2025 IEEE World Haptics Conference (WHC)
Suwon Convention Center, Suwon, Korea
July 8 ~ 11, 2025

Hands-On or Hands-Off? Active Touch Influences Multisensory Perception of Referred Haptics

Elyse D. Z. Chase
Dept. of Mechanical Engineering
Rice University
Houston, USA
ec100@rice.edu 0009-0002-3398-5048

Daziyah H. Sullivan
Dept. of Mechanical Engineering
Rice University
Houston, USA
0000-0002-5181-8668

Marcia K. O'Malley
Dept. of Mechanical Engineering
Rice University
Houston, USA
0000-0002-3563-1051

***Abstract*— Wearable haptic devices can provide haptic feedback in both active and passive touch interactions. In virtual or extended reality environments, wearable devices also enable referred haptic feedback, where touch sensation expected in one place on the body is directed to another. How people perceive the distinction between these feedback forms has been relatively under-explored, especially considering the additional role that vision plays in our multisensory understanding of the world. To explore how active and passive touch affect the perception of referred haptic feedback, we conducted an experiment in VR where participants chose the stiffer of two springs in a 2-interval, 2-alternative forced-choice design. We find that participants can be categorized into two groups based on the different strategies they employ for making the decision – those with a haptic or a visual prior, similar to related work. We also considered both active and passive feedback conditions. Notably, people are more accurate in judging haptic stiffness during the active case. Our results have implications for designers of virtual systems and simulations where users receive various sensory inputs, both via active and passive interactions, with potential mismatches due to latency, bandwidth, or design issues.**

***Index Terms*—passive touch, squeeze, force feedback, wearable devices, virtual environments, virtual reality (VR), stiffness perception, springs, yoked**

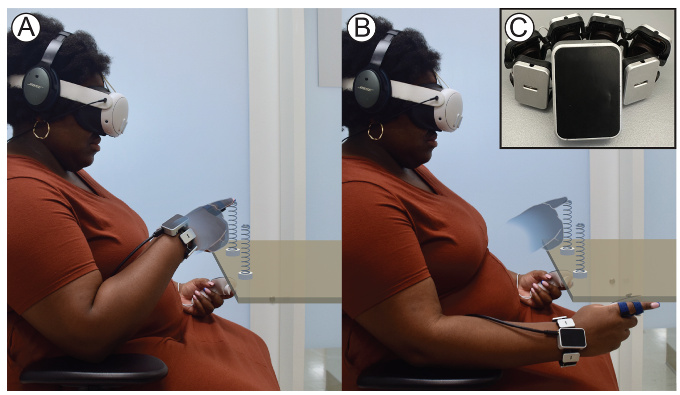

Fig. 1. *Study Setup*: A participant is seated in an empty study space. Virtual items are shown semi-transparently. Participants virtually see two springs on a table and a proxy (virtual) hand that sometimes tracks their real hand. In addition to the Quest 3 HMD, users wear noise-canceling headphones and hold two Quest Controllers. **(A) Active**: The participant reaches out to compress the springs as desired. **(B) Passive**: The participant sits with both arms on the armrests while the proxy hand moves to compress the springs. Forces are displayed via a wrist-worn haptic device. (C) The Tasbi haptic device is worn on the right wrist and provides squeezing forces.

I. Introduction

Multisensory perception and interaction are part of everyday life. Consider sending a text message using your smartphone. With each key press, you receive subtle vibrations or a "click" sound. You might think the device was malfunctioning if you received these cues while not interacting with your phone or touching anywhere else than the keypad. Vision or sound can strongly influence how we perceive the same haptic feedback. Where and when we perceive the haptic feedback influences our perception of it. Whether we actively cause or passively receive the haptic feedback also affects our sense of what is happening [1].

This work was supported by an appointment to the Intelligence Community Postdoctoral Research Fellowship Program at Rice University administered by Oak Ridge Institute for Science and Education (ORISE) through an interagency agreement between the U.S. Department of Energy and the Office of the Director of National Intelligence (ODNI) and in part by the National Science Foundation Graduate Research Fellowship Program under Grant No. (2236422). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

For haptic feedback in *active* touch scenarios, the user's actions directly result in feedback from the system. In contrast, in *passive* touch scenarios, haptic feedback is provided in a manner that is not directly dependent on the user's actions. Simply put, active haptics is *touching*, while passive haptics is *being touched* [2]. These interaction methods vary the user's agency, control over one's action, over the interaction – with active touch allowing the user to be in charge [3]. Each mode is of substantial importance, and oftentimes, we experience the world in a combinatorial fashion – with separate processes to interpret these different forms of feedback. Both types of feedback are critical to how we interpret our interactions with the world through touch [4].

Wearable haptic devices have become increasingly common and can use *sensory substitution*, where a different modality delivers information than would usually be experienced. Referred haptic feedback is a type of sensory substitution where haptic information that should be felt at one location is moved (referred) to an alternative location. For example, a haptic wristband can display haptic feedback that would normally occur at the fingertips. This method allows the entire hand (palm and fingers) to remain free for additional touching and grasping interactions – expanding the range of applications where these devices could be used because there are fewer restrictions on the user. Prior work has even shown that integrating visual pseudo-haptic illusions and referred haptic feedback can strengthen the range of perception of an object's stiffness beyond either modality alone [5].

However, we do not know much about how referred haptic feedback is perceived under active and passive touch conditions, even though these conditions occur during virtual reality (VR) interactions. There are also cases where there are slight mismatches between what the user does and what a system displays, from computational issues, like lag, to designed interactions, like purposeful retargeting of a user's movement [6]. As VR becomes increasingly prevalent and includes multisensory interactions, we must understand how these concepts contribute to our perception when combined.

Here, we test the combination of active and passive touch with referred haptics under visual mismatch (Fig. 1). In a user study, we asked participants to judge spring stiffness in a 2-interval 2-alternative forced-choice task, initially conceived by [7], and later expanded to referred haptic feedback at the wrist and VR by [8]. Participants can interact and move as desired to determine relative spring stiffness in a free exploration, active condition. We introduce a second, yoked condition where we replay their explorations and haptic feedback while they passively sit with their arm resting. Two groups emerge from the data: those with a haptic prior and those with a visual prior. We also find differences in how participants judge between the active and passive conditions, specifically in cases without visual mismatch. During the active case, participants are more accurate at judging haptic stiffness.

II. Related Work

Touch is a combination of active input and the sensation that follows. Thus, we know that the timing of active and passive touch differs; sensation requires backward information, while active touch ignites our sensorimotor control system [9]. However, in this work, though there was a difference in the timing of perception between active and passive touch, it was not significant. Several works have expressed doubt about the differences between active and passive touch, specifically doubting the importance of self-produced motion but rather finding that it was differences between the administration of cues in the passive compared to the active condition [10], [11]. Others have claimed that any perceived differences between these types of touch are only conceptual, due to a good internal understanding of the task and linkage of information gathered to those concepts [12]. More recent work has attributed the differences to how much information one can gain in the active compared to the passive case and considerations of cognitive load [13].

Beyond cognitive factors, the timing of multisensory signals plays a critical role in determining how sensory inputs are integrated into a coherent perceptual experience. Small changes in the timing of signals, including audio-tactile [14], [15] and visual-haptic [16], [17] interactions, can significantly influence perceived intensity, response accuracy, and experience. These findings emphasize the importance of synchrony for perceptual coherence. They also align with the perspective that passive touch has some inherent latency and reduced proprioceptive feedback that could lead to diminished haptic accuracy, particularly when combined with visual discrepancies.

Other research with haptic feedback has found distinct differences between active and passive touch. Vogels determined that active and passive touch affect our perception of asynchrony, with more sensitivity in the passive condition compared to the active [17]. Other work in haptics has found differences in performance between active and passive feedback [18], [19]. The relationship between haptic feedback and agency has also been studied [3] – specifically in virtual environments. Evangelou et al. found that adding mid-air haptics can increase a user's sense of embodiment in VR [20]. Bergstrom-Lehtovirta et al. found that self-skin contact (one hand touching the other hand for confirmation of a button press) had higher agency than a button or computer touchpad interface [21]. Others have explored methods to increase user response time without compromising agency [22]. Many of these works try to increase the embodiment that users feel in VR due to mismatches between the real and virtual body [23].

There is also a strong relationship between agency and causal perception. While agency answers, "did I do that?", causality answers a lower-level question, "did A make B do that?". To have agency, there must be a causal relationship in addition to an intentional action by an agent [24]. Thus, having agency is inherently active, while causal relationships may be active or passive. Prior work has demonstrated that people use kinesthetic and vibrotactile information when making causal inferences [25] and that the realism of the sensory experience (across multiple sensory modes) increases these causal linkages [26]. The type of touch, whether active or passive, is also thought to be important in how people interpret these scenarios [27]. When an observer produces motions themselves, they can better predict possible changes in the perception of movement – and thus might come to different causal conclusions.

Wearable devices have been employed to deliver haptic feedback in both active and passive touch scenarios. Wristband devices have been used to communicate passively received information [28], like language [29], guidance [30], [31], and social touch [32], to active information such as material properties [5], [8]. Zook and O'Malley [33] used a wrist-worn haptic device to explore agency and perception. They found that users' focus direction and sense of agency affected their perception of the tactile cues. Thus, especially in the case of wearable haptic devices, it is key to consider the type of touch that will be administered and the user's state.

III. Methods

We conducted an experiment in VR where participants chose the stiffer of two springs in a 2-interval, 2-alternative forced-choice design [34], [35]. Below, we introduce the necessary hardware, software, explanation of conditions, choice of experimental parameters, and haptic rendering algorithms.

A. Hardware & Software

We used the Tasbi haptic wristband [36], which has a custom tension mechanism that produces controllable squeeze around the wrist. We used position-based control during the experiment to modulate the squeeze cue felt by users. During setup with each participant, we used a pre-developed, torque-based control to tighten to a predetermined nominal low torque value, which allowed us to calibrate and provide a similar initial tightness across participants – developed by [36]. This fit ensures the bracelet will not rotate on the wrist, but is not too tight on the participant. After determining the baseline fit and, thus, the resting position, we applied maximum torque and found the spooling range of the device. All remaining interactions were calculated based on each participant's range of available spooling.

It is important to note that position-based control, compared to alternatives such as force-based control, may not provide consistent results if participants hold their wrist at flexion or extension compared to the wrist's orientation at calibration. Thus, we ensured that we performed calibration with the wrist at the correct location for the task, and we instructed participants to maintain a neutral position of their right hand and wrist while interacting with our VR environment.

In addition to the haptic bracelet, all study participants donned a Meta Quest 3 Head-Mounted Display (HMD) and held two Touch Plus Controllers. We used the controllers for simple system inputs. Participants also wore noise-canceling headphones playing pink noise to mask motor sounds (Fig. 1). Here, VR enabled us to modify the visual location of participants' hands during both the active and passive conditions while masking the location of their real hands from therm. We developed this study with the Unity Game Engine (Version 2022.3.15f1) and the Meta Interaction SDK.

B. Active & Passive Touch Conditions

Active touch is how we normally interact with the world; you are in charge of your own actions and receive feedback that is coordinated with those actions. In contrast, passive touch removes that control and allows another source to govern what information a user feels. This non-volitional information could be chosen from many sources, which can affect how users perceive that information.

We wanted to provide passive and active conditions that were equivalent for each participant – meaning that the interaction method and strategies were consistent across the two types of touch. Thus, we adopted a *yoked* condition, in which participants' own interactions are coupled to a particular trial and then played back to them. In the passive condition, participants received visual and haptic feedback corresponding to pre-recorded movements from their own active trials. While this approach provides a controlled passive experience, it may introduce residual proprioceptive cues from participants' prior active interactions, potentially influencing their perception.

Often, yoked control paradigms group together two participants, where one experimental subject performs some task, and the second yoked subject experiences the first experimental subject's performance [18], [37]. In our work, we chose a within-subjects version of yoking to ensure that differences in strategies between participants did not affect the results [38].

1) Data Saving: To replay the user's position, we recorded information about the location of the right fingertip and the base of the wrist during the active condition. We used this data to render the proxy hand and virtual spring in the passive condition. Specifically, we overrode the location of the participant's real wrist with that of the corresponding active trial. Using the right fingertip data, in conjunction with saved information about spring compression, and rendered the spring accordingly. We directly replayed the forces displayed from the previous trials to the user's wrist.

To maintain embodiment from the movement of other fingers but restrict the index finger to the same position across trials, we provided participants with a finger splint for their right index finger (Fig. 1). The splint helped them maintain their hand pose for the passive condition while supporting them and reducing fatigue throughout the experiment.

C. Parameter Choice

To render two springs for comparison, we need to select the values for a reference stiffness and several spring stiffness deltas (K_Δ). The reference stiffness value will always be assigned to one of the springs. Then we can use the reference stiffness and spring stiffness deltas to calculate our comparison stiffness values for the other spring.

We used the same reference stiffness (K_o = 50) as our prior work that used the same haptic device [8]. With this, we can calculate comparison stiffness values, K_c (Eq. 1). We selected two values for K_Δ, 0.3 and 0.6, which allowed us to compare any differences in perception that may occur with smaller and larger differences in stiffness.

$$K_c = K_o(1 + K_\Delta) \tag{1}$$

We also use a visual scaling parameter, $\lambda \in [0, 1]$, first defined by [7]. We use λ to create a dimensionless scaling factor (Eq. 2) where the denominator is a weighted average of the two spring stiffnesses (K_c and K_o).

$$k_{scaling} = \frac{K_c}{(1-\lambda)\ K_c + \lambda\ K_o} \tag{2}$$

The numerator ensures that two conditions are held. When $\lambda = 0$, the scaling factor is $\frac{K_c}{K_c} = 1$. When combined with Hooke's Law, the spring constant remains the same: $F = k_{scaling} K_o x = K_o x$. When $\lambda = 1$, $k_{scaling} = \frac{K_c}{K_o}$, which results in a switch of the spring constant: $F = k_{scaling} K_o x = \frac{K_c}{K_o} K_o x = K_c x$. Any value of λ in between will be a weighted average of the two different spring stiffnesses.

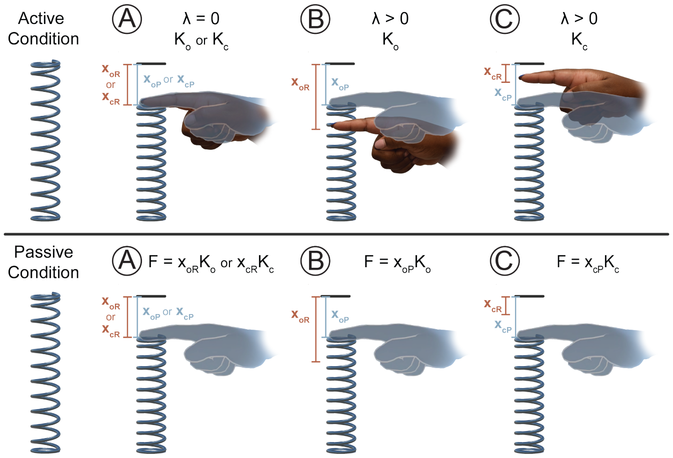

Fig. 2. *Rendering of Virtual Springs and Hands*: real, physical hand shown in color and proxy, virtual hand shown in blue. Forces are calculated the same way for both conditions. **Row 1 – Active Touch**: (A) $\lambda = 0$ either spring, the two hands move identically ($x_{oR} = x_{oP}$ or $x_{cR} = x_{cP}$); (B) $\lambda > 0$ reference spring (K_o), the real hand moves more than the proxy hand ($x_{oR} > x_{oP}$); (C) $\lambda > 0$ comparison spring which by definition ($K_c > K_o$), the real hand moves less than the proxy hand ($x_{cR} < x_{cP}$). **Row 2 – Passive Touch**: Forces relate directly to the yoked position of the user's real and proxy hand from a previous trial. (A) $x_{oR} = x_{oP}$ or $x_{cR} = x_{cP}$; (B) $x_{oR} > x_{oP}$; (C) $x_{cR} < x_{cP}$.

D. Rendering

Prior work has explored rendering the control-to-display (C/D) ratio with impedance [8] and admittance control devices [7]. Tasbi is an impedance-type device, so we use a rendering schema similar to our previous work [8] with some modifications.

Using the parameters defined in the previous section, we can visually produce three different types of interactions depending on the $0 \leq \lambda \leq 1$ of that trial (Fig. 2). These interactions are akin to modifying the C/D ratio, which modifies the relative locations of the real (x_{oR} or x_{cR}) and proxy (x_{oP} or x_{cP}) hands through retargeting [6] (here, $_o$ and $_c$ refer to the motion of the hand touching the reference and comparison springs). There is no retargeting when $\lambda = 0$, but the real and proxy hands are offset in all other cases ($\lambda \neq 0$). This difference in location can make the springs appear visually less or more stiff, depending on the direction of the change. For example, compare B and C in Fig. 2, Row 1. The parameters in B require the user to move their real hand more for a slight compression in the spring, visualized by a small movement with the proxy hand, while the reverse is true in C (small real hand movements result in large proxy hand movements).

In the passive condition (Fig. 2, Row 2), we manipulate the location of the proxy hand, but without any physical relationship to the real hand location, as the user rests their arm during this condition. While participants may infer the speed of the proxy hand's motion, this might affect interpretations when there is a visual difference ($\lambda > 0$). However, the squeeze cues presented are identical across both active and passive conditions per each trial.

Below, we define the calculations for rendering the two springs – the reference spring (K_o) and the comparison spring (K_c), whose relationship was defined above in Eq. 1, such that $K_o < K_c$. These two springs are influenced by the location of the real hand (x_{oR} or x_{cR}). However, we do not always render the real hand in its true position and thus introduce the location of a proxy hand (x_{oP} or x_{cP}), which aligns with the top of the spring during compression. We use the following rendering equations:

$$F = K_o \, x_{oP} \ , \ \ x_{oP} = \frac{K_c}{(1-\lambda)\ K_c + \lambda\ K_o} \, x_{oR} \tag{3}$$

$$F = K_c \, x_{cP} \ , \ \ x_{cP} = \frac{K_o}{(1-\lambda)\ K_o + \lambda\ K_c} \, x_{cR} \tag{4}$$

We calculate the forces output by the haptic device (F) using a combination of the proxy's displacement (x_{oP} or x_{cP}) and the virtual spring's stiffness (K_o or K_c). Similar to [8], we use force related to the proxy hand's position rather than the real hand's position (x_{oR} or x_{cR}), used by [7]. Our method results in a realistic feeling interaction compared to using the real hand's position, which can result in larger forces at the interaction surface but no additional forces at full compression – at least in cases when $\lambda \neq 0$.

In this work, we exclude the scaling factor used previously [8]. The scaling factor increased the visual displacement. We determined, through pilot testing, that including the scaling factor was unnecessary for perceptible results with our selection of parameters. Its elimination simplifies the interpretation of later results.

IV. Hypotheses

Prior work observed two different approaches among the participants in their study, which were explained by those with haptic or visual priors [8]. So, we hypothesize that we will see the same distinction since we are employing the same task. But now, with the expansion of the study to include both active and passive touch conditions, we expect that those two sensory modality groups will persist regardless of the additional conditions.

(H1) *Participants will self-organize into two sensory modality groups: those with a haptic prior and those with a visual prior. Those with the haptic prior will be more accurate in selecting the haptically stiffer spring than those with a visual prior.*

From both [7], [8], we expect that as λ increases, accuracy in picking the haptically stiffer spring will decrease for the participants with a visual prior. We expect no effect of lambda on the haptic group as they prefer to weigh the haptic information heavier than the visual information.

(H2 a) *Participants with a visual prior will be less accurate in discriminating object haptic stiffness as visual discrepancy (λ) increases.*

(H2 b) *Participants with a haptic prior will have no change in haptic accuracy as visual discrepancy (λ) changes.*

Furthermore, [8] found that there was an interaction between λ and K_Δ. In their work, this effect manifested as lower haptic accuracy when there was a larger stiffness difference (K_Δ) and as the visual difference increased ($\lambda > 0$). Here, we have removed the visual offset (ϵ), which could have influenced this result. Therefore, we hypothesize that there will be a main effect of K_Δ.

(H3) *For those participants with a visual prior, there will be a decrease in haptic accuracy with an increase of K_Δ.*

This means that people are more likely to use the visual information and will weigh that information more heavily than the haptic cues. Specifically, in the case of a larger K_Δ, there will also be larger proprioceptive differences and visual changes (λ) because of the relationship between these variables.

In the passive case, compared to the active case, people lose agency over the interaction in addition to proprioceptive information, as their hands rest during these trials. For both groups, we predict that the loss of agency will cause them to pay less attention to the information presented, specifically the squeeze cues at the wrist that are now farther away from the location where the force should occur.

(H4) *In the active condition, compared to the passive condition, participants will be more accurate in identifying the haptically stiffer spring when there is no visual difference ($\lambda = 0$).*

We have no set hypotheses for any remaining differences and instead complete exploratory analyses.

V. Study

13 participants (age: μ = 24.2, σ = 4.6, range = [19, 35]; sex: 3 female, 10 male; 12 right-handed, 1 ambidextrous) took part in the study and were compensated \$15. All participants gave informed consent, and the protocol was approved by the Rice University Institutional Review Board (IRB-FY2019-49).

A. Experimental Setup

1) Design: The study is a within-subjects, repeated measures design with three factors: condition (2 levels: active, passive), spring stiffness delta (K_Δ, 2 levels: 0.3, 0.6), and visual difference (λ, 5 levels: 0, 0.25, 0.5, 0.75, 1). There were 10 repetitions per combination, for 200 trials total[1].

The task was to determine which of the two springs was stiffer. We structured the active touch portion of this experiment as a two-interval, two-alternative forced choice paradigm (2I-2AFC) [34], [35]. We selected this study type to compare our results with those of related work [7], [8]. Additionally, it allows for free exploration in the active condition while providing a foundation for the yoked condition.

[1]Due to an experimental error, two participants completed 11 repetitions, 220 trials. However, we see no evidence of this affecting their results.

We grouped trials into two halves, with two blocks per half. The halves were the *active* and *passive* touch conditions. Because of the yoked method, all participants began with the active half and then completed the passive half. While this can introduce bias, it was necessary for the yoked design.

The two blocks in each half were the two K_Δ values, presented randomly. Within a block, we kept K_Δ constant to create a method of constant stimuli [35]. Half of all trials had the reference spring on the right, while the remaining half had the reference spring on the left to remove any effect of location. We pseudo-randomized the 10 repetitions in a block by grouping the first and second halves of the repetitions. The pseudo-randomization provided a more even distribution of trial presentation and, for the passive condition, kept trials that were earlier near the front half of the study. Presenting the earlier trials near the beginning of the block allowed for better mimicry of real time spent in the trial, as people often spend more time on earlier trials while developing their strategy.

2) Procedure: Participants received an overview of the study and then completed our consent forms. We then introduced them to the headset and controllers, reviewing necessary buttons and knobs for interaction, including the interpupillary distance (IPD) to ensure participants could see clearly. Once seated in an open area, we adjusted the armrests for comfort – resting the elbow (if desired) during the active condition and the entire arm during the passive condition. We then fit the Tasbi to each user using the fitting routine described in Methods. Participants donned the HMD, adjusted the IPD, and received the controllers. They calibrated the height of the virtual scene to be within their comfortable reach range. We attached a finger splint to their right index finger to help maintain their outstretched posture (Fig. 1) while reducing fatigue. Finally, we placed noise-canceling headphones over the HMD and played pink noise to eliminate any effects of sound from the motor in the haptic device.

Participants completed two practice trials to familiarize themselves with the task. Across both practice trials, we presented them with the largest stiffness difference ($K_\Delta = 0.6$) and no visual difference ($\lambda = 0$), which gave users the most distinct example of haptic discrimination that could occur in the study. The first practice was *active*, where users could interact with the springs freely and move on when ready. The second practice was *passive*, where participants felt and observed a virtual hand touching the springs for them. During this, we used a prerecorded file explicitly created for practice, where the hand touches each of the springs at a slow, measured rate, and similarly between the left and the right spring. We did not use a yoked condition for the practice so that people would not be alerted that hand motions would be recorded and played back later in the study. After asking any clarifying questions, participants started the experiment.

In each trial, we required that participants touch both virtual springs from the top surface with their right index finger. No visual or haptic feedback was provided if they touched the springs from any other location (e.g., laterally or finger not aligned with the top surface). We asked participants to move at a slow, measured rate and to fully decompress the spring before removing the finger to keep hand motion recordings smooth and accurate for replay in the second condition. During the *active* condition, people could interact with the springs as much as they wanted, minimally one interaction per spring, before moving on to following the question:

Which spring was stiffer?[2]

Participants then used the controller to select either *Left* or *Right*. If they moved on before being ready to make a decision, we allowed them to return and interact with the springs more.

In the *passive* condition, we instructed participants to relax onto the chair's armrest while keeping their wrist off the edge of the armrest, allowing the haptic device freedom to squeeze. Additionally, we asked them to maintain the finger-pointing pose with their right index finger from the active condition, aided by the finger splint. This request increased the embodiment between the proxy and the real hand.

Across the study, we required participants to take breaks of at least 30 seconds between the four blocks. Between study halves (the transition from *active* to *passive*), we displayed a screen to remind participants about the differences in the passive condition. Specifically, in these trials, they should relax their arm and focus as they cannot go back and replay the trial. At the end of the experiment, participants filled out a survey with demographic information and open-ended questions about their interactions. Participants completed the instructions, consent, and experiment in 60 minutes or less.

VI. Results

In this experiment, our primary dependent variable is response accuracy. This is binary as participants either selected the **haptically stiffer spring** (1) or not (0). When $\lambda \neq 0$, the other option (0) appeared to be the **visually stiffer spring**. Several independent variables could also affect accuracy, including condition (*active* v. *passive*), λ, and K_Δ.

We grouped participants into the haptic or visual prior by considering their average accuracy. Those with a haptic prior (n = 6) were $\geq$ 95% accurate across all conditions. These prior groupings were confirmed by considering all individuals' results graphically, as there are strong differences.

We fit the generalized linear mixed-effects models to the data with the appropriate family and a probit link function. We use Bayesian methods to fit models with the *brms* package in R [39]. We used approximate leave-one-out cross-validation with the *loo* function [40] to determine the model that best predicts response accuracy. Our inference criterion was that the 95% credible interval (CrI) excludes zero. When appropriate, we report β values in addition to the 95% CrI. The variable β is the regression coefficient, which is a measure of the average change in the dependent variable for a one-unit increase in the independent variable.

[2]We took this question from our prior work [8], to remove any weighing of one sensory mode above another (e.g., haptic versus visual feedback). No further clarification was given to participants.

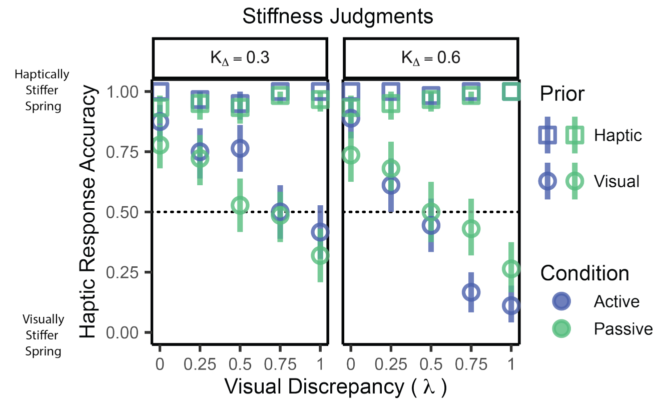

Fig. 3. Stiffness Judgments: grouped into two columns by increasing K_Δ. The x-axis shows increasing visual discrepancy (λ), where 0 is no visual change, and 1 is a visual inversion of stiffness. The y-axis shows haptic response accuracy, where 1 indicates 100% accuracy in selecting the stiffer spring and 0 indicates selecting the visually stiffer spring. The dashed line marks where people would be 50/50 guessing between which spring is haptically stiffer. Shape shows *prior*: haptic (square) and visual (circle). Color shows *condition*: active (blue) and passive (green). Shapes mark the mean, and lines show the 95% Bootstrapped CI.

As noted in the Study Section, two participants completed one extra repetition of each combination, five extra trials per block. To confirm that these extra trials did not affect the results, we ran the tests with and without the additional trials. There were no differences in the findings (although slight variations in estimates and CrI would occur). We report on results with all of the trials, as this is representative of how the participants experienced the study. This increases the weight of the graphical judgments by those two participants by 10%.

A. Confirmatory Analyses

To test H1 and H2, we fit two different versions of the model: one simple (M0) and one more complex (M1) with an interaction effect. Both models have main effects of K_Δ, λ, and *prior* and a random effect of *subject*. M1 contains the additional interaction between λ and *prior*.

M0: single-trial accuracy $\sim 1 + K_\Delta + \lambda + \text{prior} + (1 \mid \text{subject})$

M1: single-trial accuracy $\sim 1 + K_\Delta + \lambda * \text{prior} + (1 \mid \text{subject})$

Comparing M1 to M0, M1 is a better predictor of the trial accuracy data ($elpd_{loo} = 24.7$, $SE = 7.5$). Therefore, we use M1 to address H1.

For H2 and H3, we use M1 for an initial analysis. For more specific results between the prior groups, we fit additional models on the data from one prior group at a time, as follows:

M2: single-trial accuracy $\sim 1 + K_\Delta + \lambda + (1 \mid \text{subject})$

For H4, we consider an additional model (M3) that is fit only on the data when $\lambda = 0$. The main effect of *condition* is coded as either active or passive. We consider more complex models in the exploratory analysis.

M3: single-trial accuracy $\sim 1 + K_\Delta + \text{condition} + (1 \mid \text{subject})$

1) H1: Haptic vs. Visual Priors: From M1, we find a main effect of *prior*, where the visual prior group has lower haptic response accuracy compared to the haptic prior group ($\beta = -1.8$, 95% CrI $[-2.92, -0.73]$). The two groups have different strategies (Fig. 3), with the haptic prior group maintaining high accuracy across all other study conditions, while the visual prior group decreases with increasing visual discrepancies. These results support our hypothesis that sensory priors separate two groups and that the haptic prior group is more accurate in this task.

2) H2: Visual Discrepancy: With M1, we find a main effect of λ and an interaction between λ and *prior*. The interaction indicates that for the visual group, accuracy decreases as λ increases ($\beta = -1.80$, 95% CrI $[-3.17, -0.52]$). The main effect suggests that as λ increases, so does accuracy ($\beta = 1.34$, 95% CrI $[0.25, 2.47]$). To better address the sub-parts of this hypothesis, we consider models fit with just that particular group's data.

a) Visual Prior: Using M2 fit only with data for the visual prior participants, we find a main effect of λ, whereas λ increases, accuracy decreases ($\beta = -2.85$, 95% CrI = $[-3.18, -2.53]$). These results support our hypothesis that those with a visual prior will be affected by visual changes.

b) Haptic Prior: Using M2 fit only with data for the haptic prior participants, we find an effect of λ where accuracy increases with λ ($\beta = 1.33$, 95% CrI = $[0.23, 2.47]$). This result does not support our hypothesis, as it appears that people are more accurate in determining the haptically stiffer spring in cases with increasing visual discrepancy.

3) H3: Stiffness Difference: Using M1, we find a main effect of K_Δ across all data which shows an inverse relationship between K_Δ and accuracy ($\beta = -1.96$, 95% CrI = $[-2.72, -1.18]$). As our hypothesis was explicitly for those with a visual prior, we also consider M2, where we find a main effect of K_Δ. As K_Δ increases, accuracy decreases ($\beta = -2.37$, 95% CrI = $[-3.06, -1.69]$). This finding supports our hypothesis that there is a decrease in accuracy with strong stiffness differences for those participants with a visual prior. For completeness, if you consider the haptic prior group fit on M2, we find no effect of K_Δ ($\beta = 1.95$, 95% CrI = $[-0.56, 4.60]$).

4) H4: Active vs. Passive: From M3, we find a main effect of *condition*, but no effect of K_Δ ($\beta = -0.30$, 95% CrI = $[-2.28, 1.64]$). The passive condition is less accurate compared to the active condition ($\beta = -1.20$, 95% CrI = $[-1.76, -0.68]$). As this data is only fit on cases when there is no visual difference ($\lambda = 0$), this supports our hypothesis that people will be less accurate at determining the haptic stiffness in the passive compared to the active condition (Fig. 4).

B. Exploratory Analyses

1) Active vs. Passive: Our main interest was to compare the active and passive conditions, with one hypothesis about the case without any visual changes. Here, we consider several exploratory models that encompass the entire set of data and consider an additional main effect of *condition* (M4) as well as an interaction between *condition* and *prior* (M5).

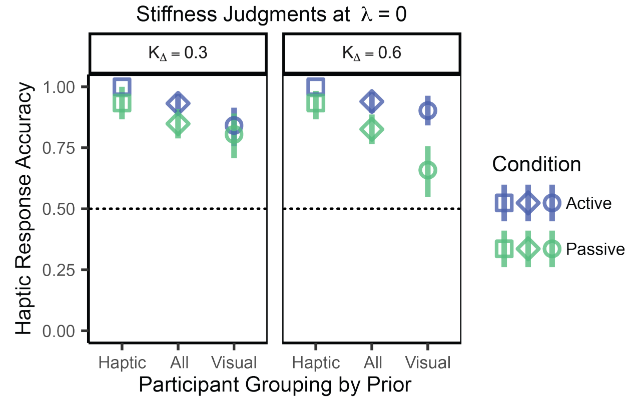

Fig. 4. Active vs. Passive Stiffness Judgments: grouped into two columns by increasing K_Δ. The x-axis shows the two prior-based groups on the left (Haptic) and the right (Visual), with all people lumped together in the middle. The y-axis shows haptic response accuracy, where 1 indicates 100% accuracy in selecting the stiffer spring and 0 indicates selecting the visually stiffer spring. The dashed line marks where people would be 50/50 guessing between which spring is haptically stiffer. Shape shows *prior*: haptic prior (square), visual prior (circle), all participants grouped (diamond). Color shows *condition*: active (blue) and passive (green). Shapes mark the mean, and lines show the 95% Bootstrapped CI.

M4: single-trial accuracy $\sim 1 + K_\Delta + \lambda + \text{condition} + \text{prior} + (1 \mid \text{subject})$

M5: single-trial accuracy $\sim 1 + K_\Delta + \lambda + \text{condition} * \text{prior} + (1 \mid \text{subject})$

M5 is a better predictor of the data than M4 ($elpd_{loo} = 4.0$, $SE = 3.4$). From M5, all effects are notable. Several results align with our confirmatory findings. The visual prior group is less accurate than the haptic prior group ($\beta = -5.12$, 95% CrI = $[-6.67, -3.65]$). There is an inverse relationship between accuracy and both K_Δ ($\beta = -1.88$, 95% CrI = $[-2.66, -1.10]$) and λ ($\beta = -2.41$, 95% CrI = $[-2.76, -2.08]$).

Additionally, we find there is an effect of *condition*, where the passive condition is less haptically accurate in judging stiffness than the active condition ($\beta = -1.52$, 95% CrI = $[-2.54, -0.62]$). The interaction effect between *condition* and *prior* is also significant. It indicates there is a positive interaction between the passive condition and the visual prior group ($\beta = 1.48$, 95% CrI = $[0.55, 2.53]$). This is likely due to the slight increase in accuracy seen for the visual prior group as visual discrepancy increases compared to the active case. This increase in accuracy is greater than the change between conditions in the haptic prior group.

2) Interaction Methods: In addition to perceptual responses, we also have data on the participants' interactions during each trial. First, for the active condition, we consider how much time participants spent interacting with either spring. Second, for the passive condition, we calculate how much each participant moved their wrist while the proxy hand was interacting with the spring. Similar to the above analyses,

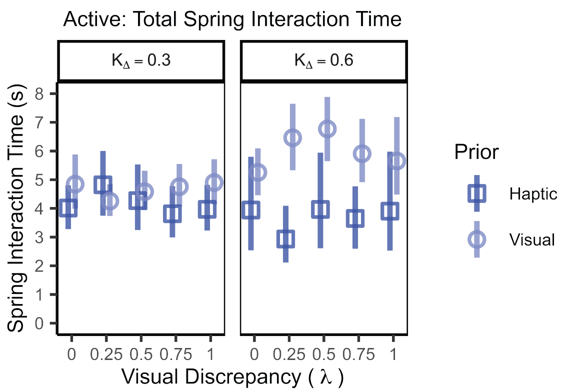

Fig. 5. Active Spring Interaction Time: grouped into two columns by increasing K_Δ. The x-axis shows increasing visual discrepancy (λ), where 0 is no visual change, and 1 is a visual inversion of stiffness. The y-axis marks the time touching the springs in seconds. Shape shows *prior*: haptic (square) and visual (circle). Shapes mark the mean, and lines show the 95% Bootstrapped CI.

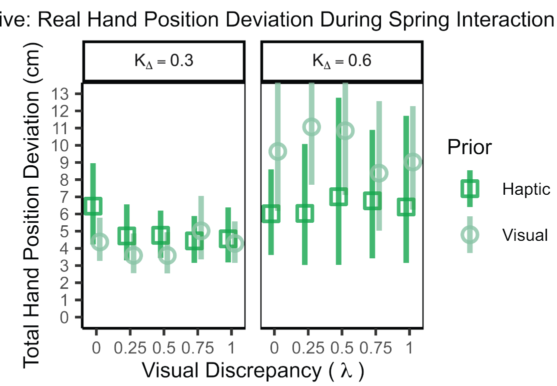

Fig. 6. Passive Hand Motions: grouped into two columns by increasing K_Δ. The x-axis shows increasing visual discrepancy (λ), where 0 is no visual change, and 1 is a visual inversion of stiffness. The y-axis marks summed hand position deviation (in centimeters) while the proxy hand compressed the spring. Shape shows *prior*: haptic (square) and visual (circle). Shapes mark the mean, and lines show the 95% Bootstrapped CI.

we fit models to see how our independent variables relate to these outputs. For both, we consider two models, each fit on either the active or passive tracked hand motion data, one simple (I1) and one more complex (I2):

I1: interaction metric $\sim 1 + \text{prior} + K_\Delta + \lambda + (1 \mid \text{subject})$

I2: interaction metric $\sim 1 * \text{prior} + K_\Delta + \lambda + (1 \mid \text{subject})$

For the active case, the *interaction metric* is the total time spent compressing either spring. In the passive case, the summed hand deviation is calculated for times when the proxy hand compresses the spring.

a) Time: Comparing I1 to I2 fit on the active condition data only, I2 is a better predictor of the total time spent compressing either spring ($elpd_{loo} = 5.2$, $SE = 5.2$). Within that model, we find a notable interaction effect between K_Δ and *prior* ($\beta = 6.13$, 95% CrI = $[3.12, 9.13]$). In Fig. 5, there are larger interaction times in the visual prior group with the haptically stiffer spring difference ($K_\Delta = 0.6$, right column) compared to the smaller spring stiffness delta.

b) Motion: Comparing I1 to I2 fit on the passive condition data, I2 better predicts the summed hand deviation when the proxy hand compresses the spring ($elpd_{loo} = 5.2$, $SE = 3.5$). In that model, we find a slight, but still notable, interaction effect of K_Δ and *prior* ($\beta = 0.14$, 95% CrI = $[0.05, 0.22]$). This result indicates that when the difference in spring stiffness was larger, those in the visual prior group moved their real hands more during the passive condition (i.e., while their proxy hand was recreating their pre-recorded movements). In Fig. 6, the visual prior group (circles) has increased real hand motion during the stiffer spring comparison trials compared to the less stiff spring.

VII. Discussion

Our user study explored how humans perceive active and passive feedback with referred haptics and visual mismatch. We gathered human perceptual and interaction data using a yoked condition where people could interact passively with springs and feel the same forces at the wrist, precisely as they had previously done in the corresponding active scenarios. Our study results support several of our hypotheses and answer questions about multisensory interaction and judgment.

First, we found that participants fell into *two distinct groups based on their sensory priors*: haptic and visual (3). The visual prior group was less accurate in selecting the haptically stiffer spring (H1). This decrease in accuracy is likely because those participants are more influenced by the visual information than by the squeezing of the haptic device. Thus, as we modified the visual discrepancy (λ), the effect on perception depended on that modification's strength. There are no clear demographic relationships (e.g., age, sex, familiarity with haptic devices) to predict who would fall into either of the prior groups. This sample size is small, especially when their interactions categorize them into two different groups. Future work should expand on these interaction paradigms with larger participant pools to ensure the same trends hold.

For changes resulting from visual manipulation, we find that those with a visual prior had decreased perceptual accuracy with more visual change, i.e., increased λ values (H2a). We did not find evidence for the haptic prior individuals to support our hypothesis (H2b) that they would be indifferent to λ values. Our results indicate that the *haptic prior group got more accurate as the visual discrepancy increased*, meaning they used the visual information in their decision-making. This result could be because haptic feedback is directly related to the speed of interaction and is even more apparent in cases with visual discrepancies. Many participants wrote in the qualitative feedback that they tried to interact with both springs similarly – therefore, when $\lambda = 1$, the stiffer spring would more quickly deliver squeeze cues to the user than the less stiff spring would.

Similar to our prior results [8], we found that spring stiffness would affect the visual prior group. Specifically, we found that in the visual prior group, a larger difference in spring stiffness ($K_\Delta = 0.6$) resulted in lower haptic accuracy than a smaller difference ($K_\Delta = 0.3$), which supports H3. Even when the haptic device squeezed those participants with larger forces, they doubled down on their visual prior. In the active condition, these results could also have been due to proprioception, as the larger K_Δ in combination with a $\lambda > 0$ results in larger hand re-targeting. However, given that these effects persist in the passive condition, proprioception is a less likely underlying cause. These results also align well with other prior work [7] in terms of both the grouping of participants based upon sensory priors and the effect of visual discrepancy on haptic accuracy.

We expanded beyond the scope of prior work to introduce the dual conditions of active and passive touch. Our central hypothesis about active versus passive feedback concerned the case without any visual manipulation ($\lambda = 0$), where we expected haptic perception (and thus accuracy as defined here) to decrease in the active compared to the passive condition. We found results that supported H4, highlighted by Fig. 4.

Our exploratory analysis considered the complete data set and its relationship to the active and passive conditions. Notably, we found an effect of condition across all the data – accuracy decreased in the passive condition compared to the active one. *There was also an interaction between condition and prior*. This is highlighted in the differences between colors and shapes (Fig. 3). For the haptic prior group, there is a slight increase in accuracy in the passive condition as λ increases. Similarly, there is an increase in accuracy for the visual prior group in the passive condition. Thus, the haptic prior group is worse at determining which spring is haptically stiffer in the passive case. Still, the visual information helps them align with their goal of picking the haptically stiffer spring. This may indicate that, for some participants, exaggerated visual scaling provides a compensatory cue, potentially leveraging memory from prior active interactions or reinforcing confidence in haptic judgments despite visual mismatch. This effect could also be related to the cognitive strategy of using consistent tactile cues as anchors in uncertain multisensory contexts.

Finally, we explored metrics beyond accuracy: interaction time and motion deviation. This additional data can give us insights into how participants make their perceptual decisions. In the active condition, participants with a visual prior spent more time touching the springs when the spring stiffness was greater (Fig. 5). This could be because the visual discrepancy was changing, resulting in larger proprioceptive movements to compress the haptically less stiff spring. It also could be a sign of taking more time to think and process the mismatch in the information they were receiving between the visual and haptic channels during these trials. The haptic group maintained similar interaction ties across the study, possibly because those participants could easily judge the squeezing forces at the wrist and ignored any visual cues.

The results from the motion deviation are slight, although apparent. When considering the total amount of wrist movement that occurred while the proxy hand compressed the spring, we find that the visual group again deviates in strategy between the two values of K_Δ. They move their wrist more during the larger K_Δ than the smaller K_Δ (Fig. 6). Thus, it is possible that the visual prior participants were moving their right wrist to give themselves additional proprioceptive cues to match and interpret the movement of the proxy hand. However, this case also has a large variation in the haptic prior group's hand deviation. Thus, additional analysis is necessary to compare the trajectories of the proxy hand to those of the real hand to support this theory.

In our study, the parameter λ is a critical factor in understanding how humans integrate visual and haptic feedback, a core focus of this work. It directly influences the perception of stiffness and the weighting of sensory inputs, which are essential for designing more realistic haptic systems. This parameter, introduced by [7], weights the relative influence of the comparison and reference stiffnesses, captured in 4 and 3. Notably, the perceptual impact of λ is not strictly unidirectional. This asymmetry arises because λ *influences the visual scaling differently depending on whether it is applied to the reference or comparison spring*. Specifically, when $\lambda > 0$ is applied to the stiffer (comparison) spring, the virtual hand appears to move further than the real hand. In contrast, when $\lambda > 0$ is applied to the softer (reference) spring, the real hand must move further to achieve the same visual compression, effectively reducing the perceived difference in stiffness. This differential scaling can lead to asymmetrical perceptual effects, where the perceived mismatch can either enhance or diminish the participant's confidence in their haptic judgment, depending on the direction of the λ scaling.

Our findings suggest that participants with a visual prior are particularly sensitive to these shifts, as the visual discrepancy can dominate their perception, effectively overriding the haptic feedback. In contrast, the haptic prior group showed a more complex interaction with λ, where increasing visual discrepancy did not uniformly degrade performance. Given that our measure of haptic accuracy inherently reflects the directional relationship between perceived and actual stiffness, *our analysis captures potential systematic biases in participants' judgments*. This may indicate that, for some participants, exaggerated visual scaling provides a compensatory cue, enhancing their ability to differentiate stiffness despite the intended haptic focus. This nuance highlights the need for future studies to carefully consider the directional effects of λ, potentially exploring real-time adaptive scaling that dynamically adjusts to individual users' perceptual strategies.

While we studied active and passive touch at a within-subjects level by modifying the traditional between-subjects yoking method, future work should study differences between how people explore and interact with scenes. Participants in this study maintained their perceptual strategies because we replayed their exact interactions, and the task was the same across both conditions. Introducing additional changes between conditions more closely matches real-world interactions, where it would be unlikely that the passive information we receive is directly measured and replayed from a prior active interaction. However, changing the task (slightly) or modifying interactions could result in large cognitive load differences and different outcomes. The ordering in the study may also have affected people's strategies as the active condition always came before the passive condition. Future work could consider modifying the yoked study design, such as between-subjects, to ensure this does not change the results.

Future work in this area should go beyond force and squeeze to explore vibration, as it is more common across consumer devices. Humans have more experience with perceiving vibration passively through smartphones and gaming controllers. However, introducing vibration comes with its own difficulties, including matching perception levels across different types of vibrotactors and people's perceptual differences as well [41]. Additional studies should also consider the role that referred haptics play in embodiment and agency, as there is an inherent mismatch between the signal display location and the intended perceptual location.

VIII. Conclusion

In this work, we provided new insights into the multisensory perception of referred haptic feedback under both active and passive touch conditions. Our findings indicate that perceptual strategies are highly individual, with participants falling into distinct sensory prior groups that significantly affect their ability to accurately judge haptic stiffness. Importantly, the loss of agency and proprioceptive input in the passive condition generally reduced haptic accuracy, a finding with implications for designing wearable haptic devices and virtual reality systems that aim to replicate real-world tactile experiences. Our work extends prior research by integrating active and passive conditions within a unified experimental framework, providing a more comprehensive understanding of the factors influencing multisensory perception.

Moreover, this study highlights the critical role of pseudo-haptic illusions, such as C/D ratio or λ, in shaping perceptual outcomes, emphasizing the importance of careful parameter selection when designing virtual interactions. Our findings suggest that future systems should account for both the feedback devices' physical properties and the users' cognitive strategies to optimize perceptual fidelity and realism. Given the rapid advancement of haptic technologies, including those employing vibrotactile and force-feedback mechanisms, there are significant opportunities for further research into how these systems can be tuned to better match the natural multisensory processing capabilities of the human brain.

Finally, while this study focused on stiffness perception, the broader implications extend to any scenario where mismatched multisensory cues could impact user experience, including virtual reality and immersive gaming. These findings are also relevant for the design of teleoperation systems, surgical robotics, and remote tactile interfaces, where precise control over multisensory feedback is essential for accurate task performance. Future work should explore how these insights apply to other haptic modalities, such as vibration and shear stretch, and consider the role of user expertise, task context, and cognitive load in shaping multisensory perception.

References

[1] R. L. Klatzky, "Haptic perception and its relation to action," *Annual Review of Psychology*, vol. 76, 2024.

[2] J. J. Gibson, "Observations on active touch.," *Psychological review*, vol. 69, no. 6, p. 477, 1962.

[3] J. W. Moore, "What is the sense of agency and why does it matter?," *Frontiers in psychology*, vol. 7, p. 1272, 2016.

[4] C. E. Chapman, "Active versus passive touch: factors influencing the transmission of somatosensory signals to primary somatosensory cortex," *Canadian journal of physiology and pharmacology*, vol. 72, no. 5, pp. 558–570, 1994.

[5] E. Pezent, A. Macklin, J. M. Yau, N. Colonnese, and M. K. O'Malley, "Multisensory pseudo-haptics for rendering manual interactions with virtual objects," *Advanced Intelligent Systems*, p. 2200303, 2023.

[6] M. Azmandian, M. Hancock, H. Benko, E. Ofek, and A. D. Wilson, "Haptic retargeting: Dynamic repurposing of passive haptics for enhanced virtual reality experiences," in *Proceedings of the 2016 chi conference on human factors in computing systems*, pp. 1968–1979, 2016.

[7] M. A. Srinivasan, G. L. Beauregard, and D. L. Brock, "The impact of visual information on the haptic perception of stiffness in virtual environments," in *ASME International Mechanical Engineering Congress and Exposition*, vol. 15281, pp. 555–559, American Society of Mechanical Engineers, 1996.

[8] E. D. Z. Chase and M. K. O'Malley, "The interplay of vision and referred haptic feedback in vr environments," in *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications*, pp. 385–397, Springer, 2024.

[9] R. Winter, V. Harrar, M. Gozdzik, and L. R. Harris, "The relative timing of active and passive touch," *Brain Research*, vol. 1242, pp. 54–58, 2008.

[10] A. S. Schwartz, A. J. Perey, and A. Azulay, "Further analysis of active and passive touch in pattern discrimination," *Bulletin of the Psychonomic Society*, vol. 6, no. 1, pp. 7–9, 1975.

[11] V. Cronin, "Active and passive touch at four age levels.," *Developmental Psychology*, vol. 13, no. 3, p. 253, 1977.

[12] J. M. Loomis and S. J. Lederman, "What utility is there in distinguishing between active and passive touch," in *Psychonomic Society meeting*, 1984.

[13] G. H. Van Doorn, V. Dubaj, D. B. Wuillemin, B. L. Richardson, and M. A. Symmons, "Cognitive load can explain differences in active and passive touch," in *Haptics: Perception, Devices, Mobility, and Communication: International Conference, EuroHaptics 2012, Tampere, Finland, June 13-15, 2012. Proceedings, Part I*, pp. 91–102, Springer, 2012.

[14] V. Occelli, C. Spence, and M. Zampini, "Audiotactile interactions in temporal perception," *Psychonomic bulletin & review*, vol. 18, pp. 429–454, 2011.

[15] D. Brahimaj, M. Ouari, A. Kaci, F. Giraud, C. Giraud-Audine, and B. Semail, "Temporal detection threshold of audio-tactile delays with virtual button," *IEEE Transactions on Haptics*, vol. 16, no. 4, pp. 491–496, 2023.

[16] D. J. Levitin, K. MacLean, M. Mathews, L. Chu, and E. Jensen, "The perception of cross-modal simultaneity (or "the greenwich observatory problem" revisited)," in *AIP Conference Proceedings*, vol. 517, pp. 323–329, American Institute of Physics, 2000.

[17] I. M. Vogels, "Detection of temporal delays in visual-haptic interfaces," *Human Factors*, vol. 46, no. 1, pp. 118–134, 2004.

[18] M. A. Symmons, B. L. Richardson, D. B. Wuillemin, and G. H. VanDoorn, "Active versus passive touch in three dimensions," in *First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics Conference*, pp. 108–113, IEEE, 2005.

[19] L. Jones, *Haptics*. MIT press, 2018.

[20] G. Evangelou, H. Limerick, and J. Moore, "I feel it in my fingers! sense of agency with mid-air haptics," in *2021 IEEE World Haptics Conference (WHC)*, pp. 727–732, IEEE, 2021.

[21] J. Bergstrom-Lehtovirta, D. Coyle, J. Knibbe, and K. Hornbæk, "I really did that: Sense of agency with touchpad, keyboard, and on-skin interaction," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–8, 2018.

[22] S. Kasahara, J. Nishida, and P. Lopes, "Preemptive action: Accelerating human reaction using electrical muscle stimulation without compromising agency," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–15, 2019.

[23] E. A. Caspar, A. Cleeremans, and P. Haggard, "The relationship between human agency and embodiment," *Consciousness and cognition*, vol. 33, pp. 226–236, 2015.

[24] E. van Aken, "Getting a grip: On causation, agency, and the meaning of "manipulation"," *Theoria*, vol. 88, no. 6, pp. 1228–1247, 2022.

[25] E. D. Z. Chase, P. Wolff, T. Gerstenberg, and S. Follmer, "A causal feeling: How kinesthetic haptics affects causal perception," in *2021 IEEE World Haptics Conference (WHC)*, pp. 347–347, IEEE, 2021.

[26] E. D. Z. Chase, T. Gerstenberg, and S. Follmer, "Realism of visual, auditory, and haptic cues in phenomenal causality," in *2023 IEEE World Haptics Conference (WHC)*, pp. 306–312, 2023.

[27] S.-J. Blakemore, D. Wolpert, and C. Frith, "Why can't you tickle yourself?," *Neuroreport*, vol. 11, no. 11, pp. R11–R16, 2000.

[28] E. D. Chase, A. Israr, P. Preechayasomboon, S. Sykes, A. Gupta, and J. Hartcher-O'Brien, "Learning vibes: Communication bandwidth of a single wrist-worn vibrotactile actuator," in *2021 IEEE World Haptics Conference (WHC)*, pp. 421–426, IEEE, 2021.

[29] M. Matscheko, A. Ferscha, A. Riener, and M. Lehner, "Tactor placement in wrist worn wearables," in *14th IEEE International Symposium on Wearable Computers (ISWC 2010), 10-13 October 2010, Seoul, Korea*, pp. 1–8, IEEE Computer Society, 2010.

[30] 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*, pp. 3683–3693, 2016.

[31] J. Hong, A. Pradhan, J. E. Froehlich, and L. Findlater, "Evaluating wrist-based haptic feedback for non-visual target finding and path tracing on a 2d surface," in *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '17, (New York, NY, USA), p. 210–219, Association for Computing Machinery, 2017.

[32] A. Israr, S. Zhao, and O. Schneider, "Exploring embedded haptics for social networking and interactions," in *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, pp. 1899–1904, 2015.

[33] Z. A. Zook and M. K. O'Malley, "Effect of focus direction and agency on tactile perceptibility," in *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications*, pp. 119–126, Springer, 2022.

[34] M. J. Hautus, N. A. Macmillan, and C. D. Creelman, *Detection theory: A user's guide*. Routledge, 2021.

[35] 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, 2012.

[36] E. Pezent, P. Agarwal, J. Hartcher-O'Brien, N. Colonnese, and M. K. O'Malley, "Design, control, and psychophysics of tasbi: A force-controlled multimodal haptic bracelet," *IEEE Transactions on Robotics*, vol. 38, no. 5, pp. 2962–2978, 2022.

[37] B. L. Richardson, D. B. Wuillemin, and G. J. MacKintosh, "Can passive touch be better than active touch? a comparison of active and passive tactile maze learning," *British Journal of Psychology*, vol. 72, no. 3, pp. 353–362, 1981.

[38] R. M. Church, "The yoked control design," in *Aversion, Avoidance, and Anxiety*, pp. 403–415, Psychology Press, 2014.

[39] P.-C. Bürkner, "brms: An r package for bayesian multilevel models using stan," *Journal of statistical software*, vol. 80, pp. 1–28, 2017.

[40] A. Vehtari, J. Gabry, M. Magnusson, Y. Yao, P.-C. Bürkner, T. Paananen, and A. Gelman, "loo: Efficient leave-one-out cross-validation and waic for bayesian models," *CRAN: Contributed Packages*, 2015.

[41] D. H. Sullivan, E. D. Z. Chase, and M. K. O'Malley, "Comparing the perceived intensity of vibrotactile cues scaled based on inherent dynamic range," *IEEE Transactions on Haptics*, 2024.