2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)

2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR) | 978-1-6654-5325-7/22/$31.00 ©2022 IEEE | DOI: 10.1109/ISMAR55827.2022.00029

ComforTable User Interfaces: Surfaces Reduce Input Error, Time, and Exertion for Tabletop and Mid-air User Interfaces

Yi Fei Cheng* Tiffany Luong* Andreas Rene Fender* Paul Streli* Christian Holz*

Department of Computer Science
ETH Zürich

ABSTRACT

Real-world work-spaces typically revolve around tables, which enable knowledge workers to comfortably perform tasks over an extended period of time during productivity tasks. Tables afford more ergonomic postures and provide opportunities for rest, which raises the question of whether they may also benefit prolonged interaction in Virtual Reality (VR). In this paper, we investigate the effects of tabletop surface presence in situated VR settings on task performance, behavior, and subjective experience. In an empirical study, 24 participants performed two tasks (selection, docking) on virtual interfaces placed at two distances and two orientations. Our results show that a physical tabletop inside VR improves comfort, agency, and task performance while decreasing physical exertion and strain of the neck, shoulder, elbow, and wrist, assessed through objective metrics and subjective reporting. Notably, we found that these benefits apply when the UI is placed on and aligned with the table itself as well as when it is positioned vertically in mid-air above it. Our experiment therefore provides empirical evidence for integrating physical table surfaces into VR scenarios to enable and support prolonged interaction. We conclude by discussing the effective usage of surfaces in situated VR experiences and provide initial guidelines.

Index Terms: Human-centered computing—Human computer interaction (HCI)—Empirical Studies in HCI

1 INTRODUCTION

Our everyday environments are filled with tabletops. Tabletops are central to social gatherings, collaborative settings, and individual productivity. Especially during work, tables fulfill two simultaneous functions: They accommodate objects, allowing workers to keep the many items they often require within arm's reach; at the same time, tables serve an ergonomic function to enable workers to support themselves during prolonged tasks.

In contrast, objects in Virtual Reality (VR) can be freely placed around the user, even creating surrounding environments that would not be feasible in the physical world. The affordances of VR interfaces can thus fully replace the use of tabletops to keep items within reach, instead offering the opportunity to dynamically show relevant items and optimize their appearance for immediate use [10].

In contrast, the benefit of tables for ergonomically supporting prolonged interaction may extend to work in VR. Interaction inside VR is dominated by bimanual input, mediated through hand-held controllers or direct input using one's bare hands. Most VR scenarios assume an empty physical space around the user during interaction, allowing free-range input through spatial gestures. Once they exceed quick interactions, however, mid-air control impacts input precision [61] and leads to fatigue [32] ('gorilla arm' [24]). These effects pose a challenge for VR use in productivity tasks that are characterized by a need for prolonged and accurate interaction.

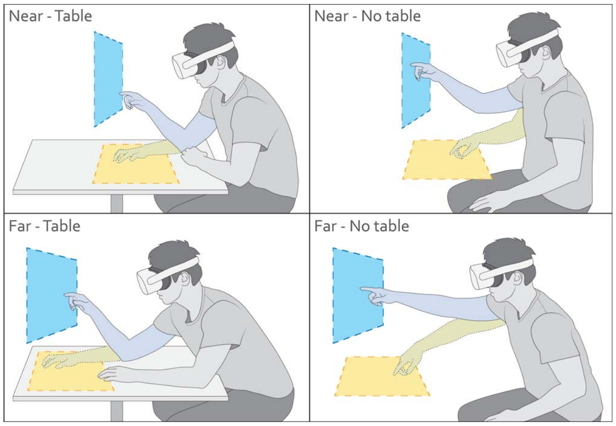

Figure 1: Our 24-participant user study investigated the impact a physical table has on input performance for virtual user interfaces, particularly accuracy, speed, and fatigue over time. Our study varied interface placement in orientation (horizontal—yellow or vertical—blue), distance (near and far), and input task (selection and docking).

In order to transfer the benefits of physical surfaces to VR, prior efforts have incorporated them into the interaction with virtual UIs to ground all input [38], especially for touch to increase input precision [54] and to support input with passive haptic feedback [2, 8]. In our own work, we have proposed adopting surfaces for *direct* interaction in VR while simultaneously offering an opportunity for rest during prolonged interaction in productivity scenarios (TapID [43]). We have also co-opted surfaces for opportunistic text ten-finger input anywhere to transfer the benefits of surface-based keyboard typing to mobile scenarios (TapType [52]). Consumer VR devices (e.g., Oculus Quest 2) have similarly begun to support bringing physical surfaces (e.g., couches or desks) into immersive environments [4].

In this paper, we present the results of a systematic study on the use of passive tabletops to support interaction in VR with a particular focus on interactions that exceed quick input. Beyond touch input, we investigated the effects and ergonomic benefits of surfaces on input performance in mid-air (i.e., above a physical table), depending on interface placement relative to them as shown in Fig. 1.

The effect of a tabletop surface on input performance

Our study of input performance with 24 participants investigated the effects of table presence, i.e., *table* versus *no table* on participant performance, behavior, and subjective preference. Participants completed two tasks with prolonged duration: *selection* and *docking*. Our study also varied the spatial configuration of the interaction space, particularly the *distance* and *UI orientation* of interface elements to determine if tabletops—even when not aligned with the virtual UI—maintain their effect on performance. We assessed user performance through quantitative metrics such as input accuracy and task completion time, and derived ergonomic metrics from motion-capture data of participants' shoulders, elbows, and wrists. Finally,

---

*e-mail: firstname.lastname@inf.ethz.ch

978-1-6654-5325-7/22/$31.00 ©2022 IEEE
DOI 10.1109/ISMAR55827.2022.00029

we evaluated participants' subjective preferences and feelings of comfort, exertion, and strain with questionnaires.

Our results show that table presence influences input accuracy, speed, and fatigue. Tabletops enabled interaction with greater accuracy and speed, even when the UI orientation was *not directly aligned* with the physical surface. Their performance benefits were complemented with support for less fatiguing interaction, which we observed through quantitative metrics and participants' self-reports. These results ultimately suggest that physical tabletops may support more prolonged and ergonomic interaction in VR. Notably, the results suggest that even for mid-air interaction (e.g. for vertically-oriented interface elements placed above the surface), providing a tabletop below the UI for resting is beneficial.

Our results challenge current guidelines that VR usage ought to mostly be hosted in empty physical spaces. On the contrary, we found that incorporating tabletop surfaces into immersive experiences may have concrete productive and ergonomic benefits. Our study showed that these benefits extend beyond simply providing passive haptic feedback. Furthermore, these insights indicate that a potential re-design of safeguard mechanisms (i.e., HTC Vive *Chaperone*, Oculus *Guardian*) may be needed in the future.

Taken together, we contribute an empirical study on the use and benefits of tabletops for virtual interfaces. Through our results, we provide a set of design guidelines for the effective usage of surfaces in VR as a supplementary input modality to current mid-air controls.

2 RELATED WORK

In the following, we survey related research on surface-supported interactions and physically situated user interfaces and experiences.

2.1 Surface-supported Interactions

There is extensive prior research on techniques for interfacing with devices and virtual interfaces. Of this body of literature, our study relates most closely to work on mid-air [33,36] and surface interactions [11,20,29,49] , especially studies empirically examining the distinctions between the two. More specifically, we are concerned with how surfaces influence performance and user experience in finger-driven direct manipulation interactions. Lindeman et al.'s early work [38], for instance, examined how the availability of hand-held and world-fixed surfaces affected task performance for 2D selection and docking in VR. Viciana-Abad et al. [57] similarly varied surface availability for a selection task in VR, and showed that passive haptic feedback improved user presence. Jakobsen et al. [31] designed a study that allowed participants choose freely between touch and mid-air gestures when performing gestures on an interactive surface. Bruder et al. [6,7] compared 2D and 3D mid-air interactions in a Fitts' law experiment for objects with varying stereoscopic parallax. Zielasko et al. [64,65] evaluated the impact of surface availability on interactions with touch-based menus with varying orientations (e.g., horizontal, vertical).

First, we draw inspiration from prior work on the influence of surface availability on interactions in our experimental design and variables (e.g. task [7,38], task orientation [65,66]). Second, while most prior work focuses on evaluating the benefits of a task-aligned surface in affording haptic feedback, we further investigate how surfaces affect performance in contexts where it is not aligned with the virtual task space. In particular, we examine how the presence of a tabletop may influence interactions with UIs on a vertical interaction plane placed above it (i.e., not aligned with the virtual UI).

Beyond experiments that explored surface support for direct touch interaction, our work also relates to prior research on the use of surfaces as a physical constraint for other input devices [53]. Several studies have examined the effects of surface availability on the use of a digital stylus [2,58] and desktop mice [35,55]. For both mouse and pen input, physical surfaces improve accuracy [2], user satisfaction [58], and reduce fatigue [35]. Prior work also revealed that the degrees of freedom of the input modality should be critically considered depending on the task requirements [26].

The work in this paper adds to prior research on surface-supported interaction with peripherals. While the literature provides insights into its utility in free-hand scenarios, there are marked differences in the characteristics of hand-driven interactions and controller-based input. A comparison of hand-driven mid-air and surface interaction could therefore more directly guide future controller-free interfaces.

Lastly, to understand the role of table surfaces in VR, we draw inspiration from work on interface ergonomics. Ergonomic factors play an important role in the design of physical spaces and interactions [13,45]. Several projects have prioritized this in their designs (e.g., Gunslinger [39] and elbow-anchored interactions [56]). Others have focused on deriving methods for evaluating ergonomic factors. RULA employed a simple heuristic for evaluating the ergonomic quality of mid-air postures based on joint angles [42]. Consumed Endurance [24] and Cumulative Arm Fatigue [32] applied a bi-mechanical model of arm pose to estimate physical exertion. Aside from heuristic and automated methods, prior work has also relied on subjective evaluations, such as custom Likert scales, the Borg CR10 scale [5], and the NASA Task Load Index [21]. We investigate the ergonomic benefits of table usage in VR tasks. To our knowledge, most quantitative metrics in HCI are designed for mid-air interactions, and focus on evaluating static poses. As such, in our study, we decided to rely primarily on self-reported metrics, such as the Borg CR10, to evaluate the ergonomic differences between surfaces-supported and mid-air interactions. We supplement these metrics with quantitative metrics on motion and pose.

2.2 Situated Virtual User Interfaces

In Augmented Reality (AR), it is typical for the virtual content that a user interacts with to be anchored to features of the physical environment [15]. In many AR applications, users interact around a virtual plane aligned to a physical surface [27,46]. Recent research has further leveraged optimization to place AR content in an automated, environment-aware manner [17].

Virtual Reality (VR) applications, by contrast, were dominantly designed in a way that isolates users from their surroundings and relies primarily on mid-air free-hand controls. Recent VR research, however, has also recognized benefits of situating fully immersive experiences in physical reality, such as improving immersion [30], task performance [23], and ergonomics during direct interaction [43,52]. A popular approach to situating VR experiences is through Substitutional Reality [50], appropriating passive affordances in the physical environment for virtual interfaces. This concept of Substitutional Reality has been explored at multiple scales, from adapting virtual environments to entire rooms to defining mappings between virtual elements to physical proxies (e.g., [8,9,23,62]). Another approach to situating VR experiences involves visualizing features of the physical environment (e.g., RealityCheck [22]). Consumer VR devices like the Oculus Quest, for instance, now enable users to bring their physical desk and couch into their immersive experience [4].

In our work, we empirically examine whether situating virtual interfaces around tables may benefit interactions. Here, we focus on table surfaces in particular because they are often readily accessible [34], provide support for prolonged interactions [3,43,52], and afford touch interaction [43,61]. We implemented an experiment platform to study user behavior and performance when completing tasks with and without a table in the absence of tracking issues. We envision our results informing systems that situate virtual content in different settings, particularly those with a table or smaller surfaces present (e.g., mobile workspaces). Our results go beyond suggesting that physical reality serves as a good way of providing passive haptic feedback, indicating additionally that they may influence interactions with virtual elements situated in mid-air around them.

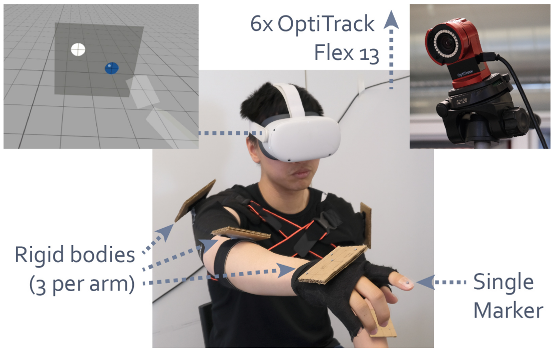

Figure 2: Apparatus of our experiment. Each participant wore a VR headset and cardboard patches strapped to their shoulder, elbow, and hand on each arm, all tracked by an OptiTrack system. We tracked the right index finger through a single facial marker on the nail.

3 EXPERIMENT

Prior work has indicated the potential benefit of tabletops for passive haptic feedback and for supporting ergonomic and restful postures [43]. To enable future VR applications to effectively co-opt tabletops, particularly to support prolonged usage, it is important to understand their merits over free-space interaction. Therefore, we investigate the following two research questions: (Q1) How does the presence of a table in VR affect users' task performance, interaction ergonomics, and preference? (Q2) To which degree do the effects of table presence depend on task type, interface distance, and orientation?

To inform our study choices, we first conducted a series of pilot studies with a total of 12 participants. Our pilot studies explored a selection of four TASKS (*selection*, *docking*, *flicking*, and *bi-manual alignment*) in *table* and *no table* conditions. We varied additional factors including task duration and placement.

Our final study involved 24 participants. We compare the effect of tabletop presence on interaction for two TASKS: *selection* and *docking*. Our procedure also varied the DISTANCE (*near* versus *far*) and ORIENTATION (*horizontal* and *vertical*) of the UI where participants performed the tasks. To represent productivity scenarios, tasks lasted for around three minutes to enable studying continued use. We quantify the effects on participants' performance (e.g., task completion time and accuracy), user behavior (i.e., ergonomic factors such as body posture and movement characteristics), and subjective preferences (e.g., Likert scale evaluations of comfort). In addition, we evaluate the effect of a tabletop surface for tasks that do *not* directly benefit from passive haptic feedback by repeating tasks in mid-air above the surface.

3.1 Apparatus

As shown in Fig. 2, our study apparatus continuously tracked participants' motions and input gestures using a 6-camera OptiTrack system with sub-mm accuracy. Each participant had a total of 29 markers attached on their their left and right arms. 24 markers were used to track the exact position and orientation of their left and right shoulders, elbow, and hand. Four markers yielded the position and orientation of a 0.7×0.9 m table surface employed in the *table* condition. Finally, one marker was used to track the position of the participant's right index finger. Using the output of the OptiTrack system, we used the poses of the markers to animate a virtual representation of the participants' arm and to place the table surface within our virtual scene. Following prior work [41, 48], we used rigid bodies to track participants' arms without body approximations through the OptiTrack system. We designed these rigid bodies from sturdy cardboard and attached them to each participant's shoulder, elbow, and hand using shoulder pads, Velcro straps, and gloves, respectively. For consistent tracking, at the beginning of each study session, the experimenter placed each marker at the same position on the straps participants wore to ensure consistency across participants. To verify that the registration with our VR apparatus had not slipped, participants also confirmed that the virtual arm shown in VR matched the location of their arm. If this was not the case, we recalibrated the whole system. We relied on OptiTrack's differentiation between labeled (either manually labeled or belonging to a rigid body) and unlabeled markers to perform single marker tracking of the user's finger [47]. As our set-up only consisted of single unlabeled marker, we could immediately associate it with the participant's index finger. Our apparatus ran on a Windows 10 PC with an Intel Core i7 processor and an NVIDIA RTX 3070.

Virtual scene

Participants produced input inside a light-grey gradient environment with a floor-aligned grid. Since prior work has found that abstract bodily representations provide a comparable or even greater sense of agency when interacting with virtual environments [1], we visualized participants' arms as three boxes aligned with their upper arm, lower arm, and hand, all calibrated in size to their bodies. Participants used their right index finger for input, which was displayed as semi-transparent white sphere measuring 1 cm in diameter. The remaining fingers were not shown.

The VR scene was rendered using Unity 2020.3.14f1 and an Oculus Quest 2 HMD with a refresh rate of 90 frames per second. Before each study, we calibrated the HMD's position as reported by the Oculus Integration package with the physical coordinate system of OptiTrack markers and arranged the setup using metric units.

3.2 Pilot Studies

To begin our explorations, we conducted a series of pilot studies with a total of 12 participants. Our objective with conducting these initial studies was primarily to determine parameters for our final study, as opposed to deriving insights. As such, we only report on findings relevant to our final study design. In our pilot studies, we explored a broader range of tasks, including *selection*, *docking*, *flicking*, and *bi-manual alignment*. *Selection* and *docking* are described in section 3.3. *Flicking* involved performing quick swipes in specified directions. Our *bi-manual alignment* task approximately followed the procedure presented in Forlines et al. [16]. Participants performed tasks in *table* and *no table* conditions. Our main insight from conducting our initial studies was that participants demonstrated a subjective preference between the two conditions. However, we did not record any metrics providing an explanation for the difference, and hence in our final study, we aimed to more systematically study dependent variables capturing user performance, behavior, and subjective experience. From our pilot studies, we also decided to focus on *selection* and *docking*. We excluded *flicking* because it was less accuracy oriented and typically used for shorter interactions. We leave *bi-manual alignment* for future work. Through our pilot studies, we also experimented with various task durations. We aimed for a time-frame that was sufficiently prolonged but not exhausting. We ultimately settled on 3 minutes for our final study. Lastly, in our pilot studies, we observed differences in user performance while varying UI orientation and distance. We therefore decided on varying the factors for a more systematic evaluation in our final experiment.

3.3 Task

We evaluated two tasks in our experiment: 2D selection and docking. The tasks were selected as representative abstractions of tasks commonly performed in graphical interfaces [16]. We focused on 2D tasks due to the pervasiveness of window-based interactions in 3D user interfaces [14, 37]. Tasks were displayed on a virtual

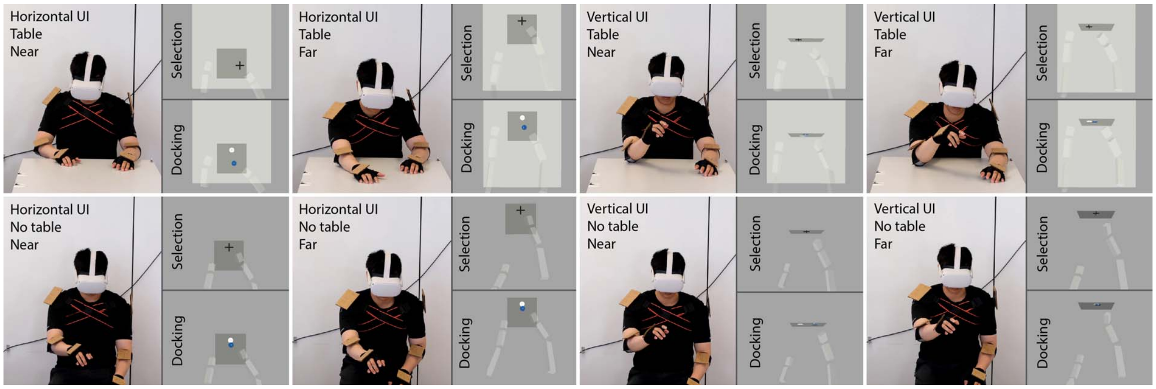

Figure 3: All conditions of our $2 \times 2 \times 2 \times 2$ mixed design. The left half shows conditions with a horizontal UI and the right half with a vertical UI (the only between-subject factor). A table is only present in the upper row. The distances for the *near* and *far* conditions depend on the participant's arm length. The virtual space is shown from a top-view (note how the UI positions depend on *near* and *far*). In the *selection* task, participants touched a sequence of crosses (enhanced in pictures for clarity). In the *docking* task, participants dragged the blue circle into the white circle.

semi-transparent rectangular canvas, and placed depending on the specific condition (Section 3.4). Participants performed all tasks in a seated position within the volume of the tracking system. For each condition, participants repeated input tasks for 3 minutes.

Selection task

We implemented a circular target arrangement following the ISO 9241-9 standard and based on prior experiments on selection performance [51,55,63]. Our interface comprised seven targets arranged on a plane, visualized as 4.5×4.5 cm black cross-hairs, with a distance of 15 cm between targets. Our apparatus registered target selections when the collider that tracked participants' right index fingers intersected with the task plane. We counted selections as successful when the touch was less than 2 cm away from the center of the cross-hair, and as unsuccessful otherwise. Note that selections were only considered successful if the participant's first point of contact with the task plane was within 2 cm; dragging one's finger from one target position to the next through the task plane without lift-off was invalid (e.g. crossing-based selection [40]). Upon a successful selection, the color of the cross-hairs turned green and audio feedback was provided. Participants were also notified of erroneous selections with audio feedback playing a different sound. Feedback appeared for 0.1 seconds before our apparatus proceeded to the next target's position.

Docking task

We used an adaptation of the ISO 9241-9 standard selection task for docking. Each trial showed a 4.5×4.5 cm diameter light blue target item and an identically sized white docking item. Both items were placed in consecutive positions of the ISO 9241-9 arrangement. Participants first selected the target similar to the selection task: using their index finger, they touched the task plane within 2 cm of the target's center. The target then followed a projected position of their index (i.e., onto the virtual plane) until it was dropped. The target is dropped if participants lifted their index finger 2 cm above the task plane. To complete the task, participants had to drag the target within a 0.5 cm distance of the dock for 0.5 seconds. We decided to require hovering to auto-activate a drop event to prevent input errors due to lift off. In addition, during input, we enabled participants to freely protrude or push through the task plane. We afforded participants this leniency to represent interaction with existing applications more realistically (e.g. interactions in Horizon Workrooms [44]). Successful target selection played a confirmation sound and the target's appearance changed to a darker shade of blue. As soon as a participant dragged the target into the accepted dock zone, its appearance switched to green. Successfully finishing docking played a success sound. Making contact with the plane outside the range of a target or dropping the target before successful docking played an error sound.

3.4 Additional independent variables

Table presence

Between each condition, we varied the TABLE PRESENCE. In the *table* conditions, the OptiTrack-registered table was displayed as a semi-transparent plane in the virtual environment. The table was initially placed relative to the seated user based on ISO 9241-5 recommendations for workspace layouts. We then adjusted the table placement and height based on the participant's individual preferences. In the *no table* conditions, no representation of the table surface was shown in the virtual environment and the physical table was shifted away outside participants' reach.

UI Distance and Orientation

Between each condition, we varied the UI DISTANCE and ORIENTATION. Here, UI refers to the virtual interaction plane on which all task items were located. We set the UI DISTANCE based on each participant's lower and upper arm lengths. Table 1 shows our placement parameters, which we determined through pilots. We intended for the *near* UI to be placed within a comfortable personal interaction range [18] and the *far* UI to be placed just slightly at the limits of the participant's each. Note that in the *table* condition, if the ORIENTATION was horizontal, the task plane was co-located with the table surface, so that the participant received passive haptic feedback upon input. In both, *vertical* conditions and *no table* × *horizontal* conditions, the task was performed in mid-air. We calibrated the table height by placing the physical table in front of the participant at a comfortable position.

3.5 Procedure

We used a mixed factorial design with 4 independent variables, each with 2 levels: TASK (*selection*, *docking*), TABLE PRESENCE (*table*, *no table*), DISTANCE (*near*, *far*), ORIENTATION (*horizontal*, *vertical*). All conditions of our study can be seen in Fig. 3 as well as in the supplementary video. ORIENTATION was evaluated as a between-subject variable. The remaining variables were evaluated within subjects. Half the participants were assigned *horizontal* conditions, the other half completed *vertical* conditions.

Table 1: Parameters for UI placement based on DISTANCE and ORIENTATION. We calibrated the UI placement distance based on each participant's lower and upper arm lengths, which are denoted here as l_{lower} and l_{upper} respectively.

DISTANCE	ORIENTATION	Height	Distance
near	*horizontal*	0.0	l_{lower}
near	*vertical*	$0.75l_{lower}$	$0.75l_{upper}$
far	*horizontal*	0.0	l_{upper}
far	*vertical*	$0.75l_{lower}$	l_{upper}

Table 2: Demographics of the analyzed study sample by between-subject groups. Continuous variables are summarized as M (SD) and ordinal variables are summarized as median.

Variable	Horizontal ($N = 12$)	Vertical ($N = 12$)
Age [20–52] (in years)	26.25 (8.55)	23.08 (2.97)
% Male	66.6% ($n = 8$)	66.6% ($n = 8$)
% Female	33.3% ($n = 4$)	33.3% ($n = 4$)
VR experience [1–5]	2	3
Hand-tracking experience [1–5]	1	1
Gaming frequency [1–5]	2	4
Level of alertness [0–7] [28]	6.5	7
Level of physical activity [1–3]	2	2

Experiments were divided into two DISTANCE blocks. Each block consisted of repetitions performed at either *near* or *far*. Within each block, participants performed each TASK twice, once under each *table* condition. Each repetition consisted of continuous task performance for exactly 3 minutes. Participants were required to rest for at least 3 minutes between each repetition to control for contamination effects. Participants could request for longer breaks if needed. Experiment session typically lasted about one hour and twenty minutes. In summary, the design was: 2 ORIENTATION × 2 DISTANCE × 2 TASK × 2 TABLE PRESENCE.

To mitigate ordering effects, we counterbalanced the DISTANCE order. In each DISTANCE block, we further counterbalanced the TASK order. This yielded eight possibilities (i.e. 2 TASK orders × 2 DISTANCE blocks × 2 DISTANCE orders). For the first 16 participants, we used all eight possibilities twice (i.e. once per ORIENTATION). For the last eight participants (i.e. four per ORIENTATION), we randomly selected four out of the eight existing possibilities, keeping the DISTANCE variable balanced (i.e. two participants started with *near* and two started with *far* per ORIENTATION). In each TASK, we randomized the TABLE PRESENCE order. The supplementary material contains a table of the per participant condition orders. We analyzed ordering effects (8 levels, i.e. DISTANCE x TASK orders, between-subject) and did not find a main effect of order on any of our dependent variables (all $p > .05$).

3.6 Participants

We recruited 24 participants via snowball sampling starting from university message groups and social networks. The participants had to be 18–70 years old, right-handed or ambidextrous. They should neither have had any COVID-19 symptoms nor have been in contact with confirmed cases in the previous 14 days. They should not have known physical disabilities nor feel strain in the arms.

In a pre-questionnaire, we asked participants to report their demographic information (age, gender, nationality), prior experience with VR technology and hand-tracking on a 5-point Likert scale (from 1–Never to 5–More than 20 times), gaming frequency (from 1–never or occasionally to 5–at least once a week), their level of alertness using the Stanford Sleepiness Scale [28] (from 0–Asleep to 7–Active, vital, alert, or wide awake), and their level of physical activity using the International Physical Activity Questionnaire - Short Form [12] (1–low, 2–moderate, 3–high; scored from 7 items). We report the participant demographics by our between-subject groups in Table 2. Of the 24 participants, 23 were right-handed and one was ambidextrous. They originated from 15 distinct countries. Participants received a small gift as gratuity for their time.

3.7 Measures

As dependent variables, we measured a range of quantities to capture participants' task performance, behavior, and subjective experience.
Selection Task Performance: We evaluate selection performance on the basis of *selection time*, *selection errors*, and *selection offset*. *Selection time* refers to the elapsed time between the start of a trial and when participants successfully selected the target. *Selection errors* counts the number of missed selections (i.e. when the participant's first point of contact with the task plane was greater than 2 cm from the target center). *Selection offset* is the distance between the participant's registered touch position and the center of the target on a successful selection.
Docking Task Performance: We evaluate docking performance on the basis of *dock time*, *drop count*, and *dock offset*. *Dock time* refers to the elapsed time between the start of a trial and when participants successfully docked the target. *Drop count* counts the number of drops. *Dock offset* is the distance between the participant's registered dock position and the actual center of the dock.
Self-reported Metrics: At the end of each condition, we ask participants to evaluate their *comfort*, *agency*, *exertion*, and *strain*. We used a custom Likert scale for eliciting participant feelings of *comfort*, *agency*, and *strain*. Participants evaluated *strain* on their neck and right shoulder, arm, and wrist/hands. We used the Borg CR10 scale for measuring *exertion*.
Behavioral Data: To characterize user behavior during conditions, we log the OptiTrack positions and orientations of the user's shoulder, arm, and wrist at a approximate rate of 90 frames per second. From raw OptiTrack positions and orientations, we compute three metrics: *upper arm movement*, *lower arm movement*, and *hand movement*. *Upper arm movement* refers to the accumulated angular difference between vectors pointing from the shoulder to the elbow. *Lower arm movement* refers to the accumulated angular difference between vectors pointing from the elbow to the wrist. *Hand movement* refers to the accumulated angular difference between the forward vector of the hand (i.e. from the bottom of the wrist to the knuckle of the middle finger).

4 RESULTS

We analyzed interval data using a mixed ANOVA. When the equal variances assumption was violated (Mauchly's test $p < .05$), we corrected the degrees of freedom using Greenhouse-Geisser. When the assumption about the normality of the residuals and homogeneity was violated (Shapiro-Wilk test $p < .05$), we either transformed the data using a log or square root function or analyzed them using the Aligned Rank Transform (ART) [60]. Ordinal data (questionnaire ratings) were analyzed using ART. For each data value, the USER was considered as a random factor, the ORIENTATION as a between-subject factor, and all the other independent variables (i.e., TABLE PRESENCE, DISTANCE, TASK) as within-subject factors. When needed, we performed pairwise post-hoc tests (Bonferroni with adjustment).The statistical analysis was performed using the R statistical software. For the sake of conciseness, we only report statistically significant main effects ($p < .05$) as well as interaction effects involving TABLE PRESENCE. We document all reported effects in Tables 3 and 4. We include additional graphs highlighting several interaction effects in our supplementary material.

4.1 Selection Task Performance

The ART analysis showed a main effect of TABLE PRESENCE ($F_{1,22} = 8.55$, $p < .01$, $\eta_p^2 = .11$) on *selection offset*. Participants

Table 3: Effect of TABLE PRESENCE in the *selection* task. The sample characteristics for the *no table* and *table* conditions across the 12 participants $\times$ 2 distances are summarized as M (SD). $\Delta Table$ represents the difference between the mean of *table* and the mean of *no table*. The p-values were obtained with pairwise post-hoc test comparisons (interaction effect between TABLE PRESENCE and ORIENTATION with Bonferroni correction). Only significant differences are presented. Significance levels: ***$p < .001$, **$p < .01$, *$p < .05$.

Variable	Horizontal				Vertical			
	no table	*table*	$\Delta Table$	*p*	*no table*	*table*	$\Delta Table$	*p*
Selection offset (mm)	5.79 (2.36)	4.82 (2.09)	-0.97	**	6.61 (2.58)	6.34 (2.48)	-0.27	
Physical Exertion	5.25 (2.36)	2.46 (1.5)	-2.79	***	4.88 (2.07)	3.33 (1.31)	-1.55	***
Comfort	3.5 (1.5)	5.29 (1.33)	1.79	***	3.63 (1.66)	4.96 (1.23)	1.33	***
Agency	4.25 (1.26)	5.71 (0.86)	1.46	***	5.38 (1.41)	5.38 (1.41)	0.00	
Shoulder strain	4.33 (1.74)	1.92 (0.97)	-2.41	***	5.29 (1.3)	4.25 (1.48)	-1.04	***
Elbow strain	3.38 (1.47)	1.83 (0.82)	-1.55	***	3.71 (1.76)	3.08 (1.74)	-0.63	
Wrist strain	3.13 (1.94)	2.04 (1.12)	-1.09	*	3.08 (1.35)	2.46 (1.06)	-0.62	
Upper arm movement (°)	3444.61 (1106.78)	3088.29 (928.86)	-356.32		3961.14 (1006.92)	3198.28 (1553.47)	-762.86	**
Lower arm movement (°)	6714.91 (2396.75)	5112.62 (1367.37)	-1602.29	***	6857.16 (1934.92)	6980.7 (2208.75)	123.54	

Table 4: Effect of TABLE PRESENCE in the *docking* task. The sample characteristics for the *no table* and *table* conditions across the 12 participants $\times$ 2 distances are summarized as M (SD). $\Delta Table$ represents the difference between the mean of *table* and the mean of *no table*. The p-values were obtained via pairwise post-hoc test comparisons (interaction effect between TABLE PRESENCE and ORIENTATION with Bonferroni correction). Only significant differences are presented. Significance levels: ***$p < .001$, **$p < .01$, *$p < .05$, $+p < .10$.

Variable	Horizontal				Vertical			
	no table	*table*	$\Delta Table$	*p*	*no table*	*table*	$\Delta Table$	*p*
Dock time (ms)	2332.8 (265.29)	1984.99 (228.9)	-347.81	***	2349.99 (460.18)	2226.5 (396.89)	-123.49	
Dock offset (mm)	1.89 (0.28)	1.26 (0.44)	-0.63	***	2.11 (0.27)	1.85 (0.25)	-0.26	***
Physical Exertion	5.13 (2.27)	2.46 (1.41)	-2.67	***	5.04 (2.16)	3.25 (1.62)	-1.79	***
Comfort	3.21 (1.72)	5.46 (1.1)	2.25	***	3.5 (1.5)	4.63 (1.47)	1.13	**
Agency	3.75 (1.23)	5.83 (0.76)	2.08	***	4.96 (1.37)	5.46 (1.29)	0.5	
Neck strain	2.29 (1.27)	1.83 (1.09)	-0.46	*	2.71 (1.52)	2.54 (1.35)	-0.17	
Shoulder strain	3.88 (1.54)	2.21 (0.88)	-1.67	***	5.42 (1.56)	3.54 (1.53)	-1.88	***
Elbow strain	3.63 (1.56)	1.83 (0.87)	-1.8	***	3.63 (1.86)	3 (2)	-0.63	
Wrist strain	3.04 (1.88)	1.92 (1.1)	-1.12	**	2.96 (1.37)	2.46 (1.44)	-0.5	
Upper arm movement (°)	2480.43 (507.74)	2714.86 (593.28)	234.43	+	2178.7 (402.64)	1602.52 (546.53)	-576.18	***
Hand movement (°)	3247.14 (1246.3)	3612.66 (1238.94)	365.52		4421.49 (2267.67)	5979.76 (2487.71)	1558.27	***

incurred a lower offset (i.e. more accurate) in *table* conditions ($M = 5.58\,mm, SD = 2.4\,mm$) than in *no table* conditions ($M = 6.2\,mm, SD = 2.5\,mm$). There was also a two-way marginal interaction effect between TABLE PRESENCE and ORIENTATION ($F_{1,22} = 3.76$, $p = .057$, $\eta_p^2 = .11$). Post-hoc tests showed that in *horizontal* conditions, participants were significantly more accurate with a *table* ($p < .01$). We did not observe significant effects on selection *time* and *error*.

4.2 Docking Task Performance

Dock time: We transformed *dock time* using a log function for the analysis. The ANOVA analysis showed a main effect of TABLE PRESENCE ($F_{1,22} = 22.98$, $p < .001$, $\eta_p^2 = .12$) and DISTANCE ($F_{1,22} = 4.67$, $p = .0'$, $\eta_p^2 = .03$). Participants were significantly faster in *table* conditions ($M = 734\,ms$, $SD = 152\,ms$) than in *no table* conditions ($M = 839\,ms$, $SD = 150\,ms$). They were also faster in *near* conditions ($M = 763\,ms$, $SD = 155\,ms$) than in *far* conditions ($M = 809\,ms$, $SD = 163\,ms$). There was also a two-way interaction effect between TABLE PRESENCE and ORIENTATION ($F_{1,22} = 6.11$, $p = .02$, $\eta_p^2 = .04$). Post-hoc tests showed that in *horizontal* conditions, participants were significantly faster with a *table* ($p < .001$).

Drop count: The ART analysis showed a main effect of ORIENTATION ($F_{1,22} = 11.26$, $p < .01$, $\eta_p^2 = .34$). Participants dropped the target significantly less frequently in *horizontal* conditions ($M = 1.3$, $SD = 1.8$) than in *vertical* conditions ($M = 3.79$, $SD = 3.6$). There was a two-way interaction effect between TABLE PRESENCE and ORIENTATION ($F_{1,22} = 5.26$, $p = .02$, $\eta_p^2 = .07$). Post-hoc tests did not reveal any significant differences.

Dock offset: The ANOVA showed a main effect of TABLE PRESENCE ($F_{1,22} = 81.31$, $p < .001$, $\eta_p^2 = .35$), ORIENTATION ($F_{1,22} = 14.48$, $p < .001$, $\eta_p^2 = .31$), and DISTANCE ($F_{1,22} = 10.57$, $p < .01$, $\eta_p^2 = .04$). Participants were significantly more accurate in *table* conditions ($M = 1.56\,mm$, $SD = .46\,mm$) than *no table* conditions ($M = 2\,mm$, $SD = .3\,mm$), in *horizontal* conditions ($M = 1.57\,mm$, $SD = .48\,mm$) than in *vertical* conditions ($M = 1.98\,mm$, $SD = .29\,mm$), and in *near* conditions ($M = 1.71\,mm$, $SD = .45\,mm$) than in *far* conditions ($M = 1.83\,mm$, $SD = .44\,mm$). There was also a two-way interaction effect between TABLE PRESENCE and ORIENTATION ($F_{1,22} = 13.24$, $p < .001$, $\eta_p^2 = .08$). Post-hoc tests showed that in both ORIENTATION conditions, participants were significantly more accurate with a *table* (all $p < .01$). Accuracy benefits of including a *table* were more pronounced in *horizontal* conditions than in *vertical* conditions.

4.3 Self-reported Metrics

The following analyses were performed using ART.

Physical exertion: The analysis showed a main effect of TABLE PRESENCE ($F_{1,22} = 158.7$, $p < .001$, $\eta_p^2 = .51$) and DISTANCE ($F_{1,22} = 22.93$, $p < .001$, $\eta_p^2 = .13$). Participants reported significantly lower *exertion* in *table* conditions (m (median) $= 3$) than *no table* conditions ($m = 5$). They also reported lower *exertion* in *near* conditions ($m = 3$) than in *far* conditions ($m = 4$). We additionally observed two-way interaction effects between TABLE PRESENCE and ORIENTATION ($F_{1,22} = 10.72$, $p < .01$, $\eta_p^2 = .07$) and between TABLE PRESENCE and DISTANCE ($F_{1,22} = 4.45$, $p = .04$, $\eta_p^2 = .03$). Post-hoc tests showed that in both ORIENTATION conditions, participants reported significantly lower exertion when a *table* was available (all $p < .01$). *table* effects on *exertion* were more pronounced in *horizontal* condition. Participants also reported significantly higher *exertion* when no *table* was available in *far* conditions ($p < .01$).

Comfort: The analysis showed a main effect of TABLE PRESENCE ($F_{1,22} = 110.88$, $p < .001$, $\eta_p^2 = .42$) and DISTANCE ($F_{1,22} = 4.55$, $p = .03$, $\eta_p^2 = .03$). Participants reported significantly higher comfort in *table* conditions ($m = 6$) compared to *no table* conditions ($m = 3$) and in *near* conditions ($m = 4$; $M = 4.48$) compared to *far* conditions ($m = 4$; $M = 4.06$). There was also a two-way interaction effect between TABLE PRESENCE and ORIENTATION ($F_{1,22} = 5.36$, $p = .02$, $\eta_p^2 = .03$). Post-hoc tests showed that in both ORIENTATION conditions participants reported significantly higher comfort when a *table* was available (all $p < .01$). Benefits of *table* on comfort were more pronounced in *horizontal* conditions than in *vertical* conditions.

Agency: The analysis showed a main effect of TABLE PRESENCE ($F_{1,22} = 79.96$, $p < .001$, $\eta_p^2 = .34$) and TASK ($F_{1,22} = 7.10$, $p < .01$, $\eta_p^2 = .04$). Participants reported significantly higher agency in *table* conditions ($m = 6$) than in *no table* conditions ($m = 5$). They also reported higher agency in *selection* conditions ($m = 6$) than in *docking* conditions ($m = 5$). There was also a two-way interaction effect between TABLE PRESENCE and ORIENTATION ($F_{1,22} = 28.78$, $p < .001$, $\eta_p^2 = .16$) and between TABLE PRESENCE and DISTANCE ($F_{1,22} = 5.44$, $p = .02$, $\eta_p^2 = .03$). Post-hoc tests showed that in both *far* and *horizontal* conditions participants reported significantly higher agency with a *table* (both $p < .001$).

Neck strain: The analysis showed a main effect of TABLE PRESENCE ($F_{1,22} = 4.02$, $p = .05$, $\eta_p^2 = .02$). Participants reported significantly less strain in their neck in *table* conditions ($m = 2$; $M = 2.22$) than in *no table* conditions ($m = 2$; $M = 2.4$).

Shoulder strain: The analysis showed a main effect of TABLE PRESENCE ($F_{1,22} = 130.96$, $p < .001$, $\eta_p^2 = .46$), ORIENTATION ($F_{1,22} = 13.90$, $p < .01$, $\eta_p^2 = .39$), and DISTANCE ($F_{1,22} = 5.17$, $p = .02$, $\eta_p^2 = .03$). Participants reported significantly less *strain* in their shoulder in *table* conditions ($m = 3$) compared to *no table* conditions ($m = 5$), in *horizontal* conditions ($m = 3$) compared to *vertical* conditions ($m = 5$), and in *near* conditions ($m = 4$; $M = 3.74$) compared to *far* conditions ($m = 4$; $M = 4$).

Elbow strain: Our analysis showed a main effect of TABLE PRESENCE ($F_{1,22} = 55.47$, $p < .001$, $\eta_p^2 = .26$) and DISTANCE ($F_{1,22} = 6.30$, $p = .01$, $\eta_p^2 = .04$). Participants reported significantly less *strain* in their elbow in *table* conditions ($m = 2$) than in *no table* tasks ($m = 3$). They also reported less *strain* in their elbow in *near* conditions ($m = 3$; $M = 2.87$) compared to *far* conditions ($m = 3$; $M = 3.16$). We also found a two-way interaction between TABLE PRESENCE and ORIENTATION ($F_{1,22} = 12.64$, $p < .001$, $\eta_p^2 = .08$). Post-hoc tests showed that participants reported significantly lower elbow *strain* when a *table* was available in both ORIENTATION conditions (all $p < .05$). *table* lowered strain more in the *horizontal* than in the *vertical* condition.

Wrist strain: The analysis showed a main effect of TABLE PRESENCE ($F_{1,22} = 33.28$, $p < .001$, $\eta_p^2 = .18$). Participants reported significantly less strain in their wrist in *table* conditions ($m = 2$) than in *no table* conditions ($m = 3$).

Table preference: Results showed that participants preferred *table* ($N = 83/96$) over *no table* ($N = 13/96$). *Table* was most preferred in the *far* condition ($N = 45/48$) compared to in the *near* condition ($N = 38/48$). Preferences between *horizontal* ($N = 42/48$) and *vertical* ($N = 41/48$) and between *selection* task ($N = 43/48$) and *docking* task ($N = 40/48$) did not differ by much. Overall, it was unanimously preferred in the *vertical* × *far* × *docking* condition ($N = 12/12$), and least preferred in the *vertical* × *near* × *docking* condition ($N = 7/12$).

4.4 Behavioral Data

We transformed behavioral data using a square root function and analyzed them using an ANOVA.

Upper arm movement: The analysis showed a main effect of TABLE PRESENCE $F_{1,22} = 11.26$, $p < .01$, $\eta_p^2 = .05$ and TASK $F_{1,22} = 49.42$, $p < .001$, $\eta_p^2 = .30$. Participants moved their upper arm significantly less in *table* conditions $M = 2651^\circ$, $SD = 1163.36^\circ$ compared to *no table* conditions $M = 3016.22^\circ$, $SD = 1079.52^\circ$, and *docking* conditions $M = 2244.13^\circ$, $SD = 659.06^\circ$ compared to *selection* conditions $M = 3423.08^\circ$, $SD = 1203.96^\circ$. There were also two-way interaction effects between TABLE PRESENCE and ORIENTATION ($F_{1,22} = 9.12$, $p < .01$, $\eta_p^2 = .04$), TABLE PRESENCE and DISTANCE ($F_{1,22} = 4.46$, $p = .05$, $\eta_p^2 = .01$), and TABLE PRESENCE and TASK ($F_{1,22} = 6.79$, $p = .02$, $\eta_p^2 = .01$). Post-hoc tests showed that in *vertical* conditions, participants moved their upper arm significantly less when a *table* was available ($p < .001$). In both DISTANCE and TASK conditions, participants likewise moved their upper arm significantly less when *table* was available (all $p < .05$). The effect of TABLE PRESENCE was, however, more pronounced in *far* and *selection* conditions.

Lower arm movement: The analysis showed a main effect of TASK $F_{1,22} = 68.07$, $p < .001$, $\eta_p^2 = .41$. Participants moved their right lower arm significantly less in *docking* conditions $M = 4049.52^\circ$, $SD = 967.64^\circ$ than in *selection* conditions $M = 6416.35^\circ$, $SD = 2124.25^\circ$. There were also two-way interaction effects between TABLE PRESENCE and ORIENTATION $F_{1,22} = 6.07$, $p = .02$, $\eta_p^2 = .02$, and TABLE PRESENCE and TASK $F_{1,22} = 15.05$, $p < .001$, $\eta_p^2 = .02$. Post-hoc tests showed that in *horizontal* conditions, participants moved their right lower arm more when there was *no table* ($p < .01$). Similarly, in the *selection* condition, participants moved their right lower arm significantly less when a *table* was available ($p < .01$).

Hand movement: The analysis showed a main effect of TABLE PRESENCE $F_{1,22} = 8.12$, $p < .01$, $\eta_p^2 = .02$, ORIENTATION $F_{1,22} = 11.28$, $p < .01$, $\eta_p^2 = .21$ and TASK $F_{1,22} = 16.12$, $p < .001$, $\eta_p^2 = .08$. Participants moved their right hand significantly less in *table* conditions $M = 4728.39^\circ$, $SD = 2679.28^\circ$ compared to *no table* conditions $M = 5337.65^\circ$, $SD = 2779.18^\circ$, in *horizontal* conditions $M = 3840.83^\circ$, $SD = 1834.57^\circ$ compared to *vertical* conditions $M = 6225.2^\circ$, $SD = 2975.19^\circ$, and in *docking* conditions $M = 4315.26^\circ$, $SD = 2146.32^\circ$ compared to *selection* conditions $M = 5750.78^\circ$, $SD = 3072.8^\circ$. There was also a two-way interaction effect between TABLE PRESENCE and ORIENTATION ($F_{1,22} = 5.37$, $p = .03$, $\eta_p^2 = .01$). Post-hoc tests showed that in *vertical* conditions, participants moved their right hand significantly more when a *table* was available ($p < .01$).

5 DISCUSSION

Our results show that the presence of a physical tabletop surface influences interaction in VR. We can largely attribute the effects to how the presence of a table enabled users to adopt more ergonomic postures. This stresses the importance of considering ergonomic factors for VR interaction design, adding to previous studies that examined mid-air interaction without surfaces (e.g., [13, 45]) as well as elbow-anchored interactions [56]. We add to the ongoing discussion of ergonomic design for VR our study of interactions that *complement* the so-far almost exclusive investigations of mid-air interactions by investigating physical tabletop presence as a distinct factor. Overall, we find numerous benefits in performance and comfort across most of our tasks, highlighting the promise of physical surfaces in supporting prolonged interactions no matter whether the UI is placed on the table itself or hovering in mid-air above it. We now discuss our results in detail with regard to performance, participants' self-reports, our analysis of limb motions and joint rotations, as well UI placement.

Task performance

Our results show that the availability of a tabletop surface provided accuracy and speed benefits. In particular, participants performed selections at higher levels of accuracy when the UI was oriented horizontally and co-located with the table surface. Likewise, when participants performed docking tasks, the table surface generally reduced their docking offsets. We note that the presence of the surface increased docking accuracy in *both* UI orientations. In the

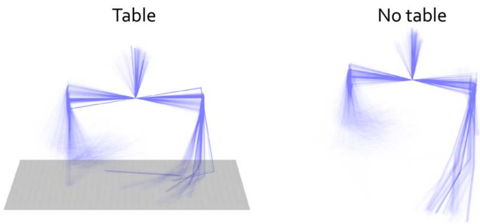

Figure 4: Postures for *vertical-near-docking* conditions. Participants always rested their dominant arm while only sometimes resting the elbow of their dominant arm. See supplemental material for the postures of all conditions.

horizontal condition, this increase in accuracy was further complemented with an increase in speed. The aforementioned findings on performance benefits of aligning UI with a physical surface corroborate prior research findings (e.g., [2, 38]). The accuracy and speed improvements we found in horizontal × table conditions can likely be attributed to the table's affordance for passive haptic feedback, which supports users in leveraging their senses of proprioception and kinesthesia [19, 25].

Interestingly, these improvements even extended to the vertical condition, where the UI is not aligned with the surface and, thus, participants do not receive any feedback on their input. We attribute the improvements in docking accuracy in vertical × table conditions to the table's affordance as a source of stability, as well as a (haptic) reference point [19, 59]. We conclude from our performance assessment that incorporating table surfaces into VR is particularly beneficial for tasks that require accuracy of input as well as those that span a prolonged duration.

Self-reports on exertion, strain, and comfort

Participants reported that the presence of a physical table helped reduce exertion and strain, as well as increase comfort when interacting in horizontal conditions. In vertical conditions, the reported benefits were less pronounced, reducing only exertion and specifically shoulder strain, while increasing comfort. In general, the suggested potential for tables to reduce exertion, strain, and discomfort speaks to its utility in supporting prolonged interaction. This also supports our initial assumption: while tables lose their practical need for storing items close-by in VR, the ergonomic benefits indicate that they may not be wholly obsolete. On the contrary, we argue that they may even be necessary depending on the required interactions, especially if VR environments are to redeem the promise of superseding current productivity work, especially for information workers. Participants' reports of reduced exertion, strain, and discomfort in vertical conditions show that tables can be beneficial even for interacting with *mid-air UIs* that offer no direct tactile feedback. Future 3D user interfaces may therefore benefit from placements not only co-located with table surfaces, but above or around.

Limb motion and joint rotation

Here, our analysis provided further support for the use of tabletops in productivity tasks. In the selection task, interactions with a table reduced lower arm movement in the horizontal condition and upper arm movement in the vertical condition. In the docking task, the effect of a table surface was only significant in the vertical condition, where the upper arm movement was reduced but hand movement was increased. A possible explanation for the differences in movement in horizontal conditions for selection is that the table enabled participants to slide their lower arm across the table and rely on moving their index up and down for executing the selection. In the vertical condition, the table enabled participants to plant their elbow on the surface as an anchor (Fig. 4). This effectively decreased upper arm movement, but required participants to rely more heavily on their wrists for executing the tasks. We believe that decreased movement could entail task execution in a more controlled and stabilized manner. Likewise, reduced movements in the upper-part of the kinematic chain could potentially contribute to differences in self-reported exertion, strain, and comfort.

UI placement

Our experiment varied UI placement to investigate the merit of a tabletop depending on the spatial context. Above, we discussed how the tabletop showed benefits for both horizontally and vertically oriented interfaces. In addition, the table was particularly appreciated by participants when interactions took place far away. A larger proportion of our participants preferred table in far conditions, even unanimously in vertical × far conditions. Furthermore, they reported that it considerably reduced exertion while increasing their sense of agency. Although guidelines typically suggest avoiding such UI placements, our results indicate that, if needed, tables could better support interactions with interface elements at a slight distance.

6 LIMITATIONS AND FUTURE WORK

Our study monitored participants' arm movements and touch events using optical markers. During calibration and piloting, we found that our apparatus accurately and reliably captured these inputs for task completion. However, while our apparatus tracked participants' arm pose, it did not explicitly verify whether their elbows made precise contact with the tabletop. To further evaluate the benefit of passive surfaces, we believe that instrumenting the surface with either a capacitive touch sensor [52] or force plate could yield more insights. For example, future research could use this data to analyze the weight participants place on their joints, contact positions, and their impact on performance.

In our experiment, we evaluated two tasks, two orientations, and two distances. While the evaluated range of interactions is substantial, future work remains needed to evaluate the utility of a table surface for different VR interactions. Future work could, for instance, investigate variables including target size, interaction scale, accuracy requirements, and task dimensionality (i.e. 2D or 3D).

Lastly, the insights from our present work aim to contribute to the design of future *situated* VR experiences, which our results show should take into consideration an understanding of users' natural postures when surfaces are available. Future work may build upon our findings to inform designs of interfaces that *opportunistically co-opt* table surfaces, as well as walls, within immersive environments. We are particularly excited about how such future work could contribute to the development of an adaptive input model for VR that fully leverages the passive affordances of the physical environment.

7 CONCLUSION

We presented the results of a 24-participant empirical study investigating the effects of incorporating a physical table surface into VR on user interaction performance, behavior, and subjective experience. We observed that table presence had benefits for task accuracy, task speed, agency, ergonomics, and comfort.

Our results suggest that incorporating a physical table surface into VR may enable more prolonged interaction, allowing VR interfaces to embrace situated interaction in mobile scenarios that promote comfort of use. Even if the task itself is in mid-air, we found that the presence of physical surfaces benefit interaction with UIs that are aligned with the passive surface as well as UIs in mid-air above the surface, quantifying and discussing their impact. We believe that our insights will support future work on leveraging physical constraints to enhance immersive interaction and, more broadly, enable and inform situated VR experiences in physical reality.

REFERENCES

[1] F. Argelaguet, L. Hoyet, M. Trico, and A. Lecuyer. The role of interaction in virtual embodiment: Effects of the virtual hand representation. In *2016 IEEE Virtual Reality (VR)*, pp. 3–10, 2016. doi: 10.1109/VR.2016.7504682

[2] R. Arora, R. H. Kazi, F. Anderson, T. Grossman, K. Singh, and G. Fitzmaurice. Experimental evaluation of sketching on surfaces in vr. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, p. 5643–5654. ACM, USA, 2017. doi: 10.1145/3025453.3025474

[3] M. Bachynskyi, G. Palmas, A. Oulasvirta, J. Steimle, and T. Weinkauf. Performance and ergonomics of touch surfaces: A comparative study using biomechanical simulation. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, p. 1817–1826. ACM, USA, 2015. doi: 10.1145/2702123.2702607

[4] H. Baker. How to set up your desk in vr on oculus quest 2, 2021.

[5] G. Borg. Psychophysical scaling with applications in physical work and the perception of exertion. *Scandinavian Journal of Work, Environment & Health*, 16:55–58, 1990.

[6] G. Bruder, F. Steinicke, and W. Stuerzlinger. Touching the void revisited: Analyses of touch behavior on and above tabletop surfaces. In P. Kotzé, G. Marsden, G. Lindgaard, J. Wesson, and M. Winckler, eds., *Human-Computer Interaction – INTERACT 2013*, pp. 278–296. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.

[7] G. Bruder, F. Steinicke, and W. Sturzlinger. To touch or not to touch? comparing 2d touch and 3d mid-air interaction on stereoscopic tabletop surfaces. In *Proceedings of the 1st Symposium on Spatial User Interaction*, SUI '13, p. 9–16. ACM, USA, 2013. doi: 10.1145/2491367.2491369

[8] L.-P. Cheng, E. Ofek, C. Holz, H. Benko, and A. D. Wilson. Sparse haptic proxy: Touch feedback in virtual environments using a general passive prop. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, p. 3718–3728. /ation for Computing Machinery, New York, NY, USA, 2017. doi: 10.1145/3025453.3025753

[9] L.-P. Cheng, E. Ofek, C. Holz, and A. D. Wilson. VRoamer: Generating On-The-Fly VR Experiences While Walking inside Large, Unknown Real-World Building Environments. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 359–366, 2019. doi: 10.1109/VR.2019.8798074

[10] Y. Cheng, Y. Yan, X. Yi, Y. Shi, and D. Lindlbauer. Semanticadapt: Optimization-based adaptation of mixed reality layouts leveraging virtual-physical semantic connections. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, UIST '21, p. 282–297. ACM, USA, 2021. doi: 10.1145/3472749.3474750

[11] A. Cockburn, D. Ahlström, and C. Gutwin. Understanding performance in touch selections: Tap, drag and radial pointing drag with finger, stylus and mouse. *International Journal of Human-Computer Studies*, 70(3):218–233, 2012. doi: 10.1016/j.ijhcs.2011.11.002

[12] C. Craig, A. Marshall, M. Sjostrom, A. Bauman, P. Lee, D. Macfarlane, T. Lam, and S. Stewart. International physical activity questionnaire-short form. *J Am Coll Health*, 65(7):492–501, 2017.

[13] J. a. M. Evangelista Belo, A. M. Feit, T. Feuchtner, and K. Grønbæk. Xrgonomics: Facilitating the creation of ergonomic 3d interfaces. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21. ACM, USA, 2021. doi: 10.1145/3411764.3445349

[14] S. Feiner, B. MacIntyre, M. Haupt, and E. Solomon. Windows on the world: 2d windows for 3d augmented reality. In *Proceedings of the 6th Annual ACM Symposium on User Interface Software and Technology*, UIST '93, p. 145–155. ACM, USA, 1993. doi: 10.1145/168642.168657

[15] A. Fender, P. Herholz, M. Alexa, and J. Müller. Optispace: Automated placement of interactive 3d projection mapping content. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, p. 1–11. ACM, USA, 2018. doi: 10.1145/3173574.3173843

[16] C. Forlines, D. Wigdor, C. Shen, and R. Balakrishnan. Direct-touch vs. mouse input for tabletop displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, p. 647–656. ACM, USA, 2007. doi: 10.1145/1240624.1240726

[17] R. Gal, L. Shapira, E. Ofek, and P. Kohli. Flare: Fast layout for augmented reality applications. In *2014 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 207–212, 2014. doi: 10.1109/ISMAR.2014.6948429

[18] S. Greenberg, N. Marquardt, T. Ballendat, R. Diaz-Marino, and M. Wang. Proxemic interactions: The new ubicomp? *Interactions*, 18(1):42–50, jan 2011. doi: 10.1145/1897239.1897250

[19] S. Gustafson, C. Holz, and P. Baudisch. Imaginary phone: Learning imaginary interfaces by transferring spatial memory from a familiar device. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, UIST '11, p. 283–292. ACM, USA, 2011. doi: 10.1145/2047196.2047233

[20] M. Hancock, O. Hilliges, C. Collins, D. Baur, and S. Carpendale. Exploring tangible and direct touch interfaces for manipulating 2d and 3d information on a digital table. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*, ITS '09, p. 77–84. ACM, USA, 2009. doi: 10.1145/1731903.1731921

[21] S. G. Hart and L. E. Staveland. Development of nasa-tlx (task load index): Results of empirical and theoretical research. In P. A. Hancock and N. Meshkati, eds., *Human Mental Workload*, vol. 52 of *Advances in Psychology*, pp. 139–183. North-Holland, 1988. doi: 10.1016/S0166-4115(08)62386-9

[22] J. Hartmann, C. Holz, E. Ofek, and A. D. Wilson. Realitycheck: Blending virtual environments with situated physical reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, p. 1–12. ACM, USA, 2019. doi: 10.1145/3290605.3300577

[23] S. Henderson and S. Feiner. Opportunistic tangible user interfaces for augmented reality. *IEEE Transactions on Visualization and Computer Graphics*, 16(1):4–16, 2010. doi: 10.1109/TVCG.2009.91

[24] J. D. Hincapié-Ramos, X. Guo, P. Moghadasian, and P. Irani. Consumed endurance: A metric to quantify arm fatigue of mid-air interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, p. 1063–1072. ACM, USA, 2014. doi: 10.1145/2556288.2557130

[25] K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell. Passive real-world interface props for neurosurgical visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '94, p. 452–458. ACM, USA, 1994. doi: 10.1145/191666.191821

[26] K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell. A survey of design issues in spatial input. In *Proceedings of the 7th Annual ACM Symposium on User Interface Software and Technology*, UIST '94, p. 213–222. ACM, USA, 1994. doi: 10.1145/192426.192501

[27] D. Hix, J. Swan, J. Gabbard, M. McGee, J. Durbin, and T. King. User-centered design and evaluation of a real-time battlefield visualization virtual environment. In *Proceedings IEEE Virtual Reality (Cat. No. 99CB36316)*, pp. 96–103, 1999. doi: 10.1109/VR.1999.756939

[28] E. Hoddes, V. Zarcone, and W. Dement. Stanford sleepiness scale. *Enzyklopädie der Schlafmedizin*, 1184, 1972.

[29] C. Holz and P. Baudisch. Understanding touch. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, p. 2501–2510. ACM, USA, 2011. doi: 10.1145/1978942.1979308

[30] B. E. Insko. *Passive haptics significantly enhances virtual environments*. PhD thesis, 2001. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2021-10-26.

[31] M. R. Jakobsen, Y. Jansen, S. Boring, and K. Hornbæk. Should i stay or should i go? selecting between touch and mid-air gestures for large-display interaction. In J. Abascal, S. Barbosa, M. Fetter, T. Gross, P. Palanque, and M. Winckler, eds., *Human-Computer Interaction – INTERACT 2015*, pp. 455–473. Springer International Publishing, Cham, 2015.

[32] S. Jang, W. Stuerzlinger, S. Ambike, and K. Ramani. Modeling cumulative arm fatigue in mid-air interaction based on perceived exertion and kinetics of arm motion. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, p. 3328–3339. ACM, USA, 2017. doi: 10.1145/3025453.3025523

[33] J. Jankowski and M. Hachet. Advances in interaction with 3d environments. *Computer Graphics Forum*, 34(1):152–190, 2015. doi: 10.1111/cgf.12466

[34] N. Joshi and D. Vogel. An evaluation of touch input at the edge of a

table. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, p. 1–12. ACM, USA, 2019. doi: 10.1145/3290605.3300476

[35] F. Kern, P. Kullmann, E. Ganal, K. Korwisi, R. Stingl, F. Niebling, and M. E. Latoschik. Off-the-shelf stylus: Using xr devices for handwriting and sketching on physically aligned virtual surfaces. *Frontiers in Virtual Reality*, 2, 2021. doi: 10.3389/frvir.2021.684498

[36] P. Koutsabasis and P. Vogiatzidakis. Empirical research in mid-air interaction: A systematic review. *International Journal of Human–Computer Interaction*, 35(18):1747–1768, 2019. doi: 10.1080/10447318.2019.1572352

[37] J. H. Lee, S.-G. An, Y. Kim, and S.-H. Bae. Projective windows: Bringing windows in space to the fingertip. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, p. 1–8. ACM, USA, 2018. doi: 10.1145/3173574.3173792

[38] R. W. Lindeman, J. L. Sibert, and J. K. Hahn. Towards usable vr: An empirical study of user interfaces for immersive virtual environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '99, p. 64–71. ACM, USA, 1999. doi: 10.1145/302979.302995

[39] M. Liu, M. Nancel, and D. Vogel. Gunslinger: Subtle arms-down mid-air interaction. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, UIST '15, p. 63–71. ACM, USA, 2015. doi: 10.1145/2807442.2807489

[40] Y. Luo and D. Vogel. Crossing-based selection with direct touch input. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, p. 2627–2636. ACM, USA, 2014. doi: 10.1145/2556288.2557397

[41] S. Mayer, V. Schwind, R. Schweigert, and N. Henze. The effect of offset correction and cursor on mid-air pointing in real and virtual environments. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, p. 1–13. ACM, USA, 2018. doi: 10.1145/3173574.3174227

[42] L. McAtamney and E. Nigel Corlett. Rula: a survey method for the investigation of work-related upper limb disorders. *Applied Ergonomics*, 24(2):91–99, 1993. doi: 10.1016/0003-6870(93)90080-S

[43] M. Meier, P. Streli, A. Fender, and C. Holz. TapID: Rapid Touch Interaction in Virtual Reality using Wearable Sensing. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*, pp. 519–528, 2021. doi: 10.1109/VR50410.2021.00076

[44] Meta. Horizon workrooms, 2022.

[45] R. A. Montano Murillo, S. Subramanian, and D. Martinez Plasencia. Erg-o: Ergonomic optimization of immersive virtual environments. In *Proceedings of ACM UIST 2017*, UIST '17, p. 759–771. ACM, USA, 2017. doi: 10.1145/3126594.3126605

[46] U. Obeysekare, C. Williams, J. Durbin, L. Rosenblum, R. Rosenberg, F. Grinstein, R. Ramamurti, A. Landsberg, and W. Sandberg. Virtual workbench-a non-immersive virtual environment for visualizing and interacting with 3d objects for scientific visualization. In *Proceedings of Seventh Annual IEEE Visualization'96*, pp. 345–349. IEEE, 1996.

[47] OptiTrack. Unity plugin, 2022.

[48] V. Schwind, S. Mayer, A. Comeau-Vermeersch, R. Schweigert, and N. Henze. Up to the finger tip: The effect of avatars on mid-air pointing accuracy in virtual reality. In *Proceedings of the 2018 Annual Symposium on Computer-Human Interaction in Play*, CHI PLAY '18, p. 477–488. ACM, USA, 2018. doi: 10.1145/3242671.3242675

[49] C. Shen, K. Ryall, C. Forlines, A. Esenther, F. D. Vernier, K. Everitt, M. Wu, D. Wigdor, M. R. Morris, M. Hancock, and E. Tse. Informing the design of direct-touch tabletops. *IEEE Computer Graphics and Applications*, 26(5):36–46, 2006. doi: 10.1109/MCG.2006.109

[50] A. L. Simeone, E. Velloso, and H. Gellersen. Substitutional reality: Using the physical environment to design virtual reality experiences. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, p. 3307–3316. ACM, USA, 2015. doi: 10.1145/2702123.2702389

[51] R. W. Soukoreff and I. S. MacKenzie. Towards a standard for pointing device evaluation, perspectives on 27 years of fitts' law research in hci. *International Journal of Human-Computer Studies*, 61(6):751–789, 2004. Fitts' law 50 years later: applications and contributions from human-computer interaction. doi: 10.1016/j.ijhcs.2004.09.001

[52] P. Streli, J. Jiang, A. R. Fender, M. Meier, H. Romat, and C. Holz. TapType: Ten-finger text entry on everyday surfaces via Bayesian inference. In *CHI Conference on Human Factors in Computing Systems*, pp. 1–16, 2022.

[53] W. Stuerzlinger and C. A. Wingrave. *The Value of Constraints for 3D User Interfaces*, pp. 203–223. Springer Vienna, Vienna, 2011. doi: 10.1007/978-3-211-99178-7_11

[54] H. B. Surale, A. Gupta, M. Hancock, and D. Vogel. Tabletinvr: Exploring the design space for using a multi-touch tablet in virtual reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, p. 1–13. ACM, USA, 2019. doi: 10.1145/3290605.3300243

[55] R. J. Teather, D. Natapov, and M. Jenkin. Evaluating haptic feedback in virtual environments using iso 9241–9. In *2010 IEEE Virtual Reality Conference (VR)*, pp. 307–308, 2010. doi: 10.1109/VR.2010.5444753

[56] R. Veras, G. Singh, F. Farhadi-Niaki, R. Udhani, P. P. Patekar, W. Zhou, P. Irani, and W. Li. Elbow-anchored interaction: Designing restful mid-air input. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21. ACM, USA, 2021. doi: 10.1145/3411764.3445546

[57] R. Viciana-Abad, A. R. Lecuona, and M. Poyade. The influence of passive haptic feedback and difference interaction metaphors on presence and task performance. *Presence*, 19(3):197–212, 2010. doi: 10.1162/pres.19.3.197

[58] P. Wang, X. Bai, M. Billinghurst, S. Zhang, D. Han, M. Sun, Z. Wang, H. Lv, and S. Han. Haptic feedback helps me? a vr-sar remote collaborative system with tangible interaction. *International Journal of Human–Computer Interaction*, 36(13):1242–1257, 2020. doi: 10.1080/10447318.2020.1732140

[59] Y. Wang and C. L. MacKenzie. The role of contextual haptic and visual constraints on object manipulation in virtual environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '00, p. 532–539. ACM, USA, 2000. doi: 10.1145/332040.332494

[60] J. O. Wobbrock, L. Findlater, D. Gergle, and J. J. Higgins. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 143–146, 2011.

[61] R. Xiao, J. Schwarz, N. Throm, A. D. Wilson, and H. Benko. Mrtouch: Adding touch input to head-mounted mixed reality. *IEEE Transactions on Visualization and Computer Graphics*, 24(4):1653–1660, 2018. doi: 10.1109/TVCG.2018.2794222

[62] J. J. Yang, C. Holz, E. Ofek, and A. D. Wilson. Dreamwalker: Substituting real-world walking experiences with a virtual reality. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, UIST '19, p. 1093–1107. ACM, USA, 2019. doi: 10.1145/3332165.3347875

[63] X. Zhang and I. S. MacKenzie. Evaluating eye tracking with iso 9241 - part 9. In J. A. Jacko, ed., *Human-Computer Interaction. HCI Intelligent Multimodal Interaction Environments*, pp. 779–788. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007.

[64] D. Zielasko, M. Krüger, B. Weyers, and T. W. Kuhlen. Menus on the desk? system control in deskvr. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 1287–1288, 2019. doi: 10.1109/VR.2019.8797900

[65] D. Zielasko, M. Krüger, B. Weyers, and T. W. Kuhlen. Passive haptic menus for desk-based and hmd-projected virtual reality. In *2019 IEEE 5th Workshop on Everyday Virtual Reality (WEVR)*, pp. 1–6, 2019. doi: 10.1109/WEVR.2019.8809589

[66] D. Zielasko, B. Weyers, and T. W. Kuhlen. A non-stationary office desk substitution for desk-based and hmd-projected virtual reality. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 1884–1889, 2019. doi: 10.1109/VR.2019.8797837