

RICE UNIVERSITY

Referred Haptic Feedback for Virtual Hand Interactions Through a Bracelet Interface

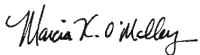
By

Evan Pezent

A THESIS SUBMITTED  
IN PARTIAL FULFILLMENT OF THE  
REQUIREMENTS FOR THE DEGREE

Doctor of Philosophy

APPROVED, THESIS COMMITTEE



Marcia K. O'Malley

Thomas Michael Panos Family Professor in  
Mechanical Engineering, Computer Science,  
and Electrical and Computer Engineering



Matthew Brake

Assistant Professor in Mechanical  
Engineering

Benjamin J. Fregly

Benjamin J. Fregly (Mar 3, 2021 10:39 CST)

BJ Fregly

Professor of Mechanical Engineering and  
Bioengineering and a CPRIT Scholar in  
Cancer Research

Philip Kortum

Philip Kortum (Feb 25, 2021 19:21 CST)

Philip Kortum

Associate Professor of Psychological  
Sciences and Chairman of the  
Undergraduate Committee

HOUSTON, TEXAS

February 2021

## ABSTRACT

### Referred Haptic Feedback for Virtual Hand Interactions Through a Bracelet Interface

by

Evan Pezent

A revolution in mobile interfaces is unfolding, as researchers and large corporations race toward all-day wearable technology for virtual and augmented reality. In a future where interacting with digital entities and objects is a facet of everyday life, an important aspect emerges: our ability to touch and feel objects that do not physically exist. Tactile feedback for virtual interaction is currently limited to either simple handheld controllers or obtrusive wearable devices for the hand, both unsuited to all-day use. This thesis presents an alternate approach: referred haptic feedback through bracelet interfaces. The design of a compact yet robust multimodal haptic bracelet is presented, along with novel control solutions for wrist squeeze force and high density vibrotactile arrays. The haptic rendering capabilities of the device are experimentally characterized and evaluated through psychophysical studies. Leveraging the multisensory combination of wrist squeeze and vibration with visual illusions, we explore the potential of providing substitutive feedback for interaction forces that would otherwise occur at the hands and fingertips. Findings from two human subject studies suggest that referred haptic feedback to the wrist provides more than just a metaphor for interaction forces, and instead invokes genuine perceptions of object stiffness.

## Acknowledgments

I would like to thank my advisor Dr. Marcia O'Malley for the opportunity to work in the MAHI Lab and her relentless support and mentorship over the years. Thank you for always allowing me to pursue my interests, as tangential as they may have been at times, and for helping me grow into the scientist and engineer I am today. I would also like to thank my committee members, Dr. Fregly, Dr. Brake, and Dr. Kortum for their time in being a part of this process. Special thanks is owed to Dr. Nick Colonnese for his unwavering support throughout my PhD. This research would not have been possible without you, and I look forward to many years of working together. Thank you to my fellow MAHI Lab members, especially Chad Rose, Craig McDonald, Janelle Clark, Dylan Losey, Nathan Dunkelberger, Brandon Cambio, Troy Dennis, Jenny Sullivan, Alix Macklin, and Zane Zook. Chad, you continue to be one of my greatest influences and supporters. Thank you for all of the help and words of wisdom you have shared with me. Craig, the summer we decided to hunker down in the basement to write *just a bit* of software is still one of my fondest memories, though I do apologize for the extra year it may have added to your degree. Janelle, our journey together though this trying process has been a wild ride full of many unforgettable memories. Your friendship and emotional support has undoubtedly helped me reach this point, and I know you aren't far behind! Dylan, our mutual respect for the sciences and of course NASCAR and college football always meant there was a heated discussion to be had. Nathan, it has been a joy watching you develop your technical skills and institutional knowledge of all things software and hardware related – do not let our efforts die in vain! Jenny, the day you left was a sad time in the MAHI Lab, but I take pleasure in knowing that the folks at NASA are now enjoying your fiery personality as much as I did. I still laugh about how much we struggled solving that CAPTCHA. Brandon, thank you for your help

with Tasbi and Syntacts and for putting up with my ridiculous tendencies on a daily basis; your dedication was *shocking* at times. Troy, you should know that since you left I have only been to the gym three times. I guess our morning workouts together were enough to get my decrepit body through these last few years. Alix, the times we spent fumbling our way through MLE will not be forgotten; I hope you can accept a rain check for our crazy EEG experiment. Zane, I am forever grateful for your help in distributing our open-source projects, and I can't wait to see where you decide to take it next. To the students I have mentored, Caz, Rod, Ahalya, and Lianne: I hope you got as much out of our time together as I did. Your creativity showed no bounds, and I was always impressed by the things you accomplished. I would be remiss not to mention my friends and colleagues at Facebook Reality Labs – Gaurav Mukherjee, Ali Israr, Priyanshu Agarwal, Jess Hartcher-O'Brien, Shea Robinson, and Majed Samad – who were all instrumental in the success of this work. My parents, brothers, and in-laws have provided no shortage of love and support. The mothers in my life, Amy and Valerie, have been absolute rocks during difficult times. Thank you coming to the rescue when Krystin and I needed you most (and for proofreading everything I ever wrote!). Jamey, it's been fun going through our PhDs together, and I can't wait to see your work shoot for the stars (literally). Finally, and above all, I thank my loving wife and best friend Krystin for her constant devotion and support. Together (and with Penny, of course) we did the impossible of completing two doctoral degrees under the same roof. Nothing stands in our way now. Here's to spending the rest of our lives together as we continue to grow, build a family, and pursue our dreams. Love you!

# Contents

Abstract	i
Acknowledgments	ii
List of Illustrations	ix
List of Tables	xxv
<b>1 Introduction</b>	<b>1</b>
1.1 State of the Art in XR Haptic Feedback	2
1.1.1 Handheld Controllers	2
1.1.2 Haptic Glove Devices	3
1.1.3 Haptic Fingertip Devices	3
1.2 Upper Limb-Grounded Haptic Devices	5
1.2.1 Where should an all-day wearable device be worn?	5
1.2.2 Which haptic modalities are both compelling and viable?	6
1.2.3 Are compact wrist devices possible with current technology?	7
1.3 Referred Haptic Feedback for XR Hand Interactions	8
1.4 Thesis Contribution and Outline	9
<b>2 Haptic Bracelets and Armbands</b>	<b>11</b>
2.1 Vibrotactile Bracelets	11
2.2 Squeeze Bracelets and Armbands	14
2.3 Multimodal Bracelets and Armbands	18
<b>3 Design of Tasbi Haptic Bracelet</b>	<b>20</b>
3.1 Design Criteria	20

3.2 Squeeze Mechanism	21
3.3 Squeeze Force Sensing	26
3.4 Vibrotactile Band	28
3.5 Power and Control Unit	30
<b>4 Squeeze Control</b>	<b>32</b>
4.1 Torque Control	33
4.2 Position Control	35
4.3 Force Control	36
4.4 Sensor Calibration	38
4.5 Controller Implementation	40
4.6 Controller Characterization	43
4.7 Discussion and Future Improvements	46
<b>5 Vibrotactile Control</b>	<b>48</b>
5.1 Background	48
5.2 Introduction to Vibrotactor Control	50
5.2.1 Function Generators	50
5.2.2 Integrated Circuits	50
5.2.3 Dedicated Controllers	51
5.2.4 Audio Output Devices	52
5.3 Hardware for Audio-Based Control	52
5.3.1 Sound Cards / Digital-to-Analog Converters	52
5.3.2 Amplifiers	54
5.3.3 Syntacts Amplifier	56
5.4 Software for Audio-Based Control	58
5.4.1 Syntacts API	59
5.4.2 Syntacts GUI	66
5.5 Comparison	67

5.5.1	Latency Benchmarking	67
5.5.2	Overall Comparison	71
5.6	Limitations and Future Work	72
<b>6</b>	<b>Psychophysics of Haptic Stimuli at the Wrist</b>	<b>74</b>
6.1	Study 1: Squeeze Difference Threshold	74
6.1.1	Subjects and Procedures	75
6.1.2	Results	76
6.2	Study 2: Vibrotactile Identification versus Squeeze Force	77
6.2.1	Subjects and Procedures	77
6.2.2	Results	79
6.3	Study 3: Vibrotactile Identification versus Amplitude	82
6.3.1	Subjects and Procedures	83
6.3.2	Results	84
6.4	Discussion	84
<b>7</b>	<b>Referred Haptic Feedback for Hand Interactions</b>	<b>86</b>
7.1	Multisensory Feedback Paradigm	86
7.1.1	Vibrotactile Feedback	87
7.1.2	Squeeze Feedback	89
7.1.3	Pseudo-Haptic Control/Display Feedback	89
7.2	Virtual Reality Implementation and Sandbox Environment	93
7.3	Interaction Themes and Examples	96
7.3.1	Primitive Motions	96
7.3.2	Surfaces and Textures	98
7.3.3	In-Hand Manipulation	99
7.3.4	Object Weight and Inertia	99
7.3.5	Force Fields	101
7.3.6	User Movement and Proprioception	102

7.3.7	Bimanual Interactions	103
7.3.8	Mutli-Paradigm Interactions	104
<b>8</b>	<b>Perception of Mid-Air Stiffness</b>	<b>105</b>
8.1	Background	106
8.1.1	Evidence Supporting Referred Squeeze	106
8.1.2	Evidence Supporting Pseudo-Haptics	107
8.1.3	Multisensory Integration	108
8.2	Experiment 1: Equating Mid-Air and Physical Stiffness	109
8.2.1	Physical Button Apparatus	110
8.2.2	Virtual Environment	110
8.2.3	Tasbi Calibration Procedure	111
8.2.4	Subjects	112
8.2.5	Experimental Procedure	112
8.2.6	Data Analysis	115
8.2.7	Results	116
8.2.8	Discussion	119
8.3	Experiment 2: Discriminating Mid-Air Stiffness	125
8.3.1	Subjects	126
8.3.2	Virtual Environment	126
8.3.3	Experimental Procedure	126
8.3.4	Survey	127
8.3.5	Data Analysis	129
8.3.6	Results	129
8.3.7	Discussion	131
8.4	Conclusion and Future Work	137
<b>9</b>	<b>Conclusion and Future Work</b>	<b>140</b>
9.1	Tasbi Self-Containment	143



9.2 Hands-Free AR Interactions . . . . .	144
9.3 Integrating Tasbi with Other Technologies . . . . .	146
9.4 Sharing Technology with the Research Community . . . . .	147
<b>Bibliography</b>	<b>149</b>

# Illustrations

1.1	Haptic devices for XR interaction can be broadly categorized as either (1) handheld controllers, (2) wearable glove devices, or (3) wearable fingertip devices. . . . .	4
1.2	Referred haptic feedback could enable a variety of interactions with the virtual world. Vibrotactile feedback may substitute for fingertip contact with virtual buttons or other types of user interfaces, while directional forces may convey the weight or stiffness associated with manipulating virtual objects. Referred haptic feedback can further add to telepresence and remote social interactions such as hand shaking or holding, and also provide immersive feedback for gaming and training environments. . . . .	8
1.3	This thesis presents Tasbi, a haptic bracelet that provides referred squeeze and vibrational feedback for hand interactions in virtual environments. . . . .	10
2.1	Vibrotactile Bracelets – (a-c) Planar vibrotactile grid arrays for the wrist. (d-g) Radial vibrotactile arrays for the wrist. (h) A slightly different but notable indentation-based tactile wristband. . . . .	13
2.2	Squeeze Bracelets and Armbands — (a-g) Electromechanical, band-based devices. (h) A device employing shape memory alloys. (i-k) Devices that squeeze only in the normal direction. (l-p) Pneumatic based squeezing devices. . . . .	17

2.3 Multimodal Bracelets and Armbands – (a) Baumman’s device integrated servo-based wrist squeeze and tapping. (b) The MISSIVE device combined servo-based squeeze and stretch with four voice coil vibrotactors. (c) Aggravi’s device combined servo-based squeeze with four ERM vibrotactors. . . . . 19

3.1 The final iteration of Tasbi as worn on the wrist. The device incorporates six vibrotactile actuators and a novel squeezing mechanism. Not pictured is the tethering cable required to power and with interface. . . . . 21

3.2 (a) Typical constricting-band approaches to squeeze produce non-uniform and tangential forces which would cause embedded factors to shift. (b) Our decoupled approach aims to produce pure, uniform normal forces. . . . 22

3.3 (a) A rudimentary bracelet was developed to test the proposed pin and cord approach to squeeze. (b) The first prototype of Tasbi was 3D printed and provided only open-loop torque control capabilities. (c) Tasbi v1.0 improved the design with metal reinforced components and integrated encoder feedback for position control of the tensioning spool. (d) Tasbi v2.0 added an improved HDMI connector interface and integrated force sensing to enable closed-loop control of normal squeeze forces. . . . . 23

3.4	Tasbi Exploded View – The squeeze mechanism consists of a 12 mm DC motor <b>(a)</b> and a 13 mm Harmonic Drive gearbox <b>(b)</b> which drives a two-sided spool <b>(c)</b> to create tension in a UHMWP cord <b>(d)</b> . Spool position feedback is provided through an optical encoder consisting of a reflective code wheel <b>(e)</b> and optoelectronic sensor <b>(f)</b> . Squeeze force feedback is measured via a force-sensing capacitor <b>(g)</b> and signal conditioning PCB <b>(h)</b> . The force sensor is held in light compression against a bottom plate <b>(i)</b> with four corner springs <b>(j)</b> . The drive assembly drops into the main housing <b>(k)</b> and is secured in place with a housing lid <b>(l)</b> . Each vibrotactor unit <b>(m)</b> contains a 10 mm LRA vibrotactor and a smooth stainless steel pin to convert cord tension into normal force. Vibrotactor units are clipped into elastic sidings <b>(n)</b> and secured with lids <b>(o)</b> . All signals and power to and from Tasbi are transmitted over a micro-HDMI cable that connects to an internal breakout PCB <b>(q)</b> . . . . .	25
3.5	Tasbi Force Sensor – A SingleTact 8mm, 10 N force-sensing capacitor is in sandwiched between the main housing and a plastic bottom plate. The plate is held in compression through four compression springs and screws so that the sensor experiences minimally pre-load. . . . .	27
3.6	Tasbi Vibrotactor Assembly – The tactor is secured into the lower plastic housing via two O-rings and a foam layer. The metal lid contains a press fit smoothpin over which the tension cord slides. Tactor power cables are embedded within the 3D printed band sides in a post-processing step. . . .	29
3.7	The Tasbi control box interfaces each Tasbi with the host PC via intermediate connections to a Quanser Q8-USB digital acquisition device and a MOTU 24Ao audio sound card. . . . .	31

3.8	Tasbi fabrication. (a) Each Tasbi is comprised of over 150 parts that require hand assembly. (b) A vibrotactile band assembly before and after tactor enclosure. (c) A tensioner housing assembly before enclosure (d) Completed Tasbi v2.0 units. Eight units have been fabricated for internal uses and to share with collaborators. . . . .	31
4.1	(a) A 3D printed rig was designed to characterize Tasbi’s distribution of squeeze load. A fixed ATI Nano17 (1) force sensor measure squeeze force under the main housing, while a relocatable Nano17 (2) allowed for measuring squeeze force under each tactor independently. (b) A Tasbi stretched onto the force rig for characterization. (c) The forces under the housing and each tactor were characterized. Colors correspond with Fig. 4.2. . . . .	34
4.2	(a) Steady state normal vs. tangential forces under each vibrotactor module and the main housing as a function of commanded motor torque. (a) Representative responses for each torque step under the main housing. .	35
4.3	The impedance of the wrist changes considerably with posture. Here, measurements of force and spool position are shown as a user transitions between wrist extension and flexion while either a sinusoidal position (left) or force trajectory (right) is controlled for. The top two plots clearly show we cannot assume that squeeze force will remain proportional to spool position. Though position control is accurately maintained during hand movement, the delivered squeeze stimulus changes drastically (bottom left). Thus, we require direct control of squeeze force (bottom right).	37

4.4	Apparati – (a) The instrumented wrist integrates a high-accuracy ATI Nano17 force sensor and silicone-based simulated tissue. It served a critical role during the force controller development and tuning phases. (b) While general purpose calibrations of Tasbi’s integrated force sensor can be obtained with the instrumented wrist, an actuated calibrator facilitated customized calibrations to individuals by applying a known load through the bracelet. It was primarily used prior to the psychophysical experiments in Chapter 6 and 8 to ensure accurate reporting of force. . . . .	39
4.5	Tasbi’s internal force sensor is calibrated against the instrumented wrist’s Nano17. The left shows the voltage-force data and fit, and the right shows the Tasbi’s measured force after the fit is applied compared to the Nano17. . . . .	40
4.6	Three controllers were tested by first closing the control loop with feedback from the Nano17 sensor, then applying the controller to Tasbi’s sensor and scaling the controller gains to stability. A feedforward and PD controller with the derivative term conditioned on <i>motor velocity</i> offered the best performance. . . . .	41
4.7	Top - The PD <sub>v</sub> +FF controller was further refined to display RMSE less than 5% for a sinusoidal trajectory, and a rise time of 70 ms. Bottom - The final controller was validated in a VR context. Here, the controller renders the interaction force as a user jiggles a virtual button (see Chapter 7). . . . .	44
4.8	Comparison of the position and force controllers. Comparable ranges of squeeze and position are first identified (a), and then the controllers are separately commanded to track an excitation signal within their determined range (b). Although the force controller exhibits more phase lag (d), its overall bandwidth defined by the 3 dB cutoff is a quite comparable 9.1 Hz (d). . . . .	45

5.1 Pipeline for vibrotactile control through audio interfaces. A host PC (a) runs the target application or virtual environment and interfaces with an audio rendering layer (b). The audio renderer interfaces with a digital-to-analog converter (DAC) sound card to output low voltage audio signals (c). The signals are power amplified (d) and sent to the vibrotactors (e) where they are felt as vibrations. Syntacts provides both the audio rendering layer (b) and amplifiers purpose made for driving vibrotactors (d). 53

5.2 Mean Windows audio driver API latencies with standard deviation. Data collection methods are described in Sec. 5.5. For reference, the dashed line indicates the perceptual threshold of visual-haptic simultaneity [117]. 55

5.3 The effect on latency due to changing audio buffer sizes. . . . . 56

5.4 The Syntacts amplifier is an open-source fully differential, linear amplifier capable of driving eight vibrotactors with minimal noise. Two variants are available: one with a single AES-59 DB25 input for connecting to high-end audio devices such as the MOTU 24Ao, and one with four standard 3.5 mm TRS headphone inputs for connecting to general audio outputs or surround sound cards. Both require a 5V power supply, and output amplified signals through a universal 0.1” pitch header. . . . . 57

5.5 The Syntacts amplifier can be used in a variety of applications, ranging from dense tactile arrays (a) to wearable devices such as Tasbi (b). Designs for the tactile array are available online as a reference implementation. . . . . 57

5.6 Signals created in Listing 5.2 . . . . . 62

5.7 Sequenced Signals created in Listing 5.3 . . . . . 63

5.8 The Spatializer created in Listing 5.4 . . . . . 65

5.9 Syntacts GUI - The left-hand side demonstrates cue design. Users drag, drop, and configure Signals from the design Palette to the Designer workspace. The Signal is visualized and can be played on individual channels of the opened device. The right-hand side shows the GUI's track-based sequencer (background) and spatializer (foreground) interfaces. Once designs are complete, they can be saved and later loaded from the programming APIs. . . . . 66

5.10 Syntacts In Use - This figure demonstrates a real-world implementation of the Syntacts amplifier, where it has been used to drive two Tasbi haptic bracelets. A professional grade audio device (MOTU 24Ao) is connected to two Syntacts amplifier boards that have been integrated into separate Tasbi control units. Amplifier output is transmitted to each Tasbi over a multi-conductor cable. Each Tasbi bracelet incorporates six Mplus 1040W LRA tactors radially spaced around the wrist, for a total of twelve utilized audio channels. The audio device interfaces with a host PC (not shown) through a USB connection. . . . . 67

5.11 Syntacts In Use - Here, the C# binding of the Syntacts API is used in Unity Engine to provide haptic effects for a virtual fan interaction designed for the Tasbi setup shown in Fig. 5.10. Two usage paradigms are in effect. The first leverages pre-designed, finite Signals for knob detents (designed in the Syntacts GUI and loaded at runtime) and button contact events (created programmatically on-the-fly, parameterized by hand approach velocity). The second paradigm uses an infinitely long Signal for the fan air stream. The volume and pitch of this Signal are modified in realtime based on the user's hand location and the fan speed, respectively. One-dimensional spatialization is used to target only the tactors which are oriented toward the fan in a continuous fashion. . . . . 68



5.12	The testing used for all latency benchmarking. An Mplus ML1040W LRA was epoxied to a 100 g ABS block, and an accelerometer measured LRA induced vibrations along the y-axis. Latency was defined as the time from calling the software APIs to command vibration to the time at which 0.015 g of acceleration was measured.	69
5.13	Latency as a function of channels rendered, measured as the time from software triggering to the detection of tactor acceleration. Only four channels are shown for the EAI control unit since this is its max.	71
6.1	The GUI subjects interacted with during the studies. (a) Interface for vibrotactile identification study. (b) Interface for squeeze difference threshold study.	75
6.2	(a) The mean psychometric function(s) experimentally determined for wrist squeeze force. Error bars represent a 95% confidence interval. (b) The mean difference threshold, or JND, was found to be 1.28 N, and (c) the PSE shows little bias from the Standard of 7 N. Importantly, we find that wrist size has no significant effect on the JND.	78

6.3	Stimuli-response confusion matrices for each squeeze force and VT stimulus duration pairing, aggregated across all subjects. The probability of subjects correctly responding are given as a function of the stimulus site. Individual columns sum to 100%. The bottom left matrix combines all conditions. The total percentage of correct responses for each condition are given in the subplot titles. The bottom right inlay displays the approximate location of each stimulus relative to wrist and forearm anatomy, particularly the radial (R) and ulnar (U) bones. In general, we observe: 1) identification rates are greatly reduced at stimulus sites located over bony areas, 2) subjects seem to perform better given a longer stimulus duration, and 3) the middle squeeze level of 5N yields the best performance, suggesting that there may exist an optimum level of preload squeeze force. . . . .	80
6.4	Percentage of correctly identified vibrotactor stimuli for both groups of wrist size as a function of the stimulus location. The data represents the mean of all conditions, and error bars are for a 95% confidence interval. . .	81
6.5	Percentage of correct responses separated by squeeze condition and vibration stimuli amplitude. High amplitude vibrations yielded best performance, and individuals with large wrists circumferences ( $\geq 168$ mm) outperform those with small wrist sizes. . . . .	83
7.1	Vibrations can be emulated using a decaying sinusoid model. Here, the model is fit to vibration data recorded while tapping wood. Adapted from [128]. . . . .	87

7.2	Example of the multisensory feedback paradigm – *a) The user approaches the virtual button simulated by a mass $m$ , stiffness $k$ , and damping $b$ . The proxy finger control (blue) and display (gray) are coupled via a virtual spring of stiffness $k_p$ and are initially co-located. When contact is made, Tasbi’s six LRAs render a vibration to simulate the event.	
	(b) The user begins to push the button downward. Tasbi squeeze force increases proportional to the button displacement $x$ . The proxy hand control continues to track the users true hand position and orientation, while the display remains on the surface of the button. The control-to-display (C/D) ratio is given by the ratio of of $x_{\text{control}}$ and $x_{\text{display}}$ . At the end of travel, squeeze reaches its maximum force level, and the C/D discrepancy is most pronounced. Note that subjects did not receive a visual representation of the control hand, and are unaware of its presence.	91
7.3	All virtual Tasbi interactions were built using Unity Engine. A custom plugin wraps the native Tasbi C++ API into accessible C# scripts, and provides Unity <i>prefab</i> objects that can be easily added to any scene. For instance, one <i>TasbiHub</i> prefab is required to launch the squeeze control server and Syntacts engine, and then any number of <i>Tasbi</i> prefabs can be added to the scene to interface with a physical Tasbi device.	94
7.4	The Tasbi sandbox environment where most haptic interaction design took place. The environment presents users with six islands that can be teleported between. Each island displays a particular idea or concept we wish to explore with Tasbi.	95

7.5 The multisensory mid-air button interaction. (a) The user approaches the button. The proxy hand control (blue) and display (black) are initially co-located. (b) The user makes initial contact with the button. Tasbi's LRAs render a vibration to simulate the contact event. (c) The user begins to push the button downward. Tasbi squeeze force increases proportional to the button displacement. The proxy hand control continues to track the users true hand position and orientation, while the display remains on the surface of the button. (d) At the end of travel, squeeze reaches its maximum force level, and the C/D discrepancy has become more pronounced. Note that the blue control hand is only shown for illustrative purposes, and users are unaware of its presence. . . . . 97

7.6 Primitive Interactions – (a) The button already described by Fig. 7.5, showing the events of squeeze and vibration. Pull handles (b) and rotary knobs (d) use an approach similar to the button paradigm for twisting and pulling motions, respectively. Both also display vibration, squeeze, and C/D haptics. These three prototypes form the basis of many of the interactions that follow. . . . . 97

7.7 Surfaces and Textures – vibration is used to render fine detail textures (a), while squeeze offers the ability to render course, geometry-based surfaces (c). Both can be combined (b) and even extended to 2D surfaces (d) . . . . 98

7.8 In-Hand Manipulation – Squeeze and vibration can be used to convey finger interaction forces when compressing visually deformable objects. . . 99

7.9 Object Weight and Inertia – The weight and inertia of rigid-body objects  
 can be rendered through squeeze. We find that the naive approach to  
 displaying mass, through applying only a static amount of squeeze  
 proportional to the object’s mass, is not effective. Instead, it is more  
 compelling to provide fluctuating squeeze proportional to the object’s  
 dynamics (e.g. its velocity or momentum) or proportional to wrench  
 torques its center of mass would otherwise impart on the hand. . . . . 100

7.10 Force Fields – Squeeze and vibration somewhat convincingly render the  
 sensation of air streams, as in this fan example where squeeze increases  
 with proximity, and noisy vibrations are rendered via the factors facing the  
 stream. The example on the right is a playful science fiction device that  
 provides visual context for Tasbi’s calibration process. . . . . 101

7.11 User Movement and Proprioception – This ladder interaction, where  
 squeeze increases as users lift themselves, explores movement and  
 enhancing proprioception. However, it is arguably not a particularly well  
 suited application for squeeze, since it impossible provide enough force  
 with Tasbi to convincingly convey someone’s own mass. . . . . 102

7.12 Bimanual Interactions – (a) In the steering wheel example, squeeze  
 increases when torques applied by the wrist oppose each other and  
 decreases when torques are applied in the same direction. (b) For the  
 bow-and-arrow, both wrists experience increasing squeeze as the bow  
 string is drawn and an instant release when the arrow is fired. Subtle  
 vibrations while drawing increase realism. . . . . 103

7.13 Mutli-Paradigm Interactions – These interactable objects combine many  
 of the rendering approaches presented so far into a single context. Despite  
 squeeze feedback being generally uniform for each point of interaction,  
 the illusion generally holds in the context of the user’s motion and the  
 rendered visuals. . . . . 104

8.1	Mitsuda showed there to be a linear relationship between pressure applied to the forearm and the sense of weight. Figure adapted from [158]. . . . .	106
8.2	(a) Samad et al. tasked users with adjusting the weight of a physical cube (blue) to a C/D modulated cube (brown). Changing the C/D ratio of the user's hand height in VR (b) altered the perceived mass of the cube (c). Figure adapted from [148]. . . . .	108
8.3	The estimation of an environment variable $S$ can be assumed to be the integration of the unimodal visual and haptic estimators $S_V$ and $S_H$ by the maximum likelihood estimate rule. Figure adapted from [161]. . . . .	109
8.4	Variable Stiffness Button (VSB) and Calibrator – (a) Image of the apparatus. (b) The apparatus in calibration mode, which serves to calibrate Tasbi's internal force sensor to the user's arm. (c) The apparatus in button mode, which was used as the physical stiffness comparison in Experiment 1. . . . .	111
8.5	Tasbi Calibration Procedure – (a,b) Tasbi is mounted to the subject's wrist using 3M mounting tape and Transpore medical tape. (c) The VSB calibrator applies a known load profile through Tasbi to which Tasbi's internal sensor is calibrated against. . . . .	112
8.6	Experiment 1 – The virtual buttons and environment as viewed from the subject's VR perspective. The left button was the real physical button, and was spatially aligned with the physical button apparatus. The right button was the mid-air button, rendered through Tasbi squeeze force and/or C/D manipulation. Depending on the current condition, buttons were either visually opaque with observable displacement, or they were transparent with no observable displacement, acting instead as a volume. . . . .	113

8.7	Experiment 1 Main Results – The top row shows the mean response and fits for individual subjects as well as the group level fit for each condition. The two identified outliers are denoted by dashed lines. Error bars are a 95% confidence interval. The bottom row provides box plots of subject means for every level in each condition, with the same outliers denoted by hollow circles.	116
8.8	Four metrics comparing the fits obtained from Experiment 1 across conditions. Error bars represent a 95% confidence interval. Significance obtained from the bimodal-unimodal t-tests are denoted with bars above the plots.	118
8.9	Six individual subjects from Experiment 1 shown in comparison to the group fits. The order in which they completed condition blocks is given in the bottom right corner. Error bars represent one standard deviation in the subject's responses for that level.	120
8.10	Experiment 2 – Subjects were presented with two mid-air buttons and forced to choose the one they thought was stiffer. Subjects were constrained to two button presses or five seconds to make their selection. The same mid-air button conditions from Experiment 1 were used: Visual only, Haptic only, or the combination of the two, Haptic-Visual.	125
8.11	The estimated stiffness of the Haptic and Visual conditions from Experiment 1 were used as a proxy to derive perceptually equivalent stimuli levels for Experiment 2. Eleven comparison levels for each condition were interpolated based on a stiffness range from 50 to 150 N/m. The bold lines indicate the central Standard level.	128
8.12	Aggregate psychometric curves obtained from Experiment 2. The Bimodal condition was found to produce a significantly lower JND compared to with Visual only condition, but not compared with Haptic only condition.	130

8.13	Four Metrics Comparing Conditions of Experiment 2 – (a) The mean just noticeable difference threshold for each condition. (b) The point of subjective equality, or bias from the Standard level of 100 (solid black line), for each condition. (c) The percentage of times subjects correctly selected the stiffer button. (d) The average number of button presses taken in each condition, with 800 being the maximum (solid black line). Error bars denote a 95% confidence interval. . . . .	132
8.14	Six subjects from Experiment 2. Their thresholds for the unimodal Visual and Haptic conditions are $T_V$ and $T_H$ . The empirical and MLE predicted thresholds for the Haptic-Visual condition are $T_{HV}$ and $\hat{T}_{HV}$ . Estimated unimodal weights are $w_v$ and $w_h$ . . . . .	135
8.15	Survey Results for Experiment 2. Asterisks indicate significant difference from a neutral response of 0.5. See Table 8.3.4 for the full questionnaire. . . . .	136
8.16	Psychometric functions from all subjects in Experiment 2 overlaid in each condition. The conformity of subjects in the bimodal condition is similar to the observations of Experiment 1 (see Fig. 8.7). . . . .	137
9.1	The controller-free, finger tracked hands of AR will present new opportunities for Tasbi research. (a) Fine motor tasks tasks such as playing a piano will require explorations into how to appropriate squeeze and vibrotactile feedback to multiple and simultaneous finger interactions. (b) The absence of controllers will remove inherent tactile feedback user receive when holding them, such as for the virtual bow interaction. However, this issue is less concerning once we realize that devices like Tasbi can turn essentially any inanimate object (e.g. this frivolous banana example) into a haptic feedback controller. . . . .	145



9.2	Integrating Tasbi with other technologies presents interesting opportunities. (a) Devices like Tasbi can provide haptic feedback for ordinary capacitive displays. (b) Emerging technologies, such as brain-machine interfaces and electromyography (EMG) might enhance the believability of Tasbi-mediated interactions in XR.	. . . . . 147
9.3	Technology developed for Tasbi is being freely shared with research communities. (a-b) Syntacts vibrotactile amplifier kits that were distributed during an IROS 2020 workshop. (c) A vibrotactile-only clone of Tasbi we have made available online. (d) A fellow lab member's wearable haptic display made possible by Syntacts and Tasbi research.	. . . 148

## Tables

<b>2.1 Vibrotactile Bracelets</b> . . . . .	13
<b>2.2 Squeeze Bracelets and Armbands</b> . . . . .	16
<b>5.1 Comparison of Tactor Control Methods Tested</b> . . . . .	72
<b>8.1 Survey Questions for Experiment 2</b> . . . . .	128

# Chapter 1

## Introduction

Within the last decade, advances in optics and displays, tracking, environment mapping, computer graphics, and audio have revolutionized technologies for extended reality (XR). XR, an umbrella term referring to all combinations of real and virtual environments and human-computer interaction, encompasses both entirely *virtual* reality (VR) as well as *mixed* and *augmented* realities (MR/AR) which merge real and digital information. Recent growth in XR can largely be attributed to the emergence of consumer grade VR head-mounted displays (HMD) developed for gaming in the 2010s. Since then, the scope of XR has exploded, with applications now spanning education, marketing, remote work, and training for medicine, industry, and military [1]. Enormous resources are being invested into XR, and spending on research and development is estimated to rise from \$4 billion in 2018 to \$121 billion by 2023. Adoption of XR technology is rapidly growing as well, with consumer spending expected to surge from \$5 to \$40 billion in that same time frame [2].

XR has the potential to dramatically change the way we interact with both the physical and digital world, as well as other people. It is widely believed that all-day wearable XR devices will reinvent computer interfaces in ways that rival the smartphone and personal computer. With this, companies such as Microsoft, Apple, Google, and Facebook are currently racing toward consumer grade XR glasses. As we move toward a future where seeing and hearing virtual entities is commonplace, we must also consider another important sensory aspect: our ability to feel and touch objects that do not actually exist. Thus, *haptic* technology is yet another major area of focus in XR development.

## 1.1 State of the Art in XR Haptic Feedback

Haptic feedback, or technology that recreates the sense of touch, spans the disciplines of robotics and human perception. Decades of research and development have given rise to numerous haptic technologies ranging from simple one degree-of-freedom (DOF) actuators that vibrate the skin, to electrostatic surfaces that render texture to the finger, to highly articulated robotic manipulators and exoskeletons that transmit kinesthetic forces to the hands or limbs. Haptic devices have been extensively studied for the purposes of enhancing realism in virtual environments [3], closing the action-confirmation loop in user interfaces [4], improving the efficacy of training programs [5], and providing feedback to augment or substitute for other senses [6]. For these reasons, the addition of haptic feedback to XR is extremely appealing, and many approaches have been taken thus far [7].

### 1.1.1 Handheld Controllers

Handheld controllers are by far the most hardware for transmitting haptic feedback in XR interactions (Fig 1.1). Currently, most commercial VR HMDs (Oculus [8], HTC Vive [9], Valve Index [10], etc.) ship with controllers that provide basic haptic feedback in the form of vibration. If the user receives a virtual notification, the controller might vibrate. If a user touches a virtual surface, the controller vibrates. If the user picks up an object, the controller vibrates. Clearly, vibrating controllers can provide only a limited experience for the type of interactions users may encounter in an XR context. Some researchers have also investigated other haptic primitives in handheld devices, such as force feedback [11], [12] for grasping and stiffness rendering, and configurable textures for surface rendering [13]. Though handheld controllers certainly have their place in short-spanned contexts such as gaming, their compatibility with all-day extendable XR glasses is limited since they encumber the hands and prevent concurrent interaction with the physical world.

### 1.1.2 Haptic Glove Devices

To expand the capability of rendering forces to the hand, significant research efforts have been invested in developing wearable haptic gloves and exoskeletons [14, 15, 16, 17] (Fig 1.1). While implementations vary, most haptic gloves involve applying kinesthetic forces to the joints of the fingers. For instance, if a user grabs a virtual object, the glove may restrict or impede finger movement to simulate touching the object's surface. Devices that target only the palm have also been explored [18]. In contrast to handhelds, only a few haptic glove systems have emerged on the commercial market, and due to their high costs, they are typically reserved to industrial sectors [19]. Glove systems greatly enhance the realism of virtual interactions compared to handheld controllers, but they do not necessarily decrease the burden on end-users. Limitations in actuation technologies result in many of these devices being bulky, motion restrictive, and encumbering. Novel actuation techniques (e.g soft-actuators [20]) may ultimately lead to compact and unintrusive designs, but the implementation of all-day wearable haptic gloves seems to be in the distant future.

### 1.1.3 Haptic Fingertip Devices

Noting the current limitations of full-hand glove systems, many researchers have investigated the prospect of simplifying designs to the finger tip only [21]. Typically, such devices render either vibrotactile or cutaneous cues to the pads of the finger. Examples include fabric displays to convey object softness [22]; compressing and shearing finger belts for contact forces [23]; and finger-pad skin stretch for directional forces [24, 25, 26]. Designs that ground reaction forces between the fingers (e.g. the index finger and thumb) simulate grasping force feedback [27, 28]. Overall, these devices often blend the simplicity of handheld controllers with the fidelity of haptic gloves. However, as can be seen in Fig. 1.1, fingertip devices can still be quite encumbering and unsuitable for extended periods of use.

### Handheld Controllers



Facebook (Oculus), 2019



Valve Corporation, 2019



Lee et al., 2019



Choi et al., 2018

### Haptic Glove Devices



HaptX Inc., 2017



CyberGlove Systems, 2007




Trinitatova et al., 2017



Polygerinos et al., 2015

### Haptic Fingertip Devices



Leonardis et al., 2015



Shorr et al., 2017



Bianchi et al., 2017



Pacchierotti et al., 2016

Figure 1.1 : Haptic devices for XR interaction can be broadly categorized as either (1) handheld controllers, (2) wearable glove devices, or (3) wearable fingertip devices.

## 1.2 Upper Limb-Grounded Haptic Devices

Limitations of the haptic devices discussed thus far include: (1) difficulty in ergonomically grounding forces to the body; (2) the need for several actuators on the hand or complex actuation routing schemes; (3) high power requirements; (4) unacceptable encumbrance; and (5) prevention of the hands from interacting with physical world objects. Each of these points is detrimental to realizing an all-day wearable haptic interface for XR interaction, and thus more practical devices are desired.

Recently, the field of haptics has begun to take a more pragmatic approach to device design. Researchers are now investigating ways of creating compelling experiences with low-cost actuators and novel modalities in *limb-grounded* formats. These generally encompass devices such as bracelets, armbands, and sleeves. Limb-grounded wearables can be generalized to many different tasks and applications. They have been successfully utilized for navigation [29, 30], prosthetics feedback [31], and text communication and notifications [32]. Depending upon the implementation, a single limb-grounded wearable could conceivably offer not only all of these capabilities, but much more. It is not hard to imagine a future where limb-grounded haptics are tightly integrated with other mobile technologies such as smartphones and smartwatches. *It is even plausible that an all-day wearable limb-grounded device could provide substantive feedback for XR interactions.* If XR glasses one day emerge as an all-day wearable technology, then an equally all-day wearable limb-grounded device could serve as a companion interface. However, before such a device can be realized, important questions and design choices must be considered.

### 1.2.1 Where should an all-day wearable device be worn?

The wrist is the sensible choice as an all-day wearable device. Although the design space is more restricted than that of an armband or sleeve, technology-integrated wristbands and

bracelets are well established commodities, socially acceptable [33], and sometimes even fashionable [34]. Importantly, wristbands leave the hands free, which is important for AR and MR as they allow the hands to manipulate the physical world unhindered. Furthermore, a precedent for haptic bracelets already exists. Several commercial wrist devices now incorporate basic haptic feedback, including the Apple Watch [35] which provides simple vibrations for notifications, and several recently developed *social touch bracelets* can effectively communicate emotional states (e.g. HEY Bracelet [36] and Bond Touch [37]). These devices can only offer limited experiences in XR due to the simplicity of their feedback modalities.

### 1.2.2 Which haptic modalities are both compelling and viable?

The majority of today's wearable haptic devices continue to leverage simple vibrotactile feedback. This is an obvious choice since vibration is a ubiquitous feedback primitive, and vibration actuators are usually inexpensive, low power, and easily controllable. For this reason, multi-actuator vibrotactile *arrays* are common in haptic research devices. More exotic skin-contact related modalities also exist, including squeeze [38], stretch [39], shear [40], twist [41, 42], thermal [43], and electrical [44] stimulation. There is strong evidence to support the *multimodal* combination of one or more of these modalities with vibrotactile feedback, and this is perhaps one way in which wristband devices can be made more applicable to XR interactions. Multimodal haptics offer the ability to target different mechanoreceptors in the skin, enabling higher rates of information transfer to users [45] or more realistic simulation of virtual events [46, 47]. However, it is not practical for a device to implement all modalities, nor advisable since they are easily confused [48, 45]. Squeeze, stretch, shear, and twist are all similar in that they apply localized topical pressures to the skin. The latter three offer bidirectional stimulation, which in theory could provide an advantage to



navigation or tasks requiring large cue sets. Unfortunately, these three modalities require consistent skin contact through friction [39] or adhesives [41], a matter which is complicated by skin moisture, the environment, and prolonged use. Thus, squeeze appears to be the most practical choice and is certainly the most researched. Squeeze feedback is thought to be less attention demanding than vibration [49], provides more intimate feedback similar to how one human might attract the attention of another [50, 38], and may elicit affective or emotional responses [51]. Squeeze has additionally been identified as a strong candidate for providing proportional kinesthetic information (e.g. grasp force [40, 52]).

### **1.2.3 Are compact wrist devices possible with current technology?**

The primary limiting factors of any wearable haptic device are usually actuators and sensors, whether in regard to their size, weight, cost, or power consumption. With the exception of vibrotactile feedback, all modalities discussed thus far typically rely on either bulky servos or pneumatic systems. In addition to size concerns, the power requirements for these types of actuators are often high, which poses a significant challenge to ultimately operating entirely on battery power. Engineers will need to resort to clever and efficient mechanisms to compactly integrate current actuators until the day that more ideal actuators and materials are readily available [53]. As this thesis will show, the control of these actuators is not a trivial matter either, as variations in limb geometry, tissue impedance, and posture degrade the consistency of haptic feedback. Advanced sensing capabilities can mitigate these issues, but potentially increase the cost and size of the device. While there are other important aspects not discussed here (e.g. wireless communication), clearly there are already significant challenges to realizing compact haptic wearables.




Figure 1.2 : Referred haptic feedback could enable a variety of interactions with the virtual world. Vibrotactile feedback may substitute for fingertip contact with virtual buttons or other types of user interfaces, while directional forces may convey the weight or stiffness associated with manipulating virtual objects. Referred haptic feedback can further add to telepresence and remote social interactions such as hand shaking or holding, and also provide immersive feedback for gaming and training environments.

### 1.3 Referred Haptic Feedback for XR Hand Interactions

Assuming a compact, multimodal haptic bracelet is even feasible, how might we leverage its capabilities to provide haptic feedback for XR interactions? The majority of work surrounding haptic bracelets and armbands has focused on delivering notifications or providing state information to enhance user performance in various tasks. While these paradigms certainly have their place within an XR context, an equally compelling application is to employ *referred feedback* for interactions with virtual objects. In this scenario, we may render substitutive haptic feedback through a bracelet interface for forces that would otherwise be expected at the hands and fingertips (Fig. 1.2).

Referred haptics have been extensively explored for providing force-feedback to myoelectric prosthesis users [54], where the haptics may convey information such as hand aperture or grip force [55, 31, 56]. Users have shown a surprising ability to reinterpret force feedback delivered to other parts of the body, thanks largely in part to the plastic nature of the brain [57]. Referred haptic feedback for virtual interaction is essentially a similar problem, yet inverted. *Instead of providing feedback for an artificial hand interacting with a real environment, we seek to provide feedback for a real hand interacting*

*with an artificial environment.* The case of referred haptic feedback for virtual hand interaction is arguably more difficult to realize. First, XR users are probably less motivated to interpret such feedback than impaired individuals seeking to regain their sense of touch and autonomy. The second difficulty is that non-impaired individuals using these devices will continue to use their hands to interact with physical objects. Context switching between real interactions and referred feedback mediated virtual interactions may simply be irreconcilable if the feedback is not intuitive.

The concept of applying referred haptic feedback for XR hand interactions has yet to be explored in depth. With the exception of publications that comprise parts of this thesis [58], [46], [59], we are only aware of two other instances. Moriyama et al. [60] presented a device that applied pressure to the forearm to convey virtual forces arising from finger interactions. Works-in-progress from Sarac et al. have investigated normal and shear forces at the wrist for virtual interaction using haptic sketches [61] and an actuated wrist band [62]. Both Pezent [58] and Sarac [62] showed that feedback delivered to the wrist enhanced user's ability to perceive stiffness when pressing virtual objects. Tinguy et al. demonstrated a similar idea [63], but with a more proximal finger worn device, thus we do not consider this referred haptic feedback. Obviously, a significant amount of work must still be performed to fully understand the potential of referred haptic feedback in XR interaction.

## 1.4 Thesis Contribution and Outline

This thesis presents significant efforts in realizing a compact wearable bracelet for referred haptic feedback in XR. We introduce **Tasbi** (Tactile And Squeeze Bracelet Interface), a multimodal haptic wristband (Fig. 1.3) that integrates traditional vibrotactile actuation with robust wrist squeeze in an extremely compact form factor to enable XR hand interactions. In Chapter 2, the current landscape of haptic bracelets is surveyed and critical




Figure 1.3 : This thesis presents Tasbi, a haptic bracelet that provides referred squeeze and vibrational feedback for hand interactions in virtual environments.

points that led to Tasbi’s development are discussed. An in-depth look into the bracelet’s design is provided in Chapter 3, with special emphasis on a novel squeeze mechanism that affords Tasbi high force output in a relatively small package size. In Chapter 4, we discuss methods and challenges to squeeze control, and demonstrate for the first time accurate and high-performance closed-loop control of wrist squeeze using inexpensive capacitive force sensors. Chapter 5 addresses control of Tasbi’s vibrotactile elements, which ultimately resulted in the distribution of an open-source software and hardware framework called Syn-tacts. Leveraging Tasbi’s force control abilities, we conduct novel psychophysical studies in Chapter 6 that highlight user sensitivity to wrist squeeze as well as its potential impact on vibrotactile identification. Chapter 7 explores Tasbi’s potential to provide referred haptic feedback for a number of unique VR hand interactions using multisensory approaches. Finally, in Chapter 8, we investigate the possibility that referred haptic feedback is more than just a metaphor for interaction forces, and instead invokes genuine perceptions of object stiffness. My concluding thoughts and vision for the future are provided in Chapter 9.

## Chapter 2

### Haptic Bracelets and Armbands

In Chapter [1](#), we postulated that the wrist might be the ideal location for an all-day wearable haptic interface. A number of possible haptic modalities were discussed, and we identified vibration and squeeze as being the most promising options available. We further hypothesized that the multimodal combination of squeeze with a vibrotactile array could provide the diversity of feedback required for XR interactions. This Chapter surveys the field of vibrotactile, squeeze, and multimodal bracelets. In the latter two cases, we extend our survey to additionally include devices for the forearm and bicep, as they are common. The insights presented here ultimately drove many decisions during the design of Tasbi.

#### 2.1 Vibrotactile Bracelets

Generating vibration cues is most often accomplished through the use of small vibration motors (vibrotactors or tactors). Types of vibrotactors include eccentric rotating masses (ERM), linear resonance actuators (LRA) or voice coils, or piezo actuators (see Chapter [5](#) for more details). Vibrotactors are commonly found in handheld devices such as controllers and smartphones. Many researchers have investigated placing vibrotactors on the wrist and arm as well [\[64\]](#). Because the wrist and arm provide ample surface area, the most interesting work usually employs *vibrotactile arrays*, i.e. two or more independently controlled vibrotactors, as a means of increasing information transfer. Two possible configurations exist – planar grid arrays or radial arrays.

Oakley et al. [65]. studied a 3x3 planar configuration on the dorsal side of the wrist. The results of their experiments showed that subjects more easily localize vibrations when they are presented perpendicular to the axis of the arm and inline with the wrist strap as opposed to along the length of the arm. They further showed that localization is improved by placing vibrotactors near bodily landmarks, such as the edges of the arm. Using a similar 3x3 planar configuration, Chen et al. [66] compared placing the tactors on the dorsal versus the volar side of the wrist. Their results agreed with the findings of Oakley's study with regards to direction and bodily landmarks, but they did not find a significant difference in localization between the dorsal and volar sides of the wrist.

Matscheko et al. [67] compared arranging four tactors in a planar grid on the dorsal wrist versus radially around the wrist. They showed that subject performance was best for the radial configuration in a memory and distraction task, and logically concluded that this was the result of spreading the tactors further apart. Following their advice, Carcedo et al. [68] tested variations of a band with 3, 4, 5, and 6 radially spaced tactors. The results showed identification rates above 90% for the 3 and 4 tactor configurations, and around 80% and 70% for the 5 and 6 tactor configurations, respectively. Gupta et al. [69] developed a device with four radial voice coil actuators to enhance manipulation feedback of a touch screen surface, and Pece et al. [70] have developed a variation of voice coil technology that indents instead of vibrates the skin. The designs presented so far did not attempt to isolate vibration transfer between adjacent motors, which probably has a non-trivial impact on localization accuracy. Hong et al. [71] addressed issues of vibration transfer by separating radially spaced tactors with thin elastic thread. They concluded that in this configuration, up to eight tactors can increase accuracy in a guidance task. All devices are summarized in Table 2.1 and Fig. 2.1, and in general, it would seem that a safe bet for vibrotactile wrist bands is to incorporate at least four tactors in a radial configuration.

Table 2.1 : Vibrotactile Bracelets

	Actuators	Configuration
Oakley, 2006 [65]	9 LRA	Planar
Chen, 2008 [66]	9 LRA	Planar
Lee, 2010 [72]	3 LRA	Planar
Matscheko, 2010 [67]	4 VC	Planar/Radial
Carcedo, 2016 [68]	3-7 ERM	Radial
Gupta, 2016 [69]	4 VC	Radial
Hong, 2016 [71]	4/8 ERM	Radial
Pece, 2017 [70]	4 VC (indenting)	Radial

ERM = ecentric rotating mass, LRA = linear resonant actuator, VC = voice coil




Figure 2.1 : Vibrotactile Bracelets – (a-c) Planar vibrotactile grid arrays for the wrist. (d-g) Radial vibrotactile arrays for the wrist. (h) A slightly different but notable indentation-based tactile wristband.

While identification accuracy is indeed important for discrete information transfer, we should not necessarily use these results as a driving factor when designing bracelets for XR interaction. For one, vibrotactors are inexpensive and consume small amounts of power, and including more than perceptually identifiable is not of high concern – we can always use a subset of the available tactors when identification accuracy is needed. The reason we may choose to include redundant tactors is to achieve smoother transitions when presenting *continuous spatial patterns*, i.e., blending the vibrations of adjacent tactors to create the illusion of motion, a topic that is discussed further in Chapter 5. One unresolved question from these studies regards the importance of mechanical coupling between tactors and skin, and whether localization accuracy improves and degrades with increasing coupling force. This thesis presents some insights toward this in Chapter 6.

## 2.2 Squeeze Bracelets and Armbands

While vibration has been extensively studied for decades, squeeze or compression feedback has only become of interest to the haptics community within the past few years. There are many reasons for investigating squeeze. First, while vibration primarily excites Pacinian nerve endings, squeeze can innervate the slower-adapting Merkel and Ruffini receptors [73, 74, 75]. Zheng et al. argued that squeeze, compared to vibration, is less attention demanding and more appropriate for ambient feedback [49]. Along these lines, Baumann et al. [38] suggested that squeeze provides intimate feedback that more closely resembles social touching behaviors, and Tsetseruko et al. [51, 76] have further used squeeze to elicit affective emotional responses from users. Important to XR applications, squeeze has been demonstrated to be more appropriate for providing *continuous* feedback, in contrast to the discrete alert-type buzzes vibration offers. As such, squeeze has been used to communicate grasping forces for prosthetic [55] and teleported applications [77].



Though some squeeze devices have been developed specifically for the wrist, many have been deployed to the forearm or bicep (Table 2.2). Most squeezing devices employ electromechanical servos or DC motors to tension flexible bands around the limb [78, 43, 77, 79], and are characterized by generating both *normal* and *tangential* forces on the skin. Noting that tangential forces provide another haptic modality, i.e. skin stretch, some designers have included two or more actuators so that squeeze and stretch can be actuated independently and simultaneously [40, 52, 80]. However, Zook and Fleck have demonstrated that squeeze and stretch cues perceptually interfere with one another and are easily confused by users [48, 81]. Consequentially, some devices attempt to isolate squeeze to only normal forces, and have used clamping mechanisms [82], linkage-based mechanisms [60], and linear actuators [62] to do so.

As can be seen in Fig. 2.2, most electromechanical squeezing devices are large in size. Gupta et al. [83] addressed the size concerns of the aforementioned devices by employing shape-memory alloys (SMA), but this approach required high power (up to 30 W) and insulation to shield users from heat. Others have resorted to pneumatic based compression. Pohl et al. [75] used pneumatically actuated bladders to create uniform compression akin to a blood pressure cuff. Young et al. [84] used eight small inflatable bellows to provide targeted squeeze around the wrist. Other instances of pneumatic actuation include Pneusleeve [85], WRAP [86], and a device from Payne et al. [87]. While these devices are both sleek and likely more comfortable due to their soft designs, pneumatic actuation is currently limited by bulky compressor technology. Some researchers have developed modular micro-compressors [88, 89] to power pneumatic devices, often worn on the waste or in a backpack. However, it seems very unlikely that consumers would be willing to wear a separate power source for an all-day wearable wristband. Thus, a traditional electromechanical approach will still be required for the foreseeable future.

Table 2.2 : Squeeze Bracelets and Armbands

	<b>Form</b>	<b>Method</b>	<b>Actuators</b>	<b>Sensors</b>	<b>Control</b>
Stanley, 2012 [78]	Wrist	Belt	1 servo	None	Position
Brown, 2017 [77]	Wrist	Belt	1 servo	None	Position
Bianchi, 2014 [80]	Wrist	Belt	2 motor	2 encoder	Position
Song, 2015 [43]	Wrist	Belt	1 servo	None	Position
Meli, 2018 [52]	Forearm	Belt	4 servo,	None	Position
Wang, 2012 [79]	Bicep	Belt	1 servo	None	Position
Casini, 2015 [40]	Bicep	Belt	2 motor	2 encoder	Position
Treadway, 2015 [55]	Bicep	Belt	1 servo	None	Position
Gupta, 2017 [83]	Wrist	SMA	1-3 SMA	None	Current
Chinello, 2014 [82]	Forearm	Clamp	3 servo	3 FSR	Force
Moriyama, 2018 [60]	Wrist	Normal	4 servo	None	Position
Sarac, 2020 [62]	Wrist	Normal	2 linear	2 encoder	Position
Pohl, 2017 [75]	Wrist	Pneumatic	Custom	None	Pressure
Raitor, 2017 [86]	Wrist	Pneumatic	4 regulator	None	Pressure
Young, 2019 [84]	Wrist	Pneumatic	8 regulator	None	Pressure
Payne, 2018 [87]	Bicep	Pneumatic	6 regulator	2 FSC	Force
Zhu, 2020 [85]	Forearm	Pneumatic	1 regulator	1 FSC	Force

SMA = shape memory alloy, FSR = force sensing resistor, FSC = force sensing capacitor



Figure 2.2 : Squeeze Bracelets and Armbands — (a-g) Electromechanical, band-based devices. (h) A device employing shape memory alloys. (i-k) Devices that squeeze only in the normal direction. (l-p) Pneumatic based squeezing devices.

An often overlooked facet of electromechanical squeeze feedback is the control implementation. It is typical for these devices to employ servo position control of the band-tensioning actuator. Thus, the amount of squeeze force delivered is inherently tied to the tissue impedance of the stimulus site. Sarac et al. [62] noted the limitations of this approach

in their study for VR interactions, and attempted to resolve the issue by estimating applied force from the average impedance of hairy-skin [90]. Control of contact forces is generally limited to pneumatic actuation where bladder pressure, and thus contact pressure, is controlled for using open-loop control. Closed-loop control on the forearm has been accomplished with custom force sensors as well [87, 85]. However, because pneumatic bladders necessarily change shape and size, there still remains some ambiguities as to what these devices are actually controlling for from a *perceptual* standpoint. Another challenge to force-control is in knowing how much force is sufficient in the first place. All psychometric analysis of wrist squeeze, thus far, has been quantified in units that are inherently tied to the device (e.g. the amount of motor displacement, band-tension, or regulator-pressure). This thesis presents answers to these unresolved issues in Chapters 4 and 6.

### 2.3 Multimodal Bracelets and Armbands

Most devices discussed so far were developed for notification type feedback, and thus can only offer limited experiences for XR hand interactions. One way in which wristband devices can be made more generalizable is by enabling *multimodal* feedback, or more specifically, integrating *both* vibration and squeeze into a single interface. Combining squeeze and vibration could not only provide a richer cue set with higher information throughput, but also the ability to convincingly depict virtual interactions that are inherently multimodal. This concept has been most prevalent is in the development of glove interfaces for XR applications, where kinesthetic mechanisms and vibration actuators are combined to convey more realistic interaction (see Chapter 1).

Despite a high volume of research into vibration and squeeze feedback wearables alone, a relatively limited number of devices have implemented both modalities into a single wearable (Fig. 2.3). Bauman et al. [38] developed a multimodal wrist device which




Figure 2.3 : Multimodal Bracelets and Armbands – (a) Baumman’s device integrated servo-based wrist squeeze and tapping. (b) The MISSIVE device combined servo-based squeeze and stretch with four voice coil vibrotactors. (c) Aggravi’s device combined servo-based squeeze with four ERM vibrotactors.

featured squeezing and low frequency tapping for emulating human attention getting practices. Dunkelberger et al. developed the MISSIVE [32], which combines separate bands for vibration and squeeze. They successfully used the device to convey language through skin, and argued that this multimodal approach allowed them to render a larger set of distinguishable cues than would have been possible under a unimodal approach. Aggravi et al. presented a forearm device [91] that incorporated squeeze and four factors into the same band, but did not address the fact their design causes factors to translate on the skin, which likely confuses users. None of these devices managed to achieve a design robust or compact enough for all day wearability. Finally, while Zook [48] showed that squeeze interferes with skin stretch, it is not currently known if the same is true for squeeze and vibration. Thus, this thesis also attempts to answer this question in Chapter 6.

## Chapter 3

### Design of Tasbi Haptic Bracelet

This chapter details the mechanical and electrical design of Tasbi (Tactile and Squeeze Bracelet Interface), a multimodal bracelet that combines vibrotactile feedback with radial squeeze haptics [58]. Tasbi's design spanned approximately two years of iteration and prototyping (Fig. 3.3) before arriving at the final compact design shown in Fig. 3.1.

#### 3.1 Design Criteria

Tasbi's design process was driven by several considerations and constraints. Our primary goal was to build a highly compact unit with minimal sacrifices to actuation output and bandwidth. Based on commercial smartwatches, our target size was 50x50x10 mm with a total mass less than 200 grams. We decided the device should emit little audible noise to avoid annoying users or interfering the haptic experience. To accommodate various modes of squeeze control, integrating both position and force sensing were desired. Actuator power consumption was constrained 2 W, so as to not dissipate an uncomfortable amount of heat and to reasonably remain within the capabilities of lithium-ion batteries. With that, it is important to note upfront that we did *not* aim to completely self contain this iteration of Tasbi. Therefore, Tasbi does not include onboard microcontrollers, batteries, or wireless communication, and relies on an external control unit. We felt that it was more important to focus on miniaturizing the mechanisms and sensing for the initial prototype, and address these concerns in a future iteration.



Figure 3.1 : The final iteration of Tasbi as worn on the wrist. The device incorporates six vibrotactile actuators and a novel squeezing mechanism. Not pictured is the tethering cable required to power and with interface.

### 3.2 Squeeze Mechanism

As already discussed, most squeezing devices use a similar scheme where one or more rotational actuators are used to directly wind a band element into an actuator housing [38, 78, 40, 43, 77, 52]. While this approach is straightforward, it presents two main issues (Fig. 3.2-a). First, directly tensioning the band itself gives rise to an unequal distribution of forces where there are concentrated tangential shear forces on the sides of the arm, and smaller normal forces on the underside. Furthermore, this results in non-trivial squeeze




Figure 3.2 : **(a)** Typical constricting-band approaches to squeeze produce non-uniform and tangential forces which would cause embedded factors to shift. **(b)** Our decoupled approach aims to produce pure, uniform normal forces.

force losses due to friction between the band and skin. Second, because this method causes the band to translate along the skin, it is not well suited to embedding vibrotactile elements in the band since they would consequently translate too. Maintaining the radial positions of the vibrotactors is key since their movement would decrease user identification rates and possibly cause discomfort. Some devices have circumvented this issue by using two separate bands: one for generating squeeze, and one for housing vibrotactors [45]. However, this approach is less than ideal for a wrist-watch form factor and complicates donning and doffing the device.

Tasbi solves these problems by decoupling squeeze actuation from the wrist band and vibrotactors. This is accomplished by means of small diameter, flexible ultra high molecular weight polyethylene (UHMWP) cord, (trade name Dyneema/Sprectra) which wraps circumferentially around the exterior of the band (Fig 3.2-b). Tensioning this cord, not the band, creates squeeze forces. Friction is minimized by separating contact between the cord and band with smooth, polished steel pins placed directly above each vibrotactor. This mechanism results in cord tension being transmitted as an inward force approximately nor-






Figure 3.3 : (a) A rudimentary bracelet was developed to test the proposed pin and cord approach to squeeze. (b) The first prototype of Tasbi was 3D printed and provided only open-loop torque control capabilities. (c) Tasbi v1.0 improved the design with metal reinforced components and integrated encoder feedback for position control of the tensioning spool. (d) Tasbi v2.0 added an improved HDMI connector interface and integrated force sensing to enable closed-loop control of normal squeeze forces.

mal to the vibrotactor, as can be understood through simple geometric inspection. Because friction between the pin and cord is small, little tangential force is transmitted to the band, and as a result each vibrotactor maintains its radial position around the circumference of the wrist. Eliminating friction and tangential forces also allows for a smaller actuator, since most power is ideally converted to purely normal squeeze force. Furthermore, tensioning the lightweight cord instead of the entire band means less mass must be moved to accomplish similar displacement, further reducing the power required of the tensioning actuator.

A rudimentary prototype of the squeeze mechanism and haptic sketching [92] suggested that approximately 10 N of tension would be required to achieve an appropriate range of squeeze stimuli. Several tensioning mechanisms and actuators were initially considered. For ease of implementation and control, an electromechanical approach, as opposed to pneumatic or other exotic approaches, was decided. Linear actuation methods were disregarded since achieving a stroke length necessary to generate enough cord takeup would necessitate a non-ideal housing length. For this reason, a rotary scheme with a winding

spool was chosen. Many hobbyist servo motors met our size and torque requirements, but generally these actuators produce a high degree of audible noise and do not provide continuous rotation. We therefore chose to use a brushed DC motor. An additional reduction stage was required to meet our torque needs. Commercially available gear units, which typically implement one or more serial stages of planetary gears, were found to also suffer audible noise issues and were unavailable in sufficiently compact sizes.

Our final solution utilizes a 12 mm Maxon DCX [93] motor coupled to a 100:1, 13 mm strain-wave gear unit from Harmonic Drive [94]. These drives offered a set of characteristics we deemed necessary to fully realize Tasbi: (1) a sufficiently high torque reduction, (2) low audible noise due to having zero mechanical backlash, and most importantly (3) compactness far superior to conventional gear units. The DC motor and Harmonic gear unit are contained within a  $50 \times 30 \times 15$  mm housing (below our target size) which rests on the dorsal side of the user's wrist. Attached to the output of the gear unit is a 10 mm diameter spool. Both ends of the cord terminate to the either side of the spool so that the take-up rate is doubled (anecdotally, it is possible to double cord tension at the expense of half the take-up rate by fixing one end of the cord to the housing, achieving a pulley-like effect). The cord is redirected internally over additional smooth pins to exit at the center of the main housing, balancing a moment arm that would otherwise cause the housing to torque about its short axis.

The DC motor is driven by an externally located 4-quadrant PWM servo controller (Maxon Escon 24/2 [93]) operating in a current control mode. Position estimation is achieved via incremental encoder feedback placed on the motor side of the mechanism. To maintain a small footprint, we used an ultra miniature optical encoder (Elesta E OI R016 [95]) featuring a reflective mirror code wheel, with the optoelectronic sensor PCB embedded into the rear connector panel of the housing. With 128 counts per revolution




Figure 3.4 : Tasbi Exploded View – The squeeze mechanism consists of a 12 mm DC motor (a) and a 13 mm Harmonic Drive gearbox (b) which drives a two-sided spool (c) to create tension in a UHMWP cord (d). Spool position feedback is provided through an optical encoder consisting of a reflective code wheel (e) and optoelectronic sensor (f). Squeeze force feedback is measured via a force-sensing capacitor (g) and signal conditioning PCB (h). The force sensor is held in light compression against a bottom plate (i) with four corner springs (j). The drive assembly drops into the main housing (k) and is secured in place with a housing lid (l). Each vibrotactor unit (m) contains a 10 mm LRA vibrotactor and a smooth stainless steel pin to convert cord tension into normal force. Vibrotactor units are clipped into elastic sidings (n) and secured with lids (o). All signals and power to and from Tasbi are transmitted over a micro-HDMI cable that connects to an internal breakout PCB (q).

and a 100:1 reduction, Tasbi can achieve  $0.007^\circ$  positional accuracy at the spool output in 4X quadrature mode. This level of accuracy is unnecessary for spool positioning, but provides smooth velocity estimation which was ultimately critical to our squeeze control implementation (see Chapter 4).

It is worth noting that the drive mechanism is not easily back-driven due to its high gearing ratio. While this does present mild safety concerns (i.e. users cannot manually loosen the device with ease), it means that Tasbi can maintain varying levels of squeeze without continuous input from the motor. For example, the motor can be used to squeeze to a desired level, turned off, and then friction in the gear unit will hold the squeeze level. This property may eventually prove beneficial to self-contained versions of Tasbi and similar devices where battery power must be conserved.

Readers familiar with tendon routing mechanism might understandably be skeptical to the long-term viability of the cord-based mechanism presented here. Indeed, early versions of Tasbi [58] were prone to internal snagging and spool disconnection. Subsequent revision eliminated these issues. Though time will tell, eight Tasbi prototypes have been fabricated and deployed at Rice University and Facebook Reality Labs, and no failures have been reported. One unit in particular has been used for well over a year for demonstrations and over 100 hours of device and subject experimentation without issue.

### 3.3 Squeeze Force Sensing

Enabling Tasbi with squeeze-force sensing was one of the more challenging aspects of the design. To simplify the process, we made the assumption that placing a single sensor in the main housing would accurately capture gross squeeze force around the entire wrist. This assumption can be made due to the way in which Tasbi equally distributes forces (see Fig. 3.2-b), which is experimentally shown in Fig. 4.2.




Figure 3.5 : Tasbi Force Sensor – A SingleTact 8mm, 10 N force-sensing capacitor is in sandwiched between the main housing and a plastic bottom plate. The plate is held in compression through four compression springs and screws so that the sensor experiences minimally pre-load.

Because a significant portion of the main housing was already consumed by the tensioning mechanism, integrating a traditional single-axis load-cell was ruled out due to space constraints. Therefore, we only considered force sensors with a thin-film or sheet-like form factor. Force-sensing resistors (FSR) are a technology that typically falls under this category, but are notorious for exhibiting signal drift and hysteresis issues. A similar, yet relatively newer technology, force sensing capacitors (FSC), has been shown to provide improved sensitivity and repeatability compared to FSRs [96] at the expense of requiring more sophisticated electronics and signal conditioning. The most readily available FSCs are the SingleTact sensors from Pressure Profile Systems (PPS).

Tasbi incorporates an 8 mm diameter, 10 N SingleTact FSC sensor, located between the underside of the main housing and a bottom plate. Ideally, the entire surface between the housing and pressure plate would have been covered with a force sensing layer so that all force is transmitted through the sensor. However, these sensors were only available in a small circular form factor. Two important design choices were made to ensure that the majority of force seen at the skin interface is transmitted through the sensor head. First, the bottom plate is held in compression with four springs and screws located at each corner of

Tasbi's housing. This allows the plate to be carefully fastened until just a slight amount of pre-load is measured by the sensor. Second was the design of the bottom plate itself, which (1) has an extended contact surface to mitigate issues that arise with skin deformation under load, and (2) is made from a flexible ABS plastic with intentional material removal to give it a leaf-spring like characteristic. The SingleTact electronics PCB is integrated directly into Tasbi's housing, and transmits the force measurement as an analog voltage through Tasbi's connector interface. Overall, our experience with SingleTact's FSCs has been good. The sensors display some manageable hysteresis, but no significant drift. Interestingly, their sensitivity is good enough that it is possible to detect some users' heartbeat from the force measurement when Tasbi is tightened beyond 5 N of normal force.

### **3.4 Vibrotactile Band**

Tasbi's wrist band contains six vibrotactor units. Each unit consists of a plastic housing for the vibrotactor. The vibrotactors are generic 2.5 VAC, 10 mm linear resonant actuators (LRA) with a nominal frequency range of 150 to 200 Hz. Because each tactor's performance is sensitive to fitting tolerances, we used two O-rings to achieve a snug but not overly tight fit into the assembly. The tactor is secured axially with a foam layer and the housing lid. Along the underside of the lid is a press-fit hole for one of the aforementioned smooth pins. The distance from skin to the pin was optimized so that the tensioning cord would clear and not rub against the user's skin.

Each tactor unit is clipped in between polyurethane rubber sidings via T-shaped joints. Tactor power cables are embedded within the rubber siding and enter the main housing through openings on both sides. The elasticity and geometry of the sidings, inspired by the commercial Myo armband, allows the band to be stretched over the user's hand during the donning and doffing process while also reducing vibration transfer between adjacent tactor




Figure 3.6 : Tasbi Vibrotactor Assembly – The tactor is secured into the lower plastic housing via two O-rings and a foam layer. The metal lid contains a press fit smoothpin over which the tension cord slides. Tactor power cables are embedded within the 3D printed band sides in a post-processing step.

units, similar to the device in [71]. Based on available anthropometric data [97, 98], Tasbi has an nominal inner circumference of approximately 150 mm, equal to the 50th percentile female wrist circumference. Thus for most users, the band provides a light amount of passive squeeze to ensure a comfortable initial fit.

Vibrotactor control is accomplished using the Syntacts vibrotactile framework [59], which leverages audio interfaces to control high-density tactile arrays with low latency. In our implementation, we interface the framework with a MOTU 24Ao sound card [99], which connects to a Syntacts amplifier board housed in the Tasbi control unit. Using Syntacts, it is possible to generate a wide variety waveforms, both discrete and continuous, that can be played on individual or multiple tactors. Possible waveforms can be composed from simple oscillators (e.g. sine, square, saw), amplitude and frequency modulation, amplitude shaping envelopes, and track-based sequences. Syntacts also provides a spatializer mode that treats the tactors in Tasbi’s band as a continuously space where amplitudes of adjacent

tactors are blended. To achieve the lowest latency possible ( $< 4$  ms), we use Steinberg's ASIO sound driver [100] with Syntacts. We point the reader to Chapter 5 for more details regarding the Syntacts' framework and amplifier board design.

### 3.5 Power and Control Unit

With multiple actuators and sensors, Tasbi requires a relatively complex control architecture. As previously mentioned, all power and control is done through a custom external unit. This control unit houses linear DC power supplies (so to avoid noise associated with common switch-mode power supplies), a PWM servo controller for the DC motor, and a multi-channel Syntacts amplifier [59] for the vibrotactors. Each Tasbi connects to its own control box via a standard micro-HDMI cable (an unconventional choice, but one of the few interfaces providing the necessary pin count and form factor). With the exception of vibrotactor input signals, all digital and analog signals between the control box and the host PC are done over a Quanser Q8-USB [101] sampled at 1000 Hz.






Figure 3.7 : The Tasbi control box interfaces each Tasbi with the host PC via intermediate connections to a Quanser Q8-USB digital acquisition device and a MOTU 24Ao audio sound card.




Figure 3.8 : Tasbi fabrication. (a) Each Tasbi is comprised of over 150 parts that require hand assembly. (b) A vibrotactile band assembly before and after factor enclosure. (c) A tensioner housing assembly before enclosure (d) Completed Tasbi v2.0 units. Eight units have been fabricated for internal uses and to share with collaborators.

## Chapter 4

# Squeeze Control

In this Chapter, we present and discuss the development of Tasbi’s squeeze feedback controller. The majority of similar devices leverage squeeze for general purpose cues, notifications, or predefined effects [43, 45, 85]. With Tasbi, we additionally aim to accommodate virtual hand and finger interactions in AR/VR (see Chapter 7), which requires more sophisticated real-time control. As such, the following requirements were specifically set forth during the development of Tasbi’s squeeze controller:

- **High Dynamic Range:** The controller should be able to produce low and high amplitude output to convey the wide range of forces that arise from virtual interactions.
- **Fast Response Times:** The controller should be robust to the unpredictability of user interaction that present scenarios where squeeze must rapidly change from low to high output in a moment’s notice.
- **Accurate Tracking:** The controller should be able to accurately track continuous, real-time inputs.
- **Smooth Operation:** The controller should be perceptually free of mechanical noise so that users do not confuse squeeze for vibration stimuli.
- **Consistent Stimuli:** The controller should produce perceptually equivalent stimuli for a given input regardless of hand posture or wrist impedance.

The following sections extend the work of [58], and detail the progressive process of developing and characterizing a closed-loop squeeze controller which fulfils these requirements.

## 4.1 Torque Control

At the lowest level, squeeze can be produced through simple open-loop torque control of Tasbi’s tensioning motor using only the ESCON servo-controller in current control mode. We first characterized Tasbi’s squeezing capabilities as a function of motor torque. We created a fixture containing two ATI Nano17 F/T sensors (Fig. 4.1); one remained fixed under the main housing for referencing, and the other was able to relocate under any factor housing. The device was stretched over the fixture with the band center axis oriented upward so that gravity would not affect normal and tangential force measurements. The tensioning mechanism, initially with the cord loose, was commanded to step to a certain percentage of the max motor torque, hold for three seconds, and then return to the loose position. Force measurements were taken underneath the main and factor housings in both the normal and tangential directions. This procedure was repeated ten times for each of ten torque levels from 10% to 100% of the maximum motor torque (3.21 mNm). The full test was repeated for each of the six factor housings.

Fig. 4.2 shows the force response under a representative factor for each torque level. For torques above 50% there is a noticeable relaxing effect, most likely due to the material properties of the plastic housings and UHMWP tensioning cord. Torque levels below 20% produce little to no force output, revealing some dead band in the tensioning mechanism due to friction in the drive components. Fig. 4.2 shows the “steady state” (i.e. the mean of the last half-second of the responses in Fig. 4.2) normal and tangential forces under each factor. Importantly, we can see that there is negligible tangential force, with the force distribution being almost entirely normal and thus satisfying our design goals. The more




Figure 4.1 : (a) A 3D printed rig was designed to characterize Tasbi's distribution of squeeze load. A fixed ATI Nano17 (1) force sensor measure squeeze force under the main housing, while a relocatable Nano17 (2) allowed for measuring squeeze force under each tactor independently. (b) A Tasbi stretched onto the force rig for characterization. (c) The forces under the housing and each tactor were characterized. Colors correspond with Fig. [4.2](#).

proximal tactor housings have a higher normal force than the distal tactor housings, which is likely due to cord tension drop-off between adjacent tactors as a result of pin friction. There is also some bias, with left-side normal forces being higher, but generally we see a linear region from 20% to 70% torque.

However, Fig. [4.2](#) also illustrates two challenges for open-loop torque control: 1) torque levels below 15% max torque produce little to no force output, revealing dead band in the tensioning mechanism due to friction in the drive components; and 2) torques above 50% display a noticeable relaxing effect, most likely due to the material properties of the tensioning cord.

Even if these issues could be overcome, torque-only control is fundamentally flawed because it presents no means to deescalate squeeze force due to non-monotonic behavior. While it is possible to produce increasing levels of squeeze force by ramping torque, it is not possible to reverse force by decreasing torque because wrist impedance is incapable of backdriving the tensioning mechanism. Thus, a closed-loop control scheme is required.




Figure 4.2 : (a) Steady state normal vs. tangential forces under each vibrotactor module and the main housing as a function of commanded motor torque. (a) Representative responses for each torque step under the main housing.

## 4.2 Position Control

Our first approach to closing the loop on squeeze was through position control of the tensioning spool. Such a controller is trivially implemented with feedback from the incremental encoder on the motor side and a proportional-derivative (PD) control law. In this mode, a range of squeeze stimuli is produced by controlling the squeeze mechanism between a minimum and maximum spool position. The range is determined in situ by recording the

steady-state position when open-loop torque is held at the minimum level (i.e. 15%) and the maximum level (i.e. 100%). This range was different for individuals, but typically between 40° and 60° of rotation.

Because motor position control is simple and practical, this is the most common approach used by servo-actuated squeeze devices. Indeed, we have made abundant use of position controlled squeeze with Tasbi, and the controller works well enough for demonstrations and short-lived uses. However, a number of issues plaguing position controlled squeeze make it unsuitable for long term and general use. First and perhaps most importantly is that controlling for spool position offers no means to provide a consistent perceptual stimulus across users because it is inherently coupled to the impedance of the user's wrist. For example, 30° of spool rotation likely feels different for a person with toned wrists than it does for a person with soft wrists. Additionally, it can also feel different to an individual if the bracelet shifts along the arm, which is unavoidable. The issue is exacerbated by drift and/or creep in the cord tensioning mechanism, such that a given spool rotation does not produce the same amount of cord deflection over time. Finally, wrist impedance changes drastically with hand and finger posture, which negatively impacts the performance of position-controlled squeeze. In the left column of Fig. 4.3, Tasbi is commanded to track a sinusoidal position trajectory between 0° and 60° as an individual cycles through various wrist poses. Although nearly perfect positioning tracking is maintained, the actual amount of delivered squeeze force changes significantly.

### 4.3 Force Control

Based on the aforementioned issues with position-based control, it is clear that squeeze should be controlled through a variable directly related to the contact mechanics between the bracelet and skin. Two choices include the amount of skin indentation or the applied force

or pressure. Because it is not yet clear which stimulus is more perceptually important at the wrist, and given the complexities of implementing sensors for the former, we chose to




Figure 4.3 : The impedance of the wrist changes considerably with posture. Here, measurements of force and spool position are shown as a user transitions between wrist extension and flexion while either a sinusoidal position (left) or force trajectory (right) is controlled for. The top two plots clearly show we cannot assume that squeeze force will remain proportional to spool position. Though position control is accurately maintained during hand movement, the delivered squeeze stimulus changes drastically (bottom left). Thus, we require direct control of squeeze force (bottom right).

explore direct control of the contact force between Tasbi and the skin. The implementation of force-based squeeze was considerably more involved than the previous methods, and so a detailed description of our approach follows.

#### 4.4 Sensor Calibration

Prior to implementing any sort of force control, the force sensing capacitor in Tasbi needed to be accurately calibrated post-installation to account for (1) sensor offset from the center-line of the main housing, and (2) non-negligible force leakage through the pressure plate compression springs. To this end, we fabricated two apparatus to perform sensor calibrations, both endogenously and exogenously (Fig. 4.4).

The first apparatus was an instrumented wrist cross-section with an integrated ATI Nano17 transducer that measured force along the axis perpendicular to Tasbi’s underside (Fig. 4.4-a). To reasonably simulate tissue mechanics, a 5 mm thick, molded silicone “skin” layer surrounds the outside of the wrist. We used SmoothOn Ecoflex silicone (#00-30 Shore hardness), which has seen widespread use in simulating tissue for medical training and research [102].

To perform the calibration, the ATI Nano17 was first zeroed without any externally applied load. Next, a Tasbi was slipped over the instrumented wrist and tightened to a base level of squeeze by setting the motor torque to 15% of its maximum. After the motor quit spinning, Tasbi was zeroed for position, switched into the position-control mode and commanded to track a compound sinusoidal trajectory for 10 seconds. Force measurements from the instrumented wrist sensor and Tasbi’s internal sensor were recorded. The force data were related with a second-order polynomial, which provided a quality fit with  $R^2$  typically greater than 0.95. Fig. 4.5 illustrates the calibration process and the accuracy of Tasbi’s internal sensor after the fit is applied.






Figure 4.4 : Apparati – (a) The instrumented wrist integrates a high-accuracy ATI Nano17 force sensor and silicone-based simulated tissue. It served a critical role during the force controller development and tuning phases. (b) While general purpose calibrations of Tasbi’s integrated force sensor can be obtained with the instrumented wrist, an actuated calibrator facilitated customized calibrations to individuals by applying a known load through the bracelet. It was primarily used prior to the psychophysical experiments in Chapter [6](#) and [8](#) to ensure accurate reporting of force.




Figure 4.5 : Tasbi’s internal force sensor is calibrated against the instrumented wrist’s Nano17. The left shows the voltage-force data and fit, and the right shows the Tasbi’s measured force after the fit is applied compared to the Nano17.

The instrumented wrist provides a general purpose calibration that can be associated with the device and subsequently used on individuals as-is. However, minor error in the calibration can be expected considering each individual’s wrist deforms slightly differently under load. Thus, to uphold the integrity of the psychophysical experiments presented later in Chapter 6, we developed an alternative calibration apparatus to generate user-specific calibrations in situ. With this procedure, Tasbi’s force sensing capacitor is calibrated to a known-load transmitted through Tasbi’s housing and onto the wrist.

## 4.5 Controller Implementation

With a calibrated force measurement from Tasbi, we next began controller design. The primary challenge during the design process was in overcoming the low-precision and noisy analog signal from Tasbi’s internal SingleTact force sensor. The noise seen on the sensor is largely due to high frequency interference from the PWM motor driver signal, though some inherent noise is associated with the sensors as well. Although we could have taken physical corrective actions (e.g. linear drivers, improved cable shielding, or leveraging the

sensor's I2C interface), we chose to attempt a software-only solution.

Much of the controller development took place with Tasbi placed on the instrumented wrist. Because the force reported by the instrumented wrist and Tasbi's internal sensor are well correlated post-calibration, we found it particularly useful to first design the controllers by closing the loop with the much higher quality instrumented wrist force signal, and then apply the prototype controller to Tasbi's force measurement (Fig. 4.6). Three separate controllers were developed and tested.




Figure 4.6 : Three controllers were tested by first closing the control loop with feedback from the Nano17 sensor, then applying the controller to Tasbi's sensor and scaling the controller gains to stability. A feedforward and PD controller with the derivative term conditioned on *motor velocity* offered the best performance.

Taking the simplest approach first, we applied a PD control law similar to what was done with the position controller:

$$e(t) = F_{ref}(t) - F_{act}(t) \quad (4.1)$$

$$\tau = K_p e(t) + K_d \frac{de(t)}{dt} \quad (4.2)$$

where  $F_{ref}(t)$  is the desired force,  $F_{act}(t)$  is the actual force measurement from either the instrumented wrist of Tasbi,  $e(t)$  is the force error,  $K_p$  and  $K_d$  are the proportional and derivative control gains, and  $\tau$  is the torque to be commanded to motor. Fig. 4.6 illustrates the design process. The controller was first roughly tuned using the instrumented wrist force measurement as the process variable (first column), then switched to use to the force measurement from Tasbi (second column). Clearly, the PD controller suffers when using the noisier feedback. The controller was stabilized by reducing the gains and filtering the force measurement with a median filter, but the tracking accuracy was limited, with the controller being incapable of reaching peak forces (third column).

The simple PD controller was next modified to include a feedforward term proportional to the desired reference force:

$$\tau = K_p e(t) + K_d \frac{de(t)}{dt} + K_{ff} F_{ref}(t) \quad (4.3)$$

The intuition here stems from knowing the relationship between squeeze force and torque (Fig. 4.2). Thus, we can predict the amount of torque required to generate a particular force and supplement the feedback partition with a portion of it. As shown in the second row of Fig. 4.6, the feedforward term greatly enhances the tracking accuracy of the simple PD controller. However, unacceptable tracking errors were still present. Ideally, we would

have simply increased the value of the proportional gain  $K_p$  to eliminate the remaining tracking errors, with complementary increases to  $K_d$  to maintain stability. Unfortunately, since derivative action is taken on the backwards differentiated force error, this approach was to susceptible to noise to be viable.

If we maintain that the role of the derivative term is simply to dampen the action of the proportional term (which we wish to increase), we need not constrain the controller to using the derivative of force error, and can substitute it with a less noisy signal that is also proportional to the rate of squeeze. Thus, our final controller replaces the derivative term with *spool velocity*:

$$\tau = K_p e(t) - K_{d,v} \frac{d\theta(t)}{dt} + K_{ff} F_{ref}(t) \quad (4.4)$$

With this simple modification, we can stably increase  $K_p$  to the point of eliminating tracking errors, provided an appropriate value of the new derivative gain  $K_{d,v}$  is set. In addition to eliminating tracking error, the herein referred to as PD<sub>v</sub>+FF controller also provides significantly smoother operation since it is not ridden with noise from the error derivative. Theoretical analysis on the stability of this “hybrid” control method is outside of the scope of this chapter, but in practice it has proven to be quite stable even at high frequency.

## 4.6 Controller Characterization

After refining the gains of PD<sub>v</sub>+FF controller and the median filter window width, the overall controller displayed a root-mean-square error (RMSE) of 0.48 N while tracking a 10 N, 1 Hz sinewave (< 5% error), and a rise time of 70 ms for a 10 N step response (Fig. 4.7). The controller was subsequently validated across a variety of VR applications and on a number of individuals displaying different wrist characteristics. The controller also

performed well regardless of the sensor calibration scheme, i.e., on the the instrumented wrist of in situ with the calibration applicator.

To further quantify the performance of the force controller, we benchmarked it against the position controller under the assumption that the latter drives the system near peak performance. Tasbi was placed on the instrumented wrist and commanded to track a si-




Figure 4.7 : Top - The  $PD_v+FF$  controller was further refined to display RMSE less than 5% for a sinusoidal trajectory, and a rise time of 70 ms. Bottom - The final controller was validated in a VR context. Here, the controller renders the interaction force as a user jiggles a virtual button (see Chapter 7).




Figure 4.8 : Comparison of the position and force controllers. Comparable ranges of squeeze and position are first identified (a), and then the controllers are separately commanded to track an excitation signal within their determined range (b). Although the force controller exhibits more phase lag (d), its overall bandwidth defined by the 3 dB cutoff is a quite comparable 9.1 Hz (d).

nusoidal force trajectory between 0 and 10 N. During this motion, the motor position was observed to consistently follow a trajectory between  $0^\circ$  and  $40^\circ$  (Fig. 4.8-a). We then performed a closed-loop system identification for both controllers, where the controller was commanded to track a 10 s Schroeder multisined excitation [103, 104] signal displaying a frequency range from 0 to 15 Hz (Fig. 4.8-b). The amplitude of the force trajectory was from 0 to 10 N, and the amplitude of the position trajectory was from  $0^\circ$  to  $40^\circ$ . Thus, both tests produced equal amounts of squeeze per Fig. 4.8-a.

Fig. 4.8-b shows both controllers attenuating near the 5 s mark. The Bode diagram in Fig. 4.8-c,d shows that the force and position controllers display a surprisingly comparable bandwidth of 9.1 Hz and 10.3 Hz, respectively. The phase diagram in Fig. 4.8-d shows that the force controller lags considerably more than the position controller. This is not particularly surprising given the amount of filtering required to sufficiently smooth the force sensor signal.

The right column of Fig. 4.3 shows the extent to which the controller rejects external disturbances, where accurate tracking is maintained as the user cycles through various wrist orientations. More importantly, in contrast to the original position controller, the force controller provides a consistent stimulus regardless of wrist orientation or impedance.

## 4.7 Discussion and Future Improvements

The  $PD_v+FF$  squeeze force controller works considerably well given the simplicity of both the sensors used and control law. Nonetheless, some limitations and room for future improvements remain. The primary challenge resides in accurately and consistently estimating the force at the bracelet-wrist interface. Despite the good fit of the sensor voltage to measured or applied force, we find that this curve does shift slightly across several calibrations, particularly if Tasbi is not worn in the same location. We also observe that



calibrations are not perfectly consistent across different individuals, thus necessitating the apparatus in Fig. 4.4-b. Ideally, we would like to eliminate this process, and perform only a single device-specific calibration. Our design made use of inexpensive and commercially available sensors, and so a custom force sensing capacitor that offers full coverage of the contact surface without force leakage and a higher resolution could significantly enhance the estimation of contact force. Other obvious improvements include better signal conditioning and electrical shielding. It is also worth mentioning that it may be possible to leverage different techniques of contact force estimation entirely, such as photoplethysmogram signals [105].

Room for a more sophisticated control law exists as well. Although our final controller improves force control using supplemental state information from the optical encoder, we suspect that further improvements could be made using more a rigorous fusion of sensor data. For example, if a dynamic model of wrist impedance could be formulated, then samples from the force-position distribution (Fig. 4.3 and Fig. 4.8-a) and/or the torque-force distribution (Fig. 4.2) combined with Kalman filtering might provide a more accurate and smoother estimation of squeeze force.

# Chapter 5

## Vibrotactile Control

In this Chapter, we present the methods by which Tasbi’s six linear resonant actuator (LRA) vibrotactors are controlled. During the course of development, it was realized that that our unique approach had the potential to be used in a number of applications well outside the scope of Tasbi. Thus, efforts were placed in generalizing the implementation, which ultimately led to an open-source software and hardware packaged called Syntacts. Syntacts has been freely distributed online and through workshops, and has been well received by the haptics community. Portions of this chapter have appeared in [59] and we acknowledge the contributions of our co-author Brandon Cambio [106].

### 5.1 Background

One of the most important and ubiquitous modes of haptic feedback is vibration. Haptic vibrations are commonly delivered to users through small actuators known as vibrotactors, or simply tactors. Vibrotactors come in many forms, such as eccentric rotating mass (ERM) actuators, linear resonant actuators (LRA), voice coil actuators, and Piezo actuators. For several decades, vibrotactile feedback has been used extensively across a wide variety of applications, most notably mobile and wearable devices [107].

The modern era of vibrotactile research is faced with a number of new needs and requirements. For instance, the field has recently begun moving away from providing users with simple alert type cues to delivering salient cues rich in information. Many

researchers are now designing devices with larger numbers of tactors integrated into single interfaces such as bracelets, armbands, and sleeves [68, 58, 108], full body suits and clothing [109, 110], and chairs [111]. Unfortunately, driving many vibrotactors simultaneously has traditionally been a difficult task for engineers and non-engineers alike due to the technical skills required, interfacing difficulty, or cost of equipment. Further, high-density arrays require more sophisticated rendering algorithms. *Spatialization*, or the manipulation of several actuators in an array-based on the placement of a virtual target location, has been explored to some extent [111].

In addition to increasing actuator counts, some vibrotactile research has recently focused on delivering complex vibration waveforms, beyond simple buzzes, to convey more meaningful information to users [112], or to more accurately simulate real-world phenomena [113]. The synthesis of such cues, however, is not a trivial task, with some researchers resorting to pre-recorded libraries [114] or high-level creation tools [115, 116]. Finally, while the advent of mainstream virtual reality (VR) systems has introduced new opportunities for vibrotactile feedback, it has also imposed additional constraints on control including low latency [117] and the need to alter cues on the fly in response to virtual events.

This chapter aims to highlight a method of vibrotactor control that accommodates many of these requirements and deserves detailed attention: *control through digital audio interfaces*. We present a new open-source software and hardware package, **Syntacts**, that lowers the technical barrier to synthesizing and rendering vibrations with audio. In Section 5.2, we discuss common vibrotactor control schemes along with their advantages and shortcomings. Section 5.3 provides an overview of the hardware requirements for audio-based control, underscoring some of the lesser known details that can have a high impact on control, and introduces the Syntacts Amplifier board. In Section 5.4, we discuss software for audio-based control and then present the Syntacts software library. Finally, in Section 5.5,

we provide comparisons between Syntacts-based audio control and other methods. Conclusions and areas for future work follow in Section 5.6. Syntacts software and hardware designs are freely available at: [www.syntacts.org](http://www.syntacts.org).

## 5.2 Introduction to Vibrotactor Control

Because vibrotactors have been a staple of haptics for a long time, there exist many scenarios and approaches for their control. A typical *research-oriented* scenario requires controlling vibrotactors from a PC that may also coordinate an experiment, record data, and/or render visuals. Within this context, we summarize a few possible control strategies.

### 5.2.1 Function Generators

The simplest implementation uses a standalone function generator connected directly to the tactor. This is easy because generators are purpose-built to output oscillating signals and envelopes, and can often meet the tactor's power requirements. However, function generators are limited in cue design, channel count, and may be challenging to integrate with custom software. For these reasons, they are a poor choice for complex control.

### 5.2.2 Integrated Circuits

To serve the mobile device market, specialized integrated circuits (IC) have been developed for vibrotactor control. These ICs often handle both signal generation and power amplification, making them an all-in-one package. A common chip, the DRV2605L from Texas Instruments (TI) [118], features a built-in library of effects that can be triggered and sequenced through I<sup>2</sup>C commands. Some ICs are capable of closed-loop control which automatically detects the tactor's resonant frequency and can provide faster response times. The utility of ICs for laboratory research, however, is restricted by the need to design and

fabricate custom PCBs, since their small package sizes make it difficult to prototype on breadboards (though preassembled PCBs and breakouts can be found in various online shops). Controlling many factors becomes complicated and usually requires additional components such as multiplexers. Finally, PCs generally do not provide an I<sup>2</sup>C interface, so a USB adapter or microcontroller (e.g., an Arduino) must be introduced to broker communication between the PC and ICs.

### 5.2.3 Dedicated Controllers

Unlike other actuators such as DC motors, there exist very few off-the-shelf, plug-and-play controllers for vibrotactors. One product marketed as such is the Universal Controller from Engineering Acoustics, Inc (EAI) [119]. It is designed to drive their ubiquitous C2 and C3 voice coil actuators, but can drive other factors with similar load impedance. The controller interfaces to a PC via USB and can output up to eight individual channels, though the datasheet and our own testing (Section 5.5) indicates that only four can be driven simultaneously. EAI provides a GUI and C API with adequate cue synthesization features, so integrating the controller with custom software is straightforward. The major downside of this controller is a very high upfront cost (approximately \$2,250) that not all researchers are willing or able to afford.

Texas Instruments also sells the DRV2605LEVM-MD, an evaluation module for the DRV2605L mentioned above, that could be considered a controller unit. The board integrates eight DRV2605L ICs, an I<sup>2</sup>C multiplexer, and a USB interface. Unlike the EAI controller, no high-level communication API is available, so either low-level serial programming or I<sup>2</sup>C brokerage is still required to integrate it with a PC. Finally, a recent startup, *Actronika* [120], aims to sell a haptic processing unit, the Tactronik; however, details are currently sparse.

## 5.2.4 Audio Output Devices

Another approach to driving tactors, and main focal point of this chapter, is through *digital audio output devices*. This approach hinges on the understanding that some vibrotactors, particularly LRA and voice coil variants, operate very similarly to headphones or loudspeakers. Like speakers, these tactors consist of an electrical coil within a magnetic field. Energizing the coil induces a magnetic force that, in the case of speakers, drives a cone to generate sound pressure, or, in the case of vibrotactors, drives a mass to generate vibrations. As such, the same hardware that drives loudspeakers can also drive vibrotactors with a few adjustments and considerations. Figure 5.1 provides a high-level overview of the audio-to-vibration rendering pipeline.

The technique of using audio to drive haptic actuators is simple yet relatively underutilized within the field. Outside of a few workshops [121, 122], the process has received limited documentation or comparison with existing control solutions. The remainder of this chapter will discuss the implementation of audio-based control while introducing a new open-source hardware and software solution, **Syntacts**. We will show that using this approach can provide a number of benefits including relatively low implementation cost, support for large channel counts, and ultra-low latency.

## 5.3 Hardware for Audio-Based Control

### 5.3.1 Sound Cards / Digital-to-Analog Converters

The most important piece of hardware for audio-based control is the digital-to-analog converter (DAC) device. The DAC is responsible for converting digitally represented waveforms, like music files, to analog signals to be played through headphones or speakers. Virtually all PCs have a DAC integrated into the motherboard that outputs two analog sig-




Figure 5.1 : Pipeline for vibrotactile control through audio interfaces. A host PC (a) runs the target application or virtual environment and interfaces with an audio rendering layer (b). The audio renderer interfaces with a digital-to-analog converter (DAC) sound card to output low voltage audio signals (c). The signals are power amplified (d) and sent to the vibrotactors (e) where they are felt as vibrations. Syntacts provides both the audio rendering layer (b) and amplifiers purpose made for driving vibrotactors (d).

nals through a headphone or line out jack (typically a 3.5mm phone jack) for left and right audio channels. If no more than two vibrotactors are needed, use of the built-in headphone jack may be sufficient for some users.

Driving more than two channels generally requires a dedicated DAC, or sound card. The least expensive options are *consumer grade* surround sound cards, which can be had in typical PCI-e or USB interfaces. Up to six tactors can be driven with 5.1 surround sound cards, while up to eight can be driven with 7.1 surround sound cards. We have found this to be a viable solution if consideration is given to differences between channel types (e.g., subwoofer channels are usually tuned for lower impedance loads than speaker channels). Offerings from Creative Soundblaster and Asus are among the most readily available choices. There also exist *professional grade* audio interfaces with more than eight outputs, such as the MOTU UltraLite-mk4 and 16A with 12 and 16 channels, respectively. For even higher channel counts, the purely analog output MOTU 24Ao is a popular choice [123, 124]. A single unit provides 24 output channels, and up to five units can be connected using Audio Video Bridging (AVB) to drive 120 vibrotactors if desired. It should be noted that some professional devices may feature other I/O channels (e.g., MIDI, S/PDIF, etc.)

that are of little use for driving factors.

An *extremely* important consideration in sound card selection is the device's driver API support. An API describes a digital audio transmission protocol, and most drivers support many different APIs. Windows standardizes at least four first-party APIs: WDM-KS, WASAPI, MME, and DirectSound. As shown in Fig. 5.2, not all APIs are created equally. Because MME, which exhibits highly perceptible latency, is usually the default API, it could be easy to conclude that audio is insufficient for realtime haptics. Steinberg's third-party ASIO driver is widely considered to be the most performant option, but it is often only implemented by professional grade equipment. Regardless, API selection is a rather opaque setting under Windows, and appropriate software is usually required to select the preferred driver API (see Section 5.4). Driver API selection is less of an issue on macOS, with CoreAudio being the universally recommended option. Another important consideration is audio buffer-size, or the number of audio samples sent on every transmission to the device. If the host PC has sufficient processing speed, smaller buffer sizes should be preferred for low latency (Fig. 5.3).

### 5.3.2 Amplifiers

Audio DACs typically output a low-power signal at what is called "line-level" because they expect that the output device will amplify the signal before it is actually played. Vibrotactors are similar to typical 8 to 16  $\Omega$  speakers, and therefore require amplification. Amplifiers are divided into different classes based on how they operate. *Digital Class D* amplifiers are the most common. They expect an analog input signal and output an amplified version of the signal with pulse-width modulation (PWM). This type of amplification tends to be very power efficient, but high-frequency PWM switching can add large amounts of electrical noise to a system. This is especially true when designing for arrays of vibro-






Figure 5.2 : Mean Windows audio driver API latencies with standard deviation. Data collection methods are described in Sec. 5.5. For reference, the dashed line indicates the perceptual threshold of visual-haptic simultaneity [117].

factors, where multiple naively implemented Class D amplifiers can create enough noise to be physically felt. Class A, B, and AB amplifiers are *linear* amplifiers. These amplifiers tend to have much lower efficiency than the Class D, which can lead to heat problems if their thermal design is overlooked. However, because they do not constantly switch at high frequencies, they introduce considerably less noise into the overall system. Finally, a stable power supply is critical to the amplifier's ability to condition the signal. Batteries or linear power supplies provide much more stable power than typical switch-mode power supplies and allow amplifiers to operate with less noise.

Noisy power amplification can have detrimental effects on the performance of haptic devices that integrate sensors. For example, the first iteration of Tasbi's [58] tactor control hardware featured three commercial stereo Class D amplifiers powered by a generic switch-mode power supply. The high-frequency content emitted by these components resulted in




Figure 5.3 : The effect on latency due to changing audio buffer sizes.

errant motor encoder readings and noisy analog force sensor measurements beyond usability. As another example, we have noticed considerable noise emission from the C2 tactors and EAI Universal Controller (which also uses switching amplifiers) in MISSIVE [45] during EEG measurements.

### 5.3.3 Syntacts Amplifier

Based on these difficulties and limited commercial options for high-density output, we designed the purpose-built, eight channel **Syntacts Amplifier** board (Fig. 5.4). It is based on the TI TPA6211A1-Q1 3.1W audio power amplifier IC, featuring a Class AB architecture and fully differential inputs and outputs that together eliminate all noise issues we have experienced with commercial options. The Syntacts amplifier can drive small to medium sized vibrotactors with load impedances above  $3 \Omega$  from a 5V power supply at typical vibrotactile frequencies, making it suitable for many applications (Fig. 5.5). We have successfully tested it with various LRAs, EAI's C2 and C3 voice coil actuators, and Nanoport's TacHammer actuators. The amplifier is not intended for use with ERM actuators, which are trivially powered with DC voltage, nor Piezo actuators, which require higher voltages

or custom controllers altogether. The Syntacts amplifier has thermal and short circuit protection and operates at voltage levels generally considered safe. However, potential users should understand that it has not undergone the testing required of commercial devices, and should take this into their safety considerations.




Figure 5.4 : The Syntacts amplifier is an open-source fully differential, linear amplifier capable of driving eight vibrotactors with minimal noise. Two variants are available: one with a single AES-59 DB25 input for connecting to high-end audio devices such as the MOTU 24Ao, and one with four standard 3.5 mm TRS headphone inputs for connecting to general audio outputs or surround sound cards. Both require a 5V power supply, and output amplified signals through a universal 0.1” pitch header.




Figure 5.5 : The Syntacts amplifier can be used in a variety of applications, ranging from dense tactile arrays **(a)** to wearable devices such as Tasbi **(b)**. Designs for the tactile array are available online as a reference implementation.

Open-source designs for two variants of the amplifier, one with four 3.5 mm phone inputs and one with a standardized AES-59 DB25 connector, are available online along with manuals and data sheets. We provide packaged CAD files and BOMs for direct submission to turn-key PCB manufactures, where the boards can be built for roughly \$50-100 USD depending on the quantity ordered and requested fabrication time. Alternatively, the PCB and components can be ordered separately and soldered by hand or in a reflow oven.

## 5.4 Software for Audio-Based Control

Software is necessary both to interface audio devices and to synthesize and render waveforms. Many commercial GUI applications provide these features for the creation of music and sound effects. While some researchers have leveraged such software (particularly MAX MSP [122]) for vibrotactor control, they tend to be overly complex, lack features useful for haptic design, and are difficult to integrate with other applications programmatically. Though a number of haptic effect software GUIs and frameworks have been developed for commercial [125] or one-off, custom hardware [126], only a few examples of general purpose, audio-based vibrotactor software exist. One example is Macaron [115], a WebAudio-based online editor where users create haptic effects by manipulating amplitude and frequency curves. The software, however, is primarily focused on ease of design, and provides little in the way of device interfacing or integration with other code.

To fill this void, we developed **Syntacts**, a complete software framework for audio-based haptics. Driven by the needs of both Tasbi [58] and MISSIVE [45], we have integrated a number of useful features, including:

- a user-friendly API that integrates with existing code
- direct access to external sound card devices and drivers

- flexible and extensive waveform synthesis mechanisms
- the ability to generate and modify cues in realtime
- spatialization of multi-channel factor arrays
- saving and loading cues from a user library
- compatibility with existing file formats and synthesizers
- a responsive GUI for cue design and playback

Each point is further detailed in the following sections. Syntacts is completely open-source, with code and binaries for Windows and macOS freely available at: [www.syntacts.org](http://www.syntacts.org).

#### **5.4.1 Syntacts API**

Syntacts' primary goal is to provide a flexible, *code-oriented* interface that can be easily integrated with existing software and applications. The library is written in C and C++ to facilitate accessing low-level drivers and maximizing performance. Additionally, bindings are currently provided for C# and Python. The former is particularly useful for integrating Syntacts with Unity Engine for creating 3D virtual environments, while the latter allows for high-level scripting and interactivity (e.g., with Jupyter notebooks). Integration with other languages is possible via C shared library (i.e., DLL) loading, and additional languages may be officially supported in the future (e.g., a MATLAB interface would be useful to many academics). Code presented in this section is taken from the Python binding, but the native C++ API and C# binding are similar in their syntax and usage.

## Interfacing Devices

Syntacts will interface with virtually any audio card on the commercial market. The API allows users to enumerate and select devices based on specific drivers, a feature typically reserved to professional commercial software. While Syntacts can open devices under any audio API, users should be mindful of the considerations discussed in Section [5.2](#), favoring low latency options such as ASIO. Library usage begins with creating an audio context, or Session. A Session opens communication with a requested audio device and starts an output stream to it in a separate processing thread.

```
# create an audio context
session = Session()

# enumerate connected hardware
for dev in session.available_devices:
    print(dev.index)          # e.g., 6
    print(dev.name)          # e.g., MOTU Pro Audio
    print(dev.max_channels)  # e.g., 24
    print(dev.api_name)      # e.g., ASIO
    ...                       # etc.

# open device 6 with 24 channels at 48 kHz
session.open(6,24,48000)
```

Listing 5.1: Querying hardware information and opening devices

## Creating Effects with Signals

Vibration waveforms are represented abstractly by one or more Signals. Signal classes define a temporal sampling behavior and length, which may be finite or infinite. A variety of built-in Signals are available in Syntacts. For example, the classes Sine, Square, Saw,

and Triangle implement typical oscillators with normalized amplitude and infinite duration, while Envelope and ASR (Attack, Sustain, Release) define amplitude modifiers with finite duration. Signals can be mixed using basic arithmetic. The act of multiplying and adding Signals can be thought of as an element-wise operation between two vectors. Multiplying two Signals yields a new Signal of duration equal to the shortest operand, while adding two Signals yields a new Signal of duration equal to the longest operand. Gain and bias can be applied to Signals with scalar operands as well.

In Listing 5.2 and Fig. 5.6, the Signals `sqr` and `sin` are implicitly of infinite length, while `asr` has a length of 0.3 s. Multiplying `sqr` by `sin` yields another infinite Signal with a 100 Hz square carrier wave, amplitude modulated with a 10 Hz sine wave (`sig1`). This Signal can further be given shape and duration by multiplication with `asr` to yield the finite Signal `sig2`. The Signal `sig3` represents another form of modulation through summation instead of multiplication. While the examples here only demonstrate passing scalar arguments to Signal constructors, some Signals can accept other Signals as their input arguments. For instance, it is possible to pass `sin` as the frequency argument to `sqr`'s constructor, yielding a form of frequency modulation. The modularity of the API allows users to create a wide variety of effects with minimal code. Syntacts can also be easily extended with custom user-defined Signals simply by creating classes which define the functions `sample` and `length`.

```
sqr = Square(100)      # 100 Hz square
sin = Sine(10)         # 10 Hz sine
asr = ASR(0.1,0.1,0.1) # attack , sustain , release
# basic examples mixing the Signals above
sig1 = sqr * sin
sig2 = sig1 * asr
```

```

sig3 = 0.5 * (sqr + sin) * asr
# play Signals on channel 0 and 1
session.play(0, sig1) # plays until stopped
session.play(1, sig2) # plays for 0.3 seconds
...
session.stop(0) # stop sig1

```

Listing 5.2: Creating, mixing, and playing Signals




Figure 5.6 : Signals created in Listing 5.2

### Sequencing Signals

Multiple Signals can be concatenated or sequenced temporally to create patterns of effects using the insertion, or left-shift, operator. Consider the examples in Listing 5.3 and Fig. [5.7](#). First, two finite Signals `sigA` (0.3 s) and `sigB` (0.4 s) are created. Signal `sig4` demonstrates their direct concatenation, resulting in a 0.7 second long vibration where



sigB is rendered immediately after sigA. Delay and pause can be achieved through the insertion of positive scalar operands, as shown in sig5. Inserting negative scalars moves the insertion point backward in time, allowing users to overlay or fade Signals into each other as in sig6. Sequences of Signals can also be sequenced as in sig7.

```
sigA = Sine(100) * ASR(0.1,0.1,0.1)      # 0.3 s
sigB = Sine(50) * ADSR(0.1,0.1,0.1,0.1) # 0.4 s
sig4 = sigA << sigB                      # 0.7 s
sig5 = 0.1 << sigA << 0.2 << sigB       # 1.0 s
sig6 = sigA << -0.1 << sigB             # 0.6 s
sig7 = sig4 << sig5 << sig6             # 2.3 s
session.play(2, sig7)
```

Listing 5.3: Sequencing Signals in time




Figure 5.7 : Sequenced Signals created in Listing 5.3

### Spatialization and Realtime Modifications

In addition to playing Signals on discrete channels, multiple channels can be mapped to a normalized continuous 1D or 2D spatial representation with the Spatializer class. Similar to the Mango editor from Schneider et al. [111], users can configure a virtual grid to

match the physical layout of a tactor array, and then set a virtual *target* coordinate and radius to seamlessly play and blend multiple tactors at once. Channel positions can be set individually or as uniformly spaced grids. Only channels within a *target* radius are played, and their volume is scaled according to a specified drop-off law (e.g., linear, logarithmic, etc.) based on their proximity to the target location. By moving the target location, for example, in a `while` loop or in response to changes in device orientation, developers can create sweeping motions and the illusion of continuous space with their tactile arrays (Listing 5.4, Fig. 5.8).

Other parameters, such as master volume and pitch, can be modified in realtime for Spatializers or individual channels. This offers developers the ability to move beyond playing discrete, pre-designed cues, to instead modifying continuous cues in response to conditions within the application. For example, consider the VR application in Fig. 5.11. In addition to pre-designed haptic effects that are triggered for specific events (such as button clicks), a continuous haptic effect is rendered when the player's hand is inside the fan air stream. Volume, the spatializer target, and pitch are changed based on hand proximity, wrist orientation, and the fan speed, respectively.

```

spatial = Spatializer(session)      # 2D Spatializer
spatial.create_grid(4,6)           # 4 rows X 6 cols
spatial.set_position(18,(0.1,0.8)) # move channel 18
spatial.radius = 0.3               # effect radius
spatial.target = (0.2, 0.1)        # target location
spatial.roll_off = 'linear'        # roll off law
spatial.play(sig1)                 # play inf Signal

# modification in a loop
while condition:
```

```

...
spatial.target = (x,y)
spatial.volume = v
spatial.pitch = p
spatial.stop()

```

Listing 5.4: Spatializing factor arrays and modifying parameters in realtime




Figure 5.8 : The Spatializer created in Listing 5.4

### Saving and Loading Signals

User-created Signals can be saved to disk and reloaded at a later time using the functions `saveSignal` and `loadSignal`. The default file format is a binary representation of the serialized Signal. That is, instead of saving all individual audio samples, only the parameters needed to reconstruct the Signal at runtime are saved. This results in considerably




Figure 5.9 : Syntacts GUI - The left-hand side demonstrates cue design. Users drag, drop, and configure Signals from the design Palette to the Designer workspace. The Signal is visualized and can be played on individual channels of the opened device. The right-hand side shows the GUI's track-based sequencer (background) and spatializer (foreground) interfaces. Once designs are complete, they can be saved and later loaded from the programming APIs.

smaller files which can be loaded more quickly on the fly than typical audio file formats. Nonetheless, Syntacts can still export and import WAV, AIFF, and CSV file formats for interoperability with existing haptic libraries.

## 5.4.2 Syntacts GUI

In addition to the raw APIs, Syntacts ships with a feature-rich GUI (Fig. 5.9). The GUI includes a drag-and-drop interface for designing Signals from built-in configurable primitives. The resulting Signal is immediately visualized to facilitate the design process. A track-based sequencer and spatialization editor are also included. Signals can be tested on a selected device's output channels, and then saved to the user's library for later use. Leveraging library features, users employ the GUI to rapidly tune haptic effects being loaded and played from Syntacts code in an separate application (e.g., iteratively tuning the effects for the buttons and knobs of the fan in Fig. 5.11). The GUI application is available as a precompiled executable or in source code format.



Figure 5.10 : Syntacts In Use - This figure demonstrates a real-world implementation of the Syntacts amplifier, where it has been used to drive two Tasbi haptic bracelets. A professional grade audio device (MOTU 24Ao) is connected to two Syntacts amplifier boards that have been integrated into separate Tasbi control units. Amplifier output is transmitted to each Tasbi over a multi-conductor cable. Each Tasbi bracelet incorporates six Mplus 1040W LRA factors radially spaced around the wrist, for a total of twelve utilized audio channels. The audio device interfaces with a host PC (not shown) through a USB connection.

## 5.5 Comparison

In this Section, we evaluate Syntacts against two of the commercially available control options discussed in Section 5.2: the EAI Universal Controller, and the TI DRV2605LEVM-MD evaluation board. Each controller was implemented with the manufacturer-recommended configuration so as to best compare them with the Syntacts framework.

### 5.5.1 Latency Benchmarking

Latency is a critical measure of a system's ability to render cues, especially for time sensitive applications like VR. For high-density tactile arrays, latency can increase with the number of channels simultaneously played since each subsequent channel adds more processing or transmission time. If multiple channels are played at once, the last actuated channel may lag the first actuated channel by several milliseconds depending on the over-




Figure 5.11 : Syntacts In Use - Here, the C# binding of the Syntacts API is used in Unity Engine to provide haptic effects for a virtual fan interaction designed for the Tasbi setup shown in Fig. 5.10. Two usage paradigms are in effect. The first leverages pre-designed, finite Signals for knob detents (designed in the Syntacts GUI and loaded at runtime) and button contact events (created programmatically on-the-fly, parameterized by hand approach velocity). The second paradigm uses an infinitely long Signal for the fan air stream. The volume and pitch of this Signal are modified in realtime based on the user's hand location and the fan speed, respectively. One-dimensional spatialization is used to target only the tactors which are oriented toward the fan in a continuous fashion.




Figure 5.12 : The testing used for all latency benchmarking. An Mplus ML1040W LRA was epoxied to a 100 g ABS block, and an accelerometer measured LRA induced vibrations along the y-axis. Latency was defined as the time from calling the software APIs to command vibration to the time at which 0.015 g of acceleration was measured.

all implementation. For this reason, we chose to benchmark latency as a function of the number of channels played at once.

We defined latency as the time from calling the functions to create and play a cue on  $n = [1, 8]$  factors until an appreciable acceleration (0.015 g) was measured on the last actuated factor. To perform the test, we constructed an apparatus (Fig. 5.12) based on the factory testing rig for the Mplus ML1040W LRA vibrotactors that were used. An accelerometer (TE Connectivity 4000A-005) was attached perpendicular to gravity on a 100g block of acrylonitrile butadiene styrene (ABS). The block rested on a layer of polyurethane foam to mitigate external vibrations. A C++ testing application, also available online, controlled the experiments and ran 100 trials for each device. Data was collected with a Quanser QPID digital acquisition device polled at 50 kHz. All systems rendered a 178 Hz sine wave between  $\pm 5V$  with a duration of 1,000 ms.

Syntacts software was configured to control a MOTU 24Ao under the ASIO driver API and a buffer size of 16, with power amplification being performed by the Syntacts amplifier board. Syntacts and EAI systems were controlled through their respective APIs, called directly from the testing application. As the datasheet for the EAI Universal Controller notes, it can only play four factors at full amplitude simultaneously and its API imposes this limit, so its testing concluded there. Due to the nature of the TI chip, I<sup>2</sup>C brokerage was required to interface with the testing application. We used an Arduino Uno for this purpose, under the assumption that it represented the most likely use case. The EAI and TI drivers were programmed to use manufacturer recommended methods to minimize cue latency.

Accelerometer data were reduced to find the mean and standard deviation of the latency for each system and number of channels played (Fig. 5.13). The Texas Instruments system has the highest latency for a single factor, but does not increase latency through four factors. After the fourth factor, the average latency and standard deviation increase, possibly due to I<sup>2</sup>C multiplexer components, but again stays constant after five factors. The Arduino likely contributes most to this latency, but since it represents a very plausible implementation, we consider it a fair comparison. The EAI system has lower latency than the TI system for one and two factors, but the latency linearly increases with number of channels played to greater than the TI system, and as noted cannot play more than four channels. The Syntacts system has significantly lower latency than either of the commercially available systems tested and does not seem to be a function of channels played, so the system could expand to larger factor arrays without delays. Though not shown, we measured similar latency values for the MOTU 24Ao with 24 channels played simultaneously.






Figure 5.13 : Latency as a function of channels rendered, measured as the time from software triggering to the detection of factor acceleration. Only four channels are shown for the EAI control unit since this is its max.

### 5.5.2 Overall Comparison

Whole-system comparisons of the vibrotactile control methods tested are summarized in Table 5.1. The different programming APIs show the extent to which hardware can be integrated within software. The GUI column lists the different functionality of the included graphical user interfaces. Synthesizers are able to create cues, Sequencers have the ability to organize cues in time on one or more channels, and Spatializers allow users to specify the center of vibration for an array of tactors. The audio hardware listed only represents a small subset of the possible options, but as can be seen Syntacts allows users to select audio devices based on output needs and cost. For around \$125 USD, researchers can interface a 7.1 surround sound card with Syntacts and the Syntacts amplifier to achieve a complete 8 channel setup comparable in performance to the \$2,250 USD EAI Universal Controller. Though rendering more than 8 channels with audio comes at a cost, it can still be done for

Table 5.1 : Comparison of Tactor Control Methods Tested

Method	Interface	Open Source	API / Language	GUI	Hardware	Max. Channels	Avg. Latency (ms)	Approx. Cost (USD)
Audio	Syntacts	Yes	C, C++, C#, Python	Synthesizer Sequencer Spatialization	Headphone Jack	2	8.88	\$75 <sup>†</sup>
					SB Audigy RX7.1 (PCI-e)	8	4.22	\$125 <sup>†</sup>
					MOTU mk4 (USB)	12	5.20	\$750 <sup>†</sup>
					MOTU 24Ao (USB)	24	2.97	\$1,225 <sup>†</sup>
Controller	EAI	No	C	Synthesizer Sequencer	EAI Universal Controller	8 <sup>‡</sup>	5.35	\$2,250
IC	I <sup>2</sup> C	Yes	N/A	Synthesizer	DRV2605LEVM-MD	8	6.65	\$150

<sup>†</sup>Includes the cost of the number of Syntacts amplifiers (at \$75 USD ea.) to accommodate the maximum available channels of the audio interface.

<sup>‡</sup>While the EAI Universal Controller supports eight channels, only four can be played simultaneously.

much less that the cost of multiple EAI controllers and is considerably more manageable than implementing an integrated circuit based design.

## 5.6 Limitations and Future Work

Syntacts is not without its limitations, and may not be the perfect tool for all researchers. For one, Syntacts requires a host PC and tethering hardware to audio interfaces. Therefore, Syntacts is not well suited for mobile or wireless haptic devices. Second, the Syntacts amplifier, while being compatible with large portion of commercial tactors, is not designed to power ERM or Piezo-actuators, and may have difficulty driving large and/or higher power actuators.

Future work for Syntacts involves both improvements on the usability of the software as well as understanding the use space more fully. In particular, immediate work will focus on extending realtime Signal modification features for VR applications. We aim to integrate more haptically oriented tools as well, perhaps eventually favoring tactile perceptual models

over audio centric concepts such as volume and pitch. Iteration on the API and GUI from user feedback would further increase the usability of the program. In closing, given the open-source nature of Syntacts, we welcome and hope that the haptics community will also contribute to its continued development.

## Chapter 6

### Psychophysics of Haptic Stimuli at the Wrist

In this Chapter, we characterize Tasbi in terms of human perceptual performance. The main contribution of these studies stems from Tasbi's unique ability to control directly for wrist squeeze force. We quantify the fundamental squeeze force thresholds for the wrist (Study 1) as well as vibrotactor identification rates at varying levels of squeeze load (Study 2) and stimulus amplitude (Study 3). These studies have direct implications to not only Tasbi applications, but also to future device designs.

#### 6.1 Study 1: Squeeze Difference Threshold

Our first study sought to characterize users' perception of wrist squeeze stimuli. Two common measures of haptic perceptual performance are the absolute and difference thresholds [127]. Here, we choose to focus on the latter, and attempt to quantify the just noticeable difference (JND) for wrist squeeze force. Although JND studies have been conducted for squeeze on the wrist and arm, they are typically quantified in units that are indirectly related to the perceived stimulus and are inherently tied to the device with which the study was performed (e.g. the angular displacement of a motor used to produce squeeze [48], the linear displacement of a squeezing belt [91], or the axial load in a squeeze inducing shape memory alloy [83]). Because Tasbi has the ability to control directly for uniform squeeze force, we can quantify wrist squeeze in practical units of normal force against the skin. To our knowledge, this is the first reporting of such. Whether squeeze perception is linked




Figure 6.1 : The GUI subjects interacted with during the studies. (a) Interface for vibrotactile identification study. (b) Interface for squeeze difference threshold study.

to contact force, tissue displacement, or something else entirely remains an open question, but the results presented here should provide future designers and scientists with a more general understanding of wrist squeeze perception than what currently exists.

### 6.1.1 Subjects and Procedures

In accordance with Rice University IRB Protocol #IRB-FY2020-43, we recruited 12 subjects (5 female, ages 21 to 33, mean 26). Each subject participated in both Study 1 and Study 2 in a single 90 minute session divided by a short break. All subjects were naive to wearable haptics and had no prior training or experience with Tasbi.

Subjects completed the study by interacting with a on-screen GUI (Fig. 6.1) using a mouse with their right hand. Tasbi was worn on the right wrist, and each subject's arm was supported such that Tasbi was suspended over free space and not inadvertently resting on any surfaces. To prevent use of visual or auditory information, a curtain occluded subjects' view of Tasbi, and pink noise played over headphones throughout the experiments.

Following the experiments, each subject had their wrist dimensions measured at the Tasbi stimulus site (approximately 6 mm behind the styloid process). The means and

standard deviations for wrist circumference  $C$ , width  $W$ , and height  $H$  were found to be  $185.7 \pm 13.6$  mm,  $55.7 \pm 5.1$  mm, and  $53.9 \pm 5.6$  mm, respectively. To analyze the effect of wrist size, subjects were evenly binned into either a *small wrist* group ( $C < 185$  mm) or *large wrist* group ( $C \geq 185$  mm).

To determine the squeeze threshold, the method of constant stimuli with a two interval forced choice (2IFC) procedure was used. On each trial, the subject was presented with two sequential squeezes cues and tasked with choosing the cue that squeezed harder, either the first or second. One cue was always the Standard force level (7 N), and the other cue was one of 11 Comparison levels (2 to 12 N in 1 N intervals). The Standard was pseudo-randomly presented either first or second in a counterbalanced manner to mitigate the so-called *time error* of 2IFC procedures [127]. To eliminate reliance on temporal information, each squeeze cue was rate controlled to last one second regardless of the target force level, where Tasbi ramped up to the force level over  $1/3$  s, held the force for another  $1/3$  s, and then ramped down to no squeeze force over the remaining  $1/3$  s. A  $1/4$  s delay was placed between the first and second cue. Subjects made their selection in the GUI, and performed 550 trials, or 50 repetitions of each comparison level. The trials were evenly divided into 5 windows, separated by a 60 second break.

### 6.1.2 Results

Fig. 6.2 shows the proportion of times subjects indicated that each Comparison squeeze level was greater than the Standard level. The data was fit to a general linear model with a logit link function to estimate the psychometric function for each subject. The JND is defined as the difference between the 75% (or 25%) threshold and the 50% threshold (i.e. the point of subjective equality or PSE). Across all subjects, the JND for wrist squeeze force was found to be  $1.28 \pm 0.46$  N (mean + SD). Given the Standard of 7 N, the Weber

fraction was approximately 0.18. The PSE was  $7.10 \pm 0.24$  N, and corresponds well with the Standard level.

The boxplot of JNDs for all subjects in Fig. 6.2-b suggests that some subjects are far more perceptive to wrist squeeze than others, with JNDs ranging from as low as 0.72 N to as high as 2.07 N. In contrast to Study 1, we cannot attribute this difference to wrist size, which was found to have no significant effect on JND ( $t(10)=-.48$ ,  $p=.65$ ). This further evidenced in Fig. 6.2-a by comparing the mean psychometric functions both the small and large wrist groups.

## 6.2 Study 2: Vibrotactile Identification versus Squeeze Force

The second study aimed to characterize user's ability to successfully identify each of Tasbi's six vibrotactors when presented at random. The study was further designed to test if identification rates would be affected by different levels of static squeeze force and vibration stimuli duration.

### 6.2.1 Subjects and Procedures

The same subjects that performed Study 1 were used in this study. They completed Study 2 directly after completing Study 1.

The experiment was divided into three blocks conditioned on the level of preload squeeze force (0.5 N, 5 N, or 10 N) for that block. At the beginning of each block, Tasbi tensioned to the target force and held that force for the remainder of the block. The block order presentation was randomized between subjects so that each of the six possible orders were equally represented. Within each block, 240 vibration stimuli were presented. Each stimulus was characterized by the individual vibrotactor actuated, or the stimulus location (T1, T2, T3, T4, T5 or T6), and the duration of the stimulus (50 ms or 250 ms). The excita-




Figure 6.2 : **(a)** The mean psychometric function(s) experimentally determined for wrist squeeze force. Error bars represent a 95% confidence interval. **(b)** The mean difference threshold, or JND, was found to be 1.28 N, and **(c)** the PSE shows little bias from the Standard of 7 N. Importantly, we find that wrist size has no significant effect on the JND.



tion signal was held at a constant frequency (170 Hz) and amplitude (2.5 V<sub>rms</sub>), consistent with the nominal operating conditions of Tasbi’s LRAs. The stimuli conditions were evenly distributed, and thus 20 repetitions of each actuator-duration pairing were presented in each of the three blocks. Subjects indicated the vibrotactor they identified as being played via the GUI, which displayed a 2D schematic of Tasbi’s vibrotactor wrist layout similar to that shown in the bottom right of Fig. 6.3. Subjects were given approximately two minutes to self-explore the Tasbi’s vibrotactors using the GUI at the beginning of the experiment to help them internalize the GUI schematic in relation to the tactile stimuli.

### 6.2.2 Results

The main results are shown in Fig. 6.3 where the proportion of all subjects’ responses under each condition are plotted as confusions matrices. The percentage of correct responses  $P_C$  is read along the diagonal of each matrix. Overall, we see an identification rate of approximately 67.8% across all conditions, consistent with the findings in [68] for a 6-tactor design. A few notable differences should be taken into consideration. First, we tested much shorter stimulus durations (50 ms and 250 ms versus 600 ms) which we felt represented a more likely range of stimulus duration. Second, because the 12 o’clock position on Tasbi is occupied by the squeeze tensioner housing, Tasbi’s tactors are more densely packed than the device in [68]. However, the difference is likely offset by the fact that our test was conducted more distally from the styloid process (6 cm versus 3.5 cm) where circumference is larger. Indeed, a rough calculation of tactor-to-tactor spacing for both studies is a comparable 26 mm.

To further analyze the data, a three-way repeated measures ANOVA (6 locations  $\times$  3 force levels  $\times$  2 stimulus durations) was conducted with  $P_C$  being the dependent measure. The majority of groups passed the Shapiro–Wilk test for normality, and all groups passed




Figure 6.3 : Stimuli-response confusion matrices for each squeeze force and VT stimulus duration pairing, aggregated across all subjects. The probability of subjects correctly responding are given as a function of the stimulus site. Individual columns sum to 100%. The bottom left matrix combines all conditions. The total percentage of correct responses for each condition are given in the subplot titles. The bottom right inlay displays the approximate location of each stimulus relative to wrist and forearm anatomy, particularly the radial (R) and ulnar (U) bones. In general, we observe: 1) identification rates are greatly reduced at stimulus sites located over bony areas, 2) subjects seem to perform better given a longer stimulus duration, and 3) the middle squeeze level of 5N yields the best performance, suggesting that there may exist an optimum level of preload squeeze force.




Figure 6.4 : Percentage of correctly identified vibrotactor stimuli for both groups of wrist size as a function of the stimulus location. The data represents the mean of all conditions, and error bars are for a 95% confidence interval.

Mauchly's test for sphericity. The main effect of stimulus location was found to be significant ( $F(5,55)=3.72$ ,  $p=.006$ ) and is evident in Fig. 6.3-g, where identification accuracy varies greatly from 53.0% to 74.9%. Collapsing the data across factors of force and duration and performing pairwise comparisons between stimulus locations with a Bonferroni correction shows that subjects perform significantly worse for the T1 and T5 stimuli locations ( $p<.005$  in all relevant comparisons). Unsurprisingly, feedback from subjects during the self-exploration phase suggested that these vibrotactors were most difficult to identify. The general consensus was that vibrotactors over bony areas of the wrist were more difficult to localize than those over soft tissue. Fig. 6.3-h illustrates this phenomenon, where the approximate location the radial (R) and ulnar (U) bones can be seen.

Although the total percent correct in Fig. 6.3-e,f, along with qualitative feedback from subjects, suggests that identifications rates decrease with higher levels of squeeze force, we find no significant difference for the main effect of squeeze force. While it seems probable that an effect could be found with a more nuanced study, it is reassuring to find that squeeze

does not drastically interfere with the perception of vibration as experiments have shown for other combinations of multimodal cues (e.g. skin stretch and squeeze [48]).

The main effect of duration was also found to be significant, with the 250 ms stimulus providing more accurate responses ( $F(1,11)=5.63$ ,  $p=0.037$ ). This is evident when comparing the total percent correct between columns in Fig. 6.3. Interestingly, we note that the mean accuracy difference between the long and short stimuli increases as a function of squeeze force (1.6%, 2.5%, and 4.2% for 0.5 N, 5 N, and 10 N, respectively.) However, this may only be a trend as we find no significant interaction between squeeze force and duration. Regarding other interactions, the ANOVA revealed a significant interaction between location, force, and duration ( $F(10,110)=2.07$ ,  $p=.033$ ) and between location and force ( $F(10,110)=2.13$ ,  $p=.028$ ).

Finally, we note that subjects with large wrists ( $C \geq 185$  mm) perform significantly better than those with small wrists when analyzed across all conditions ( $t(430)=2.97$ ,  $p=.003$ ). Fig. 6.4 shows that the large wrist group outperforms the small wrist group at 5 of the 6 stimulus locations. This phenomenon was also observed in [68], and is not particularly surprising considering larger wrists spread adjacent factors further apart. This may also be a function of tissue impedance, as individuals with large wrists tend exhibit higher concentrations of adipose tissue. A study that correlates identification accuracy with body mass index (BMI) would be an interesting follow up.

### 6.3 Study 3: Vibrotactile Identification versus Amplitude

The third study investigated the effect that vibration stimulus intensity (i.e. amplitude) would effect vibrotactile identification rates and test three separate squeeze levels.




Figure 6.5 : Percentage of correct responses separated by squeeze condition and vibration stimuli amplitude. High amplitude vibrations yielded best performance, and individuals with large wrists circumferences ( $\geq 168$  mm) outperform those with small wrist sizes.

### 6.3.1 Subjects and Procedures

We first identified three conditions of squeeze to test: (1) a “loose” amount of only passive squeeze from the elastic band, i.e. no active squeeze, (2) a “fitted” level of active squeeze, corresponding to the amount of squeeze users felt was comfortable for all day wearability, and (3) a “tight” level of squeeze which users felt would be acceptable for only a few minutes. Subjective pilot testing among subjects determined the “fitted” level to be approximately 0.5 N of normal squeeze force, and the “tight” level at 5 N of normal force. We displayed two vibration levels, one with an amplitude at 20 dB SL and the other at 30 dB SL. Both stimuli were rendered at the factors’ 170 Hz resonance frequency for 500 ms.

In accordance with IRB Protocol #20182617, we recruited 13 subjects (7 males, ages 23 to 44, mean 30). The mean wrist circumference was measured as 168.7 mm (SD = 16.6 mm). Each subject received all three conditions as three separate blocks of one-interval, six-alternative forced choice (1I-6AFC) trials. In each trial, a single random factor rendered

one of the two vibration stimuli at random, and the user then identified the stimuli location using a mouse and GUI interface. Each factor/stimuli pair was repeated 5 times for a total of 60 trials per block. After the block was over, the device was tensioned to the squeeze level for the next block and the procedure was repeated. The order in which the squeeze condition blocks were presented was randomized among subjects.

### 6.3.2 Results

The main results are shown in Fig. 6.5 where the percentages of trials correctly identified are shown. The effect of vibration stimuli amplitude was significant ( $F(1,11)=34$ ;  $p<0.01$ ), as users showed higher levels of performance for the higher amplitude stimulus. By binning users into two groups of wrist circumference ( $n=7<168$  mm,  $n=6\geq 168$  mm), the effect of wrist size was found to be significant ( $F(1,11)=8.0$ ;  $p=0.01$ ) with small wrist subjects displaying lower performance. This result mirrors the finding in Study 2.

## 6.4 Discussion

Using Tasbi's ability to control directly for squeeze force, this chapter presented three perceptual studies. The first tested for the just noticeable difference of wrist squeeze force, which we found to be 1.28 N averaged across all subjects. The threshold represents approximately 10% of Tasbi's 15 N rendering range, and suggests that around 10 unique levels of squeeze may be deliverable to users. We find that wrist size has no significant effect on this threshold. This is a fortunate result since it means that we do not require specific squeeze designs for different individuals.

Our second study tested subjects' ability to identify stimuli from Tasbi's six vibrotactors under varying levels of static squeeze force, and offered the following insights: 1) vibrotactor identification rates are affected by radial location, with vibrotactors over bony areas

of the wrist being significantly more difficult to discern; 2) subjects with large wrists perform significantly better than those with small wrists; and, 3) the amount of static squeeze has no significant effect on identification rates. The first and second points make it clear that perception of Tasbi's vibration stimuli is affected by limb geometry. Future designs may wish to address this issue with user configurable vibrotactor spacing options. The final point is particularly important for multimodal devices seeking to integrate squeeze and vibrotactile feedback. The third study showed that the amplitude of vibration at the wrist has a significant effect on identification as well, with higher amplitude vibrations proving to be more discernable.

## Chapter 7

### Referred Haptic Feedback for Hand Interactions

This chapter presents methods for and examples of using Tasbi to provide referred haptic feedback for hand interactions in virtual reality. Our aim is to deliver haptic stimuli to the wrist for interaction forces that would otherwise be felt at the hands and fingertips. Because this idea is relatively unexplored in the literature, there is limited theoretical basis on which we can rely to build such interactions. Thus, the work presented here is mostly exploratory in nature, and sets the stage for more focused research questions that are described in Chapter 8. The content of this chapter demonstrates the wide range of possibilities for referred haptic feedback. Many examples first appeared in [46]. Where appropriate, comments and insights on the interactions presented are provided, based on personal experience and subjective feedback that has been gathered during numerous demonstrations.

#### 7.1 Multisensory Feedback Paradigm

The task of rendering convincing referred feedback for hand interactions requires clever use of all available haptic channels in a way that matches user's expectations about a particular interaction. For Tasbi, we must develop a framework to associate both squeeze and vibration with particular phenomena. We should also consider the possibility that haptic feedback does not exist in isolation from other sensory modalities, and may only be as effective as the visual and audio stimuli that coincide with it. For this reason, special care has been taken in appropriately match haptic stimuli with audio and visual stimuli. In fact,



visual stimuli, or rather visual-haptic illusions, play a large role in our approach to referred haptic rendering. Sections 7.1.1 to 7.1.3 describe our three part method to rendering many of the interactions the follow in this chapter.

### 7.1.1 Vibrotactile Feedback

The first element to our rendering approach is vibration. Vibrotactile feedback is most appropriate for *discrete* events such as finger contact, hand collisions, impacts, and other high frequency content. There are a number of models that have been developed to realistically render vibrotactile feedback for various surface types [3, 129]. Data driven approaches, which leverage recordings of physical vibrations, are also frequently used [130]. Often, it is sufficient to emulate contact events with a simple exponentially decaying sinusoid of the form  $Ae^{-Bt} \sin(2\pi\omega t)$  (Fig. 7.1). Okamura et al. provides amplitude  $A$ , decay  $B$ , and frequency  $\omega$  parameters for wood, rubber, and metal materials [128]. This is the approach




Figure 7.1 : Vibrations can be emulated using a decaying sinusoid model. Here, the model is fit to vibration data recorded while tapping wood. Adapted from [128].

that we take.

All of the implementations above were specifically developed for hand-held devices. Relocating feedback from the hand to the wrist requires special consideration toward the effect of spatial incongruence, i.e. how users will react when contact visually perceived at the hands or fingertips is haptically perceived at the wrist. One option might be to take corrective action for these spatial discrepancies. For instance, Shao et al [131] investigated and modeled how vibration waveforms originating at the fingertip propagate through the hand and wrist. Though they did not explicitly suggest it, it might be possible to leverage such a model for referred haptic feedback. In this case, the input to the model would be the vibration occurring at the fingertip, and the output would be the propagated waveform at the wrist, which in turn could be rendered through Tasbi.

Another other option is to simply ignore the effects of spatial mismatches altogether. It is well known that the brain is highly malleable in terms of vision, tactile perception, and proprioception. The Rubber Hand Illusion is an often referenced and studied phenomena, whereby subjects will reassign ownership of entire limbs in the presence of conflicting visual and haptic stimuli [132, 133]. Phenomena such as these have been used as the basis for providing referred feedback for prosthetic applications with success [134, 135]. Thus, spatial mismatches between the fingertips and wrist may not matter all too much. Tasbi offers some additional assurances. Usually, we actuate all six radial vibrotactors during contact events. The effect that this has is that it is difficult for subjects to localize the source of vibration to the surface of the wrist. Thus, they are more inclined to attribute it to the contact event occurring at the fingertip.

The effect of vibrotactile feedback is greatly enhanced when rendered in concert with audio stimuli. If audio is to be played, it is necessary in many cases to render both simultaneously or within the just noticeable threshold of temporal asynchrony (around 24 ms

[136]), though in some circumstances audio should slightly lag visual stimuli for the most realistic presentation [137]. Additionally, it is important for vibration stimuli to be temporally synced with visual stimuli, as research has shown the perceptual limit of visual-haptic simultaneity to be 50 ms [117] when visual contact precedes haptic feedback, and as low as 15 ms when haptic feedback precedes visual contact. There is however some leeway in tactile-audio and tactile-visual synchrony, as there is evidence that the brain can recalibrate to temporal incongruencies of up to around 100 ms between all modalities [138].

### 7.1.2 Squeeze Feedback

The second part of our approach is squeeze feedback. We leverage proportional squeeze for *continuous* interaction forces. For example, we can ramp squeeze as a virtual button is pressed to render its stiffness, or squeeze when objects are picked up to convey their mass or inertia. In contrast to vibration, there exists very little literature to suggest how one might map virtual forces to wrist squeeze. In other words, if one would like to convey a particular stiffness or mass through wrist squeeze, how much squeeze force should be provided? In Chapter 8 we investigate this question more deeply through controlled human subject studies. For the sake of this chapter, we use a simple ad-hoc approach whereby squeeze force is controlled to be arbitrarily proportional to the virtual forces to be rendered. Though subjective, these initial explorations provide some insight into how effective squeeze may be for conveying continuous interaction forces originating at the hands and fingertips.

### 7.1.3 Pseudo-Haptic Control/Display Feedback

With vibration and squeeze providing haptic feedback, the third and final element to our approach is entirely visual. We begin by first employing the *god-object* haptic rendering method [139, 140, 141, 142]. Here, the god-object is the avatar finger displayed to the

user in VR, and is constrained by a collision simulation such that it can never penetrate virtual objects, even if the user’s physical finger location would otherwise place them in an overlapping state. The god-object is coupled to the user’s true finger position through a virtual spring. Thus, when a user contacts virtual object, we can compute the force to apply to the object from the displacement between the god-object and actual finger positions.

We extend this idea through a concept known as *psuedo-haptics* [143], where discrepancies between the real (control) hand, and the virtual (display) hand are exaggerated to suggest that certain actions require more physical effort. We have previously shown that manipulating the control-to-display (C/D) ratio congruently with squeeze improves users’ perception of virtual “stiffness” [58].

Historically, C/D manipulation has been used in 2D graphical user interfaces to change the rate a which control input affects the displayed output. For example, consider a mouse interaction where the C/D ratio  $\lambda$  is defined as the ratio of mouse to cursor movement:

$$\lambda = \frac{x_{\text{mouse}}}{x_{\text{cursor}}} \quad (7.1)$$

In this case, the C/D ratio visually scales the cursor position such that C/D values greater than 1 compress the user’s mouse movements, while values less than 1 extrapolate their movements. Lécuyer and others have famously demonstrated that altering the C/D ratio can influence users’ perception of the physical properties associated with virtual objects, including mass, stiffness, and friction [143].

Here, we use a C/D approach that was loosely described in our previous paper [58] and subsequently formalized by our co-authors in [144]. This approach differs from the traditional method in that C/D is not simply a visual scaling from the control to the display, but rather the manifestation of changing the parameters of the god-object rendering method.

Consider the object shown in Figure 7.2. The object is a simulated second-order system




Figure 7.2 : Example of the multisensory feedback paradigm – \*a) The user approaches the virtual button simulated by a mass  $m$ , stiffness  $k$ , and damping  $b$ . The proxy finger control (blue) and display (gray) are coupled via a virtual spring of stiffness  $k_p$  and are initially co-located. When contact is made, Tasbi's six LRAs render a vibration to simulate the event. (b) The user begins to push the button downward. Tasbi squeeze force increases proportional to the button displacement  $x$ . The proxy hand control continues to track the users true hand position and orientation, while the display remains on the surface of the button. The control-to-display (C/D) ratio is given by the ratio of  $x_{control}$  and  $x_{display}$ . At the end of travel, squeeze reaches its maximum force level, and the C/D discrepancy is most pronounced. Note that subjects did not receive a visual representation of the control hand, and are unaware of its presence.

parameterized by a mass  $m$ , a damping coefficient  $b$ , a stiffness  $k$ , and a displacement  $x$ :

$$m\ddot{x} + b\dot{x} + kx = F \quad (7.2)$$

Interaction forces applied to the object are computed from the displacement of the god-object spring and its stiffness  $k_p$ . Thus, greater penetration depths of  $x_p$  result in larger forces being applied to the object:

$$F = k_p x_p \quad (7.3)$$

We can achieve a desired steady-state C/D by changing the ratio of object stiffness  $k$  to proxy hand stiffness  $k_p$ . Let the C/D ratio  $\lambda$  be defined as:

$$\lambda = \frac{x_{\text{control}}}{x_{\text{display}}} = \frac{x + x_p}{x} \quad (7.4)$$

Ignoring the dynamic contributions of  $m$  and  $b$ , the force balance equation is simplified to:

$$F = k_p x_p = kx \quad (7.5)$$

Finally, combining [7.4](#) and [7.5](#), we define the C/D ratio  $\lambda$  in terms of the two stiffnesses:

$$\lambda = \frac{\frac{F}{k} + \frac{F}{k_p}}{\frac{F}{k}} = \frac{k + k_p}{k_p} \quad (7.6)$$

Thus, to achieve a desired C/D ratio  $\lambda$ , the implementation can compute either  $k$  or  $k_p$  while holding the other constant. Although either approach is valid, the most logical implementation holds the proxy hand stiffness  $k_p$  constant for all interactions, and lets the desired C/D ratio drive the calculation of object stiffness  $k$ .

$$k = k_p(\lambda - 1) \quad (7.7)$$

The choice of the free variables  $k_p$ ,  $m$ , and  $b$  are context dependent, and our choices for these values are discussed where necessary. It is important to note that because we have chosen to ignore the object mass and damping in the force balance equation ([7.5](#)), the *actual*

C/D ratio will slightly fluctuate during dynamic interactions, and level out to the true value at steady-state (i.e. when the user's finger movement ceases). The amount of C/D error in the dynamic case can be minimized by carefully tuning the object's natural frequency and damping to the anticipated speed of interaction. A more robust method may choose to include the object's mass and damping terms in the force balance equation, and/or add a damping term to the proxy finger impedance. Thus, a trade-off between accuracy and simplicity exists and should be considered carefully.

Pseudo-haptic and C/D methods are also available for rendering weight in virtual reality. More often than not, these methods are accomplished purely through visual offsets or by injecting lag into the rendering system [145, 146, 147, 148], though in some cases physically based approaches have been taken [149]. Though we are not the first to augment haptic feedback with pseudo-haptic effects, [150, 151, 152, 153], this may be the first instance where pseudo-haptics have been used to enhance the efficacy of referred haptic feedback on the arm or wrist during XR hand interactions.

## 7.2 Virtual Reality Implementation and Sandbox Environment

A 3D rendering environment was required to begin developing virtual reality interactions. There exist many game engines that support video output to virtual reality head mounted displays (HMD), with the two most popular options being Unreal Engine [154] and Unity [155]. More haptic oriented engines, such as Chai3D [156] provide both a visual rendering context as well abstracted APIs for various commercial haptic devices, but not much in the way of custom hardware such as Tasbi. Ultimately, we chose to use Unity due to its high-level C# scripting engine and a vibrant community of VR developers and asset makers.

Another important point factoring into decision to use Unity was the availability of SteamVR [157]. The SteamVR SDK provides many necessary utilities for VR develop-




Figure 7.3 : All virtual Tasbi interactions were built using Unity Engine. A custom plugin wraps the native Tasbi C++ API into accessible C# scripts, and provides Unity *prefab* objects that can be easily added to any scene. For instance, one *TasbiHub* prefab is required to launch the squeeze control server and Syntacts engine, and then any number of *Tasbi* prefabs can be added to the scene to interface with a physical Tasbi device.

ment, notably a hand posing and animation system. This was essential since our implementation relies on handheld controllers to track hand location and receive user inputs. Most VR controllers incorporate capacitive sensors throughout the grip area and provide developers with a general estimate of the hand pose the user is making (e.g. a fist, pointing, thumbs-up, etc.). It is then the responsibility of the developer to animate the user's avatar hands in a convincing manner. For instance, if the user reaches out toward an object and squeezes the grip of the controller, the avatar hand should animate from its current state to a state where it is holding the object. This is a very challenging animation task, and the SteamVR hand animation system greatly simplifies the development process. The SteamVR SDK also conveniently wraps hardware APIs for various commercial HMDs, so that headsets can be easily swapped without needing to update API calls. We primarily used an Oculus Rift CV1 and an Oculus Quest HMD [8] during development.

Tasbi is integrated into Unity through a custom plugin interface. Several C# scripts wrap the native Tasbi C++ API and Syntacts framework in a modular fashion so that individual elements can be easily interfaced (e.g. separate scripts for squeeze control, vibro-






Figure 7.4 : The Tasbi sandbox environment where most haptic interaction design took place. The environment presents users with six islands that can be teleported between. Each island displays a particular idea or concept we wish to explore with Tasbi.

tactile control, 3D model control, etc.). These scripts are further combined into reusable assets called *prefabs*, such that integrating Tasbi(s) into an existing VR application is as simple as dragging and dropping prefabs into the scene hierarchy (Fig. 7.3).

A sandbox environment was developed to explore interaction design (Fig. 7.4). The environment consists of six separate islands that the user can teleport between. Each island focuses on a particular “theme” of referred haptic feedback, each discussed in the following sections. The environment intentionally takes place within an ocean setting that presents users with the audio of crashing waves. This was done to provide contextual white noise, so to speak, as means of drowning out all mechanical noises emanating from Tasbi.

## 7.3 Interaction Themes and Examples

### 7.3.1 Primitive Motions

Our multisensory approach is easily understood through the example of a virtual push button (Fig. 7.5). When the user’s virtual finger first makes contact with its surface, a low amplitude vibration is rendered. As the button is pressed, squeeze increases proportional to its displacement. Changing the C/D ratio to increase the amount users must extend for a given button displacement gives the impression that the button is harder to press, or “stiffer”. Additional vibrations can be played along the button’s range of motion and end of travel for heightened realism. We consider the button an example for the primitive interaction of *pushing* objects. We have also explored other primitive motions, such as *twisting* in the form of knobs, and *pulling* in the form of pull handles (Fig. 7.6). Although the haptic feedback for each of these is exactly the same, we have found that users generally interpret each primitive interaction differently given the context of the associated visuals. As such, Tasbi is flexible enough to accommodate many types of interactions.




Figure 7.5 : The multisensory mid-air button interaction. **(a)** The user approaches the button. The proxy hand control (blue) and display (black) are initially co-located. **(b)** The user makes initial contact with the button. Tasbi's LRAs render a vibration to simulate the contact event. **(c)** The user begins to push the button downward. Tasbi squeeze force increases proportional to the button displacement. The proxy hand control continues to track the users true hand position and orientation, while the display remains on the surface of the button. **(d)** At the end of travel, squeeze reaches its maximum force level, and the C/D discrepancy has become more pronounced. Note that the blue control hand is only shown for illustrative purposes, and users are unaware of its presence.




Figure 7.6 : Primitive Interactions – (a) The button already described by Fig. [7.5](#), showing the events of squeeze and vibration. Pull handles (b) and rotary knobs (d) use an approach similar to the button paradigm for twisting and pulling motions, respectively. Both also display vibration, squeeze, and C/D haptics. These three prototypes form the basis of many of the interactions that follow.

### 7.3.2 Surfaces and Textures

Because Tasbi is equipped with six LRAs, it can implement many existing texture rendering algorithms. Vibration in general is mostly confined to rendering *fine* surface details – bumps, cracks, ridges, etc. Rendering surfaces with *coarse* details related to geometry is typically reserved to grounded kinesthetic devices that physically restrict or guide hand motion. Tasbi, as an ungrounded wearable, cannot exert net forces/torques, yet, we find squeeze provides a surprisingly compelling way to render surface normals.

Each block in Fig. 7.7-a,b,c has a wavy geometry with a unique frequency and amplitude, and employs the previously mentioned "god-object" rendering technique, where squeeze is proportional to the proxy hand penetration. As the user sweeps their finger across the surface, penetration depth and thus squeeze force change. For the large amplitude block, this change in squeeze conveys surface height reasonably well. For the small amplitude block, typical vibration effects are employed at each peak. The middle block balances both squeeze and vibration based rendering, and is arguably the most compelling.




Figure 7.7 : Surfaces and Textures – vibration is used to render fine detail textures (a), while squeeze offers the ability to render coarse, geometry-based surfaces (c). Both can be combined (b) and even extended to 2D surfaces (d)




Figure 7.8 : In-Hand Manipulation – Squeeze and vibration can be used to convey finger interaction forces when compressing visually deformable objects.

### 7.3.3 In-Hand Manipulation

More often than not, users pick up and interact with objects in XR. Squeeze can be used to convey the net forces their fingers apply to these objects. This is especially convincing when the objects are visually deformable (such as the ball in Fig. 7.8-a) and/or display high-frequency mechanical properties that can be rendered with vibration effects (such as the friction in the hand gripper in Fig. 7.8-b).

### 7.3.4 Object Weight and Inertia

A challenge for all ungrounded haptic devices is realistically conveying the sense of object weight and inertia. Of course, it is impossible to render the downward force of weight, or the forces/torques of linear/angular inertia, without being connected to the physical world. Some glove-type devices may attempt rendering a type of pseudo-weight by applying forces to the fingers in the direction of gravity or in the opposite direction of acceleration. The challenge here is that the applied forces/torques must be grounded to other parts of the hand, where undesired reaction forces may result in perceptual discontinuities. In contrast, radial squeeze offers an interesting approach because it is *self grounding*, and




Figure 7.9 : Object Weight and Inertia – The weight and inertia of rigid-body objects can be rendered through squeeze. We find that the naive approach to displaying mass, through applying only a static amount of squeeze proportional to the object’s mass, is not effective. Instead, it is more compelling to provide fluctuating squeeze proportional to the object’s dynamics (e.g. its velocity or momentum) or proportional to wrench torques its center of mass would otherwise impart on the hand.

forces are only felt as localized inward pressure. A simple approach would map a range of object weight and inertia to a range of squeeze levels. We have found this *static* approach to be rather ineffective. Instead, rendering *dynamic* effects result is a much more compelling interaction. In Fig. [7.9](#), a base amount of squeeze is applied when the tennis ball or racket is picked up, but additionally, the inertia of these objects is rendered through small fluctuations in squeeze as the user waves them in mid-air. For the tennis racket, we also increase squeeze proportional to the tilt angle, conveying a sense of the moment arm torque it would impart to the wrist. One challenge, however, is appropriately scaling each phenomenon’s contribution to total squeeze since only one degree of freedom is available.




Figure 7.10 : Force Fields – Squeeze and vibration somewhat convincingly render the sensation of air streams, as in this fan example where squeeze increases with proximity, and noisy vibrations are rendered via the factors facing the stream. The example on the right is a playful science fiction device that provides visual context for Tasbi’s calibration process.

### 7.3.5 Force Fields

We have also explored rendering force fields for phenomena that have no physical points of interaction. The fan in Fig. 7.10 demonstrates this (see also Section 5.4). When the users’ hand moves in front of the air stream, squeeze increases to convey air pressure, and changes proportionally with increases proportional to the distance between the fan and hand. Light vibrations are rendered on vibrotactors facing the fan to convey localized wind effects. Since the vibrations are easily localized when played separately, this interaction can feel strange as the sensations are expected on the surface of the hand and not the wrist.



Figure 7.11 : User Movement and Proprioception – This ladder interaction, where squeeze increases as users lift themselves, explores movement and enhancing proprioception. However, it is arguably not a particularly well suited application for squeeze, since it impossible provide enough force with Tasbi to convincingly convey someone’s own mass.

### 7.3.6 User Movement and Proprioception

Many VR games solve the locomotion problem by allowing players to grab nearby surroundings and pull themselves in a desired direction. However, if the player is not fixated on their target, or is also focused on other tasks, proprioception can break down and the sense of locality will be lost. Squeeze can mitigate this by conveying distance traveled along a path. In the ladder example of Fig. [7.11](#), vibrations are rendered with each grab, and squeeze increases as users pull themselves up each rung, resetting on release, and thereby providing a sense of climbing. When both hands are in contact, only half the total available squeeze is provided on each wrist, and when only one is in contact, the full amount is provided. Overall, the effect is intuitive but the amount of squeeze Tasbi provides is not sufficient to actually convey a grown person’s mass.






Figure 7.12 : Bimanual Interactions – (a) In the steering wheel example, squeeze increases when torques applied by the wrist oppose each other and decreases when torques are applied in the same direction. (b) For the bow-and-arrow, both wrists experience increasing squeeze as the bow string is drawn and an instant release when the arrow is fired. Subtle vibrations while drawing increase realism.

### 7.3.7 Bimanual Interactions

The most compelling Tasbi interactions are those that render feedback on *both wrists*. This is most apparent when the interaction simulates reactive forces between hands. For example, when using a bow and arrow (Fig. 7.12), squeeze is rendered on the bow hand to simulate stabilizing forces, while increasing squeeze on the string hand simulates the build-up of bow tension. As in most examples, the experience is substantially enhanced with subtle vibration and C/D manipulation. The steering wheel in Fig. 7.12 presents the opportunity to increase and decrease squeeze depending on whether the user is applying convergent or divergent torques to the wheel.




Figure 7.13 : Mutli-Paradigm Interactions – These interactable objects combine many of the rendering approaches presented so far into a single context. Despite squeeze feedback being generally uniform for each point of interaction, the illusion generally holds in the context of the user’s motion and the rendered visuals.

### 7.3.8 Mutli-Paradigm Interactions

Even more complex handheld interactables are possible, as shown in Fig. [7.13](#). We consider these interactions to span multiple rendering paradigms as they leverage squeeze and vibration for various motion types. For example, the toy on the left combines the primitive motions of Fig. [7.6](#), push, pull, and twist, into a single context. The pistol features unimanual haptic effects, such as squeeze for trigger pull and vibration for firing, and bimanual effects such as squeeze forces when the player uses the opposite hand to cock the slide or twist off the silencer. One would expect the illusion of wrist squeeze to break down when presented in a way that forces users to make frequent context switches in a short amount of time. However, we find that users generally do not mind or notice that the haptic feedback, regardless of their motion or the visuals, is the same for each point of interaction.

## Chapter 8

### Perception of Mid-Air Stiffness

Chapter 7 demonstrated the versatility of referred haptic feedback for hand and finger interactions in VR. By leveraging Tasbi's vibrotactile and squeeze feedback in conjunction with visual pseudo-haptic control-to-display (C/D) illusions, we were able to convey linear and rotational stiffness, mass, textures, and others effects. In this chapter, we take a deeper look into the button interaction in Fig. 7.5 and the perception of mid-air stiffness.

The button interaction has been widely demonstrated internally and publicly at conferences to experienced hapticians and novices alike. Qualitative feedback we have received from users has been positive, with many suggesting that the interaction feels eerily similar to a physical button despite there being no button at all. One user describes the interaction nicely: *"It doesn't feel exactly like a real button, but my brain is able to immediately make sense of it. This is just how VR buttons feel!"* With most individuals expressing similar experiences, our initial assumption that squeeze and C/D pseudo-haptics simply provide a *metaphor* for stiffness shifted to the idea that these stimuli may induce a genuine perception of stiffness. Thus, the goals of this chapter are twofold:

1. To identify if and how individuals *map* stiffness conveyed through wrist squeeze and/or C/D manipulation to actual, physical stiffness.
2. To identify if individuals *integrate* both wrist squeeze and C/D manipulation in their estimation of stiffness when simultaneously presented, and to more broadly determine if there is a benefit to providing both stimuli.

## 8.1 Background

### 8.1.1 Evidence Supporting Referred Squeeze

Few studies have explored referred squeeze to convey virtual physical quantities. Perhaps the only example that can offer any guidance is from Mitsuda [158], who investigated how well pneumatic pressure feedback can be used to communicate the sense of holding objects. The author found that providing between 2 to 6 kPa pressure to the forearm resulted in a sensation that was equivalent to supporting between 4 and 10 N of weight. Though he only tested three pressures with very few repetitions, the relationship between pressure and perceived weight was found to follow a linear trend (Fig. 8.1). His paper also explored using pressure to convey the deformation of virtual objects when pushed, however he did not compare this to a physical reference and the results were only supported by questionnaires.

These results do not suggest how users may perceive stiffness, but they do support the idea that individuals can remap physical sensations to referred squeeze stimuli. Other researchers have explored stiffness rendering with referred squeeze feedback. Sarac et al. [62] performed a virtual stiffness discrimination task where contact forces at the virtual fingertip were conveyed through normal forces at the wrist. They used a direct mapping




Figure 8.1 : Mitsuda showed there to be a linear relationship between pressure applied to the forearm and the sense of weight. Figure adapted from [158].

from the simulated contact force to the applied wrist force. However, because their device was position controlled, they had to make broad assumptions about the impedance of wrist tissue to estimate the amount of displacement needed to render the desired force. Further, their experiment only tested discrimination accuracy of the cue, which says nothing about how subjects actually perceived the stimuli and its relation to stiffness. Moriyama et al. [60] used normal wrist squeeze to convey contact forces when pushing virtual objects, but the results of their study are not particularly useful since they chose to evaluate the interaction using their own metric of *strangeness*.

### 8.1.2 Evidence Supporting Pseudo-Haptics

Lécuyer et al. showed that subjects can discriminate stiffness of two virtual objects when a visual pseudo-haptic technique is applied, as well as discriminate stiffness between a pseudo-haptic stiffness and a physical stiffness [159]. Their study did not necessarily quantify a range of pseudo-haptic stimuli that can be used to simulate particular stiffnesses, and their implementation was not a mid-air interaction, requiring the user to physically interact with an isometric input device to displace the stiffness element. In a similar study, Srinivasan et al. [160] showed that when the visual deformation of a spring was discrepant with finger movement, vision altered subjects' estimate of the spring's stiffness.

Samad et al. provided the first quantification of the range of C/D ratio that can be used to simulate weight in VR [148]. In their experiments, they tasked subjects with adjusting an amount of physical weight until it was perceptually equivalent to that of a C/D modulated cube. Their results showed an almost linear relationship between the change in C/D ratio and the perceived change in mass of the cube (Fig. 8.2). Though this study also dealt in the perception of weight like that of Mitsuda's squeeze study, it provides evidence that users can map pseudo-haptic stimuli to physical analogues.




Figure 8.2 : (a) Samad et al. tasked users with adjusting the weight of a physical cube (blue) to a C/D modulated cube (brown). Changing the C/D ratio of the user's hand height in VR (b) altered the perceived mass of the cube (c). Figure adapted from [148].

### 8.1.3 Multisensory Integration

To our knowledge, no one has yet investigated using both referred squeeze feedback and visual pseudo-haptics to convey virtual stiffness, so it is not known how individuals may integrate both stimuli. However, many studies have explored how humans integrate haptic and visual stimuli in general. Famously, Ernst and Banks demonstrated that individuals integrate visual and haptic estimations of height by way of maximum likelihood estimation (MLE) [161]. If  $\hat{S}_V$  and  $\hat{S}_H$  are the visual and haptic estimates of an environment variable, then the MLE  $\hat{S}$  is their weighted sum, where the visual and haptic weights  $w_V$  and  $w_H$  can be determined from the variance in the unimodal estimates,  $\sigma_V^2$  and  $\sigma_H^2$ :

$$\hat{S} = w_V \hat{S}_V + w_H \hat{S}_H \quad \text{with} \quad w_i = \frac{1/\sigma_i^2}{1/\sigma_V^2 + 1/\sigma_H^2} \quad (8.1)$$

The final estimate  $\hat{S}$  is thus situated somewhere in between the visual and haptic estimates, as shown in Fig. 8.3, with lower variance than the unimodal variances given by:

$$\sigma_{HV}^2 = \frac{\sigma_V^2 \sigma_H^2}{\sigma_V^2 + \sigma_H^2} \quad (8.2)$$




Figure 8.3 : The estimation of an environment variable  $S$  can be assumed to be the integration of the unimodal visual and haptic estimators  $S_V$  and  $S_H$  by the maximum likelihood estimate rule. Figure adapted from [161].

Others have shown that MLE applies to other visuo-haptic or visuo-proprioceptive tasks [162, 163, 164, 165, 166], but not in all cases, such as when stimuli are incongruent [167] or affected by long-term influences (i.e. priors) [168]. Nonetheless, integration of haptic and visual stimuli by MLE provides a reasonable initial hypothesis on how users may integrate referred haptic feedback and visual pseudo-haptic stimuli. Since both stimuli could be considered noisy proxies for stiffness, then we hypothesize that users will make a better estimation when both are combined.

## 8.2 Experiment 1: Equating Mid-Air and Physical Stiffness

Our first experiment tasked subjects with adjusting the stiffness of a physical button until it was perceptually equivalent to the stiffness conveyed by a mid-air button across three conditions: Haptic-only, H (i.e. squeeze-only), Visual-only, V (i.e. C/D-only), and their combination, Haptic-Visual, HV. The experiment was loosely modeled after Mitsuda's [158] and Samad's [148] studies.

### 8.2.1 Physical Button Apparatus

A physical, variable stiffness button (VSB) was constructed and served as the physical stiffness comparison (Fig. 8.4). The VSB was driven by a Maxon RE-25 motor and capstan cable mechanism, similar to that found on Phantom haptic displays [169], with a transmission ratio of 0.105 mm/deg. Closed loop current control was accomplished via an Advanced Motion Controls AB15A100 servo drive and a Quanser QPIDE DAQ interface sampled at 2 kHz on the host PC. After gravity (0.59 N) and kinetic friction (0.18 N) of the button were appropriately compensated for with feed-forward control, a proportional-derivative (PD) position controller allowed for setting the desired stiffness  $k_p$  and damping  $k_d$  of the button. The physical button was able to simulate stiffnesses ranging from  $k_p = 5$  N/m to 400 N/m before over-drawing current from the power supply. Throughout the experiment, the button damping was computed such that the button was always critically damped:  $k_d = 2\sqrt{k_p m}$ , where the button mass  $m$  was determined to be 0.06 kg. It is worth noting that this apparatus served an auxiliary role, as it was also used to calibrate the force sensor in Tasbi to each user (explained below and shown in Fig. 8.4-b).

### 8.2.2 Virtual Environment

Subjects performed the experiment in a VR environment created inside of the Unity sandbox environment described in Chapter 7. An Oculus Rift CV1 served as the VR HMD, and Oculus Touch controllers were used to track subjects' hand position and gestures. Within the environment, subjects were presented with two visually identical buttons placed side-by-side. The button on the left represented the physical stiffness and was visuo-spatially aligned with the physical button apparatus so that when subjects pressed the button in VR, they felt the button in reality. The button on the right represented the mid-air stiffness, and displayed one of the three haptic rendering methods as detailed in the following section.






Figure 8.4 : Variable Stiffness Button (VSB) and Calibrator – (a) Image of the apparatus. (b) The apparatus in calibration mode, which serves to calibrate Tasbi’s internal force sensor to the user’s arm. (c) The apparatus in button mode, which was used as the physical stiffness comparison in Experiment 1.

Both buttons were 50 mm in diameter and 35 mm tall. Their color was nominally a shade of green, but would change to pink if displaced beyond the table surface, indicating to the subject to stop pushing.

### 8.2.3 Tasbi Calibration Procedure

Each session began by fitting and calibrating the Tasbi bracelet to the user (Fig. [8.5](#)). To ensure consistent stimulus delivery, 3M Transpore and double-sided mounting tape were used to secure the bracelet to the subject’s arm approximately 6 cm behind the ulnar styloid process. This location is slightly further back than that for typical wrist-watches, but offers a more circular cross-section and less dense stimulation site, which enhance the performance and perception of Tasbi’s squeeze feedback, respectively.

Once Tasbi was securely mounted, the subject placed their wrist and Tasbi under the calibration applicator. The apparatus delivered a sinusoidal force profile sweeping from 0 to 15 N to the topside of Tasbi. Voltage measurements from Tasbi’s internal capacitive force sensor were calibrated against the applied force. Following this, subjects removed their arm




Figure 8.5 : Tasbi Calibration Procedure – (a,b) Tasbi is mounted to the subject’s wrist using 3M mounting tape and Transpore medical tape. (c) The VSB calibrator applies a known load profile through Tasbi to which Tasbi’s internal sensor is calibrated against.

from the calibrator, and Tasbi was verified to render the full range of 0 to 15 N to the user’s arm. More information regarding Tasbi force sensing, calibration, and closed-loop control can be found in Chapters [3](#) and [4](#).

#### 8.2.4 Subjects

A total of 12 subjects (age:  $M = 22$ ,  $SD = 2.9$ , 8 female) completed the experiment. Subjects were recruited from the Rice University undergraduate and graduate student bodies under Rice University IRB protocol #IRB-FY2020-43. With the exception of a single subject, none had any experience with squeezing haptic displays, and all reported no or very limited experience with VR systems.

#### 8.2.5 Experimental Procedure

Using a method of adjustments, the experiment tasked subjects with adjusting the stiffness of the physical button on the left until it was perceptually equivalent to stiffness depicted by the mid-air button on the right. The experiment was divided into three experimental




Figure 8.6 : Experiment 1 – The virtual buttons and environment as viewed from the subject's VR perspective. The left button was the real physical button, and was spatially aligned with the physical button apparatus. The right button was the mid-air button, rendered through Tasbi squeeze force and/or C/D manipulation. Depending on the current condition, buttons were either visually opaque with observable displacement, or they were transparent with no observable displacement, acting instead as a volume.

blocks, conditioned on the haptic rendering or stimuli method of the mid-air button. Each conditional block tested four stimuli levels.

**Visual Condition (V):** The stiffness of the mid-air button was conveyed only through a unimodal C/D stimulus as described in Section 7.1.3. The four stimuli levels tested were  $C/D = 3, 4, 5,$  and  $6$ . With regards to the equations in Section 7.1.3, the proxy hand stiffness  $k_v$  was set to a constant  $50 \text{ N/m}$ , and the mass of the button  $m_b$  was set to  $0.06 \text{ kg}$ , equal to the mass of the physical button. Like the physical button, the mid-air button was always critically damped given the desired C/D ratio and resulting computation of the button stiffness  $k_b$ . The choice to critically damp both buttons was made so that subjects would not inadvertently use oscillatory motion to assess stiffness similarities.

**Haptic Condition (H):** The stiffness of the mid-air button was conveyed only through

a unimodal wrist squeeze stimulus and the C/D simulation was disabled. In this condition, both the mid-air button and the physical button were made visually transparent and static, so that subjects could see their fingers move through the button volumes, but *could not* see the buttons displace (see Fig. 8.6-c,d). The effect minimized subjects' usage of visual information and forced them make the comparison based on what they felt. The squeeze stimulus force was proportional to the amount of finger penetration into the button volume, reaching a maximum force at  $x_{\text{proxy}} = 35$  mm. The four stimuli levels were defined by this maximum squeeze force,  $F_{\text{squeeze,max}} = 3, 7, 11, \text{ and } 15$  N. Through self-piloting, these levels were identified to be roughly equivalent to the C/D levels in terms of perceived stiffness.

**Haptic-Visual Condition (HV):** The stiffness of the mid-air button was conveyed through the bimodal combination a C/D stimulus and a wrist squeeze stimulus. The same stimuli levels from unimodal conditions were used and presented congruently (e.g. C/D = 4 with  $F_{\text{squeeze,max}} = 7$  N, etc.). In this condition,  $F_{\text{squeeze}}$  was proportional to the C/D button displacement  $x_b$  instead of  $x_{\text{proxy}}$ , as was necessary in the Haptic condition.

Each conditional block consisted of 64 trials. In each trial, subjects were presented with a mid-air button displaying a fixed stimulus level. The starting stiffness of the physical button was randomly set near the low end of its rendering range at 25 N/m, or the high end at 375 N/m. Subjects assessed both buttons, and then used the joystick of the Oculus Touch controller to change the stiffness of the physical button until it was perceptually equal to the mid-air button. Subjects were allowed to transition between buttons freely, but were required to complete the adjustment in 25 seconds. Subjects were given the option to advance to the next trial once they were confident both buttons were equivalent.

The first 16 trials of each block were practice trials and presented the same stimulus level, which was taken from the center of the stimuli ranges above (e.g. C/D = 4.5 and/or

$F_{\text{squeeze}} = 9 \text{ N/m}$ ). The remaining 48 trials randomly presented one of the 4 tested stimuli levels, with each repeated 12 times. Subjects were given a short break in between each conditional block. Importantly, the presentation order of the three blocks was randomized between subjects so that each of the 6 possible orders were equally represented. Overall, the experiment lasted approximately two hours for each subject.

### 8.2.6 Data Analysis

The responses for each subject across stimuli level and condition were averaged to find their mean estimated stiffness values. To identify subjects' mapping from mid-air stiffness to physical stiffness, their stiffness response to the four stimuli levels within each condition were fit with a linear regression. Fits were applied at both the individual subject and group levels.

To determine if a difference existed between the three conditions, subject means were compared between the bimodal HV and unimodal H and V conditions separately using 2-way repeated measures ANOVAs with factors of Condition and Level. In addition to mean values, a similar set of ANOVAs were performed for the measure of subject standard deviation to assess whether subjects were more consistent with their response in any particular condition. Sphericity violations were treated with a Huynh-Feldt correction where needed.

Four level-independent metrics were used to compare group differences between conditions using paired samples student's t-tests. The first three metrics were simply the parameters obtained from subject linear regressions: slope  $m$ , bias  $b$ , and quality of fit  $R^2$ . The last metric was the Fréchet distance between each subject fit and the mean *group* fit. Fréchet distance is a measure of similarity between two discrete curves [170], and was used to assess whether subjects behaved more similar to each other in any one condition.

A subject was identified as an outlier in any one condition if two or more of their Level




Figure 8.7 : Experiment 1 Main Results – The top row shows the mean response and fits for individual subjects as well as the group level fit for each condition. The two identified outliers are denoted by dashed lines. Error bars are a 95% confidence interval. The bottom row provides box plots of subject means for every level in each condition, with the same outliers denoted by hollow circles.

responses deviated more than 1.5 from the group inner-quartile range (IQR) for that Level. With this criteria, two subjects were identified as outliers; one in the V condition and one in the HV condition (as denoted by the dashed lines in Fig. 8.7). Their responses for all four levels in the respective condition were replaced with the group mean for that condition. Thus, 8 of 144 means (12 subjects x 3 conditions x 4 levels) were replaced (5.6% of the total data).

## 8.2.7 Results

The main results are shown in Fig. 8.7. Each plot in the top row depicts individual subject means and fits, as well as the group means and fit for all three conditions. Error bars are a 95% confidence interval.

We observe greater mean stiffness responses in the HV condition than either unimodal condition for most of the presented Levels. The ANOVA between the HV and H conditions on the measure of subject means showed a significant effect of Level ( $F(3,33) = 164$ ,  $p < .001$ ) but not Condition ( $F(1,11) = 4.14$ ,  $p = .067$ ). However, the ANOVA between the HV and V conditions shows significant effects of both Level ( $F(3,33) = 184$ ,  $p < .001$ ) and Condition ( $F(1,11) = 13.6$ ,  $p = .004$ ). Interaction contrasts show a significant difference in the linear slope across Level in both ANOVA comparisons ( $F(1,11) = 8.84$ ,  $p = .013$  for HV-to-H and  $F(1,11) = 7.20$ ,  $p = 0.02$  for HV-to-V). When comparing subject standard deviations, no significant effect of Condition could be observed, though significant effects of Level were present ( $p < .001$  in both unimodal-bimodal comparisons).

Fig. 8.8 compares the fit metrics across all three conditions. The mean slopes of all subjects were found to be 37.3, 47.2, and 58.8 N/m/Level for the H, V, and HV conditions, respectively (Fig. 8.8-a). The paired t-tests between the bimodal and unimodal conditions shows a significant difference in slope for both the HV-to-H comparison ( $t(11) = 3.0$ ,  $p = .01$ ) and the HV-to-V comparison ( $t(11) = 2.7$ ,  $p = .02$ ). Conversely, no significant differences of bias were found between conditions, as evidenced by the degree of variability among subjects (Fig. 8.8-b). The mean  $R^2$  values of subject fits were 0.81, 0.95, and 0.98 (Fig. 8.8-c). A significant difference was found between the H and HV conditions ( $t(11) = 2.7$ ,  $p = 0.02$ ) but not between the Visual and HV condition. The Fréchet distance was found to be significantly less in the HV condition than in both the H condition ( $t(11) = 3.2$ ,  $p = .009$ ) and V condition ( $t(11) = 3.2$ ,  $p = .002$ ), indicating more conformity among subjects in the HV condition.




Figure 8.8 : Four metrics comparing the fits obtained from Experiment 1 across conditions. Error bars represent a 95% confidence interval. Significance obtained from the bimodal-unimodal t-tests are denoted with bars above the plots.



## 8.2.8 Discussion

### Were subjects able to map physical stiffness to mid-air stiffness?

The results indicate that most subjects were able to titrate the stiffness of the physical button to the mid-air stimuli in all three conditions. This is supported by both the observation that stiffness responses monotonically increase with the stimuli levels (i.e. positive slope) and the quality of the fits obtained from the data *for each subject*. The black lines in the top row of Fig. 8.7 show the linear fit to the group level data, which can be equally obtained by fitting to either all of the subject means (colored dots), or the group means (black dots). When fit to the group mean, we obtain comparable  $R^2$  values greater than 0.9 in all cases (actual values are provided in the legend). However, this is misleading since the same regression applied to the subject means yields far inferior fit quality for the unimodal conditions (0.408 and 0.584 for the H and V conditions, respectively). Clearly there are great differences in how subjects responded to the unimodal stimuli.

Interestingly, we note far less dispersion between subjects in the bimodal condition. This is supported not only by visual inspection, but by the significant reduction in Fréchet distance in the Bimodal case. To further appreciate this, consider the six individual subjects shown in Fig. 8.9. In each plot, the subjects' means and fit are shown along with the *group* level fit obtained from Fig. 8.7. Subjects 6, 9, and 10 show a similar phenomenon in that their unimodal responses are relatively similar to the group level responses (i.e. the black line), as well as to each other (i.e. the stiffness reported for the four stimuli levels in the H condition are comparable to those reported in the V condition). Further, their bimodal response almost agrees with the group level response, and in the case of 6 and 9, there is less variance than the unimodal conditions, which is surprising considering the slope and thus the range of the bimodal condition is largest. Where things become interesting




Figure 8.9 : Six individual subjects from Experiment 1 shown in comparison to the group fits. The order in which they completed condition blocks is given in the bottom right corner. Error bars represent one standard deviation in the subject's responses for that level.

is when we examine individuals whose unimodal responses do not agree, either with the group response or with each other. Consider subjects 4 and 11. The Visual response of subject 4 exhibits a similar slope, but substantially more bias than the group level response, while their haptic response is well described by the group level trend. Subject 11 on the other hand, displays more bias in their haptic response, and their ability to make sense of the unimodal visual condition is questionable. Surprisingly, however, the Haptic-Visual response of both of these subjects ultimately conforms to the group level response. Subject 3 is arguably even more unique, in that their unimodal responses are both below the group trend, and yet still their bimodal response conforms to the group response.

### **Did subjects *actually* map stiffness?**

Any easy criticism of this experiment might be that subjects did not map the actual stimuli of the mid-air button to the physical button, but rather mapped a *percentage* of the stimuli range to a percentage of physical button's stiffness range. This is often a valid concern for experiments that employ the method of adjustments as we have done here. However, we do not believe this to be the case for three reasons.

First, we have no anecdotal evidence of this occurring. We casually asked subjects after the experiment if this was the approach that they took, and none indicated so. We further probed subjects on how many stimuli levels of the mid-air button they thought there to be. Answers varied wildly, some suspecting as many as 10 total levels. And so, it seems unlikely that that subjects did or would have been able to assess the stimuli range of the conditions. This alone, however, is not sufficient evidence to support our argument.

The second point that can be made is that we very rarely see stiffness responses above 250 to 300 N/m. If subjects had been matching ranges, then it would seem probable that we would see the mean response for the maximum stimuli level of each condition near the

400 N/m mark, i.e. the maximum stiffness the physical button was capable of rendering. Since we do not, then it is evident that subjects did not arbitrarily associate the maximum stimuli level with the maximum stiffness that was possible.

Finally, recall that each condition block began with 16 practice trials that repeatedly presented a constant stimulus level centered at the range of stimuli presented in the main experiment. If these practice trials are plotted against the fits obtained from the experimental trials, we find that they are almost perfectly described by the fit at the group level (green markers in Fig. 8.7), and typically well described even at the individual level (green markers in Fig. 8.9). In the case of the first presented block, whichever it may have been, it is impossible for the subject to have made any assessment about the range of stimuli that would be subsequently provided. It could be argued that subjects may have learned the stimuli levels between the bimodal condition and the unimodal conditions, and vice versa, so as to influence their practice trials in subsequent blocks. This would have been quite challenging to do, however, since 1) a timeout period was placed between blocks to allow for the residuals to expire, and 2) subjects were not aware of the fact that the levels in the unimodal and bimodal conditions were the same. Thus, it seems quite probable that subjects did in fact associate the stimuli provided by the mid-air button with their perception of the physical button's stiffness.

### **What were subjects comparing to make their assessment?**

The experiment was not specifically designed to determine *how* subjects make comparisons, and so further experimentation would be required to form a complete theory. However, we can make a few logical conclusions and assess qualitative feedback from subjects.

In the haptic only condition, proprioception between the mid-air and physical button was held constant (i.e. the subjects' fingers displaced the same amount for both buttons).

Therefore, they could only have used a cue of either force or pressure to make their assessment. In questioning subjects on their methods, two common approaches arose. Logically, some subjects suggested that they were able to compare the force felt on their fingertip when pressing the physical button to the force delivered to the wrist by Tasbi when pressing the mid-air button. Unexpected though, are the subjects that suggested making comparisons of the muscle tension they felt in their wrist while exerting force on the physical button to the gross pressure delivered by Tasbi to those same muscles.

In the visual only condition, no haptic feedback was provided for the midair button so subjects could not use force or pressure information to make their assessment. The responses from subjects would indicate that they themselves did not know how they were able to make the comparison. The most common suggestion was that they used temporal information, i.e. adjusting the physical button until it took the same amount of time to press as the mid-air button. The perception of time is generally considered poor and susceptible to illusions [171, 172], even for short durations, [173]. If subjects did use temporal information, then this might explain some of the variability we see in our data. Some subjects suggested that they leveraged visual information, so that buttons *appeared* to displace at equal rates. Others suggest they compared how difficult it was to displace the buttons, or the amount of *effort* it took. And at least two subjects explicitly said they compared the *force* it took to displace the buttons, which clearly indicates a lack of terminology to describe their experience.

The last two points may indicate that subjects leverage their perception of *work* done on the buttons, as suggested in a similar study on pseudo-haptic weight by Samad et al. [148]. They presented an equation that accurately predicted the perceived mass of lifting a cube while under a visual C/D manipulation in VR. If  $m_{prop}$  is the proprioceptive sensation of mass,  $h_{prop}$  is the proprioceptive change in height,  $h_{vis}$  is the visual change in height,

work done is  $W$ , and  $\lambda$  is the C/D ratio, the perceived mass can be derived assuming that the signals from vision and proprioception are integrated in an statistically optimal fashion [161]:

$$h_{per} = \alpha h_{prop} + \beta h_{vis} , \alpha + \beta = 1 \quad (8.3)$$

$$m_{per} = \frac{W}{gh_{per}} = \frac{m_{prop}h_{prop}}{h_{per}} \quad (8.4)$$

$$m_{per} = \frac{m_{prop}}{\alpha + \beta\lambda} \quad (8.5)$$

Similarly, we can derive an equation for perceived C/D modulated stiffness,  $k_{per}$ , based on work done, from proprioceptive stiffness  $k_{prop}$ , proprioceptive displacement  $x_{prop}$  and visual displacement  $x_{vis}$ :

$$x_{per} = \alpha x_{prop} + \beta x_{vis} , \alpha + \beta = 1 \quad (8.6)$$

$$k_{per} = \frac{2W}{x_{per}^2} = \frac{k_{prop}x_{prop}^2}{x_{per}^2} \quad (8.7)$$

$$k_{per} = \frac{k_{prop}}{(\alpha + \beta\lambda)^2} \quad (8.8)$$

Unfortunately, this is where drawing parallels between our experiment and theirs ends, because in our experiment there is no proprioceptive sense of stiffness  $k_{prop}$  (i.e. we are not modulating the stiffness of a real button with C/D; we are creating the sense of stiffness in mid-air). The authors did not provide insight on how this approach may work on entirely virtual objects such as ours, and if we simply assume that  $k_{prop}$  is zero, then by equation 8.8 the perceived stiffness  $k_{per}$  should also be zero, which clearly it is not. It is also not obvious how the haptic stimuli from Tasbi would be integrated into the idea of work done by the participant in the case of the bimodal condition. Nonetheless, the idea of using work as the foundation for answering this question is interesting, and should be investigated further.

### 8.3 Experiment 2: Discriminating Mid-Air Stiffness

Although the results of Experiment 1 indicated that subjects were more consistent at the group level in the bimodal condition compared to the unimodal conditions, it did not necessarily show that subjects were independently more consistent in their estimation of stiffness when receiving both C/D and squeeze haptics. Furthermore, we would like to compare the relative performance of subjects in all three conditions. One metric of performance we can compare across conditions is the just noticeable difference (JND) of button stiffness. The smaller the JND, the more attuned individuals are to the stimuli presented, and thus their ability to make accurate judgments of stiffness is increased. We designed a follow-up experiment that tested subjects' ability to discriminate the stiffness of *two mid-air buttons* in the same three conditions as the first experiment: Visual (V), Haptic (H), Haptic-Visual (HV).




Figure 8.10 : Experiment 2 – Subjects were presented with two mid-air buttons and forced to choose the one they thought was stiffer. Subjects were constrained to two button presses or five seconds to make their selection. The same mid-air button conditions from Experiment 1 were used: Visual only, Haptic only, or the combination of the two, Haptic-Visual.

### 8.3.1 Subjects

Twelve new subjects were recruited for Experiment 2 (age:  $M = 25$ ,  $SD = 4.8$ , 3 female). One subject had to be excluded from all analyses for failing to follow our directions. Subjects were pooled from the Rice University undergraduate and graduate student bodies under Rice University IRB protocol #IRB-FY2020-43. Only one subject reported any experience with squeezing haptic displays, and none had significant prior experience with virtual reality devices.

### 8.3.2 Virtual Environment

The same virtual button environment from in Experiment 1 was used, but the left physical button was replaced with a second mid-air button. The color and opacity of the mid-air buttons behaved in the same manner as in Experiment 1, where the button changed from green to red to indicate that the subject should stop pressing, and was a transparent volume in the haptic-only condition. Tasbi was calibrated to subjects' arms using the method previously described in Section [8.2.3](#).

### 8.3.3 Experimental Procedure

Using the method of constant stimuli and a two alternative forced-choice (2AFC) procedure, the experiment tasked subjects with the selecting the stiffer of two mid-air buttons. As before, each of the three mid-air button conditions (V, H, and HV) was presented in a separate experimental block. Each block consisted of 8 practice trials and 220 test trials. In each trial, two visually identical buttons were presented on the left and right, and subjects were allowed either five seconds or two presses of each button, whichever came first, to decide on which button was the hardest to press. Subjects were instructed to alternate between the buttons and to make their selection as soon as they were confident. Subjects



made their selection by moving the joystick on the Oculus controller to the left or right (Fig. 8.10). One button, the Standard, presented the same stimulus in every trial and was evenly randomized to appear on either the left or right side. The other button, the Comparison, displayed one of 11 stimuli levels: 5 below the Standard level, 5 above the Standard level, and the Standard level itself. Each Comparison level was repeated 20 times, and the presentation order was randomized for each subject.

The stimuli levels for the H and V conditions were derived from the results of Experiment 1 by using the group mean estimated stiffness as a proxy to perceptually match the squeeze and C/D levels to each other. Noting the similarities between the first three levels of the H and V conditions in Experiment 1, we chose to test a proxy stiffness range of 50 to 150 N/m (Fig. 8.11). Thus, the chosen comparison max squeeze levels for the H condition in this experiment were 3.0, 3.8, 4.6, 5.4, 6.2, 7.0, 7.8, 8.6, 9.4, 10.2, 11 N with a Standard level of 7.0 N. The Comparison C/D levels for the Visual condition were 3.0, 3.2, 3.4, 3.6, 3.8, 4.0, 4.2, 4.4, 4.6, 4.8, 5.0 with a Standard level of 4.0. The HV condition presented the unimodal levels above congruently. Subjects were given a 5 minute break in between condition blocks, and the presentation order of blocks was randomized between subjects in a counterbalanced manner. The experiment lasted 2 hours for each subject.

#### 8.3.4 Survey

A short survey was presented at the end of each block. The questions, listed in Table 8.3.4, highlighted subjects' disposition toward the buttons presented in that block. Subjects responded to these questions using a continuous slider on a scale of Strongly Disagree (0.0) to Strongly Agree (1.0). The same questions were presented after each block, with the order randomized.




Figure 8.11 : The estimated stiffness of the Haptic and Visual conditions from Experiment 1 were used as a proxy to derive perceptually equivalent stimuli levels for Experiment 2. Eleven comparison levels for each condition were interpolated based on a stiffness range from 50 to 150 N/m. The bold lines indicate the central Standard level.

Table 8.1 : Survey Questions for Experiment 2

Number	Label	Question
Q1	Believable	The buttons were believable
Q2	Realistic	The buttons were realistic
Q3	Immaterial	The buttons were immaterial
Q4	Pleasant	The buttons were pleasant
Q5	Natural	My interaction with the buttons felt natural
Q6	Intuitive	My interaction with the buttons was intuitive
Q7	Location	The virtual hands appeared in the same location as my hands
Q8	Body	The virtual hands seemed to belong to my body
Q9	Confident	I was confident in my selections

### 8.3.5 Data Analysis

The psychometric curves for each subject in each condition were determined by fitting a general linear model with a logit link function to the proportion of times the subject indicated that each Comparison level was greater than the Standard level. From this, we obtained estimates of the just noticeable difference (JND) and point of subjectively equality (PSE). Two additional metrics were used to assess performance: the percentage of times subjects correctly identified the stiffer button (excluding Standard-Standard comparisons) and the total number of buttons presses. Few button presses would indicate that subjects were more confident during the selection process. One-way repeated measures ANOVAs were used to identify differences between all three conditions for each of the four metrics. No sphericity violations were observed. Survey responses were analyzed using a Wilcoxon signed-rank test, and tested whether the median differed significantly from 0.5 (a neutral response).

### 8.3.6 Results

The main results are shown in Fig. 8.12, where an aggregate psychometric curve is given for the entire study group. By visual analysis alone, we can see that the slope of the HV condition is greater than the slopes of the two unimodal conditions, indicating a smaller difference threshold. The ANOVA for JND (Fig. 8.13-a) reveals a significant effect of condition ( $F(2,20) = 5.5$ ,  $p = .012$ ). A post-hoc using a Bonferroni correction shows a significant difference between the Visual and HV conditions ( $p = 0.008$ ), but not between the H and HV ( $p = .076$ ) or V and H conditions ( $p = 1.000$ ). The PSE appears to be biased for the unimodal conditions, and less so for the HV condition (Fig. 8.13-a). However, we find no significant effect of condition on PSE ( $F(2,20) = 2.2$ ,  $p = .089$ ). A significant effect of condition was found for the percentage of correct responses ( $F(2,20) = 8.879$ ,




Figure 8.12 : Aggregate psychometric curves obtained from Experiment 2. The Bimodal condition was found to produce a significantly lower JND compared to with Visual only condition, but not compared with Haptic only condition.

$p = .002$ ). The post-hoc showed that subjects achieved a significantly higher percentage of correct responses (Fig. 8.13-c) in the HV condition compared to the V condition ( $p = .002$ ), but not compared to the H condition ( $p = .071$ ). Finally, we note a significant effect of condition on the number of presses subjects made ( $F(2,20) = 4.2$ ,  $p = .03$ ), with subjects making fewer presses in the Bimodal and Visual conditions (Fig. 8.13-d). However, the post-hoc failed to show a significant difference between conditions ( $p = .076$  for HV-to-H,  $p = 1.000$  for HV-to-V, and  $p = .078$  for V-to-H), likely due to insufficient statistical power.

### 8.3.7 Discussion

#### Did subjects perform best in the Haptic-Visual condition?

At the group level, the results show that both JND and percentage of correct responses were significantly better in the HV condition compared to the visual-only condition, but not when compared to the Haptic only condition. A larger sample size with greater statistical power might ultimately reveal that the HV condition produces significantly better results than the Haptic only condition as well, considering the degree of separation in the means of difference thresholds shown in Fig. 8.13-a. Overall though, it is clear that subjects were able to integrate both the visual and haptic stimuli to improve their estimation of mid-air button stiffness. Because the HV-to-V comparison is significantly different and the HV-to-H comparison is not, it could be argued that subjects relied more on the haptic stimuli to make their assessment of button stiffness. That subjects showed slightly better performance in the haptic condition versus the visual condition lends credit to this theory. It is interesting, though, that subjects required more button presses to make a determination of stiffness in the haptic condition (Fig. 8.13-d) compared to the visual condition, indicating a lack in confidence. It would seem that subjects either underestimated their ability to use squeeze feedback, and/or overestimated their ability to use visual feedback.




Figure 8.13 : Four Metrics Comparing Conditions of Experiment 2 – (a) The mean just noticeable difference threshold for each condition. (b) The point of subjective equality, or bias from the Standard level of 100 (solid black line), for each condition. (c) The percentage of times subjects correctly selected the stiffer button. (d) The average number of button presses taken in each condition, with 800 being the maximum (solid black line). Error bars denote a 95% confidence interval.

### How did subjects integrate the haptic and visual stimuli?

Based on [16], we expected to find that subjects integrated the visual and haptic stimuli by way of the MLE rule. In this case, the difference threshold of the bimodal condition can be estimated from the unimodal difference threshold through maximum likelihood estimation:

$$\hat{T}_{HV}^2 = \frac{T_V^2 T_H^2}{T_V^2 + T_H^2} \quad (8.9)$$

According to this rule, the bimodal estimates should have lower variance and thus smaller discrimination thresholds than either of the unimodal conditions. Indeed, if the calculation is applied to the difference thresholds obtained from the aggregate psychometric curves in Fig. 8.12 using values of  $T_H = 14.85$  N/m and  $T_V = 16.80$  N/m, we estimate the bimodal threshold as  $\hat{T}_{HV} = 11.13$  N/m which agrees well with the imperial value of  $T_{HV} = 11.02$  N/m. However, integration through maximum likelihood estimation does not perfectly describe all of our subjects individually.

Fig. 8.14 shows the responses from six subjects in each condition. Subjects 6 and 9 represent subjects who showed similar perceptual responses to the haptic and visual unimodal conditions, as we had expected to be the common case. In their bimodal condition, it is clear that they integrated the visual and haptic stimuli. Although the MLE predicted bimodal threshold does not perfectly match the empirical value, it is either very close, or in the case of 9, higher than their actual response. The same could be said for subjects 3 and 5, however, here we see a more noticeable discrepancy between in unimodal responses. Given by the estimated weights of vision  $w_V$  and haptics  $w_H$ , we can see that subject 3 weighed haptics greater, while subject 5 weighed visuals greater. Overall, we find that 7 out of 11 subjects weigh the haptic stimuli greater than the visual stimuli, a surprising finding in light of similar visuo-haptic studies [168]. Where MLE seems to break down, is when subjects

are very perceptive in *both* the haptic and visual conditions. Subjects 8 and 10 show no improvement in the bimodal conditions, as the MLE predicted threshold would have their response more closely resembling a step function.

MLE further fails to describe the PSE of many subjects. According to the rule, we should expect the PSE of the bimodal condition to be located in between the PSEs of the unimodal conditions, favoring that of the higher weighed modality. However, subjects 6, 9, 3, and 5 (as well as the aggregate curves from Fig. 8.10) clearly show the PSE of the bimodal condition *outside* the range of the unimodal PSEs with less or even no overall bias.

Though more analysis will be required, our results may be described by probability summation models. For example, Kuschel et. al [167] showed that subjects behave suboptimally when visual and haptic information is incongruent. Even though we matched haptic and visual levels based on means from Experiment 1, it is very likely that these levels were not well matched for all subjects. Indeed, some subjects explicitly said the visual and haptic stimuli did not seem to agree. Other possible explanations include causal inference models [174, 175, 133]. Causal inference models take into account bayesian priors (i.e. subjects' prior knowledge or expectations of stiffness), in contrast to MLE which assumes uniform priors (i.e. no prior expectations). Of course, one of the primary challenges with this type of approach is in knowing what the priors for mid-air stiffness even are.

### **What were subjects' dispositions toward the mid-air buttons?**

The survey results in Fig. 8.15 show that subjects rated the bimodal button higher in terms of its believability, pleasantness, and naturalness of interaction. Subjects indicated that the bimodal and visual-only buttons were roughly equivalent in terms of realism and higher than the haptic-only button. The bimodal and visual-only buttons were rated less immaterial than the haptic button, though this question may have been flawed considering the






Figure 8.14 : Six subjects from Experiment 2. Their thresholds for the unimodal Visual and Haptic conditions are  $T_V$  and  $T_H$ . The empirical and MLE predicted thresholds for the Haptic-Visual condition are  $T_{HV}$  and  $\hat{T}_{HV}$ . Estimated unimodal weights are  $w_v$  and  $w_h$ .




Figure 8.15 : Survey Results for Experiment 2. Asterisks indicate significant difference from a neutral response of 0.5. See Table [8.3.4](#) for the full questionnaire.

haptic button was in fact visually immaterial. Intuitiveness of interaction was positive for all three buttons, with the bimodal button rated slightly higher. The questions regarding location of hand and body ownership seem to indicate that subjects were largely unphased by the C/D manipulation and the discrepancy between their actual hand location and the rendered hand location caused. Finally, it would seem that the bimodal button inspired far more confidence in the selection process than either of the unimodal buttons.

### Similarities to Experiment 1

It is worth pointing out some parallels between Experiment 2 and Experiment 1. First, we note a similar a phenomenon where subjects seem to become more conformal in the bimodal condition. Fig. [8.16](#) shows the clustering of psychometric fits for all subjects in Experiment 2, where the bimodal fits are arguably more densely packed than either of the




Figure 8.16 : Psychometric functions from all subjects in Experiment 2 overlaid in each condition. The conformity of subjects in the bimodal condition is similar to the observations of Experiment 1 (see Fig. 8.7).

unimodal conditions. This is quite similar to what we see in Fig. 8.7. If this is indeed a real phenomenon, then it has major implications on the feasibility of providing consistent feedback to all users. It is also interesting to note that the mean fit biases of Experiment 1 (21.4, 7.1, 1.3 N/m for H, V, and HV conditions, respectively) to the biases of Experiment 2 (5.0, 2.9, and 0.9 N/m) show a similar trend where the haptic only condition is most biased, the visual-only trend shows less bias, and the bimodal condition has almost no bias.

## 8.4 Conclusion and Future Work

To summarize, the first study showed that subjects can make comparisons between actual stiffness and mid-air stiffness when conveyed through haptic-only wrist squeeze, visual-only C/D pseudo-haptics, and their bimodal combination. We find that subjects differ substantially in the unimodal cases, but become more conformal in the bimodal case. The second study showed that many subjects integrate both vision and haptics in their estimation of button stiffness, with the bimodal presentation of squeeze and C/D pseudo-haptics resulting in significantly smaller JNDs compared to the unimodal cases. Similar to the

first study, we find that subjects weigh the haptic and visual stimuli differently from one another. Contrary to our initial hypothesis, it does not seem that subjects integrate signals through classical assumption of maximum likelihood estimation, and so further exploration of integration models will be required.

The experiments presented here were fairly restricted in scope, and so further studies should be conducted. First, it will be important to determine if the relationships between the mid-air stimuli and physical stiffness found in Experiment 1 hold under different button conditions. For example, what will happen if the stroke length of the buttons is increased or decreased? It is likely that the observed fits are only valid within a narrow range of displacement. This is especially true in the visual-only C/D case where only a finite amount of visual-spatial incongruence can be tolerated. If it is found that the perception of stiffness is a function of displacement, then perhaps this could be accurately predicted by an overarching model. Similarly, it will be interesting to see if these models can be applied to stiffness in other contexts and, in particular, if they hold in the case of *tensile* stiffness (e.g. when applied to the pull-knob in Fig 7.6-b).

A second area of interest is in determining how to provide subject-specific visual and haptic stimuli so that they are perceptually matched. While Experiment 1 showed that some subjects felt the haptic and visual stimuli levels were equivalent in terms of stiffness (e.g. subjects 6, 9, and 10 in Fig. 8.7), others did not (subjects 4, 11, and 3). Despite our efforts to refine the levels for Experiment 2, its probable that the issue persisted. The effect that such sensory conflicts have is that they confuse users to the point that they must pick one modality to focus on instead of naturally integrating both. In such a case, the user's perception can only be as precise as their perception in the unimodal case.

Finally, our experiment ignored the effect that vibration may have on stiffness perception, as the vibrotactile stimuli from Tasbi during surface contact was the same for all

conditions and stimuli levels. If we wish to extend the ideas presented here toward rendering deformable materials with different properties, then vibration will obviously play an important role. Unfortunately, the complexity of designing an experiment around squeeze, vibration, and C/D manipulation is significant. Entirely different experimental paradigms will likely be needed to fully explore the perceptual space. Bayesian adaptive procedures seem like a promising path forward [176, 177].

## Chapter 9

### Conclusion and Future Work

This thesis has presented Tasbi, a haptic wristband featuring multimodal squeeze and vibration. The design makes significant strides toward realizing a compact, all-day wearable wrist interface, and to our knowledge is the smallest of its kind with an overall footprint of 50x30x15 mm and 120 g. As detailed in Chapter 3, we accomplished this through a novel tensioning mechanism that is both mechanically robust and squeeze efficient, producing peak normal forces up to 15 N without significant tangential losses. This mechanism allowed us to use a miniature DC motor which was important to both Tasbi's overall size and potential to be battery powered in the future.

In Chapter 4, we presented a responsive and accurate squeeze force control solution that makes use of a low-cost force sensing capacitor as well as incremental encoder velocity estimation. Across Tasbi's nominal force range of 0 to 10 N, the controller shows less than 5% tracking error, a 70 ms step response, and an overall bandwidth of 9.1 Hz. Unlike traditional methods of squeeze that leverage position control, force control offers the ability to provide consistent squeeze cues regardless of wrist impedance, size, or posture.

Chapter 5 detailed Syntacts, Tasbi's vibrotactile control framework. Development ultimately led to an open-source suite of vibrotactile software and hardware based on audio principles. The framework is purpose-built to control arrays of vibrotactors with extremely low latency. In addition, Syntacts includes a graphical user interface designed to synthesize and sequence cues, and spatialize them on tactile arrays. The Syntacts Amplifier easily integrates with the audio hardware to provide high-quality analog signals to the tactors without

adding excess noise to the system. Importantly, neither Syntacts software nor the Syntacts amplifier are required by each other; users can choose to mix Syntacts software with their own amplifiers, or use the Syntacts amplifier with their own software. We benchmarked the Syntacts system and showed that audio plus Syntacts is as, if not more, effective and flexible than commercially available solutions.

With Tasbi's unique ability to control directly for squeeze force, Chapter 6 presented three psychophysical studies of haptic perception. The first study tested for the just noticeable difference of wrist squeeze force, which we found to be 1.28 N averaged across all subjects. In contrast to vibration, we find that wrist size has no significant effect on this threshold. Our second study tested subjects' ability to identify stimuli from Tasbi's six vibrotactors under varying levels of static squeeze force, and offered the following insights: 1) vibrotactor identification rates are affected by radial location, with vibrotactors over bony areas of the wrist being significantly more difficult to discern; 2) subjects with large wrists perform significantly better than those with small wrists; and, 3) the amount of static squeeze has no significant effect on identification rates. The latter is a particularly important point for multimodal devices seeking to integrate squeeze and vibrotactile feedback. The final study showed that the amplitude of vibration at the wrist has a significant effect on identification as well, with higher amplitude vibrations proving to be more discernable.

With finalized Tasbi hardware and control and a basic understanding of how individuals perceive haptic cues at the wrist, we next explored using Tasbi for referred haptic feedback. In Chapter 7, we presented a novel multisensory rendering approach for VR hand interactions that leverages squeeze for continuous forces, vibration for discrete forces, and visual control-to-display illusions for proprioceptive information. A number of interaction examples were described, ranging in complexity from simple button presses, to bimanual

manipulations, to complex interactables with multiple rendering paradigms in play.

Finally, in Chapter 8, we took a deeper look into how individuals perceive referred haptic feedback for hand interactions, with a focus toward stiffness perception. Using the method of adjustments, we showed that individuals can map the true stiffness of physical buttons to the stiffness of mid-air Tasbi buttons as conveyed either through squeeze-only, visuals-only, or their bimodal combination. The results indicated that while individuals weigh each modality differently and show large variations between each other in the unimodal cases, they ultimately converge to a point of uniformity when both modalities are provided. Following this, we examined the possibility that the bimodal presentation of cues can increase stiffness discrimination at the individual level. Having subjects compare two mid-air buttons, we found that the just-noticeable difference threshold for stiffness is significantly higher in the bimodal case. Although the classical assumption of multisensory integration by maximum-likelihood-estimation (MLE) described our data at the group level, like the first experiment, we found that subjects vary drastically in their approaches to integrating haptic and visual information. MLE did not describe all subjects at the individual level, and we suspect that users may ultimately use approaches similar to probability summation. Interestingly, unlike many visuo-haptic experiments, our results suggest that individuals weigh haptic cues more importantly than visual cues during mid-air interaction, though both clearly have an overall impact on perception. Overall, the results of these studies as well as qualitative survey responses show that the bimodal presentation of squeeze and visual-illusions for referred haptic feedback greatly enhance user performance and the believability of interactions.



## 9.1 Tasbi Self-Containment

Though Tasbi is both compact and power efficient, great strides will need to be taken to realize a completely self-contained, untethered version of the device. The addition of on-board microcontrollers and wireless communication modules will obviously be required. However these additions pale in comparison to the challenge of introducing battery power. Several steps provide a plausible path forward. First, the efficiency of Tasbi's squeeze mechanism must be thoroughly analyzed and refined. In particular, it is highly probable that the power efficiency of the tensioning cord routing inside of Tasbi's main housing is suboptimal. Here, the cord makes several tight turns before exiting the housing, and by the capstan cable equation [178], we know that friction losses exponentially increase with wrap angle. Thus, the addition of small bearings or eliminating these inflection points entirely should greatly increase the tension output of the mechanism. This in turn will allow for a reduction in motor size, which has a two benefits to battery power: 1) smaller motors consume less power allowing for smaller batteries, and 2) the reduction in space consumption by the motor opens up room for a battery compartment. Although we initially chose to forgo a linearly tensioning mechanism (instead, favoring a rotary spool), a folded lead-screw mechanism may further reduce the total size of the torque reduction stage, opening up more space for batteries. Next, we note that each of Tasbi's vibrotactor modules contain significant dead-space. If an efficient distributed power system can be devised, each of Tasbi's linear resonant actuators could be powered by their own independent battery located in these housings. Research into novel battery technologies may also provide a path forward as well, as both printable [179] and multifunctional structural [180] batteries could allow for storing energy in the exterior housing elements.

## 9.2 Hands-Free AR Interactions

For technical reasons, the work presented in this thesis only focused on using Tasbi for hand interactions in VR with controller based tracking. Although the interaction approaches presented in Chapter 7 should more or less translate to AR and MR contexts, there are some points to consider. First, we must note that there are two current approaches to AR: either holographic projection onto transparent lenses or video pass through. The ability to use visual and proprioceptive illusions to enhance Tasbi haptic feedback will depend on which of these two approaches the industry decides to adopt. In the former case, users will see the real world through optically transparent lenses while digital objects are projected onto the surface of the lenses. Thus, we have no means of distorting the user's visual perception of their hand location or world objects, and so control-to-display techniques will become intractable. In the latter case, externally facing cameras transmit video to stereoscopic screens similar to how current VR headsets work. In this case, we have the ability to intercept the video and digitally alter the user's hand position, so control-to-display techniques may still be viable. However, some researchers have expressed concerns that AR *feels* inherently softer than VR [181], so either way the transition may not be completely straightforward.

We must also note that the interactions in this thesis relied on hand tracking through handheld controllers. Eventually, sophisticated hand and finger tracking will be provided by externally facing cameras and machine learning algorithms [182, 183] (indeed, some HMDs already support experimental hand tracking [8]). Again, the transition should be overcomable but there are points to consider. First, interactions will need to manage the complexity that multi-degree-of-freedom virtual hands bring. No longer will we be able to build feedback algorithms around pre-defined hand poses and gestures, but instead will need to generalize our approach to handle the unpredictabilities of users. Consider, for




Figure 9.1 : The controller-free, finger tracked hands of AR will present new opportunities for Tasbi research. (a) Fine motor tasks such as playing a piano will require explorations into how to appropriate squeeze and vibrotactile feedback to multiple and simultaneous finger interactions. (b) The absence of controllers will remove inherent tactile feedback user receive when holding them, such as for the virtual bow interaction. However, this issue is less concerning once we realize that devices like Tasbi can turn essentially any inanimate object (e.g. this frivolous banana example) into a haptic feedback controller.

example, the piano in Fig. 9.1. Each piano key could be considered an extension of the button interaction discussed in Chapters 7 and 8. However, we now have a scenario where each finger can independently and simultaneously be in a state of contact. How should we distribute vibration and squeeze cues for all of the fingers? A naive approach might be to assign one vibrotactor to each finger for contact events, and render only a percentage (i.e. 20%) of squeeze when each finger depresses a key. In practice we will likely find this to be a much more nuanced problem, and so other mapping strategies will need to be considered. The implementation of the physics simulation will need to drastically change as well, as accurate simulation of virtual hands is an *extremely* challenging problem that many researchers are actively pursuing [184, 185, 186]. Nonetheless, these problems offer very compelling areas for future work, as one might consider fine motor skill tasks such as playing a piano or typing on a mid-air keyboard [187] to be the *holy grail* for referred haptic feedback.

We should also consider that in many cases, the act of simply holding a controller provides tactile information that will vanish when interactions become completely mid-air based. Consider the bow-and-arrow interaction presented in Chapter 7. While Tasbi provides haptic feedback for drawing the bow string, the controller provides the tactile sensation of actually holding the bow handle. However, consider this: while the controller may ultimately be removed, there is nothing preventing the user from picking up an inanimate and readily available object to substitute for the bow handle (e.g. a banana, Fig. 9.1). In this way, Tasbi combined with sophisticated machine-vision can be used to turn almost any object into a haptic feedback controller at a moment's notice.

### 9.3 Integrating Tasbi with Other Technologies

The possibility of integrating Tasbi with other technologies is of course another important area of future work. In one perspective, devices like Tasbi can be used to enhance already existing technologies. Take for example the capacitive touch-screen displays that are now prevalent in almost every facet of human-computer interaction – even vehicles. Gone are the days where a driver could simply reach over and adjust the stereo volume without taking their eyes off the road, because traditional tactile knobs and switches have been replaced with flat, slick touch interfaces. Tasbi, wirelessly paired to the vehicle, could potentially restore some of this tactile feedback without requiring manufactures to adopt their own haptic technology. This of course is true for any touch screen display, or any *surface* for that matter.

There are also opportunities for leveraging other technologies to enhance Tasbi. Brain-machine interfaces are one such technology. Recent advances in electromyography (EMG) have the potential to radically alter human inputs and interaction with virtual interfaces, from simple estimations of effort [188] to full hand pose reconstruction [189]. Mixing



Figure 9.2 : Integrating Tasbi with other technologies presents interesting opportunities. (a) Devices like Tasbi can provide haptic feedback for ordinary capacitive displays. (b) Emerging technologies, such as brain-machine interfaces and electromyography (EMG) might enhance the believability of Tasbi-mediated interactions in XR.

Tasbi into the picture presents an interesting possibility where EMG activity prescribes a user's intent in a virtual context, and Tasbi provides confirmation that their intent was received. Consider the example in Fig. [9.2](#), where the user cannot push a virtual block until sufficient EMG activity is measured. Once this has been accomplished, Tasbi renders back the physics of the interaction. Closing the sensorimotor loop in this fashion might further blur the line between real and virtual worlds.

## 9.4 Sharing Technology with the Research Community

Many researchers have expressed interest in obtaining a device like Tasbi, as not all scientists and labs have the resources or skill-set required to implement such devices. Thus, sharing the fruits of this thesis, whether it be hardware or software, is a worthwhile effort. With the release of Tasbi's vibrotactile rendering engine, Syntacts, we have already made progress toward this goal. We have freely distributed over 40 Syntacts amplifier kits




Figure 9.3 : Technology developed for Tasbi is being freely shared with research communities. (a-b) Syntacts vibrotactile amplifier kits that were distributed during an IROS 2020 workshop. (c) A vibrotactile-only clone of Tasbi we have made available online. (d) A fellow lab member's wearable haptic display made possible by Syntacts and Tasbi research.

(9.3-a,b) to the haptics community during workshops and conferences [190], and made available an open-source, vibrotactile-only rendition of Tasbi [191] (Fig 9.3-b). Already, we have seen researchers begin to utilize our technology for their own wearable haptic projects, both internally (Fig. 9.3-d) and externally. With Syntacts being selected as the tool-of-choice for the 2021 World Haptics Conference Student Innovation Challenge, our technology will soon be in the hands of students and aspiring scientists as well. As we move forward, we are excited to see what the research communities will create and hope to see a growth in interest toward referred haptic feedback for XR interactions.

## Bibliography

- [1] S. H.-W. Chuah, “Why and who will adopt extended reality technology? literature review, synthesis, and future research agenda,” *Literature Review, Synthesis, and Future Research Agenda (December 13, 2018)*, 2018.
- [2] L. Morvan, F. Hintermann, and A. Ovanessoff, “Preparing for the risky world of extended reality,” *MIT Sloan Management Review*, vol. 61, no. 2, pp. 1–4, 2020.
- [3] K. J. Kuchenbecker, J. Fiene, and G. Niemeyer, “Improving contact realism through event-based haptic feedback,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 2, pp. 219–230, 2006.
- [4] V. Hayward, O. Ashley, C. Hernandez, D. Grant, and G. Robles-De-La-Torre, “Haptic interfaces and devices,” *Sensor Review*, vol. 24, pp. 16–29, 03 2004.
- [5] R. Sigrist *et al.*, “Augmented visual, auditory, haptic, and multimodal feedback in motor learning: A review,” *Psychonomic Bulletin & Review*, vol. 20, no. 1, pp. 21–53, 2013.
- [6] Y. Visell, “Tactile sensory substitution: Models for enaction in HCI,” *Interacting with Computers*, vol. 21, pp. 38–53, 08 2008.
- [7] C. Bermejo and P. Hui, “A survey on haptic technologies for mobile augmented reality,” *CoRR*, vol. abs/1709.00698, 2017.

- [8] Oculus, “VR Headsets and Equipment.” <https://www.oculus.com/>. Accessed: 2021-01-02.
- [9] H. Corporation, “HTC Vive.” <https://www.vive.com/us/>. Accessed: 2021-01-02.
- [10] Valve, “Valve Index.” <https://www.valvesoftware.com/en/index>. Accessed: 2021-01-02.
- [11] 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*, p. 1–13. New York, NY, USA: Association for Computing Machinery, 2018.
- [12] 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*, pp. 1–13, 2019.
- [13] 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 *ACM Conference on Human Factors in Computing Systems (CHI)*, (New York, NY, USA), pp. 86:1–86:12, ACM, 2018.
- [14] G. Burdea, P. Richard, and P. Coiffet, “Multimodal virtual reality: Input-output devices, system integration, and human factors,” *International Journal of Human-Computer Interaction*, vol. 8, no. 1, pp. 5–24, 1996.
- [15] M. Bouzit, G. Burdea, G. Popescu, and R. Boian, “The Rutgers master ii-new design force-feedback glove,” *IEEE/ASME Transactions on Mechatronics*, vol. 7, no. 2, pp. 256–263, 2002.



- [16] J. Blake and H. B. Gurocak, "Haptic glove with mr brakes for virtual reality," *IEEE/ASME Transactions on Mechatronics*, vol. 14, no. 5, pp. 606–615, 2009.
- [17] K. Koyanagi, Y. Fujii, and J. Furusho, "Development of vr-stef system with force display glove system," in *Proceedings of the 2005 international conference on Augmented tele-existence*, pp. 91–97, 2005.
- [18] D. Trinitatova, D. Tsetserukou, and A. Fedoseev, "Touchvr: A wearable haptic interface for vr aimed at delivering multi-modal stimuli at the user's palm," in *SIGGRAPH Asia 2019 XR, SA '19*, (New York, NY, USA), p. 42–43, Association for Computing Machinery, 2019.
- [19] J. Perret and E. Vander Poorten, "Touching virtual reality: a review of haptic gloves," in *ACTUATOR Conference*, pp. 1–5, June 2018.
- [20] P. Polygerinos, Z. Wang, K. C. Galloway, R. J. Wood, and C. J. Walsh, "Soft robotic glove for combined assistance and at-home rehabilitation," *Robotics and Autonomous Systems*, vol. 73, pp. 135–143, 2015. Wearable Robotics.
- [21] 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 Transactions Haptics*, vol. 10, pp. 580–600, Oct 2017.
- [22] M. Bianchi, E. Battaglia, M. Poggiani, S. Ciotti, and A. Bicchi, "A wearable fabric-based display for haptic multi-cue delivery," in *IEEE Haptics Symposium*, pp. 277–283, April 2016.
- [23] C. Pacchierotti, G. Salvietti, I. Hussain, L. Meli, and D. Prattichizzo, "The hring: A wearable haptic device to avoid occlusions in hand tracking," in *2016 IEEE Haptics Symposium (HAPTICS)*, pp. 134–139, IEEE, 2016.

- [24] D. Leonardis, M. Solazzi, I. Bortone, and A. Frisoli, “A wearable fingertip haptic device with 3 dof asymmetric 3-rsr kinematics,” in *2015 IEEE World Haptics Conference (WHC)*, pp. 388–393, IEEE, 2015.
- [25] M. Maisto, C. Pacchierotti, F. Chinello, G. Salvietti, A. De Luca, and D. Prattichizzo, “Evaluation of wearable haptic systems for the fingers in augmented reality applications,” *IEEE Transactions on Haptics*, vol. 10, no. 4, pp. 511–522, 2017.
- [26] S. B. Schorr and A. M. Okamura, “Three-dimensional skin deformation as force substitution: Wearable device design and performance during haptic exploration of virtual environments,” *IEEE Transactions on Haptics*, vol. 10, no. 3, pp. 418–430, 2017.
- [27] I. Choi, E. W. Hawkes, D. L. Christensen, C. J. Ploch, and S. Follmer, “Wolverine: A wearable haptic interface for grasping in virtual reality,” in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 986–993, 2016.
- [28] 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, UIST '17*, (New York, NY, USA), p. 119–130, Association for Computing Machinery, 2017.
- [29] F. Chinello, C. Pacchierotti, J. Bimbo, N. G. Tsagarakis, and D. Prattichizzo, “Design and evaluation of a wearable skin stretch device for haptic guidance,” *IEEE Robotics and Automation Letters*, vol. 3, no. 1, pp. 524–531, 2018.
- [30] D. Dobbstein, E. Stemasov, D. Besserer, I. Stenske, and E. Rukzio, “Movelet: A

self-actuated movable bracelet for positional haptic feedback on the user’s forearm,” in *Proceedings of the 2018 ACM International Symposium on Wearable Computers*, ISWC ’18, (New York, NY, USA), p. 33–39, Association for Computing Machinery, 2018.

- [31] E. Battaglia, J. P. Clark, M. Bianchi, M. G. Catalano, A. Bicchi, and M. K. O’Malley, “Skin stretch haptic feedback to convey closure information in anthropomorphic, under-actuated upper limb soft prostheses,” *IEEE Transactions on Haptics*, vol. 12, no. 4, pp. 508–520, 2019.
- [32] N. Dunkelberger, J. L. Sullivan, J. Bradley, I. Manickam, G. Dasarathy, R. G. Baraniuk, and M. K. O’Malley, “A multi-sensory approach to present phonemes as language through a wearable haptic device,” *IEEE Transactions on Haptics*, pp. 1–1, 2020.
- [33] K. J. Kim and D. Shin, “An acceptance model for smart watches: Implications for the adoption of future wearable technology,” *Internet Research: Electronic Networking Applications and Policy*, vol. 25, 08 2015.
- [34] J. Choi and S. Kim, “Is the smartwatch an it product or a fashion product? a study on factors affecting the intention to use smartwatches,” *Computers in Human Behavior*, vol. 63, pp. 777 – 786, 2016.
- [35] Apple Inc., “Apple Watch.” <https://www.apple.com/watch/>. Accessed: 2021-01-02.
- [36] House of Haptics, “HEY Bracelet.” <https://www.feelhey.com>. Accessed: 2021-01-02.

- [37] Bond Touch, “Bond Touch Bracelet.” <https://www.bond-touch.com/>. Accessed: 2021-01-02.
- [38] M. A. Baumann, K. E. MacLean, T. W. Hazelton, and A. McKay, “Emulating human attention-getting practices with wearable haptics,” in *IEEE Haptics Symposium*, pp. 149–156, 2010.
- [39] J. P. Clark, S. Y. Kim, and M. K. O’Malley, “The rice haptic rocker: Altering the perception of skin stretch through mapping and geometric design,” in *2018 IEEE Haptics Symposium (HAPTICS)*, pp. 192–197, 2018.
- [40] S. Casini *et al.*, “Design and realization of the cuff-clenching upper-limb force feedback wearable device for distributed mechano-tactile stimulation of normal and tangential skin forces,” in *IEEE International Conference Intell. Robots and Syst. (IROS)*, pp. 1186–1193, 2015.
- [41] K. Bark, J. Wheeler, P. Shull, J. Savall, and M. Cutkosky, “Rotational skin stretch feedback: A wearable haptic display for motion,” *IEEE Transactions on Haptics*, vol. 3, no. 3, pp. 166–176, 2010.
- [42] J. Wheeler, K. Bark, J. Savall, and M. Cutkosky, “Investigation of rotational skin stretch for proprioceptive feedback with application to myoelectric systems,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 18, no. 1, pp. 58–66, 2010.
- [43] S. Song, G. Noh, J. Yoo, I. Oakley, J. Cho, and A. Bianchi, “Hot & tight: Exploring thermo and squeeze cues recognition on wrist wearables,” in *International Symposium on Wearable Computers (ISWC)*, (New York, NY, USA), pp. 39–42, ACM, 2015.

- [44] P. Lopes, S. You, L.-P. Cheng, S. Marwecki, and P. Baudisch, “Providing haptics to walls &#38; heavy objects in virtual reality by means of electrical muscle stimulation,” in *ACM Conference on Human Factors in Computing Systems (CHI)*, (New York, NY, USA), pp. 1471–1482, ACM, 2017.
- [45] J. L. Sullivan, N. Dunkelberger, J. Bradley, J. Young, A. Israr, F. Lau, K. Klumb, F. Abnousi, and M. K. O’Malley, “Multi-sensory stimuli improve distinguishability of cutaneous haptic cues,” *IEEE Transactions on Haptics*, vol. 13, no. 2, pp. 286–297, 2020.
- [46] E. Pezent, M. K. O’Malley, A. Israr, M. Samad, S. Robinson, P. Agarwal, H. Benko, and N. Colonnese, “Explorations of wrist haptic feedback for ar/vr interactions with tasbi,” *CHI EA ’20*, (New York, NY, USA), p. 1–4, Association for Computing Machinery, 2020.
- [47] D. Wang, K. Ohnishi, and W. Xu, “Multimodal haptic display for virtual reality: A survey,” *IEEE Transactions on Industrial Electronics*, vol. 67, no. 1, pp. 610–623, 2019.
- [48] Z. A. Zook, J. J. Fleck, T. W. Tjandra, and M. K. O’Malley, “Effect of interference on multi-sensory haptic perception of stretch and squeeze,” in *2019 IEEE World Haptics Conference (WHC)*, pp. 371–376, 2019.
- [49] Y. Zheng and J. B. Morrell, “Haptic actuator design parameters that influence affect and attention,” in *IEEE Haptics Symposium*, pp. 463–470, March 2012.
- [50] R. Wang, F. Quek, D. Tatar, K. S. Teh, and A. Cheok, “Keep in touch: Channel, expectation and experience,” in *ACM Conference on Human Factors in Computing Systems (CHI)*, (New York, NY, USA), pp. 139–148, ACM, 2012.

- [51] D. Tsetserukou, “Haptihug: A novel haptic display for communication of hug over a distance,” in *Haptics: Generating and Perceiving Tangible Sensations* (A. M. L. Kappers, J. B. F. van Erp, W. M. Bergmann Tiest, and F. C. T. van der Helm, eds.), (Berlin, Heidelberg), pp. 340–347, Springer Berlin Heidelberg, 2010.
- [52] L. Meli, I. Hussain, M. Aurilio, M. Malvezzi, M. K. O’Malley, and D. Prattichizzo, “The hbracelet: A wearable haptic device for the distributed mechanotactile stimulation of the upper limb,” *IEEE Robotics and Automation Letters*, vol. 3, pp. 2198–2205, July 2018.
- [53] S. Biswas and Y. Visell, “Emerging material technologies for haptics,” *Advanced Materials Technologies*, vol. 4, no. 4, p. 1900042, 2019.
- [54] P. B. Shull and D. D. Damian, “Haptic wearables as sensory replacement, sensory augmentation and trainer—a review,” *Journal of neuroengineering and rehabilitation*, vol. 12, no. 1, pp. 1–13, 2015.
- [55] E. Treadway, B. Gillespie, D. Bolger, A. Blank, M. O’Malley, and A. Davis, “The role of auxiliary and referred haptic feedback in myoelectric control,” in *2015 IEEE World Haptics Conference (WHC)*, pp. 13–18, 2015.
- [56] J. D. Brown, A. Paek, M. Syed, M. K. O’Malley, P. A. Shewokis, J. L. Contreras-Vidal, A. J. Davis, and R. B. Gillespie, “An exploration of grip force regulation with a low-impedance myoelectric prosthesis featuring referred haptic feedback,” *Journal of neuroengineering and rehabilitation*, vol. 12, no. 1, pp. 1–17, 2015.
- [57] P. Bach-y Rita, “Tactile sensory substitution studies,” *Annals of the New York Academy of Sciences*, vol. 1013, no. 1, pp. 83–91, 2004.

- [58] E. Pezent, A. Israr, M. Samad, S. Robinson, P. Agarwal, H. Benko, and N. Colonese, “Tasbi: Multisensory squeeze and vibrotactile wrist haptics for augmented and virtual reality,” in *IEEE World Haptics Conference (WHC)*, (Tokyo, Japan), pp. 1–6, IEEE, July 2019.
- [59] E. Pezent, B. Cambio, and M. K. O’Malley, “Syntacts: Open-source software and hardware for audio-controlled haptics,” *IEEE Transactions on Haptics*, pp. 1–1, 2020.
- [60] T. Moriyama, T. Nakamura, and H. Kajimoto, “Development of a wearable haptic device that presents the haptic sensation corresponding to three fingers on the forearm,” in *Proceedings of the Symposium on Spatial User Interaction, SUI ’18*, (New York, NY, USA), p. 158–162, Association for Computing Machinery, 2018.
- [61] M. Sarac, A. M. Okamura, and M. D. Luca, “Haptic sketches on the arm for manipulation in virtual reality,” in *IEEE World Haptics Conference (WHC)*, pp. Work-in-Progress, 2019.
- [62] M. Sarac, A. M. Okamura, and M. D. Luca, “Effects of haptic feedback on the wrist during virtual manipulation,” in *IEEE Haptics Symposium*, pp. Work-in-Progress, 2020.
- [63] X. de Tinguy, C. Pacchierotti, M. Marchal, and A. Lécuyer, “Enhancing the stiffness perception of tangible objects in mixed reality using wearable haptics,” in *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 81–90, 2018.
- [64] R. W. Cholewiak and A. A. Collins, “Vibrotactile localization on the arm: Effects of place, space, and age,” *Perception & psychophysics*, vol. 65, no. 7, pp. 1058–1077, 2003.

- [65] I. Oakley, Y. Kim, J. Lee, and J. Ryu, “Determining the feasibility of forearm mounted vibrotactile displays,” in *Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pp. 27–34, March 2006.
- [66] H.-Y. Chen, J. Santos, M. Graves, K. Kim, and H. Z. Tan, “Tactor localization at the wrist,” in *EuroHaptics*, pp. 209–218, 2008.
- [67] M. Matscheko, A. Ferscha, A. Riener, and M. Lehner, “Tactor placement in wrist worn wearables,” in *International Symposium on Wearable Computers (ISWC)*, pp. 1–8, Oct 2010.
- [68] M. G. Carcedo, S. H. Chua, S. Perrault, P. Wozniak, R. Joshi, M. Obaid, M. Fjeld, and S. Zhao, “Hapticolor: Interpolating color information as haptic feedback to assist the colorblind,” in *ACM Conference on Human Factors in Computing Systems (CHI)*, pp. 3572–3583, ACM, 2016.
- [69] A. Gupta, T. Pietrzak, N. Roussel, and R. Balakrishnan, “Direct manipulation in tactile displays,” in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI ’16, (New York, NY, USA), p. 3683–3693, Association for Computing Machinery, 2016.
- [70] F. Pece, J. J. Zarate, V. Vechev, N. Besse, O. Gudozhnik, H. Shea, and O. Hilliges, “Magtics: Flexible and thin form factor magnetic actuators for dynamic and wearable haptic feedback,” in *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, UIST ’17, (New York, NY, USA), p. 143–154, Association for Computing Machinery, 2017.
- [71] J. Hong, L. Stearns, J. Froehlich, D. Ross, and L. Findlater, “Evaluating angular accuracy of wrist-based haptic directional guidance for hand movement,” in *Graph-*



- ics Interface Conference*, pp. 195–200, Canadian Human-Computer Comm. Society, 2016.
- [72] S. C. Lee and T. Starner, *BuzzWear: Alert Perception in Wearable Tactile Displays on the Wrist*, p. 433–442. New York, NY, USA: Association for Computing Machinery, 2010.
- [73] S. J. Bolanowski, G. A. Gescheider, R. T. Verrillo, and C. M. Checkosky, “Four channels mediate the mechanical aspects of touch.,” *J. Acoustical Society of America*, vol. 84 5, pp. 1680–94, 1988.
- [74] S. J. Lederman and R. L. Klatzky, “Haptic perception: A tutorial,” *Attention, Perception, & Psychophysics*, vol. 71, pp. 1439–1459, Oct 2009.
- [75] H. Pohl, P. Brandes, H. Ngo Quang, and M. Rohs, “Squeezeback: Pneumatic compression for notifications,” in *ACM Conference on Human Factors in Computing Systems (CHI)*, (New York, NY, USA), pp. 5318–5330, ACM, 2017.
- [76] D. Tsetserukou, A. Neviarouskaya, H. Prendinger, N. Kawakami, and S. Tachi, “Affective haptics in emotional communication,” in *International Conference on Affective Computing and Intelligent Interaction and Workshops*, pp. 1–6, Sep. 2009.
- [77] J. D. Brown, J. N. Fernandez, S. P. Cohen, and K. J. Kuchenbecker, “A wrist-squeezing force-feedback system for robotic surgery training,” in *2017 IEEE World Haptics Conference (WHC)*, pp. 107–112, 2017.
- [78] A. A. Stanley and K. J. Kuchenbecker, “Evaluation of tactile feedback methods for wrist rotation guidance,” *IEEE Transactions Haptics*, vol. 5, no. 3, pp. 240–251, 2012.

- [79] R. Wang, F. Quek, D. Tatar, K. S. Teh, and A. Cheok, “Keep in touch: channel, expectation and experience,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 139–148, 2012.
- [80] M. Bianchi, G. Valenza, A. Serio, A. Lanata, A. Greco, M. Nardelli, E. P. Scilingo, and A. Bicchi, “Design and preliminary affective characterization of a novel fabric-based tactile display,” in *2014 IEEE Haptics Symposium (HAPTICS)*, pp. 591–596, IEEE, 2014.
- [81] J. J. Fleck, Z. A. Zook, T. W. Tjandra, and M. K. O’Malley, “A cutaneous haptic cue characterization testbed,” in *2019 IEEE World Haptics Conference (WHC)*, pp. 319–324, 2019.
- [82] F. Chinello, M. Aurilio, C. Pacchierotti, and D. Prattichizzo, “The HapBand: A cutaneous device for remote tactile interaction,” in *Haptics: Neuroscience, Devices, Modeling, and Applications* (M. Auvray and C. Duriez, eds.), (Berlin, Heidelberg), pp. 284–291, Springer, 2014.
- [83] A. Gupta, A. A. R. Irudayaraj, and R. Balakrishnan, “HapticClench: Investigating squeeze sensations using memory alloys,” in *ACM Symposium on User Interface Soft. and Tech. (UIST)*, (New York, NY, USA), pp. 109–117, ACM, 2017.
- [84] 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)*, pp. 55–60, 2019.
- [85] M. Zhu, A. H. Memar, A. Gupta, M. Samad, P. Agarwal, Y. Visell, S. J. Keller, and N. Colonnese, “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*, CHI '20, (New York, NY, USA), p. 1–12, Association for Computing Machinery, 2020.
- [86] M. Raitor, J. M. Walker, A. M. Okamura, and H. Culbertson, “Wrap: Wearable, restricted-aperture pneumatics for haptic guidance,” in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 427–432, IEEE, 2017.
- [87] C. Payne, E. Gallardo Hevia, N. Phipps, A. Atalay, O. Atalay, B. Seo, D. Mooney, and C. Walsh, “Force control of textile-based soft wearable robots for mechanotherapy,” in *IEEE International Conference on Robotics and Automation (ICRA)*, (Brisbane, Australia, May 21-25), 2018.
- [88] S. J. Kim, H. Chang, J. Park, and J. Kim, “Design of a portable pneumatic power source with high output pressure for wearable robotic applications,” *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 4351–4358, 2018.
- [89] M. Wehner, M. T. Tolley, Y. Mengüç, Y.-L. Park, A. Mozeika, Y. Ding, C. Onal, R. F. Shepherd, G. M. Whitesides, and R. J. Wood, “Pneumatic energy sources for autonomous and wearable soft robotics,” *Soft robotics*, vol. 1, no. 4, pp. 263–274, 2014.
- [90] T. Diller, “Frequency response of human skin in vivo to mechanical stimulation,” Master’s thesis, Massachusetts Institute of Technology, 2001.
- [91] M. Aggravi, F. Pausé, P. R. Giordano, and C. Pacchierotti, “Design and evaluation of a wearable haptic device for skin stretch, pressure, and vibrotactile stimuli,” *IEEE Robotics and Automation Letters*, vol. 3, pp. 2166–2173, July 2018.
- [92] C. Moussette, *Simple haptics: Sketching perspectives for the design of haptic interactions*. PhD thesis, Umeå Universitet, 2012.

- [93] Maxon Group, “Precision Drive Systems by Maxon.” <https://www.maxongroup.com/>. Accessed: 2021-01-02.
- [94] Harmonic Drive, “Harmonic Drive High Precision Gear.” <https://www.harmonicdrive.net/>. Accessed: 2021-01-02.
- [95] Elesta GmbH, “Elesta GmbH.” <https://www.elesta-gmbh.com/en/>. Accessed: 2021-01-02.
- [96] M. Park, B.-G. Bok, J.-H. Ahn, and M.-S. Kim, “Recent advances in tactile sensing technology,” *Micromachines*, vol. 9, no. 7, p. 321, 2018.
- [97] J. P. Mogk and P. J. Keir, “Wrist and carpal tunnel size and shape measurements: effects of posture,” *Clinical biomechanics*, vol. 23, no. 9, pp. 1112–1120, 2008.
- [98] NASA, “NASA-STD-3000 Main-Systems Integration Standards, Volume 1, Section 3: Anthropometry and Biomechanics.” <https://msis.jsc.nasa.gov/sections/section02.htm>, 1995. Accessed: 2018-06-01.
- [99] MOTU, “Mark of the Unicorn Audio Interfaces.” <https://motu.com/en-us/>. Accessed: 2021-01-02.
- [100] Steinberg, “Audio Stream Input/Output.” <https://www.steinberg.net/en/company/developers.html>. Accessed: 2021-01-02.
- [101] Quanser, “Quanser - Innovate - Educate.” <https://www.quanser.com/>. Accessed: 2021-01-02.
- [102] J. Sparks, N. Vavalle, K. Kasting, B. Long, M. Tanaka, P. Sanger, K. Schnell, and T. Conner-Kerr, “Use of silicone materials to simulate tissue biomechanics as related to deep tissue injury,” *Advances in skin and wound care*, vol. 28, pp. 59–68, 02 2015.

- [103] M. Schroeder, “Synthesis of low-peak-factor signals and binary sequences with low autocorrelation (corresp.),” *IEEE Transactions on Information Theory*, vol. 16, pp. 85–89, Sept. 2006.
- [104] J. B. Hoagg, S. L. Lacy, V. Babuska, and D. S. Bernstein, “Sequential multisine excitation signals for system identification of large space structures,” in *2006 American Control Conference*, pp. 418–423, June 2006.
- [105] P. E. Fortin, J. R. Blum, A. Weill-Duflos, and J. R. Cooperstock, “Contact force estimation from raw photoplethysmogram signal,” in *2020 IEEE SENSORS*, pp. 1–4, 2020.
- [106] B. Cambio, *Ultra-Low Latency Control of Large Vibrotactile Arrays for Haptic Interactions*. PhD thesis, Rice University, 2020.
- [107] S. Choi and K. J. Kuchenbecker, “Vibrotactile display: Perception, technology, and applications,” *Proceedings of the IEEE*, vol. 101, pp. 2093–2104, Sep. 2013.
- [108] H. Culbertson, C. M. Nunez, A. Israr, F. Lau, F. Abnoui, and A. M. Okamura, “A social haptic device to create continuous lateral motion using sequential normal indentation,” in *2018 IEEE Haptics Symposium (HAPTICS)*, pp. 32–39, 2018.
- [109] Y. Konishi, N. Hanamitsu, B. Outram, K. Minamizawa, T. Mizuguchi, and A. Sato, “Synesthesia suit: The full body immersive experience,” in *ACM SIGGRAPH 2016 VR Village, SIGGRAPH ’16*, (New York, NY, USA), Association for Computing Machinery, 2016.
- [110] A. Israr, S. Zhao, K. Schwalje, R. Klatzky, and J. Lehman, “Feel effects: Enriching storytelling with haptic feedback,” *Transactions on Applied Perception (TAP)*, vol. 11, pp. 11:1–11:17, Sept. 2014.

- [111] O. S. Schneider, A. Israr, and K. E. MacLean, “Tactile animation by direct manipulation of grid displays,” in *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology*, UIST ’15, (New York, NY, USA), p. 21–30, Association for Computing Machinery, 2015.
- [112] H. Seifi, “Crowdsourcing haptic data collection,” in *Personalizing Haptics: From Individuals’ Sense-Making Schemas to End-User Haptic Tools*, pp. 111–133, Cham: Springer International Publishing, 2019.
- [113] R. Rosenkranz and M. E. Altinsoy, “Tactile design: Translating user expectations into vibration for plausible virtual environments\*,” in *IEEE World Haptics Conference (WHC)*, pp. 307–312, July 2019.
- [114] Y. Ochiai, T. Hoshi, J. Rekimoto, and M. Takasaki, “Diminished haptics: Towards digital transformation of real world textures,” in *Haptics: Neuroscience, Devices, Modeling, and Applications*, (Berlin, Heidelberg), pp. 409–417, Springer Berlin Heidelberg, 2014.
- [115] O. S. Schneider and K. E. MacLean, “Studying design process and example use with macaron, a web-based vibrotactile effect editor,” in *2016 IEEE Haptics Symposium*, pp. 52–58, April 2016.
- [116] O. Schneider, S. Zhao, and A. Israr, *FeelCraft: User-Crafted Tactile Content*, pp. 253–259. Tokyo: Springer Japan, 2015.
- [117] M. Di Luca and A. Mahnan, “Perceptual limits of visual-haptic simultaneity in virtual reality interactions,” in *2019 IEEE World Haptics Conference*, pp. 67–72, July 2019.

- [118] Texas Instruments, “Embedded Processing.” <https://www.ti.com/>. Accessed: 2021-01-02.
- [119] EAI, “Engineering Acoustics, Inc..” <https://www.eaiinfo.com/>. Accessed: 2021-01-02.
- [120] Aktronica, “Aktronica - Do It, Sense It.” <https://www.actronika.com/>. Accessed: 2021-01-02.
- [121] C. Frisson, T. Pietrzak, S. Zhao, Z. Schwemier, and A. Israr, “WebAudioHaptics: Tutorial on haptics with web audio,” in *Web Audio Conference (WAC)*, 2016.
- [122] A. Israr, S. Zhao, Z. Schwemler, and A. Fritz, “Stereohaptics toolkit for dynamic tactile experiences,” in *HCI International 2019 – Late Breaking Papers* (C. Stephanidis, ed.), (Cham), pp. 217–232, Springer International Publishing, 2019.
- [123] J. Chen, P. Castillo, R. Turcott, A. Israr, and F. Lau, “Feeling speech on the arm,” in *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI EA ’18, (New York, NY, USA), Association for Computing Machinery, 2018.
- [124] C. M. Reed, H. Z. Tan, Z. D. Perez, E. C. Wilson, F. M. Severgnini, J. Jung, J. S. Martinez, Y. Jiao, A. Israr, F. Lau, K. Klumb, R. Turcott, and F. Abnoui, “A phonemic-based tactile display for speech communication,” *IEEE Transactions on Haptics*, vol. 12, pp. 2–17, Jan 2019.
- [125] C. Swindells, S. Pietarinen, and A. Viitanen, “Medium fidelity rapid prototyping of vibrotactile haptic, audio and video effects,” in *IEEE Haptics Symposium*, pp. 515–521, Feb 2014.

- [126] M. J. Enriquez and K. E. MacLean, “The hapticon editor: a tool in support of haptic communication research,” in *Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pp. 356–362, March 2003.
- [127] L. A. Jones and H. Z. Tan, “Application of psychophysical techniques to haptic research,” *IEEE Transactions on Haptics*, vol. 6, no. 3, pp. 268–284, 2013.
- [128] A. M. Okamura, M. R. Cutkosky, and J. T. Dennerlein, “Reality-based models for vibration feedback in virtual environments,” *IEEE/ASME Transactions on Mechatronics*, vol. 6, pp. 245–252, Sep. 2001.
- [129] H. Culbertson, J. Unwin, and K. J. Kuchenbecker, “Modeling and rendering realistic textures from unconstrained tool-surface interactions,” *IEEE Transactions on Haptics*, vol. 7, no. 3, pp. 381–393, 2014.
- [130] H. Culbertson, J. J. López Delgado, and K. J. Kuchenbecker, “One hundred data-driven haptic texture models and open-source methods for rendering on 3d objects,” in *2014 IEEE Haptics Symposium (HAPTICS)*, pp. 319–325, 2014.
- [131] Y. Shao, V. Hayward, and Y. Visell, “Spatial patterns of cutaneous vibration during whole-hand haptic interactions,” *Proceedings of the National Academy of Sciences*, vol. 113, no. 15, pp. 4188–4193, 2016.
- [132] S. J. Lederman and L. A. Jones, “Tactile and haptic illusions,” *IEEE Transactions on Haptics*, vol. 4, no. 4, pp. 273–294, 2011.
- [133] M. Samad, A. J. Chung, and L. Shams, “Perception of body ownership is driven by bayesian sensory inference,” *PLOS ONE*, vol. 10, pp. 1–23, 02 2015.



- [134] M. D'Alonzo and C. Cipriani, "Vibrotactile sensory substitution elicits feeling of ownership of an alien hand," *PLOS ONE*, vol. 7, pp. 1–9, 11 2012.
- [135] C. Cipriani, M. D'Alonzo, and M. C. Carrozza, "A miniature vibrotactile sensory substitution device for multifingered hand prosthetics," *IEEE Transactions on Biomedical Engineering*, vol. 59, no. 2, pp. 400–408, 2012.
- [136] B. D. Adelstein, D. R. Begault, M. R. Anderson, and E. M. Wenzel, "Sensitivity to haptic-audio asynchrony," in *Proceedings of the 5th International Conference on Multimodal Interfaces, ICMI '03*, (New York, NY, USA), p. 73–76, Association for Computing Machinery, 2003.
- [137] M. Keetels and J. Vroomen, "Perception of synchrony between the senses," in *The Neural Bases of Multisensory Processes*, CRC Press, 2012.
- [138] M. Keetels and J. Vroomen, "Temporal recalibration to tactile–visual asynchronous stimuli," *Neuroscience Letters*, vol. 430, no. 2, pp. 130 – 134, 2008.
- [139] C. B. Zilles and J. K. Salisbury, "A constraint-based god-object method for haptic display," in *IEEE International Conference Intell. Robots and Syst. (IROS)*, vol. 3, pp. 146–151 vol.3, Aug 1995.
- [140] M. Ortega, S. Redon, and S. Coquillart, "A six degree-of-freedom god-object method for haptic display of rigid bodies with surface properties," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 3, pp. 458–469, 2007.
- [141] M. Sagardia and T. Hulin, "A fast and robust six-dof god object heuristic for haptic rendering of complex models with friction," in *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology*, pp. 163–172, 2016.

- [142] A. Talvas, M. Marchal, and A. Lécuyer, “The god-finger method for improving 3d interaction with virtual objects through simulation of contact area,” in *2013 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 111–114, IEEE, 2013.
- [143] A. Lécuyer, “Simulating haptic feedback using vision: A survey of research and applications of pseudo-haptic feedback,” *Presence: Teleoper. Virtual Environ.*, vol. 18, pp. 39–53, Jan. 2009.
- [144] P. Preechayasomboon, A. Israr, and M. Samad, “Chasm: A screw based expressive compact haptic actuator,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI ’20, (New York, NY, USA), p. 1–13, Association for Computing Machinery, 2020.
- [145] M. Rietzler, F. Geiselhart, J. Gugenheimer, and E. Rukzio, *Breaking the Tracking: Enabling Weight Perception Using Perceivable Tracking Offsets*, p. 1–12. New York, NY, USA: Association for Computing Machinery, 2018.
- [146] D. A. Gómez Jáuregui, F. Argelaguet Sanz, A.-H. Olivier, M. Marchal, F. Multon, and A. Lécuyer, “Toward pseudo-haptic avatars: Modifying the visual animation of self-avatar can simulate the perception of weight lifting,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, pp. 654–661, April 2014.
- [147] L. Dominjon, A. Lécuyer, J. Burkhardt, P. Richard, and S. Richir, “Influence of control/display ratio on the perception of mass of manipulated objects in virtual environments,” in *IEEE Proceedings. VR 2005. Virtual Reality, 2005.*, pp. 19–25, 2005.
- [148] M. Samad, E. Gatti, A. Hermes, H. Benko, and C. Parise, “Pseudo-haptic weight: Changing the perceived weight of virtual objects by manipulating control-display

- ratio,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, (New York, NY, USA), p. 1–13, Association for Computing Machinery, 2019.
- [149] R. Yu and D. A. Bowman, “Pseudo-haptic display of mass and mass distribution during object rotation in virtual reality,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 5, pp. 2094–2103, 2020.
- [150] M. Achibet, B. Le Gouis, M. Marchal, P.-A. Leziart, F. Argelaguet, A. Girard, A. Lécuyer, and H. Kajimoto, “Flexifingers: Multi-finger interaction in vr combining passive haptics and pseudo-haptics,” in *2017 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 103–106, IEEE, 2017.
- [151] T. Hachisu, G. Cirio, M. Marchal, A. Lécuyer, and H. Kajimoto, “Pseudo-haptic feedback augmented with visual and tactile vibrations,” in *2011 IEEE International Symposium on VR Innovation*, pp. 327–328, IEEE, 2011.
- [152] I. Jang and D. Lee, “On utilizing pseudo-haptics for cutaneous fingertip haptic device,” in *2014 IEEE Haptics Symposium (HAPTICS)*, pp. 635–639, IEEE, 2014.
- [153] M. Li, J. Konstantinova, E. L. Secco, A. Jiang, H. Liu, T. Nanayakkara, L. D. Seneviratne, P. Dasgupta, K. Althoefer, and H. A. Wurdemann, “Using visual cues to enhance haptic feedback for palpation on virtual model of soft tissue,” *Medical & biological engineering & computing*, vol. 53, no. 11, pp. 1177–1186, 2015.
- [154] Epic Games, “Unreal Engine: The most powerful real-time 3D creation platform.” <https://www.unrealengine.com/>. Accessed: 2021-01-02.
- [155] Unity Technologies, “Unity Real-Time Development Platform.” <https://unity.com/>. Accessed: 2021-01-02.

- [156] F. Conti, F. Barbagli, R. Balaniuk, M. Halg, C. Lu, D. Morris, L. Sentis, J. Warren, O. Khatib, and K. Salisbury, “The chai libraries,” in *Proceedings of Eurohaptics 2003*, (Dublin, Ireland), pp. 496–500, 2003.
- [157] Valve, “Steam VR Unity Plugin.” [https://valvesoftware.github.io/steamvr\\_unity\\_plugin/](https://valvesoftware.github.io/steamvr_unity_plugin/). Accessed: 2021-01-02.
- [158] T. Mitsuda, “Pseudo force display that applies pressure to the forearms,” *Presence*, vol. 22, no. 3, pp. 191–201, 2013.
- [159] A. Lécuyer, S. Coquillart, A. Kheddar, P. Richard, and P. Coiffet, “Pseudo-haptic feedback: can isometric input devices simulate force feedback?,” in *Proceedings IEEE Virtual Reality 2000 (Cat. No. 00CB37048)*, pp. 83–90, IEEE, 2000.
- [160] M. A. Srinivasan, G. L. Beauregard, D. L. Brock, and K. Danai, “The impact of visual information on the haptic perception of stiffness in virtual environments, session, asme; dynamic systems and control division,” in *ASME; Dynamic Systems and Control Division, ASME -PUBLICATIONS- DSC, Session, ASME; Dynamic Systems and Control Division*, vol. 58, pp. 555–560, ASME;, 1996.
- [161] M. O. Ernst and M. S. Banks, “Humans integrate visual and haptic information in a statistically optimal fashion,” *Nature*, vol. 415, no. 6870, pp. 429–433, 2002.
- [162] S. Gepshtein, J. Burge, M. O. Ernst, and M. S. Banks, “The combination of vision and touch depends on spatial proximity,” *Journal of vision*, vol. 5, no. 11, pp. 7–7, 2005.
- [163] H. B. Helbig and M. O. Ernst, “Optimal integration of shape information from vision and touch,” *Experimental brain research*, vol. 179, no. 4, pp. 595–606, 2007.

- [164] H. B. Helbig and M. O. Ernst, “Visual-haptic cue weighting is independent of modality-specific attention,” *Journal of vision*, vol. 8, no. 1, pp. 21–21, 2008.
- [165] R. J. van Beers, A. C. Sittig, and J. J. van der Gon Denier, “How humans combine simultaneous proprioceptive and visual position information,” *Experimental brain research*, vol. 111, no. 2, pp. 253–261, 1996.
- [166] R. J. Van Beers, A. C. Sittig, and J. J. D. v. d. Gon, “Integration of proprioceptive and visual position-information: An experimentally supported model,” *Journal of neurophysiology*, vol. 81, no. 3, pp. 1355–1364, 1999.
- [167] M. Kuschel, M. Di Luca, M. Buss, and R. L. Klatzky, “Combination and integration in the perception of visual-haptic compliance information,” *IEEE Transactions on Haptics*, vol. 3, no. 4, pp. 234–244, 2010.
- [168] C. Cellini, L. Kaim, and K. Drewing, “Visual and haptic integration in the estimation of softness of deformable objects,” *i-Perception*, vol. 4, no. 8, pp. 516–531, 2013.
- [169] 3D Systems, “Phantom Premium Haptic Devices.” <https://www.3dsystems.com/haptics-devices/3d-systems-phantom-premium>. Accessed: 2021-01-02.
- [170] T. Eiter and H. Mannila, “Computing discrete fréchet distance,” tech. rep., Technische Universität Wien, 1994.
- [171] D. M. Eagleman, “Human time perception and its illusions,” *Current opinion in neurobiology*, vol. 18, no. 2, pp. 131–136, 2008.
- [172] W. J. Matthews and W. H. Meck, “Time perception: the bad news and the good,” *Wiley Interdisciplinary Reviews: Cognitive Science*, vol. 5, no. 4, pp. 429–446, 2014.

- [173] N. Cellini, M. Fabbri, M. Martoni, L. Tonetti, and V. Natale, “Discontinuity in the perception of sub-second intervals,” *Procedia-Social and Behavioral Sciences*, vol. 126, pp. 222–223, 2014.
- [174] L. Shams and U. R. Beierholm, “Causal inference in perception,” *Trends in cognitive sciences*, vol. 14, no. 9, pp. 425–432, 2010.
- [175] K. P. Körding, U. Beierholm, W. J. Ma, S. Quartz, J. B. Tenenbaum, and L. Shams, “Causal inference in multisensory perception,” *PLoS one*, vol. 2, no. 9, p. e943, 2007.
- [176] J. Browder, S. Bochereau, F. van Beek, and R. King, “Stiffness in virtual contact events: A non-parametric bayesian approach,” in *2019 IEEE World Haptics Conference (WHC)*, pp. 515–520, IEEE, 2019.
- [177] A. Gotovos, N. Casati, G. Hitz, and A. Krause, “Active learning for level set estimation,” *IJCAI ’13*, p. 1344–1350, AAAI Press, 2013.
- [178] X. Gao, L. Wang, and X. Hao, “An improved capstan equation including power-law friction and bending rigidity for high performance yarn,” *Mechanism and Machine Theory*, vol. 90, pp. 84–94, 2015.
- [179] R. R. Kohlmeyer, A. J. Blake, J. O. Hardin, E. A. Carmona, J. Carpena-Núñez, B. Maruyama, J. D. Berrigan, H. Huang, and M. F. Durstock, “Composite batteries: a simple yet universal approach to 3d printable lithium-ion battery electrodes,” *Journal of Materials Chemistry A*, vol. 4, no. 43, pp. 16856–16864, 2016.
- [180] L. E. Asp, M. Johansson, G. Lindbergh, J. Xu, and D. Zenkert, “Structural battery composites: a review,” *Functional Composites and Structures*, vol. 1, no. 4, p. 042001, 2019.

- [181] Y. Gaffary, B. Le Gouis, M. Marchal, F. Argelaguet, B. Arnaldi, and A. Lécuyer, “Ar feels “softer” than vr: Haptic perception of stiffness in augmented versus virtual reality,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 11, pp. 2372–2377, 2017.
- [182] H. Chen, A. Dey, M. Billinghamurst, and R. W. Lindeman, “Exploring the design space for multi-sensory heart rate feedback in immersive virtual reality,” in *Proceedings of the 29th Australian conference on computer-human interaction*, pp. 108–116, 2017.
- [183] A. Aristidou and J. Lasenby, “Motion capture with constrained inverse kinematics for real-time hand tracking,” in *2010 4th International Symposium on Communications, Control and Signal Processing (ISCCSP)*, pp. 1–5, IEEE, 2010.
- [184] M. Höll, M. Oberweger, C. Arth, and V. Lepetit, “Efficient physics-based implementation for realistic hand-object interaction in virtual reality,” in *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 175–182, IEEE, 2018.
- [185] J.-S. Kim and J.-M. Park, “Physics-based hand interaction with virtual objects,” in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3814–3819, IEEE, 2015.
- [186] J. Jacobs and B. Froehlich, “A soft hand model for physically-based manipulation of virtual objects,” in *2011 IEEE Virtual Reality Conference*, pp. 11–18, IEEE, 2011.
- [187] A. Gupta, M. Samad, K. Kin, P. O. Kristensson, and H. Benko, “Investigating remote tactile feedback for mid-air text-entry in virtual reality,” in *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 350–360, IEEE, 2020.
- [188] M. Rietzler, G. Haas, T. Dreja, F. Geiselhart, and E. Rukzio, “Virtual muscle force: Communicating kinesthetic forces through pseudo-haptic feedback and muscle in-

put,” in *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, UIST '19, (New York, NY, USA), p. 913–922, Association for Computing Machinery, 2019.

[189] E. F. Melcer, M. T. Astolfi, M. Remaley, A. Berenzweig, and T. Giurgica-Tiron, “Ctrl-labs: Hand activity estimation and real-time control from neuromuscular signals,” in *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–4, 2018.

[190] E. Pezent, “IROS 2020 Haptics Tutorial - Introduction to Syntacts.” <http://iros-haptics-tutorial.org/>. Accessed: 2021-01-02.

[191] Syntacts, “Syntacts: The Tactor Synthesizer.” <https://www.syntacts.org/>. Accessed: 2021-01-02.