ($Z = -1.8823, p < 0.05$), indicating that participants using the LFF reported greater fatigue post-experiment. Preference data showed that 10 out of 14 participants (71.4%) favoured the modified controller with the LFF attachment in passive mode (Question 4).

The findings reveal that while tactile augmentation using the LFF attachment in passive mode did not lead to expected improvements in performance, it did not adversely affect performance either. Participants using the LFF reported increased perceived fatigue, yet their aiming accuracy remained consistent throughout the tests, with no significant differences observed across various performance metrics. A majority of users demonstrated a preference for the passive tactile augmentation provided by the LFF, aligning with existing research that suggests such features can enhance user enjoyment and immersion. However, this preference might be influenced by the novelty of the device. Given that the passive implementation of the LFF did not affect performance outcomes, further investigation is necessary to determine if activating the LFF feature could introduce additional benefits, particularly regarding the impact force generated by the pneumatic system.

6.5.3 User Study 2: Linear Kinesthetic feedback vs Passive

In this second experiment, the objective was to explore the effects of incorporating active Linear Kinesthetic Feedback (LFF) into a controller that mimics the shape and weight of a real firearm on user performance. While the impact of LFF in virtual reality environments has been explored previously, the specific effects of integrating LFF into a controller designed to replicate the actual form and weight of a pistol have not been thoroughly investigated.

6.5.3.1 Methodology

The virtual environment and tasks were consistent with those used in the first study, with the exception that participants were tested under four different conditions: passive LFF

attachment, and three levels of active kinesthetic feedback (LOW, MED, and HIGH). Each condition was assessed through 72 trials, and performance was measured using the same variables as before: Time Per Shot (TPS) and Hit Rate.

To ensure more uniform measurement across various target distances, the error angle (EA) metric was utilised instead of distance to the centre of the target, following the methodology outlined by Kopper et al. [253]. A full Latin square design was used to counterbalance the order of the conditions for all participants.

The training and introductory procedures closely mirrored those from Study 1, with slight adjustments to accommodate the different feedback conditions. The questionnaire for this study included specific questions focusing on the effects of the active LFF:

- (Q5) asked participants whether the action in VR matched the feedback received.
- (Q6) asked participants to rate their enjoyment of the experience.
- (Q7) asked participants to rate the perceived realism of the feedback.

Additionally, a final forced-choice question (Q8) was included, where participants were asked to state their preferred feedback condition. To monitor physiological responses, hand grip strength was measured before and after each shooting block using the methodology described by Mathiowetz et al. [201]. This measurement was repeated eight times throughout the experiment, providing data to assess potential fatigue or other physiological effects on the participants' hands due to the various feedback conditions

6.5.3.2 Results

The analysis of error angle (EA) revealed minimal differences, with the largest discrepancy between conditions showing a mean average of 0.07 degrees (Figure 108 A). Statistical tests showed no significant main effects on EA ($F(3,13) = 1.01, p > 0.1$). Time per shot (TPS) and

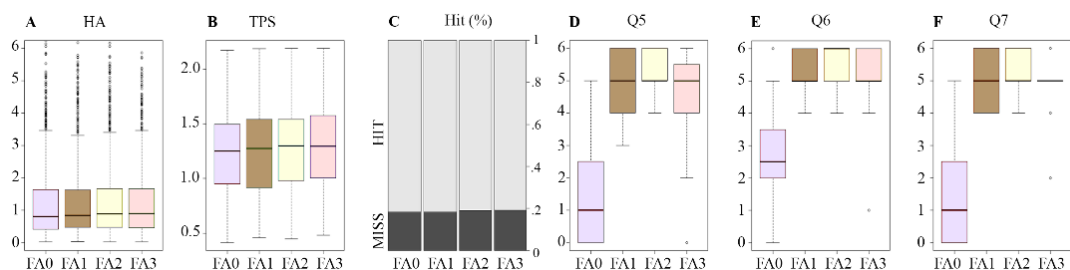


Figure 107. From left to right Distribution of Hit Distance results for experiment one. Time per shot for the study 1, % of hit and miss per condition. Right boxplot shows result from questions 1, 2 and 3

hit rate, illustrated in Figure 108B and Figure 108C respectively, also remained consistent across conditions, without significant differences. A slight trend in TPS was observed, increasing from an average of 1.238 seconds (SD = 0.383s) in the LFFa condition to 1.281 seconds (SD = 0.377s) in the HIGH condition, suggesting that higher impulse levels might slightly slow down shooting times. However, hand grip strength measurements showed no significant variations across conditions.

Question 6 results indicated a significant difference ($p < 0.005$) between the three LFF conditions and the passive attachment, suggesting that temporal delays in feedback did not affect perception adversely. Delays between trigger action and the receipt of LFF are known to potentially impact the sense of agency [255], [256], but the force magnitude (LOW, MED, HIGH) did not significantly affect participants' perceptions, as indicated by non-significant differences in paired tests. This finding implies that while some level of LFF can enhance the perception of action-consequence matching, the specific force level does not critically alter this perception. Participant enjoyment (Q7) and perceived realism (Q8) analyses also showed statistical differences ($p < 0.01$), with higher ratings for conditions involving LFF, according to paired analysis. This underscores that LFF can significantly enhance the user experience.

Finally, responses to the forced-choice question (Q9) suggest that high levels of LFF are not necessary for positive user experience, with six participants preferring the LOW condition, three preferring MED, and only two favouring the highest LFFa setting. This preference distribution indicates that the benefits of tactile augmentation may not require high forces of LFF to be perceived effectively.

In conclusion, incorporating LFF into a controller that mimics the shape and weight of a real firearm can enhance user performance, enjoyment, and realism. However, the specific level of LFF does not seem relevant, as lower levels were preferred, indicating that high forces of LFF may not be necessary to realize the desired benefits of tactile augmentation.

6.5.4 User Study 3: Assessing the Impact of Kinesthetic Feedback Levels

The aim of this experiment was to explore potential carryover effects and changes in participants' learning curves. The experiment's design was tailored to progressively increase difficulty based on each participant's skill level to determine the maximum distance they

could accurately target under each condition. The virtual environment and hardware setup were maintained consistently with previous studies, with modifications only to the experimental method.

6.5.4.1 Methodology

To assess the impact of Linear Kinesthetic Feedback (LFF) on user performance, the target presentation within the virtual environment was adjusted using a staircase design. Targets initially appeared at a distance of 25 meters from the participant, a starting point based on the high hit rate (96%) observed at this distance in Study 1, ensuring accessibility for all skill levels. If a target was missed at this initial position, the distance was reduced to a minimum of 8 meters to accommodate varying skill levels.

A "three up, one down" rule was implemented for adjusting target distances, where hitting three consecutive targets prompted an increase in difficulty by moving the next target 4 meters farther. Conversely, a single missed target resulted in a distance decrease. This adaptive design aimed to push participants to their skill limits, with the maximum distance constrained only by their abilities. The current distance was continuously displayed in the top right corner of the virtual environment to motivate optimal performance [202]. A pilot study with 4 participants determined the most challenging and feasible distance increments within the task duration for skilled participants.

As in previous experiments, four conditions (LFF attachment, LOW, MED, HIGH) were tested with 72 trials per condition, resulting in a final set of 72 distances (d1-d72). The experiment was introduced at the beginning and participants' hand grip was measured before and after each condition. No questionnaires were included in this study. Hand grip forces were measured, but their analysis confirmed previous experiment results and will not be discussed further.

6.5.4.2 Results

The results of User Study 3 were analysed using linear regression to understand the impact of Linear Kinesthetic Feedback (LFF) on user performance across different levels of feedback intensity. The regression model plotted the normalised shooting distance against the number of shots taken, effectively quantifying how users adapted to increasing distances based on the kinesthetic feedback provided.

Figure 107 illustrates the linear regression outcomes for the four feedback conditions, with each curve showing the rate at which participants' performance changed with distance. A higher slope indicates a quicker adjustment to the shooting task, representing a faster learning curve. The fitting of each model was robust, with a coefficient of determination (R^2) consistently above 0.93, suggesting that the models accurately explained the variation in user performance. The table below outlines the coefficients of the models for each feedback condition:

- HIGH: Displayed the steepest slope, indicating participants reached their maximum skill level more rapidly than in other conditions.
- MEDIUM and LOW: Showed moderate learning curves.
- PASSIVE: Exhibited the gentlest slope, suggesting slower adaptation to distance increases.

A Friedman test comparing the effects of different LFF levels on the mean normalised distances (from d1 to d72) revealed significant differences. The passive (PA) condition was notably less effective than both the LOW/MED ($p < 0.05$) and the HIGH ($p < 0.001$) settings. These findings suggest that LFF assists users in reaching their peak performance more swiftly than passive feedback, particularly at higher feedback intensities.




Figure 108. Linear model regression for the 4 conditions and r-square value, shows the models obtained and the fitting model for each. The X axis represents the normalized distance from 0 to 10 while the Y axis represents the number of shots taken.

Despite observable performance improvements, physiological measurements of hand grip strength showed no significant variations across conditions, indicating that the observed learning enhancements were not due to physical fatigue or strain.

6.6 Conclusions

This study explored the effects of shape, weight, and Low-Frequency Feedback (LFF) on 3D shooting tasks within virtual reality (VR) first-person shooter (FPS) games. Results and recommendations apply to devices with similar properties such as weight, impulse response, and peak forces. Key insights and recommendations derived from the research:

- **Enhanced Learning with Physical Properties:** Incorporating shape and weight analogous to conventional FPS controllers improved the learning curve. The integration of LFF enabled participants to achieve their maximum target distance skills more efficiently, underscoring LFF's role in facilitating task learning.
- **Subjective Experience and Immediate Performance:** While LFF did not directly affect immediate aiming outcomes such as hit rates and shooting time, replicating realistic gun properties (shape, weight, LFF) significantly influenced users' subjective experiences during the tasks.
- **Physical Exertion and Fatigue:** Hand grip strength tests showed no significant differences across different conditions, indicating that the LFF levels employed did not lead to notable fatigue. This contradicted expectations considering the experiment's length and the subjective fatigue reported by participants, especially in comparison with the standard 550g controller.
- **Participant Disposition towards LFF:** A minority of the participants (three out of forty-two) expressed a dislike for LFF, citing a general disinterest in shooting games. This preference might stem from anticipated reactions or individual biases.
- **Perceptions of Realism and Feedback Levels:** Post-experiment feedback indicated that many participants likened the sensation of LFF to using an airsoft gun. The consensus was that lower levels of LFF were sufficient to mimic real actions, and higher levels did not necessarily enhance realism or engagement.
- **Implications for VR Controller Design:** Impulse responses and peak forces did not detract from performance inform future VR controller designs in FPS gaming

and eGames. Future designs should isolate feedback mechanisms from sensitive tracking components like accelerometers to avoid adverse effects such as tracking drift.

- **Design Comparisons and Limitations:** The controller was designed to mimic commercially available VR controllers. However, notable differences include weight distribution, recoil force, trigger quality, materials, and grip shape. While these factors were not the focus, these variances can influence realism and user experience.

This study offers important insights for VR controller development, particularly for FPS gaming and eGames. By accurately replicating the shape and weight of real pistol controllers and judiciously applying LFF, user experience and learning outcomes are enhanced while better meeting individual preferences. Future designs should balance realism and user comfort while considering the limitations in replicating real pistol controllers.

7 Proxy Touch




Figure 109. (a) Hand gesture during knob interaction in virtual space. (b) Real hand gesture in real space. (c) Retargeting representation from virtual to real space the user perceives an object pin.

This chapter presents the initial steps towards developing a retargeting approach that capitalises on visual dominance over tactile and proprioceptive cues to create the illusion of pinching a virtual object between the user’s fingers in Virtual Reality (VR). By leveraging the robust visual dominance in VR environments [203], [204], [205], this study investigates mutual finger contact during pinch gestures as a means of providing haptic feedback.

The main objective is to explore whether the sense of agency effect can be measured for pinch interactions using a modified version of the Force Matching Paradigm (FMP). The FMP measures Somatosensory Attenuation (SA), a psychological effect where the perceived strength of a tactile sensation reduces when the brain identifies the sensation as self-produced rather than external.

The research employs an experimental setup with a modified force-matching device to investigate the SA effect in pinch interactions. The results confirm the existence of SA in pinch interactions and demonstrate that the central comparator mechanism responsible for attenuating predictable self-generated tactile signals extends beyond bimanual interactions.

The FARE-KE framework guides the characterisation process of the force-matching device. The principles and methodologies outlined in the framework adapt to investigate the psychophysical aspects of haptic perception, ensuring accurate measurement of force profiles during user studies.

The findings represent the first step towards a full implementation of a self-touch haptic illusion that requires no external proxies, utilising the participant’s body to recreate haptic sensations within a virtual environment.

7.1 Introduction

Previous approaches to provide haptic feedback in VR have avoided self-touch for object gripping. Most common methods employ mechanical or electrical systems creating a gap between index and thumb fingers for kinesthetic feedback [68], [206], or utilised cutaneous feedback [207], [208] to provide tactile sensations on the skin.

This chapter presents a novel retargeting approach capitalising on visual dominance over tactile and proprioceptive cues. Although Ban et al. previously suggested this approach in an augmented reality setting [209], the feasibility of users perceiving touch on a virtual object rather than their own fingers remains unestablished.

The results of this study confirm the SA effect in pinch interactions and demonstrate that the central comparator mechanism responsible for attenuating predictable self-generated tactile signals extends beyond bimanual interactions. These findings represent the first step towards a self-touch haptic illusion without external proxies. Moreover, the future work outlined sets the stage for designing a user study incorporating finger redirection techniques to further validate this novel haptic approach in virtual environments.

7.2 Related Work

7.2.1 Technologies for Kinesthetic feedback in VR

The main methods for delivering haptic feedback in virtual environments include kinesthetic [70] and tactile approaches [ref], with some systems using both [ref]. Haptic interfaces are typically controlled through devices like controllers, but there are also specialised devices like exo-gloves [210] and wearables for fingers [154], [211], [212], [213] that provide haptic sensations.

Recent research has explored controllers that dynamically change to influence weight perception [40], offer multiple interfaces to the user [181], or have shape-changing features to engage the user's fingertips [34], [207]. Many controllers designed for grasping tasks [36], [38], [214], [215] require the user to hold the controller, which limits free hand movement.

Exo-gloves enhance haptic interaction by freeing the user from having to hold a device, providing a more unencumbered and natural experience. These devices are divided into two main categories: passive and active. Passive exo-gloves use mechanical [ref], electrical [68], or

pneumatic jamming [216] to restrict movement by employing various forms of resistance mechanisms like brakes or jams, anchored to the user's body. Active exo-gloves, on the other hand, not only restrict movement but also feature sophisticated closed-loop control systems. These systems measure the force exerted by the user and dynamically adjust to offer a more realistic manipulation of virtual objects. While exo-gloves are promising for their immersive capabilities, they also come with challenges such as bulkiness and the high costs associated with their complex electromechanical systems.

A less intrusive method for delivering haptic feedback utilises finger interfaces, which mimic the sensation of handling an object by creating fingertip indentations that replicate pressure [207], [213], [217], [218], [219]. Other techniques include providing asymmetric vibrotactile stimuli to the fingers to simulate virtual forces or textures [101], [179]. Additionally, methods involving heat [220] and electro-stimulation [221] have been developed to enhance the capabilities of finger interfaces. While these interfaces are less cumbersome than exo-gloves, they still involve some equipment that can restrict finger movement, potentially reducing the sense of immersion.

Until now, using self-touch to provide haptic feedback in virtual reality (VR) environments has not been explored. The closest related work is the AirHaptics technique introduced by Ban et al. [209], which began to explore a similar idea. In their approach, a video see-through method allows the user's fingers to visually correspond with the surface of a virtual object when they touch each other. However, their study focused on the user's ability to judge the size of virtual objects (visual task) and did not explore the psychophysical impact of this illusion. This thesis aims to fill this gap by evaluating the potential of self-touch as a method for delivering haptic feedback in VR.

7.2.2 Techniques exploiting visual dominance in VR.

As Debarba et al. put it, "a better understanding of the limits of human perception can lead to new avenues for effective VR interaction" [222]. Essentially, by deepening the understanding of human perception, we can develop methods that use less physical technology yet still facilitate effective interaction in VR environments. This strategy stands in contrast to traditional methods that rely on specialised haptic technology for delivering feedback in VR.

The dominance of visual input in VR environments is well-documented, particularly when multiple sensory channels provide conflicting information, leading the brain to prioritize visual data. Research in VR interaction has shown that visual stimuli can influence how we perceive our own motor actions. This means that differences between what we see and what we do can affect both movement and tactile sensations, underscoring the predominance of visual feedback over proprioceptive senses [204]. This principle of visual dominance has been successfully leveraged to enhance VR experiences through redirection techniques. These methods subtly adjust the user’s physical movements to create convincing perceptual illusions.

Redirection techniques that utilise visual dominance have been integrated into various VR interactions, including walking [223], [224], [225], [226], arm and hand movements [227], and object manipulation [205], [228]. These methods involve subtle translational or rotational adjustments or the use of distractors to create the illusion that the user is moving through a large virtual space while confined to a smaller physical area. Additionally, these techniques have been used to simulate the sensations of grasping [228] and interacting with diverse objects [205], as well as perceiving different shapes [229], complex geometries [230], textures [231], and weights [232], all within a simplified physical setup compared to electromechanically actuated technologies.

The proposed approach leverages visual dominance to reduce the complexity of the physical setup required for haptic feedback in VR. Unlike previous studies, this method has not been explored in the context of pinch interactions or using the user’s own body as a passive prop. The approach was evaluated by analysing the psychophysical responses typically observed during self-touch and non-self-touch scenarios.

7.2.3 Somatosensory Attenuation

This study investigates the application of Somatosensory Attenuation (SA) as a psychophysical measure of subjective touch experiences. The goal is to assess whether SA can distinguish between the perception of touching an external object and the sensation of touching oneself




Figure 110. Depiction of the Force Matching Paradigm. (a) Presented Force, characterized by constant forces applied to the subject's left index finger using a motor actuator. (b) Matching Force produced by the subject's right index finger pressing a sensor directly above the left index finger (0 cm), which controls the force output on the opposing finger. (c) Matching Force generated by the subject's right index finger pressing a sensor positioned 25 cm away from their left index finger, controlling the force output on the opposite finger. (d) Attenuation index calculated as the difference between matched force and presented force.

Somatosensory Attenuation refers to the reduction of tactile sensations experienced during self-touch [233], [234]. This phenomenon is thought to serve as a protective mechanism, allowing the central nervous system to differentiate between self-generated signals and potentially harmful external inputs [235]. For example, when one touches their own arm, the sensation is generally milder compared to the touch delivered by another person or an object. This effect is believed to be due to the brain's prediction of self-generated contact, which reduces the intensity of the anticipated sensation [236], [237], [238]. This mechanism helps explain phenomena such as why individuals cannot tickle themselves [239]. Although this effect has been well-documented in psychological studies, its potential applications in Human-Computer Interaction (HCI) have yet to be fully explored.

7.2.3.1 The Force Matching Paradigm

The Force Matching Paradigm (FMP) is a widely used psychophysical method for quantifying Somatosensory Attenuation (SA) [240], [241], [242], [243]. The FMP involves applying a constant force to a subject's relaxed left index finger using a motor actuator (Figure 111a). The subject replicates this force with their right index finger on a force sensor positioned directly above (but not touching) the left index finger (Figure 111b). Subsequently, the user applies a matching force on a sensor placed at a distance from the hand (Figure 111c). The "attenuation index" is calculated as the percentage by which the matched force exceeds the presented force (Figure 111d, red line). This attenuation significantly reduces (Figure 111d, green line) when the matching force is applied from a distance, as self-touch is unlikely under this conditions. This effect relies on visual information, as actual skin contact is unnecessary to produce the self-touch conditions; rather, the position and distance between the fingers

creates this effect. This attenuation has also been observed using a rubber fake hand to create visuomotor conflicts that affect the sense of body ownership [235].

Hence, it is technically feasible to create visually induced conditions that create perceptions of self-touch and non-self-touch based on the distance between the index fingers and the point of contact. For example, if the right and left index fingers are separated by 25cm (indicating unlikely self-touch), but a fake hand visually suggests that the fingers are on top of each other, an illusion of self-touch can be created, leading to increased SA. Conversely, if the right and left index fingers are directly touching each other (actual self-touch), but a fake hand visually suggests that the fingers are 25cm apart, an illusion of non-self-touch can be created, leading to decreased SA [235].

Previous studies have shown that visual information can significantly impact physiological responses, to the extent that a discrepancy between observed and actual hand positions can lead to the perception of "alien actions" as one's own [204]. This study explores this effect further by examining the use of haptic feedback for pinch interactions employing self-touch in virtual reality (VR) and retargeting techniques. To adapt the Force Matching Paradigm (FMP) for pinch interactions, we noted several differences from the traditional setup. Prior research in this area indicates:

- The internal forward models that predict and mitigate sensory outcomes of movements is based on bimanual interactions involving both index fingers.
- The modulation of Somatosensory Attenuation (SA) is influenced by a defined distance between the hands, with distances such as 25 cm used to indicate non-self-touch and 0 cm to represent actual self-touch [235].
- This specified distance alone is enough to distinguish between self-touch and non-self-touch, showing that actual skin contact is not required for this differentiation. In the experimental setup, the matched force is generated by pressing a sensor positioned directly above (but not touching) the left index finger, as depicted in **Figure 2b**.

To examine the effect of SA in pinch interactions, it is necessary to determine if SA is observable under the following conditions: (i) the matched force is generated by the same hand, using the index and thumb fingers, (ii) the distance between the fingers is less evident, about 5 cm, set by a physical object positioned between the fingers, and (iii) actual skin




Figure 111. Device used for the Force Matching Paradigm.

contact occurs between the fingers, as opposed to a scenario where no skin contact occurs (i.e., a physical proxy is placed between the fingers). In the subsequent section, a user evaluation is conducted to investigate these three aspects.

7.3 FARE-KE & Force Measuring Setup

The development of the force measurement system for the Proxy-Touch project presented unique challenges that contributed to understanding of precise force measurement, informing several aspects of the FARE-KE framework.

When designing this experimental setup, the initial consideration was selecting an appropriate force sensor. While force sensing resistors (FSRs) offered a thin profile, our early explorations with these sensors revealed significant inconsistencies in readings due to variations in finger size and pressure application techniques across participants. Additionally, introducing an extra layer between skin contact points would have potentially biased the somatosensory attenuation measurements, a critical variable in our experiment. These limitations led us to explore more precise measurement alternatives.

We ultimately selected a Honeywell model FSG15N1A button load cell due to its compact size, higher precision. This decision reflected our growing understanding of sensor selection

based on measurement context rather than merely technical specifications—a principle that would later become central to the framework.

For signal conditioning and data acquisition, we employed an HX711 amplifier connected to an Arduino Nano microcontroller. This represented our first experience with this digital amplifiers, revealing their 80 SPS sampling rate limitation. While sufficient for this particular experiment where participants maintained consistent pressure for periods of seconds, this limitation informed our later understanding of sampling rate requirements for different haptic modalities. The implementation also required optimising microcontroller code to maximise effective sampling rate and data transmission, challenges that directly informed the framework’s data acquisition recommendations and provided implementations of the DAQ.

The experimental setup incorporated wrist and forearm supports to maintain consistent hand positioning and alignment between the force vector and sensor, reducing measurement variability. This experience highlighted the importance of mechanical stability in force measurement setups (particularly for human-subject experiments) and influenced the framework’s emphasis on proper alignment and support structures.

Signal processing included implementing a 10Hz Butterworth Lowpass Filter to eliminate high-frequency noise, a technique adapted from our literature review. This approach to filtering became a standard recommendation in the framework, particularly for measurements involving residual vibration elements.

Drawing inspiration from Kiltner et al. [235], we designed a closed-loop feedback system with a motor-controlled lever that could apply controlled forces to participants’ fingers. Significantly, we implemented a fully integrated control system where a single microcontroller managed lever activation, visual cue presentation, and force data acquisition. This unified approach enabled precise synchronisation between stimulus presentation and measurement, allowing us to account for system latency and accurately correlate applied forces with participant responses. This integration of timing-critical functions within a single control system later influenced the framework’s recommendations for data acquisition and experimental control.

7.4 Proxy-Touch SA Evaluation

In this section, we explore the possibility of using mutual contact between the index and thumb fingers to simulate the sensation of holding (pinching) an object, focusing on the implications for Somatosensory Attenuation (SA). We begin by detailing the adaptation of the setup for the Force Matching Paradigm, enabling the assessment of SA in pinch interactions. We then present a study validating the occurrence of the SA effect in pinch interactions, not only between different hands.

7.4.1 Experiment Design and Testing Procedure

In this study, we utilised a modified force-matching setup, as described in the previous section, to assess whether Somatosensory Attenuation (SA) occurs in pinch interactions. The experimental method was adapted from the procedures outlined in [235].

Participants were instructed to sit in front of the setup and place their dominant hand on the measuring device, which was adjusted for each participant to ensure comfort before beginning the experiment. In each trial, the setup initially exerted a Presented Force on the participant's relaxed thumb for 3 seconds (see Figure a). In the self-touch condition, the lever was promptly removed after presenting the force, allowing direct skin contact. Conversely, in the proxy-touch condition, the DC motor was disengaged, leaving the lever above the thumb without exerting any force, thus enabling participants to attempt to replicate the force without additional assistance from the lever. After the force presentation, a red light signalled participants to replicate the Presented Force using their index finger. The Matched Force was either directly applied by the participant's index finger in the self-touch condition or indirectly through pressing a 5 cm physical proxy placed between their fingers in the proxy-touch condition.

Participants were given 3 seconds, indicated by a blinking green light, to match the force in each trial. To ensure accuracy in the analysis, the average force was calculated during a specific time window from 1.5 to 2 seconds to exclude the initial period of force adjustment. The setup, including the presence or absence of proxy objects, was always visible to participants. Each participant underwent trials involving six different Presented Forces (1 N, 1.5 N, 2 N, 2.5 N, 3 N, and 3.5 N), presented seven times in a counterbalanced order, totalling 84 trials per participant. The entire experiment lasted approximately 40 minutes.

7.4.2 Participants

The study included 18 right-handed participants (5 females, mean age = 28 years old, SD = 4.39) with normal or corrected-to-normal vision and no hand or finger impairments. The local ethics committee approved this research, and participants received no compensation for their involvement.

7.4.3 Results

As we review the results of our study, we investigate the connection between the forces generated by the participants (Matched Force) and those applied externally by the lever (Presented Force). Figure d includes a dotted line indicating the ideal scenario where the Matched Force perfectly aligns with the Presented Force. If the actual force applied by participants exceeds this line, it indicates an attenuation effect, demonstrating that participants apply more force in the self-touch condition (self-related signals) compared to when an object is between their fingers (non-self-related signals). To further analyse the attenuation effect, Figure e shows the average difference between the Matched Forces and the Presented Forces across each experimental condition.

We applied the Shapiro-Wilk test to determine the normality of the data and found a non-normal distribution, prompting the use of non-parametric tests for further analysis. To compare the Matched Forces and attenuation indices under two different conditions (self-touch and proxy-touch), we conducted a Wilcoxon Signed-Ranks Test. The results indicated that the Matched Force was statistically significantly stronger during the self-touch condition than during the proxy-touch condition ($Z = -4.07, p < 0.001$). Figure 31d displays the mean values of Matched Forces relative to Presented Forces, illustrating a greater attenuation when participants matched the forces without an object between their fingers. Further analysis using the Wilcoxon Signed-Ranks Test on the mean attenuation index showed that the attenuation index in the self-touch condition was statistically significantly higher than in the proxy-touch condition ($Z = -3.74, p < 0.001$). Figure e presents the mean values of attenuation. In prior research, the Wilcoxon Signed-Ranks Test has often been utilised for

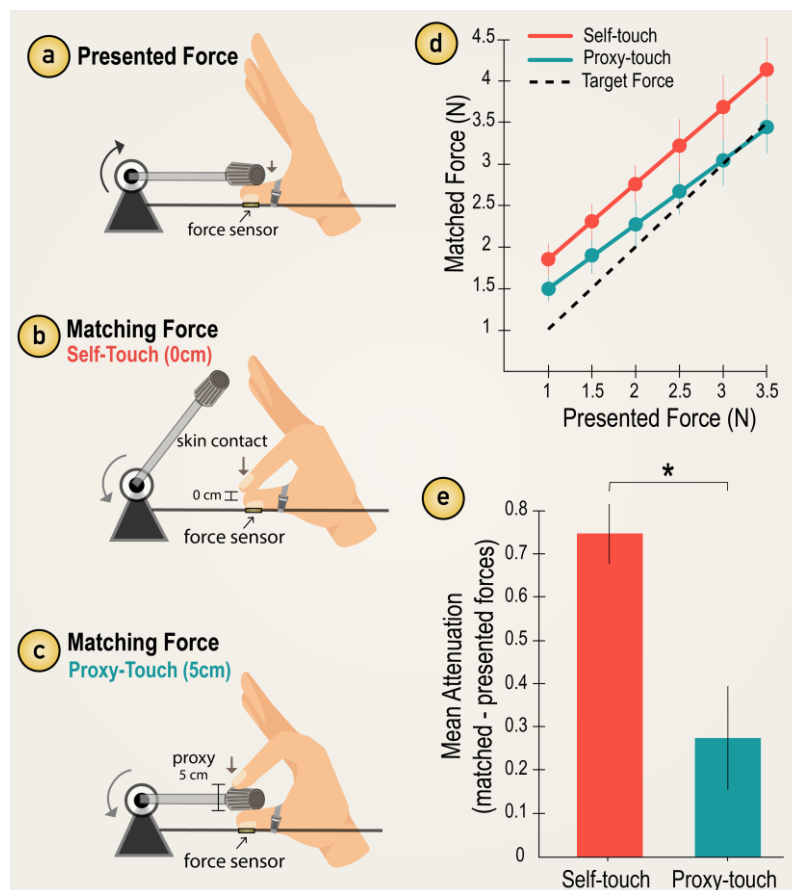


Figure 112. (a) Presented Force applied by a lever (1N, 1.5N, 2N, 2.5N, 3N, 3.5N) on participants fixed thumb. (b) Direct Matching Force applied by participants' index finger (skin contact). (c) Indirect Matching Force applied by the index finger with a physical proxy in between (5cm). In both b and c, the setup was either visible (directly in front) or not visible (from top view) by participants. (d) Plot for comparison of Presented Vs Matching Forces. (e) Mean attenuation Matched Force - Presented Force.

analysing non-parametric data associated with the Force Matching Paradigm (FMP) as referenced in [235]. Consistent with this approach, in the study, we used the Standard Error of Mean (SEM) to depict error bars in the plots when comparing experimental conditions with controls, adhering to Rule 4 specified in [244].

7.5 Conclusions

The results align with those of Kilteni et al. (2017), showing a more pronounced Somatosensory Attenuation (SA) effect when participants pressed against their own skin rather than against a physical proxy (see Figures d and e). Participants consistently applied stronger forces in the self-touch condition across all presented forces (1N, 1.5N, 2N, 2.5N, 3N, and 3.5N) compared to the proxy-touch condition. As depicted in Figure d, participants achieved higher accuracy in matching forces greater than 2.5N in the proxy-touch condition (green line), suggesting stronger forces were more readily discernible. However, in the self-touch condition (red line), more pronounced differences were observed between the presented and matched forces.

The modulation of the SA effect is affected by the distance between the fingers, as established by a physical object. Unlike previous studies utilising the Force Matching Paradigm (FMP) that typically did not involve actual skin contact, these results demonstrate that the distance between the fingers suffices to trigger the SA effect. Furthermore, when comparing matched forces with and without direct skin contact, variations in the SA effect were noted, thus supporting earlier research in this area.

7.5.1 Future work

This research establishes the foundation for a novel haptic technique integrating redirection strategies with self-touch, facilitating haptic illusions for virtual reality (VR) users without additional mechanical or proxy elements. Building on these findings, the next phase will involve designing and implementing a user study utilising a finger redirection technique termed "self-proxy." This method leverages visual cues to create the illusion that users are pinching a virtual object when their fingers contact one another.

The planned VR study aims to assess the level of sensory attenuation between the self-proxy condition and the conventional self-touch condition. The objective is to establish that the self-proxy method in VR can induce an SA response comparable to physically interacting with a

real object, confirming the efficacy of this innovative haptic approach. Ultimately, this work aims to enhance understanding of the SA effect within self-touch haptic techniques and contribute to developing more immersive and realistic VR experiences that integrate haptic feedback without extensive hardware or external devices.

8 HapticWhirl

Most haptic actuators currently available generate a single type of stimulus, restricting kinesthetic haptic controllers' ability to provide varied feedback. HapticWhirl addresses this limitation as an untethered VR controller delivering torque feedback and supporting a spectrum of haptic effects including collisions, vibrations, high-impact torque, and viscous kinesthetic feedback. The design features a flywheel generating moment of inertia, encased within a two-axis gimbal with slip rings enabling unrestricted rotations. The gimbal design allows flywheel position adjustment, modifying torque orientation and ensuring continuous rotation. A mechanical brake enables precise flywheel speed control, either slowed gradually to reduce gyroscope-related effects or stopped abruptly to create impact forces. By amalgamating these effects, the controller enriches the tactile dimension of VR environments.

This study evaluated a HapticWhirl prototype, focusing on modelling operational dynamics and torque output characteristics. The findings indicate potential to improve simulation of virtual object interactions and inertia effects in VR environments. This chapter details the design, implementation, and characterisation of the device.

Controllers generating torque at a distance from the hand present significant characterisation challenges, typically requiring 6-axis load cells difficult to access and implement. This chapter explores alternative methods for characterising these advanced devices.

The FARE-KE framework guides the characterisation process of the HapticWhirl controller. The framework enables accurate measurement of complex force and torque outputs. This chapter highlights connections to key elements of the framework, demonstrating its application to devices with multiple degrees of freedom and intricate actuation mechanisms.

8.1 Introduction

VR technology has advanced significantly toward lightweight, portable wearable headsets offering immersive virtual environments at high resolutions and refresh rates. Despite the pivotal role of haptic feedback in how users experience a Virtual Environment (VE), haptic controllers have not experienced a similar advancement. This lag stems from the limitation that most actuators generate only a single haptic modality, restricting the expressiveness and depth of tactile experiences they can deliver.

Kinesthetic haptics presents significant challenges, with limited integration into commercial consumer devices [245], [246]. This limitation stems from designing haptic devices for specific use cases, resulting in lack of broader applicability. These challenges become more pronounced for hand-held interfaces, which are naturally suited for untethered VR experiences but require more sophisticated feedback mechanisms.

A common challenge across various haptic devices is their ability to produce only a single haptic modality. Controllers like CLAW [250] and CapstanCrunch [247] deliver controllable stiffness during pinching interactions, while Thor’s Hammer [248] and Aero- plane [249] use propellers to generate forces and torques. Active force feedback exogloves such as CyberGrasp [250] and Dexmo [251] limit user movements to simulate collisions through mechanical links or Bowden cables. Hand-held devices with indirect actuation, including Shifty [252], Transcalibur [253], and Drag:on [254], provide feedback based on controller inertia but are limited to specific modalities. These examples highlight how existing haptic controllers produce a singular type of haptic stimulus, limiting the variety of haptic experiences offered.

This chapter introduces HapticWhirl: an untethered handheld device utilising a flywheel-based system capable of delivering multiple haptic effects through a single actuator. The device generates five distinct haptic modalities:

1. **Torque Feedback:** Direct forces created through controlled gimbal orientation changes
2. **Impact Simulation:** High-magnitude forces generated through rapid flywheel braking
3. **Vibration Feedback:** Low-frequency, high-amplitude vibrations created through oscillatory gimbal movements
4. **Inertial Effects:** Sensations of weight and resistance through controlled transparency delay
5. **Whirl Effects:** Complex rotational forces created through continuous rotational precession

The flywheel system operates within two 360-degree gimbals, allowing precise orientation in any direction at various speeds. Additionally, a clamp brake enables sudden stopping of the flywheel, enhancing the dynamic range of haptic feedback and creating impact sensations not possible with gimbal actuation alone.

This chapter provides a comprehensive account of the device’s design, implementation, and characterisation, establishing groundwork for future exploration. The investigation includes

integration within virtual environments and potential applications in force simulation, vibration recreation, and impact effects. Measurements of output torque, vibrations created using the brake, and vibrations generated by the inner flywheel validate the effectiveness of the proposed solution. The device follows open source and open hardware principles, enabling construction by anyone interested in haptic interfaces experimentation [255].

The following sections provide detailed exploration of the device’s design and characterisation. In response to constraints of existing haptic technologies, this work introduces a novel approach to multi-modal haptic feedback. The contributions include:

1. A design and implementation approach for a haptic controller delivering various haptic effects, enhancing the tactile dimension of VR environments and improving user immersion.
2. Detailed analysis of flywheel and gimbal combination operation and control, supporting further device development.
3. Thorough characterisation of the proposed device, supported by extensive measurements of controller output torque and vibrations, validating the solution’s effectiveness.

8.2 Related Work

This section discusses related work on gyroscope feedback devices and force feedback implementations for VR..

8.2.1 Torque Gyroscope Feedback Devices

Inertia can be generated with the angular momentum of a flywheel spinning at high speed. Changes in speed or adjustments to the rotation axis orientation create an output force orthogonal to the plane of the flywheel. The direction of the output torque depends on the wheel’s rotational direction and its re-orientation angle and speed. The device developed by Yano et al. offers two degrees of freedom and has been evaluated for object detection, guidance, and enhancing the user’s sense of movement [256]. Recent studies have further investigated various applications of the gyro effect, demonstrating how the output torque can induce a wrist twist during arm swings [257].

Altering the flywheel speed to create output torque around the axis of rotation has been explored as a means of delivering directional cues. GyroCube [258], [259] presents a method using three individual flywheels arranged in a hand-held cube. Directional force feedback can

be experienced by sequentially switching one actuator at a time. Some studies explored adding a mechanical brake to the flywheel to generate directional force feedback on wearable devices for directional guidance [260], [261]. Other examples use multiple gyros to counteract undesired torque effects when returning the gyro to the home position [262], [263]. The results from Amemiya et al. [260] suggest that an angular velocity change with sudden-start profile is more effective compared to sudden-stops.

For mobile applications, TorqueScreen [264] and GyroTab [265] provide force feedback for tablet devices. GyroTab uses a flywheel to provide torque feedback reactive to the user's movements, while the flywheel that is added to TorqueScreen can impose angular momentum on the tablet without user movement by rotating over one axis. A handheld controller developed by Winfree et al. [266] integrated the flywheel in a two-degree-of-freedom cage similar to the design from Yano et al. [256], but with improved control design. They provided thorough evaluation of the torque generated with the device and in-depth operation. Building on their work, this research explores how such forces can be applied in VR, evaluating the torque generated by the device and the impact of the mechanical brake.

Gyroscope designs focusing on VR interactions include GyroVR [267], which explored head and body worn flywheels to support motions like flying, diving or floating. Recent work explored how feedback created by a flywheel can correct posture of a golf club while swinging it using a single pivot axis point for the flywheel [268].

These devices typically employ the gyroscopic effect in a singular modality. HapticWhirl diverges by exploring multiple modalities through control of the flywheel and its orientation. This approach enables a broader spectrum of haptic experiences in VR, ranging from inertia simulation to directional feedback, expanding potential applications and user experiences within virtual environments.

8.2.2 Haptic Interfaces for VR

8.2.2.1 Motor Actuated Devices

For immersion in virtual environments, force feedback has been implemented to mimic different kinaesthetic experiences. Head mounted attachments like FacePush [269], use servo motors attached to the VR headset to generate push and pull forces. Thor's Hammer [248] uses propellers to induce ungrounded force feedback in a handheld haptic device, combining

six propellers to generate force feedback in three dimensions. With the strong and continuous force feedback generated, interactions like moving an object through water or simulating weight can be supported. The light weight of propellers makes them more suitable for VR applications and wearables compared to flywheels. Wind-Blaster [270] exemplifies a wrist-worn device providing force feedback using two propellers. PaCaPa [271] offers a simple mechanical handheld device to make indirect interaction with handheld tools in VR more realistic.

8.2.2.2 Weight Shifting

Several studies utilise weight shifting to generate haptic experiences for VR. SWISH [272] reproduces the physical sensation of fluid vessels using a mechanical construction to change the horizontal and vertical position of a mass in a cylinder. Shifty [252] controls only the vertical position of a weight proxy to generate weight-shifting feedback for handheld devices. Cheng et al. [273] used liquid-based haptic feedback to simulate realistic weight for virtual objects using pumps and water tanks. Drag:on [254] adjusts its surface area to vary resistance when the controller is moved, simulating the weight of objects moved in VR. Transcalibur [253] alters its shape and centre of mass to mimic the haptic feel of different weapon-type controllers.

8.2.2.3 Elastic Interfaces

Force feedback can be generated by varying the tension of elastic bands. ElasticVR [274] and ElastImpact [275] are wearable devices that generate impact forces instantly by releasing power stored in elastic bands in advance. The same principle applies in ElastOscillation [276] a device providing damped oscillation effect in 3 degrees-of-freedom. Similarly, Elastic-Arm [277] uses an elastic body attachment between the shoulder and user hand to provide passive feedback.

8.2.2.4 Exogloves

An alternative untethered design for limb force feedback uses the user's body as an anchor point for actuators. Cable driven designs such as Wireality [278] use a spring loaded cable system mounted on the shoulder, with cables attaching to multiple joints on the user's hand allowing for locking action when a collision in the VE is detected. Some exosuits [279], [280] use similar cable/clutch design integrating bowden cables on the suit or soft fabric creating motion constraints via passive actuation. Wearables with more advanced control designs [281]

include brushless motor actuators to the body attachment allowing for more realistic active force feedback effects. The controller presented in this chapter aims to provide feedback for continuous motion and impact by combining a flywheel that can rotate in 2 degrees-of-freedom with a brake mechanism.

8.2.3 Multimodal Kinaesthetic Feedback

Handheld devices for virtual reality, while providing broad workspace and ease of use, typically offer single mode haptic feedback such as vibration [282]. To enhance user immersive experience, multimodal haptic feedback is required, often necessitating various types of actuators to simulate different tactile sub-modalities like force, shape, texture, shear, or temperature.

NormalTouch and TextureTouch devices by Benko et al. [283] provide force feedback and simulate shape and texture of virtual objects using a movable platform and an array of actuated pins to create an experience of interacting with 3D shapes. Choi et al. [284] designed a device combining a gripping mechanism with asymmetric skin stretch to simulate touch, grasp, gravity, and inertia. It creates a sense of holding a weighty object and offers wide range of motion. A unidirectional brake and two voice coil actuators provide perception of gravitational and inertial forces.

The Claw controller [285], another innovation by Choi et al., is a multipurpose controller providing haptic feedback to the index finger, simulating various textures, holding soft or rigid objects, and triggering a gun. Haptic Revolver [286] by Whitmire et al. simulates touch contact, pressure, shear forces, textures, and shapes through a rotating wheel beneath the user's index finger, which rises and falls to imitate contact with virtual surfaces.

Tasbi [119] is a multimodal haptic wristband that provides radial squeeze forces around the wrist along with vibrotactile feedback at six discrete locations. The device uses a DC motor to drive the squeezing mechanism, minimising tangential forces between the band's points of contact with the skin. Force sensing capacitors enable closed-loop control of the squeeze force, while linear resonant actuators provide vibrotactile feedback.

While many handheld devices offer a single perceptual haptic sub-modality, achieving more complete, immersive user experience requires different types of actuators to generate multimodal haptic feedback.




Figure 113. The physical components of the HapticWhirl controller. Blue color is used to highlight the components that form part of the Z axis gimbal, green for the Y axis and yellow for all the components of the braking mechanism.

8.3 Controller Design

This section discusses the development of HapticWhirl, a haptic device utilising a flywheel actuator to generate multimodal 3-DOF force feedback.

8.3.1 Hardware Implementation

The main actuator of HapticWhirl is a brushless motor coupled with a flywheel disk repurposed from a standard hard disk drive. The motor's flat package design aids in aligning it with the gimbal's centre of mass. The rotating disk and motor's combined weight is 96 grams, with the disk having a diameter of 95mm. Motor speed is regulated by a generic 30A electronic speed controller (ESC), enabling rotational speeds up to 12,000 RPM.

The key mechanism for generating torque involves changing the orientation of the spinning flywheel. The disk/motor assembly is mounted on a mechanically actuated spherical gimbal frame, offering two degrees of rotation. When the controller is held vertically in its rest position, the disk remains horizontally oriented. The gimbal mechanisms are driven by Maxwell 110147 DC motors through a pulley system, achieving gear ratios of 10:1 for yaw rotation and 11.5:1 for pitch. This setup amplifies the motor's output torque while maintaining rotational speed up to approximately 10.45 rad/s. A compact PCB featuring a 12-bit resolution digital encoder (AS5047) is positioned opposite the capstan pulley equipped with a magnet. Direction and speed of the gimbal motors are controlled using an H-bridge controller (LN298). The gimbals can rotate over 360 degrees with slip rings installed on both axes. The motors for each axis utilise 3D-printed Miniature Extra Light (MXL) belts made




Figure 114. The controller's inputs and outputs. Left components represent the gimbal control sensors and actuators, on right flywheel components and communication.

from Ninjabflex 95A material and MXL gears printed using PLA, allowing for continuous rotation speeds of 8.35 rad/s (about 800rpm).

To enhance haptic feedback, the design incorporates a scissor drum brake mechanism on one side of the gimbal, allowing rapid cessation of the flywheel. This mechanism, actuated by an SG90 micro servo, can stop the flywheel in 0.8 seconds. Full details and links for downloading all 3D parts and designs are freely available [255].

The final gimbal design has an outer diameter of 125mm, and the controller's total height from the bottom of the handle to the top is 260mm. The total weight, including an HTC Vive tracker mounted at the base for 3D tracking, is 720g.

The flywheel motor and gimbal motors are powered by a 12V supply, while other electronics operate at 3.3V. A Teensy 3.2 microcontroller runs a control loop at 500 Hz and communicates with the PC via serial communication at a baud rate of 115200. Real-time data from the encoders determines the orientation and angular speed of the flywheel, essential for the closed-loop control system.

8.3.2 Working Principle

HapticWhirl generates two primary kinaesthetic forces, originating from its flywheel and dual actuated gimbals. The device utilises the gyroscopic effect by adjusting the orientation of the spinning flywheel through these gimbals. As the flywheel rotates, its orientation change encounters resistance, perceived as torque ($\vec{\tau}$) acting perpendicular to the rotation plane. This torque creates a tactile sensation of push or pull for the user. The gyroscopic effect's operation can be quantified with a set of equations. The torque generated by the flywheel is the rotational force arising from the flywheel's change in angular momentum (\vec{L}). The output torque is a combination of the flywheel's angular momentum:

$$\vec{\tau} = \frac{d\vec{L}}{dt} = \mathbf{I} \frac{d}{dt} \vec{\omega}^{\text{disk}} = \mathbf{I} \dot{\vec{\omega}}^{\text{disk}}$$

In this equation, the torque ($\vec{\tau}$) is a product of the flywheel's moment of inertia (\mathbf{I}) and the rate of change of its angular velocity ($\dot{\vec{\omega}}^{\text{disk}}$). This is the fundamental torque generated by rapidly altering the speed of the spinning flywheel using the brake mechanism. When the flywheel's speed changes abruptly, a reactive torque manifests according to the law of conservation of angular momentum, acting in the direction opposite to the disk's rotation during acceleration or in the same direction during deceleration using the drum brake. Building on this principle, the steered momentum wheel concept, introduced in the equation:

$$\vec{\tau} = \frac{d\vec{L}}{dt} = \vec{\omega}^{\text{gimbal}} \times \mathbf{I} \vec{\omega}^{\text{disk}}$$

represents the dynamics of the gimbal changing the orientation of the flywheel. Here, the torque ($\vec{\tau}$) results from the interaction between the gimbal's angular velocity ($\vec{\omega}^{\text{gimbal}}$) and the flywheel's angular momentum ($\vec{L}^{\text{flywheel}}$). This interaction explains the additional momentum generated when both components move simultaneously. The combination of both equations describes the overall system dynamics within frame of reference A (origin at the centre of the flywheel):

$$\vec{\mathbf{M}}_A^{\text{gyro}} = \mathbf{I} \cdot \dot{\vec{\omega}}_A^{\text{Disk}} + \vec{\omega}_A^{\text{Gimbal}} \times (\mathbf{I} \cdot \vec{\omega}_A^{\text{Disk}})$$

Each component can be expanded, combined and simplify to present the operation of the controller with a single equation for each of the 3D components, as a factor of pitch (ω), yaw (ψ) and disk angle (ρ), full step by step details of the evaluation of the mathematical model are available online with schematics and visual aids [255]:

$$\vec{\mathbf{M}}_A^{\text{gyro}} = \begin{pmatrix} \frac{1}{4} \mathbf{I} (2\dot{\theta}\dot{\rho} - \ddot{\psi} \sin \theta) \\ \frac{1}{4} \mathbf{I} (\ddot{\theta} + \ddot{\psi} \sin \theta (2\dot{\rho} + \dot{\psi} \cos \theta)) \\ \frac{1}{2} \mathbf{I} (\ddot{\rho} + \ddot{\psi} \cos \theta - \dot{\theta} \dot{\psi} \sin \theta) \end{pmatrix}$$

This mathematical model was used to determine the output of the controller across the different actuation modalities during testing.

8.3.3 Controller Haptic Features

During normal operation, the HapticWhirl controller functions in transparent mode, where the flywheel is stabilised by the gimbals' actuation. This stabilisation ensures that the user does not feel gyroscopic forces from the flywheel while moving the controller. However, the controller can manipulate its angular momentum to transition into five distinct haptic modes, each offering a unique interactive experience. These modes are illustrated in Figure 116 and described in this section.

In Torque Mode, as seen in Figure 116 a, the controller actively generates feedback forces that the user perceives by manipulating the flywheel's stabilised state through gimbal actuation. This action shifts the rotation plane of the flywheel, utilising the law of conservation of angular momentum to produce the desired torque. Although the device can operate in a transparent mode without influencing the user's movements, in Torque Mode it intentionally introduces torque. This haptic feedback, which engages the user with resistive or assisting forces, has been explored in previous works such as iTorqu [266] and GyroVR [267].




Figure 115. Operational Modes of the HapticWhirl Controller. (a) Torque Mode: Actuation of gimbals alters the flywheel's orientation, generating torque. (b) Vibration Mode: Rapid pitch-axis oscillation of the flywheel produces tactile vibrations. (c) Impact Mode: Sudden halting of the flywheel simulates an impact. (d) Inertial Mode: Delayed flywheel response creates a sensation of drag. (e) Whirl Mode: Continuous yaw-axis actuation induces a stirring sensation.

In Vibration Mode (Figure 116b), HapticWhirl uses dynamic actuation of the gimbals to create rapid back-and-forth motion along the pitch axis. This movement results in distinct low-frequency vibrations of higher magnitude than those produced by traditional vibrotactile actuators. When the flywheel is static, these vibrations are felt as single-axis vector push-pull forces within the plane of the flywheel. However, when the flywheel rotates, the vibration pattern becomes more complex due to an added tangential component. This complexity arises from the rotational movement, enhancing the basic push-pull oscillation with an additional directional force that increases the vibration's magnitude.

In Impact Mode (Figure 116c), the device utilises its drum brake mechanism to suddenly stop the flywheel's rotation, offering a different kind of force feedback compared to the more gradual approach in Torque Mode. The quick deceleration of the wheel results in a high-magnitude torque that the user feels as a stark, striking force. This immediate stop more closely mimics the sensation of a real-world collision or impact, making it suitable for simulating scenarios like hitting or crashing in virtual environments. While similar concepts have been explored in other areas such as self-assembly robotics [287], the use of a gimbal actuation for kinesthetics feedback is a novel implementation.

Inertial Mode (Figure 116d) is achieved by introducing a delay into the controller's response, altering the usual transparent interaction. This delay allows the gyroscopic forces of the flywheel to be partially perceived by the user, simulating the sensation of inertia. As a result, when the user moves the device, they experience resistance, akin to moving an object through a medium that opposes their motion. It replicates the feeling of pushing against a physical mass that resists immediate change, providing a realistic sensation of weight and drag. This intentional delay transforms the haptic experience to accurately simulate scenarios where inertia is significant, such as navigating through viscous environments or handling heavy objects.

The Whirl Effect (Figure 116e) leverages the controller's unique design to generate a sustained 'stirring' sensation or rotational forces. This is achieved by maintaining the flywheel at a specific rotation angle and actuating the yaw axis to spin continuously, altering the orientation of the angular momentum. This action induces gyroscopic precession, where the rotational axis itself experiences rotation, producing a persistent rotational force. Users experience this as a smooth rotational motion, enhancing simulations that involve circular

dynamics, such as stirring a liquid or performing rotational movements with drag. The ability to simulate sustained rotation without physical stops, by capitalising on gyroscopic precession, showcases a novel approach within the Whirl Mode.

8.4 FARE-KE & HapticWhirl

The HapticWhirl controller presented the most complex characterization challenge of the three prototypes presented, ultimately leading to the development of the FARE-KE framework. The multi-axis force and torque outputs generated by the gyroscopic effect required measurement approaches beyond what we had previously implemented, necessitating innovative solutions particularly during COVID-19 lockdown restrictions.

Unlike the single-axis devices in previous projects, the gyroscopic effects of the flywheel produced counter-intuitive forces/torques that were difficult to predict and verify without measuring equipment. The traditional approach of using commercial 6-axis force-torque sensors was not feasible due to limited access to laboratory facilities during the pandemic. This constraint became the main motivation at the time for developing our own multi-axis measurement system, a solution that would later evolve into a core component of the FARE-KE framework.

Drawing on our experience with beam load cells from previous projects, we designed and constructed a custom 6-axis force-torque plate using an array of four 3-axis load cells, each assembled from three single-axis load cells positioned orthogonally. This represented our first implementation of a multi-axis measurement system using affordable components. The 3D-printed mounting components we used ensured proper alignment between the load cells, which was essential for accurate force and torque measurements.

The calibration process for this custom sensor system required developing methodologies beyond the single-axis approaches used previously. We implemented the calibration procedure that would later be formalized in Appendix 23, achieving an accuracy of approximately 2.3% for forces up to 20N.

Mounting the HapticWhirl controller on this custom force-torque plate allowed us to simultaneously compute theoretical output based on our mathematical model while measuring actual forces and torques. This comparison provided crucial insights into both the controller's operational dynamics and the effectiveness of our measurement approach. The testing

revealed discrepancies between theoretical predictions and measured outputs, highlighting the need for refined calibration methods to address alignment imperfections and mechanical tolerances.

Once access to laboratory equipment was available, we were able to characterise Hapticwrl using an industrial ATI Mini40 force-torque sensor as detailed in the published work [255]. The comparative testing demonstrated that our custom-built system provided measurements that generally aligned with those from the commercial sensor, with minor discrepancies attributed to alignment imperfections and manufacturing tolerances in our custom plate.

These challenges and solutions encountered during the HapticWhirl development formed the foundation for what would become the FARE-KE framework. The necessity of designing and validating our own measurement systems, coupled with the integration of signal processing techniques to address noise issues, established the practical knowledge base that informed the framework's comprehensive approach to kinesthetic device characterization.

8.5 Experimental Validation

This section details the experiments conducted to validate the forces produced by the HapticWhirl device.

8.5.1 Torque Feedback

To determine the maximum output torque of HapticWhirl, the flywheel was programmed to rotate at its maximum speed of 12,000 RPM, while concurrently actuating the pitch gimbal at its peak speed of 800 RPM. To minimise potential external disturbances during the experiment, the yaw axis (Z-axis) was secured and rotation allowed only on the pitch axis (Y-axis). The flywheel's orientation was varied between -60 and 60 degrees, where the home position (0 degrees) corresponds to the disk rotating horizontally while the controller remains upright. The controller was actuated ten times under these conditions, with the collected data subsequently averaged before analysis. The readings were processed through a Butterworth




Figure 116. (a) Measurement of the magnitude of HapticWhirl's output on a single axis. (b) A breakdown of forces for each axis - X (red), Y (green), and Z (blue).

low-pass filter with a cutoff frequency of 50 Hz to eliminate high-frequency noise. Using the following equation:

$$F_{magnitude} = \sqrt{F_x^2 + F_y^2 + F_z^2}$$

The force magnitude was calculated, delivering a peak total output of 4.34 N, as shown in Figure 117a. The output force felt by a user holding the controller is perceived as a torque with two principal components, mainly exerted along the X and Y planes, as shown in Figure 117 b.

8.5.2 Impact Simulation

In the impact simulation mode, the flywheel brake created torque around the axis of rotation through a sudden stop of the flywheel. The drum brake could stop the flywheel in under 300ms which delivered a sudden peak of $\sim 1.8\text{Nm}$, as shown in Figure 118. However, such intense actuations could diminish the lifespan of the HapticWhirl due to increased mechanical




Figure 117. Output torque resulting from suddenly stopping the flywheel while it is rotating at full speed.

wear and stress on the brake system. Consequently, the next iteration of HapticWhirl will feature a redesigned mechanism and a faster motor. These improvements aim not only at enhancing durability but also at enabling varied patterns of on/off braking. Currently, the slow response of the servo motor restricts actuation to a maximum frequency of 10Hz. The proposed faster motor and robust mechanism will allow for higher frequency actuations, improving the dynamic range and responsiveness of the impact simulations.

8.5.3 Vibration Feedback

Oscillating the pitch axis around the home position at high speed generated back-and-forth vibrations. However, these oscillations were constrained by the gimbal motors. In the experiments, oscillation frequencies of up to 170 Hz were achieved without flywheel actuation, and up to 25 Hz when the flywheel was operating at maximum speed.

The first method, operating without flywheel actuation, provided advantages in control and the ability to reach higher vibration frequencies before experiencing motor drift. This vibration had a clear directional characteristic, aligned tangentially to the rotation axis, effectively providing distinct directional cues, as evidenced by the output shown in Figure 119a. Nonetheless, the magnitude of feedback from this method was relatively low compared to the second approach.

The second method involved activating the flywheel during the generation of vibration patterns, which produced higher magnitude output force. Yet, the motor limitations became apparent as maintaining vibration patterns above 25 Hz was unfeasible when the flywheel was at full speed due to motor drift. Additionally, this approach resulted in vibrations that demonstrated multi-component torque dynamics, as detailed in Figure 119b.




Figure 118. (a) Output from back-and-forth rotation of the flywheels gimbal (pitch) with the flywheel at rest. (b) Output the same rotation, but with the flywheel rotating at full speed. In both scenarios, the gimbal was actuated at a frequency of 25Hz.

8.5.4 Whirl Effect (Rotational Forces)

The output torque of the rotating flywheel is orthogonal to the rotation plane. This means that when the controller remains upright and the disk rotates horizontally, their Z-axes are aligned. However, if the disk is rotated to align with any axis other than the upward axis and spun around that axis, the flywheel outputs a sinusoidal torque, as illustrated in Figure 120. To demonstrate this, the rotation of the flywheel around each of the handle axes was measured. Observations revealed that when the flywheel is aligned with the handle's Y-axis and continuously rotates, it generates an output torque on the XZ plane. Similarly, if the flywheel is aligned with the handle's X-axis, the output torque occurs along the YZ plane, and if rotated around the handle's Z-axis, the output is exclusively on the XY plane.

To gather relevant data, measurements were performed with the gimbal rotating at 450 RPM, with the disk aligned to each handle axis. Under these conditions, a consistent torque of approximately 0.2 Nm was recorded across each axis of the plane, while the axis to which the

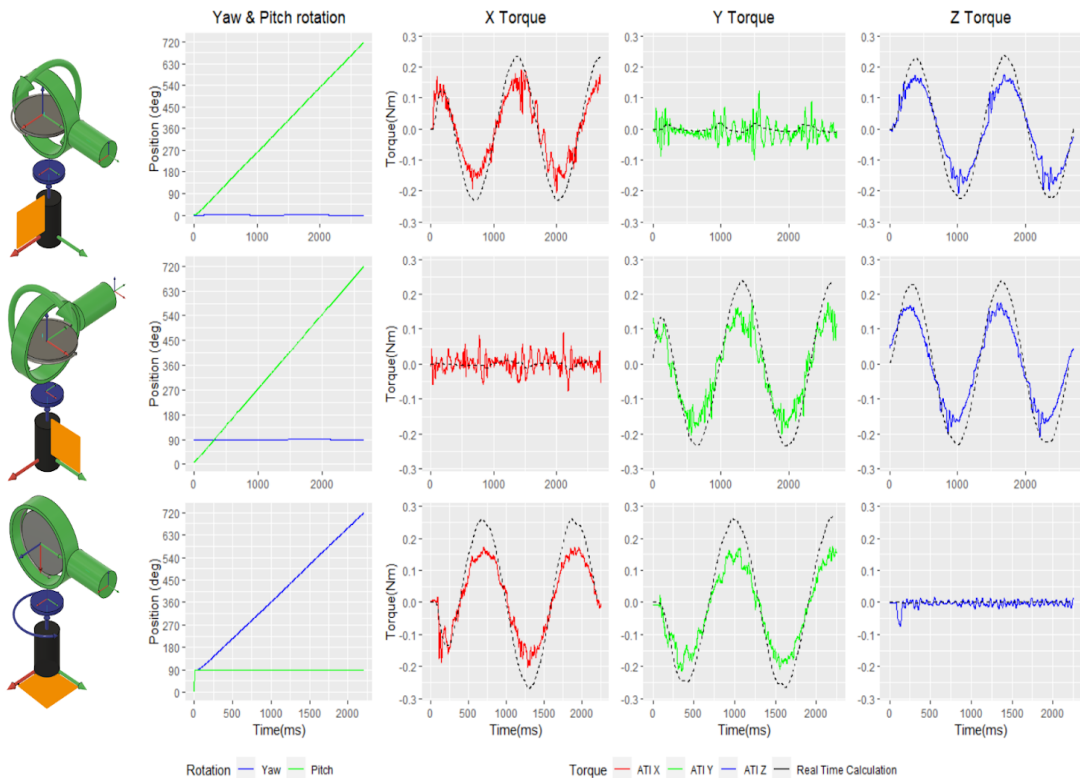


Figure 119. Measured outputs of the HapticWhirl when creating a whirl effect. The top row shows the rotation of the flywheel around its center over the handle. The middle row displays the outputs and position when rotated with alignment to the handle's X-axis. The bottom row illustrates the rotation of the flywheel around the handle's Z-axis. In all instances, the active rotation axis was rotated at approximately 450 RPM.

flywheel was aligned showed no significant torque, only registering noise. Figure 120 presents the measured torque and rotation values for this effect.

This test also compared the mathematical model to the real-time output torque, depicted in the graphs as a black dotted line overlapping each axis in Figure 120. The estimated output torque closely aligns with the readings, though it consistently offsets the peaks by approximately 0.05 Nm. This discrepancy is expected, as the mathematical model does not account for friction or mechanical losses. Finally, all forces generated are within the ranges perceivable by the hand-arm system and are comparable to those used in other systems and experiments, such as references [266], [288].

8.6 Applications

In virtual reality (VR) environments, interactions and collisions are typically experienced by users as torque through a handheld controller or "proxy tool". The HapticWhirl controller's design and features enable a wide array of force feedback simulations to enhance the user's experience in such scenarios.

The controller can generate torque through manipulation of the flywheel and gimbal orientations. This characteristic enables realistic replication of impact effects, such as the feedback experienced when holding a racket and striking a ball. This tangible force feedback amplifies the realism of interaction within a VR space and could be used in various applications in gaming, training, and simulation.

HapticWhirl can create distinctive vibration patterns by oscillating the gimbal. These pronounced vibrations, coupled with the ability to direct forces along different axes, can simulate diverse effects, including shockwaves of an explosion or recoil from firing a weapon with decreasing magnitude in a VR game, as well as texture effects and vibrations from air and fluid flows in simulations or remote-control applications.

The level of torque exerted by the flywheel is directly proportional to the gimbal axle's rotation speed. This enables fine-grained control over haptic feedback, from simulating hard impacts to subtle guidance forces. Such nuanced manipulations can guide users through a VR space, adjust perceived weight of virtual objects, or correct a user's hand position to align with the VR representation.

By utilising the controller’s ability to continuously rotate on the yaw axis, a unique effect akin to gyroscopic precession can be generated. This capability is particularly useful for simulating the experience of stirring a fluid, where similar haptic forces are commonly encountered. The diverse range of feedback modalities available with HapticWhirl extends the possibilities for designing and creating haptic effects, surpassing those offered by existing systems. This expansion enhances the scope for potential applications, empowering designers to innovate new haptic experiences.

8.7 Conclusions

This chapter introduced HapticWhirl, an innovative handheld haptic controller designed for virtual reality applications. The HapticWhirl exploits the gyroscopic effects of a flywheel, utilising it in diverse ways to create a range of haptic sensations. Unique to gyroscopic devices, HapticWhirl leverages gyroscopic precession, enabling torque generation without a fixed grounding point, and uses the flywheel’s rotational energy for dynamic torque output.

The core of this research has been the exploration of novel operational methods that utilise gyroscopic principles. The use of a controlled delay in the device’s transparency to simulate inertia offers a realistic representation of weight and resistance in virtual environments. This aspect will be further explored in user studies in future work. Additionally, the implementation of a braking mechanism enables the simulation of sudden impacts, extending beyond the capabilities of gimbal actuation to maintain the flywheel’s orientation.

HapticWhirl’s ability to support unlimited rotation in both gimbals introduces new possibilities for expanding haptic interactions. This includes capabilities to reorient the disk and maintain operational transparency, aspects that have not been previously explored in prior work.

Previous research has introduced various gyroscope-based haptic devices [265] [267], yet the complex and often unintuitive nature of gyroscopic principles poses significant challenges, particularly for multidisciplinary haptic researchers without extensive backgrounds in mathematical modelling. Recognising this barrier, this work has aimed to demystify the operational mechanics of the HapticWhirl. This chapter provided an online comprehensive, step-by-step guide to the mathematical modelling and operational mechanisms of gyroscopes used in the HapticWhirl [259]. The aim is to bridge the gap between intricate gyroscopic

theory and practical application, enabling a broader range of researchers to engage with and innovate in the field of haptic technology.

Future work will focus on enhancing the structural integrity of the device, specifically improving the axis of rotation for enhanced resilience under high-speed and high-torque conditions. This refinement aims to balance the robustness of the device with a more compact and user-friendly form factor. Additionally, the utilisation of a faster motor and a smaller disk is anticipated to reduce the overall size of the device, thereby improving control and responsiveness.

Furthermore, future work will concentrate on user evaluations of the device and exploring applications, as well as validating the five operational modes in VR environments. This exploration will include investigating potential retargeting applications, enabling subtle repositioning of the user's hand in VR without perceptible detection by gradually increasing the torque as the user extends their arm.

9 Conclusion

This thesis presents a framework for characterising kinesthetic haptic devices. The Framework for Affordable, Reliable Kinesthetic Evaluation (FARE-KE) addresses the need for standardised, accessible methods that align with contemporary haptic device development practices.

9.1 FARE-KE Framework Validation and Summary

The FARE-KE framework incorporates validated hardware components, software tools and measurement methodologies. Single-axis beam load cells (TAL220, Degraw, and CZL635) demonstrated linear behaviour with R_{ffl} values exceeding 0.999 across all tested calibration methods (Table 17). The polynomial regression calibration approach yielded the lowest RMSE values (6.98-9.00 gf), confirming the viability of low-cost beam load cells for characterising simple haptic devices.

The custom-built 3-axis load cell design (Section 5.5.2) achieves approximately 2.3% accuracy for forces up to 20 N when calibrated using polynomial regression. Button and S-type load cells (Section 5.3) proved useful for inline force measurements in space-constrained applications. The 6-axis force-torque plate (Section 5.5.3) showed significant accuracy improvements following calibration, with mean absolute percentage error reducing from 2.81% to 1.85% (Table 26).

The custom data acquisition board (Chapter 4) supports simultaneous sampling of up to 12 load cells at rates exceeding 10 kSPS. The evaluation framework's analysis methods, including ramp and step response analyses (Section 3.9), provide insights into device performance metrics such as maximum force/torque, hysteresis, sensitivity and response times.

The three case studies demonstrated the framework's adaptability to diverse kinesthetic haptic devices and its effectiveness in standardising evaluation procedures.

9.2 Research Contributions

This thesis addresses six research questions outlined in Chapter 1:

RQ1: How can a standardised framework be developed to make haptic device characterisation more widely adopted?

The FARE-KE framework provides a standardised approach with open-source hardware and software components and step-by-step instructions suitable for researchers with varying expertise levels. This promotes collaboration and knowledge sharing within the haptic research community.

RQ2: What components and methodologies establish a foundation for standardised measurements?

The thesis identifies essential elements for haptic device characterisation, including hardware selection guidelines (Section 3.4), data acquisition methodologies (Section 3.6), performance analysis techniques (Section 3.9), and sensor validation protocols (Chapter 5). The framework establishes consistent metrics and procedures for evaluating haptic devices across different modalities.

The thesis also introduces innovative multi-axis load cell designs and provides detailed calibration and performance testing protocols. The literature review synthesises common characterisation approaches across tactile and kinesthetic haptic technologies.

RQ3: How can the framework's versatility be demonstrated across different kinesthetic feedback modalities?

The case studies (Chapters 6-8) demonstrate the framework's versatility by applying it to different kinesthetic feedback modalities: linear kinesthetic feedback (VR Recoil), body-based kinesthetic feedback (Proxy Touch), and continuous rotational flywheel-based haptic effects (HapticWhirl).

RQ4: What is the impact of linear kinesthetic feedback on user performance in virtual reality shooting scenarios?

The VR Recoil study (Chapter 6) showed that linear kinesthetic feedback enhanced realism and immersion without compromising performance. Lower feedback levels proved sufficient for enhancing user experience, suggesting manufacturers could optimise user satisfaction while minimising production costs.

RQ5: Can kinesthetic feedback be recreated using the user's own body, without external proxies?

The Proxy Touch study (Chapter 7) demonstrated somatosensory attenuation effects in pinch interactions, confirming the potential for creating haptic illusions using only the user’s body without external mechanical devices.

RQ6: How can continuous rotational flywheels simulate fluid motion haptic effects?

The HapticWhirl study (Chapter 8) validated a controller capable of delivering five distinct haptic modalities: torque feedback, impact simulation, vibration feedback, inertial effects, and whirl effects. The FARE-KE framework provided insights into the complex force and torque outputs generated by the device.

9.3 Limitations and Considerations

The FARE-KE framework and proposed load cell configurations have several limitations:

First, the lack of extensive comparisons between the custom-built 6-axis force-torque plate and industrial-grade sensors limits validation of the results’ comparability. The 3D-printed components and affordable load cells result in reduced force measurement ranges compared to industrial-grade sensors, potentially limiting applicability for high-magnitude forces or wide dynamic ranges.

The 6-axis force-torque plate’s larger footprint makes it suitable for bench-top testing rather than direct integration, potentially limiting characterisation in actual operating conditions and failing to capture effects of factors such as weight distribution on user experience.

The framework focuses on physical evaluation without extensively addressing psychophysical aspects of haptic interaction. As haptic technology evolves, the framework may require adaptation to accommodate new actuation methods, materials, and design paradigms.

Finally, strain gauge technology is susceptible to temperature fluctuations, electromagnetic interference, and mechanical stress, which may impact long-term measurement stability despite calibration and data processing techniques.

9.4 Future Work

Future development areas include:

Improving the data acquisition system by incorporating an external clock to reduce noise, improve signal quality, and enable higher sampling rates (Section 4.3). This would enhance the capture of fast transient events and improve temporal resolution.

Validating and refining the framework through extensive user studies and collaborations with practitioners would enhance its applicability. COVID-19 restrictions limited device testing during development, so workshops at conferences and practitioner feedback would provide valuable insights.

Developing a dedicated website repository with hardware designs, software tools, documentation, and case studies (Section 3.1) would facilitate adoption, designed for accessibility to ensure researchers with varying expertise levels can utilise the resources effectively.

Future work should explore integrating psychophysical evaluation methods into the framework and investigating applications to emerging haptic technologies as new actuation methods, materials, and design paradigms emerge.

9.5 Concluding Remarks

The FARE-KE framework represents an advancement in haptic device characterisation, offering a standardised, accessible approach that promotes consistency, reproducibility, and comparability in haptic research. By addressing the lack of uniformity in evaluation methodologies, this work contributes to accelerating the development of effective haptic technologies. The load cell analysis and validation provide a foundation for characterising kinesthetic haptic devices within the framework. While the custom-built load cell configurations have limitations compared to industrial sensors, they offer viable alternatives for researchers with limited resources.

Through engagement with the haptic research community, promoting framework adoption, and continuous refinement of methods and tools, this work aims to advance haptic device characterisation and shape future development of haptic technology.

10 Bibliography

- [1] E. Samur, *Performance Metrics for Haptic Interfaces*. Springer Science & Business Media, 2012.
- [2] V. Hayward and O. R. Astley, “Performance Measures for Haptic Interfaces,” *Robotics Research*, pp. 195–206, 1996, doi: 10.1007/978-1-4471-1021-7_22.
- [3] T. A. Kern, C. Hatzfeld, and A. Abbasimoshaei, Eds., *Engineering Haptic Devices*. in Springer Series on Touch and Haptic Systems. Cham: Springer International Publishing, 2023. doi: 10.1007/978-3-031-04536-3.
- [4] H. Seifi *et al.*, “Haptipedia: Accelerating Haptic Device Discovery to Support Interaction & Engineering Design,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, in CHI '19. New York, NY, USA: ACM, 2019, p. 558:1-558:12. doi: 10.1145/3290605.3300788.
- [5] J. R. Blum *et al.*, “Getting Your Hands Dirty Outside the Lab: A Practical Primer for Conducting Wearable Vibrotactile Haptics Research,” *IEEE Trans. Haptics*, vol. 12, no. 3, pp. 1–1, 2019, doi: 10.1109/TOH.2019.2930608.
- [6] S. Papetti, M. Fröhlich, F. Fontana, S. Schiesser, and F. Avanzini, “Implementation and Characterisation of Vibrotactile Interfaces,” *Musical Haptics*, pp. 257–282, 2018, doi: 10.1007/978-3-319-58316-7_13.
- [7] A. L. Ager *et al.*, “Proprioception: How is it affected by shoulder pain? A systematic review,” *Journal of Hand Therapy*, vol. 33, no. 4, pp. 507–516, Oct. 2020, doi: 10.1016/j.jht.2019.06.002.
- [8] S. A. Khonsary, “Guyton and Hall: Textbook of Medical Physiology,” *Surg Neurol Int*, vol. 8, no. 1, p. 275, Nov. 2017, doi: 10.4103/sni.sni_327_17.
- [9] J. Dargahi and S. Najarian, “Human tactile perception as a standard for artificial tactile sensing—a review,” *The international journal of medical robotics and computer assisted surgery*, vol. 1, no. 1, pp. 23–35, 2004, doi: 10.1002/rcs.3.
- [10] K. Myles and M. S. Binseel, “The tactile modality: a review of tactile sensitivity and human tactile interfaces,” 2007.
- [11] R. S. Johansson, “Tactile sensibility in the human hand: receptive field characteristics of mechanoreceptive units in the glabrous skin area.,” *The Journal of physiology*, vol. 281, no. 1, pp. 101–125, 1978, doi: 10.1113/jphysiol.1978.sp012411.
- [12] R. S. Johansson and A. B. Vallbo, “Tactile sensibility in the human hand: relative and absolute densities of four types of mechanoreceptive units in glabrous skin.,” *The Journal of physiology*, vol. 286, no. 1, pp. 283–300, 1979, doi: 10.1113/jphysiol.1979.sp012619.
- [13] R. D. Hamer, R. T. Verrillo, and J. J. Zwislocki, “Vibrotactile masking of Pacinian and non-Pacinian channels,” *The Journal of the Acoustical Society of America*, vol. 73, no. 4, pp. 1293–1303, 1983, doi: 10.1121/1.389278.
- [14] M. A. Lyle and T. R. Nichols, “Evaluating intermuscular Golgi tendon organ feedback with twitch contractions,” *The Journal of Physiology*, vol. 597, no. 17, pp. 4627–4642, Sep. 2019, doi: 10.1113/JP277363.
- [15] S. J. Bolanowski, G. A. Gescheider, and R. T. Verrillo, “Hairy skin: psychophysical channels and their physiological substrates,” *Somatosens Mot Res*, vol. 11, no. 3, pp. 279–290, 1994, doi: 10.3109/08990229409051395.

- [16] J. Bell, S. Bolanowski, and M. H. Holmes, "The structure and function of pacinian corpuscles: A review," *Progress in Neurobiology*, vol. 42, no. 1, pp. 79–128, Jan. 1994, doi: 10.1016/0301-0082(94)90022-1.
- [17] L. A. Jones and N. B. Sarter, "Tactile displays: guidance for their design and application," *Hum Factors*, vol. 50, no. 1, pp. 90–111, Feb. 2008, doi: 10.1518/001872008x250638.
- [18] S. Choi and K. J. Kuchenbecker, "Vibrotactile Display: Perception, Technology, and Applications," *Proc. IEEE*, vol. 101, no. 9, pp. 2093–2104, Sep. 2013, doi: 10.1109/jproc.2012.2221071.
- [19] H. Culbertson, S. B. Schorr, and A. M. Okamura, "Haptics: The Present and Future of Artificial Touch Sensation," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 1, no. 1, pp. 385–409, May 2018, doi: 10.1146/annurev-control-060117-105043.
- [20] J. Yin, R. Hinchet, H. Shea, and C. Majidi, "Wearable Soft Technologies for Haptic Sensing and Feedback," *Adv Funct Materials*, vol. n/a, no. n/a, p. 2007428, Sep. 2021, doi: 10.1002/adfm.202007428.
- [21] V. Hayward and K. E. Maclean, "Do it yourself haptics: part I," *IEEE Robot. Automat. Mag.*, vol. 14, no. 4, pp. 88–104, Dec. 2007, doi: 10.1109/m-ra.2007.907921.
- [22] "Adafruit DRV2605L Haptic Controller Breakout," Adafruit Learning System. Accessed: Nov. 17, 2023. [Online]. Available: <https://learn.adafruit.com/adafruit-drv2605-haptic-controller-breakout/overview>
- [23] R. H. Gault, "Recent developments in vibro-tactile research," *Journal of the Franklin Institute*, vol. 221, no. 6, pp. 703–719, Jun. 1936, doi: 10.1016/S0016-0032(36)90521-5.
- [24] H. M. Culbertson, "Data-Driven Haptic Modelling and Rendering of Realistic Virtual Textured Surfaces."
- [25] M. Ito, D. Wakuda, S. Inoue, Y. Makino, and H. Shinoda, "High Spatial Resolution Midair Tactile Display Using 70kHz Ultrasound," in *Proceedings, Part I, of the 10th International Conference on Haptics: Perception, Devices, Control, and Applications - Volume 9774*, in EuroHaptics 2016. Berlin, Heidelberg: Springer-Verlag, Jul. 2016, pp. 57–67. doi: 10.1007/978-3-319-42321-0_6.
- [26] J. Lee, "Wind Tactor: An Airflow-based Wearable Tactile Display," in *Adjunct Publication of the 30th Annual ACM Symposium on User Interface Software and Technology*, Québec City QC Canada: ACM, Oct. 2017, pp. 91–94. doi: 10.1145/3131785.3131838.
- [27] Y. Sato and R. Ueoka, "Investigating Haptic Perception of and Physiological Responses to Air Vortex Rings on a User's Cheek," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, Denver Colorado USA: ACM, May 2017, pp. 3083–3094. doi: 10.1145/3025453.3025501.
- [28] C. Jiang *et al.*, "Douleur: Creating Pain Sensation with Chemical Stimulant to Enhance User Experience in Virtual Reality," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 5, no. 2, p. 66:1-66:26, Jun. 2021, doi: 10.1145/3463527.
- [29] J. Lu, Z. Liu, J. Brooks, and P. Lopes, "Chemical Haptics: Rendering Haptic Sensations via Topical Stimulants," in *The 34th Annual ACM Symposium on User Interface Software and Technology*, Virtual Event USA: ACM, Oct. 2021, pp. 239–257. doi: 10.1145/3472749.3474747.

- [30] J. Brooks, S. Nagels, and P. Lopes, “Trigeminal-based Temperature Illusions,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, in CHI ’20. Honolulu, HI, USA: Association for Computing Machinery, Apr. 2020, pp. 1–12. doi: 10.1145/3313831.3376806.
- [31] A. Adilkhanov, M. Rubagotti, and Z. Kappassov, “Haptic Devices: Wearability-Based Taxonomy and Literature Review,” *IEEE Access*, vol. 10, pp. 91923–91947, 2022, doi: 10.1109/ACCESS.2022.3202986.
- [32] I. Choi, E. Ofek, H. Benko, M. Sinclair, and C. Holz, “CLAW: A Multifunctional Handheld Haptic Controller for Grasping, Touching, and Triggering in Virtual Reality,” in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, Montreal QC Canada: ACM Press, 2018, pp. 1–13. doi: 10.1145/3173574.3174228.
- [33] M. Sinclair, E. Ofek, M. Gonzalez-Franco, and C. Holz, “CapstanCrunch: A Haptic VR Controller with User-supplied Force Feedback,” in *Proceedings of the 32Nd Annual ACM Symposium on User Interface Software and Technology*, in UIST ’19. New York, NY, USA: ACM, 2019, pp. 815–829. doi: 10.1145/3332165.3347891.
- [34] H. Benko, C. Holz, M. Sinclair, and E. Ofek, “NormalTouch and TextureTouch: High-fidelity 3D Haptic Shape Rendering on Handheld Virtual Reality Controllers,” in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, in UIST ’16. New York, NY, USA: ACM, 2016, pp. 717–728. doi: 10.1145/2984511.2984526.
- [35] G. Ko and S. Nam, “Bstick: Handheld Virtual Reality Haptic Controller for Hand Rehabilitation,” *Systems*, vol. 10, no. 3, p. 54, Apr. 2022, doi: 10.3390/systems10030054.
- [36] Y. Sun, S. Yoshida, T. Narumi, and M. Hirose, “PaCaPa: A Handheld VR Device for Rendering Size, Shape, and Stiffness of Virtual Objects in Tool-based Interactions,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, in CHI ’19. New York, NY, USA: ACM, 2019, p. 452:1-452:12. doi: 10.1145/3290605.3300682.
- [37] D. K. Y. Chen, J.-B. Chossat, and P. B. Shull, “HaptiVec: Presenting Haptic Feedback Vectors in Handheld Controllers using Embedded Tactile Pin Arrays,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, in CHI ’19. New York, NY, USA: Association for Computing Machinery, May 2019, pp. 1–11. doi: 10.1145/3290605.3300401.
- [38] S.-Y. Teng *et al.*, “PuPoP: Pop-up Prop on Palm for Virtual Reality,” in *The 31st Annual ACM Symposium on User Interface Software and Technology*, in UIST ’18. New York, NY, USA: ACM, 2018, pp. 5–17. doi: 10.1145/3242587.3242628.
- [39] S.-P. Hu and J.-H. Hou, “Pneu-Multi-Tools: Auto-Folding and Multi-Shapes Interface by Pneumatics in Virtual Reality,” in *The Adjunct Publication of the 32nd Annual ACM Symposium on User Interface Software and Technology*, New Orleans LA USA: ACM, Oct. 2019, pp. 36–38. doi: 10.1145/3332167.3357107.
- [40] A. Zenner and A. Kruger, “Shifty: A Weight-Shifting Dynamic Passive Haptic Proxy to Enhance Object Perception in Virtual Reality,” *IEEE Trans. Visual. Comput. Graphics*, vol. 23, no. 4, pp. 1285–1294, Apr. 2017, doi: 10.1109/tvcg.2017.2656978.
- [41] A. Zenner and A. Krüger, “Drag:on: A Virtual Reality Controller Providing Haptic Feedback Based on Drag and Weight Shift,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, in CHI ’19. New York, NY, USA: ACM, 2019, p. 211:1-211:12. doi: 10.1145/3290605.3300441.

- [42] J. Shigeyama *et al.*, “Transcalibur: Weight Moving VR Controller for Dynamic Rendering of 2D Shape Using Haptic Shape Illusion,” in *ACM SIGGRAPH 2018 Emerging Technologies*, in SIGGRAPH ’18. New York, NY, USA: ACM, 2018, p. 19:1-19:2. doi: 10.1145/3214907.3214923.
- [43] H.-R. Tsai, C.-W. Hung, T.-C. Wu, and B.-Y. Chen, “ElastOscillation: 3D Multilevel Force Feedback for Damped Oscillation on VR Controllers,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, Honolulu HI USA: ACM, 2020, p. 12. doi: 10.1145/3313831.3376408.
- [44] N. Ryu, W. Lee, M. J. Kim, and A. Bianchi, “ElaStick: A Handheld Variable Stiffness Display for Rendering Dynamic Haptic Response of Flexible Object,” in *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, in UIST ’20. New York, NY, USA: Association for Computing Machinery, Oct. 2020, pp. 1035–1045. doi: 10.1145/3379337.3415862.
- [45] S. Heo, C. Chung, G. Lee, and D. Wigdor, “Thor’s Hammer: An Ungrounded Force Feedback Device Utilising Propeller-Induced Propulsive Force,” in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, in CHI ’18. New York, NY, USA: ACM, 2018, p. 525:1-525:11. doi: 10.1145/3173574.3174099.
- [46] K.-H. Liu, T. Sasaki, H. Kajihara, A. Hiyama, M. Inami, and C.-H. Chen, “Designing WindCage- Unpacking the Thinking and Prototyping a Propeller-Based Haptic Unit,” in *2022 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, Christchurch, New Zealand: IEEE, Mar. 2022, pp. 131–135. doi: 10.1109/VRW55335.2022.00038.
- [47] T. Sasaki, R. S. Hartanto, K.-H. Liu, K. Tsuchiya, A. Hiyama, and M. Inami, “Leviopole: mid-air haptic interactions using multirotor,” in *ACM SIGGRAPH 2018 Emerging Technologies*, Vancouver British Columbia Canada: ACM, Aug. 2018, pp. 1–2. doi: 10.1145/3214907.3214913.
- [48] T. Sasaki, K.-H. Liu, T. Hasegawa, A. Hiyama, and M. Inami, “Virtual Super-Leaping: Immersive Extreme Jumping in VR,” in *Proceedings of the 10th Augmented Human International Conference 2019*, in AH2019. New York, NY, USA: ACM, 2019, p. 18:1-18:8. doi: 10.1145/3311823.3311861.
- [49] S. Je, H. Lee, M. J. Kim, and A. Bianchi, “Wind-blaster: A Wearable Propeller-based Prototype That Provides Ungrounded Force-feedback,” in *ACM SIGGRAPH 2018 Emerging Technologies*, in SIGGRAPH ’18. New York, NY, USA: ACM, 2018, p. 23:1-23:2. doi: 10.1145/3214907.3214915.
- [50] S. Je *et al.*, “Aero-plane: A Handheld Force-Feedback Device That Renders Weight Motion Illusion on a Virtual 2D Plane,” in *Proceedings of the 32Nd Annual ACM Symposium on User Interface Software and Technology*, in UIST ’19. New York, NY, USA: ACM, 2019, pp. 763–775. doi: 10.1145/3332165.3347926.
- [51] Y.-W. Wang *et al.*, “JetController: High-speed Ungrounded 3-DoF Force Feedback Controllers using Air Propulsion Jets,” in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, Yokohama Japan: ACM, May 2021, pp. 1–12. doi: 10.1145/3411764.3445549.
- [52] Y. Tanaka, S. Masataka, K. Yuka, Y. Fukui, J. Yamashita, and N. Nakamura, “Mobile Torque Display and Haptic Characteristics of Human Palm,” p. 6.
- [53] H. Ando, K. Obana, M. Sugimoto, and T. Maeda, “A wearable force display based on brake change in angular momentum,” p. 6.
- [54] N. Nakamura and Y. Fukui, “An Innovative Non-grounding Haptic Interface ‘GyroCubeSensuous’ displaying Illusion Sensation of Push, Pull and Lift,” in *ACM*

SIGGRAPH 2005 Posters on - SIGGRAPH '05, Los Angeles, California: ACM Press, 2005, p. 1. doi: 10.1145/1186954.1187059.

- [55] J. M. Rigby, D. P. Brumby, S. J. J. Gould, and A. L. Cox, "Development of a Questionnaire to Measure Immersion in Video Media: The Film IEQ," in *Proceedings of the 2019 ACM International Conference on Interactive Experiences for TV and Online Video*, in TVX '19. New York, NY, USA: ACM, 2019, pp. 35–46. doi: 10.1145/3317697.3323361.
- [56] T. Nakamura and H. Koike, "Golf Club-Type Device with Force Feedback for Modifying Club Posture," in *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, in CHI EA '20. Honolulu, HI, USA: Association for Computing Machinery, Apr. 2020, pp. 1–7. doi: 10.1145/3334480.3383024.
- [57] K. N. Winfree, J. M. Romano, J. Gewirtz, and K. J. Kuchenbecker, "Control of a high fidelity ungrounded torque feedback device: The iTorqU 2.1," in *2010 IEEE International Conference on Robotics and Automation*, Anchorage, AK: IEEE, May 2010, pp. 1347–1352. doi: 10.1109/ROBOT.2010.5509485.
- [58] K. N. Winfree, J. Gewirtz, T. Mather, J. Fiene, and K. J. Kuchenbecker, "A high fidelity ungrounded torque feedback device: The iTorqU 2.0," in *World Haptics 2009 - Third Joint EuroHaptics conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, Salt Lake City, UT, USA: IEEE, 2009, pp. 261–266. doi: 10.1109/WHC.2009.4810866.
- [59] J. Chiu and A. Goswami, "Design of a Wearable Scissored-Pair Control Moment Gyroscope (SP-CMG) for Human Balance Assist," in *Volume 5A: 38th Mechanisms and Robotics Conference*, Buffalo, New York, USA: American Society of Mechanical Engineers Digital Collection, Jan. 2015, p. V05AT08A023. doi: 10.1115/DETC2014-35539.
- [60] M. Antolini, M. Bordegoni, and U. Cugini, "A haptic direction indicator using the gyro effect," in *2011 IEEE World Haptics Conference*, Istanbul: IEEE, Jun. 2011, pp. 251–256. doi: 10.1109/WHC.2011.5945494.
- [61] J. M. Walker, M. Raitor, A. Mallery, H. Culbertson, P. Stolka, and A. M. Okamura, "A dual-flywheel ungrounded haptic feedback system provides single-axis moment pulses for clear direction signals," in *2016 IEEE Haptics Symposium (HAPTICS)*, Philadelphia, PA, USA: IEEE, Apr. 2016, pp. 7–13. doi: 10.1109/HAPTICS.2016.7463148.
- [62] J. M. Walker, H. Culbertson, M. Raitor, and A. M. Okamura, "Haptic orientation guidance using two parallel double-gimbal control moment gyroscopes," *IEEE Trans. Haptics*, vol. 11, no. 2, pp. 267–278, Apr. 2018, doi: 10.1109/TOH.2017.2713380.
- [63] X. Gu, Y. Zhang, W. Sun, Y. Bian, D. Zhou, and P. O. Kristensson, "Dexmo: An inexpensive and lightweight mechanical exoskeleton for motion capture and force feedback in VR," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, San Jose California USA: ACM, 2016, pp. 1991–1995. doi: 10.1145/2858036.2858487.
- [64] "Dexta Robotics." Accessed: Nov. 26, 2018. [Online]. Available: <https://www.dextarobotics.com/>
- [65] D. Ryu *et al.*, "Micro hydraulic system using slim artificial muscles for a wearable haptic glove," in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Nice: IEEE, Sep. 2008, pp. 3028–3033. doi: 10.1109/IROS.2008.4651159.

- [66] M. Hosseini, A. Sengul, Y. Pane, J. De Schutter, and H. Bruyninck, “ExoTen-Glove: A Force-Feedback Haptic Glove Based on Twisted String Actuation System,” in *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, Nanjing: IEEE, Aug. 2018, pp. 320–327. doi: 10.1109/roman.2018.8525637.
- [67] I. Choi and S. Follmer, “Wolverine: A Wearable Haptic Interface for Grasping in VR,” in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, in UIST ’16 Adjunct. New York, NY, USA: ACM, 2016, pp. 117–119. doi: 10.1145/2984751.2985725.
- [68] R. Hinchet, V. Vechev, H. Shea, and O. Hilliges, “DextrES: Wearable Haptic Feedback for Grasping in VR via a Thin Form-Factor Electrostatic Brake,” in *The 31st Annual ACM Symposium on User Interface Software and Technology*, in UIST ’18. New York, NY, USA: ACM, 2018, pp. 901–912. doi: 10.1145/3242587.3242657.
- [69] H.-R. Tsai and J. Rekimoto, “ElasticVR: Providing Multi-level Active and Passive Force Feedback in Virtual Reality Using Elasticity,” in *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, in CHI EA ’18. New York, NY, USA: ACM, 2018, p. D300:1-D300:4. doi: 10.1145/3170427.3186540.
- [70] M. Achibet, A. Girard, A. Talvas, M. Marchal, and A. Lécuyer, “Elastic-Arm: Human-scale passive haptic feedback for augmenting interaction and perception in virtual environments,” in *2015 IEEE Virtual Reality (VR)*, Arles, Camargue, Provence, France: IEEE, Mar. 2015, pp. 63–68. doi: 10.1109/VR.2015.7223325.
- [71] K. Klamka and R. Dachsel, “Elasticcon: Elastic Controllers for Casual Interaction,” in *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services - MobileHCI ’15*, Copenhagen, Denmark: ACM Press, 2015, pp. 410–419. doi: 10.1145/2785830.2785849.
- [72] A. T. Bukkapatnam, P. Depalle, and M. M. Wanderley, “Defining a vibrotactile toolkit for digital musical instruments: characterizing voice coil actuators, effects of loading, and equalization of the frequency response,” *J Multimodal User Interfaces*, vol. 14, no. 3, pp. 285–301, Sep. 2020, doi: 10.1007/s12193-020-00340-0.
- [73] P. E. Fortin, J. R. Blum, and J. R. Cooperstock, “Towards Consistent Haptic Coupling with HaptiStrap: Doing Better Than ‘Tight Yet Comfortable,’” in *The Adjunct Publication of the 32Nd Annual ACM Symposium on User Interface Software and Technology*, in UIST ’19. New York, NY, USA: ACM, 2019, pp. 69–71. doi: 10.1145/3332167.3357118.
- [74] C. Gaudeni, L. Meli, L. Jones, and D. Prattichizzo, “Presenting surface features using a haptic ring: A psychophysical study on relocating vibrotactile feedback,” *IEEE Trans. Haptics*, vol. 12, no. 4, pp. 1–1, 2019, doi: 10.1109/TOH.2019.2938945.
- [75] R. Gourishetti and K. J. Kuchenbecker, “Evaluation of Vibrotactile Output from a Rotating Motor Actuator,” *IEEE Trans. Haptics*, vol. 15, no. 1, pp. 1–1, 2021, doi: 10.1109/TOH.2021.3138867.
- [76] J. M. Muñoz, “A Vibrotactile Prototyping Toolkit for Virtual Reality and Videogames.”
- [77] K. J. Kuchenbecker, W. McMahan, N. Landin, and J. M. Romano, “Dimensional Reduction of High-Frequency Accelerations for Haptic Rendering”.
- [78] P. Preechayasomboon, A. Israr, and M. Samad, “Chasm: A Screw Based Expressive Compact Haptic Actuator,” p. 13. doi: 10.1145/3313831.3376512.

- [79] W. T. A. Frier, "Rendering Spatiotemporal Mid-Air Tactile Patterns."
- [80] B. Dandu, Y. Shao, A. Stanley, and Y. Visell, "Spatiotemporal Haptic Effects from a Single Actuator via Spectral Control of Cutaneous Wave Propagation," in *2019 IEEE World Haptics Conference (WHC)*, Tokyo, Japan: IEEE, Jul. 2019, pp. 425–430. doi: 10.1109/WHC.2019.8816149.
- [81] S.-Y. Teng, P. Li, R. Nith, J. Fonseca, and P. Lopes, "Touch&Fold: A Foldable Haptic Actuator for Rendering Touch in Mixed Reality," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, New York, NY, USA: Association for Computing Machinery, May 2021, pp. 1–14. Accessed: Jul. 06, 2021. [Online]. Available: <https://doi.org/10.1145/3411764.3445099>
- [82] D. Giovanelli and E. Farella, "Force Sensing Resistor and Evaluation of Technology for Wearable Body Pressure Sensing," *Journal of Sensors*, vol. 2016, p. e9391850, Feb. 2016, doi: 10.1155/2016/9391850.
- [83] J. A. Flórez and A. Velásquez, "Calibration of force sensing resistors (fsr) for static and dynamic applications," in *2010 IEEE ANDESCON*, Bogota: IEEE, Sep. 2010, pp. 1–6. doi: 10.1109/ANDESCON.2010.5633120.
- [84] R. S. Hall, G. T. Desmoulin, and T. E. Milner, "A technique for conditioning and calibrating force-sensing resistors for repeatable and reliable measurement of compressive force," *Journal of Biomechanics*, vol. 41, no. 16, pp. 3492–3495, Dec. 2008, doi: 10.1016/j.jbiomech.2008.09.031.
- [85] Z. Del Prete, L. Monteleone, and R. Steindler, "A novel pressure array sensor based on contact resistance variation: Metrological properties," *Review of Scientific Instruments*, vol. 72, no. 2, pp. 1548–1553, 2001, doi: 10.1063/1.1340561.
- [86] M.-A. Lacasse, V. Duchaine, and C. Gosselin, "Characterisation of the electrical resistance of carbon-black-filled silicone: Application to a flexible and stretchable robot skin," presented at the 2010 IEEE International Conference on Robotics and Automation, IEEE, 2010, pp. 4842–4848.
- [87] K. Oppermann, B. R. Arminger, and B. G. Zagar, "Smart hysteresis compensation of a magneto-elastic force sensor based on Terfenol-D," presented at the 2010 IEEE Instrumentation & Measurement Technology Conference Proceedings, IEEE, 2010, pp. 662–667.
- [88] E. R. Komi, J. R. Roberts, and S. Rothberg, "Evaluation of thin, flexible sensors for time-resolved grip force measurement," *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, vol. 221, no. 12, pp. 1687–1699, 2007, doi: 10.1243/09544062JMES700.
- [89] R. Bonnaire, M. Verhaeghe, J. Molimard, P. Calmels, and R. Convert, "Characterisation of a pressure measuring system for the evaluation of medical devices," *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, vol. 228, no. 12, pp. 1264–1274, 2014, doi: 10.1177/0954411914562871.
- [90] "FlexiForce A201 Datasheet," Tekscan. Accessed: Oct. 04, 2022. [Online]. Available: <https://tekscan.com/resources/datasheets-guides/flexiforce-a201-datasheet>
- [91] I. E. Inc, "FSR 400 Series." Accessed: Oct. 04, 2022. [Online]. Available: <https://www.interlinkelectronics.com/fsr-400-series>
- [92] L. Paredes-Madrid, A. Matute, and A. Peña, "Framework for a Calibration-Less Operation of Force Sensing Resistors at Different Temperatures," *IEEE Sensors J.*, vol. 17, no. 13, pp. 4133–4142, Jul. 2017, doi: 10.1109/JSEN.2017.2706697.

- [93] L. Paredes-Madrid, C. A. Palacio, A. Matute, and C. A. Parra Vargas, "Underlying Physics of Conductive Polymer Composites and Force Sensing Resistors (FSRs) under Static Loading Conditions," *Sensors*, vol. 17, no. 9, Art. no. 9, Sep. 2017, doi: 10.3390/s17092108.
- [94] A. Hollinger and M. M. Wanderley, "Evaluation of Commercial Force-Sensing Resistors," p. 4.
- [95] L. Paredes-Madrid, J. Fonseca, A. Matute, E. Gutiérrez Velásquez, and C. Palacio, "Self-Compensated Driving Circuit for Reducing Drift and Hysteresis in Force Sensing Resistors," *Electronics*, vol. 7, no. 8, p. 146, Aug. 2018, doi: 10.3390/electronics7080146.
- [96] A. Talhan, H. Kim, and S. Jeon, "Tactile Ring: Multi-Mode Finger-Worn Soft Actuator for Rich Haptic Feedback," in *IEEE Access*, 2020, pp. 957–966. doi: 10.1109/ACCESS.2019.2961430.
- [97] S. H. Yoon *et al.*, "HapSense: A Soft Haptic I/O Device with Uninterrupted Dual Functionalities of Force Sensing and Vibrotactile Actuation," in *Proceedings of the 32Nd Annual ACM Symposium on User Interface Software and Technology*, in UIST '19. New York, NY, USA: ACM, 2019, pp. 949–961. doi: 10.1145/3332165.3347888.
- [98] S.-Y. Teng *et al.*, "TilePoP: Tile-type Pop-up Prop for Virtual Reality," in *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, New Orleans LA USA: ACM, Oct. 2019, pp. 639–649. doi: 10.1145/3332165.3347958.
- [99] Y. Mo, A. Song, and H. Qin, "Analysis and Performance Evaluation of a 3-DOF Wearable Fingertip Device for Haptic Applications," *IEEE Access*, vol. 7, pp. 170430–170441, 2019, doi: 10.1109/ACCESS.2019.2955601.
- [100] T. Amemiya, "Asymmetric Gravitational Oscillation on Fingertips Increased the Perceived Heaviness of a Pinched Object," in *Human Interface and the Management of Information. Information Presentation and Visualization*, vol. 12765, S. Yamamoto and H. Mori, Eds., in Lecture Notes in Computer Science, vol. 12765. , Cham: Springer International Publishing, 2021, pp. 247–256. Accessed: Mar. 04, 2025. [Online]. Available: https://link.springer.com/10.1007/978-3-030-78321-1_19
- [101] D. Chen, A. Song, L. Tian, L. Fu, and H. Zeng, "FW-Touch: A Finger Wearable Haptic Interface with a MR Foam Actuator for Displaying Surface Material Properties on a Touch Screen," *IEEE Trans. Haptics*, pp. 1–1, 2019, doi: 10.1109/toh.2019.2920349.
- [102] H.-R. Tsai, C. Tsai, Y.-S. Liao, Y.-T. Chiang, and Z.-Y. Zhang, "FingerX: Rendering Haptic Shapes of Virtual Objects Augmented by Real Objects using Extendable and Withdrawable Supports on Fingers," in *CHI Conference on Human Factors in Computing Systems*, in CHI '22. New York, NY, USA: Association for Computing Machinery, Apr. 2022, pp. 1–14. doi: 10.1145/3491102.3517489.
- [103] E. Strasnick, C. Holz, E. Ofek, M. Sinclair, and H. Benko, "Haptic Links: Bimanual Haptics for Virtual Reality Using Variable Stiffness Actuation," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, Montreal QC Canada: ACM Press, 2018, pp. 1–12. doi: 10.1145/3173574.3174218.
- [104] "Salter Brecknell ES-22 ElectroSamson Digital Hanging Scales, 22 x 0.2 lb - Coupons and Discounts May be Available." Accessed: Feb. 06, 2023. [Online]. Available: <https://www.scalesgalore.com/product/Salter-Brecknell-ES22-ElectroSamson-Digital-Hanging-Scales-22-x-002-lb-px7561.cfm>
- [105] C.-J. Lee, H.-R. Tsai, and B.-Y. Chen, "HairTouch: Providing Stiffness, Roughness and Surface Height Differences Using Reconfigurable Brush Hairs on a VR Controller,"

- in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, Yokohama Japan: ACM, May 2021, pp. 1–13. doi: 10.1145/3411764.3445285.
- [106] H.-R. Tsai, Y.-S. Liao, and C. Tsai, “ImpactVest: Rendering Spatio-Temporal Multilevel Impact Force Feedback on Body in VR,” in *CHI Conference on Human Factors in Computing Systems*, in CHI '22. New York, NY, USA: Association for Computing Machinery, Apr. 2022, pp. 1–11. doi: 10.1145/3491102.3501971.
- [107] H.-R. Tsai and B.-Y. Chen, “ElastImpac: 2.5D Multilevel Instant Impact Using Elasticity on Head-Mounted Displays,” in *Proceedings of the 32Nd Annual ACM Symposium on User Interface Software and Technology*, in UIST '19. New York, NY, USA: ACM, 2019, pp. 429–437. doi: 10.1145/3332165.3347931.
- [108] T.-Y. Wei *et al.*, “ElastiLinks: Force Feedback between VR Controllers with Dynamic Points of Application of Force,” in *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, in UIST '20. New York, NY, USA: Association for Computing Machinery, Oct. 2020, pp. 1023–1034. doi: 10.1145/3379337.3415836.
- [109] H.-R. Tsai *et al.*, “GuideBand: Intuitive 3D Multilevel Force Guidance on a Wristband in Virtual Reality,” in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, New York, NY, USA: Association for Computing Machinery, May 2021, pp. 1–13. Accessed: Jul. 06, 2021. [Online]. Available: <https://doi.org/10.1145/3411764.3445262>
- [110] P. Ke, S. Cai, H. Gao, and K. Zhu, “PropelWalker: A Leg-Based Wearable System With Propeller-Based Force Feedback for Walking in Fluids in VR,” *IEEE Trans. Visual. Comput. Graphics*, vol. 29, no. 12, pp. 1–14, 2022, doi: 10.1109/TVCG.2022.3205181.
- [111] C.-W. Hung, H.-R. Tsai, C.-C. Su, J.-C. Chiu, and B.-Y. Chen, “OsciHead: Simulating Versatile Force Feedback on an HMD by Rendering Various Types of Oscillation,” *Proc. ACM Hum.-Comput. Interact.*, vol. 6, no. MHCI, p. 180:1-180:23, Sep. 2022, doi: 10.1145/3546715.
- [112] Y.-S. Ye, H.-Y. Chen, and L. Chan, “Pull-Ups: Enhancing Suspension Activities in Virtual Reality with Body-Scale Kinesthetic Force Feedback,” in *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, New Orleans LA USA: ACM, Oct. 2019, pp. 791–801. doi: 10.1145/3332165.3347874.
- [113] D. Chiaradia, L. Tiseni, M. Xiloyannis, M. Solazzi, L. Masia, and A. Frisoli, “An Assistive Soft Wrist Exosuit for Flexion Movements With an Ergonomic Reinforced Glove,” *Front. Robot. AI*, vol. 7, p. 595862, 2021, doi: 10.3389/frobt.2020.595862.
- [114] M. A. Messerschmidt *et al.*, “ANISMA: A Prototyping Toolkit to Explore Haptic Skin Deformation Applications Using Shape-Memory Alloys,” *ACM Trans. Comput.-Hum. Interact.*, vol. 29, no. 3, p. 19:1-19:34, Jan. 2022, doi: 10.1145/3490497.
- [115] T. Yoshida, X. Shen, K. Yoshino, K. Nakagaki, and H. Ishii, “SCALE: Enhancing Force-based Interaction by Processing Load Data from Load Sensitive Modules,” in *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, New Orleans LA USA: ACM, Oct. 2019, pp. 901–911. doi: 10.1145/3332165.3347935.
- [116] N. S. Usevitch and A. A. Stanley, “Cutting the Cord: Soft Haptic Devices without a Pressure Source,” in *2019 2nd IEEE International Conference on Soft Robotics (RoboSoft)*, Seoul, Korea (South): IEEE, Apr. 2019, pp. 49–55. doi: 10.1109/robosoft.2019.8722740.

- [117] M. Zhu *et al.*, “PneuSleeve: In-fabric Multimodal Actuation and Sensing in a Soft, Compact, and Expressive Haptic Sleeve,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, in CHI ’20. Honolulu, HI, USA: Association for Computing Machinery, Apr. 2020, pp. 1–12. doi: 10.1145/3313831.3376333.
- [118] E. M. Young, A. H. Memar, P. Agarwal, and N. Colonnese, “Bellowband: A Pneumatic Wristband for Delivering Local Pressure and Vibration,” in *2019 IEEE World Haptics Conference (WHC)*, Tokyo, Japan: IEEE, Jul. 2019, pp. 55–60. doi: 10.1109/WHC.2019.8816075.
- [119] E. Pezent *et al.*, “Tasbi: Multisensory Squeeze and Vibrotactile Wrist Haptics for Augmented and Virtual Reality,” in *2019 IEEE World Haptics Conference (WHC)*, Tokyo, Japan: IEEE, Jul. 2019, pp. 1–6. doi: 10.1109/WHC.2019.8816098.
- [120] A. Badshah, S. Gupta, D. Morris, S. Patel, and D. Tan, “GyroTab: a handheld device that provides reactive torque feedback,” in *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI ’12*, Austin, Texas, USA: ACM Press, 2012, p. 3153. doi: 10.1145/2207676.2208731.
- [121] H. Miyahara, Y. Makino, and H. Shinoda, “Inducing Wrist Twist During Arm Swing by Using Gyro Effect,” *Haptic Interaction*, vol. 432, pp. 205–209, 2018, doi: 10.1007/978-981-10-4157-0_36.
- [122] J. Gugenheimer, D. Wolf, E. R. Eiriksson, P. Maes, and E. Rukzio, “GyroVR: Simulating Inertia in Virtual Reality Using Head Worn Flywheels,” in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, in UIST ’16. New York, NY, USA: ACM, 2016, pp. 227–232. doi: 10.1145/2984511.2984535.
- [123] J. T. Muth *et al.*, “Embedded 3D printing of strain sensors within highly stretchable elastomers,” *Advanced materials*, vol. 26, no. 36, pp. 6307–6312, 2014, doi: 10.1002/adma.201400334.
- [124] J. Yao, D. Cai, H. Zhang, H. Wang, D. Wu, and Y. Zhao, “Task-oriented design method and research on force compliant experiment of six-axis wrist force sensor,” in *Mechatronics*, Elsevier, 2016, pp. 109–121. doi: 10.1016/j.mechatronics.2016.01.007.
- [125] Y. Li, C. Yang, G. Wang, H. Zhang, H. Cui, and Y. Zhang, “Research on the parallel load sharing principle of a novel self-decoupled piezoelectric six-dimensional force sensor,” *ISA transactions*, vol. 70, pp. 447–457, 2017, doi: 10.1016/j.isatra.2017.07.008.
- [126] F. Becker, R. Jäger, F. Schmidt, B. Lapatki, and O. Paul, “Miniaturized six-degree-of-freedom force/moment transducers for instrumented teeth,” in *IEEE Sensors Journal*, IEEE, 2017, pp. 3644–3655. doi: 10.1109/JSEN.2017.2696035.
- [127] P. Estevez *et al.*, “6 DOF force and torque sensor for micro-manipulation applications,” *Sensors and Actuators A: Physical*, vol. 186, pp. 86–93, 2012, doi: 10.1016/j.sna.2012.02.037.
- [128] R. Bekhti, V. Duchaine, and P. Cardou, “Miniature capacitive three-axis force sensor,” in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Chicago, IL, USA: IEEE, Sep. 2014, pp. 3939–3946. doi: 10.1109/IROS.2014.6943116.
- [129] U. Kim, Y. B. Kim, D.-Y. Seok, J. So, and H. R. Choi, “A surgical palpation probe with 6-axis force/torque sensing capability for minimally invasive surgery,” *IEEE Transactions on Industrial Electronics*, vol. 65, no. 3, pp. 2755–2765, 2017, doi: 10.1109/TIE.2017.2739681.

- [130] D. Kim *et al.*, “Six-axis capacitive force/torque sensor based on dielectric elastomer,” presented at the Electroactive Polymer Actuators and Devices (EAPAD) 2013, SPIE, 2013, pp. 688–696.
- [131] J. Huang *et al.*, “Design of a Novel Six-Axis Force/Torque Sensor based on Optical Fibre Sensing for Robotic Applications:,” in *Proceedings of the 15th International Conference on Informatics in Control, Automation and Robotics*, Porto, Portugal: SCITEPRESS - Science and Technology Publications, 2018, pp. 517–524. doi: 10.5220/0006911705170524.
- [132] Y.-L. Park, K. Chau, R. J. Black, and M. R. Cutkosky, “Force sensing robot fingers using embedded fiber Bragg grating sensors and shape deposition manufacturing,” presented at the Proceedings 2007 IEEE International Conference on Robotics and Automation, IEEE, 2007, pp. 1510–1516.
- [133] N. Hendrich, F. Wasserfall, and J. Zhang, “3D Printed Low-Cost Force-Torque Sensors,” *IEEE Access*, vol. 8, pp. 140569–140585, 2020, doi: 10.1109/ACCESS.2020.3007565.
- [134] S. B. Kesner and R. D. Howe, “Design Principles for Rapid Prototyping Forces Sensors Using 3-D Printing,” *IEEE/ASME Trans. Mechatron.*, vol. 16, no. 5, pp. 866–870, Oct. 2011, doi: 10.1109/TMECH.2011.2160353.
- [135] Y. Noh *et al.*, “A three-axial body force sensor for flexible manipulators,” in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, Hong Kong, China: IEEE, May 2014, pp. 6388–6393. doi: 10.1109/ICRA.2014.6907802.
- [136] Y. Noh *et al.*, “A new miniaturised multi-axis force/torque sensors based on optoelectronic technology and simply-supported beam,” in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Daejeon, South Korea: IEEE, Oct. 2016, pp. 202–207. doi: 10.1109/IROS.2016.7759056.
- [137] G. Palli and S. Pirozzi, “An optical torque sensor for robotic applications,” *International Journal of Optomechatronics*, vol. 7, no. 4, pp. 263–282, 2013, doi: 10.1080/15599612.2013.879500.
- [138] G. Barnaby and A. Roudaut, “Mantis: A Scalable, Lightweight and Accessible Architecture to Build Multifunction Force Feedback Systems,” in *Proceedings of the 32Nd Annual ACM Symposium on User Interface Software and Technology*, in UIST '19. New York, NY, USA: ACM, 2019, pp. 937–948. doi: 10.1145/3332165.3347909.
- [139] Z. Chua and A. M. Okamura, “A Modular 3-Degree-of-Freedom Force Sensor for Robot-assisted Minimally Invasive Surgery Research,” Nov. 10, 2022, *arXiv:arXiv:2211.05428*. doi: 10.48550/arXiv.2211.05428.
- [140] C. Hatzfeld and T. A. Kern, Eds., *Engineering Haptic Devices: A Beginner's Guide*. in Springer Series on Touch and Haptic Systems. London: Springer London, 2014. doi: 10.1007/978-1-4471-6518-7.
- [141] C. Pacchierotti, S. Sinclair, M. Solazzi, A. Frisoli, V. Hayward, and D. Prattichizzo, “Wearable Haptic Systems for the Fingertip and the Hand: Taxonomy, Review, and Perspectives,” *IEEE Trans. Haptics*, vol. 10, no. 4, pp. 580–600, Oct. 2017, doi: 10.1109/toh.2017.2689006.
- [142] M. Moreyra and B. Hannaford, “A practical measure of dynamic response of haptic devices,” in *Proceedings. 1998 IEEE International Conference on Robotics and Automation (Cat. No.98CH36146)*, Leuven, Belgium: IEEE, May 1998, pp. 369–374 vol.1. doi: 10.1109/ROBOT.1998.676430.

- [143] B. Wu, R. L. Klatzky, and R. L. Hollis, "Force, Torque, and Stiffness: Interactions in Perceptual Discrimination," *IEEE Trans. Haptics*, vol. 4, no. 3, pp. 221–228, Jul. 2011, doi: 10.1109/toh.2011.3.
- [144] "HX711 noise-free resolution tests," Arduino Forum. Accessed: Mar. 03, 2023. [Online]. Available: <https://forum.arduino.cc/t/hx711-noise-free-resolution-tests/351959>
- [145] T. Brooks, "Telerobot Response: Requirements (Tech. Rep. No. STX/ROB/90-03)," *Lanham, MD: STX Robotics*, 1990.
- [146] T. R. Coles, D. Meglan, and N. W. John, "The Role of Haptics in Medical Training Simulators: A Survey of the State of the Art," *IEEE Trans. Haptics*, vol. 4, no. 1, pp. 51–66, Jan. 2011, doi: 10.1109/TOH.2010.19.
- [147] G. C. Burdea, *Force and touch feedback for virtual reality*. USA: John Wiley & Sons, Inc., 1996.
- [148] S. Booth, F. Angelis, and T. Schmidt-Tjarksen, "The influence of changing haptic refresh-rate on subjective user experiences-lessons for effective touch-based applications," presented at the Proceedings of eurohaptics, Citeseer, 2003, pp. 374–383.
- [149] M. Scandola, M. Vicentini, L. Gasperotti, D. Zerbato, and P. Florini, "Force feedback in psychophysics research: even low performance algorithms may lead to realistic perceptual experience," *Proceedings of Fechner Day*, vol. 27, pp. 267–272, 2011.
- [150] M. Scandola, M. Vicentini, and P. Fiorini, "How force perception changes in different refresh rate conditions," in *2011 15th International Conference on Advanced Robotics (ICAR)*, Tallinn, Estonia: IEEE, Jun. 2011, pp. 322–327. doi: 10.1109/ICAR.2011.6088597.
- [151] E. F. Austin, "A Fabricated Force Glove That Measures Hand Forces during Activities of Daily Living," *Sensors*, vol. 22, no. 4, p. 10, 2022, doi: 10.3390/s22041330.
- [152] C. Rognon, M. Koehler, C. Duriez, D. Floreano, and A. M. Okamura, "Soft Haptic Device to Render the Sensation of Flying Like a Drone," *IEEE Robot. Autom. Lett.*, vol. 4, no. 3, pp. 2524–2531, Jul. 2019, doi: 10.1109/LRA.2019.2907432.
- [153] J. Realmuto and T. Sanger, "A robotic forearm orthosis using soft fabric-based helical actuators," in *2019 2nd IEEE International Conference on Soft Robotics (RoboSoft)*, Seoul, Korea (South): IEEE, Apr. 2019, pp. 591–596. doi: 10.1109/robosoft.2019.8722759.
- [154] H. Nagano, K. Sase, M. Konyo, and S. Tadokoro, "Wearable suction haptic display with spatiotemporal stimulus distribution on a finger pad," in *2019 IEEE world haptics conference (WHC)*, IEEE, pp. 389–394. doi: 10.1109/whc.2019.8816156.
- [155] S. Schreven, P. J. Beek, and J. B. J. Smeets, "Optimising filtering parameters for a 3D motion analysis system," *Journal of Electromyography and Kinesiology*, vol. 25, no. 5, pp. 808–814, Oct. 2015, doi: 10.1016/j.jelekin.2015.06.004.
- [156] R. Bartlett, *Introduction to sports biomechanics: analysing human movement patterns*, 2nd ed. London: Routledge, 2007.
- [157] J. L. A, "Kinesthetic sensing," *Proc. Workshop on Human and Machine Haptics, Pacific Grove, CA, 1997*, pp. 1–10, 1997.
- [158] H.-Y. Chang, W.-J. Tseng, C.-E. Tsai, H.-Y. Chen, R. L. Peiris, and L. Chan, "FacePush: Introducing Normal Force on Face with Head-Mounted Displays," in *The*

- 31st Annual ACM Symposium on User Interface Software and Technology - UIST '18*, Berlin, Germany: ACM Press, 2018, pp. 927–935. doi: 10.1145/3242587.3242588.
- [159] C. Fang, Y. Zhang, M. Dworman, and C. Harrison, “Wireality: Enabling Complex Tangible Geometries in Virtual Reality with Worn Multi-String Haptics,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, Honolulu HI USA: ACM, 2020, p. 10. doi: 10.1145/3313831.3376470.
- [160] “The Arena Infinity — Product — Striker VR,” Advanced force feedback — United States — Striker VR. Accessed: Jul. 08, 2018. [Online]. Available: <https://www.strikervr.com/arena-infinity>
- [161] “MAG P90 Gun Controller VR Rifle Xbox One— PS4PRO— HTC Vive— Oculus— Windows Mixed Reality— BeswinVR.” Accessed: Jan. 04, 2019. [Online]. Available: <https://www.magp90.com/>
- [162] “VR Gun Controller: Buy The Best - The Delta Six,” Avenger. Accessed: Nov. 22, 2017. [Online]. Available: <https://www.avengercontroller.com/vr-gun-controller/>
- [163] “haptech — products,” haptech. Accessed: Nov. 22, 2017. [Online]. Available: <https://www.haptech.co/products>
- [164] “Nintendo NES Zapper - Computing History.” Accessed: Jan. 26, 2019. [Online]. Available: <http://www.computinghistory.org.uk/det/37107/Nintendo-NES-Zapper/>
- [165] N. Cooper, F. Milella, I. Cant, C. Pinto, M. White, and G. Meyer, “Augmented Cues Facilitate Learning Transfer from Virtual to Real Environments,” in *Mixed and Augmented Reality (ISMAR-Adjunct), 2016 IEEE International Symposium on*, Merida, Yucatan, Mexico: IEEE, 2016, pp. 194–198. doi: 10.1109/ismar-adjunct.2016.0075.
- [166] Y. Tokuda, M. A. Norasikin, S. Subramanian, and D. Martinez Plasencia, “MistForm: Adaptive Shape Changing Fog Screens,” in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, Denver Colorado USA: ACM Press, 2017, pp. 4383–4395. doi: 10.1145/3025453.3025608.
- [167] R. Tamborini and N. D. Bowman, “Presence in video games,” *Immersed in media: Telepresence in everyday life*, pp. 87–109, 2010.
- [168] R. Arsenault and C. Ware, “Eye-hand co-ordination with force feedback,” in *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, in CHI '00. New York, NY, USA: Association for Computing Machinery, Apr. 2000, pp. 408–414. doi: 10.1145/332040.332466.
- [169] L. R. Krol, D. Aliakseyeu, and S. Subramanian, “Haptic feedback in remote pointing,” in *CHI'09 Extended Abstracts on Human Factors in Computing Systems*, Boston MA USA: ACM, 2009, pp. 3763–3768. doi: 10.1145/1520340.1520568.
- [170] F. Morelli *et al.*, “Shooter–System Performance Variability as a Function of Recoil Dynamics,” *Hum Factors*, vol. 59, no. 6, pp. 973–985, Sep. 2017, doi: 10.1177/0018720817700537.
- [171] R. J. Spine, “Recoil in shoulder fired weapons: a review of the literature,” *US Army Human Eng. Laboratory (HEL), US Army Armament Res. Develop. Command (ARDEC), Dover, NJ, USA, Technical report*, vol. 296, 1982.
- [172] K. Blankenship, R. Evans, S. Allison, M. Murphy, and H. Isome, “Shoulder-Fired Weapons with High Recoil Energy - Quantifying Injury and Shooting,” May 2004, Accessed: May 09, 2018. [Online]. Available: <http://www.dtic.mil/docs/citations/ADA425518>

- [173] “3D Systems Phantom Premium,” 3D Systems. Accessed: Jan. 06, 2018. [Online]. Available: <https://uk.3dsystems.com/haptics-devices/3d-systems-phantom-premium>
- [174] “HAPTION SA - Virtuose™ 6D TAO.” Accessed: Jan. 06, 2018. [Online]. Available: <https://www.haption.com/en/products-en/virtuose-6d-tao-en.html>
- [175] A. Mastmeyer, T. Hecht, D. Fortmeier, and H. Handels, “Ray-casting based evaluation framework for haptic force feedback during percutaneous transhepatic catheter drainage punctures,” *Int J Comput Assist Radiol Surg*, vol. 9, no. 3, pp. 421–431, May 2014, doi: 10.1007/s11548-013-0959-7.
- [176] P. Trier, K. Ø. Noe, M. S. Sørensen, and J. Mosegaard, “The visible ear surgery simulator,” *Stud Health Technol Inform*, vol. 132, pp. 523–525, 2008.
- [177] C. Preusche, T. Ortmaier, and G. Hirzinger, “Teleoperation concepts in minimal invasive surgery,” *Control Engineering Practice*, vol. 10, pp. 1245–1250, Nov. 2002, doi: 10.1016/S0967-0661(02)00084-9.
- [178] H. Culbertson, J. M. Walker, M. Raitor, and A. M. Okamura, “WAVES: A Wearable Asymmetric Vibration Excitation System for Presenting Three-Dimensional Translation and Rotation Cues,” in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, Denver Colorado USA: ACM Press, 2017, pp. 4972–4982. doi: 10.1145/3025453.3025741.
- [179] J. Rekimoto, “Traxion: A Tactile Interaction Device with Virtual Force Sensation,” in *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*, in UIST '13. St. Andrews Scotland, United Kingdom: ACM, 2013, pp. 427–432. doi: 10.1145/2501988.2502044.
- [180] P. Lopes, A. Ion, and P. Baudisch, “Impacto: Simulating Physical Impact by Combining Tactile Stimulation with Electrical Muscle Stimulation,” in *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, in UIST '15. Charlotte NC USA: ACM, 2015, pp. 11–19. doi: 10.1145/2807442.2807443.
- [181] E. Whitmire, H. Benko, C. Holz, E. Ofek, and M. Sinclair, “Haptic Revolver: Touch, Shear, Texture, and Shape Rendering on a Reconfigurable Virtual Reality Controller,” ACM Press, 2018, pp. 1–12. doi: 10.1145/3173574.3173660.
- [182] K. Mononen, “The Effects of Augmented Feedback on Motor Skill Learning in Shooting,” p. 64, 2007.
- [183] H. M. Rao *et al.*, “Sensorimotor Learning during a Marksmanship Task in Immersive Virtual Reality,” *Front Psychol*, vol. 9, p. 58, 2018, doi: 10.3389/fpsyg.2018.00058.
- [184] S. Nisar, M. O. Martinez, T. Endo, F. Matsuno, and A. M. Okamura, “Effects of Different Hand-Grounding Locations on Haptic Performance With a Wearable Kinesthetic Haptic Device,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 351–358, Apr. 2019, doi: 10.1109/lra.2018.2890198.
- [185] J. T. Viitasalo, P. Era, N. Konttinen, H. Mononen, K. Mononen, and K. Norvapalo, “Effects of 12-week shooting training and mode of feedback on shooting scores among novice shooters,” *Scandinavian Med Sci Sports*, vol. 11, no. 6, pp. 362–368, 2001, doi: 10.1034/j.1600-0838.2001.110608.x.
- [186] “The information capacity of the human motor system in controlling the amplitude of movement. - PsycNET.” Accessed: Dec. 08, 2022. [Online]. Available: <https://psycnet.apa.org/record/1955-02059-001>
- [187] L. Bainbridge and M. Sanders, “The generality of Fitts’s Law,” *Journal of Experimental Psychology*, vol. 96, no. 1, pp. 130–133, 1972, doi: 10.1037/h0033458.

- [188] M. Akamatsu, I. S. MacKenzie, and T. Hasbroucq, "A comparison of tactile, auditory, and visual feedback in a pointing task using a mouse-type device," *Ergonomics*, vol. 38, no. 4, pp. 816–827, 1995, doi: 10.1080/00140139508925152.
- [189] M. Akamatsu and I. S. MacKenzie, "Movement characteristics using a mouse with tactile and force feedback," *International Journal of Human-Computer Studies*, vol. 45, no. 4, pp. 483–493, Oct. 1996, doi: 10.1006/ijhc.1996.0063.
- [190] A. Cockburn and S. Brewster, "Multimodal feedback for the acquisition of small targets," *Ergonomics*, vol. 48, no. 9, pp. 1129–1150, 2005, doi: 10.1080/00140130500197260.
- [191] K. K. Bhagat, W.-K. Liou, and C.-Y. Chang, "A cost-effective interactive 3D virtual reality system applied to military live firing training," *Virtual Reality*, vol. 20, no. 2, pp. 127–140, Jun. 2016, doi: 10.1007/s10055-016-0284-x.
- [192] S. Bateman, R. L. Mandryk, C. Gutwin, and R. Xiao, "Analysis and comparison of target assistance techniques for relative ray-cast pointing," *International Journal of Human-Computer Studies*, vol. 71, no. 5, pp. 511–532, May 2013, doi: 10.1016/j.ijhcs.2012.12.006.
- [193] B. C. Witmer, J. H. Bailey, and B. W. Knerr, "Training Dismounted Soldiers in Virtual Environments: Route Learning and Transfer.," ARMY RESEARCH INST FOR THE BEHAVIOURAL AND SOCIAL SCIENCES ORLANDO FL ORLANDO FIELD UNIT*, ARI-TR-1022, Feb. 1995. Accessed: Jan. 12, 2019. [Online]. Available: <https://apps.dtic.mil/docs/citations/ADA292900>
- [194] S. Bateman, R. Mandryk, C. Gutwin, and R. Xiao, "Investigation of Targeting-Assistance Techniques for Distant Pointing with Relative Ray Casting," p. 10.
- [195] B. Canfield-Hershkowitz, T. Foster, and W. Meijer, "Rifle and Shotgun Recoil Test System," p. 79.
- [196] M. J. Hall, "Measuring felt recoil of sporting arms," in *International Journal of Impact Engineering*, Jun. 2008, pp. 540–548. doi: 10.1016/j.ijimpeng.2007.03.007.
- [197] A. Koptuyug and M. Ainegren, "Experimental Measurement of Rifle Dynamics During the Range Shooting of Biathlon Weapons," *Procedia Engineering*, vol. 112, pp. 349–354, 2015, doi: 10.1016/j.proeng.2015.07.261.
- [198] A. G. Cleary, H. McKendrick, and J. A. Sills, "Hand-arm vibration syndrome may be associated with prolonged use of vibrating computer games," *BMJ*, vol. 324, no. 7332, p. 301, Feb. 2002, doi: 10.1136/bmj.324.7332.301a.
- [199] "Measuring the Recoil Force from a Gun." Accessed: Sep. 19, 2018. [Online]. Available: <http://www.loadstarsensors.com/blog/recoil-force-from-a-gunshot.html>
- [200] R. Vicencio-Moreira, R. L. Mandryk, C. Gutwin, and S. Bateman, "The effectiveness (or lack thereof) of aim-assist techniques in first-person shooter games," ACM Press, 2014, pp. 937–946. doi: 10.1145/2556288.2557308.
- [201] V. Mathiowetz, K. Weber, G. Volland, and N. Kashman, "Reliability and validity of grip and pinch strength evaluations," *Journal of Hand Surgery*, vol. 9, no. 2, pp. 222–226, 1984, doi: 10.1016/s0363-5023(84)80146-x.
- [202] E. D. Mekler, F. Brühlmann, K. Opwis, and A. N. Tuch, "Disassembling Gamification: The Effects of Points and Meaning on User Motivation and Performance," in *CHI '13 Extended Abstracts on Human Factors in Computing Systems*, in CHI EA '13. New York, NY, USA: ACM, 2013, pp. 1137–1142. doi: 10.1145/2468356.2468559.

- [203] M. Botvinick and J. Cohen, "Rubber hands 'feel' touch that eyes see," *Nature*, vol. 391, no. 6669, p. 756, 1998, doi: 10.1038/35784.
- [204] R. Salomon *et al.*, "Changing motor perception by sensorimotor conflicts and body ownership," *Sci Rep*, vol. 6, p. 25847, 2016, doi: 10.1038/srep25847.
- [205] M. Azmandian, M. Hancock, H. Benko, E. Ofek, and A. D. Wilson, "Haptic Retargeting: Dynamic Repurposing of Passive Haptics for Enhanced Virtual Reality Experiences," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, in CHI '16. New York, NY, USA: ACM, 2016, pp. 1968–1979. doi: 10.1145/2858036.2858226.
- [206] Y. Zhang, D. Wang, Z. Wang, Y. Zhang, and J. Xiao, "Passive Force-feedback Gloves with Joint-based Variable Impedance Using Layer Jamming," *IEEE Trans. Haptics*, vol. 12, no. 3, pp. 1–1, 2019, doi: 10.1109/toh.2019.2908636.
- [207] S. B. Schorr and A. M. Okamura, "Fingertip Tactile Devices for Virtual Object Manipulation and Exploration," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, in CHI '17. New York, NY, USA: ACM, 2017, pp. 3115–3119. doi: 10.1145/3025453.3025744.
- [208] V. Hayward, O. R. Astley, M. Cruz-Hernandez, D. Grant, and G. Robles-De-La-Torre, "Haptic interfaces and devices," in *Sensor Review*, 2004, pp. 16–29. doi: 10.1108/02602280410515770.
- [209] Y. Ban, T. Narumi, T. Tanikawa, and M. Hirose, "Air haptics: displaying feeling of contact with AR object using visuo-haptic interaction," in *ACM SIGGRAPH 2015 Emerging Technologies on - SIGGRAPH '15*, Los Angeles, California: ACM Press, 2015, pp. 1–1. doi: 10.1145/2782782.2792498.
- [210] I. Oakley, M. R. McGee, S. Brewster, and P. Gray, "Putting the feel in 'look and feel'," in *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '00*, The Hague, The Netherlands: ACM Press, 2000, pp. 415–422. doi: 10.1145/332040.332467.
- [211] T. Han, F. Anderson, P. Irani, and T. Grossman, "HydroRing: Supporting Mixed Reality Haptics Using Liquid Flow," in *The 31st Annual ACM Symposium on User Interface Software and Technology*, in UIST '18. New York, NY, USA: ACM, 2018, pp. 913–925. doi: 10.1145/3242587.3242667.
- [212] D. Prattichizzo, F. Chinello, C. Pacchierotti, and M. Malvezzi, "Towards wearability in fingertip haptics: a 3-dof wearable device for cutaneous force feedback," *IEEE Transactions on Haptics*, vol. 6, no. 4, pp. 506–516, 2013, doi: 10.1109/toh.2013.53.
- [213] R. Scheibe, M. Moehring, and B. Froehlich, "Tactile Feedback at the Finger Tips for Improved Direct Interaction in Immersive Environments," *2007 IEEE Symposium on 3D User Interfaces*, Mar. 2007, doi: 10.1109/3dui.2007.340784.
- [214] J. Lee, M. Sinclair, M. Gonzalez-Franco, E. Ofek, and C. Holz, "TORC: A Virtual Reality Controller for In-Hand High-Dexterity Finger Interaction," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, in CHI '19. New York, NY, USA: ACM, 2019, p. 71:1-71:13. doi: 10.1145/3290605.3300301.
- [215] T. Handa, M. Azuma, T. Shimizu, and S. Kondo, "GraspForm: Shape and Hardness Rendering on Handheld Device Toward Universal Interface for 3D Haptic TV," in *Adjunct Publication of the 30th Annual ACM Symposium on User Interface Software and Technology*, in UIST '17. New York, NY, USA: ACM, 2017, pp. 147–149. doi: 10.1145/3131785.3131824.

- [216] Y. Lee and D. Ryu, “Wearable haptic glove using micro hydraulic system for control of construction robot system with VR environment,” in *2008 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, Seoul: IEEE, Aug. 2008, pp. 638–643. doi: 10.1109/mfi.2008.4648016.
- [217] F. Chinello, M. Malvezzi, C. Pacchierotti, and D. Prattichizzo, “Design and development of a 3RRS wearable fingertip cutaneous device,” in *2015 IEEE International Conference on Advanced Intelligent Mechatronics (AIM)*, Busan, South Korea: IEEE, Jul. 2015, pp. 293–298. doi: 10.1109/aim.2015.7222547.
- [218] K. Minamizawa, S. Fukamachi, H. Kajimoto, N. Kawakami, and S. Tachi, “Wearable haptic display to present virtual mass sensation,” in *ACM SIGGRAPH 2007 sketches*, San Diego California: ACM Press, 2007, pp. 43-es. doi: 10.1145/1278780.1278833.
- [219] N. G. Tsagarakis, T. Horne, and D. G. Caldwell, “SLIP AESTHEASIS: A Portable 2D Slip/Skin Stretch Display for the Fingertip,” in *Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, in WHC ’05. Washington, DC, USA: IEEE Computer Society, 2005, pp. 214–219. doi: 10.1109/whc.2005.117.
- [220] T. Murakami, T. Person, C. L. Fernando, and K. Minamizawa, “Altered Touch: Miniature Haptic Display with Force, Thermal and Tactile Feedback for Augmented Haptics,” in *ACM SIGGRAPH 2017 Emerging Technologies*, in SIGGRAPH ’17. New York, NY, USA: ACM, 2017, p. 2:1-2:2. doi: 10.1145/3084822.3084836.
- [221] V. Yem and H. Kajimoto, “Wearable tactile device using mechanical and electrical stimulation for fingertip interaction with virtual world,” in *2017 IEEE Virtual Reality (VR)*, Los Angeles, CA, USA: IEEE, Mar. 2017, pp. 99–104. doi: 10.1109/vr.2017.7892236.
- [222] H. G. Debarba, J.-N. Khoury, S. Perrin, B. Herbelin, and R. Boulic, “Perception of redirected pointing precision in immersive virtual reality,” 2018. doi: 10.1109/vr.2018.8448285.
- [223] S. Razzaque, Z. Kohn, and M. C. Whitton, “Redirected walking,” Citeseer, 2005.
- [224] B. Williams, G. Narasimham, T. P. McNamara, T. H. Carr, J. J. Rieser, and B. Bodenheimer, “Updating orientation in large virtual environments using scaled translational gain,” in *Proceedings of the 3rd symposium on Applied perception in graphics and visualization*, ACM, pp. 21–28. doi: 10.1145/1140491.1140495.
- [225] R. A. Montano Murillo, E. Gatti, M. Oliver Segovia, M. Obrist, J. P. Molina Masso, and D. Martinez Plasencia, “NaviFields: Relevance fields for adaptive VR navigation,” in *Proceedings of the 30th annual ACM symposium on user interface software and technology*, ACM, pp. 747–758. doi: 10.1145/3126594.3126645.
- [226] T. C. Peck, H. Fuchs, and M. C. Whitton, “Improved redirection with distractors: A large-scale-real-walking locomotion interface and its effect on navigation in virtual environments,” in *2010 IEEE virtual reality conference (VR)*, IEEE, pp. 35–38. doi: 10.1109/vr.2010.5444816.
- [227] R. A. Montano Murillo, S. Subramanian, and D. Martinez Plasencia, “Erg-O: ergonomic optimization of immersive virtual environments,” *Proceedings of the 30th annual ACM symposium on user interface software and technology*. ACM, pp. 759–771. doi: 10.1145/3126594.3126605.
- [228] J. (Junrui) Yang, H. Horii, A. Thayer, and R. Ballagas, “VR Grabbers: Ungrounded Haptic Retargeting for Precision Grabbing Tools,” in *The 31st Annual ACM Symposium on User Interface Software and Technology*, in UIST ’18. New York, NY, USA: ACM, 2018, pp. 889–899. doi: 10.1145/3242587.3242643.

- [229] P. Abtahi and S. Follmer, “Visuo-haptic illusions for improving the perceived performance of shape displays,” in *Proceedings of the 2018 CHI conference on human factors in computing systems*, ACM, p. 150. doi: 10.1145/3173574.3173724.
- [230] L.-P. Cheng, E. Ofek, C. Holz, H. Benko, and A. D. Wilson, “Sparse Haptic Proxy: Touch Feedback in Virtual Environments Using a General Passive Prop,” in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, Denver Colorado USA: ACM Press, 2017, pp. 3718–3728. doi: 10.1145/3025453.3025753.
- [231] Y. Zhao and S. Follmer, “A functional optimization based approach for continuous 3D retargeted touch of arbitrary, complex boundaries in haptic virtual reality,” in *Proceedings of the 2018 CHI conference on human factors in computing systems*, ACM, p. 544. doi: 10.1145/3173574.3174118.
- [232] M. Rietzler, F. Geiselhart, J. Gugenheimer, and E. Rukzio, “Breaking the Tracking: Enabling Weight Perception Using Perceivable Tracking Offsets,” in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, in CHI '18. New York, NY, USA: ACM, 2018, p. 128:1-128:12. doi: 10.1145/3173574.3173702.
- [233] S.-J. Blakemore, C. D. Frith, and D. M. Wolpert, “Spatio-temporal prediction modulates the perception of self-produced stimuli,” *Journal of cognitive neuroscience*, vol. 11, no. 5. pp. 551–559, 1999. doi: 10.1162/089892999563607.
- [234] L. Weiskrantz, J. Elliott, and C. Darlington, “Preliminary observations on tickling oneself,” in *Nature*, 1971, p. 598. doi: 10.1038/230598a0.
- [235] K. Kilteni and H. H. Ehrsson, “Body ownership determines the attenuation of self-generated tactile sensations,” *Proc. Natl. Acad. Sci. U.S.A.*, vol. 114, no. 31, pp. 8426–8431, Aug. 2017, doi: 10.1073/pnas.1703347114.
- [236] D. M. Wolpert, R. C. Miall, and M. Kawato, “Internal models in the cerebellum,” *Trends in cognitive sciences*, vol. 2, no. 9, pp. 338–347, 1998, doi: 10.1016/s1364-6613(98)01221-2.
- [237] D. M. Wolpert and M. Kawato, “Multiple paired forward and inverse models for motor control,” in *Neural networks*, 1998, pp. 1317–1329. doi: 10.1016/s0893-6080(98)00066-5.
- [238] J. R. Flanagan, M. C. Bowman, and R. S. Johansson, “Control strategies in object manipulation tasks,” *Current opinion in neurobiology*, vol. 16, no. 6. pp. 650–659, 2006. doi: 10.1016/j.conb.2006.10.005.
- [239] “Why can’t you tickle yourself?,” 2000.
- [240] P. M. Bays, D. M. Wolpert, E. P. Haggard, Y. Rosetti, and M. Kawato, “Predictive attenuation in the perception of touch,” *Sensorimotor foundations of higher cognition*, vol. 22. pp. 339–358, 2008.
- [241] S. S. Shergill, G. Samson, P. M. Bays, C. D. Frith, and D. M. Wolpert, “Evidence for Sensory Prediction Deficits in Schizophrenia,” *AJP*, vol. 162, no. 12, pp. 2384–2386, Dec. 2005, doi: 10.1176/appi.ajp.162.12.2384.
- [242] S. S. Shergill, P. M. Bays, C. D. Frith, and D. M. Wolpert, “Two Eyes for an Eye: The Neuroscience of Force Escalation,” *Science*, vol. 301, no. 5630, pp. 187–187, Jul. 2003, doi: 10.1126/science.1085327.
- [243] M. Voss, P. M. Bays, J. C. Rothwell, and D. M. Wolpert, “An improvement in perception of self-generated tactile stimuli following theta-burst stimulation of primary

- motor cortex,” *Neuropsychologia*, vol. 45, no. 12, pp. 2712–2717, 2007, doi: 10.1016/j.neuropsychologia.2007.04.008.
- [244] G. Cumming, F. Fidler, and D. L. Vaux, “Error bars in experimental biology,” *The Journal of cell biology*, vol. 177, no. 1. pp. 7–11, 2007. doi: 10.1083/jcb.200611141.
- [245] “Logitech G29 Driving Force Steering Wheels & Pedals.” Accessed: Nov. 29, 2023. [Online]. Available: <https://www.logitechg.com/en-eu/products/driving/driving-force-racing-wheel.html>
- [246] “Microsoft SideWinder,” *Wikipedia*. May 13, 2022. Accessed: Nov. 29, 2023. [Online]. Available: [https://en.wikipedia.org/w/index.php?title=Microsoft`SideWinder&oldid=1087594017](https://en.wikipedia.org/w/index.php?title=Microsoft%27SideWinder&oldid=1087594017)
- [247] M. Sinclair, E. Ofek, M. Gonzalez-Franco, and C. Holz, “CapstanCrunch: A Haptic VR Controller with User-supplied Force Feedback,” in *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, New Orleans LA USA: ACM, Oct. 2019, pp. 815–829. doi: 10.1145/3332165.3347891.
- [248] S. Heo, C. Chung, G. Lee, and D. Wigdor, “Thor’s Hammer: An Ungrounded Force Feedback Device Utilising Propeller-Induced Propulsive Force,” in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, Montreal QC Canada: ACM, Apr. 2018, pp. 1–11. doi: 10.1145/3173574.3174099.
- [249] S. Je *et al.*, “Aero-plane: A Handheld Force-Feedback Device that Renders Weight Motion Illusion on a Virtual 2D Plane,” in *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, New Orleans LA USA: ACM, Oct. 2019, pp. 763–775. doi: 10.1145/3332165.3347926.
- [250] “CyberGrasp — CyberGlove Systems LLC.” Accessed: Nov. 29, 2023. [Online]. Available: <http://www.cyberglovesystems.com/cybergasp/>
- [251] X. Gu, Y. Zhang, W. Sun, Y. Bian, D. Zhou, and P. O. Kristensson, “Dexmo: An Inexpensive and Lightweight Mechanical Exoskeleton for Motion Capture and Force Feedback in VR,” in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, San Jose California USA: ACM, May 2016, pp. 1991–1995. doi: 10.1145/2858036.2858487.
- [252] A. Zenner and A. Kruger, “Shifty: A Weight-Shifting Dynamic Passive Haptic Proxy to Enhance Object Perception in Virtual Reality,” *IEEE Trans. Visual Comput. Graphics*, vol. 23, no. 4, pp. 1285–1294, Apr. 2017, doi: 10.1109/TVCG.2017.2656978.
- [253] J. Shigeyama *et al.*, “Transcalibur: weight moving VR controller for dynamic rendering of 2D shape using haptic shape illusion,” in *ACM SIGGRAPH 2018 Emerging Technologies*, Vancouver British Columbia Canada: ACM, Aug. 2018, pp. 1–2. doi: 10.1145/3214907.3214923.
- [254] A. Zenner and A. Krüger, “Drag:on: A Virtual Reality Controller Providing Haptic Feedback Based on Drag and Weight Shift,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, Glasgow Scotland Uk: ACM, May 2019, pp. 1–12. doi: 10.1145/3290605.3300441.
- [255] “HapticWhril,” HapticWhril. Accessed: Nov. 29, 2023. [Online]. Available: <https://telenaco.github.io/HapticWhril/>
- [256] H. Yano, M. Yoshie, and H. Iwata, “Development of a non-grounded haptic interface using the gyro effect,” in *11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2003. HAPTICS 2003. Proceedings.*, Los

- Angeles, CA, USA: IEEE Comput. Soc, 2003, pp. 32–39. doi: 10.1109/HAPTIC.2003.1191223.
- [257] H. Miyahara, Y. Makino, and H. Shinoda, “Inducing Wrist Twist During Arm Swing by Using Gyro Effect,” in *Haptic Interaction*, vol. 432, S. Hasegawa, M. Konyo, K.-U. Kyung, T. Nojima, and H. Kajimoto, Eds., in Lecture Notes in Electrical Engineering, vol. 432, Singapore: Springer Singapore, 2018, pp. 205–209. doi: 10.1007/978-981-10-4157-0_36.
- [258] M. Sakai, Y. Fukui, and N. Nakamura, “Effective Output Patterns for Torque Display,” in *Online Proceeding of the 13th International Conference on Artificial Reality and Telexistence, December 3-5, 2003, Keio University, Tokyo, Japan, 2003*. [Online]. Available: <http://www.vrsj.org/ic-at/ICAT2003/php/upload/116%20camera%20b039e8e145a858b48ca6d93123d21f38.pdf>
- [259] Y. Tanaka, S. Masataka, K. Yuka, Y. Fukui, J. Yamashita, and N. Nakamura, “Mobile torque display and haptic characteristics of human palm,” in *Proc. ICAT*, 2001, pp. 115–120.
- [260] T. Amemiya and H. Gomi, “Directional Torque Perception with Brief, Asymmetric Net Rotation of a Flywheel,” *IEEE Trans. Haptics*, vol. 6, no. 3, pp. 370–375, Jul. 2013, doi: 10.1109/TOH.2012.38.
- [261] H. Ando, K. Obana, M. Sugimoto, and T. Maeda, “A wearable force display based on brake change in angular momentum,” *Proc Artificial Reality and Telexistence 2002*, pp. 16–21, 2002.
- [262] M. Antolini, M. Bordegoni, and U. Cugini, “A haptic direction indicator using the gyro effect,” in *2011 IEEE World Haptics Conference*, Istanbul: IEEE, Jun. 2011, pp. 251–256. doi: 10.1109/WHC.2011.5945494.
- [263] J. M. Walker, H. Culbertson, M. Raitor, and A. M. Okamura, “Haptic orientation guidance using two parallel double-gimbal control moment gyroscopes,” *IEEE Trans. Haptics*, vol. 11, no. 2, pp. 267–278, Apr. 2018, doi: 10.1109/TOH.2017.2713380.
- [264] M. Murer, B. Maurer, H. Huber, I. Aslan, and M. Tscheligi, “TorqueScreen: Actuated Flywheels for Ungrounded Kinaesthetic Feedback in Handheld Devices,” in *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction*, Stanford California USA: ACM, Jan. 2015, pp. 161–164. doi: 10.1145/2677199.2680579.
- [265] A. Badshah, S. Gupta, D. Morris, S. Patel, and D. Tan, “GyroTab: a handheld device that provides reactive torque feedback,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Austin Texas USA: ACM, May 2012, pp. 3153–3156. doi: 10.1145/2207676.2208731.
- [266] K. N. Winfree, J. Gewirtz, T. Mather, J. Fiene, and K. J. Kuchenbecker, “A high fidelity ungrounded torque feedback device: The iTorqU 2.0,” in *World Haptics 2009 - Third Joint EuroHaptics conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, Salt Lake City, UT, USA: IEEE, 2009, pp. 261–266. doi: 10.1109/WHC.2009.4810866.
- [267] J. Gugenheimer, D. Wolf, E. R. Eiriksson, P. Maes, and E. Rukzio, “GyroVR: Simulating Inertia in Virtual Reality using Head Worn Flywheels,” in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, Tokyo Japan: ACM, Oct. 2016, pp. 227–232. doi: 10.1145/2984511.2984535.
- [268] T. Nakamura and H. Koike, “Golf Club-Type Device with Force Feedback for Modifying Club Posture,” in *Extended Abstracts of the 2020 CHI Conference on*

- Human Factors in Computing Systems*, Honolulu HI USA: ACM, Apr. 2020, pp. 1–7. doi: 10.1145/3334480.3383024.
- [269] H.-Y. Chang, W.-J. Tseng, C.-E. Tsai, H.-Y. Chen, R. L. Peiris, and L. Chan, “FacePush: Introducing Normal Force on Face with Head-Mounted Displays,” in *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, Berlin Germany: ACM, Oct. 2018, pp. 927–935. doi: 10.1145/3242587.3242588.
- [270] S. Je, H. Lee, M. J. Kim, and A. Bianchi, “Wind-blaster: a wearable propeller-based prototype that provides ungrounded force-feedback,” in *ACM SIGGRAPH 2018 Emerging Technologies*, Vancouver British Columbia Canada: ACM, Aug. 2018, pp. 1–2. doi: 10.1145/3214907.3214915.
- [271] Y. Sun, S. Yoshida, T. Narumi, and M. Hirose, “PaCaPa: A Handheld VR Device for Rendering Size, Shape, and Stiffness of Virtual Objects in Tool-based Interactions,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, Glasgow Scotland Uk: ACM, May 2019, pp. 1–12. doi: 10.1145/3290605.3300682.
- [272] S. Sagheb, A. Bahremand, R. LiKamWa, and B. Lahey, “SWISH: Shifting Weight-based Interfaces for Simulated Hydrodynamics in Mixed-Reality Fluid Vessels,” in *Proceedings of the Thirteenth International Conference on Tangible, Embedded, and Embodied Interaction*, Tempe Arizona USA: ACM, Mar. 2019, pp. 213–221. doi: 10.1145/3294109.3300988.
- [273] C.-H. Cheng *et al.*, “GravityCup: a liquid-based haptics for simulating dynamic weight in virtual reality,” in *Proceedings of the 24th ACM Symposium on Virtual Reality Software and Technology*, Tokyo Japan: ACM, Nov. 2018, pp. 1–2. doi: 10.1145/3281505.3281569.
- [274] H.-R. Tsai and J. Rekimoto, “ElasticVR: Providing Multi-level Active and Passive Force Feedback in Virtual Reality Using Elasticity,” in *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, Montreal QC Canada: ACM, Apr. 2018, pp. 1–4. doi: 10.1145/3170427.3186540.
- [275] H.-R. Tsai and B.-Y. Chen, “ElastImpact: 2.5D Multilevel Instant Impact Using Elasticity on Head-Mounted Displays,” in *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, New Orleans LA USA: ACM, Oct. 2019, pp. 429–437. doi: 10.1145/3332165.3347931.
- [276] H.-R. Tsai, C.-W. Hung, T.-C. Wu, and B.-Y. Chen, “ElastOscillation: 3D Multilevel Force Feedback for Damped Oscillation on VR Controllers,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, Honolulu HI USA: ACM, Apr. 2020, pp. 1–12. doi: 10.1145/3313831.3376408.
- [277] M. Achibet, A. Girard, A. Talvas, M. Marchal, and A. Lecuyer, “Elastic-Arm: Human-scale passive haptic feedback for augmenting interaction and perception in virtual environments,” in *2015 IEEE Virtual Reality (VR)*, Arles, Camargue, Provence, France: IEEE, Mar. 2015, pp. 63–68. doi: 10.1109/VR.2015.7223325.
- [278] C. Fang, Y. Zhang, M. Dworman, and C. Harrison, “Wireality: Enabling Complex Tangible Geometries in Virtual Reality with Worn Multi-String Haptics,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, Honolulu HI USA: ACM, Apr. 2020, pp. 1–10. doi: 10.1145/3313831.3376470.
- [279] M. Canesi, M. Xiloyannis, A. Ajoudani, A. Bicchi, and L. Masia, “Modular one-to-many clutchable actuator for a soft elbow exosuit,” in *2017 International Conference on Rehabilitation Robotics (ICORR)*, London: IEEE, Jul. 2017, pp. 1679–1685. doi: 10.1109/ICORR.2017.8009489.

- [280] J. Realmuto and T. Sanger, “A robotic forearm orthosis using soft fabric-based helical actuators,” in *2019 2nd IEEE International Conference on Soft Robotics (RoboSoft)*, Seoul, Korea (South): IEEE, Apr. 2019, pp. 591–596. doi: 10.1109/ROBOSOFT.2019.8722759.
- [281] G. Barnaby and A. Roudaut, “Mantis: A Scalable, Lightweight and Accessible Architecture to Build Multifunctional Force Feedback Systems,” in *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, New Orleans LA USA: ACM, Oct. 2019, pp. 937–948. doi: 10.1145/3332165.3347909.
- [282] K. E. MacLean, O. S. Schneider, and H. Seifi, “Multisensory haptic interactions: understanding the sense and designing for it,” in *The Handbook of Multimodal-Multisensor Interfaces: Foundations, User Modeling, and Common Modality Combinations - Volume 1*, S. Oviatt, B. Schuller, P. R. Cohen, D. Sonntag, G. Potamianos, and A. Krüger, Eds., ACM, 2017, pp. 97–142. doi: 10.1145/3015783.3015788.
- [283] H. Benko, C. Holz, M. Sinclair, and E. Ofek, “NormalTouch and TextureTouch: High-fidelity 3D Haptic Shape Rendering on Handheld Virtual Reality Controllers,” in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, Tokyo Japan: ACM, Oct. 2016, pp. 717–728. doi: 10.1145/2984511.2984526.
- [284] I. Choi, H. Culbertson, M. R. Miller, A. Olwal, and S. Follmer, “Gravity: A Wearable Haptic Interface for Simulating Weight and Grasping in Virtual Reality,” in *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, Québec City QC Canada: ACM, Oct. 2017, pp. 119–130. doi: 10.1145/3126594.3126599.
- [285] I. Choi, E. Ofek, H. Benko, M. Sinclair, and C. Holz, “CLAW: A Multifunctional Handheld Haptic Controller for Grasping, Touching, and Triggering in Virtual Reality,” in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, Montreal QC Canada: ACM, Apr. 2018, pp. 1–13. doi: 10.1145/3173574.3174228.
- [286] E. Whitmire, H. Benko, C. Holz, E. Ofek, and M. Sinclair, “Haptic Revolver: Touch, Shear, Texture, and Shape Rendering on a Reconfigurable Virtual Reality Controller,” in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, Montreal QC Canada: ACM, Apr. 2018, pp. 1–12. doi: 10.1145/3173574.3173660.
- [287] J. W. Romanishin, K. Gilpin, and D. Rus, “M-blocks: Momentum-driven, magnetic modular robots,” in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Tokyo: IEEE, Nov. 2013, pp. 4288–4295. doi: 10.1109/iros.2013.6696971.
- [288] M. Vicentini, S. Galvan, D. Botturi, and P. Fiorini, “Evaluation of force and torque magnitude discrimination thresholds on the human hand-arm system,” *ACM Trans. Appl. Percept.*, vol. 8, no. 1, pp. 1–16, Oct. 2010, doi: 10.1145/1857893.1857894.
- [289] “Load Cells,” Loadstar Sensors. Accessed: Oct. 17, 2022. [Online]. Available: <https://www.loadstarsensors.com/products/sensors/load-cells.html>
- [290] R. P. Khurshid, N. T. Fitter, E. A. Fedalei, and K. J. Kuchenbecker, “Effects of Grip-Force, Contact, and Acceleration Feedback on a Teleoperated Pick-and-Place Task,” *IEEE Trans. Haptics*, vol. 10, no. 1, pp. 40–53, Jan. 2017, doi: 10.1109/toh.2016.2573301.
- [291] “Micro Load Cell (0-780g) - CZL616C - 3132’0 at Phidgets.” Accessed: Oct. 17, 2022. [Online]. Available: <https://www.phidgets.com/?tier=3&catid=9&pcid=7&prodid=223>

- [292] G. C. Bettelani, A. Moscatelli, and M. Bianchi, “Contact with Sliding over a Rotating Ridged Surface: the Turntable Illusion,” in *2019 IEEE World Haptics Conference (WHC)*, Tokyo, Japan: IEEE, Jul. 2019, pp. 562–567. doi: 10.1109/WHC.2019.8816119.
- [293] A. Moscatelli *et al.*, “Touch as an auxiliary proprioceptive cue for movement control,” *Sci. Adv.*, vol. 5, no. 6, p. 11, 2019, doi: 10.1126/sciadv.aaw3121.
- [294] “Load Cell - 10kg, Straight Bar (TAL220).” Accessed: Oct. 17, 2022. [Online]. Available: <https://robosavvy.co.uk/load-cell-10kg-straight-bar-tal220.html>
- [295] “Load Cell - 5kg, Straight Bar (TAL220B) - SEN-14729 - SparkFun Electronics.” Accessed: Oct. 17, 2022. [Online]. Available: <https://www.sparkfun.com/products/14729>
- [296] E. Gonzalez, M. Rittikaidachar, and Y. Wang, “Admittance Control of a Pin-Based Shape Display for Haptic Feedback,” p. 6.
- [297] “Micro Load Cell (0-5kg) - CZL635 - 3133`0 at Phidgets.” Accessed: Oct. 17, 2022. [Online]. Available: <https://www.phidgets.com/?tier=3&catid=9&pcid=7&prodid=224>
- [298] *Musical haptics*. in Springer Series on Touch and Haptic Systems. New York, NY: Springer Berlin Heidelberg, 2018. doi: 10.1007/978-3-319-58316-7.
- [299] “SMT Overload Protected S-Type Load Cell,” Durham Instruments. Accessed: Oct. 17, 2022. [Online]. Available: <https://disensors.com/product/smt-s-type-overload-protected-load-cell/>
- [300] Y. Liu, R. I. Davidson, P. M. Taylor, J. D. Ngu, and J. M. C. Zarraga, “Single cell magnetorheological fluid based tactile display,” *Displays*, vol. 26, no. 1, pp. 29–35, Jan. 2005, doi: 10.1016/j.displa.2004.10.002.
- [301] “Miniature S-Beam Jr. Load Cell LSB200: FSH00097.” Accessed: Oct. 17, 2022. [Online]. Available: <https://www.futek.com/store/legacy-sensors-and-instruments/miniature-s-beam-LSB200/FSH00097>
- [302] S. Park *et al.*, “User-Driven Functional Movement Training with a Wearable Hand Robot after Stroke,” *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 28, no. 10, pp. 2265–2275, Nov. 2019, doi: 10.1109/TNSRE.2020.3021691.
- [303] “2519 Series S-beam Static Load Cells.” Accessed: Oct. 17, 2022. [Online]. Available: <http://www.instron.com/en/products/testing-accessories/load-cells/static/2519-series-s-beam>
- [304] S. Mintchev, M. Salerno, A. Cherpillod, S. Scaduto, and J. Paik, “A portable three-degrees-of-freedom force feedback origami robot for human–robot interactions,” *Nat Mach Intell*, vol. 1, no. 12, pp. 584–593, Dec. 2019, doi: 10.1038/s42256-019-0125-1.
- [305] “HSFPAR Series — Pressure Sensors — Sensors — Product Search — Products/Technology — Alps Alpine.” Accessed: Oct. 17, 2022. [Online]. Available: <https://tech.alpsalpine.com/e/products/category/sensor/sub/04/series/hsfpar/>
- [306] H. Kajimoto and L. Jones, “Wearable Tactile Display Based on Thermal Expansion of Nichrome Wire,” in *IEEE Transactions on Haptics*, 2019, pp. 1–1. doi: 10.1109/toh.2019.2912960.
- [307] “FS2050-0000-1500-G - Load Cell, Compression, FS20 Series, 1500 gf, 1 to 4 V Output.” Accessed: Oct. 17, 2022. [Online]. Available: <https://uk.farnell.com/sensor-solutions-te-connectivity/fs2050-0000-1500-g/compression-load-cell-1-5kg-5vdc/dp/2748855>

- [308] K. Ito, S. Okamoto, Y. Yamada, and H. Kajimoto, "Tactile Texture Display with Vibrotactile and Electrostatic Friction Stimuli Mixed at Appropriate Ratio Presents Better Roughness Textures," *ACM Trans. Appl. Percept.*, vol. 16, no. 4, pp. 1–15, Sep. 2019, doi: 10.1145/3340961.
- [309] "FC2231-0000-0050-L — TE Connectivity Load Cell, 22.67kg Range, Compression Measure — RS." Accessed: Oct. 17, 2022. [Online]. Available: https://uk.rs-online.com/web/p/strain-gauges/8937395?cm_mmc=UK-PPC-DS3A-_-google-_-DSA`UK`EN`Automation+%26+Control+Gear`Index`-`Strain+Gauges%7C+Products`-`DYNAMIC+SEARCH+ADS&matchtype=&dsa-1592138815139&ccq`src=google`ads&ccq`cmp=15904458419&ccq`term=&ccq`plac=&ccq`net=g&ccq`plt=gp&gclsrc=ds&gclsrc=ds
- [310] "LMA-A Small-sized Compression Load Cell — KYOWA." Accessed: Oct. 17, 2022. [Online]. Available: <https://www.kyowa-ei.com/eng/product/category/sensors/lma-a/index.html>
- [311] T. Mouri, H. Kawasaki, and K. Koketsu, "Compact torque sensor for a robot hand," in *2011 8th International Conference on Ubiquitous Robots and Ambient Intelligence (URAI)*, Incheon: IEEE, Nov. 2011, pp. 439–441. doi: 10.1109/URAI.2011.6145859.
- [312] "MLP Series mini low profile Load Cell." Accessed: Oct. 17, 2022. [Online]. Available: <https://www.transducertechniques.com/mlp-load-cell.aspx>
- [313] Y. Liu, S. Lu, and H. Culbertson, "Texture Classification by Audio-Tactile Crossmodal Congruence," in *2022 IEEE Haptics Symposium (HAPTICS)*, Santa Barbara, CA, USA: IEEE, Mar. 2022, pp. 1–7. doi: 10.1109/HAPTICS52432.2022.9765614.
- [314] G. Tholey, A. Pillariseti, W. Green, and J. P. Desai, "Design, Development, and Testing of an Automated Laparoscopic Grasper with 3-D Force Measurement Capability," in *Medical Simulation*, vol. 3078, S. Cotin and D. Metaxas, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 38–48. Accessed: Dec. 06, 2018. [Online]. Available: http://link.springer.com/10.1007/978-3-540-25968-8_5
- [315] "Miniature Threaded Load Button LLB210 : FSH02708." Accessed: Oct. 17, 2022. [Online]. Available: <https://www.futek.com/store/load-cells/load-button/miniature-threaded-load-button-LLB210/FSH02708>
- [316] D. Chiaradia, M. Xiloyannis, C. W. Antuvan, A. Frisoli, and L. Masia, "Design and embedded control of a soft elbow exosuit," in *2018 IEEE International Conference on Soft Robotics (RoboSoft)*, Livorno: IEEE, Apr. 2018, pp. 565–571. doi: 10.1109/robosoft.2018.8405386.
- [317] "Lutron FORCE GAUGE FG-5000A-232 Suppliers, Service / Repair Centres, Distributors, Dealers, Catalogs and Manuals in Chennai." Accessed: Oct. 17, 2022. [Online]. Available: <http://www.multilab.in/lutron/force-gauge/fg5000a-232-forcegauge-lutron.html>
- [318] J. Forsslund, M. Yip, and E.-L. Sallnäs, "WoodenHaptics: A Starting Kit for Crafting Force-Reflecting Spatial Haptic Devices," in *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction - TEI '14*, Stanford, California, USA: ACM Press, 2015, pp. 133–140. doi: 10.1145/2677199.2680595.
- [319] "Shimpo FG-3000 Digital Force Gauge." Accessed: Oct. 17, 2022. [Online]. Available: <http://shimpoinstruments.com/product/FG-3000>
- [320] "Sauter FK 100 Force Gauge 1000Hz, Range: 100N, Resolution: 0.05 N — RS." Accessed: Oct. 17, 2022. [Online]. Available: <https://ie.rs-online.com/web/p/force-gauges/0559763>

- [321] T. Roumen, J. Shigeyama, J. C. R. Rudolph, F. Grzelka, and P. Baudisch, “SpringFit: Joints and Mounts that Fabricate on Any Laser Cutter,” in *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, New Orleans LA USA: ACM, Oct. 2019, p. 13. doi: 10/gr9sm5.
- [322] “475044: High Capacity Force Gauge — Extech Instruments.” Accessed: Oct. 17, 2022. [Online]. Available: <https://www.extech.com/products/475044>
- [323] A. F. Siu, M. Sinclair, R. Kovacs, E. Ofek, C. Holz, and E. Cutrell, “Virtual Reality Without Vision: A Haptic and Auditory White Cane to Navigate Complex Virtual Worlds,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, in CHI '20. Honolulu, HI, USA: Association for Computing Machinery, Apr. 2020, pp. 1–13. doi: 10.1145/3313831.3376353.
- [324] “ZTS Digital Force Gauge - Imada Inc.” Accessed: Oct. 17, 2022. [Online]. Available: <https://imada.com/products/zts-digital-force-gauge/>
- [325] K. Yasu, “Magnetact: Magnetic-sheet-based Haptic Interfaces for Touch Devices,” *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, p. 240:1-240:8, 2019, doi: 10.1145/3290605.3300470.
- [326] C.-Y. Tsai *et al.*, “AirRacket: Perceptual Design of Ungrounded, Directional Force Feedback to Improve Virtual Racket Sports Experiences,” in *CHI Conference on Human Factors in Computing Systems*, in CHI '22. New York, NY, USA: Association for Computing Machinery, Apr. 2022, pp. 1–15. doi: 10.1145/3491102.3502034.
- [327] “DS2 Digital Force Gauge - Imada Inc. — Get Best Price Here.” Accessed: Oct. 17, 2022. [Online]. Available: <https://imada.com/products/ds2-digital-force-gauge/>
- [328] “Digital Force Gauge.” Accessed: Oct. 17, 2022. [Online]. Available: <https://www.monotaro.sg/p/37993724.html>
- [329] M. J. Kim, N. Ryu, W. Chang, M. Pahud, M. Sinclair, and A. Bianchi, “SpinOcchio: Understanding Haptic-Visual Congruency of Skin-Slip in VR with a Dynamic Grip Controller,” in *CHI Conference on Human Factors in Computing Systems*, in CHI '22. New York, NY, USA: Association for Computing Machinery, Apr. 2022, pp. 1–14. doi: 10.1145/3491102.3517724.
- [330] M. Kono, T. Miyaki, and J. Rekimoto, “In-pulse: inducing fear and pain in virtual experiences,” in *Proceedings of the 24th ACM Symposium on Virtual Reality Software and Technology - VRST '18*, Tokyo Japan: ACM, 2018, pp. 1–5. doi: 10.1145/3281505.3281506.
- [331] D. Groeger, M. Feick, A. Withana, and J. Steimle, “Tactlets: Adding Tactile Feedback to 3D Objects Using Custom Printed Controls,” p. 14. doi: 10.1145/3332165.3347937.
- [332] “FSE103 - Tension/compression load cell by Variense — DirectIndustry.” Accessed: Feb. 07, 2023. [Online]. Available: <https://www.directindustry.com/prod/variense/product-181951-1794141.html>
- [333] “株式会社レプトリノ,” 株式会社レプトリノ. Accessed: Feb. 07, 2023. [Online]. Available: <https://www.leptrino.co.jp/product/6axis-force-sensor>
- [334] “3 axis load cell 0-100N Multi-axis force sensor.” Accessed: Feb. 07, 2023. [Online]. Available: <http://www.forsentek.com/prodetail/279.html>
- [335] T. Yoshida, J. Ogawa, K. Y. Choi, S. Bushnaq, K. Nakagaki, and H. Ishii, “inDepth: Force-based Interaction with Objects beyond A Physical Barrier,” in *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction*, in TEI '21. New York, NY, USA: Association for Computing Machinery, Feb. 2021, pp. 1–6. doi: 10.1145/3430524.3442447.

- [336] "USL06-H5 Force Sensor — Products — Tec Gihan Co.,Ltd." Accessed: Feb. 07, 2023. [Online]. Available: <https://www.tecgihan.co.jp/en/products/forcesensor-amplifier-interface/forcesensor-3-axis/usl06-h5'series/>
- [337] T. Okada, S. Okamoto, and Y. Yamada, "Passive haptics: greater impact presented by pulsive damping brake of DC motor and physical indices for perceived impact," *Virtual Reality*, vol. 25, no. 1, pp. 233–245, Mar. 2021, doi: 10.1007/s10055-020-00452-8.
- [338] T. Agcayazi, M. Foster, H. Kausche, M. Gordon, and A. Bozkurt, "Multi-axis stress sensor characterisation and testing platform," *HardwareX*, vol. 5, p. e00048, Apr. 2019, doi: 10.1016/j.ohx.2018.e00048.
- [339] "ATI Industrial Automation: F/T Sensor gamma." Accessed: Feb. 07, 2023. [Online]. Available: https://www.ati-ia.com/products/ft/ft_models.aspx?id=gamma
- [340] "ATI Industrial Automation: F/T Sensor nano17." Accessed: Feb. 07, 2023. [Online]. Available: https://www.ati-ia.com/products/ft/ft_models.aspx?id=nano17
- [341] C. M. Reed *et al.*, "A Phonemic-Based Tactile Display for Speech Communication," *IEEE Trans. Haptics*, vol. 12, no. 1, pp. 1–1, 2018, doi: 10.1109/toh.2018.2861010.
- [342] Abhishek Gupta, Marcia K. O'Malley, Volkan Patoglu, and Charles Burgar, "Design, Control and Performance of RiceWrist: A Force Feedback Wrist Exoskeleton for Rehabilitation and Training," *The International Journal of Robotics Research*, vol. 27, no. 2, pp. 233–251, Feb. 2008, doi: 10.1177/0278364907084261.
- [343] J. Jiao, Y. Zhang, D. Wang, X. Guo, and X. Sun, "HapTex: A Database of Fabric Textures for Surface Tactile Display," in *2019 IEEE World Haptics Conference (WHC)*, Tokyo, Japan: IEEE, Jul. 2019, pp. 331–336. doi: 10.1109/WHC.2019.8816167.
- [344] "ATI Industrial Automation: F/T Sensor mini45." Accessed: Feb. 07, 2023. [Online]. Available: https://www.ati-ia.com/products/ft/ft_models.aspx?id=mini45
- [345] M. Laghi, M. G. Catalano, G. Grioli, and A. Bicchi, "A wearable wrist haptic display for motion tracking and force feedback in the operational space," *Wearable Technol.*, vol. 2, p. e5, 2021, doi: 10.1017/wtc.2021.4.
- [346] "HE6x6," AMTI. Accessed: Feb. 07, 2023. [Online]. Available: <https://www.amti.biz/product/he6x6/>
- [347] "Capacitive 6-axis force sensor (500N / DynPick) / WACOH-TECH." Accessed: Feb. 07, 2023. [Online]. Available: <https://wacoh-tech.com/en/products/dynpick/200n'500n'rcdb.html>
- [348] D. Schrand, "Cross-Talk Compensation Using Matrix Methods," *Sensors & Transducers*, vol. 79, no. 5, pp. 1157–1163, May 2007.
- [349] "Matrix compensation for multi-component-transducers," HBM. Accessed: May 11, 2023. [Online]. Available: <https://www.hbm.com/en/4507/matrix-compensation/>

11 Appendix – Commercial Load Cells Comparison Table

Model	Load Cell Type	Brand	Cost	Details	Publications that use this sensor
Loadstar: iLoad Analog	Beam Load cell	Loadstar		[289]	[290]
Micro Load Cell (0-780g) _CZL616C	Beam Load cell	Phidgets	\$6.00	[291]	[292], [293]
Load Cell_ 10kg,Straight_Bar_(TAL220)	Beam Load cell	Generic	\$9.50	[294]	[43], [69], [105], [109]
Degraw 5kg Load Cell and HX711 Combo Pack Kit	Beam Load cell	Generic	\$10.00	[295]	[68], [68], [296]
Micro Load Cell (0-5kg) CZL635	Beam Load cell	Phidgets	\$10.20	[297]	[45], [148], [298]
SMT-50 Overload Protected S-Type Load Cell	S-Type	Interface	\$450.00	[299]	[300]
Miniature S-Beam Jr. Load Cell LSB200 : FSH00097	S-Type	Futek	\$650.00	[301]	[113], [302], [114]
Instron 2519 Series S- beam Static Load Cell	S-Type	Instrom	\$2,950.00	[303]	[68], [304], [206]
HSFPAR Series - Micro Force Sensor	Buttom Load Cell	Alps	\$8.00	[305]	[306]
FS20500000-1500-G : MEAS Force Sensor Elements	Buttom Load Cell	MEAS—TE	\$100.00	[307]	[308]
FC2231- 0000-0100-L LOAD CELL (MSP6951- ND)	Buttom Load Cell	MEAS—TE	\$108.00	[309]	[180]
LMA-A-50N Small-sized Compression Load Cell	Buttom Load Cell		\$300.00	[310]	[311]
MLP-10 Series mini low profile Load Cell	Buttom Load Cell		\$350.00	[312]	[313], [314]
Miniature Threaded Load Cell Button LLB210	Buttom Load Cell	Futek	\$850.00	[315]	[316]
ElectroSamson - Brecknell	Force Gauge	Bracknell	\$60.00	[104]	[103]
Lutron FORCE GAUGE FG- 5000A-232	Force Gauge	Lutron	\$180.00	[317]	[318]
Shimpo FG- 3000 Digital Force Gauge	Force Gauge	Shimpo	\$290.00	[319]	[33]
Sauter FK 100 Force Gauge	Force Gauge	Sauter	\$400.00	[320]	[321]
475044: High Capacity Force Gauge	Force Gauge	Extech	\$600.00	[322]	[323]
IMADA ZTS- 2N Digital force gauge (2N/0,01N)	Force Gauge	Imada	\$650.00	[324]	[325]
IMADA ZTS- 20N Digital force gauge (20N/0,01N)	Force Gauge	Imada	\$650.00	[324]	[326]
IMADA ZTS- 50N Digital force gauge (50N/0,01N)	Force Gauge	Imada	\$650.00	[324]	[51], [97], [101]
DS2 Digital Force Gauge	Force Gauge	Imada	\$750.00	[327]	[308]
FGJN-5 Digital Force Gauge	Force Gauge	Shimpo	\$800.00	[328]	[329]
Undefined	Force Gauge				[159], [330], [331]
VARIENSE FSE103	Multi-Axis Force Sensor	Variense	\$800.00	[332]	[49], [50]
6-axis force sensor (Leptrino FFS055YA501U6)	Multi-Axis Force Sensor	Leptrino	\$1,000.00	[333]	[48]
FNZ100N Forsentek 3 axis load cell 0-100N Multi-axis force sensor	Multi-Axis Force Sensor	Forsentek	\$1200.00	[334]	[115], [335]
USL06-H5 Force Sensor	Multi-Axis Force Sensor	Tec Gihan	\$1,800.00	[336]	[337]
ATI Gamma SI- 65-5	Multi-Axis Force Sensor	ATI	\$5,200.00		[338], [339]
F/T Sensor ATI Nano17	Multi-Axis Force Sensor	ATI	\$4,800.00	[340]	[37], [118], [304], [341], [342], [343]
F/T Sensor mini45	Multi-Axis Force Sensor	ATI	\$4,600.00	[344]	[61], [62], [345]
AMTI HE6X6-105 6-DOF grounded force/torque plate	Multi-Axis Force Sensor	AMTI	-	[346]	[57], [58]
Capacitive 6-axis force sensor (500N / DynPick) / WACOH-TECH	Multi-Axis Force Sensor	Wacoh	\$3500.00	[347]	[56]

12 Appendix – VR Recoil Modelling

To determine the average force of impact (F_{avg}) as a piston actuates a weight, calculate the force exerted by the pneumatic pressure on the piston. This force (F) is computed by multiplying the pressure (P) by the piston's cross-sectional area (A), given by the formula $A = \pi(D/2)^2$, where D represents the piston diameter. This initial calculation establishes the foundation for understanding the dynamics of the piston's motion.

The total mass (M_{total}) of the system includes the mass of the piston head (m_{piston}) and the mass of the attached weight (m_{weight}). This total mass, calculated as $M_{total} = m_{piston} + m_{weight}$, determines how the system responds to the applied force.

With force and total mass known, the acceleration (a) of the piston and weight system is determined using Newton's second law: $a = F / M_{total}$. This acceleration predicts how quickly the piston moves.

The time (t) for the piston to travel a fixed distance (d_{travel}) is found by rearranging the equation of motion. Starting from rest, the time required to reach the end of the travel path is $t = \sqrt{2 \times d_{travel} / a}$. This calculation determines how long the piston moves before impacting the enclosure.

Using time and acceleration, calculate the final speed (v) of the piston at the end of travel with $v = a \times t$, indicating how fast the piston moves just before impact. This speed is critical for understanding the change in momentum.

The change in momentum is calculated next. Since the piston starts from rest, the change in velocity (Δv) equals the final speed (v). Therefore, the change in momentum is $M_{total} \times v$.

Finally, to determine the average force of impact (F_{avg}), apply the impulse-momentum theorem, which connects the change in momentum to the force exerted over the impact time. The average force is calculated using $F_{avg} = M_{total} \times \Delta v / \Delta t$, where Δt is the brief duration of the collision, estimated at 1 millisecond.

The script below calculates the impact force on collision of the piston at the end of travel at various pressures:

```
import math
import matplotlib.pyplot as plt

# Inputs
d_travel = 0.06 # piston travel distance in m (6cm)
m_piston = 0.05 # piston mass in kg (50g)
m_weight = 0.24 # weight mass in kg (240g)
D = 0.012 # piston diameter in m (12mm)
A = math.pi * (D/2)**2 # piston area in m^2

# Range of pressures to calculate for
pressure_range = [x * 0.2 for x in range(10, 41)]

# Lists to hold pressure and average force of impact values
pressure_values = []
force_values = []

# Calculate force impact for each pressure in range
for P in pressure_range:
    # Step 1: Calculate force exerted by pneumatic pressure
    F = P * 10**5 * A # Convert pressure from bar to Pa

    # Step 2: Calculate total mass
    M_total = m_piston + m_weight

    # Step 3: Calculate acceleration
    a = F / M_total

    # Step 4: Find time
    t = math.sqrt(2 * d_travel / a)

    # Step 5: Calculate final speed
    v = a * t

    # Step 6: Calculate change in momentum
    delta_v = v
    change_in_momentum = M_total * delta_v

    # Step 7: Calculate average force of impact
    delta_t = 0.001 # Assume duration of collision is 1ms
    F_avg = change_in_momentum / delta_t

    # Print the values
    print("{:.1f}\t\t{:.2f}".format(P, F_avg))

    # Append pressure and force values to lists
    pressure_values.append(P)
```

```
    force_values.append(F_avg)

# Create line plot of force vs. pressure
plt.plot(pressure_values, force_values)
plt.xlabel("Pressure (bar)")
plt.ylabel("Average Force of Impact (N)")
plt.title("Force of Impact vs. Pressure")
plt.show()
```

13 Appendix – HapticWhirl Modelling

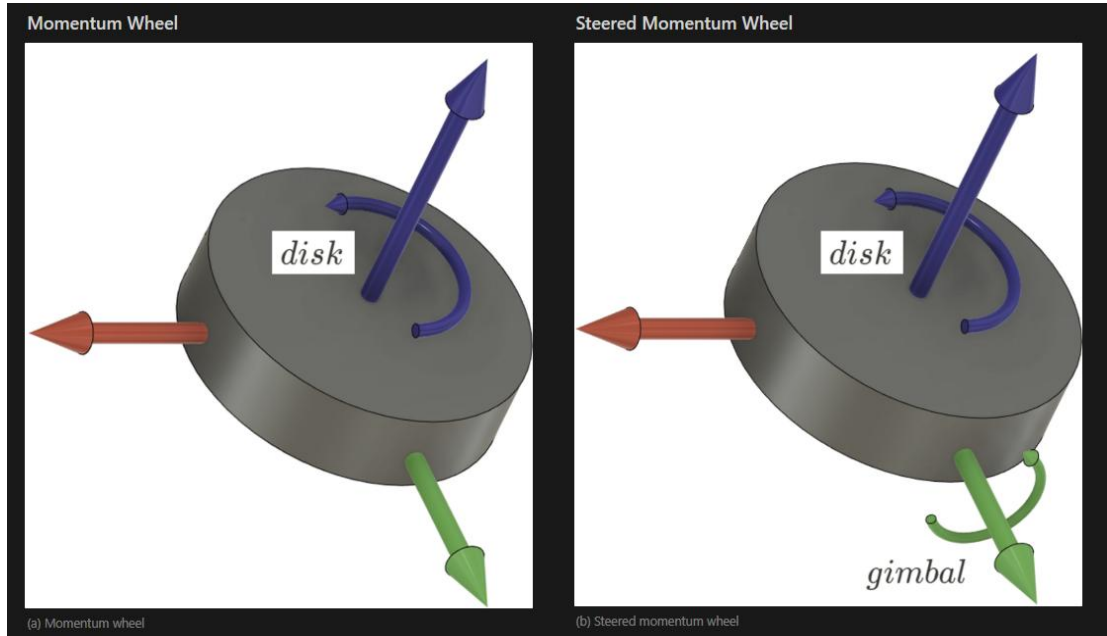


Figure 120. Left 3D model of a momentum wheel, illustrating the angular velocity. Right 3D model of a steered momentum wheel, showcasing the rotation vectors

The HapticWhirl device generates torque through the actuation of both a flywheel and a gimbal.

Momentum Wheel

The momentum of the flywheel, which keeps it rotating, always acts perpendicular to its rotation plane. The torque exerted by the flywheel can be calculated using the equation:

$$\vec{\tau} = \frac{d\vec{L}}{dt} = \mathbf{I} \frac{d}{dt} \vec{\omega}^{\text{disk}} = \mathbf{I} \dot{\vec{\omega}}^{\text{disk}}$$

This formula shows that the torque is a result of multiplying the flywheel's moment of inertia (resistance to changes in its rotation) by the rate at which its angular velocity changes.

Steered Momentum Wheel

The steered momentum wheel generates additional momentum when both the disk and the gimbal are moving. This is described by the equation:

$$\vec{\tau} = \frac{d\vec{L}}{dt} = \vec{\omega}^{\text{gimbal}} \times \mathbf{I} \vec{\omega}^{\text{disk}}$$

This can be combined in the following equation, which considers the effects of the disk's angular acceleration and the interaction (cross product) of the gimbal's angular velocity with the disk's angular momentum, detailing the overall momentum in the system's frame of reference (A).

$$\vec{\mathbf{M}}_A^{gyro} = \mathbf{I} \cdot \dot{\vec{\omega}}_A^{Disk} + \vec{\omega}_A^{Gimbal} \times (\mathbf{I} \cdot \vec{\omega}_A^{Disk})$$

Variables definition

The controller can be divided into different objects that compose it. The corresponding coordinate spaces and position variables are shown in the table below. The origin of all spaces is located at the centre of mass of the flywheel. Since the disk is mounted on the inner gimbal, the coordinate spaces A and B are always in the same orientation.

Object	Space	Angle
disk	A	ρ
inner (pitch)	B	θ
outer (yaw)	C	ψ
handle	D	




Figure 121. Illustration of the HapticWhirl kinematic model showing the A, B, C, and D coordinate frames with axes labeled in red, green, and blue, representing the different components of the system including the disk and gimbals

The variables used in this document are denoted with the following subscript and superscript notation:

$$\text{variable}_{\text{space}}^{\text{object}}$$

For example, in the following equation, we refer to the angular velocity of gimbal2 (the outer gimbal) in the A frame:

$$\vec{\omega}_A^{outer}$$

Angular Velocity Disk ($\vec{\omega}_A^{Disk}$)

This is the sum of the angular velocity of the disk and the contributions to it from the two gimbals in frame A :

$$\vec{\omega}_A^{Disk} = \vec{\omega}_A^{disk} + \vec{\omega}_A^{inner} + \vec{\omega}_A^{outer}$$

Rotation matrices must be applied to the gimbal; the inner gimbal is rotated from frame B to A , and the outer gimbal from frame C to A :

$$\vec{\omega}_A^{Disk} = \begin{pmatrix} 0 \\ 0 \\ \dot{\rho} \end{pmatrix} + (R_{B \rightarrow A} \cdot \vec{\omega}_B^{inner}) + (R_{C \rightarrow A} \cdot \vec{\omega}_C^{outer})$$

The disk and inner gimbal share the same coordinate frame, hence $A = B$. Therefore, the rotation matrix $R_{B \rightarrow A}$ is an identity matrix. The rotation matrix $R_{C \rightarrow A}$ rotates over the Y -axis.




Figure 122. Rotation about the Y -axis illustrating the rotation matrix $(R_{C \rightarrow A})$.

The rotation matrix $R_{A \rightarrow B}$ is a rotation over the Y -axis, but since the rotation in this case is $R_{B \rightarrow A}$, it must be inverted.

$$\vec{\omega}_A^{Disk} = \begin{pmatrix} 0 \\ 0 \\ \dot{\rho} \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 0 \\ \dot{\theta} \\ 0 \end{pmatrix} + \begin{pmatrix} \cos \theta & 0 & -\sin \theta \\ 0 & 1 & 0 \\ \sin \theta & 0 & \cos \theta \end{pmatrix} \begin{pmatrix} 0 \\ 0 \\ \dot{\psi} \end{pmatrix}$$

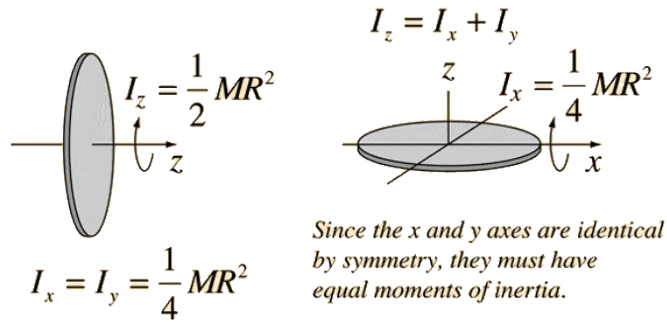
$$\vec{\omega}_A^{Disk} = \begin{pmatrix} 0 + 0 + -\dot{\psi} \sin \theta \\ 0 + \dot{\theta} + 0 \\ \dot{\rho} + 0 + \dot{\psi} \cos \theta \end{pmatrix}$$

$$\vec{\omega}_A^{Disk} = \begin{pmatrix} -\dot{\psi} \sin \theta \\ \dot{\theta} \\ \dot{\rho} + \dot{\psi} \cos \theta \end{pmatrix}$$

The inertia torque matrix of the disk, denoted by \mathbf{I} , is defined as the inertial tensor representing the moment of inertia about the different axes.

$$\mathbf{I} = \begin{pmatrix} I_{xx} & 0 & 0 \\ 0 & I_{yy} & 0 \\ 0 & 0 & I_{zz} \end{pmatrix}$$

For a disk, this can be simplified to a diagonal matrix. Replace I_{xx} , I_{yy} , and I_{zz} with the disk's moment of inertia. Hence, \mathbf{I} becomes:



Hence, \mathbf{I} becomes:

$$\mathbf{I} = \begin{pmatrix} \frac{1}{4} MR^2 & 0 & 0 \\ 0 & \frac{1}{4} MR^2 & 0 \\ 0 & 0 & \frac{1}{2} MR^2 \end{pmatrix}$$

Angular Velocity Gimbal ($\vec{\omega}_A^{Gimbal}$)

The other component in equation (1) corresponds to the sum of the angular velocities of both axes in the gimbal with respect to frame A .

$$\vec{\omega}_A^{Gimbal} = R_{B \rightarrow A} \cdot \vec{\omega}_B^{inner} + R_{C \rightarrow A} \cdot \vec{\omega}_C^{outer}$$

$$\vec{\omega}_A^{Gimbal} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} 0 \\ \dot{\theta} \\ 0 \end{pmatrix} + \begin{pmatrix} \cos \theta & 0 & -\sin \theta \\ 0 & 1 & 0 \\ \sin \theta & 0 & \cos \theta \end{pmatrix} \cdot \begin{pmatrix} 0 \\ 0 \\ \dot{\psi} \end{pmatrix}$$

$$\vec{\omega}_A^{Gimbal} = \begin{pmatrix} -\dot{\psi} \sin \theta \\ \dot{\theta} \\ \dot{\psi} \cos \theta \end{pmatrix}$$

Angular Acceleration Disk($\dot{\omega}_A^{Disk}$)

Derivative of applying the chain rule:

$$\dot{\omega}_A^{Disk} = \begin{pmatrix} -\ddot{\psi} \sin \theta - \dot{\psi} \dot{\theta} \cos \theta \\ \ddot{\theta} \\ \ddot{\rho} + \ddot{\psi} \cos \theta - \dot{\psi} \dot{\theta} \sin \theta \end{pmatrix}$$

Equation Expanded

The equation for the gyroscopic torque can be expanded by considering the inertial properties of the disk and the angular velocities involved. We apply the inertia tensor to the angular acceleration vector, then add the cross product of the angular velocity vector and the angular momentum vector. This approach considers the non-commutative nature of the cross product and matrix multiplication in rotational dynamics. The expanded form of equation (1) is as follows:

$$\vec{M}_A^{gyro} = \begin{pmatrix} \frac{1}{4}MR^2 & 0 & 0 \\ 0 & \frac{1}{4}MR^2 & 0 \\ 0 & 0 & \frac{1}{2}MR^2 \end{pmatrix} \cdot \begin{pmatrix} -\ddot{\psi} \sin \theta - \dot{\psi} \dot{\theta} \cos \theta \\ \ddot{\theta} \\ \ddot{\rho} + \ddot{\psi} \cos \theta - \dot{\psi} \dot{\theta} \sin \theta \end{pmatrix} + \begin{pmatrix} -\dot{\psi} \sin \theta \\ \dot{\theta} \\ \dot{\psi} \cos \theta \end{pmatrix} \times \begin{pmatrix} \frac{1}{4}MR^2 & 0 & 0 \\ 0 & \frac{1}{4}MR^2 & 0 \\ 0 & 0 & \frac{1}{2}MR^2 \end{pmatrix} \cdot \begin{pmatrix} -\dot{\psi} \sin \theta \\ \dot{\theta} \\ \dot{\rho} + \dot{\psi} \cos \theta \end{pmatrix}$$

We can combine the members of the equation by grouping the different equations to solve in three steps.

$$\vec{M}_A^{gyro} = \begin{pmatrix} \frac{1}{4}MR^2 & 0 & 0 \\ 0 & \frac{1}{4}MR^2 & 0 \\ 0 & 0 & \frac{1}{2}MR^2 \end{pmatrix} \cdot \begin{pmatrix} -\ddot{\psi} \sin \theta - \dot{\psi} \dot{\theta} \cos \theta \\ \ddot{\theta} \\ \ddot{\rho} + \ddot{\psi} \cos \theta - \dot{\psi} \dot{\theta} \sin \theta \end{pmatrix} + \begin{pmatrix} -\dot{\psi} \sin \theta \\ \dot{\theta} \\ \dot{\psi} \cos \theta \end{pmatrix} \times \begin{pmatrix} \frac{1}{4}MR^2 & 0 & 0 \\ 0 & \frac{1}{4}MR^2 & 0 \\ 0 & 0 & \frac{1}{2}MR^2 \end{pmatrix} \cdot \begin{pmatrix} -\dot{\psi} \sin \theta \\ \dot{\theta} \\ \dot{\rho} + \dot{\psi} \cos \theta \end{pmatrix}$$

First replace the diagonal values in I with I_x, I_y and I_z , where $I_x = I_y$.

$$\mathbf{I}_x = \mathbf{I}_y = \frac{1}{4}MR^2 \quad \mathbf{I}_z = \frac{1}{2}MR^2$$

Solving group 1:

$$\begin{pmatrix} \mathbf{I}_x (-\ddot{\psi} \sin \theta - \dot{\psi} \dot{\theta} \cos \theta) \\ \mathbf{I}_x \ddot{\theta} \\ \mathbf{I}_z (\ddot{\rho} + \ddot{\psi} \cos \theta - \dot{\psi} \dot{\theta} \sin \theta) \end{pmatrix}$$

Solving group 2 :

$$\begin{pmatrix} \mathbf{I}_x (-\dot{\psi} \sin \theta) \\ \mathbf{I}_x \dot{\theta} \\ \mathbf{I}_z (\dot{\rho} + \dot{\psi} \cos \theta) \end{pmatrix}$$

Solving group 3. First, compute the cross product:

$$\begin{pmatrix} -\dot{\psi} \sin \theta \\ \dot{\theta} \\ \dot{\psi} \cos \theta \end{pmatrix} \times \begin{pmatrix} \mathbf{I}_x (-\dot{\psi} \sin \theta) \\ \mathbf{I}_x \dot{\theta} \\ \mathbf{I}_z (\dot{\rho} + \dot{\psi} \cos \theta) \end{pmatrix}$$

This can be expanded using the determinant of the following matrix:

$$\begin{pmatrix} i & j & k \\ -\dot{\psi} \sin \theta & \dot{\theta} & \dot{\psi} \cos \theta \\ \mathbf{I}_x (-\dot{\psi} \sin \theta) & \mathbf{I}_x \dot{\theta} & \mathbf{I}_z (\dot{\rho} + \dot{\psi} \cos \theta) \end{pmatrix}$$

Resulting in:

$$\begin{pmatrix} \dot{\theta} \mathbf{I}_z (\dot{\rho} + \dot{\psi} \cos \theta) - \dot{\psi} \cos \theta \mathbf{I}_z (\dot{\rho} + \dot{\psi} \cos \theta) \\ (\dot{\psi} \cos \theta)(-\mathbf{I}_x \dot{\psi} \sin \theta) - (\mathbf{I}_z (\dot{\rho} + \dot{\psi} \cos \theta)(-\dot{\psi} \sin \theta)) \\ -\dot{\psi} \sin \theta \mathbf{I}_x \dot{\theta} - \mathbf{I}_x (-\dot{\psi} \sin \theta) \dot{\theta} \end{pmatrix}$$

Grouping the terms:

$$\begin{pmatrix} -\mathbf{I}_x \dot{\theta} \dot{\psi} \cos \theta + \mathbf{I}_z \dot{\theta} \dot{\rho} \\ -\mathbf{I}_x \dot{\psi}^2 \cos \theta \sin \theta + \mathbf{I}_z \dot{\psi}^2 \cos \theta \sin \theta + \mathbf{I}_z \dot{\rho} \dot{\psi} \sin \theta \\ 0 \end{pmatrix}$$

This results on the equation:

$$\bar{\mathbf{M}}_A^{gyro} = \begin{pmatrix} \mathbf{I}_x (-\dot{\psi} \sin \theta - \dot{\psi} \dot{\theta} \cos \theta) \\ \mathbf{I}_x \ddot{\theta} \\ \mathbf{I}_z (\dot{\rho} + \dot{\psi} \cos \theta - \dot{\psi} \dot{\theta} \sin \theta) \end{pmatrix} + \begin{pmatrix} -\mathbf{I}_x \dot{\theta} \dot{\psi} \cos \theta + \mathbf{I}_z \dot{\theta} \dot{\rho} \\ -\mathbf{I}_x \dot{\psi}^2 \cos \theta \sin \theta + \mathbf{I}_z \dot{\psi}^2 \cos \theta \sin \theta + \mathbf{I}_z \dot{\rho} \dot{\psi} \sin \theta \\ 0 \end{pmatrix}$$

Combining terms, the gyroscopic moment becomes:

$$\bar{\mathbf{M}}_A^{gyro} = \begin{pmatrix} \mathbf{I}_x (-\dot{\theta} \dot{\psi} \cos \theta - \dot{\psi} \sin \theta) - \mathbf{I}_x \dot{\theta} \dot{\psi} \cos \theta + \mathbf{I}_z \dot{\theta} \dot{\rho} \\ \mathbf{I}_x \ddot{\theta} - \mathbf{I}_x \dot{\psi}^2 \cos \theta \sin \theta + \mathbf{I}_z \dot{\psi}^2 \cos \theta \sin \theta + \mathbf{I}_z \dot{\rho} \dot{\psi} \sin \theta \\ \mathbf{I}_z (\dot{\rho} + \dot{\psi} \cos \theta - \dot{\theta} \dot{\psi} \sin \theta) \end{pmatrix}$$

To further simplify the equation, we substitute the \mathbf{I} terms across the equation:

$$\text{If } \mathbf{I} = MR^2, \quad \mathbf{I}_x = \frac{1}{4} MR^2 = \frac{1}{4} \mathbf{I}, \quad \mathbf{I}_z = \frac{1}{2} MR^2 = \frac{1}{2} \mathbf{I}$$

$$\bar{\mathbf{M}}_A^{gyro} = \begin{pmatrix} \frac{1}{4} \mathbf{I}(-\dot{\theta}\dot{\psi} \cos \theta - \ddot{\psi} \sin \theta) - \frac{1}{4} \mathbf{I}\dot{\theta}\dot{\psi} \cos \theta + \frac{1}{2} \mathbf{I}\dot{\theta}\dot{\psi} \cos \theta + \frac{1}{2} \mathbf{I}\dot{\theta}\dot{\rho} \\ \frac{1}{4} \mathbf{I}\ddot{\theta} - \frac{1}{4} \mathbf{I}\dot{\psi}^2 \cos \theta \sin \theta + \frac{1}{2} \mathbf{I}\dot{\psi}^2 \cos \theta \sin \theta + \frac{1}{2} \mathbf{I}\dot{\rho}\dot{\psi} \sin \theta \\ \frac{1}{2} \mathbf{I}(\ddot{\rho} + \dot{\psi} \cos \theta - \dot{\theta}\dot{\psi} \sin \theta) \end{pmatrix}$$

After collecting common terms and simplifying the equation. The final equation for the momentum of the flywheel in the VR controller is:

$$\bar{\mathbf{M}}_A^{gyro} = \begin{pmatrix} \frac{1}{4} \mathbf{I}(2\dot{\theta}\dot{\rho} - \ddot{\psi} \sin \theta) \\ \frac{1}{4} \mathbf{I}(\ddot{\theta} + \dot{\psi} \sin \theta(2\dot{\rho} + \dot{\psi} \cos \theta)) \\ \frac{1}{2} \mathbf{I}(\ddot{\rho} + \dot{\psi} \cos \theta - \dot{\theta}\dot{\psi} \sin \theta) \end{pmatrix}$$

Torque applied on the handle

The output for the momentum of the flywheel in the VR controller represents the momentum in frame A . To express this relative to the handle, it needs to translate it to the corresponding system of reference by applying the appropriate rotation matrix:

$$\bar{\mathbf{M}}_D^{gyro} = R_{A \rightarrow D} \cdot \bar{\mathbf{M}}_A^{gyro} = R_{C \rightarrow D} \cdot R_{A \rightarrow C} \cdot \bar{\mathbf{M}}_A^{gyro}$$

The transformation matrix from A to D is composed of two rotation matrices. The first step from A to C involves a rotation about the Y -axis, and the second step from C to D involves a rotation about the Z -axis:




Figure 123. Left Rotation from frame A to frame D through intermediate frame C . Right Rotation matrix for Z -axis transformation from frame C to D

Rotation from frame A to frame C :

$$R_{A \rightarrow C} = \begin{pmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{pmatrix}$$

Rotation from frame C to frame D :

$$R_{C \rightarrow D} = \begin{pmatrix} \cos \psi & -\sin \psi & 0 \\ \sin \psi & \cos \psi & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

By combining both rotation matrices $R_{A \rightarrow D} = R_{A \rightarrow C} \cdot R_{C \rightarrow D}$, we obtain the following rotation matrix:

$$R_{A \rightarrow D} = \begin{pmatrix} \cos \theta \cos \psi & -\sin \psi & \sin \theta \cos \psi \\ \cos \theta \sin \psi & \cos \psi & \sin \theta \sin \psi \\ -\sin \theta & 0 & \cos \theta \end{pmatrix}$$

Finally, the torque applied on the handle is computed by applying the rotation to the result of the momentum equation as follows:

$$\overline{\mathbf{M}}_D^{gyro} = R_{A \rightarrow D} \cdot \overline{\mathbf{M}}_A^{gyro}$$

14 Appendix – Previous Amplifier Designs Used (alternatives to HX711)

For many of the testbeds developed during my PhD, I used a range of amplifiers directly feeding the ADC on the microcontroller. This is an alternative approach when a small number of channels is required, and provides unlimited sampling rate which is constrained by the ADC. The following amplifiers are compared on the following list - AD623, AD8237, AD8420, INA317, INA333, INA125, INA129. The following is a comparison of their main properties to determine the most suitable design.

- 1. Supply Voltage:** The operational power supply voltage ensures compatibility with the available power supply in the given application. Amplifiers such as the INA125 and INA129, with a broad supply voltage range of 2.7 to 36V, offer flexibility across various power supply contexts. However, the INA333 and AD8237, operating in a lower range of 1.8 to 5.5V, are more suitable for low-voltage, battery-powered, or USB-powered applications.
- 2. Gain Range:** The gain range of an amplifier should ideally align with the output of the sensor and the input requirements of the following stage in the signal processing chain, as its name states it determines how much the input signal can be amplified, this is normally controlled via a resistor value. Amplifiers like the INA125 and INA129, offering a gain range of 4-10000 and 1-10000 respectively, provide a wider adjustable gain range compared to others. However, a gain range of 1-1000, as provided by the AD623, INA317, INA333, AD8237, and AD8420, is typically sufficient for load cell applications.
- 3. Noise:** The noise generated by an amplifier can critically influence signal quality, especially in high-precision applications. The INA129, with a noise level of $8\text{nV}/\sqrt{\text{Hz}}$, stands out as the best performer in this category, offering the lowest noise level among the compared amplifiers. However, all the amplifiers have noise levels considered very low.
- 4. Power Consumption:** The INA317 and INA333, with a low power consumption of just $50\mu\text{A}$, provide the most energy-efficient solutions.
- 5. Common-Mode Rejection Ratio (CMRR):** measures an amplifier's ability to minimise the effects of noise and interference. The AD8420, with a CMRR of 130dB, provides the best common-mode noise rejection capability among the considered amplifiers.

6. **Bandwidth:** This refers to the range of signal frequencies that the amplifier can handle without significant attenuation. The INA129, with a bandwidth of 1300kHz at a gain of 1, offers the highest bandwidth at low gain among the amplifiers compared. While this high bandwidth is beneficial for various sensor applications, it is not critical for amplifying a DC signal from a Wheatstone bridge, which is the primary use case in this context.
7. **Slew Rate:** The slew rate, representing the maximum rate of change of output voltage per unit time, can influence the amplifier's fidelity in reproducing rapidly changing signals. The INA129, with a slew rate of $4V/\mu s$, offers the fastest rate of change in output voltage.
8. **Temperature Drift:** Changes in temperature can impact the accuracy and stability of an amplifier. The INA333, with a temperature drift of just $0.1\mu V/^\circ C$, offers the highest stability over a range of operating temperatures, making it a reliable choice for applications requiring precise measurements under varying environmental conditions.
9. **Differential Input Impedance:** A high input impedance minimises the loading effect on the load cell sensor, the amplifier will draw a minimal amount of current from the load cell, which is essential for accurate measurements. The INA125, INA317, and INA333, with a differential input impedance of $100G\Omega \parallel 2pF$, $100G\Omega \parallel 3pF$, and $100G\Omega \parallel 3pF$ respectively, provide a high input impedance, thereby minimising loading effects.
10. **Common Mode Input Impedance:** Like differential input impedance, a high common mode input impedance is desirable to minimise interference with the sensor's output. The INA125 and INA129, both with a common mode input impedance of $100G\Omega \parallel 9pF$, stand out in this category.
11. **Price:** Budget constraints often dictate the final decision in hardware selection. Among the considered amplifiers, the INA317, priced at €2.59, is the most economical choice, while still offering a competitive set of specifications.

The INA333, with its balanced specifications, is a viable choice for load cell applications. It offers a suitable supply voltage and gain range, coupled with low power consumption. Although its noise level is not the lowest, it is sufficient for the proposed application. Its moderate bandwidth, stable performance over a wide temperature range, and high input impedance ensure accurate load cell measurements. Priced reasonably at €5.72, it provides a cost-effective solution for a DAQ system interfacing with a load cell.

Amplifier	Supply Voltage (V)	Gain range	Noise (nV/ $\sqrt{\text{Hz}}$)	Power Consumption (μA)	CMRR	Bandwidth (for different Gain) (kHz)	Slew Rate (V/ μs)	Temperature Drift ($\mu\text{V}/^\circ\text{C}$)	Differential Input Impedance ($\text{G}\Omega \parallel \text{pF}$)	Common Mode Input Impedance ($\text{G}\Omega \parallel \text{pF}$)	Price (EUR)
AD623	2.7 to 12	1-1000	35	550	80-100	G=1: 800; G=10: 100; G=100: 10; G=1000: 2	0.3	2 min; 12 max	2 \parallel 2	2 \parallel 2	3.08
INA125	2.7 to 36	4-10000	38	460	86-100	G=4: 150; G=10: 45; G=100: 4.5; G=500: 0.9	0.2	2 max	100 \parallel 2	100 \parallel 9	6.97
INA129	4.5 to 36	1-10000*	8*	2000	100-113	G=1: 1300; G=10: 700; G=100: 200; G=1000: 20	4*	Not specified	10 \parallel 2	100 \parallel 9	9.69
INA317	1.8 to 5.5	1-1000	50	50*	85-100	G=1: 150; G=10: 35; G=100: 3.5; G=1000: 0.35	0.16	0.3	100 \parallel 3	100 \parallel 3	2.59
INA333	1.8 to 5.5	1-1000	50	50*	100	G=1: 150; G=10: 35; G=100: 3.5; G=1000: 0.35	0.16	0.1*	100 \parallel 3	100 \parallel 3	5.72
AD8237	1.8 to 5.5	1-1000	68	115	100-114	Low Bandwidth Mode – G=1: 200; G=10: 20; G=100: 2; G=1000: 0.2; High Bandwidth Mode – G=10: 100; G=100: 10; G=1000: 1	Low Bandwidth Mode: 0.05; High Bandwidth Mode: 0.15	0.3	0.1 \parallel 5	0.8 \parallel 10	2.90*
AD8420	2.7 to 36	1-1000	55	90	130*	G=1: 250; G=10: 25; G=100: 2.5; G=1000: 0.25	1	1	0.13 \parallel 3	1 \parallel 3	3.08

15 Appendix – FARE-KE Metrics Table Comparison

Variable	Samur proposal	FARE-KE	metrics
Workspace	✓	✓	100×50× 50mm
DOFs	✓	✓	2x rot., 1x transl.
Structure	✓	×	
Elastostatic Dynamics	Elastostatics Structural stiffness	Elastostatics Structural stiffness	25N/mm is consider a rigid object
	Global conditioning index	×	
	Dynamics Structural dynamics	×	
	Generalised inertia ellipsoid	×	
	Acceleration radius	×	
Dexterity	✓	×	
Supply Voltage	×	✓	
Device Dimensions	×	✓	
Weight	×	✓	
Static			
peak force	✓	✓	
continious force	✓	✓	5N
Minimum force	✓	✓	5N
Hysteresis	✓	✓	
Sensitivity	✓	✓	
Output Force Resolution	✓	✓	
Dynamic Range	✓	✓	
Power Consumption	×	✓	
Noise – sound	×	✓	
Noise – vibrations	×	✓	
Dynamic			
Force bandwidth	✓	×	
Useful frequency range	✓	×	
Amplifier bandwidth	✓	×	
Output impedance	✓	×	
Force fidelity	✓	×	
Transient			
Rise time	✓	✓	
Deceleration time	×	✓	
Settling time	✓	✓	
Overshoot	✓	✓	
Output force accuracy	✓	✓	
Force precision	✓	✓	
Performance			
Peak speed	✓	×	
peak acceleration	✓	×	
Structural deformation ratio	✓	×	

16 Appendix – FARE-KE

Proposed Table Of Metrics

Measurement Section	Variable	Metrics	Relevance
Physical Properties	Workspace	XYZ mm	Determines the physical range in which the device can operate
	Degrees of Freedom	2x rot., 1x transl.	Identifies the independent directions in which the device can move or rotate
	Electrical Properties	Voltage, Power	Defines the electrical requirements for the device to function properly
	Device Dimensions	XYZ mm	Specifies the physical size of the device, important for compatibility with use cases
	Weight	grams	Indicates the heaviness of the device which can impact user comfort and device portability
	Operational Noise	decibels	Provides an idea of the acoustic disturbance caused by the device during operation
	Vibrations	m/s ²	Measures the mechanical vibrations generated by the device, which can affect user experience
Ramp Analysis	Max Force/Torque	Newtons	Determines the maximum force or torque that the device can exert
	Min Force/Torque	Newtons	Specifies the smallest detectable force or torque that the device can exert
	Hysteresis	Newtons	Describes the lag in response exhibited by the device when subjected to changing inputs
	Sensitivity	Newtons/Volt	Measures the change in force output per unit change in input signal
	Output Force Resolution	Newtons	Identifies the smallest detectable change in force output by the device
	Dynamic Range	dB	Captures the range of force output the device can provide, from the minimum to maximum force
Step Response Analysis	Stiffness	Newtons/mm	Quantifies the resistance of the device to deformation in response to an applied force
	Peak Force (overshoot)	Newtons	Records the maximum force output by the device in response to a sudden change in input
	Continuous Force	Newtons	Measures the force that the device can maintain continuously over a prolonged period
	Rise time	Seconds	Notes the time it takes for the device to reach the desired output from an initial state
	Settling Time	Seconds	Measures the time the device takes to stabilise after a change in input
	Output Error	%	Provides an estimation of the error between the desired and actual output of the device

17 Appendix – FARE-KE DAQ MCP3654 PCB design

BOM and PCB design and kicad files available on the [Github Repository](#)

18 Appendix – FARE-KE Calibrating Single Axis Load Cell

Calibration is crucial to ensure accurate weight or force measurements from a single-axis load cell. This appendix provides a step-by-step guide on calibrating a single-axis load cell using the proposed Data Acquisition (DAQ) system. Data analysis and calibration values are calculated using Python, with the following necessary steps:

1. Hardware Setup

Connect the load cell to one of the channels of the ADC MCP3564. Once calibrated, always use the same port on the DAQ. Ensure the load cell is securely mounted on one end and oriented in a manner representative of its typical mounting when measuring forces.

2. Data Collection

Load the Arduino sketch “8_mcp_calibration.cpp” onto the microcontroller. Use a program such as TyCommander or telemetryViewer to read the serial command and receive the raw data sent over the serial port (USB). Begin with no load on the load cell and record the zero-value reading. Gradually add known weights to the load cell, covering the expected range of forces to be measured. Record the digital readings for each weight in a CSV file, along with the corresponding weight in grams. If the load cell is expected to measure push/pull forces, invert the orientation of the load cell and repeat the process to calibrate for negative weights. While increasing the number of readings enhances the accuracy of the model derived for the load cell, most load cells tested, including low-cost versions, exhibit a highly linear response. Unless the load cell exhibits non-linear behaviour, obtaining two readings in each direction typically results in an R^2 value very close to 1. Example CSV format:

```
`` `CSV
Weight,Reading
-1500.52,-1591000
-1056.84,-858400
...
1500.52,3379500
1901.97,4045800
`` `
```

3. Data Analysis and Calibration

Use the Python script `Beam_loadCell_calibration.py` to calculate and compare various calibration techniques based on the collected data. The script implements three calibration

techniques: Linear Regression, Polynomial Regression, and Scaling Factor Calibration. The goal is to find the best-fit equation that relates the digital readings to the actual weights.

- **Scaling Factor Calibration:**

$$y = kx$$

where y is the actual weight, x is the digital reading, and k is the scaling factor.

- **Linear Regression:**

$$y = mx + b$$

where y is the actual weight, x is the digital reading, m is the slope, and b is the y -intercept.

- **Polynomial Regression (2nd degree):**

$$y = ax^2 + bx + c$$

where y is the actual weight, x is the digital reading, and a , b , and c are the coefficients.

The script applies an offset correction to the readings by subtracting the reading corresponding to zero weight. Then it visualises the calibration results, comparing the linear regression, polynomial regression, and scaling factor calibration fits. Finally, the script calculates the R^2 score for each calibration technique to assess their performance.

```
```python
r2_linear = r2_score(y_corrected, y_pred_linear)
r2_poly = r2_score(y_corrected, y_pred_poly)
r2_scaling = r2_score(y_corrected[non_zero_mask], y_pred_scaling[non_zero_mask])
```
```

The calibration equations for linear regression, polynomial regression, and the scaling factor are provided. The results and analysis are saved in a Markdown file ('calibration_results.md') for easy review and comparison.

```
```python
md_content = f"""
...
Regression Equations
- **Linear Regression Equation**: $y = \{linear_regressor.coef_0\}x + \{linear_regressor.intercept\}$
- **Polynomial Regression Equation**: $y = \{poly_regressor.coef_2\}x^2 + \{poly_regressor.coef_1\}x + \{poly_regressor.intercept\}$
- **Scaling Factor**: $\{average_scaling_factor\}$
...
"""
```
```

4. Converting Weight to Force

To convert the calibrated weight (in grams) to force (in Newtons), multiply the weight by the gravitational acceleration constant.

$$F = mg$$

where F is the force in Newtons, m is the mass in kilograms, and g is the gravitational acceleration constant (9.80665 m/s^2).

5. Ensuring Accurate Measurements

To ensure accurate calibration of the load cell, maintain the object being weighed in a stationary position and ensure even weight distribution on the load cell during calibration. If the object is suspended, verify that it hangs vertically without any swaying. Temperature can impact the performance of the load cell; therefore, it is advisable to operate the load cell in a controlled environment or undertake periodic recalibration if a controlled environment is not feasible.

5. Calibration Code

```
python
# Beam_loadCell_calibration.py

import numpy as np
import pandas as pd
import matplotlib.pyplot as plt
from sklearn.linear_model import LinearRegression
from sklearn.preprocessing import PolynomialFeatures
from sklearn.metrics import mean_squared_error, r2_score

# Load the data
data = pd.read_csv('calibrationWeightSingleBeamMCP.csv')
x = data['Weight'].values.reshape(-1, 1)
y = data['Reading'].values

# Find the offset by finding the reading corresponding to zero weight
offset = data.loc[data['Weight'] == 0, 'Reading'].values[0]

# Apply offset to the readings
y_corrected = y - offset

# Linear Regression
linear_regressor = LinearRegression()
linear_regressor.fit(x, y_corrected)
y_pred_linear = linear_regressor.predict(x)

# Polynomial Regression (2nd Degree)
poly = PolynomialFeatures(degree=2)
x_poly = poly.fit_transform(x)
poly_regressor = LinearRegression()
poly_regressor.fit(x_poly, y_corrected)
y_pred_poly = poly_regressor.predict(x_poly)

# Scaling Factor Calibration
# Exclude zero readings to prevent division by zero
```

```

non_zero_mask = x.flatten() != 0
scaling_factors = y_corrected[non_zero_mask] / x[non_zero_mask].flatten()
average_scaling_factor = np.mean(scaling_factors)
y_pred_scaling = x.flatten() * average_scaling_factor

# Plotting the results
plt.figure(figsize=(10, 6))
plt.scatter(x, y_corrected, color='red', label='Actual Data Points', marker='o')
plt.plot(x, y_pred_linear, color='blue', label='Linear Regression', linestyle='-',
marker='>')
plt.plot(x, y_pred_poly, color='green', label='Polynomial Regression',
linestyle=':', marker='s')
plt.scatter(x[non_zero_mask], y_pred_scaling[non_zero_mask], color='orange',
label='Scaling Factor Calibration', marker='^', zorder=5)
plt.axhline(0, color='black', linewidth=0.5) # Add a line at zero for reference
plt.title('Calibration Techniques Comparison')
plt.xlabel('Weight')
plt.ylabel('Corrected Reading')
plt.legend()
plt.savefig("calibration_comparison.png")
plt.close()

r2_linear = r2_score(y_corrected, y_pred_linear)
r2_poly = r2_score(y_corrected, y_pred_poly)
r2_scaling = r2_score(y_corrected[non_zero_mask], y_pred_scaling[non_zero_mask])

# Prepare Markdown content
md_content = f"""
# Calibration Results for Single-Axis Load Cell

This document presents the results of applying three different calibration
techniques to sensor readings obtained from a single-axis load cell.

## Calibration Techniques
- **Linear Regression**
- **Polynomial Regression**
- **Scaling Factor Calibration**

![Calibration Techniques Comparison](calibration_comparison.png)

## Statistical Analysis

Technique	R^2
Linear Regression	{r2_linear:.4f}
Polynomial Regression	{r2_poly:.4f}
Scaling Factor	{r2_scaling:.4f}

## Regression Equations

- **Linear Regression Equation**:  $y = \{linear\_regressor.coef\_ [0]:.4f\}x + \{linear\_regressor.intercept\_:.4f\}$ 
- **Polynomial Regression Equation**:  $y = \{poly\_regressor.coef\_ [2]:.4f\}x^2 + \{poly\_regressor.coef\_ [1]:.4f\}x + \{poly\_regressor.intercept\_:.4f\}$ 
- **Scaling Factor**:  $\{average\_scaling\_factor:.4f\}$ 

The analysis includes R^2 scores for each calibration method to help determine the
most accurate approach for converting sensor readings into weight measurements.
"""

# Save Markdown content to file
with open('calibration_results.md', 'w') as md_file:
    md_file.write(md_content)

print("Calibration results and analysis saved to calibration_results.md.")

```

19 Appendix – FARE-KE 3-Axis Load Cell Assembly Guide

Proper alignment of the load cell centres is critical during the assembly process to ensure accurate measurements. The 3D-printed parts play a crucial role in facilitating the alignment of the load cells during assembly and contribute to the overall structural integrity of the assembly.





Figure 124. Left final 3 axis load cell, Right layout and overlap of the single load cell bars.



If manufacturing capabilities allow, using a stiffer material such as aluminium or another material with higher stiffness than 3D-printed materials can significantly enhance the sensor's performance. The increased stiffness of these materials can improve the overall stability and accuracy of the load cell, resulting in more reliable and precise measurements.

Assembly Steps

1. Before inserting the first load cell, apply aluminium epoxy into the load cell holder.



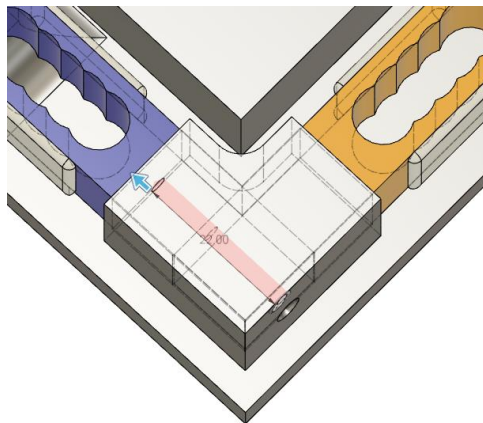
Mount the load cells using the 3D-printed parts, ensuring the orientation aligns with the illustrations



Once the parts are assembled, mark the holes on the aluminium and drill a 2.5mm hole through the load cells.

Use a tapping tool to create a 3mm thread on the load cell.

Secure the joint using an M3 screw.



20 Appendix – FARE-KE 3-Axis Single Point Calibration

Single-point calibration methods are widely utilised for calibrating multi-axis load cells due to their simplicity and straightforward implementation. These methods involve applying a known load on one of the axes and using the resultant sensor readings to establish the relationship between the digital reading and the weight. Each load cell channel exhibits a direct response to the load applied along its respective axis, which is described by a transfer function. For a 3-axis load cell, the primary outputs when a load is applied along the x , y , and z axes are:

$$\begin{bmatrix} O_x \\ O_y \\ O_z \end{bmatrix} = \begin{bmatrix} k_{xx} \cdot F_x \\ k_{yy} \cdot F_y \\ k_{zz} \cdot F_z \end{bmatrix}$$

Where O_i is the output reading, k_{ii} is the sensitivity of the load cell for the respective axis, and F_i is the applied force. When a force is applied along one axis, the other axes register a response due to crosstalk. The transfer functions can be expanded to incorporate crosstalk as follows:

$$\begin{bmatrix} O_x \\ O_y \\ O_z \end{bmatrix} = \begin{bmatrix} k_{xx} \cdot F_x + k_{xy} \cdot F_y + k_{xz} \cdot F_z \\ k_{yx} \cdot F_x + k_{yy} \cdot F_y + k_{yz} \cdot F_z \\ k_{zx} \cdot F_x + k_{zy} \cdot F_y + k_{zz} \cdot F_z \end{bmatrix}$$

The calibration matrix \mathbf{K} is constructed using the main sensitivity coefficients along the diagonal and the crosstalk coefficients as the off-diagonal elements:

$$\mathbf{K} = \begin{bmatrix} k_{xx} & k_{xy} & k_{xz} \\ k_{yx} & k_{yy} & k_{yz} \\ k_{zx} & k_{zy} & k_{zz} \end{bmatrix}$$

To correct for sensor inaccuracies and crosstalk, the inverse of the calibration matrix is calculated. This inverse matrix transforms the output readings (previously denoted as (\mathbf{O})) into actual force values. The mathematical representation for calculating the actual forces \mathbf{F} from the output readings using the inverse calibration matrix \mathbf{K}^{-1} is as follows:

$$\mathbf{F} = \mathbf{K}^{-1}\mathbf{O}$$

Here, \mathbf{O} represents the vector of output readings from the sensor, corresponding to the raw data previously mentioned. The inverse matrix serves two key purposes: it establishes a relationship between the digital reading and the weight/load applied to the sensor and compensates for crosstalk, which is essential for ensuring accurate force measurement.

Numerical Example Based on Calibration Data

1. Measurements under Known Loads: With known loads applied along each axis (150g in this example), we record the corresponding output readings for all three axes.
2. Calculating the Sensitivities: The sensitivity along each axis is calculated as the ratio of the output reading when a load is applied along that axis. For example:

$$Sensitivity_x = \frac{O_x \text{ with } F_x}{F_x}$$

3. Calculating Crosstalk Coefficients: The crosstalk coefficients are similarly ratios of the output reading on the non-loaded axes when a load is applied on one axis. For instance:

$$Crosstalk_{xy} = \frac{O_y \text{ with } F_x}{F_x}$$

4. Constructing the Calibration Matrix: Using the calculated sensitivities and crosstalk coefficients, we construct the calibration matrix \mathbf{K} .
5. Solving for Actual Forces: To find the actual forces from any given set of load cell readings, we multiply the inverse calibration matrix \mathbf{K}^{-1} by the raw sensor readings vector \mathbf{O} , obtaining \mathbf{F} .

Implementing the Calibration in Practice

The following numerical example is part of the calibration process for one of the assembled 3-axis load cells build. The load cell was loaded with a 150g align along each one of the axis and this data is used to construct the calibration matrix \mathbf{K} and its inverse. The data for a 3-axis load cell, which contains the weights suspended from the load cell aligned to each of the axes and the digital readings from the ADC, is as follow:

| load_x | load_y | load_z | reading_x | reading_y | reading_z |
|--------|--------|--------|-----------|-----------|-----------|
| 0 | 0 | 0 | 251370 | -549405 | 1150070 |
| 150 | 0 | 0 | 503955 | -544585 | 1147510 |
| 0 | 150 | 0 | 249710 | -292290 | 1147780 |
| 0 | 0 | 150 | 254260 | -543310 | 897411 |

Table 29. Readings from the 3axis load cell with a 150g load aligned on each of the axis.

Calculate the sensitivities and crosstalk coefficients as follows, sensitivity for the x-axis:

$$Sensitivity_x = \frac{Reading_{x \text{ with } F_x} - Reading_{x \text{ with no force}}}{F_x}$$

$$Sensitivity_x = \frac{503955 - 251370}{151 - 0} \approx 1672.55$$

The same process is applied to the other two axes. Then for the crosstalk coefficients, for the x-axis with force applied on the y-axis:

$$Crosstalk_{xy} = \frac{Reading_{x \text{ with } F_y} - Reading_{x \text{ with no force}}}{F_y}$$

$$Crosstalk_{xy} = \frac{249710 - 251370}{151 - 0} \approx -11.06$$

For the x-axis with force applied on the z-axis:

$$Crosstalk_{xz} = \frac{Reading_{x \text{ with } F_z} - Reading_{x \text{ with no force}}}{F_z}$$

$$Crosstalk_{xz} = \frac{254260 - 251370}{151 - 0} \approx -19.21$$

These calculations compute the sensitivity coefficients that reflect how much the sensor's output changes per unit of force applied, and the crosstalk coefficients that reflect the sensor's response on one axis when force is applied on another. Now, we use the calculated sensitivities and crosstalk coefficients to construct the calibration matrix \mathbf{K} :

$$\mathbf{K} = \begin{bmatrix} Sensitivity_x & Crosstalk_{xy} & Crosstalk_{xz} \\ Crosstalk_{yx} & Sensitivity_y & Crosstalk_{yz} \\ Crosstalk_{zx} & Crosstalk_{zy} & Sensitivity_z \end{bmatrix} \rightarrow \mathbf{K} = \begin{bmatrix} 1672.55 & -11.06 & -19.21 \\ 31.82 & 1703.01 & 40.30 \\ -17.02 & -15.17 & 1672.65 \end{bmatrix}$$

The inverse of the calibration matrix \mathbf{K}^{-1} is then computed, which is used to determine the actual forces from the raw sensor readings compensating for sensitivity and crosstalk:

$$\mathbf{K}^{-1} = \begin{bmatrix} 5.9788 \times 10^{-4} & 3.9432 \times 10^{-6} & 6.7715 \times 10^{-6} \\ -1.1313 \times 10^{-5} & 5.8699 \times 10^{-4} & -1.4273 \times 10^{-5} \\ 5.9811 \times 10^{-6} & 5.3638 \times 10^{-6} & 5.9779 \times 10^{-4} \end{bmatrix}$$

This inverse matrix can be calculated using numerical methods such as NumPy's 'linalg.inv()' function in Python. Implementation of this calculation is integrated into the provided Arduino library, specifically within the class 'MCP356x3axis.h', designed for reading from the custom-built PCB. Although the Single-Point Calibration Method is straightforward, it has inherent limitations. It assumes linearity in the sensor's response and does not accommodate non-linear behaviours or complex interactions between axes that may arise from misalignment or other assembly-related factors.

Code:

```

python
import numpy as np

# Given data for each load case - Updated with new readings
measured_no_load = np.array([251370, -549405, 1150070]) # No load case
measured_load_x = np.array([503955, -544585, 1147510]) # Load 1 (Applying load_x = 150g)
measured_load_y = np.array([249710, -292290, 1147780]) # Load 2 (Applying load_y = 150g)
measured_load_z = np.array([254260, -543310, 897411]) # Load 3 (Applying load_z = 150g)

# Known loads applied along each axis
load_x = 150 # Known force applied for load case x
load_y = 150 # Known force applied for load case y
load_z = 150 # Known force applied for load case z

# Sensitivity Calculation (counts/N)
sensitivity_x = (measured_load_x[0] - measured_no_load[0]) / load_x
sensitivity_y = (measured_load_y[1] - measured_no_load[1]) / load_y
sensitivity_z = (measured_load_z[2] - measured_no_load[2]) / load_z

# Crosstalk Calculation (counts/N)
# Assuming no crosstalk when no load is applied
crosstalk_xy = (measured_load_x[1] - measured_no_load[1]) / load_x
crosstalk_xz = (measured_load_x[2] - measured_no_load[2]) / load_x
crosstalk_yx = (measured_load_y[0] - measured_no_load[0]) / load_y
crosstalk_yz = (measured_load_y[2] - measured_no_load[2]) / load_y
crosstalk_zx = (measured_load_z[0] - measured_no_load[0]) / load_z
crosstalk_zy = (measured_load_z[1] - measured_no_load[1]) / load_z

# Calibration Matrix
calibration_matrix = np.array([
    [sensitivity_x, crosstalk_xy, crosstalk_xz],
    [crosstalk_yx, sensitivity_y, crosstalk_yz],
    [crosstalk_zx, crosstalk_zy, sensitivity_z]
])

# Inverse of the Calibration Matrix
calibration_matrix_inverse = np.linalg.inv(calibration_matrix)

print("Calibration matrix:")
print(calibration_matrix)
print("Inverse of the calibration matrix:")
print(calibration_matrix_inverse)

```

21 Appendix – FARE-KE 3-Axis Multipoint Calibration Method

The Multipoint Calibration Method addresses the limitations of the Single-Point Calibration Method by considering multiple load cases, accounting for non-linearity and crosstalk effects. This method is particularly relevant for the custom-built 3-axis load cell used in the project. Due to the nature of the custom design, the load cell may exhibit more crosstalk and non-linearity compared to precision-manufactured, single solid-body load cells. The Multipoint Calibration Method helps mitigate these issues and provides a more accurate calibration matrix. This section presents two different approaches: the least-squares method and polynomial regression.

Least-Squares Regression

The least-squares regression is a method that finds the calibration matrix \mathbf{K} by minimising the sum of squared differences between the observed sensor outputs readings \mathbf{R} and the model's predictions. This approach inherently accounts for crosstalk since it treats the calibration as a system of linear equations that considers all interactions between axes:

$$\mathbf{K} = \operatorname{argmin} \|\mathbf{F} - \mathbf{R}\mathbf{K}\|^2$$

Here, \mathbf{F} is the matrix of applied forces and \mathbf{R} is the matrix of sensor readings. The calibration matrix \mathbf{K} obtained from this method compensates for crosstalk by incorporating it into the model.

Polynomial Regression

The polynomial regression calibration model also aims to compensate for non-linear relationships. It fits a polynomial equation to the data points, offering a more nuanced model of sensor behaviour:

$$F = K_2 \times R^2 + K_1 \times R + K_0$$

Where K_2 , K_1 , and K_0 represent the coefficients for the quadratic, linear, and constant terms.

Implementation and Validation

1. Data Collection: As with the single point calibration of the single axis load cell, a number of known weights are suspended from the sensor, aligning with the axis being loaded, and the digital value is recorded. These values are then saved in CSV format. An example from the calibration of one of the assemble 3-axis load cells looks is:

| $Load_x$ | $Load_y$ | $Load_z$ | $read_x$ | $read_y$ | $read_z$ |
|----------|----------|----------|----------|----------|----------|
| 0 | 0 | 0 | 251370 | -549405 | 1150070 |
| 150 | 0 | 0 | 503955 | -544585 | 1147510. |
| 300 | 0 | 0 | 730660 | -540480 | 1145120 |
| 600 | 0 | 0 | 1236265 | -531180 | 1139075 |
| 900 | 0 | 0 | 1709310 | -522370 | 1128960 |
| 0 | 0 | 0 | 251370 | -549405 | 1150070 |
| 0 | 150 | 0 | 249710 | -292290 | 1147780 |
| 0 | 300 | 0 | 248770 | -60985 | 1143355 |
| 0 | 600 | 0 | 248065 | 454290 | 1143280 |
| 0 | 900 | 0 | 244730 | 937005 | 1132765 |
| 0 | 0 | 0 | 251370 | -549405 | 1150070 |
| 0 | 0 | 150 | 254260 | -543310 | 897411 |
| 0 | 0 | 300 | 256340 | -538270 | 670606 |
| 0 | 0 | 600 | 262460 | -526805 | 164996 |
| 0 | 0 | 900 | 264660 | -515265 | -308384 |

Table 30 Digital reading for the loading of each of the 3 axis to compute calibration matrix

The axes are loaded with weights (in grams), and the readings are captured in digital count form. Once the data is ready, the next steps are handled by a Python script that computes the calibration matrix and compares the performance of both methods. To assess performance, a separate set of readings from the load cell with known weights, which have not been used for calibration, are captured for validation of the load cell performance.

2. Calibration Matrix Computation: The calibration matrix is derived from the collected data using least-squares regression, which considers the linear behaviour of the sensor.

3. Curve Fitting: Polynomial regression is used to fit curves to the data for each axis, capturing non-linear behaviours.

4. Validation: The models are validated using a separate dataset, not involved in the calibration process.

5. Error Analysis: The model's accuracy is assessed through MAE and RMSE, allowing for model refinement.

Code

```

import pandas as pd
import numpy as np
from numpy.linalg import lstsq
from numpy.polynomial.polynomial import Polynomial
from sklearn.metrics import mean_absolute_error, mean_squared_error

# Load the calibration data from a CSV file into a pandas DataFrame
load_cell_data = pd.read_csv('3axisCell2CalibrationData.csv')

# Selects only the relevant columns for calibration
data = load_cell_data[['force_x', 'force_y', 'force_z', 'corrected_x',
'corrected_y', 'corrected_z']]
data = data.dropna() # Drop rows with NaN values

# Extracts the force measurements and the corresponding corrected sensor readings
forces = data[['force_x', 'force_y', 'force_z']].values
corrected_readings = data[['corrected_x', 'corrected_y', 'corrected_z']].values

# Identifies the no-load (zero force) sensor readings to use for offset correction
no_load_readings = corrected_readings[forces[:,0] == 0, :][0]
corrected_readings_offset_corrected = corrected_readings - no_load_readings

# Calculates the calibration matrix using the Least Squares method
K_offset_corrected, _, _, _ = lstsq(corrected_readings_offset_corrected, forces,
rcond=None)

# Initialises Polynomial Regression for each axis (x, y, z) using the corrected
sensor readings
coefficients_poly = {}
for axis, axis_name in enumerate(['x', 'y', 'z']):
    poly_fit = Polynomial.fit(corrected_readings[:, axis], forces[:, axis], 2)
    coefficients_poly[axis_name] = poly_fit.convert().coef

# Load the validation data
validation_data = pd.read_csv('validationDataCell2.csv')

# Extract corrected sensor readings from the validation data
validation_corrected_readings = validation_data[['corrected_x', 'corrected_y',
'corrected_z']].values

# Perform offset correction on the validation data
validation_corrected_readings_offset_corrected = validation_corrected_readings -
no_load_readings

# Use the calibration matrix to estimate forces from validation data
validation_estimated_forces_ls =
np.dot(validation_corrected_readings_offset_corrected, K_offset_corrected.T)

# Use polynomial regression coefficients to estimate forces from validation data
validation_estimated_forces_poly = np.column_stack([
    np.polyval(coefficients_poly['x'][:, :-1], validation_corrected_readings[:, 0]),
    np.polyval(coefficients_poly['y'][:, :-1], validation_corrected_readings[:, 1]),
    np.polyval(coefficients_poly['z'][:, :-1], validation_corrected_readings[:, 2])
])

# Calculate MAE and RMSE for Least Squares and Polynomial Regression methods
mae_ls = mean_absolute_error(forces, validation_estimated_forces_ls)
rmse_ls = np.sqrt(mean_squared_error(forces, validation_estimated_forces_ls))
mae_poly = mean_absolute_error(forces, validation_estimated_forces_poly)
rmse_poly = np.sqrt(mean_squared_error(forces, validation_estimated_forces_poly))

# Outputs
K_offset_corrected, coefficients_poly, mae_ls, rmse_ls, mae_poly, rmse_poly

```

Validation Results

To assess the performance of the three calibration methods (single point, least squares, and polynomial regression) a separate validation dataset was used. This dataset consisted of sensor readings and corresponding known force measurements that were not used during the calibration process. The validation phase highlighted the performance of the Polynomial Regression method in terms of both Mean Absolute Error (MAE) and Root Mean Square Error (RMSE), as summarised in the following table:

| Method | MAE (g) | RMSE (g) | Max Error (g) |
|-----------------------|---------|----------|---------------|
| Single-Point | 8.91 | 15.24 | 70.43 |
| Least Squares | 8.70 | 14.72 | 67.29 |
| Polynomial Regression | 7.54 | 11.28 | 44.93 |

Table 31. Validation Performance Comparison across different modelling techniques for load prediction. The table quantifies the models' accuracy using Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Maximum Error, all measured in grams (g).

The Polynomial Regression method demonstrated consistent performance across different datasets and exhibited robustness against crosstalk. The Single-Point method showed slightly higher errors compared to the Least Squares method in the validation phase, indicating its limitations in compensating complex sensor behaviours. However, the improvements from the Least Squares method are limited. The validation data was processed using each calibration method to estimate the forces and evaluate the performance of each technique. These estimated forces were then compared to the actual force measurements to calculate the absolute errors. The subsequent figure present the absolute errors for each method, across

Error Values (grams)

| Load(g) | | | Least Squares Error | | | Polynomial Error | | | Single Point Error | | |
|---------|--------|---------|---------------------|------|------|------------------|------|------|--------------------|------|------|
| x | y | z | x | y | z | x | y | z | x | y | z |
| -60.0 | 0.0 | 0.0 | 6.1 | 1.0 | 3.2 | 7.2 | 7.6 | 1.3 | 6.1 | 1.2 | 3.3 |
| -285.0 | 0.0 | 0.0 | 8.0 | 0.6 | 8.9 | 9.8 | 6.1 | 2.3 | 7.8 | 1.2 | 9.4 |
| -465.0 | 0.0 | 0.0 | 8.3 | 0.6 | 14.1 | 10.8 | 5.3 | 5.7 | 8.0 | 1.6 | 14.9 |
| -1055.0 | 0.0 | 0.0 | 12.5 | 1.3 | 32.8 | 18.3 | 3.3 | 18.8 | 11.8 | 3.7 | 34.5 |
| -1925.0 | 0.0 | 0.0 | 18.7 | 3.2 | 67.3 | 31.9 | 1.0 | 44.9 | 17.3 | 7.4 | 70.4 |
| 0.0 | 60.0 | 0.0 | 5.8 | 0.8 | 1.1 | 5.6 | 4.5 | 1.3 | 5.9 | 0.8 | 1.1 |
| 0.0 | 285.0 | 0.0 | 9.4 | 1.9 | 0.7 | 5.0 | 0.8 | 2.0 | 9.5 | 2.1 | 1.0 |
| 0.0 | 465.0 | 0.0 | 12.4 | 1.5 | 1.9 | 4.6 | 2.3 | 5.1 | 12.6 | 1.9 | 2.3 |
| 0.0 | 1055.0 | 0.0 | 22.8 | 5.4 | 10.0 | 4.2 | 2.0 | 10.8 | 23.4 | 6.2 | 10.9 |
| 0.0 | 1925.0 | 0.0 | 41.0 | 15.5 | 36.3 | 6.4 | 22.5 | 4.7 | 42.2 | 16.9 | 37.9 |
| 0.0 | 0.0 | -60.0 | 4.8 | 0.3 | 3.1 | 4.8 | 5.9 | 0.1 | 4.6 | 0.3 | 3.0 |
| 0.0 | 0.0 | -285.0 | 5.4 | 1.6 | 4.6 | 2.5 | 2.2 | 4.2 | 4.5 | 2.0 | 4.5 |
| 0.0 | 0.0 | -465.0 | 6.3 | 2.0 | 4.8 | 1.0 | 0.2 | 5.3 | 4.9 | 2.7 | 4.5 |
| 0.0 | 0.0 | -1055.0 | 8.5 | 7.4 | 8.0 | 4.7 | 12.0 | 5.9 | 5.4 | 9.0 | 7.3 |
| 0.0 | 0.0 | -1925.0 | 9.2 | 16.2 | 12.8 | 15.7 | 30.3 | 10.4 | 3.4 | 19.0 | 11.6 |

Figure 125. Heatmap table comparison for the three calibration methods presented, showing the error on grams when loading the sensor after calibrated.

all the validating measurement visualising the performance across a range of load conditions.

As illustrated in the heatmaps, error rates are notably lower for the Least Squares and Polynomial Regression methods compared to the Single-Point method, with the Polynomial Regression method demonstrating the lowest error percentages overall.

The mean absolute error, even the highest performance calibration method can have a MAE of up to 7.5 grams. This affects measurements under the 100-gram range particularly negatively, where this error value of 7.5 grams could represent up to an 7.5% error. However, as we measure higher forces, this noise becomes less relevant and tend to remain at lower than 2% error over the loading reading.

The calibration and validation process for the custom-built 3-axis load cell provides a comprehensive overview of the performance of the three calibration methods. The Polynomial Regression method was found to be the most effective, offering slightly improved accuracy compared to the Least Squares and Single-Point methods. However, it is crucial to acknowledge that even the best-performing method has its limitations, particularly when measuring low forces, as demonstrated by the data in. Future work will focus on addressing these challenges by enhancing the ADC noise immunity and increasing the sampling rate to acquire more data for better noise filtering and compensation.

Additionally, regular validation and performance monitoring of the calibrated sensor are essential to ensure consistent and accurate force measurements over time. The long-term performance of these sensors has not been extensively studied. Establishing a systematic

Error Percentages (%)

| Load(g) | | | Least Squares Error | | | Polynomial Error | | | Single Point Error | | |
|---------|------|-------|---------------------|------|------|------------------|-------|------|--------------------|------|------|
| x | y | z | x | y | z | x | y | z | x | y | z |
| -60 | 0 | 0 | 10.17 | 1.67 | 5.33 | 12.00 | 12.67 | 2.17 | 10.17 | 2.00 | 5.50 |
| -285 | 0 | 0 | 2.81 | 0.21 | 3.12 | 3.44 | 2.14 | 0.81 | 2.74 | 0.42 | 3.30 |
| -465 | 0 | 0 | 1.78 | 0.13 | 3.03 | 2.32 | 1.14 | 1.23 | 1.72 | 0.34 | 3.20 |
| -1055 | 0 | 0 | 1.18 | 0.12 | 3.11 | 1.73 | 0.31 | 1.78 | 1.12 | 0.35 | 3.27 |
| -1925 | 0 | 0 | 0.97 | 0.17 | 3.50 | 1.66 | 0.05 | 2.33 | 0.90 | 0.38 | 3.66 |
| 0 | 60 | 0 | 9.67 | 1.33 | 1.83 | 9.33 | 7.50 | 2.17 | 9.83 | 1.33 | 1.83 |
| 0 | 285 | 0 | 3.30 | 0.67 | 0.25 | 1.75 | 0.28 | 0.70 | 3.33 | 0.74 | 0.35 |
| 0 | 465 | 0 | 2.67 | 0.32 | 0.41 | 0.99 | 0.49 | 1.10 | 2.71 | 0.41 | 0.49 |
| 0 | 1055 | 0 | 2.16 | 0.51 | 0.95 | 0.40 | 0.19 | 1.02 | 2.22 | 0.59 | 1.03 |
| 0 | 1925 | 0 | 2.13 | 0.81 | 1.89 | 0.33 | 1.17 | 0.24 | 2.19 | 0.88 | 1.97 |
| 0 | 0 | -60 | 8.00 | 0.50 | 5.17 | 8.00 | 9.83 | 0.17 | 7.67 | 0.50 | 5.00 |
| 0 | 0 | -285 | 1.89 | 0.56 | 1.61 | 0.88 | 0.77 | 1.47 | 1.58 | 0.70 | 1.58 |
| 0 | 0 | -465 | 1.35 | 0.43 | 1.03 | 0.22 | 0.04 | 1.14 | 1.05 | 0.58 | 0.97 |
| 0 | 0 | -1055 | 0.81 | 0.70 | 0.76 | 0.45 | 1.14 | 0.56 | 0.51 | 0.85 | 0.69 |
| 0 | 0 | -1925 | 0.48 | 0.84 | 0.66 | 0.82 | 1.57 | 0.54 | 0.18 | 0.99 | 0.60 |

Figure 126. Heatmap table comparison for the three calibration methods presented, showing the error as a percentage of the load applied to the axis.

approach to recalibration and performance evaluation will be crucial in maintaining the sensor's precision and reliability across various applications.

22 Appendix – FARE-KE 6-Axis Load Cell Assembly

This appendix complements the discussions in Chapter 8, specifically addressing the assembly of the 6-axis force and torque sensor used in the HapticWhirl project. The aim is to provide a detailed inventory of parts, along with step-by-step assembly and calibration instructions, enabling individuals to independently assemble a F/T plate. The equations below are fundamental for calculating the forces and torques (moments) applied by haptic devices:


$$\begin{bmatrix} F_x \\ F_y \\ F_z \\ M_x \\ M_y \\ M_z \end{bmatrix} = \begin{bmatrix} F_{x_A} + F_{x_B} + F_{x_C} + F_{x_D} \\ F_{y_A} + F_{y_B} + F_{y_C} + F_{y_D} \\ F_{z_A} + F_{z_B} + F_{z_C} + F_{z_D} \\ (W/2)(-F_{z_A} - F_{z_B} + F_{z_C} + F_{z_D}) \\ (L/2)(-F_{z_A} + F_{z_B} + F_{z_C} - F_{z_D}) \\ (L/2)(F_{y_A} - F_{y_B} - F_{y_C} + F_{y_D}) + (W/2)(F_{x_A} + F_{x_B} - F_{x_C} - F_{x_D}) \end{bmatrix}$$

Comprehensive guidelines for assembling this open-source sensor are provided, with the latest version of this document and source materials accessible via the project's GitHub repository. This document includes visual aids for assembly and calibration, with general information about component fabrication and procurement detailed in the "BOM.md" file on GitHub. Further details, including CAD files, are available in the Source/CAD folder.

The document is structured to first list all components necessary for assembling the force/torque plate, followed by step-by-step assembly instructions. Separate documentation covers the connection of electronics, firmware configuration, and calibration procedures.

23 Appendix – FARE-KE 6-Axis Load Cell Multipoint Calibration

Multi-axis force-torque sensors are integral in diverse applications such as robotics, biomechanics, and industrial automation. These sensors provide crucial data by measuring forces and torques along multiple axes, aiding in control, monitoring, and analysis. However, ensuring accurate and reliable measurements necessitates thorough calibration.



The calibration process defines a precise relationship between raw digital sensor readings and the actual force/torque values they represent. This relationship is typically encapsulated by a calibration matrix that transforms the raw sensor data into quantifiable force and torque measurements, taking into account factors such as sensor sensitivity, crosstalk effects, and nonlinear response characteristics.

Calibrating multi-axis force-torque sensors involves addressing challenges like nonlinearities (where the sensor's output does not linearly correlate with the applied force or torque) and crosstalk, which occurs when a force or torque applied on one axis affects the measurements on other axes. These phenomena can significantly impact the accuracy and reliability of sensor measurements if not carefully addressed during calibration.

Single-Point Calibration Methods

Calibration methods for multi-axis force-torque sensors often employ single-point calibration techniques. These methods involve applying a single known force/weight and torques and using the corresponding sensor readings to compute the calibration matrix. The calibration process determines the elements of the calibration matrix that best fit the relationship between the applied forces/torques and the sensor readings. The single-point calibration matrix is typically obtained by solving a system of linear equations:

$$\begin{bmatrix} F_x \\ F_y \\ F_z \\ \tau_x \\ \tau_y \\ \tau_z \end{bmatrix} = \begin{bmatrix} k_{11} & k_{12} & \cdots & k_{16} \\ k_{21} & k_{22} & \cdots & k_{26} \\ \vdots & \vdots & \ddots & \vdots \\ k_{61} & k_{62} & \cdots & k_{66} \end{bmatrix} \begin{bmatrix} r_1 \\ r_2 \\ \vdots \\ r_6 \end{bmatrix}$$

Here, F_x, F_y, F_z denote the applied forces, τ_x, τ_y, τ_z the applied torques, k_{ij} the elements of the calibration matrix, and r_1, r_2, \dots, r_6 the raw sensor readings.

Single-point calibration methods have limitations in capturing the full range of sensor behaviour and accounting for nonlinearities across the entire measurement range. However, on industrial sensors, this approach has been proven to reduce the effects of cross-talk to less than 0.5% of the sensor's full-scale capacities [348], [349].

The force/torque (F/T) plate under study is constructed using four pre-calibrated 3-axis load cells. These load cells are rigidly assembled to form a 3-axis four-cell torque measurement plate. By combining 3-axis load cells in this configuration, we can provide orthogonal force measurements along the x, y, and z axes, representing a total of 12 independent channels of forces on each axis. In this case, the cells are calibrated before mounting. Once mounted on the measuring plate, the readings from these 12 channels are combined to calculate the force and torque on each axis as follows:

$$\begin{bmatrix} F_x \\ F_y \\ F_z \\ M_x \\ M_y \\ M_z \end{bmatrix} = \begin{bmatrix} F_{x_A} + F_{x_B} + F_{x_C} + F_{x_D} \\ F_{y_A} + F_{y_B} + F_{y_C} + F_{y_D} \\ F_{z_A} + F_{z_B} + F_{z_C} + F_{z_D} \\ (W/2)(-F_{z_A} - F_{z_B} + F_{z_C} + F_{z_D}) \\ (L/2)(-F_{z_A} + F_{z_B} + F_{z_C} - F_{z_D}) \\ (L/2)(F_{y_A} - F_{y_B} - F_{y_C} + F_{y_D}) + (W/2)(F_{x_A} + F_{x_B} - F_{x_C} - F_{x_D}) \end{bmatrix}$$

Here, $F_x, F_y,$ and F_z represent the net forces in the x, y, and z directions, respectively. Similarly, $M_x, M_y,$ and M_z denote the net moments about the corresponding axes. The subscripts $A, B, C,$ and D identify the individual load cells. W and L signify the width and length of the measurement plate, respectively.

The coordinate origin is defined at the centre of the plate, where the load cells are rigidly mounted. This 6-axis load cell design, unlike highly precise industrial builds, exhibits higher crosstalk and lower linearity. Consequently, it is crucial to employ calibration approaches that mitigate these limitations. We employ multipoint calibration, which involves collecting data at multiple force and torque levels across the sensor's measurement range. This method allows us to capture the sensor's nonlinear behaviour and enhance the calibration matrix's capability to compensate for crosstalk effects. Benefits of multipoint calibration include:

- **Improved accuracy:** Multipoint calibration captures the sensor's nonlinear behaviour, thereby providing more accurate force and torque estimates compared to single-point calibration methods.
- **Enhanced robustness:** This calibration approach is less sensitive to measurement noise and outliers, utilising multiple data points to compute a more reliable calibration matrix.
- **Better compensation for crosstalk effects:** By considering interactions between different measurement axes, multipoint calibration significantly reduces errors attributable to crosstalk.

In the following sections, we explore the methodology, experimental setup, and results of the multipoint calibration approach, assessing its performance on accuracy and reliability of multi-axis force-torque sensor measurements.

Methodology

Our proposed multipoint calibration approach aims to enhance the accuracy and robustness of multi-axis force-torque sensor calibration by using data collected at various force/torque levels. The key steps in the methodology include:

1. Data Collection:

Apply known forces and torques at multiple calibration points across the sensor's measurement range. Record the corresponding raw sensor readings for each point. The collected data is organised into matrices: \mathbf{F} for applied forces/torques and \mathbf{R} for raw sensor

readings. The reading vector can be visualized as a 6x6 matrix for each calibration point i , denoted as \mathbf{R}_i :

$$\mathbf{R}_i = \begin{bmatrix} F_x(\text{Primary}) & F_y(\text{Secondary}) & \cdots & M_z(\text{Secondary}) \\ F_x(\text{Secondary}) & F_y(\text{Primary}) & \cdots & M_z(\text{Secondary}) \\ \vdots & \vdots & \ddots & \vdots \\ F_x(\text{Secondary}) & F_y(\text{Secondary}) & \cdots & M_z(\text{Primary}) \end{bmatrix}$$

In this matrix, the primary effects are on the diagonal, while the off-diagonal terms reflect crosstalk effects.

2. Calibration Matrix Computation

The objective is to derive a calibration matrix \mathbf{K} that converts uncalibrated force/torque readings into precise measurements. The least squares method is utilised to compute \mathbf{K} , minimising the discrepancy between the product of raw sensor readings \mathbf{R} and the calibration matrix \mathbf{K} , and the known applied forces and torques \mathbf{F} . This is mathematically represented as:

$$\min_{\mathbf{K}_i} \sum_{i=1}^n \|\mathbf{F}_i - \mathbf{R}_i \cdot \mathbf{K}_i\|^2$$

where xn is the total number of calibration points.

The least squares method seeks to find the calibration matrix \mathbf{K} that best fits the relationship between the raw sensor readings and the known applied forces and torques. To compute \mathbf{K} using the least squares method, the raw sensor readings \mathbf{R} and the corresponding known forces and torques \mathbf{F} from all calibration points are organised into matrices. Let \mathbf{R}_{cal} denote the matrix of raw sensor readings from the calibration points, and \mathbf{F}_{cal} denote the matrix of corresponding known forces and torques. The least squares solution for the calibration matrix \mathbf{K} is given by:

$$\mathbf{K} = (\mathbf{R}_{\text{cal}}^T \mathbf{R}_{\text{cal}})^{-1} \mathbf{R}_{\text{cal}}^T \mathbf{F}_{\text{cal}}$$

This equation identifies the calibration matrix that minimises the sum of squared differences between the predicted forces and torques ($\mathbf{R}_{\text{cal}} \mathbf{K}$) and the actual known forces and torques

(\mathbf{F}_{cal}) across all calibration points. The resulting calibration matrix \mathbf{K} can be used to estimate the forces and torques for a given set of raw sensor readings \mathbf{R} :

$$\mathbf{F}_{\text{estimated}} = \mathbf{R}\mathbf{K}$$

Since the least squares method incorporates all calibration points in computing the calibration matrix, no additional optimization is required. This method effectively minimises the overall error across all points, accounting for the sensor's nonlinear behaviour and crosstalk effects. By leveraging data from multiple points across the measurement range, the least squares method offers a robust and efficient calibration for the multi-axis force-torque sensor.

In the subsequent sections, we will explore the experimental setup, results, and performance evaluation of the multipoint calibration strategy using the least squares method on the custom-designed 6-axis load cell.

Experimental Setup

The calibration was conducted using a 6-axis F/T plate assembled from four 3-axis load cells. Recorded readings were obtained from the 12 channels to compute the forces and torques. The force plate was carefully positioned to ensure proper alignment with the desired measurement axes. For force measurements, the plate was placed sideways and levelled to maintain a flat and vertical orientation, aligning the applied forces with the axes being measured.

To measure forces, weights were suspended from the centre of the origin, aligning the applied force with the x and y axes. For the z-axis, the force plate rested flat, and weights were placed on a platform mounted over the plate's origin, ensuring even weight distribution. The 3-axis load cells, already calibrated in Newtons, recorded the loaded weights as follows: 0.59 N, 1.42 N, 2.01 N, 2.80 N, 3.92 N, 6.72 N, 11.33 N, 15.89 N, and 21.68 N. For torque measurements, a 115 mm lever-arm was attached to the plate's centre. Weights were suspended from the arm's end to generate torques of 0.09 Nm, 0.21 Nm, 0.30 Nm, 0.42 Nm, 0.59 Nm, 1.01 Nm, 1.70 Nm, 2.38 Nm, and 3.25 Nm. Zero values were also recorded, resulting in a total of 10 calibration points for each force and torque component.

| Fx | Fy | Fz | Mx | My | Mz | known_Fx | known_Fy | known_Fz | Known_Mx | Known_My | known_Mz |
|--------|--------|--------|--------|-------|--------|----------|----------|----------|----------|----------|----------|
| -0.026 | -0.049 | 0.013 | 0.001 | 0.002 | 0.002 | 0 | 0 | 0 | 0 | 0 | 0 |
| -0.006 | 0.563 | 0 | -0.01 | 0.002 | -0.002 | 0 | 0.5886 | 0 | 0 | 0 | 0 |
| 0.042 | 2.04 | -0.031 | -0.045 | 0.002 | -0.011 | 0 | 2.0111 | 0 | 0 | 0 | 0 |
| 0.026 | 0.007 | -0.076 | 0.001 | 0 | -0.005 | 0 | 0 | 0 | 0 | 0 | 0 |
| 0.048 | 0.608 | -0.107 | -0.093 | 0.002 | -0.012 | 0 | 0.5886 | 0 | -0.0883 | 0 | 0 |
| 0.072 | 2.055 | -0.166 | -0.32 | 0.005 | -0.026 | 0 | 2.0111 | 0 | -0.3017 | 0 | 0 |

Table 32. shows an example of the calibration data collected for the force and torque measurements.

It is important to note that when applying a load to measure torque, the weights will also be measured as a force aligned with that axis. For example, when measuring torque on the x-axis, the y-axis force component will also experience load, as illustrated in Table 32. In the following sections, we will explore data processing, calibration matrix computation using Python, and performance evaluation based on the collected calibration data.

Data Preparation and Calibration Matrix Computation

To obtain the calibration matrix using the method previously described, a Python script is provided that handles all functionality and includes some data validation. This script segments the collected data into calibration and validation sets to evaluate the calibration matrix's performance on alternative data points.

```

python
calibration_rows = []
validation_rows = []

for i in range(0, len(data), 10):
    axis_data = data.iloc[i:i+10]
    calibration_rows.extend(axis_data.index[1:10:2]) # odd-indexed rows starting
    from index 1
    validation_rows.extend(axis_data.index[2:10:2]) # even-indexed rows starting
    from index 2

# Specify calibration data based on indices
calibration_data = data.loc[calibration_rows]
validation_data = data.loc[validation_rows]

```

The script processes data in chunks of 10 rows, each chunk representing calibration points for each axis. Odd-indexed rows (starting from index 1) are designated for calibration, while even-indexed rows (starting from index 2) are used for validation. This segregation ensures that both data sets are distinct. Calibration data is then selected using `data.loc[calibration_rows]`, and validation data is accessed via `data.loc[validation_rows]`. Next, the script extracts the sensor readings (R) and the corresponding known forces/torques (F) that have been applied to each axis from the calibration data:

```

python

```

```
R_cal = calibration_data.iloc[:, :6].values
F_cal = calibration_data.iloc[:, 6:].values
```

```

The `iloc` function selects the first 6 columns of the calibration data, corresponding to sensor readings ( $\mathbf{R}$ ), with the subsequent columns representing the known forces/torques ( $\mathbf{F}$ ). The `values` attribute then converts this selected data into NumPy arrays for processing. The calibration matrix is computed using the least squares method:

```
python
K_cal, _, _, _ = lstsq(R_cal, F_cal, rcond=None)
```

```

The `lstsq` function from the `numpy.linalg` module addresses the linear least squares problem by taking sensor readings (\mathbf{R}_{cal}) and the known forces/torques (\mathbf{F}_{cal}) as inputs, returning the calibration matrix (\mathbf{K}_{cal}) that minimises the sum of squared residuals between the predicted and actual forces/torques. This calibration matrix represents the linear transformation that maps raw sensor readings to calibrated force/torque values, incorporating adjustments for the sensor's sensitivity, crosstalk effects, and inherent biases.

Calibration Matrix Results

Using the calibration data from the sensor and the Python script described previously, we obtained the following calibration matrix:

$$\mathbf{K}_{\text{cal}} = \begin{bmatrix} 0.9996 & 0.0102 & 0.0126 & -0.0004 & -0.0217 & 0.0051 \\ -0.0303 & 0.9946 & 0.0150 & 0.0258 & 0.0010 & 0.0049 \\ -0.0091 & -0.0159 & 0.9946 & 0.0011 & -0.0094 & -0.0000 \\ -0.1138 & -0.1117 & -0.1168 & 1.1472 & 0.0196 & -0.0137 \\ -0.4016 & -0.0569 & 0.5885 & 0.0120 & 1.1564 & 0.0466 \\ -0.0173 & 0.0068 & 0.1061 & 0.0170 & -0.0043 & 1.0059 \end{bmatrix}$$

To assess the calibration matrix's performance, we conducted various tests comparing the corrected readings' accuracy to the original readings obtained during calibration. First, we evaluated calibration effectiveness using Mean Absolute Percentage Error (MAPE) and R-squared metrics. The results indicated that the readings obtained during calibration, which combined data from the twelve readings of the calibrated four 3-axis load cells, already achieved an R^2 value of 0.98. However, after applying the calibration matrix, the sensor accuracy further improved to an R^2 value of 0.999.

| Metric | Raw Readings | Calibration |
|-----------|--------------|-------------|
| MAPE (%) | 2.814737 | 1.849263 |
| R-squared | 0.983804 | 0.999691 |

Table 33 Mean Absolute Percentage Error (MAPE) and R-squared metrics before and after using the calibration matrix.

Next, histograms and box plots illustrated the distribution of errors between the raw and corrected data Figure 127. The histograms shows a notable convergence of the calibrated readings towards zero error, signifying a substantial reduction in error magnitude relative to the raw data. The box plots confirmed these results, showing that the interquartile range (IQR) of the calibrated data was significantly narrower and the median closer to zero, suggesting enhanced accuracy.




Figure 127. Error Distribution for Raw and Calibrated Readings histogram and Error Distribution (Box Plot) for Raw and Calibrated Readings

To visualise the error across the loading range and identify any outliers, we plotted scatter plots illustrating raw and calibrated readings across the six axes in comparison to the known applied force. For all axes, the calibrated values were closer to the linear model representative of the sensor's response, indicating higher accuracy after calibration.

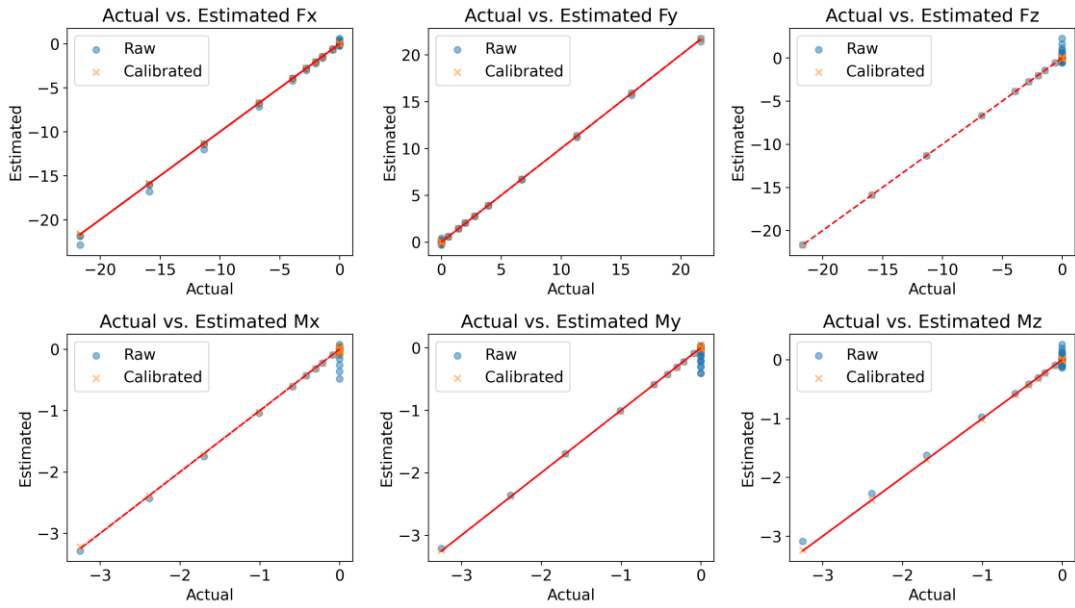


Figure 130. Scatter plots for Actual vs. Estimated Fx, Fy, Fz, Mx, My, and Mz, with readings on Newton and Newton/meter respectively.




Figure 128. MAE by Axis for Raw vs. Calibrated bar chart.




Figure 129. As axis are actuated, they have an undesired effect on other axis called crosstalk. This plot shows the crosstalk effect on the Fz axis above and the Mz axis below when another axis are actuated. On orange are the raw values when compensated with the calibration matrix K_{cal} .

each axis in Newtons for the force axes and Newton-meters for the torque axes. Overall, calibrated readings offered lower error readings compared to the raw readings.

However, this analysis primarily assessed the calibration matrix's effect on overall accuracy without specifically addressing its impact on crosstalk compensation. Consequently, additional tests were conducted to evaluate the calibration matrix's influence on crosstalk. In the crosstalk behaviour in the 6-axis load cell design, we plotted the crosstalk errors and compared them to the readings after applying the compensation matrix across all calibration points. The calibration sequence, depicted in the graph, began with applying a force along F_y , followed by M_x , F_x , M_y , M_z , and finally F_z . This plot highlights the crosstalk on the F_z and M_z axes and demonstrates how crosstalk intensity escalates with increased load on the principal axis. The compensated readings, shown in orange, indicate a substantial reduction in crosstalk across these axes, with similar outcomes observed for the other axes.

| Raw Force/Torque Crosstalk | | | | | | Calibrated Force/Torque Crosstalk (N) | | | | | |
|----------------------------|----------|----------|-----------|-----------|-----------|---------------------------------------|----------|----------|-----------|-----------|-----------|
| $F_x(N)$ | $F_y(N)$ | $F_z(N)$ | $M_x(Nm)$ | $M_y(Nm)$ | $M_z(Nm)$ | $F_x(N)$ | $F_y(N)$ | $F_z(N)$ | $M_x(Nm)$ | $M_y(Nm)$ | $M_z(Nm)$ |
| 0.026 | 0.049 | 0.013 | 0.001 | 0.002 | 0.002 | 0.026 | 0.049 | 0.013 | 0.000 | 0.003 | 0.002 |
| 0.006 | 0.000 | 0.000 | 0.01 | 0.002 | 0.002 | 0.023 | 0.000 | 0.011 | 0.003 | 0.003 | 0.001 |
| 0.015 | 0.000 | 0.018 | 0.029 | 0.001 | 0.002 | 0.025 | 0.000 | 0.008 | 0.004 | 0.002 | 0.006 |
| 0.042 | 0.000 | 0.031 | 0.045 | 0.002 | 0.011 | 0.015 | 0.000 | 0.006 | 0.001 | 0.003 | 0.000 |
| 0.055 | 0.000 | 0.048 | 0.059 | 0.000 | 0.002 | 0.022 | 0.000 | 0.001 | 0.004 | 0.001 | 0.013 |
| 0.107 | 0.000 | 0.074 | 0.093 | 0.000 | 0.022 | 0.000 | 0.000 | 0.005 | 0.006 | 0.000 | 0.001 |
| 0.180 | 0.000 | 0.129 | 0.156 | 0.000 | 0.037 | 0.004 | 0.000 | 0.011 | 0.006 | 0.001 | 0.001 |
| 0.301 | 0.000 | 0.22 | 0.258 | 0.001 | 0.062 | 0.01 | 0.000 | 0.021 | 0.004 | 0.000 | 0.001 |
| 0.453 | 0.000 | 0.283 | 0.362 | 0.001 | 0.089 | 0.017 | 0.000 | 0.004 | 0.007 | 0.000 | 0.003 |
| 0.594 | 0.000 | 0.363 | 0.483 | 0.003 | 0.116 | 0.002 | 0.000 | 0.015 | 0.003 | 0.001 | 0.000 |

Table 34. Crosstalk Error From Raw Readings and calibrated Values after applying the calibration matrix to the loading values along the F_y axis.

After calibration, there was a clear reduction in crosstalk error across all axes. The calibration proved particularly effective in mitigating crosstalk on the M_x axis, reducing errors to near zero, on other axes, the crosstalk magnitudes were significantly lowered. Overall, the calibration matrix substantially reduced crosstalk, particularly along the torque axes, where the standard deviation decreased by a factor of 10 for all metrics used to assess crosstalk errors.

| Statistic | Fx | Fy | Fz | Mx | My | Mz |
|-------------|--------|--------|--------|--------|--------|--------|
| Std Dev raw | 0.1139 | 0.0821 | 0.3907 | 0.0846 | 0.0979 | 0.0528 |
| Std Dev Cal | 0.0224 | 0.0377 | 0.0343 | 0.0154 | 0.0085 | 0.0043 |
| Mean Raw | 0.0683 | 0.0541 | 0.2465 | 0.0332 | 0.0538 | 0.0369 |
| Mean Cal | 0.0172 | 0.0231 | 0.0310 | 0.0072 | 0.0049 | 0.0045 |
| Max Raw | 0.5940 | 0.4120 | 2.2360 | 0.4830 | 0.4100 | 0.2580 |
| Max Cal | 0.0861 | 0.1700 | 0.1287 | 0.0783 | 0.0517 | 0.0158 |

Table 35. Standard deviation, mean values and max error for the error on the raw readings, and the errors after applying the calibration matrix to the data.

In conclusion, the assessment of the calibration matrix, as applied to the sensor data, highlights its effectiveness in improving measurement accuracy and reducing crosstalk across multiple axes. The statistical analysis, demonstrating decreases in standard deviation, mean errors, and maximum error values post-calibration, confirms the calibration matrix's role in enhancing sensor performance and the capabilities of the proposed 6-axis design. It is important to note that the calibration and subsequent error evaluations were conducted with forces and torques applied separately to individual axes. This method ensures precise calibration but does not consider scenarios where multiple axes are actuated simultaneously. While no adverse effects on calibration accuracy have been observed under these conditions, the possibility remains. Future work will compare these results with those obtained using industrial-grade sensors under multi-axis loading conditions to provide deeper insights into the robustness of the calibration matrix and the sensor's reliability in more complex operational environments.

24 Appendix – FARE-KE Analysis Implementations

Step and ramp response analysis can be found on the directory

[Haptic Device Testing](#)

25 Appendix – FARE-KE ADC Libraries

All the files for the ADC microcontroller, libraries, dependencies and examples can be found on the repository ([MCP356xScale](#), [MCP356x3axis](#) and [MCP356x6axis](#))

<https://github.com/telenaco/FARE->

[KE/tree/main/MCP356x Arduino Library/lib/mcp356x/src](https://github.com/telenaco/FARE-KE/tree/main/MCP356x%20Arduino%20Library/lib/mcp356x/src)