A Faster VR Body to Speed Up Choices

Difeng Yu
Department of Computer Science
University of Copenhagen
Copenhagen, Denmark
diyu@di.ku.dk

Joanna Bergström
Department of Computer Science
University of Copenhagen
Copenhagen, Denmark
joanna@di.ku.dk

Kasper Hornbæk
Department of Computer Science
University of Copenhagen
Copenhagen, Denmark
kash@di.ku.dk

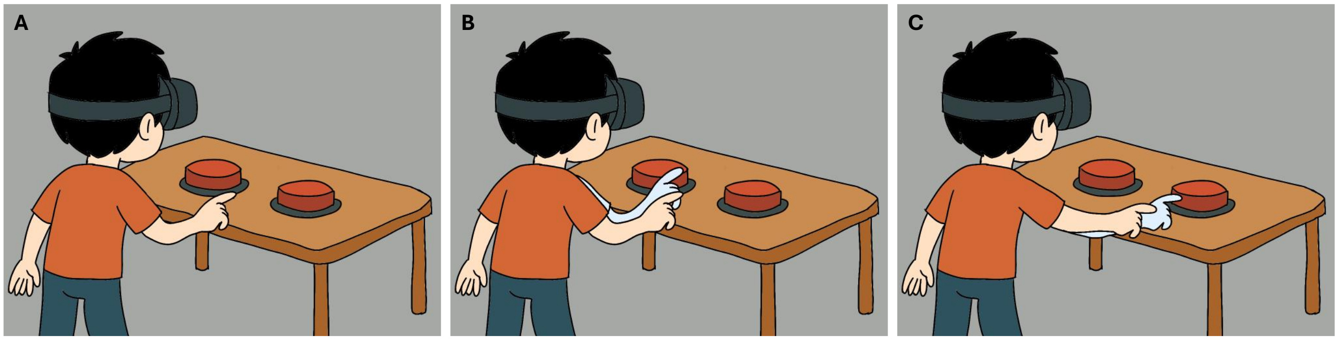

Figure 1: (A) A user is choosing between two targets. (B) A faster VR body, focusing on an accelerated hand and arm, starts moving towards one of the targets before the user has decided which target to reach for. (C) Once the user starts moving, the faster VR body quickly adjusts its movement to align with the user's physical movement, ensuring a congruent decision.

Abstract

A VR body can move faster than its user, making actions like reaching more efficient. We propose a VR body that not only moves faster during reaching, but also starts moving before the user has decided which target to reach for. However, it is unclear whether such a VR body would speed up choices, since moving towards a wrong target might cause confusion, or influence users' choices. To explore these questions, we built a faster VR body prototype, focusing on an accelerated hand and arm, in a choice task. Thirty-four participants viewed random-dot displays to judge the overall motion direction and indicated their choice by pressing the corresponding button. Task difficulty was varied to influence choice uncertainty. Results showed that choice time decreased when the virtual body was 0.1 seconds ahead of the physical body, but increased when it was 0.3 seconds ahead. Users' choice distributions showed no significant differences.

CCS Concepts

• **Human-centered computing → Virtual reality**; *Interaction techniques.*

Keywords

Virtual body, interaction design, perceptual decision making, sensorimotor control

CHI '26, Barcelona, Spain

ACM ISBN 979-8-4007-2278-3/26/04
https://doi.org/10.1145/3772318.3790543

ACM Reference Format:
Difeng Yu, Joanna Bergström, and Kasper Hornbæk. 2026. A Faster VR Body to Speed Up Choices. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3772318.3790543

1 Introduction

A virtual reality (VR) body can move faster than its user. The most common way to achieve this is increasing the control-display gain. The control-display gain defines the ratio between displayed action (e.g., VR body movement) and user action (e.g., physical body movement), where a higher gain causes the VR body to move faster than the physical body. Existing techniques have scaled up this gain to improve efficiency (e.g., quicker task completion) in interaction tasks such as target reaching [12, 54–56].

In addition to VR bodies that move faster in response to user actions, we propose a VR body that moves faster in time and could make "faster decisions" than users. More concretely, when a user is deciding on which target to choose, the VR body could start moving towards one of the targets. The goal is to help users make their decision response (by selecting the target) even more quickly. However, there are two initial questions regarding this proposal.

First, will such a VR body speed up task completion? Given that a system most likely can only "guess" which target the user intends to select before they act, moving towards a wrong target might cause confusion, thus delaying task completion [61]. Second, will such a VR body change users' choices? By moving closer to a target, VR bodies can reduce the user's perceived effort to reach it and potentially bias their choices [17]. Moreover, users might inattentively "follow" virtual body movements [15], influencing their choices.

To explore these questions, we built a prototype called *VR Luke*, named after the main character in the comic series *Lucky Luke*, who shoots faster than his shadow. *VR Luke* focuses on speeding up the virtual hand and arm movements of a VR body by following three key mechanisms. First, it maintains a temporal offset (τ) between the virtual body and the user's physical body, so that the virtual body is at the location where the physical body will typically be τ seconds later. Second, the virtual body starts moving to one of the target τ seconds before the average human reaction time. Third, once the user starts moving, the virtual body quickly shifts to the moving direction of the user's physical body, ensuring a congruent decision (see Figure 1). Notably, *VR Luke* is different from control-display gain approaches, which only respond to user actions.

We deployed *VR Luke* in a user study (N=34) that involved a choice task based on random-dot displays [25, 46, 47]. In the task, participants were asked to observe a set of moving dots, some moving randomly, others moving systematically towards either left or right. They indicated their choice of the overall direction by moving their hands to press the left or right button. With this task paradigm, users' uncertainties in their choices were manipulated by the proportion of random vs. systematic dots. *VR Luke*, in this case, moved the virtual hand towards one of the buttons, which might be correct or incorrect, before an initial user decision.

The results show that choice completion was faster when the virtual body was 0.1 seconds ahead of the user's physical body, regardless of whether it moved toward the correct or incorrect target. However, the choice completion was slower when the virtual body attempted to be 0.3 seconds ahead. Furthermore, our analyses did not find a significant impact of *VR Luke* on users' choice distribution, indicating that such VR bodies could speed up choice completion without having a significant impact on user choices. We discuss how this may open up the potential of using a faster VR body to assist and accelerate everyday decision-making.

2 Background and Related Work

In this section, we begin by differentiating two types of faster VR bodies: spatially faster and temporally faster. We then introduce the difference between decision-making and action within a choice task. Lastly, we introduce how a faster VR body may influence and be influenced by other factors in interaction.

2.1 Faster VR Bodies

A VR body can move faster than its user. We here distinguish two kinds of "faster" bodies: one in the *spatial* sense, the other in the *temporal* sense.

A spatially faster VR body works by amplifying the virtual body's movements compared to a user's physical movement, essentially increasing the control-display gain. This gain can be increased by extending the position or velocity of the virtual body. For instance, Go-Go [15] scales up the virtual hand position using a quadratic function, while Scaled HOMER [56] adjusts a velocity-dependent scaling factor so that the virtual hand's speed can be at most 1.2 times that of the physical hand. Similar techniques have been applied in object reaching, selection, and manipulation [12, 33]. Notably, increasing the gain does not always improve user performance (e.g., efficiency), as the control can become too sensitive for interaction [5]. Additional techniques [12, 55], such as slowing down virtual movement to enhance precision, need to be combined with such spatially faster VR bodies to fully leverage their potential.

A temporally faster VR body works by presenting the virtual body's movements t seconds ahead of a user. This can be conceptualized as *negative latency* in the system: whereas positive latency introduces delays between input and output, negative latency displays output based on the anticipated future inputs [19, 34]. While such a faster body is not commonly used for VR interaction, Kasahara et al. [22] pioneered a system that extrapolates movements of a physical body to predict the virtual body movements 25-100 ms in the future. This led users to perceive their virtual bodies as being lighter. Notably, temporally faster VR bodies are predicting future movements, while spatially faster VR bodies neither require nor intend to do so.

In this paper, we refer to a faster VR body as being temporally faster, with which we predict a user's future movements in a controlled two-button choice task. One key difference of our temporally faster VR body, as compared to the previous ones, is that such a prediction can not only happen during user actions, but also during their initial decision-making. Therefore, our VR body can start moving before a user has decided which target to reach for, which is a novel capability.

2.2 Decision-Making vs. Action

We distinguish two main processes when selecting alternative targets in a choice task. One is decision-making, which is about determining a target (i.e., "*what* is the object of interest?"). The other involves translating a decision into action which includes motor planning (i.e., "*how* will the target be attained?") and execution (e.g., performing a target reaching movement) [27, 57, 59].

One view is that decision-making happens only before an action [57]. The most prominent models of such a view, including race models [50] and drift-diffusion models [32, 44], consider decision-making as accumulating evidence over time to a decision threshold. Once a decision threshold is met, an action then begins.

However, previous findings suggest that people can revise their decisions during ongoing actions, when the actions are not highly ballistic (e.g., button presses or saccadic eye movements) [46, 57]. In other words, in scenarios where actions take time to complete (e.g., reaching), users can adjust their decisions as they move. To account for such "change of mind", previous research has extended the drift-diffusion model to include an initial decision phase along with continuous evidence updating during movements [46]. Other views, such as the affordance competition hypothesis [6], propose that decision-making and action run in parallel, with competing actions encoded simultaneously before one is selected [9, 36, 58].

We propose a VR body that is faster not only during actions but also during (initial) decision-making. Because such a VR body modified the virtual representation of a user during these processes, it may or may not influence users' choices.

2.3 Making Choices with A Faster VR Body

Decision-making happens when users need to choose one target from a set of alternatives [10]. In many cases, users may not be completely certain about the correct choice and must instead select the target that seems most likely based on the available evidence. To study and manipulate the level of uncertainty in such cases, previous research has developed several experimental paradigms, including random-dot displays [47] (detailed later). Our empirical study considered how a faster VR body may influence user choices under different levels of choice uncertainties.

Choice-making under uncertainties can be influenced by many factors [10, 42, 57]. We identify two causes that may bias a user's choice with the proposed faster VR body. The first is *movement effort* [35, 37, 49]. For instance, previous research shows that decision-making under uncertainty is biased away from choice that require, or are perceived to require, more movement effort [17]. Since our faster VR body moves towards one of the targets first, users may perceive that reaching it requires less effort. The second cause is the *self-avatar follower effect.* When introducing an offset between the virtual and the physical body, as with our faster VR body, users could automatically act to reduce that offset [4, 15]. Because a faster VR body may reduce perceived movement effort and induce a self-avatar follower effect, our study explores whether it will influence users' choices.

An important consideration when manipulating the VR body in choice tasks is preserving the user's sense of agency in making the choice, the feeling of "I made that choice" [16]. Previous research has explored *computer-driven touch*, where they used electrical muscle stimulation (EMS) to actuate the user's to select one of the choice involuntarily [23, 51]. The findings suggest that when the EMS's selection was misaligned with the user's choice, the user's sense of agency dropped significantly. One key difference in our study is that, while the faster VR body may initially make either an aligned or misaligned movement, it quickly drifts toward the user's chosen direction when they start moving. Further, the user has full agency in making the choice by pressing down a virtual button.

3 Study Framework

We built a prototype called *VR Luke* to understand (1) whether a faster VR body will speed up choices that involve reaching actions to complete, and (2) whether such a VR body will change users' choices. The prototype was tested in a choice task that involved a random-dot motion display. This section describes the prototype and the choice task.

3.1 VR Luke

VR Luke is a prototype that introduces a VR body that makes faster choices than its users. It is named after the main character in a comic series called *Lucky Luke*, known as the gunslinger who shoots faster than his own shadow. Analogously, the core idea of *VR Luke* is to move the virtual body, specifically the virtual hand and arm, ahead of a user's physical body in time. To achieve this purpose, the prototype relies on three key mechanisms. First, *VR Luke* tries to maintain a temporal offset ahead of a user physical body from when a task begins. Second, it forces the virtual body to move towards one of the target before the user moves. Third, once the user starts moving, the virtual body quickly shifts to the moving direction of the user's physical hand, to avoid misalignment in the "choices" between the virtual and the physical body.

3.1.1 Mechanism 1: temporal offset. Temporal offset describes a time parameter τ, meaning the virtual body is τ seconds ahead of the user's physical body. Therefore, we need to predict where the physical body might be in a short future. To achieve this, *VR Luke* maintains the temporal offset in a typical movement profile (see Figure 2A).

In a choice task, a user first makes an initial decision about which target to reach for. They then move towards a target, so that the distance between their hand and the starting point (i.e., the movement distance) increases. Their movement slows down as the hand approaches the target, and they instantiate their choice by pressing a button. To make this process computable, we parameterize it using values derived based on pilot tests. We assume the initial decision time is approximately 300 ms. The user's movement is modeled as a minimum-jerk trajectory [30], taking 1 second to complete the 0.4-meter target distance in the task (describe later). With these parameters, we can determine the movement distance of the physical hand given any timestamp t in the movement profile, and conversely, infer t based on a movement distance.

Based on this typical movement profile, we can relocate the virtual body (specifically, the virtual hand and arm) τ seconds ahead of the user by adjusting the corresponding spatial offsets between the virtual and physical body[1]. Figure 2B-C demonstrate two examples. When the user is in the initial decision phase (i.e., has not moved), the virtual hand will start moving automatically. In this case, the maximum movement distance for the virtual hand is constrained to the distance corresponding to 300 ms (the initial decision time) + τ. Once the user starts moving, the virtual hand maintains the temporal offset τ based on the movement distance of the physical hand. If the user moves before the virtual hand reaches the distance of 300 ms + τ (i.e., the initial decision time is smaller than 300 ms), the virtual hand will snap to a corresponding position that maintains the temporal offset. Therefore, the virtual body always tries to be τ seconds ahead of the physical body.

3.1.2 Mechanism 2: forced initial decision. Temporal offset (mechanism 1) only considers movement distances, not movement directions. This and the next mechanism determine the movement direction of the virtual hand. When the user is in the initial decision phase, the virtual body has to "decide" which target to reach for. In this prototype, the faster virtual hand moves straightly towards one of the targets. The probability of moving towards each target is controlled evenly in our empirical study.

3.1.3 Mechanism 3: initial decision correction. Once the user starts moving, a mismatch between the virtual body's initial decision and the user's initial decision will result in movement incongruencies. To resolve this, we use the following equation to weight the virtual hand's initial direction (D_i) and the physical hand's direction (D_p) differently over a correction phase to compute a new virtual hand direction (D_v).

$$D_v = (1 - p) \cdot D_i + p \cdot D_p \qquad (1)$$

Here, p represents the progress within the correction phase, where $p \in [0, 1]$. This ensures the smooth transition from the forced initial movement direction of the faster VR body to the user's actual

---

[1] Interested readers can find a more precise definition of spatial and temporal offsets in Appendix A.

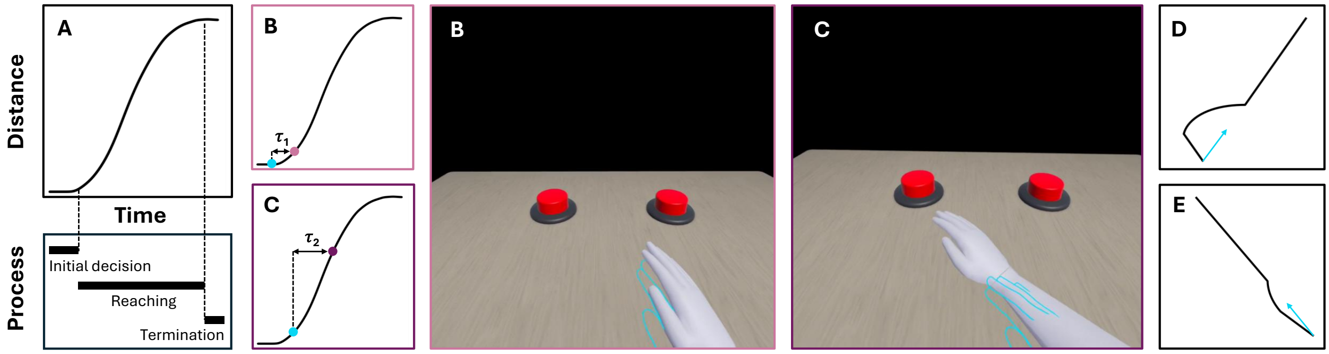

Figure 2: (A) A typical movement profile in choice tasks: as time progresses, the user's hand movement distance initially remains stable (reflecting the initial decision phase), then increases during the reaching movement, and finally stabilizes again when the user is instantiating the choice (e.g., via button pressing). (B-C) Two different temporal offsets ($\tau_2 > \tau_1$). The blue dots indicate the position of the physical hand (outlined in blue), while the pink and purple dots indicate the position of the virtual hand (shown in white). (D-E) Two top-down views on how the trajectory of the virtual hand (the black line) can be influenced by the movement direction of the physical hand (the blue arrow line), when the initial decision is either largely misaligned (D) or roughly aligned (E).

movement direction. Figure 2D-E show two top-down views on how the virtual hand position changes over task progress, when the initial decision is either largely misaligned or roughly aligned. The correction phase was confined to the initial 20% of the physical hand movement distance to reduce disruption [8].

3.1.4 VR body implementation. The VR body implementation is based on the MetaAvatarSDK (v24.0.1), using a plain white avatar. Arm posture is computed through inverse kinematics.

3.2 The Choice Task

The choice task employed a random-dot display [25, 46, 47]. Unlike other choice task paradigms such as a flanker task [40] or a Stroop test [51], where choice uncertainty cannot be easily manipulated, random-dot displays have been used as one of the standard stimulus types for introducing uncertainties in perceptual decision-making. The general idea is that a set of dots can move systematically towards one of the target directions (left/right) or randomly within the display. The percentage of systematic vs. random, also called *motion coherence*, can be manipulated. A participant watches the display for a short time and determines the general dot-moving direction (left/right) by pressing a button. Intuitively, lower motion coherence (i.e., more randomly moved dots) makes this judgment more difficult. In the following, we first overview the task procedure (see Figure 3) and then detail the implementation of the random-dot display.

3.2.1 Procedure. Before starting a task trial, a user sees a white button positioned 0.3 m in front of them and a green dot which indicates the center of a random-dot display. After pressing the button, the random dot display (detailed in the following section) appears for 400 ms, where users need to judge the overall movement directions of the random dots. Once the display extinguishes, two red buttons (radius = 0.05 m) appear 0.4 m from the (now invisible) white button, positioned 30° to the left and right. The user then indicates their choice of the overall movement direction of the dots (left or right) by pressing the corresponding button. Note manipulations of *VR Luke* occur only in the horizontal plane and are therefore independent of the user's choice (pressing a button vertically). In other words, users maintain full control over their choices.

3.2.2 Random-dot display. There are different ways to place and move the random dots. We followed a previous work where the dots performed Brownian Motion (BM) [41]. In their experiment, BM was the most effective and resulted in the least amount of directional ambiguity compared to other implementations.

More specifically, for each frame, a new set of dots is randomly picked to move towards a target direction, while the rest of the (noise) dots move towards a random direction. All dots have the same size of 0.16° and move at the same speed of 12°/s. In total, there are 50 dots within an invisible circle with 18° diameter, which makes the dot density as 16.7 dots $\deg^{-2}s^{-1}$. Note that we have used the angular representation to describe the size, where we report the subtended visual angle if placed in the forward-looking direction. These parameters follow a previous implementation [41], except that the dots were twice as wide to be easily visible in a VR headset. The center of the display was place 0.2 m above a virtual table. The users can adjust the height of the virtual table during calibration (detail later).

4 Empirical Study

We conducted an empirical study with 34 participants[2] to test *VR Luke* on the choice task based on the random-dot display. We then analyzed the collected data to understand the effect of a faster VR body on choice time, choice distribution, and movement characteristics.

---

[2] Sample size was determined based on a power analysis of paired t-test (i.e., with or without *VR Luke* in a within-subject design) using the following settings: Tails = Two, effect size $d_z = 0.5$ (medium size), $\alpha = 0.05$, and power = 0.8.

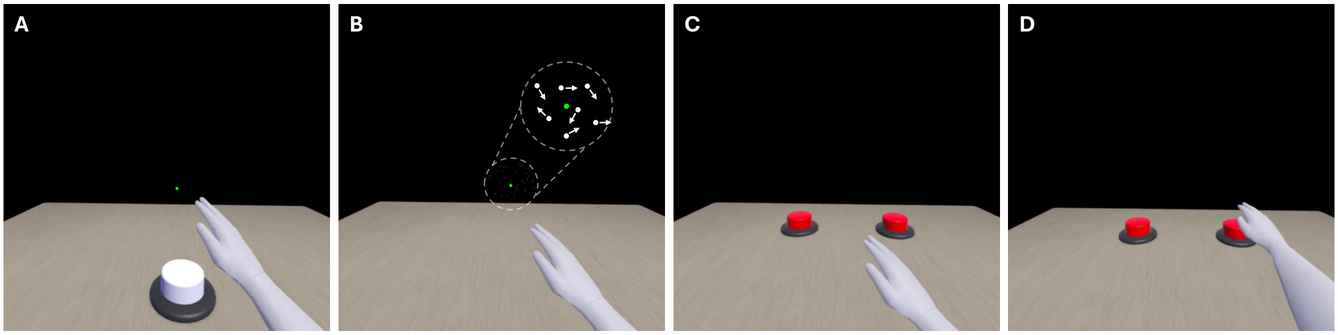

Figure 3: (A) A user presses a white button to start a task trial. (B) After that, the white button disappears, and a random-dot display is presented for 400 ms. (C) Once the display disappears, two red buttons appear, each the same distance from the white starting button, which is now invisible. (D) The user indicates the overall movement direction of the dots by pressing the left or right button.

4.1 Study Design

The study follow a within-subject design. The main variables are (1 Baseline + 2 Temporal Offset × 2 Offset Correctness) × 3 Motion Coherence.

- Baseline is a 1:1 physical-virtual mapping of users' hand movements, where *VR Luke* does not apply any manipulation. This is in contrast to Luke Conditions, where *VR Luke* applies a VR body manipulation. Baseline and Luke Conditions are combined into Body Conditions.
- Temporal Offset suggests how much "faster" the virtual body is compared to the physical body. We set two levels in Luke Conditions: 0.1 seconds and 0.3 seconds, meaning the VR body begins moving 0.1 s or 0.3 s before a user's typical reaction time (300 ms). The 0.1 s level represents a relatively small offset: once the user starts moving, the spatial distance ranges from a minimum of 0.003 meters to a maximum of 0.075 meters, given a target distance of 0.4 meters. It should maintain a higher level of avatar embodiment and thus may lead to the self-avatar follower effect [1, 15]. The 0.3 s level represents a larger offset (spatially from 0.065 m to 0.212 m), where participants could observe that the distance to one target is shortened with the faster VR body, potentially leading them to minimize movement effort [17]. The Baseline condition has zero temporal offset.
- Offset Correctness are only relevant to Luke Conditions. It has two levels: the VR body can either move to the correct or the incorrect button at the beginning of its movement.
- Motion coherence represents the proportion of the dots moving to a systematic direction (left or right) within the random-dot display. As motion coherence increases, the task becomes easier, participants' accuracy improves, and their confidence in their choices generally increases [25, 41]. We set three levels: 2%, 6%, and 20%. According to our pilot tests, users performed slightly better than random guessing in the 2% condition, were almost always correct in the 20% condition, and showed intermediate performance (around 75% accuracy) in the 6% condition.

Each experimental condition was repeated 8 times: 4 with the target on the left and 4 on the right (Target Direction). This resulted in a total of 120 trials per participant. The trials were presented in a randomized order. For each trial, we recorded performance measures, including choice time (i.e., the time interval from target appearance to button press), the correctness of a user's choice, and movement trajectory (at 50 Hz).

4.2 Procedure

Participants first accepted the task in Prolific. They were then redirected to a information sheet and downloaded the VR application through SideQuest. They used their own VR device (Oculus Quest 2, 3, or Pro) to complete the task. Participants first saw a welcome screen, and then completed a calibration phrase to ensure the virtual table (where the buttons were placed) was comfortably positioned. After that, they were introduced the task and practiced the Baseline condition for 10 trials where Motion Coherence of the random dot display decreased from 80% to 30%, to help participants to get used to the task. A short sound clip that indicated the correctness of the choice was played during this practicing phase to help with learning. However, the sound clips are not played during the formal experiment. Participants could choose to go through the practice phase again if they wanted. After that, they entered the formal experiment. They were instructed to complete the task as fast and as accurately as possible and were told that their bonus payments would be scaled based on performance. After completing the task, a redeem code was displayed in VR for them to retrieve their compensation (£6). Bonus payments (up to an additional £6) were awarded to participants with high accuracy and reasonable speed after data collection was complete. A full walkthrough of the study procedure in VR can be found in the video supplementary material.

4.3 Participants

We recruited 36 online participants from Prolific. The prescreening criteria were an approval rate above 98% and ownership of a VR headset. Two participants were removed from the analysis because their data were incomplete or their overall accuracy was worse than a random guess (50%). Of the remaining 34 participants, 5 were female, 28 were male, and 1 preferred not to say. The age ranged from 21 to 67 (*mean* = 34.0, *std* = 13.1). Among the 34 participants, 25 were right-handed, 7 were left-handed, and 2 were mixed-handed. Their overall accuracy in the task was at least 60%. We did not detect any abnormal hand-trajectory behavior (e.g., moving towards targets in a non-human manner).

4.4 Analysis and Results

We collected the data from the 34 participants, resulting in 4080 trials in total. We first removed 74 trials (1.8%) with choice time more than three standard deviations above the mean, leaving us with 4006 trials. For the remaining trials, the average accuracies at the three Motion Coherence levels were 58.3%, 71.1%, and 94.1%, from the most difficult to the easiest, which aligned with our expectations of the task difficulties. Their average choice time for a trial was 1.34 s, 1.32 s, and 1.22 s, correspondingly.

In the following, we present a set of (generalized) linear mixed models to explore a set of research questions (**Q**). They can be summarized as

- How would temporal offset (**Q1**), and further the correctness of the initial offset (**Q2**), affect choice time?
- How would an initially correct or incorrect temporal offset affect users' choices? (**Q3**) Users' choices may be influenced by their tendency to minimize movement effort or by the self-avatar follower effect.
- How would temporal offset affect when users initiate a movement (**Q4**) and their action characteristics, including movement distance (**Q5**) and movement smoothness (**Q6**)? These questions help us to understand why temporal offsets may or may not be effective in speeding up choices.

For readability, we present only the key analyses and results; detailed statistics are provided in the supplementary material.

4.4.1 Choice time. Our first question (**Q1**) was: *Did temporal offsets speed up or slow down choice completion?* To answer this, we first took a log transformation of the choice time data, which followed log-normal distribution. We then built the following linear mixed model:

$$\log(\textsc{Time}) \sim \textsc{Motion Coherence} \times \textsc{Temporal Offset} + (1|\textsc{Target Direction}) + (1|\textsc{Participant}) \tag{2}$$

In this model, Motion Coherence and Temporal Offset were fixed effects, with a potential interaction. Target Direction and Participant were treated as random effects. The model showed that, relative to Baseline, a Temporal Offset of 0.1 s was associated with significantly lower log-transformed choice time ($p <$ .001), while a Temporal Offset of 0.3 s was associated with higher choice time ($p <$ .001). Post-hoc pairwise comparisons confirmed these differences and indicated small to medium effect sizes (both $p <$ .001; $d = -0.358$ for 0.1 s *vs.* Baseline; $d = 0.220$ for 0.3 s *vs.* Baseline).[3] Results were consistent across different levels of Motion Coherence (all $p <$ .033). Overall, they indicated that when the VR body was 0.1 s ahead of the user's physical body, choice completion was faster, whereas a 0.3 s offset slowed it down, across all the motion coherence levels.

Our follow-up question (**Q2**) was then: *Did the correctness of the initial offset direction influence choice completion?* The assumption was that an incorrect initial offset direction might slow down choice completion, while a correct one might speed up choice completion. We then built a linear mixed model:

$$\log(\textsc{Time}) \sim \textsc{Motion Coherence} \times \textsc{Body Condition} + (1|\textsc{Target Direction}) + (1|\textsc{Participant}) \tag{3}$$

The model indicated that, compared to Baseline, all levels of Luke Condition were associated with significantly different log-transformed choice time (all $p <$ .001), except for *IncorrectLuke*_0.3 (incorrect initial offset with Temporal Offset = 0.3 s, p = .241). A post-hoc test showed that *CorrectLuke*_0.1 and *IncorrectLuke*_0.1 had significantly lower log-transformed choice time than Baseline (both $p <$ .001), while *CorrectLuke*_0.3 and *IncorrectLuke*_0.3 had significantly higher log-transformed choice time (both $p <$ .003). There was no significant difference between *CorrectLuke*_0.1 and *IncorrectLuke*_0.1 (p = .928) or between *CorrectLuke*_0.3 and *IncorrectLuke*_0.3 (p = .539). The conclusion was that a 0.1 s temporal offset sped up choice completion, while a 0.3 s offset slowed it down, regardless of whether the initial virtual hand movement was towards a correct or incorrect target.

4.4.2 Choice distribution. Our question (**Q3**) was: *Did VR Luke significantly influenced users' choices?* As *VR Luke* reduced the users' perceived movement effort and could induce the self-follower effect, we hypothesized that users' choices would be influenced significantly. Because we could not determine users' initial choices for each trial, as their decisions might have been influenced by *VR Luke* moving before them, we analyzed the aggregated choice distributions across all trials. We used the following generalized linear mixed model, assuming the response variable followed a binomial distribution, which was modeled with logistic regression:

$$\textsc{Correctness} \sim \textsc{Motion Coherence} \times \textsc{Body Condition} + (1|\textsc{Target Direction}) + (1|\textsc{Participant}) \tag{4}$$

The binary response variable Correctness indicated whether the user chose the correct or incorrect target for each trial. If users were more likely to follow *VR Luke*, a correct initial offset (i.e., *CorrectLuke*_0.1 and *CorrectLuke*_0.3) should systematically increase Correctness, while an incorrect initial offset (including *IncorrectLuke*_0.1 and *IncorrectLuke*_0.3) should systematically decrease Correctness, compared to Baseline. However, the model and a following post-hoc test under each Motion Coherence level showed no significant different between all Luke Condition and Baseline (all $p >$ .506). Thus, *VR Luke* did not seem to significantly influence choice distribution.

4.4.3 Movement initiation. Since temporal offsets could influence how quickly tasks were completed, it might have affected initiation time, defined as the interval between the button's appearance and the initiation of a movement. Movement initiation was approximated as the point when the hand had moved 10% of the target distance (i.e., 0.04 m) since the target appeared.[4] In this case, the averaged initiation time was 0.348 s (median = 0.308 s). The question (**Q4**) was: "*Did temporal offsets speed up or slow down movement initiation?*" We used the following linear mixed model where the

---

[3] Cohen's d was calculated on the transformed data; the same applies below.

[4] When discussing movement characteristics, including position, speed, and jerk, we only considered movements relevant to the target direction (e.g., vertical movement during button pressing was not included). Same for below.

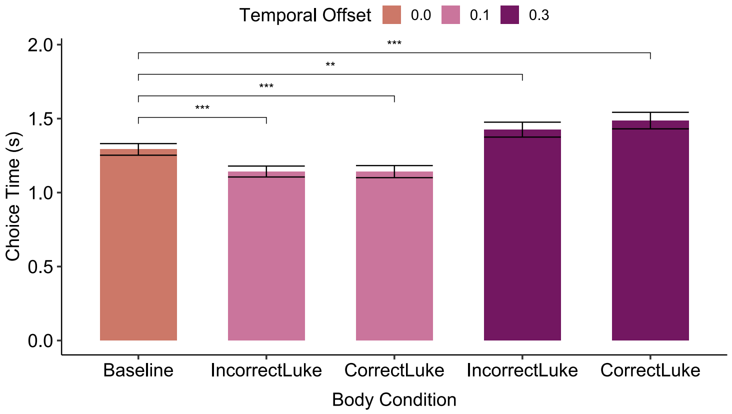

Figure 4: Bar plots of choice time for each Body Condition (Baseline + 2 Offset Correctness × 2 Temporal Offset). Error bars represent the 95% confidence intervals. Statistical significant difference between Baseline and other conditions are marked (** = p < .01 and *** = p < .001).

initiation time (InitT) was transformed through Yeo-Johnson transformation ($\lambda = -2.766$).[5]

$$\text{norm}(\textsc{InitT}) \sim \textsc{Motion Coherence} \times \textsc{Temporal Offset} + (1|\textsc{Target Direction}) + (1|\textsc{Participant}) \tag{5}$$

The model indicated significant interaction effects of Temporal Offset 0.1 : Motion Coherence 0.06 ($p = .020$) and Temporal Offset 0.3 : Motion Coherence 0.06 ($p = .003$). A post-hoc test showed that when Motion Coherence = 0.06, Temporal Offset 0.3 s resulted in significantly shorter normalized initiation time than Baseline ($p = .019$, $d = -0.192$), while Temporal Offset 0.1 s did not reach significance ($p = .098$, $d = -0.145$). It was not the case for any other level of Motion Coherence (all $p > .186$, and the d values were between 0.090 and 0.123). These results suggested that temporal offsets shortened the time for users to initiate a movement, when Motion Coherence = 0.06.

4.4.4 Action characteristics. Because temporal offsets can either accelerate or delay choice completion depending on their magnitude (0.1 s *vs.* 0.3 s), we investigated the underlying causes by analyzing movement distance and smoothness. The questions were "*Did temporal offset shortened movement distance as intended?*" (**Q5**) and "*Did temporal offset influenced movement smoothness, possibly due to unexpected VR body behavior?*" (**Q6**). Movement distance was computed by aggregating the frame-by-frame distance for each trial. Movement smoothness was quantified through normalized jerk [20, 52], a dimensionless measure of changes in movement acceleration, following a previous implementation [60]. We applied Yeo-Johnson transformation ($\lambda = -3.798$) to the movement distance data and Box-Cox transformation ($\lambda = -0.173$) to the normalized jerk data. We applied linear mixed models in the same form as Equation 5, changing only the response variable.

Based on the model results, we found both Temporal Offset 0.1 s and Temporal Offset 0.3 s were associated with significant less movement distance than Baseline (both $p < .001$). These results were confirmed by a post-hoc test showing small to medium effect sizes (both $p < .001$; $d = -0.515$ for 0.1 s *vs.* Baseline; $d = -0.348$ for 0.3 s *vs.* Baseline). The results were consistent across different levels of Motion Coherence (all $p < 0.002$). Moreover, Temporal Offset 0.1 s led to significantly shorter movement distance than Temporal Offset 0.3 s ($p < .001$, $d = -0.167$).

The model results also showed that Temporal Offset 0.1 s was associated with significantly lower normalized jerk (i.e., smoother movements) compared to Baseline, whereas Temporal Offset 0.3 s was associated with significantly higher normalized jerk than Baseline (both $p < .001$). Post-hoc pairwise tests indicated the same results with small to large effect sizes (all $p < .001$; $d = -0.278$ for 0.1 s *vs.* Baseline; $d = 0.244$ for 0.3 s *vs.* Baseline; $d = -0.605$ for 0.1 s *vs.* 0.3 s), similarly across Motion Coherence (all $p < .002$).

Overall, the results suggest that both temporal offsets (0.1 s and 0.3 s) reduced movement distance, but a 0.3 s offset may decrease movement smoothness, likely due to the unexpected VR body movements caused by the large discrepancies between the virtual and physical bodies.

5 Discussion

In this section, we discuss implications of applying a faster VR body in choice tasks. We also present variations of the *VR Luke* implementation and future research directions.

5.1 A Faster VR Body

We found that making the VR body 0.1 s faster improved choice completion time, but making it 0.3 s faster slowed down choice completion time, compared to the baseline. The results were consistent across the three coherence levels (i.e., choice uncertainty) and were independent of whether the virtual body initially moved towards the correct or incorrect target. The improvement with the 0.1 s faster VR body was expected, as it significantly reduced the movement distance compared to the normal VR body, leading to an average choice completion time that was 0.14 s faster (i.e., around

---

[5] The normalization technique was determined by the `BestNormalize` package in R. Same for the transformations below.

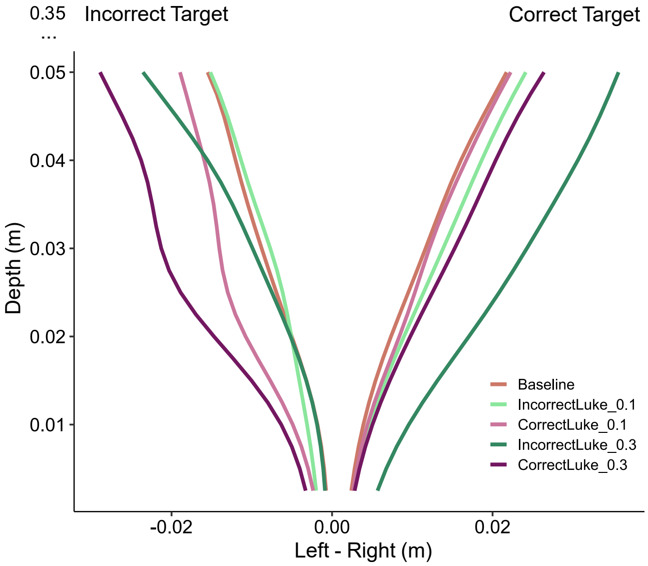

Figure 5: The averaged trajectory plot of Body Condition at the onset of movements. The horizontal axis represents left-right movement: right corresponds to the correct target, and left to the incorrect target. For trials where the actual target was on the left, the movement trajectory was mirrored to maintain consistency. The vertical axis represents depth (movement towards the target), with the target positioned at around 0.35 m. Notably, when participants chose the correct target (i.e., all the trajectories moving towards the right), their trajectory in *IncorrectLuke_0.3* differed clearly from other conditions. A similar pattern was observed when participants chose the incorrect target in *CorrectLuke_0.3*. This suggests that participants resisted following *VR Luke* when it conflicted with their choices.

11% faster). It seemed that, because the 0.1 s faster VR body began with a small spatial offset (see Figure 2B for a visual approximation), participants could easily correct directional errors when it moved towards the wrong target.

However, we did not expect that the 0.3 s faster VR body would slow down task completion, especially when it was going towards a correct target. Note that the temporal offset does not make the control more sensitive as with spatial offset techniques [5, 55, 56], as the VR body would naturally slow down when approaching a target. The results showed that although the 0.3 s faster VR body reduced movement distance, it significantly increased the physical hand's movement jerk. This could indicate that participants were "surprised", leading to abrupt changes in acceleration (i.e., higher movement jerk). From our results, it remains to be seen what is the optimal temporal offset that leads to the shortest choice completion time, and where the effect breaks down between 0.1 s and 0.3 s.

It would also be interesting to see whether users could adapt to a VR body that was consistently 0.3 s faster with systematic practice. Note that we have randomized the order of temporal offsets in our study to discourage participants from preparing explicit strategies for specific offsets (e.g., intentionally waiting for the VR body to finish moving and then beginning to move physically). However, users may learn to adjust their movements optimally once the VR body behavior becomes more consistent, similar to how they learn to use spatially faster VR bodies [60].

One mysterious finding is that when the coherence level is 6% (which resulted in an averaged accuracy of 71.1%), participants initiated their movement earlier with temporal offsets than without. This could mean that they took less time for making the initial decision, and somehow this effect was not present in the other two coherence levels. We noticed that the movement initiation time for 6% coherence level was the longest in NoLuke (on average, 2%: 372 ms, 6%: 385 ms, and 20%: 294 ms). Our suspicion was that the VR body movement could have been considered as additional evidence or an urgency signal for decision-making, which accelerated its convergence towards the decision boundary [7, 48, 63]. Future work might want to replicate this effect, as it may be due to chance, and investigate its underlying mechanisms.

5.2 Influence on Choices

We found that speeding up the VR body, whether it was 0.1 s or 0.3 s faster, did not change the choice distribution significantly, across different choice uncertainty levels. Thus, the two causes we had assumed might influence users' choices did not actually do so: the perceived movement effort was insufficient to influence a user's decision, and the self-avatar follower effect did not occur. There are several potential explanations for these results.

Participants seemed to have "strong opinions" on which targets they wanted to select, which could not be easily influenced by visual manipulations of the faster VR body. Such "strong opinions" could be reflected in their movement trajectories [10]. As shown in Figure 5, at the onset of their movements, participants attempted to resist the VR body moving in the opposite direction of their final choices. This may be because we repeatedly emphasized in the study that participants should complete the task as quickly and as accurately as possible, and that their rewards would be scaled based on performance. As a result, they did not "trust" the VR body from the beginning.

Another possible reason is that participants could have already accumulated enough evidence for a decision within the 400 ms interval when the random dot display was shown. Thus, when the target appeared, the task was more similar to a goal-directed reaching action [8]. If this was the case, it may explain why the 0.3 s faster VR body led to longer choice completion time: it caused significant discrepancies between the users' expected and actual movements, requiring additional movement corrections, similar to snapping a cursor [62]. However, our analysis on a subset of trials with longer movement initiation time produced the same results of no influence on choice distributions (see the supplementary materials). Future research could modify task instructions and procedure (e.g., starting moving the virtual hand during when the random dot display is shown) to investigate whether user choices will be influenced.

The boundary condition at which perceived movement effort begins to influence choices remains unclear, as previous work has demonstrated that such an influence exists [17]. If we take the idea to the extreme, where the VR body is 1 second faster in our task (i.e., the hand will be directly on a button before the user moves), would this affect users' choices? Furthermore, the self-avatar follower effect did not change users' choices, even when the choice uncertainty was supposed to be high with a large portion of randomly moved dots [2]. It could be because the participants were explicit about the target they wanted to reach for, even though the choice could be wrong [15, 40]. It remains interesting to apply *VR Luke* to choice tasks that are less time- or performance-critical.

5.3 Implementation Variations

In*VR Luke*, we estimated where the user's hand might be τ seconds later using a typical movement profile based on a minimum-jerk model [30]. The idea was to maintain the temporal offset in the typical movement profile by adjusting the associated spatial offsets. As shown in Appendix B, the decision and movement time which we used to parametrize the minimum-jerk model aligned reasonably well with the data from the empirical study.

An alternative implementation of temporal offsets could incorporate real-time movement kinematics [3, 11, 18, 22]. For example, Kasahara et al. [22] used hand velocities to extrapolate future hand positions. As a result, if a user abruptly stops their movement mid-reach, the virtual hand would begin to drift back towards the physical hand position, rather than remaining in position as with *VR Luke*. In *VR Luke*, we chose to maintain temporal offsets in a typical movement profile instead of predicting hand kinematics in real-time because those action characteristics did not allow us to move the VR body before users make up their decisions. Nevertheless, future work could explore such alternative implementations, which may lead to different implications than those of our study.

Future studies could also apply other predictive models [14, 39] or fitting movement recordings [13] to better align the typical movement profile with actual user movements. However, while *VR Luke* estimates where the physical hand will typically be τ seconds later, it does not allow true prediction of the future physical hand position once the "future" virtual hand is displayed. The reason is that once users observe the "future" virtual hand, they would start adjusting their physical hand movements (e.g., decelerating earlier) [8], which would result in their physical hand ending up at a different position than the previous prediction (see Appendix C). While the temporal offset is still maintained within the typical movement profile, the "prediction error" may never be fully eliminated regardless of which model is used.

Moreover, future research could generalize our approach to interaction tasks that contain more than two targets or move the faster VR body between targets (e.g., towards their midpoint) during decision-making rather than towards a single target. Additionally, the temporal offset can change over time. For example, the temporal offset can increase as the user's hand moves towards the target and as the model's predictions become more accurate [61]. Finally, although we did not evaluate combinations of temporally and spatially faster VR bodies (see Appendix A.2 for examples), nor compare these variants empirically, future work could investigate their optimal combinations.

5.4 Additional Future Directions and Applications

5.4.1 Subjective experiences and perceptions.

Although assessing subjective experiences was not feasible in our study with randomized trial order, they remain an important direction for future research. One key question is whether users regard such a faster VR body as their own or treat it only as an external tool for completing tasks. This question relates to the concept of embodiment and body ownership [21, 26, 38]. Previous research has shown, for example, that the level of embodiment and body ownership might influence action characteristics (e.g., self-follower effect) [4, 45]. The users' feeling of having a faster VR body is also worth investigating; whether it may feel "creepy" or "uncanny" as with delayed or pre-recorded VR bodies [21] or it may feel "energetic" or "too much motion" as with other faster VR body implementations [22]. Moreover, it can also be interesting to explore the VR body's influence on time perception [29, 53].

5.4.2 Beyond hand and arm.

Our current research focused on the hand and arm of the VR body, since the choice task primarily involves these two parts. Future research could consider speeding up other body parts. One example is to manipulate virtual finger movements for tasks such as clicking. Since triggering the click requires very little finger movement, the acceptable temporal offset may work on a much short timescale. Such a faster VR body may also risk compromising users' sense of agency [51]. Another example is to manipulate the full VR body, including a user's viewpoint, in tasks such as locomotion. However, such a VR body may be prone to motion sickness. Additionally, it can be interesting to apply temporal offsets to choice tasks in non-immersive environments, such as moving virtual pointers in augmented reality (AR) or cursors in desktops during decision-making, to investigate whether they can be beneficial.

5.4.3 Towards everyday decision-making.

Future research could explore applications of a faster VR body in real-world situations, given that many of them also involve interactions between decision-making and action [10]. For instance, when a user is choosing which items to purchase in a VR shopping scenario or deciding which path to take in VR navigation, the VR body can start moving towards one option before the user has made an initial decision. If aligned with our study results, the user should be faster in completing these decisions while their decision outcomes should not be significantly influenced. However, such scenarios are much more complex than the perceptual decision-making task we used, as they involve more trade-offs and are more than just a simple correct/incorrect outcome [24]. Ethical implications should also be considered if the faster VR body starts changing users' choices in those cases [28].

6 Conclusion

We have introduced a faster VR body to speed up choices. The VR body relies on temporal offsets, which reposition the virtual body τ seconds ahead of the physical body based on a typical movement profile. Therefore, the VR body can start moving to one of the targets before the user has decided which target to reach for. We explored the use of such a faster VR body through a prototype called *VR Luke*, focusing on an accelerated hand and arm, and deployed it in a choice task based on random-dot displays. A user study with 34 participants suggested that choice time decreased when the virtual body was 0.1 seconds faster than the physical body, but increased when it was 0.3 seconds faster, regardless of whether it initially went to the correct or incorrect target. Users' choice distributions showed no significant differences across different choice uncertainty levels. We envision our research opening up new opportunities for applying a faster VR body through temporal offsets and offering new ways to speed up choices and decision-making.

Acknowledgments

This work was supported by the Pioneer Centre for AI (DNRF P1), the Independent Research Fund Denmark (DFF 1134-00012B), the European Research Council (Grant agreement No. 101231721), and the Villum Fonden (VIL73783, TRACTION).

References

[1] Brett Benda, Benjamin Rheault, Yanna Lin, and Eric D Ragan. 2024. Examining Effects of Technique Awareness on the Detection of Remapped Hands in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 30, 5 (2024), 2651–2661. https://doi.org/10.1109/TVCG.2024.3372054

[2] Loën Boban, Ronan Boulic, and Bruno Herbelin. 2024. In Case of Doubt, One Follows One's Self: The Implicit Guidance of the Embodied Self-Avatar. *IEEE Transactions on Visualization and Computer Graphics* 30, 5 (2024), 2109–2118. https://doi.org/10.1109/TVCG.2024.3372042

[3] Mattias Bratt, Christian Smith, and Henrik I Christensen. 2007. Minimum jerk based prediction of user actions for a ball catching task. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2710–2716. https://doi.org/10.1109/IROS.2007.4398989

[4] Dalila Burin, Konstantina Kilteni, Marco Rabuffetti, Mel Slater, and Lorenzo Pia. 2019. Body ownership increases the interference between observed and executed movements. *PloS one* 14, 1 (2019), e0209899. https://doi.org/journal.pone.0209899

[5] Géry Casiez, Daniel Vogel, Ravin Balakrishnan, and Andy Cockburn. 2008. The impact of control-display gain on user performance in pointing tasks. *Human–computer interaction* 23, 3 (2008), 215–250. https://doi.org/10.1080/07370020802278163

[6] Paul Cisek and John F Kalaska. 2010. Neural mechanisms for interacting with a world full of action choices. *Annual review of neuroscience* 33, 1 (2010), 269–298. https://doi.org/10.1146/annurev.neuro.051508.135409

[7] Paul Cisek, Geneviève Aude Puskas, and Stephany El-Murr. 2009. Decisions in changing conditions: the urgency-gating model. *Journal of Neuroscience* 29, 37 (2009), 11560–11571. https://doi.org/10.1523/JNEUROSCI.1844-09.2009

[8] Digby Elliott, James Lyons, Spencer J Hayes, James J Burkitt, James W Roberts, Lawrence EM Grierson, Steve Hansen, and Simon J Bennett. 2017. The multiple process model of goal-directed reaching revisited. *Neuroscience & Biobehavioral Reviews* 72 (2017), 95–110. https://doi.org/10.1016/j.neubiorev.2016.11.016

[9] Jason P Gallivan and Craig S Chapman. 2014. Three-dimensional reach trajectories as a probe of real-time decision-making between multiple competing targets. *Frontiers in neuroscience* 8 (2014), 215. https://doi.org/10.3389/fnins.2014.00215

[10] Jason P Gallivan, Craig S Chapman, Daniel M Wolpert, and J Randall Flanagan. 2018. Decision-making in sensorimotor control. *Nature Reviews Neuroscience* 19, 9 (2018), 519–534. https://doi.org/10.1038/s41583-018-0045-9

[11] Nisal Menuka Gamage, Deepana Ishtaweera, Martin Weigel, and Anusha Withana. 2021. So Predictable! Continuous 3D Hand Trajectory Prediction in Virtual Reality. In *The 34th Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) *(UIST '21)*. Association for Computing Machinery, New York, NY, USA, 332–343. https://doi.org/10.1145/3472749.3474753

[12] Zihan Gao, Huiqiang Wang, Hongwu Lv, Moshu Wang, and Yifan Qi. 2020. Evaluating the effects of non-isomorphic rotation on 3d manipulation tasks in mixed reality simulation. *IEEE Transactions on Visualization and Computer Graphics* 28, 2 (2020), 1261–1273. https://doi.org/10.1109/TVCG.2020.3010247

[13] Eric J. Gonzalez, Parastoo Abtahi, and Sean Follmer. 2019. Evaluating the Minimum Jerk Motion Model for Redirected Reach in Virtual Reality. In *Adjunct Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) *(UIST '19 Adjunct)*. Association for Computing Machinery, New York, NY, USA, 4–6. https://doi.org/10.1145/3332167.3357096

[14] Eric J Gonzalez, Elyse D. Z. Chase, Pramod Kotipalli, and Sean Follmer. 2022. A Model Predictive Control Approach for Reach Redirection in Virtual Reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 39, 15 pages. https://doi.org/10.1145/3491102.3501907

[15] Mar Gonzalez-Franco, Brian Cohn, Eyal Ofek, Dalila Burin, and Antonella Maselli. 2020. The self-avatar follower effect in virtual reality. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. Ieee, 18–25. https://doi.org/10.1109/VR46266.2020.00019

[16] Patrick Haggard. 2017. Sense of agency in the human brain. *Nature Reviews Neuroscience* 18, 4 (2017), 196–207. https://doi.org/10.1038/nrn.2017.14

[17] Nobuhiro Hagura, Patrick Haggard, and Jörn Diedrichsen. 2017. Perceptual decisions are biased by the cost to act. *Elife* 6 (2017), e18422. https://doi.org/10.7554/eLife.18422

[18] Rorik Henrikson, Tovi Grossman, Sean Trowbridge, Daniel Wigdor, and Hrvoje Benko. 2020. Head-Coupled Kinematic Template Matching: A Prediction Model for Ray Pointing in VR. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/3313831.3376489

[19] Niels Henze, David Halbhuber, and Jannik Wiese. 2023. Faster Might Not be Fast Enough: Improving Users' Experience and Performance through Negative Latency. In *Mensch und Computer 2023-Workshopband*. GI, 10–18420. https://doi.org/10.18420/muc2023-mci-ws05-405

[20] Neville Hogan and Dagmar Sternad. 2009. Sensitivity of smoothness measures to movement duration, amplitude, and arrests. *Journal of motor behavior* 41, 6 (2009), 529–534. https://doi.org/10.3200/35-09-004-RC

[21] Olga Iarygina, Kasper Hornbæk, and Aske Mottelson. 2025. Does Random Movements mean Random Results? Why Asynchrony in Experiments on Body Ownership does not Work as Intended. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 369, 19 pages. https://doi.org/10.1145/3706598.3713506

[22] Shunichi Kasahara, Keina Konno, Richi Owaki, Tsubasa Nishi, Akiko Takeshita, Takayuki Ito, Shoko Kasuga, and Junichi Ushiba. 2017. Malleable Embodiment: Changing Sense of Embodiment by Spatial-Temporal Deformation of Virtual Human Body. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '17)*. Association for Computing Machinery, New York, NY, USA, 6438–6448. https://doi.org/10.1145/3025453.3025962

[23] Shunichi Kasahara, Jun Nishida, and Pedro Lopes. 2019. 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* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–15. https://doi.org/10.1145/3290605.3300873

[24] Ralph L Keeney and Howard Raiffa. 1993. *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge university press.

[25] Roozbeh Kiani, Leah Corthell, and Michael N Shadlen. 2014. Choice certainty is informed by both evidence and decision time. *Neuron* 84, 6 (2014), 1329–1342. https://doi.org/10.1016/j.neuron.2014.12.015

[26] Martin Kocur, Alexander Kalus, Johanna Bogon, Niels Henze, Christian Wolff, and Valentin Schwind. 2022. The Rubber Hand Illusion in Virtual Reality and the Real World - Comparable but Different. In *Proceedings of the 28th ACM Symposium on Virtual Reality Software and Technology* (Tsukuba, Japan) *(VRST '22)*. Association for Computing Machinery, New York, NY, USA, Article 31,

12 pages. https://doi.org/10.1145/3562939.3565614

[27] John W Krakauer, Alkis M Hadjiosif, Jing Xu, Aaron L Wong, and Adrian M Haith. 2019. Motor learning. *Compr Physiol* 9, 2 (2019), 613–663. https://doi.org/10.1002/cphy.c170043

[28] Veronika Krauß, Pejman Saeghe, Alexander Boden, Mohamed Khamis, Mark McGill, Jan Gugenheimer, and Michael Nebeling. 2024. What Makes XR Dark? Examining Emerging Dark Patterns in Augmented and Virtual Reality through Expert Co-Design. *ACM Trans. Comput.-Hum. Interact.* 31, 3, Article 32 (Aug. 2024), 39 pages. https://doi.org/10.1145/3660340

[29] Maximilian Landeck, Fabian Unruh, Jean-Luc Lugrin, and Marc Erich Latoschik. 2024. Object Motion Manipulation and time perception in virtual reality. *Frontiers in Virtual Reality* 5 (2024), 1390703. https://doi.org/10.3389/frvir.2024.1390703

[30] Edward Lank, Yi-Chun Nikko Cheng, and Jaime Ruiz. 2007. Endpoint prediction using motion kinematics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '07)*. Association for Computing Machinery, New York, NY, USA, 637–646. https://doi.org/10.1145/1240624.1240724

[31] Jialei Li, Isaac Cho, and Zachary Wartell. 2018. Evaluation of Cursor Offset on 3D Selection in VR. In *Proceedings of the 2018 ACM Symposium on Spatial User Interaction* (Berlin, Germany) *(SUI '18)*. Association for Computing Machinery, New York, NY, USA, 120–129. https://doi.org/10.1145/3267782.3267797

[32] Steve W Link. 1975. The relative judgment theory of two choice response time. *Journal of Mathematical Psychology* 12, 1 (1975), 114–135. https://doi.org/10.1016/0022-2496(75)90053-X

[33] Xiaolong Liu, Lili Wang, Shuai Luan, Xuehuai Shi, and Xinda Liu. 2022. Distant object manipulation with adaptive gains in virtual reality. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 739–747. https://doi.org/10.1109/ISMAR55827.2022.00092

[34] Yasutoshi Makino, Yutaro Shioi, Yuuki Horiuchi, and Hiroyuki Shinoda. 2019. Interference of projected future self. In *2019 IEEE International Conference on Systems, Man and Cybernetics (SMC)*. IEEE, 4113–4117. https://doi.org/10.1109/SMC.2019.8914259

[35] Encarni Marcos, Ignasi Cos, Benoît Girard, and Paul FMJ Verschure. 2015. Motor cost influences perceptual decisions. *PLoS One* 10, 12 (2015), e0144841. https://doi.org/10.1371/journal.pone.0144841

[36] Manuel Molano-Mazón, Alexandre Garcia-Duran, Jordi Pastor-Ciurana, Lluís Hernández-Navarro, Lejla Bektic, Debora Lombardo, Jaime de la Rocha, and Alexandre Hyafil. 2024. Rapid, systematic updating of movement by accumulated decision evidence. *Nature communications* 15, 1 (2024), 10583. https://doi.org/10.1038/s41467-024-53586-7

[37] Pierre Morel, Philipp Ulbrich, and Alexander Gail. 2017. What makes a reach movement effortful? Physical effort discounting supports common minimization principles in decision making and motor control. *PLoS biology* 15, 6 (2017), e2001323. https://doi.org/10.1371/journal.pbio.2001323

[38] Aske Mottelson, Andreea Muresan, Kasper Hornbæk, and Guido Makransky. 2023. A Systematic Review and Meta-analysis of the Effectiveness of Body Ownership Illusions in Virtual Reality. *ACM Trans. Comput.-Hum. Interact.* 30, 5, Article 76 (Sept. 2023), 42 pages. https://doi.org/10.1145/3590767

[39] Jörg Müller, Antti Oulasvirta, and Roderick Murray-Smith. 2017. Control Theoretic Models of Pointing. *ACM Trans. Comput.-Hum. Interact.* 24, 4, Article 27 (Aug. 2017), 36 pages. https://doi.org/10.1145/3121431

[40] Gonçalo Padrao, Mar Gonzalez-Franco, Maria V Sanchez-Vives, Mel Slater, and Antoni Rodriguez-Fornells. 2016. Violating body movement semantics: Neural signatures of self-generated and external-generated errors. *Neuroimage* 124 (2016), 147–156. https://doi.org/10.1016/j.neuroimage.2015.08.022

[41] Praveen K Pilly and Aaron R Seitz. 2009. What a difference a parameter makes: A psychophysical comparison of random dot motion algorithms. *Vision research* 49, 13 (2009), 1599–1612. https://doi.org/10.1016/j.visres.2009.03.019

[42] Timothy J Pleskac and Jerome R Busemeyer. 2010. Two-stage dynamic signal detection: a theory of choice, decision time, and confidence. *Psychological review* 117, 3 (2010), 864. https://doi.org/10.1037/a0019737

[43] Ivan Poupyrev, Mark Billinghurst, Suzanne Weghorst, and Tadao Ichikawa. 1996. The go-go interaction technique: non-linear mapping for direct manipulation in VR. In *Proceedings of the 9th Annual ACM Symposium on User Interface Software and Technology* (Seattle, Washington, USA) *(UIST '96)*. Association for Computing Machinery, New York, NY, USA, 79–80. https://doi.org/10.1145/237091.237102

[44] Roger Ratcliff and Jeffrey N Rouder. 1998. Modeling response times for two-choice decisions. *Psychological science* 9, 5 (1998), 347–356. https://doi.org/10.1111/1467-9280.00067

[45] Natasha Ratcliffe and Roger Newport. 2017. The effect of visual, spatial and temporal manipulations on embodiment and action. *Frontiers in human neuroscience* 11 (2017), 227. https://doi.org/10.3389/fnhum.2017.00227

[46] Arbora Resulaj, Roozbeh Kiani, Daniel M Wolpert, and Michael N Shadlen. 2009. Changes of mind in decision-making. *Nature* 461, 7261 (2009), 263–266. https://doi.org/10.1038/nature08275

[47] Jamie D Roitman and Michael N Shadlen. 2002. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of neuroscience* 22, 21 (2002), 9475–9489. https://doi.org/10.1523/JNEUROSCI.22-21-09475.2002

[48] Michael N Shadlen and Roozbeh Kiani. 2013. Decision making as a window on cognition. *Neuron* 80, 3 (2013), 791–806. https://doi.org/10.1016/j.neuron.2013.10.047

[49] Reza Shadmehr, Helen J Huang, and Alaa A Ahmed. 2016. A representation of effort in decision-making and motor control. *Current biology* 26, 14 (2016), 1929–1934. https://doi.org/10.1016/j.cub.2016.05.065

[50] Philip L Smith and Douglas Vickers. 1988. The accumulator model of two-choice discrimination. *Journal of Mathematical Psychology* 32, 2 (1988), 135–168. https://doi.org/10.1016/0022-2496(88)90043-0

[51] Daisuke Tajima, Jun Nishida, Pedro Lopes, and Shunichi Kasahara. 2022. Whose Touch is This?: Understanding the Agency Trade-Off Between User-Driven Touch vs. Computer-Driven Touch. *ACM Trans. Comput.-Hum. Interact.* 29, 3, Article 24 (Jan. 2022), 27 pages. https://doi.org/10.1145/3489608

[52] Hans-Leo Teulings, José L Contreras-Vidal, George E Stelmach, and Charles H Adler. 1997. Parkinsonism reduces coordination of fingers, wrist, and arm in fine motor control. *Experimental neurology* 146, 1 (1997), 159–170. https://doi.org/10.1006/exnr.1997.6507

[53] Fabian Unruh, David Vogel, Maximilian Landeck, Jean-Luc Lugrin, and Marc Erich Latoschik. 2023. Body and time: Virtual embodiment and its effect on time perception. *IEEE Transactions on Visualization and Computer Graphics* 29, 5 (2023), 2626–2636. https://doi.org/10.1109/TVCG.2023.3247040

[54] Haopeng Wang, Ludwig Sidenmark, Florian Weidner, Joshua Newn, and Hans Gellersen. 2025. HeadShift: Head Pointing with Dynamic Control-Display Gain. *ACM Trans. Comput.-Hum. Interact.* 32, 1, Article 2 (April 2025), 28 pages. https://doi.org/10.1145/3689434

[55] Kieran Waugh, Mark McGill, and Euan Freeman. 2025. Everything to Gain: Combining Area Cursors with increased Control-Display Gain for Fast and Accurate Touchless Input. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 726, 16 pages. https://doi.org/10.1145/3706598.3714021

[56] Curtis Wilkes and Doug A Bowman. 2008. Advantages of velocity-based scaling for distant 3D manipulation. In *Proceedings of the 2008 ACM symposium on Virtual reality software and technology*. 23–29. https://doi.org/10.1145/1450579.1450585

[57] Nathan J Wispinski, Jason P Gallivan, and Craig S Chapman. 2020. Models, movements, and minds: bridging the gap between decision making and action. *Annals of the New York Academy of Sciences* 1464, 1 (2020), 30–51. https://doi.org/10.1111/nyas.13973

[58] Aaron L Wong and Adrian M Haith. 2017. Motor planning flexibly optimizes performance under uncertainty about task goals. *Nature communications* 8, 1 (2017), 14624. https://doi.org/10.1038/ncomms14624

[59] Aaron L Wong, Adrian M Haith, and John W Krakauer. 2015. Motor planning. *The Neuroscientist* 21, 4 (2015), 385–398. https://doi.org/10.1177/1073858414541484

[60] Difeng Yu, Mantas Cibulskis, Erik Skjoldan Mortensen, Mark Schram Christensen, and Joanna Bergström. 2024. Metrics of Motor Learning for Analyzing Movement Mapping in Virtual Reality. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 724, 18 pages. https://doi.org/10.1145/3613904.3642354

[61] Difeng Yu, Ruta Desai, Ting Zhang, Hrvoje Benko, Tanya R. Jonker, and Aakar Gupta. 2022. Optimizing the Timing of Intelligent Suggestion in Virtual Reality. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) *(UIST '22)*. Association for Computing Machinery, New York, NY, USA, Article 6, 20 pages. https://doi.org/10.1145/3526113.3545632

[62] Difeng Yu, James Roberts, Kasper Hornbæk, and Joanna Bergström. 2025. Deriving Selection Techniques for GUIs based on the Multiple Process Model. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 851, 16 pages. https://doi.org/10.1145/3706598.3713089

[63] Ariel Zylberberg, Christopher R Fetsch, and Michael N Shadlen. 2016. The influence of evidence volatility on choice, reaction time and confidence in a perceptual decision. *Elife* 5 (2016), e17688. https://doi.org/10.1037/xhp0000404

A Spatial vs. Temporal Offsets

For interested readers, we distinguish spatial and temporal offsets more precisely in this section. We also illustrate potential interplays between spatially and temporally faster VR bodies.

A.1 Definitions and Properties

We first describe a physical input movement in VR (e.g., performed by a physical hand) through $\mathbf{x}_t$. Here, $\mathbf{x}_t$ stands for location information of the physical input, such as distances or positions, at a

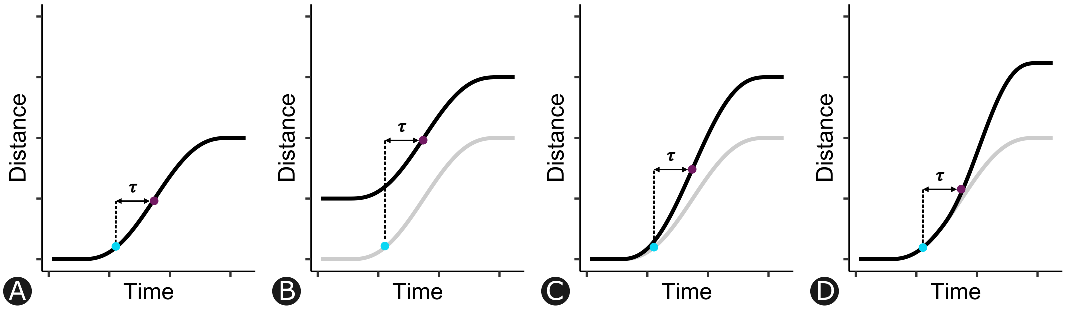

Figure 6: Temporal offsets τ can be integrated into techniques that enable spatially faster VR bodies. The blue dots indicate the distance of the physical hand, while the purple dots indicate the distance of the virtual hand. (A) Physical/virtual input movement under a 1:1 physical-virtual mapping. (B) Virtual input movement (black) by applying a constant spatial offset to the physical input movement (gray). (C) Virtual input movement with a linear gain. (D) Virtual input movement with the Go-Go technique, which applies quadratic amplification during the latter portion of the movement.

timestamp t, where $t \in \{t_0, t_1, t_2, ..., t_n\}$, from when the movement begins (t_0) to when it ends (t_n). The corresponding virtual input movement (e.g., virtual hand distances or positions) is $\mathbf{x}'_t$.

The physical and virtual input movements are connected by a *mapping function* $f(\cdot)$, so that

$$\mathbf{x}'_t = f(\mathbf{x}_t) \tag{6}$$

Equation 6 means for a given t, the virtual input movement $\mathbf{x}'_t$ is mapped from the physical input movement $\mathbf{x}_t$ through $f(\cdot)$. A common mapping function is $f(\mathbf{x}_t) = \mathbf{x}_t$. In this case, we always have $\mathbf{x}'_t = \mathbf{x}_t$, which is the feature of a 1:1 physical-virtual mapping.

A *spatial offset* occurs if $\mathbf{x}'_t \neq \mathbf{x}_t$, as the physical input differs from the virtual input at t. The offset can be depicted by $\Delta\mathbf{x}_t = \mathbf{x}'_t - \mathbf{x}_t$. For example, a mapping function based on a linear gain k ($k > 1$) can be described as $f(\mathbf{x}_t) = k \cdot \mathbf{x}_t$, which leads to the spatial offsets being $\Delta\mathbf{x}_t = k \cdot \mathbf{x}_t - \mathbf{x}_t$. A similar idea can be used to characterize other techniques that introduce new mapping functions $f(\cdot)$ to induce spatial offsets, such as Go-Go [43].

Now, we define a *temporal offset* τ in the subscript of $\mathbf{x}$. We here consider $\tau \in \mathbb{R}$ for simplicity (note that an actual computer system logs timestamps in a discrete manner). This results in

$$\mathbf{x}'_t = f(\mathbf{x}_{t+\tau}) \tag{7}$$

When $\tau = 0$, there is no temporal offset. A delay or lag in a computer system leads to $\tau < 0$, which means that a previous physical input position determines the current virtual input position. In contrast, *VR Luke* attempts to maintain $\tau > 0$, which means that an expected future physical input position determines the current virtual input position.

Based on the definition, we note that a temporal offset τ can inevitably result in a spatial offset. For instance, in the case of a 1:1 mapping, where $f(\mathbf{x}_t) = \mathbf{x}_t$, the spatial offset at t is described by $\Delta\mathbf{x}_t = \mathbf{x}_{t+\tau} - \mathbf{x}_t$. Therefore, a spatial offset occurs as long as $\mathbf{x}_{t+\tau} \neq \mathbf{x}_t$ during a physical input movement. This property also enables us to approximate the behavior of a temporal offset through spatial offsets. *VR Luke* achieved this approximation by preserving the spatial offsets associated with the temporal offset in a typical movement profile.

A.2 Interplay Between Spatially and Temporally Faster VR Bodies

One implication of the above definition is that we could integrate a temporal offset τ into techniques that enable spatially faster VR bodies, which modifies the mapping function $f(\cdot)$. We demonstrate three examples (also see Figure 6).

- Constant spatial offset [31] introduces a fixed spatial offset c between the physical and virtual input movement. Its mapping function can be written as $f(\mathbf{x}_t) = \mathbf{x}_t + c$. A temporal offset τ can be integrated to form $\mathbf{x}'_t = \mathbf{x}_{t+\tau} + c$.
- A linear gain [60] scales up the physical input movement with a constant k. Its mapping function can be expressed as $f(\mathbf{x}_t) = k \cdot \mathbf{x}_t$. A temporally and spatially faster VR body can be expressed as $\mathbf{x}'_t = k \cdot \mathbf{x}_{t+\tau}$.
- The Go-Go technique [43] applies a 1:1 mapping $f(\mathbf{x}_t) = \mathbf{x}_t$ at the start of the movement and includes a quadratic term $f(\mathbf{x}_t) = \mathbf{x}_t + k(\mathbf{x}_t - D)^2$ during the latter part of the motion. Depending on whether $\mathbf{x}_{t+\tau}$ might be within the first or second portion of the movement, a temporally and spatially faster VR body can be expressed as either $\mathbf{x}'_t = \mathbf{x}_{t+\tau}$ or $\mathbf{x}'_t = \mathbf{x}_{t+\tau} + k(\mathbf{x}_{t+\tau} - D)^2$.

Notably, if we follow the idea of *VR Luke*, we must first collect typical movement profiles of these spatially faster VR bodies when reaching for different targets. Therefore, introducing temporal offsets adapts the original target-agnostic technique, which requires no information about the target, to a target-aware one.

B Minimum-Jerk Model Parameters

To parametrize our minimum-jerk model, we assumed that a typical user movement profile (without applying any temporal offset) had an initial decision time of 0.3 s and a movement time of 1 s in the task, which led to a total choice time of 1.3 s. The empirical data in the BASELINE condition had an average choice time of 1.29 s (std = 0.57), which closely aligned with the assumed choice time of 1.3 s. The decision time was difficult to determine precisely because the movement data contained tracking inaccuracies and participants might not keep their hands perfectly still after pressing down the start button. As an approximation, we used the movement initiation

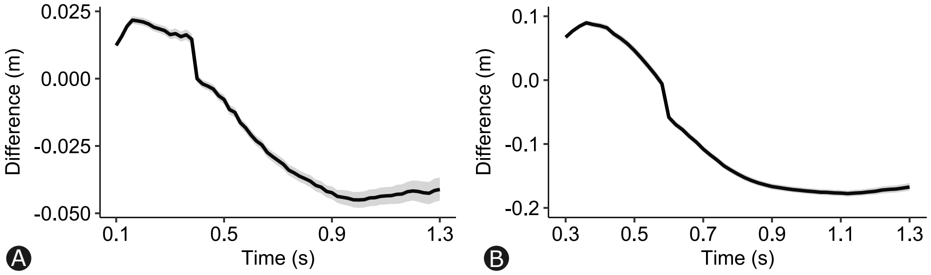

Figure 7: The average difference between the physical hand distance and the estimated virtual hand distance 0.1 s before (when the temporal offset $\tau = 0.1$, A) or 0.3 s before (when $\tau = 0.3$, B). The gray error region represents the 95% bootstrap confidence interval.

time (i.e., when the user has moved 0.04 m as in Section 4.4.3), which had an average of 0.35 s ($std = 0.26$). Note that this might be a conservative estimation (i.e., an upper boundary), so the actual decision time could be smaller. As a result, the averaged movement time, which represented the lower boundary of the actual movement time, was 0.94 s ($std = 0.47$). These values also aligned reasonably well with our initial assumption of a typical user movement.

C Expected vs. Actual Movements with Temporal Offsets

In *VR Luke*, we maintained the temporal offset in a typical movement profile by adjusting the corresponding spatial offset. Figure 7 showed the average difference between the physical hand distance and the estimated virtual hand distance 0.1 s before (when the temporal offset $\tau = 0.1$) or 0.3 s before (when $\tau = 0.3$). When $\tau = 0.1$, the most positive difference (i.e., the physical hand moved more) was 0.02 m, and the most negative difference (i.e., the virtual hand moved more) was -0.05 m. When $\tau = 0.3$, the most positive and the most negative differences were 0.09 m and -0.18 m.

Notably, *VR Luke* appeared to be underestimating the virtual hand distance at the beginning of the movement while overestimating it near the end of the trial, particularly when the temporal offset was 0.3 s. The underestimation could be due to the reaction time being shorter than expected in easier trials when motion coherence was high. The overestimation could be because once users observed that the virtual hand was much closer to the target than expected (as the VR body was faster), they slowed down their physical hand earlier (i.e., early stopping) [8]. As a result, the physical hand ended up moving a shorter distance than the model's estimation when pressing the target button with the virtual hand. This indicates that maintaining the temporal offset in a typical movement profile cannot provide true prediction for the physical hand position τ seconds later as users adjust their movement in a closed-loop manner.