



Swarm Manipulation in Virtual Reality

Xiang Li
xl529@cam.ac.uk
University of Cambridge
Cambridge, United Kingdom

John J. Dudley
jld50@cam.ac.uk
University of Cambridge
Cambridge, United Kingdom

Jin-Du Wang
jinduwang@hotmail.com
Xi'an Jiaotong University
Xi'an, China

Per Ola Kristensson
pok21@cam.ac.uk
University of Cambridge
Cambridge, United Kingdom




Figure 1: We introduce the Swarm Manipulation technique in Virtual Reality (VR) and compare it with two conventional manipulation techniques: Hand and Controller (Ray-Casting). We evaluate these techniques based on three tasks: (1) Selection, (2) Rotation, and (3) Resizing. During object selection, the swarm particles transition from their original color to blue.

ABSTRACT

The theory of swarm control shows promise for controlling multiple objects, however, its scalability is limited by costs, such as hardware and infrastructure needs. Virtual Reality (VR) can overcome these limitations, but research on swarm interaction in VR is limited. This paper introduces a novel Swarm Manipulation interaction technique and compares it with two baseline techniques: Virtual Hand and Controller (ray-casting). We evaluated these techniques in a user study ($N = 12$) in three tasks (selection, rotation, and resizing) across five conditions. Overall, our results show that the swarm manipulation technique did result in good performance, exhibiting significantly faster speeds compared with at least one of the other two techniques across most conditions for these three tasks. Furthermore, this technique notably reduced resizing size deviations in the resizing task. However, we also observed a trade-off between speed and accuracy in the rotation task. The results demonstrate the potential of the Swarm Manipulation technique to enhance the usability and user experience in VR compared to conventional manipulation techniques. In future studies, we aim

to understand and improve swarm interaction, user control, and internal particle cooperation.

CCS CONCEPTS

• Human-centered computing → Gestural input; Virtual reality; Empirical studies in HCI.

KEYWORDS

Swarm interaction, swarm manipulation, virtual reality

ACM Reference Format:

Xiang Li, Jin-Du Wang, John J. Dudley, and Per Ola Kristensson. 2023. Swarm Manipulation in Virtual Reality. In *The 2023 ACM Symposium on Spatial User Interaction (SUI '23)*, October 13–15, 2023, Sydney, NSW, Australia. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3607822.3614519>

1 INTRODUCTION

The realm of robotics and drones has seen the effective implementation of swarm control theory [16, 27, 44, 56]. Swarm control enables the coordination of numerous entities as a collective to achieve a common goal, promoting efficient simultaneous selection and manipulation of multiple objects [38, 39, 51]. The benefits are substantial, including scalability, fault tolerance, and robustness [9], which allows systems to adapt to varying numbers of entities, while fault tolerance ensures system functionality despite individual agent failures [8]. Swarm systems also exhibit dynamic adaptability, adjusting collectively to evolving conditions [9]. Despite these advantages, there are hindrances to the widespread implementation of swarm control, such as high costs associated



This work is licensed under a Creative Commons Attribution International 4.0 License.

SUI '23, October 13–15, 2023, Sydney, NSW, Australia
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0281-5/23/10.
<https://doi.org/10.1145/3607822.3614519>

with hardware, communication infrastructure, and computational resources, as well as the need for significant maintenance and expertise [19]. However, within the virtual reality (VR) environment, many of these limitations of swarm control implementation can potentially be mitigated. This makes the prospect of swarm interactions in VR a compelling area of exploration, with the potential to allow us to reimagine the way we interact with virtual objects [24].

In addition, exploring novel manipulation techniques and providing a simulated environment for intuitive and immersive interactions with virtual objects has emerged as a promising direction for improving user experience in VR [3]. Traditional techniques of object selection and manipulation in VR rely on avatar hand movements, which mimic the movements of the user's real hands, or controller-based ray-casting for object selection and allow for adjustment of movement gain to reach distant objects [4, 18, 33]. For instance, users have the flexibility to utilize both hand controllers for near and far field interactions on the Quest. Similarly, HoloLens incorporates a ray-casting technique from the hand to facilitate far-field interaction. While these methods are practical, there is a trade-off between intuitiveness and efficiency. An emergent manipulation technique is *Ninja Hand* [45], where the user's real hands in VR are mapped to multiple virtual hands. This approach holds promise in reducing time spent moving and decreasing overall workload. However, it comes with its own challenges, such as causing visual clutter due to the presence of numerous hands in the virtual space, which can limit and even reverse its benefits with increasing hand count [45].

This paper explores an alternative way of disassociating the virtual from the physical by exploring a novel Swarm Manipulation interaction technique (see Figure 1) to bestow users with the ability to select and manipulate multiple virtual objects simultaneously. To realize this idea we present the results of a user study ($N = 12$) that aims to provide us with an indication of the effectiveness of a Swarm Manipulation technique compared to other popular manipulation techniques in VR. Our study focuses on evaluating three manipulation techniques: Virtual Hand, Controller (ray-casting), and Swarm Manipulation, which are assessed in three tasks: selection, rotation, and resizing, across five conditions: single-target close-distance, single-target long-distance, dual-target close-distance, dual-target long-distance, and all-target. We aim to present empirical findings and envision the potential future of object manipulation in VR, fostering an understanding of how best to utilize the powerful capabilities of swarm control in this dynamic and immersive field. The contributions of this paper can be summarized as follows:

- To our knowledge, this is the first exploratory user study using the proof-of-concept swarm interaction in VR. We provide an in-depth analysis of the effectiveness and user experience of the Swarm Manipulation technique compared to the other two baseline techniques in VR: Hand and Controller.
- We investigate the advantages and limitations of the Swarm Manipulation technique in terms of task performance (i.e., the task completion time and deviation), perceived workload, usability, and novelty.

- Our study results offer valuable insights for developers and designers seeking to integrate Swarm Manipulation techniques into future VR applications. Furthermore, we envision the potential future developments and possibilities of swarm interactions.

2 RELATED WORK

In this section, we review the existing literature and research pertaining to object manipulation in VR and swarm interaction systems. The goal is to identify the gaps and opportunities for integrating swarm interactions into manipulation tasks in VR.

2.1 Object Manipulation in VR

Object manipulation in VR involves interacting with and manipulating virtual objects within a simulated environment [12, 34, 52]. Numerous techniques and approaches have been explored to enhance the object manipulation experience in VR [18, 33, 54, 55]. For example, Virtual Hand, a widely used mid-air interaction paradigm in modern VR systems, allows users to manipulate objects in virtual environments [33, 53]. To enhance this interaction technique, several approaches have been developed. One such approach is *Go-Go* [42] and its extensions [13, 35], which enables users to reach distant targets by using a non-linear mapping between the controlled motion and the effected motion [4]. Additionally, techniques like Ray-Casting and scaling down the virtual world have been employed to interact with out-of-reach objects [15, 41, 50]. Dewez et al. developed “avatar-friendly” manipulation techniques [10] to strengthen the bond between users and their virtual avatars. The researchers argue that a more intuitive and efficient way of manipulating objects in VR can lead to more immersive experiences. A decade earlier, Slater et al. [48] conducted a study on the body transfer phenomenon in VR, demonstrating that if users can manipulate a virtual body as their own, it significantly enhances the sense of presence and embodiment in VR. This research highlights the critical role of efficient object manipulation in creating immersive VR experiences.

2.2 Swarm Interactions

Swarm interactions involve studying human-swarm interaction (HSI) and identifying fundamental principles and invariants. Brown et al. [7] propose two invariants for geometric-based swarms: the collective state and the balance between span and persistence. Brown et al. [8] emphasize managing attractors to individual abstract agents and focus on the collective behavior. Kolling et al. [27] survey human-swarm interaction, while Kolling et al. [26] compared intermittent and environmental interaction types. Dietz et al. [11] explored the human perception of swarm robot motion.

Swarm user interfaces also introduce innovative concepts and architectures for interactive systems. Nakagaki et al. presented *HERMITS* [38] and *(Dis)Appearables* [39], an architecture that enabled dynamic reconfiguration of self-propelled Tangible User Interfaces (TUIs) using mechanical shell add-ons and actuated these swarm TUIs to appear and disappear dynamically. Yu et al. [56] introduced *AeroRigUI*, an actuated TUI for 3D spatial interaction using controlled strings attached to ceiling surfaces. Le Goc et al. [28]

presented *Zooids*, an open-source platform for swarm user interfaces, while Suzuki et al. [51] explored *Reactile*, an approach to programming swarm user interfaces through direct physical manipulation. These studies offer valuable insights into the potential and design considerations of swarm user interfaces, contributing to advancements in interactive systems and human-robot interaction.

2.3 Mapping from One to Many

The mapping from one to many in VR has been explored in previous research. Some studies have suggested using multiple limbs or additional fingers to enhance the VR experience, focusing on aspects like body acceptance and ownership [21]. However, it remains unclear whether these additions improve performance in interactive tasks. *Ninja Cursors* [25] addresses this issue by mapping input from a single mouse to multiple virtual cursors distributed across a desktop display. This approach improves target acquisition efficiency for large 2D displays. Only one cursor can actively hover over a target at a time, while the others wait in a queue. Gaze tracking has also been incorporated into this technique, as seen in the rake cursor and the work by R ih a and  pakov [43], to choose which cursor is active. Lubos et al. [30] used head tracking in VR to disambiguate between two sets of virtual hands but did not investigate the impact of manipulating the number of hands on shortest-distance gains. *Ninja Hand* [45] maps the user’s real hand to multiple virtual hands in VR, and this approach holds promise in terms of reduced movement times and lower overall workload.

However, each of these methods visually alters the basic shape of the body, and the presence of a large number of cursors or hands requires users to adapt and learn how to use these techniques, resulting in an increased workload [29, 45]. This can potentially make these methods less user-friendly and more challenging to use efficiently.

3 SWARM MANIPULATION




Figure 2: The figure illustrates the Swarm Hand concept consisting of two key components. The first component, the Swarm Hand, is composed of a swarm of particles, controlled by users through hand gestures. The second component, the non-dominant hand depicted as the left hand (L) in the figure, enables adjustment of distribution levels. This feature determines the size of the Swarm Hand and allows users to grasp or select multiple virtual objects based on their preferences.

In this paper, we present *Swarm Manipulation*, a novel interaction technique for object manipulation in virtual environments. This technique consists of two main components, each serving a specific purpose. The first component is the dominant Swarm Hand, represented by the right “hand” in Figure 2. This hand is composed of a swarm of particles that users can control through discrete hand movements. The Swarm Hand incorporates the Go-Go technique [42], which enables non-linear mapping between the user’s hand movements and the behavior of the swarm. This technique allows for intuitive and flexible manipulation of virtual objects within the VR environment. By leveraging discrete hand movements, users can easily navigate and interact with the swarm to perform various tasks and actions.

The second component of the Swarm Manipulation technique is the non-dominant hand, depicted by the left hand in Figure 2. This hand plays a crucial role in adjusting the level of distribution within the swarm. The distribution level defines the volume of the swarm hand, determining how many proximate objects it can grasp. Users can modify the swarm’s distribution level by rotating their wrist [49], providing a simple and intuitive means to control the graspability of objects within the virtual environment. To provide visual feedback and facilitate interaction, a semicircular panel surrounding the user’s wrist displays the current level of swarm distribution. This visual indicator allows users to perceive and monitor the volume and graspability of the swarm hand, aiding them in making informed manipulation decisions.

4 USER STUDY

We carried out a user study to investigate the effectiveness of the Swarm Manipulation technique compared to other popular manipulation techniques in VR. The study aimed to evaluate the performance and user experience of three manipulation techniques: Virtual Hand, Controller (ray-casting), and Swarm Manipulation. The primary objective was to examine the advantages and limitations of the Swarm Manipulation technique compared to the other techniques. This user study contributes to the existing knowledge of manipulation techniques in VR, specifically investigating the effectiveness of the Swarm Manipulation technique.

4.1 Study Design

The user study utilized a repeated measures factorial design with two independent variables: (a) *Manipulation Technique* (Hand, Controller, and Swarm), and (b) *Condition* (single-target close-distance, single-target long-distance, dual-target close-distance, dual-target long-distance, and all-target) (see Figure 3). The experiment consisted of three tasks: Selection, Rotation, and Resizing, as suggested by Bowman et al. [5]. We differentiate between the close-distance and long-distance conditions based on the overall range of interaction observed when the arm is bent (i.e., around 40 cm) and the complete extension of the arm when it is straight (i.e., around 70 cm) [53].

Dependent variables included: (a) *Perceived Workload*, measured by raw NASA-TLX [17], (b) *Usability*, measured by System Usability Scale [6], (c) *User Experience*, measured by UEQ-Short [46], and (d) *Task Performance*, measured by task-completion time and deviations in rotation and resizing. Additionally, participants were asked to

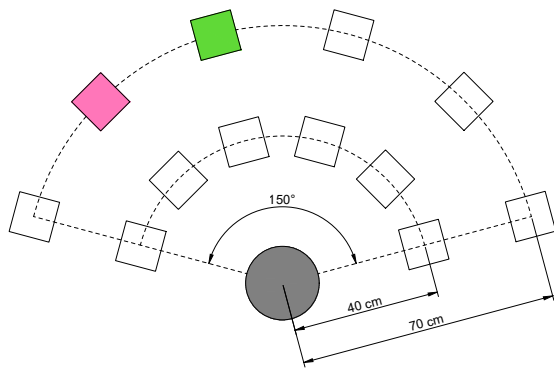


Figure 3: There are twelve objects that should be manipulated during the study. The selected objects are highlighted in red, while the target object is highlighted in green.

provide a final comparison and rate their preference and ease of use for the three manipulation techniques.

4.2 Procedure

After a brief introduction, participants were provided with a 5-minute period to familiarize themselves with the manipulation techniques and tasks that would be undertaken. To minimize the influence of learning effects, the order of manipulation techniques assigned to participants was counterbalanced using a Latin square design. All participants successfully completed the designated tasks utilizing each manipulation technique. The presentation of conditions within each assigned task was randomized, and each condition was repeated 10 times. Consequently, the overall study comprised a total of 5,400 trials, calculated as 3 (Technique) \times 3 (Task) \times 5 (Condition) \times 10 (Repeat) \times 12 (Participant).

After completing the tasks, participants were asked to fill out questionnaires evaluating their experience with the different manipulation techniques. Further, a final comparison was included, where participants were asked to rate their preference and the ease of use for each manipulation technique. Finally, a semi-structured interview was conducted to gather qualitative insights, allowing participants to share their strategies and provide suggestions. On average, participants spent approximately 30 minutes completing the study, including the tasks and questionnaire. Participants received a £5 reward for their participation.

4.3 Manipulation Techniques

Our study employed the following interactions for each manipulation technique:

4.3.1 Hand. Participants used Meta Quest hand tracking and could see their hand models in the VR environment. They touched the target object with their right hand for selection. In the Rotation task, the right hand was used to rotate the objects, with a grabbing gesture with the left hand confirming the action. The Resizing task followed the same procedure for selection, with scaling achieved by




Figure 4: A screenshot of the Swarm Manipulation technique in VR. It showcases the non-dominant hand with a distribution bar encircling the left wrist, while the dominant hand interacts with the swarm. In the field of view, twelve targets are positioned ahead.

changing the distance between the right index fingertip and thumb tip.

4.3.2 Controller. Participants used the Meta Quest 2 handheld controllers. The controller emitted a ray that could be manipulated in the VR environment. The selection was confirmed by pressing the pinch button on the controller when the ray pointed at the target object. The Rotation task involved the use of the ray and pinch buttons for object selection. The objects were rotated by rotating the controller, and the grip button was used for confirmation. The Resizing task followed a similar process for selection, with scaling achieved by dragging the ray outside or inside the objects.

4.3.3 Swarm Manipulation. Participants used the Meta Quest hand-tracking feature and could see their left-hand model, an indicator of Swarm Hand dispersion, and the Swarm Hand controlled by their right hand in the VR environment (see Figure 4). The Selection task involved touching the target objects with any particle of the Swarm Hand (see Figure 1 (1)). The Rotation task followed the same procedure for selection, with object rotation achieved by rotating the right hand (see Figure 1 (2)). The Resizing task followed the same selection procedure, with scaling achieved by changing the distance between the right index fingertip and thumb tip (see Figure 1 (3)).

4.4 Tasks

4.4.1 Selection. Participants used the assigned manipulation technique to select the target objects, which were no longer highlighted in red upon selection.

4.4.2 Rotation. Participants used the assigned manipulation technique to select target objects. A demonstration object appeared at the target angle (i.e., 45 degrees) once all target objects were selected. Participants were asked to rotate the selected target objects to match the demonstration object as closely as possible.

4.4.3 Resizing. Participants used the assigned manipulation technique to select the target objects. A demonstration object appeared at the target size (1.2 times or 0.8 times the original object size) once all target objects were selected. Participants were asked to scale the selected target objects to match the demonstration object as closely as possible.

4.5 Measures

4.5.1 Selection Task. (a) **Selection Task Completion Time:** The time taken by participants to complete each selection task, indicating the duration from target object generation to the selection of all target objects.

4.5.2 Rotation Task. (a) **Rotation Task Completion Time:** The time taken by participants to complete each rotation task, representing the duration from target object selection to the participant's perception of task completion.

(b) **Rotation Angle Deviation:** The deviation ($^{\circ}$) between the final angle of the target object and the target angle set by the demonstration object in rotation tasks. This measure considered various rotation directions and angles chosen by the participants to achieve visually similar results.

4.5.3 Resizing Task. (a) **Resizing Task Completion Time:** The time taken by participants to complete each resizing task, indicating the duration from target object selection to the participant's perception of task completion.

(b) **Resizing Size Deviation:** The absolute deviation between the current size of the target object and the target size set by the demonstration object in resizing tasks.

4.6 Participants

A total of 12 participants (8 males and 4 females) were recruited for the study. The age range of the participants was between 20 and 23 years ($M = 21.17, SD = .83$). All participants were students at a local university. All participants reported previous experience with VR, with familiarity ratings ranging from 1 to 6 on a 7-point Likert scale, where 1 indicated no experience in VR, and 7 indicated expertise ($M = 2.42, SD = 1.73$). All participants were right-handed habitual users.

4.7 Apparatus

The user study took place in a university laboratory equipped with a desktop computer, display devices, and an area for participants to engage in VR interactions using Meta Quest 2. The application used in the study was implemented in Unity 2021.3.23 and ran on a desktop computer.

5 RESULTS

5.1 Task 1: Selection

5.1.1 Task Completion Time. A two-way Analysis of Variance (ANOVA) was conducted to examine the main effects of technique and condition, as well as their interaction effect, on the task completion time. The main effect of the technique was significant with a partial Eta squared (η_p^2) of .067 ($p < .001$). The main effect of the condition was also significant with a η_p^2 of .334 ($p < .001$).




Figure 5: Mean task-completion times for each technique across five conditions of Task 1.

The interaction effect between the technique and condition was significant as well, with a η_p^2 of .039 ($p < .001$).

A post-hoc test using the Tukey HSD method was performed to make pairwise comparisons between the three techniques. In the close-distance single-target selection condition, Swarm was significantly faster than both Hand ($p < .001$) and Controller ($p = .001$). In the close-distance dual-target selection condition, Swarm was significantly faster than Hand ($p < .001$), but not significantly different from Controller ($p = .148$). In the long-distance single-target selection condition, Swarm was significantly faster than Hand ($p = .006$), but not significantly different from Controller ($p = .900$). In the long-distance dual-target selection condition, Swarm was significantly faster than Hand ($p < .001$), but not significantly different from Controller ($p = .898$). In the all selection condition, Swarm was significantly faster than both Hand ($p < .001$) and Controller ($p = .016$). The mean times for each technique across five conditions can be found in Figure 5.

5.2 Task 2: Rotation

5.2.1 Task Completion Time. A two-way ANOVA was conducted to examine the main effects of technique and condition, as well as their interaction effect, on the task completion time. The main effect of the technique was significant with a η_p^2 of .022 ($p < .001$). The main effect of the condition was also significant with a η_p^2 of .210 ($p < .001$). The interaction effect between the technique and condition was significant as well, with a η_p^2 of .036 ($p < .001$).

A post-hoc test using the Tukey HSD method was performed to make pairwise comparisons between the three techniques. In the close-distance single-target rotation condition, Swarm was significantly slower than Controller ($p = .001$), but there was no significant difference between Swarm and Hand ($p = .072$). In the close-distance dual-target rotation condition, Swarm was significantly slower than both Hand ($p = .044$) and Controller ($p = .001$). There were no significant differences between Swarm and the other two techniques in the long-distance single-target rotation and long-distance dual-target rotation conditions. In the all-rotation condition, Swarm was significantly faster than Hand ($p < .001$), but there was no significant difference between Swarm and Controller. The mean times for each technique across five conditions of Task 2 can be found in Figure 6.




Figure 6: Mean task-completion times for each technique across five conditions of Task 2.




Figure 7: Mean rotation angle deviations for each technique across five conditions of Task 2.

5.2.2 Rotation Angle Deviation. A two-way ANOVA was conducted to examine the main effects of technique and condition and their interaction effect on the rotation angle deviation. The main effect of the technique was significant with a η_p^2 of .011 ($p < .001$). The main effect of the condition was also significant with a η_p^2 of .014 ($p < .001$). The interaction effect between the technique and condition was significant as well, with a η_p^2 of .015 ($p = .001$).

A post-hoc test using the Tukey HSD method was performed to make pairwise comparisons between the three techniques. In the close-distance single-target rotation condition, Swarm showed a significantly larger difference in rotation angle deviation than Controller ($p = .001$). There was also a significant difference between Swarm and Hand ($p = .020$). In the close-distance dual-target rotation condition, there was a significant difference between Swarm and Hand ($p = .020$), but no significant difference between Swarm and Controller. There were no significant differences between Swarm and the other two techniques in the long-distance single-target rotation and long-distance dual-target rotation conditions. In the all-rotation condition, there were no significant differences between Swarm and the other two techniques in terms of rotation angle deviation. The mean rotation angle deviations for each technique across five conditions of Task 2 can be found in Figure 7.




Figure 8: Mean task-completion times for each technique across five conditions of Task 3.

5.3 Task 3: Resizing

5.3.1 Task Completion Time. A two-way ANOVA was conducted to examine the main effects of technique and condition, as well as their interaction effect, on the task completion time. The main effect of the technique was significant with a η_p^2 of .032 ($p < .001$). The main effect of the condition was also significant with a η_p^2 of .164 ($p < .001$). The interaction effect between the technique and condition was significant as well, with a η_p^2 of .024 ($p < .001$).

A post-hoc test using the Tukey HSD method was performed to make pairwise comparisons between the three techniques. In the close-distance single-target resizing condition, there were no significant differences in task completion times between Swarm and the other two techniques. In the close-distance dual-target resizing condition, Swarm was significantly faster than Controller ($p = .001$), but there was no significant difference between Swarm and Hand ($p = .044$). In the long-distance single-target resizing condition, Swarm was significantly faster than Controller ($p = .049$), but there was no significant difference between Swarm and Hand ($p = .072$). There were no significant differences in task completion times between Swarm and the other two techniques in the long-distance dual-target resizing and all-resizing conditions. In the all-resizing condition, the Swarm Manipulation technique was significantly faster than both the Hand and Controller techniques ($p < .001$). The mean times for each technique across five conditions of Task 3 can be found in Figure 8.

5.3.2 Resizing Size Deviation. A two-way ANOVA was conducted to examine the main effects of technique and condition and their interaction effect on the resizing size deviation. The main effect of the technique was significant with a η_p^2 of .008 ($p < .001$). The main effect of the condition was not significant ($p = .175$), and the interaction effect between the technique and condition was significant with a η_p^2 of .012 ($p = .0049$).

A post-hoc test using the Tukey HSD method was performed to make pairwise comparisons between the three techniques. In the close-distance single-target resizing condition, there were no significant differences between Swarm and the other two techniques. In the close-distance dual-target resizing condition, Swarm had significantly lower resizing size deviations than Controller ($p = .007$), but there was no significant difference between Swarm and Hand. In




Figure 9: Mean resizing size deviations for each technique across five conditions of Task 3.

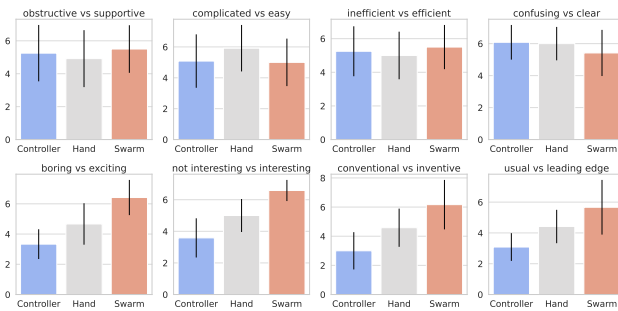


Figure 10: Bar chart illustrating the average ratings for eight UEQ dimensions across three techniques, with error bars representing standard deviation.

the long-distance single-target resizing condition, Swarm had significantly lower resizing size deviations than both Hand ($p = .001$) and Controller ($p < .001$). There were no significant differences between Swarm and the other two techniques in the long-distance dual-target resizing and all-resizing conditions. In the all-resizing condition, Swarm had significantly lower resizing size deviations than Controller ($p = .044$). However, there was no significant difference between Swarm and Hand. The mean resizing size deviations for each technique across five conditions of Task 3 can be found in Figure 9.

5.4 Ratings and Preferences

5.4.1 User Experience Questionnaire. A Friedman test was conducted to compare the UEQ scores among the Controller, Hand, and Swarm techniques (see Figure 10). Significant differences were found in the following items: “boring vs. exciting” ($\chi^2(2) = 20.83$, $p < .001$), “not interesting vs. interesting” ($\chi^2(2) = 15.83$, $p < .001$), “conventional vs. inventive” ($\chi^2(2) = 14.22$, $p < .001$), and “usual vs. leading edge” ($\chi^2(2) = 12.33$, $p = .002$). However, for the aspects “obstructive vs. supportive” ($\chi^2(2) = .35$, $p = .839$), “complicated vs. easy” ($\chi^2(2) = 1.95$, $p = .378$), “inefficient vs. efficient” ($\chi^2(2) = 1.72$, $p = .423$), and “confusing vs. clear” ($\chi^2(2) = 3.19$, $p = .203$), there were no significant differences between the techniques.

A post-hoc analysis was conducted using pairwise Wilcoxon Rank-Sum Test (Mann-Whitney U test) with Bonferroni correction for multiple comparisons. The post-hoc analysis showed significant differences in most comparisons across the three techniques for the “boring vs. exciting”, “not interesting vs. interesting”, “conventional vs. inventive”, and “usual vs. leading edge” items, except for “not interesting vs. interesting” between Controller and Hand techniques ($W = 5$, $p = .052$), and “usual vs. leading edge” between Hand and Swarm techniques ($W = 7$, $p = .103$).

5.4.2 NASA Task Load Index. A Friedman test was conducted to compare the effect of different techniques on six measures of the NASA Task Load Index: Mental Demand, Physical Demand, Temporal Demand, Performance, Effort, and Frustration (see Figure 11). The results showed significant differences in “Physical Demand” ($\chi^2(2) = 15.24$, $p < .001$), “Effort” ($\chi^2(2) = 11.35$, $p = .0034$), and “Frustration” ($\chi^2(2) = 11.53$, $p = .0031$) across the techniques. However, no significant differences were found in “Mental Demand” ($\chi^2(2) = 4.33$, $p = .115$), “Temporal Demand” ($\chi^2(2) = 5.25$, $p = .072$), or “Performance” ($\chi^2(2) = 2.11$, $p = .347$) across the techniques.

A post-hoc analysis was conducted using pairwise Wilcoxon Rank-Sum Tests (Mann-Whitney U tests) with Bonferroni correction for multiple comparisons. The post-hoc analysis revealed significant differences in the “Physical Demand” aspect. Specifically, for “Physical Demand,” the Hand technique exhibited significantly higher scores compared to both the Swarm and Controller techniques ($W = 1.0$, $p = .003$), and the Hand technique also had significantly higher scores compared to the Swarm technique ($W = .0$, $p = .010$). Regarding “Effort,” we observed significant differences between the Controller and Hand techniques. The Effort scores were significantly higher for the Hand technique compared to the Controller technique ($W = 2.5$, $p = .018$), and similarly, the Hand technique had significantly higher scores compared to the Swarm technique ($W = 3.0$, $p = .021$). For the “Frustration” aspect, significant differences were found between the Controller and Hand techniques. The Frustration scores were significantly higher for the Hand technique compared to the Controller technique ($W = 5.5$, $p = .041$), and also significantly higher for the Hand technique compared to the Swarm technique ($W = 2.0$, $p = .026$).

5.4.3 System Usability Scale. A Friedman test was conducted to compare the SUS scores among the Controller ($M = 87.08$, $SD = 14.18$), Hand ($M = 83.96$, $SD = 13.75$), and Swarm techniques ($M = 81.25$, $SD = 14.16$). The results revealed no significant difference in SUS scores across the three techniques ($\chi^2(2) = 1.22$, $p = .544$). These results suggest that there is no significant difference in usability, as measured by the SUS, between the three techniques.

5.4.4 Preference. Regarding the performance ranking, eight participants ranked Swarm Manipulation as the best technique, while Controller was ranked best by three participants, and Hand was only ranked as the best technique by one participant (see Figure 12). In terms of ease of use, Swarm Manipulation received the highest average rating of 4.5 out of 5 ($SD = .90$) among the three techniques. The Controller was rated 3.83 ($SD = .94$), while Hand received an average rating of 3.75 ($SD = .97$).

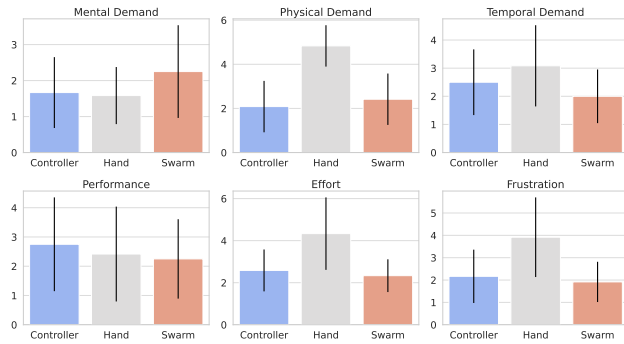


Figure 11: Bar chart showcasing the average ratings for six NASA-TLX dimensions across three techniques, with error bars indicating standard deviation.




Figure 12: The bar chart depicts the rankings of three interaction techniques—Hand, Swarm, and Controller—based on participants' preferences.

6 DISCUSSION

6.1 Task Performance

The task performance results shed light on the effectiveness and potential advantages of the Swarm Manipulation technique compared to the Hand and Controller techniques in different VR manipulation tasks. In the selection task, the Swarm Manipulation technique demonstrated significantly faster completion times compared to the Hand technique in various conditions and mixed results compared to Controller. This suggests that the swarm-based approach can enhance efficiency and speed in target selection, though the advantages over Controller may vary depending on the conditions.

In the rotation task, the Swarm Manipulation technique outperformed the Hand technique in terms of task completion time only in the all-target rotation condition and was significantly slower in both close-distance dual-target condition. This finding highlights the complexity of the rotation task and suggests that the Swarm Manipulation technique's efficiency may be context-dependent. The lower rotation angle deviation in the close-distance dual-target condition indicates a trade-off between speed and accuracy, possibly due to the distributed and collective nature of swarm entities.

In the resizing task, the Swarm Manipulation technique was significantly faster than Controller in many conditions but was only significantly faster than Hand in the all-target resizing condition. However, the Swarm manipulation approach did demonstrate significantly smaller resizing size deviations compared to Hand in the long-distance single-target resizing condition, implying higher accuracy in certain scenarios. Swarm Manipulation technique also demonstrated significantly smaller resizing size deviations than Controller in the close-distance dual-target condition, long-distance single-target condition, and all-target resizing condition. It is essential to recognize the outstanding performance of the Swarm Manipulation technique in resizing tasks, as it excelled in terms of speed and accuracy under different conditions.

Overall, the Swarm Manipulation technique shows promising advantages in speed and accuracy in certain tasks and conditions. Still, the performance is not uniformly superior across all scenarios, especially in the Rotation task. Further investigation and refinement could focus on understanding the underlying factors that contribute to these variations in performance. The collective control theory and coordination of swarm entities might help compensate for the potential challenges introduced by distance, resulting in more accurate resizing manipulations.

6.2 User Preference

Based on our analysis of the user preference questionnaires we gathered, it is evident that the Swarm Manipulation technique was generally well-received by the participants. The majority of participants not only ranked it as the most effective technique ($N = 9$) but also considered it the most user-friendly ($N = 10$).

Nevertheless, it is important to acknowledge the presence of diverse opinions among the participants. For instance, *P1* commented that the Swarm Manipulation technique provided an expanded range of finger control, enabling the manipulation of objects without the need for repetitive hand movements. However, this user also pointed out certain difficulties with gesture recognition, indicating room for improvement in terms of accuracy. Conversely, *P2* found the Swarm Manipulation technique innovative and valued its potential for adjusting reach length. Yet, they also emphasized the necessity for improved gesture recognition and stability due to issues encountered with the system misinterpreting their gestures. Additionally, *P11* suggested that incorporating a visual representation of the hand's outline within the Swarm Manipulation technique could be advantageous for novice VR users, aiding their orientation in the virtual environment. *P1*, *P2*, and *P10* also emphasized the challenge of learning the manipulation mechanisms, highlighting the need for enhanced user guidance and support.

While the overall reception of the Swarm Manipulation technique was positive, the feedback received clearly indicates areas in need of improvement, namely gesture recognition accuracy, system stability, and user guidance. Further research endeavors could delve into these specific areas to explore them more comprehensively, ultimately enhancing the usability and overall user experience of the Swarm Manipulation technique in the context of VR.

6.3 Limitations and Future Work

The study represents an initial exploration and does not attempt to cover the entire spectrum of factors influencing the manipulation of swarm entities in VR environments. Our forthcoming endeavors will be directed towards further investigations along this trajectory. A promising starting point is a comprehensive analysis of the control models [1, 2, 14] underpinning swarm interactions.

Additionally, the primary emphasis of this paper lies in the evaluation of interaction tasks using a prototype of this novel interaction technology, with limited exploration of the underlying interaction theory [20]. In our future work, we plan to employ computational interaction methods [40] to gain insights into users' understanding of swarm interaction and optimise the associated control model, aiming to further enhance performance. However, it is worth highlighting that despite our limited understanding of the control principles governing Swarm Manipulation, our research has already demonstrated significant benefits in manipulation tasks, thereby underscoring the potential of this field to define new interaction paradigms in the future [36, 37].

Finally, we outline a vision for our future research. We firmly believe that in addition to investigating how users control swarm interaction, an essential topic for subsequent studies, exploring how swarm particles *collaborate* and exchange signals among themselves (that is, "self-organising"), and consequently applying automation theory, is central to advancing research on swarm interaction. This represents a novel interpretation of Pattie Maes' vision of Intelligent Agents [31, 32], while retaining the concept of Direct Manipulation [47].

Moreover, these swarm particles may serve as the fundamental building blocks of future VR/AR interaction interfaces. Similar to Hiroshi Ishii's Tangible Bits [22] in the realm of tangible user interfaces and concepts such as Radical Atoms [23] in the future materials field, swarm particles can be regarded as the smallest units for user control and interaction, as well as for constructing virtual worlds. The advantage of using swarm particles for display and interaction is that they resemble pixel dots in the digital world, are less expensive to maintain than tangible user interfaces [19], and allow developers to design adaptive user interfaces that balance visuals with immersive experiences. Consequently, swarm particles have the potential to become the elemental units that aid researchers in understanding how to interact with virtual information in the future.

7 CONCLUSION

We have introduced the Swarm Manipulation technique for VR interactions and evaluated its performance compared with the conventional Hand and Controller manipulation techniques in a user study with 12 participants. Our findings unveil a nuanced picture: Overall, our results show the swarm manipulation technique did reveal good performance, exhibiting significantly faster speeds compared with at least one of the other two techniques across most conditions for these three tasks. Notably, the Swarm technique's advantage in rotation tasks is context-dependent, displaying increased swiftness in a close-distance single-target scenario, but also revealing an inherent trade-off between speed and accuracy, likely attributed to the distributed and collaborative nature of swarm

particles. In resizing tasks, the Swarm Manipulation technique's performance was varied, demonstrating increased speed in specific conditions and minimized size deviations in the long-distance single-target resizing condition. However, these advantages were not uniformly observed across all resizing conditions.

Additionally, the subjective user experience feedback gathered from participants unequivocally positions the Swarm Manipulation technique as the most preferred and perceived user-friendly approach among the manipulation techniques studied. This feedback may further underscore the Swarm Manipulation technique as a promising approach for VR object manipulation. This study not only establishes an initial proof-of-concept for the potential of Swarm Manipulation to enhance usability in certain scenarios but also highlights areas where performance varies, necessitating further investigation.

Our forthcoming research endeavors will delve deeper into the realm of swarm interaction in mixed reality and explore the intricacies of user control dynamics and collaborative particle behavior within the swarm paradigm. These efforts will consider the complex performance patterns revealed in this study and seek to refine and optimize the Swarm Manipulation technique for a broader range of applications and conditions.

ACKNOWLEDGMENTS

Xiang Li is supported by the China Scholarship Council (CSC) International Cambridge Scholarship (No. 202208320092). John J. Dudley and Per Ola Kristensson are supported by the EPSRC (grants EP/S027432/1 and EP/W02456X/1).

REFERENCES

- [1] Parastoo Abtahi, Sidney Q. Hough, James A. Landay, and Sean Follmer. 2022. Beyond Being Real: A Sensorimotor Control Perspective on Interactions 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 358, 17 pages. <https://doi.org/10.1145/3491102.3517706>
- [2] Parastoo Abtahi, Benoit Landry, Jackie (Junrui) Yang, Marco Pavone, Sean Follmer, and James A. Landay. 2019. Beyond The Force: Using Quadcopters to Appropriate Objects and the Environment for Haptics in Virtual Reality. 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, 1–13. <https://doi.org/10.1145/3290605.3300589>
- [3] Joanna Bergström, Tor-Salve Dalsgaard, Jason Alexander, and Kasper Hornbæk. 2021. How to evaluate object selection and manipulation in vr? guidelines from 20 years of studies. In *proceedings of the 2021 CHI conference on human factors in computing systems*. 1–20.
- [4] Doug A. Bowman and Larry F. Hodges. 1997. An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments. In *Proceedings of the 1997 symposium on Interactive 3D graphics (I3D '97)*. Association for Computing Machinery, New York, NY, USA, 35–ff. <https://doi.org/10.1145/253284.253301>
- [5] Doug A Bowman, Donald B Johnson, and Larry F Hodges. 1999. Testbed evaluation of virtual environment interaction techniques. In *Proceedings of the ACM symposium on Virtual reality software and technology*. 26–33.
- [6] John Brooke et al. 1996. SUS-A quick and dirty usability scale. *Usability evaluation in industry* 189, 194 (1996), 4–7.
- [7] Daniel S. Brown, Michael A. Goodrich, Shin-Young Jung, and Sean Kerman. 2016. Two invariants of human-swarm interaction. *Journal of Human-Robot Interaction* 5, 1 (2016), 1–31. <https://doi.org/10.5898/JHRI.5.1.Brown>
- [8] Daniel S. Brown, Sean C. Kerman, and Michael A. Goodrich. 2014. Human-swarm interactions based on managing attractors. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction (HRI '14)*. Association for Computing Machinery, New York, NY, USA, 90–97. <https://doi.org/10.1145/2559636.2559661>
- [9] Wu Chen, Jijia Liu, Hongzhi Guo, and Nei Kato. 2020. Toward Robust and Intelligent Drone Swarm: Challenges and Future Directions. *IEEE Network* 34,

- 4 (July 2020), 278–283. <https://doi.org/10.1109/MNET.001.1900521> Conference Name: IEEE Network.
- [10] Diane Dewez, Rebecca Fribourg, Ferran Argelaguet, Ludovic Hoyet, Daniel Mestre, Mel Slater, and Anatole Lécuyer. 2019. Influence of Personality Traits and Body Awareness on the Sense of Embodiment in Virtual Reality. In *2019 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 123–134. <https://doi.org/10.1109/ISMAR.2019.00-12> ISSN: 1554-7868.
- [11] Griffin Dietz, Jane L. E. Peter Washington, Lawrence H. Kim, and Sean Follmer. 2017. Human Perception of Swarm Robot Motion. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17)*. Association for Computing Machinery, New York, NY, USA, 2520–2527. <https://doi.org/10.1145/3027063.3053220>
- [12] John J Dudley, Hendrik Schuff, and Per Ola Kristensson. 2018. Bare-handed 3D drawing in augmented reality. In *Proceedings of the 2018 Designing Interactive Systems Conference*. 241–252.
- [13] Clifton Forlines, Daniel Vogel, and Ravin Balakrishnan. 2006. Hybridpointing: fluid switching between absolute and relative pointing with a direct input device. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*. 211–220.
- [14] Eric J Gonzalez and Sean Follmer. 2023. Sensorimotor Simulation of Redirected Reaching Using Stochastic Optimal Feedback Control. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (Hamburg, Germany) (CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 776, 17 pages. <https://doi.org/10.1145/3544548.3580767>
- [15] Tovi Grossman and Ravin Balakrishnan. 2006. The Design and Evaluation of Selection Techniques for 3D Volumetric Displays. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology (Montreux, Switzerland) (UIST '06)*. Association for Computing Machinery, New York, NY, USA, 3–12. <https://doi.org/10.1145/1166253.1166257>
- [16] Iman Haghighi, Sadra Sadraddini, and Calin Belta. 2016. Robotic swarm control from spatio-temporal specifications. In *2016 IEEE 55th Conference on Decision and Control (CDC)*. 5708–5713. <https://doi.org/10.1109/CDC.2016.7799146>
- [17] Sandra G. Hart. 2006. Nasa-Task Load Index (NASA-TLX); 20 Years Later. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 50, 9 (Oct. 2006), 904–908. <https://doi.org/10.1177/154193120605000909> Publisher: SAGE Publications Inc.
- [18] Ken Hinckley, Randy Pausch, John C. Goble, and Neal F. Kassell. 1994. A survey of design issues in spatial input. In *Proceedings of the 7th annual ACM symposium on User interface software and technology (UIST '94)*. Association for Computing Machinery, New York, NY, USA, 213–222. <https://doi.org/10.1145/192426.192501>
- [19] Lars Erik Holmquist. 2023. Bits Are Cheap, Atoms Are Expensive: Critiquing the Turn Towards Tangibility in HCI. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (Hamburg, Germany) (CHI EA '23)*. Association for Computing Machinery, New York, NY, USA, Article 409, 8 pages. <https://doi.org/10.1145/3544549.3582744>
- [20] Kasper Hornbæk and Antti Oulasvirta. 2017. What is interaction?. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 5040–5052.
- [21] Ludovic Hoyet, Ferran Argelaguet, Corentin Nicole, and Anatole Lécuyer. 2016. “Wow! I Have Six Fingers!”: Would You Accept Structural Changes of Your Hand in VR? *Frontiers in Robotics and AI* 3 (2016), 27.
- [22] Hiroshi Ishii. 2008. Tangible Bits: Beyond Pixels. In *Proceedings of the 2nd International Conference on Tangible and Embedded Interaction (Bonn, Germany) (TEI '08)*. Association for Computing Machinery, New York, NY, USA, xv–xxv. <https://doi.org/10.1145/1347390.1347392>
- [23] Hiroshi Ishii, Dávid Lakatos, Leonardo Bonanni, and Jean-Baptiste Labrune. 2012. Radical Atoms: Beyond Tangible Bits, toward Transformable Materials. *Interactions* 19, 1 (jan 2012), 38–51. <https://doi.org/10.1145/2065327.2065337>
- [24] Hala Khodr, Ulysse Ramage, Kevin Kim, Arzu Guneyso Ozgur, Barbara Bruno, and Pierre Dillenbourg. 2020. Being Part of the Swarm: Experiencing Human-Swarm Interaction with VR and Tangible Robots. In *Proceedings of the 2020 ACM Symposium on Spatial User Interaction (SUI '20)*. Association for Computing Machinery, New York, NY, USA, 1–2. <https://doi.org/10.1145/3385959.3422695>
- [25] Masatomo Kobayashi and Takeo Igarashi. 2008. Ninja Cursors: Using Multiple Cursors to Assist Target Acquisition on Large Screens. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Florence, Italy) (CHI '08)*. Association for Computing Machinery, New York, NY, USA, 949–958. <https://doi.org/10.1145/1357054.1357201>
- [26] Andreas Kolling, Katia Sycara, Steven Nunnally, and Michael Lewis. 2013. Human-swarm interaction: an experimental study of two types of interaction with foraging swarms. *Journal of Human-Robot Interaction* 2, 2 (2013), 103–129. <https://doi.org/10.5898/JHRI.2.2.Kolling>
- [27] Andreas Kolling, Phillip Walker, Nilanjan Chakraborty, Katia Sycara, and Michael Lewis. 2016. Human Interaction With Robot Swarms: A Survey. *IEEE Transactions on Human-Machine Systems* 46, 1 (Feb. 2016), 9–26. <https://doi.org/10.1109/THMS.2015.2480801> Conference Name: IEEE Transactions on Human-Machine Systems.
- [28] Mathieu Le Goc, Lawrence H. Kim, Ali Parsaei, Jean-Daniel Fekete, Pierre Dragicevic, and Sean Follmer. 2016. Zooids: Building Blocks for Swarm User Interfaces. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. Association for Computing Machinery, New York, NY, USA, 97–109. <https://doi.org/10.1145/2984511.2984547>
- [29] Xiang Li, Yuzheng Chen, Rakesh Patibanda, and Florian ‘Floyd’ Mueller. 2021. VrCAPTCHA: Exploring CAPTCHA Designs in Virtual Reality. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems (Yokohama, Japan) (CHI EA '21)*. Association for Computing Machinery, New York, NY, USA, Article 167, 4 pages. <https://doi.org/10.1145/3411763.3451985>
- [30] Paul Lubos, Gerd Bruder, and Frank Steinicke. 2014. Are 4 hands better than 2? bimanual interaction for quadmanual user interfaces. In *Proceedings of the 2nd ACM symposium on Spatial user interaction*. 123–126.
- [31] Pattie Maes. 1995. Agents that reduce work and information overload. In *Readings in human-computer interaction*. Elsevier, 811–821.
- [32] Pattie Maes. 1995. Artificial Life Meets Entertainment: Lifelike Autonomous Agents. *Commun. ACM* 38, 11 (nov 1995), 108–114. <https://doi.org/10.1145/219717.219808>
- [33] D. Mendes, F. M. Caputo, A. Giachetti, A. Ferreira, and J. Jorge. 2019. A Survey on 3D Virtual Object Manipulation: From the Desktop to Immersive Virtual Environments. *Computer Graphics Forum* 38, 1 (2019), 21–45. <https://doi.org/10.1111/cgf.13390> eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13390>
- [34] Daniel Mendes, Filipe Relvas, Alfredo Ferreira, and Joaquim Jorge. 2016. The benefits of DOF separation in mid-air 3D object manipulation. In *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology (VRST '16)*. Association for Computing Machinery, New York, NY, USA, 261–268. <https://doi.org/10.1145/2993369.2993396>
- [35] Mark R Mine, Frederick P Brooks Jr, and Carlo H Sequin. 1997. Moving objects in space: exploiting proprioception in virtual-environment interaction. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 19–26.
- [36] Florian ‘Floyd’ Mueller, Rakesh Patibanda, Richard Byrne, Zhuying Li, Yan Wang, Josh Andres, Xiang Li, Jonathan Marquez, Stefan Greuter, Jonathan Duckworth, and Joe Marshall. 2021. Limited Control Over the Body as Intriguing Play Design Resource. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (Yokohama, Japan) (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 435, 16 pages. <https://doi.org/10.1145/3411764.3445744>
- [37] Florian ‘Floyd’ Mueller, Nathan Semertzidis, Josh Andres, Joe Marshall, Steve Benford, Xiang Li, Louise Matjeka, and Yash Mehta. 2023. Towards Understanding the Design of Intertwined Human-Computer Integrations. *ACM Trans. Comput.-Hum. Interact.* (apr 2023). <https://doi.org/10.1145/3590766> Just Accepted.
- [38] Ken Nakagaki, Joanne Leong, Jordan L. Tappa, João Wilbert, and Hiroshi Ishii. 2020. HERMITS: Dynamically Reconfiguring the Interactivity of Self-propelled TUIs with Mechanical Shell Add-ons. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology (UIST '20)*. Association for Computing Machinery, New York, NY, USA, 882–896. <https://doi.org/10.1145/3379337.3415831>
- [39] Ken Nakagaki, Jordan L Tappa, Yi Zheng, Jack Forman, Joanne Leong, Sven Koenig, and Hiroshi Ishii. 2022. (Dis)Appearables: A Concept and Method for Actuated Tangible UIs to Appear and Disappear based on Stages. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3491102.3501906>
- [40] Antti Oulasvirta, Per Ola Kristensson, Xiaojun Bi, and Andrew Howes. 2018. *Computational Interaction*. Oxford University Press.
- [41] Jeffrey S. Pierce, Brian C. Stearns, and Randy Pausch. 1999. Voodoo dolls: seamless interaction at multiple scales in virtual environments. In *Proceedings of the 1999 symposium on Interactive 3D graphics (I3D '99)*. Association for Computing Machinery, New York, NY, USA, 141–145. <https://doi.org/10.1145/300523.300540>
- [42] Ivan Poupyrev, Mark Billinghurst, Suzanne Weghorst, and Tadao Ichikawa. 1996. The go-go interaction technique: non-linear mapping for direct manipulation in VR. (Nov. 1996), 79–80. <https://doi.org/10.1145/237091.237102> MAG ID: 2042751882.
- [43] Kari-Jouko Räihä and Oleg Špakov. 2009. Disambiguating ninja cursors with eye gaze. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1411–1414.
- [44] Fabrice Saffre, Hanno Hildmann, and Hannu Karvonen. 2021. The Design Challenges of Drone Swarm Control. In *Engineering Psychology and Cognitive Ergonomics (Lecture Notes in Computer Science)*, Don Harris and Wen-Chin Li (Eds.). Springer International Publishing, Cham, 408–426. https://doi.org/10.1007/978-3-030-77932-0_32
- [45] Jonas Schjerlund, Kasper Hornbæk, and Joanna Bergström. 2021. Ninja Hands: Using Many Hands to Improve Target Selection in VR. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3411764.3445759>
- [46] Martin Schrepp, Jörg Thomaschewski, and Andreas Hinderks. 2017. Construction of a benchmark for the user experience questionnaire (UEQ). (2017).

- [47] Ben Shneiderman and Pattie Maes. 1997. Direct manipulation vs. interface agents. *interactions* 4, 6 (1997), 42–61.
- [48] Mel Slater, Bernhard Spanlang, Maria V. Sanchez-Vives, and Olaf Blanke. 2010. First person experience of body transfer in virtual reality. *PLOS ONE* 5, 5 (May 2010). <https://doi.org/10.1371/journal.pone.0010564> MAG ID: 2163290657.
- [49] Zhaomou Song, John J Dudley, and Per Ola Kristensson. 2022. Efficient Special Character Entry on a Virtual Keyboard by Hand Gesture-Based Mode Switching. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 864–871.
- [50] Richard Stoakley, Matthew J. Conway, and Randy Pausch. 1995. Virtual reality on a WIM: interactive worlds in miniature. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '95)*. ACM Press/Addison-Wesley Publishing Co., USA, 265–272. <https://doi.org/10.1145/223904.223938>
- [51] Ryo Suzuki, Jun Kato, Mark D. Gross, and Tom Yeh. 2018. Reactile: Programming Swarm User Interfaces through Direct Physical Manipulation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3173773>
- [52] Manuel Veit, Antonio Capobianco, and Dominique Bechmann. 2009. Influence of degrees of freedom's manipulation on performances during orientation tasks in virtual reality environments. In *Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology (VRST '09)*. Association for Computing Machinery, New York, NY, USA, 51–58. <https://doi.org/10.1145/1643928.1643942>
- [53] Wenge Xu, Hai-Ning Liang, Yuzheng Chen, Xiang Li, and Kangyou Yu. 2020. Exploring Visual Techniques for Boundary Awareness During Interaction in Augmented Reality Head-Mounted Displays. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 204–211. <https://doi.org/10.1109/VR46266.2020.00039> ISSN: 2642-5254.
- [54] Difeng Yu, Xueshi Lu, Rongkai Shi, Hai-Ning Liang, Tilman Dingler, Eduardo Velloso, and Jorge Goncalves. 2021. Gaze-Supported 3D Object Manipulation in Virtual Reality. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (Yokohama, Japan) (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 734, 13 pages. <https://doi.org/10.1145/3411764.3445343>
- [55] Difeng Yu, Qiushi Zhou, Tilman Dingler, Eduardo Velloso, and Jorge Goncalves. 2022. Blending On-Body and Mid-Air Interaction in Virtual Reality. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 637–646. <https://doi.org/10.1109/ISMAR55827.2022.00081>
- [56] Lilith Yu, Chenfeng Gao, David Wu, and Ken Nakagaki. 2023. AeroRigUI: Actuated TUIs for Spatial Interaction using Rigging Swarm Robots on Ceilings in Everyday Space. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–18. <https://doi.org/10.1145/3544548.3581437>