Bimanual Mid-Air Multi-Object Manipulation of Graph Data with Gesture Phrasing

Pantea Habibi
H. Milton Stewart School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, Georgia, USA
Computer Science Department
University of Illinois Chicago
Chicago, Illinois, USA
phabibi6@gatech.edu

Debaleena Chattopadhyay
Computer Science Department
University of Illinois Chicago
Chicago, Illinois, USA
debchatt@uic.edu

Abstract

Exploring graph data in virtual reality (VR) often involves identifying subgraphs, comparing clusters, tracing paths, and making localized edits—tasks that require efficient selection and manipulation of multiple objects. While multi-selection techniques for immersive environments have been recently established, the subsequent manipulation of selected subgraphs remains underexplored for fluid graph exploration, particularly in how they can benefit from *gesture phrasing*—structuring gestures to align with users' cognitive chunking of subtasks. This paper investigates how bare-handed bimanual gesture phrasing influences multi-object dragging. We evaluate three dragging techniques with varying levels of kinesthetic tension ($n = 21$). Our results indicate that the half-tension technique was on par with high tension in terms of efficiency, but yields higher accuracy and user preference. The findings suggest that moderate kinesthetic coupling provides sufficient spatial anchoring without the fatigue associated with sustained tension, offering design implications for bimanual interaction in VR.

CCS Concepts

• **Human-centered computing → Gestural input**; *Empirical studies in interaction design*; **Empirical studies in HCI**; *Laboratory experiments*.

Keywords

Bimanual interaction, mid-air, virtual reality, multi-object manipulation, dragging, graph data

ACM Reference Format:
Pantea Habibi and Debaleena Chattopadhyay. 2026. Bimanual Mid-Air Multi-Object Manipulation of Graph Data with Gesture Phrasing. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3772363.3798391

ACM ISBN 979-8-4007-2281-3/26/04
https://doi.org/10.1145/3772363.3798391

1 Introduction

Graphs are widely used to represent relationships in domains such as citation networks, biological pathways, and knowledge graphs, often comprising hundreds or thousands of nodes [12, 19, 23]. In immersive visualization systems, users frequently need to identify subgraphs, compare clusters, and reorganize node groups, making multi-object selection and manipulation fundamental interaction primitives [3, 7, 20]. Virtual reality (VR) supports embodied navigation and spatial perception, enabling direct mid-air interaction with node-link structures that would be difficult to manage in 2D interfaces [9, 22].

Bimanual interaction is well known to support such spatial tasks by distributing control across the hands, often following asymmetric roles where the non-dominant hand establishes a spatial frame of reference and the dominant hand performs precise actions [10, 14]. Prior work across pen, touch, and mid-air modalities has shown that dependent asymmetric techniques can improve coordination and reduce cognitive load during manipulation tasks [15, 25, 27]. In XR, bimanual input has been applied to navigation, selection, and manipulation, with benefits for precision and expressiveness when supported by visual or embodied feedback [1, 5, 9].

A key design mechanism for structuring bimanual input is *gesture phrasing*: organizing actions into temporally and kinesthetically coherent units [4, 8, 26]. Phrased gestures often rely on sustained or overlapping muscle activation—referred to as *kinesthetic tension*—to anchor one hand while the other performs task-relevant motion [13, 27]. While sustained tension can stabilize multi-step input in pen and touch contexts, its role in mid-air bimanual manipulation remains unclear, especially given the lack of haptic support.

Multi-object selection in VR has been studied using volumetric brushes, lasso techniques, and slicing volumes [18, 21, 24], as well as gaze-assisted serial methods for distant targets [16]. However, fewer studies have examined how bimanual techniques support *continuous workflows* that integrate selection with subsequent manipulation for structured datasets such as graphs. Recent work investigated sphere-based bimanual mid-air techniques for selecting dense node clusters in immersive graphs [11]. That study focused on selection performance and user preference, but did not examine how different bimanual coordination strategies affect the subsequent manipulation of selected objects.

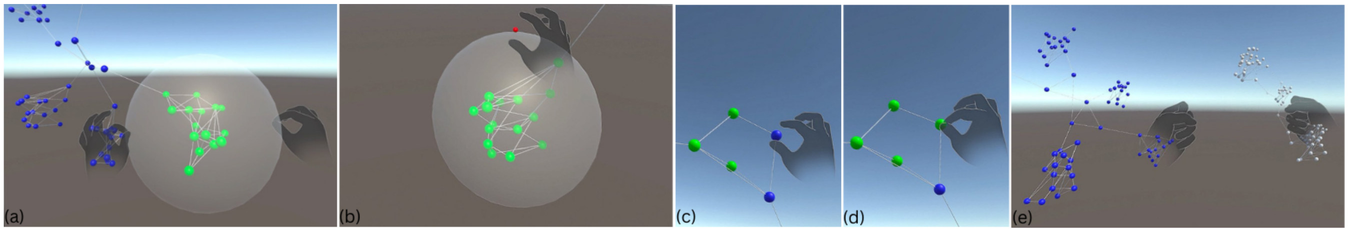

Figure 1: (a) Sphere selection; (b) the red "delete" object for sphere-deselect; (c) Single-select; (d) Single-deselect; (e) rotate and scale.

Building on this gap, we focus on manipulation—how gesture phrasing and inter-hand tension shape multi-object dragging after selection. We compare bimanual dragging techniques that vary how long the non-dominant hand maintains an anchoring gesture relative to the dominant hand's motion, and measure performance and preference in graph-based tasks.

2 Interaction Techniques

We adopt the "Sphere-SS" technique [11], where the user defines a selection volume by performing pinch gestures with both hands and then move to render a sphere (Figure 1). We selected this technique because it has been shown to be significantly more effective than serial selection for graph clusters [11]. The goal of this experiment was to investigate how specific design parameters influence user performance and preference during multi-object manipulation.

To investigate how varying levels of kinesthetic tension affect multi-object manipulation in VR, we designed three bimanual mid-air dragging techniques. All techniques assume that the user has already selected a group of objects using the Sphere-SS method and now intends to drag them to a new location. Each technique varies in how long the non-dominant hand maintains a fist gesture relative to the dominant hand's drag motion.

In the low-tension technique, the non-dominant hand makes a fist and then releases it. Only after this release can the dominant hand begin dragging. Dragging is not registered if the dominant hand moves before the release is completed (Figure 2). This configuration introduces minimal kinesthetic coupling between the two hands. We also designed a variant of the low-tension technique, low-tension with feedback, which adds a visual feedback element to the low-tension configuration (Figure 2). A textual indicator appears above the dominant hand to signal when the non-dominant hand's release has been successfully registered, clarifying when dragging may begin.

In the half-tension technique, the non-dominant hand initiates and holds the fist gesture until the dominant hand starts dragging. Once the drag begins, the non-dominant hand is free to release (Figure 2). This creates a brief period of overlapping hand activity and introduces moderate tension between the hands.

In the high-tension technique, the non-dominant hand holds the fist gesture for the entire duration of the drag motion performed by the dominant hand. The gesture ends only after the drag is completed, resulting in the highest level of sustained tension and coordination between the two hands (Figure 2).

3 Method

3.1 Tasks and Procedure

The task in this experiment involved selecting a specific subgraph from a larger graph and dragging the selected subgraph to a corresponding location in a target graph. Each trial featured a unique graph containing a simple "house" pattern—composed of five connected nodes—embedded either near the center or along the periphery of the graph. Five graphs were used in total, varying in the number of nodes from the house pattern that were attached to separate, randomly generated distractor graphs. In one trial, one node of the house was attached to a distractor graph; in another, two nodes were attached to two distinct graphs, and so on, up to five graphs (Figure 3).

During each trial, participants were required to identify and select all nodes that constituted the house pattern without selecting any distractor nodes. Once selected, participants used the assigned dragging technique to move the selected subgraph to a matching "target" location in a duplicated graph. A placement was considered successful only if all nodes in the dragged pattern precisely aligned with their corresponding nodes in the target graph. If the placement was incorrect or misaligned, the copied nodes disappeared, and the participant had to repeat the dragging process. The selection state persisted across attempts unless explicitly modified.

Participants could rotate and scale the graph, and could deselect nodes via single-deselect or by removing the selection sphere (Figure 1). Before starting, participants completed a demographic questionnaire and received verbal instructions and a demonstration of the techniques. We used an Oculus Quest 2 with hand-tracking in a Unity application (Oculus SDK). All procedures were approved by the university's Institutional Review Board (IRB).

3.2 Experimental Design

This study employed a within-subjects design with interaction technique as the primary independent variable, comprising four levels. A secondary independent variable was graph type, with five levels based on the number of nodes from the house pattern attached to distractor graphs (Figure 3). The order of interaction techniques and graph types was counterbalanced using a balanced Latin square. Each participant completed a total of 40 trials: 4 techniques × 5 graph types × 2 blocks.

Participants reported their subjective preferences for each interaction technique at the end of the study. To evaluate performance,

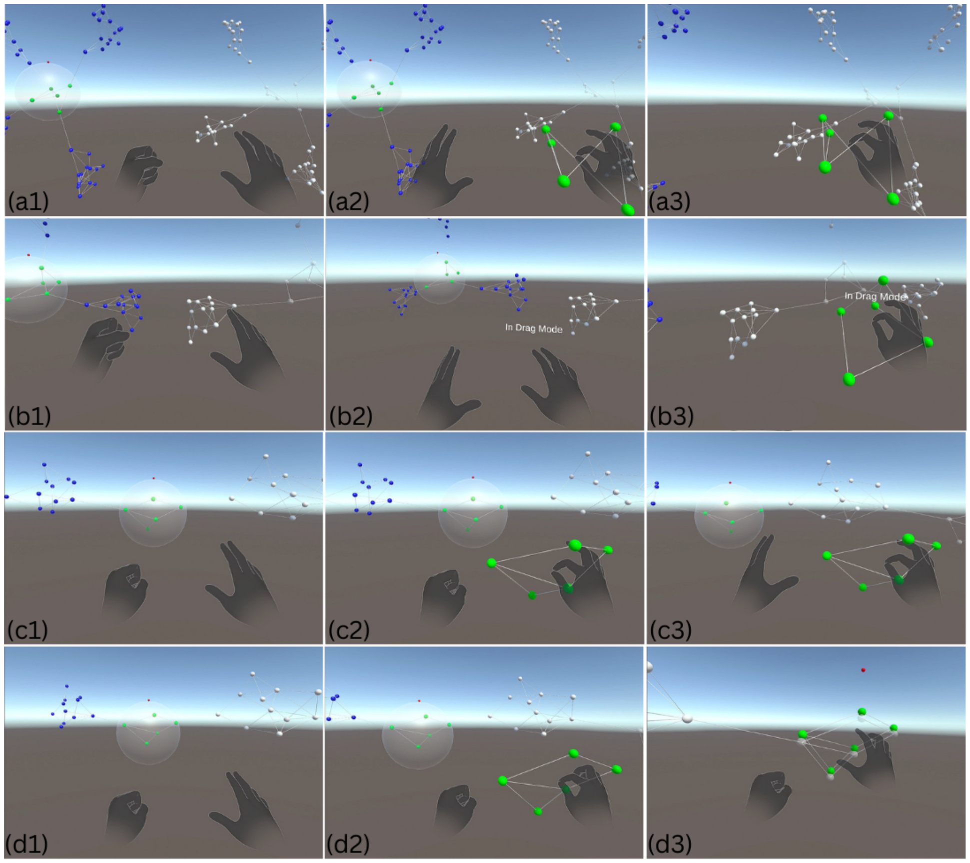

Figure 2: Bimanual dragging techniques: (a) low-tension; (b) low-tension with visual feedback; (c) half-tension; (d) high-tension.

we analyzed four quantitative metrics that captured both efficiency and accuracy in multi-object manipulation.

Total Completion Time (TCT) measured the total duration of a trial, from the start of selection to the successful placement of the subgraph. Lower values reflect faster task execution.

Successful Drag Time (SDT) captured the duration of the final, successful drag attempt. SDT reflects the efficiency of the goal-directed action alone.

Drag Efficiency Index (DEI) assessed how effectively users allocated their total drag time by comparing the successful drag to all drag attempts:

$$\mathrm{DEI} = \frac{T_{\text{last}}}{\sum_{i=1}^{N} T_i} = \frac{\text{Duration of final (successful) drag}}{\text{Total duration of all drag attempts}}$$

Higher values indicate that most of the drag effort contributed directly to task completion, while lower values suggest time lost to failed or corrective attempts.

Error Rate (ER) quantified the proportion of unsuccessful drag attempts per trial:

$$ER = \frac{D_{\text{fail}}}{D_{\text{fail}} + 1} = \frac{\text{Number of failed drag attempts}}{\text{Total number of drag attempts}}$$

This normalized metric enables fair comparison across trials with different numbers of attempts, with higher values reflecting lower accuracy.

3.3 Hypotheses

Drawing on prior research on gesture phrasing and hand coordination, we formulated the following hypotheses to examine how varying levels of inter-hand tension affect multi-object manipulation in VR.

H1. Dragging with high-tension gesture phrasing will be **(i)** more efficient and **(ii)** more accurate than dragging with **(a)** half tension or **(b)** low tension. This hypothesis is motivated by findings from pen-based interaction research, which demonstrated that sustained muscular tension—particularly when maintained by the non-dominant hand—can help users structure and stabilize complex multi-stroke input [13].

H2. Dragging with low-tension gesture phrasing will be **(i)** more accurate but **(ii)** less efficient when augmented with visual feedback compared to without. We hypothesized that visual feedback would help users disambiguate selection from manipulation, thereby improving accuracy and reducing mode confusion [6]. However, attending to this feedback might slow performance, resulting in reduced task efficiency.

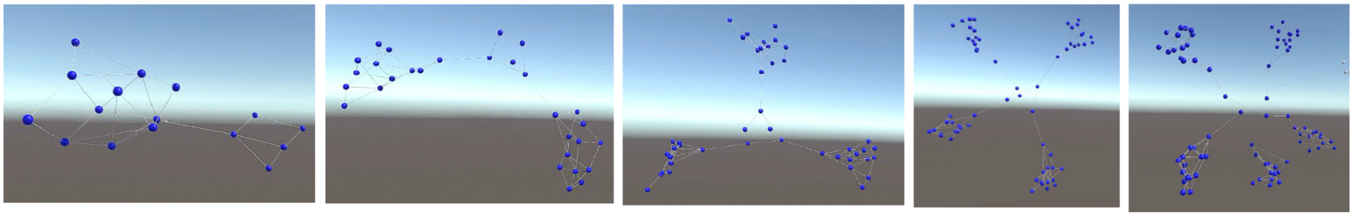

Figure 3: Five graphs used, each varying in the number of distractor graphs (1–5, left to right).

3.4 Data Analysis

Dependent variables did not meet the assumptions of normality; therefore, medians were used to summarize repeated measures for each experimental condition. We conducted a Linear Mixed Model (LMM) analysis using restricted maximum likelihood (REML), treating participants as a random factor. The analysis was performed using the `lme4` package in R [2]. We report F-statistics from Type III ANOVA with Satterthwaite approximation, and conducted pairwise comparisons using pooled variance with Holm-Bonferroni correction. The initial significance level was set at $\alpha = 0.05$.

4 Results

4.1 Participants

A total of 21 right-handed participants (10 women; *Mdn* age = 28, range = 21–36 years) took part in the study. Approximately 52% reported prior experience with mid-air input, primarily in gaming contexts. Participants were recruited via university mailing lists and word of mouth, and compensated with $15 gift cards.

4.2 Overview

Linear mixed-effects models (LMMs) revealed no significant main effect of graph type (Figure 3) on any performance measures. However, the interaction technique had a significant main effect across all dependent variables: total completion time (TCT), $F(2, 921) = 6.65, p = .001$; successful drag time (SDT), $F(2, 921) = 19.00, p < .0001$; drag efficiency index (DEI), $F(2, 921) = 5.21, p = .006$; and error rate (ER), $F(2, 921) = 4.23, p = .014$.

No significant differences were found between techniques for the selection time prior to the drag task, indicating that observed effects were specific to the manipulation phase. These results suggest that the level of tension in gesture phrasing significantly influenced the efficiency and quality of drag execution. In particular, the significant differences in DEI and ER indicate that tension modulated both the precision and the efficiency of users' multi-object dragging behavior in VR. Furthermore, no significant learning effect was found across the techniques.

4.3 Testing H1

The Drag Efficiency Index (DEI) was significantly higher for the half-tension technique ($M = 0.95$, $SE = 0.02$) compared to the high-tension technique ($M = 0.90$, $SE = 0.02$), $p = 0.017$, with a small effect size ($r = 0.22$). Additionally, the error rate (ER) was significantly lower for the half-tension technique ($M = 0.07$, $SE = 0.02$) than for the high-tension technique ($M = 0.13$, $SE = 0.02$), $p = 0.024$, also with a small effect size ($r = 0.22$). No significant differences were observed between the two techniques on the remaining performance measures.

Total Completion Time (TCT) was significantly higher for the low-tension technique ($Mdn = 14.79$, $SE = 1.08$) compared to the high-tension technique ($Mdn = 12.78$, $SE = 0.56$), $p = 0.0007$, with a medium effect size ($r = 0.35$). Similarly, the Successful Drag Time (SDT) was significantly higher for the low-tension technique ($Mdn = 7.22$, $SE = 0.34$) than for the high-tension technique ($Mdn = 5.65$, $SE = 0.37$), $p < 0.0001$, with a medium effect size ($r = 0.45$). However, no significant differences were observed between the two techniques in DEI or ER.

Taken together, these results partially support H1. High-tension gesture phrasing improved efficiency over low tension, as reflected in significantly lower TCT and SDT values. However, it was less accurate than half-tension, as indicated by a significantly higher error rate (ER).

4.4 Half- vs. High Tension

We found that the half-tension technique was significantly more accurate than the high-tension technique (lower ER) and also more efficient on one measure (higher DEI). To assess whether the remaining efficiency measures—SDT and TCT—were practically equivalent between the two techniques, we conducted two one-sided tests (TOST) for equivalence [17], using equivalence bounds of $d = \pm 0.8$ (a large effect) and $\alpha = .05$. Results showed statistically significant equivalence for both SDT, $t(104) = 7.09$, $p < .0001$, and TCT, $t(104) = 8.16$, $p < .0001$. These results allow us to reject the presence of effects larger than $|d| = 0.8$, indicating that the efficiency of the half- and high-tension techniques was statistically equivalent on these two measures.

4.5 Testing H2

TCT was significantly higher for low-tension ($Mdn = 14.79$, $SE = 1.08$) compared to the low-tension with feedback technique ($Mdn = 13.51$, $SE = 0.77$), $p = 0.026$, with a small effect size ($r = 0.22$). However, DEI was significantly lower for low-tension ($M = .88$, $SE = .02$) compared to the low-tension with feedback technique ($M = .94$, $SE = 0.02$), $p = 0.017$, with a small effect size ($r = 0.16$). No significant differences were observed between the two techniques on the remaining performance measures. Taken together, these results partially support H2. Visual feedback increased total completion time (TCT), but reduced inefficiencies—more of the drag time went toward successful actions, as reflected in a higher DEI.

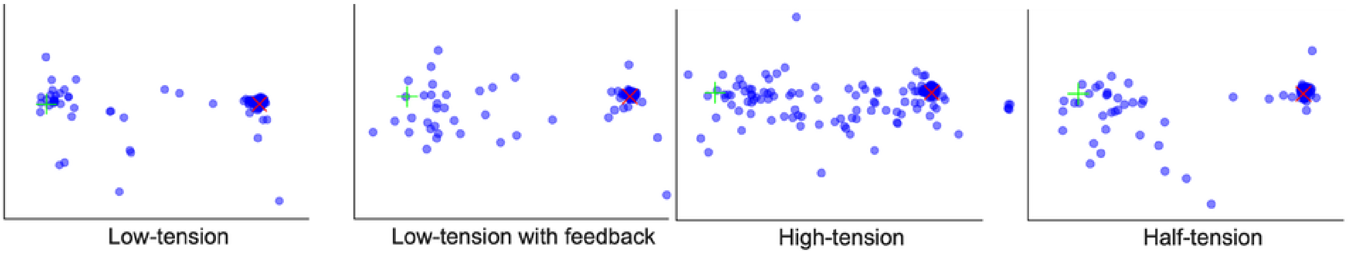

Figure 4: During high tension, dragging errors were distributed along the entire movement path (left to right), while they were more clustered near the origin during low tension.

4.6 Subjective Preference Rankings

The type of tension had a significant effect on participants' preference rankings, $\chi^2(3) = 17.74$, $p < .001$. On average, participants ranked the half-tension condition highest ($M = 1.86$, $SD = 0.85$; Figure 2), followed by low-tension with visual feedback ($M = 2.14$, $SD = 0.85$), high-tension ($M = 2.57$, $SD = 1.29$), and low-tension without feedback ($M = 3.43$, $SD = 0.81$) (Figure 2). A chi-square goodness-of-fit test confirmed that the order of technique exposure did not significantly influence participants' top preference, $\chi^2(3) = 1.29$, $p = .733$.

4.7 Anchoring in Dragging Movements

In the half-tension dragging condition, participants could choose to maintain the left-hand anchoring gesture (a fist), relax it, or release it entirely. We analyzed anchoring behavior to determine whether participants showed a preference for maintaining or releasing the fist gesture. Across all half-tension trials, participants held their left fist for the entire duration in only a small proportion of cases ($M = 0.20$, $SD = 0.37$, $Mdn = 0$). Based on per-trial behavior, 56.7% of trials were classified as mostly release, 11.2% as mixed, and 32.1% as always hold. A Wilcoxon signed-rank test showed that participants were significantly less likely than chance to hold their left fist for the entire drag, $V = 480, p < .001$. The median proportion was 0, indicating that most opted to release their hand.

4.8 Dragging Error Distribution

We analyzed the spatial coordinates where dragging movements ended erroneously across tension conditions (Figure 4). In the high-tension condition, errors were distributed along the entire movement path, suggesting that users were typically able to initiate dragging but occasionally lost control mid-motion. In contrast, errors in the half-tension and low-tension with feedback conditions clustered in the first half of the path, indicating early instability or uncertainty when initiating the gesture. In the low-tension condition without feedback, errors were highly concentrated near the starting point, likely due to mode confusion or failure to initiate the dragging gesture.

5 Discussion

XR-based tools for visualizing large and complex graph and network data are increasingly gaining traction [3, 20]. In this work, we investigated a fundamental interaction—*selecting* multiple nodes and *dragging* them to a target location—common in graph exploration tasks such as reviewing query results [20], comparing subgraphs [3, 28], or highlighting nodes for edge clustering [7]. Our approach builds on the observation that graph exploration often involves reviewing spatial clusters of nodes, either in isolation or in relation to others (e.g., snapshots over time). This structure gives rise to frequent multi-object selection and manipulation needs, where node clusters are visually discrete rather than densely entangled, as in point clouds [24].

In our Experiment, participants selected a subgraph (Figure 3) and then used different bimanual mid-air techniques to drag it to a target location. The **asymmetric asynchronous technique with half tension yielded the highest accuracy and matched the high-tension technique in efficiency** (Figure 2). The poorer accuracy of the high-tension technique may have stemmed from increased physical effort: maintaining both a non-dominant hand fist and a dominant-hand pinch throughout the drag gesture—without haptic support—likely introduced motor instability and fatigue. This is supported by the distribution of drag errors throughout the movement path, suggesting breakdowns during gesture execution rather than at initiation or completion [8]. While prior studies found high-tension techniques more accurate for structured tasks using pen, keyboard, or touch input [13], our results suggest that these findings may not generalize to mid-air input due to the absence of haptic feedback.

The half-tension technique was also rated most preferred by participants. Compared to high-tension, it enabled users to allocate more of their total drag time toward successful attempts, as reflected by a higher Drag Efficiency Index (DEI). Interaction logs showed that participants frequently released their non-dominant hand during the gesture, suggesting a natural strategy for reducing fatigue. Most errors in this condition occurred near the origin, potentially indicating uncertainty in gesture initiation or mode confusion rather than execution issues. A similar error pattern was observed in the low-tension without feedback condition. While introducing visual feedback improved accuracy in the low-tension condition, it potentially introduced a visual processing cost that reduced overall efficiency [6]. Future work could investigate whether auditory or pseudohaptic feedback offers a more effective trade-off—supporting gesture stability and mode clarity without increasing visual demand. Integrating other input modalities such as gaze, haptics, or pen-based interaction with mid-air input may further enhance control and reduce fatigue. For example, multimodal phrasing that leverages gaze or auditory cues for state transitions may enable more expressive and scalable multi-object workflows across immersive applications.

5.1 Limitations

This study was limited to right-handed participants, reducing generalizability to left-handed or ambidextrous users. Experiments were conducted in controlled lab settings with static graph datasets, which may not fully reflect real-world complexity. We used a single VR platform (Oculus Quest 2), and results may vary with other hardware. Tasks were brief and may not capture fatigue effects over prolonged use, especially under high-tension conditions. Future work should explore more diverse user populations, task types, and dynamic or occluded environments to assess broader applicability.

6 Conclusion

This paper presented a controlled study investigating gesture-phrased bimanual mid-air techniques for multi-object manipulation of selected objects in immersive graph exploration. While high-tension phrasing enhanced efficiency, half-tension offered the best trade-off across all performance measures in multi-object dragging. These findings underscore the importance of hand-role coordination and gesture phrasing in shaping effective mid-air interactions.

References

[1] Huidong Bai, Alaeddin Nassani, Barrett Ens, and Mark Billinghurst. 2017. Asymmetric Bimanual Interaction for Mobile Virtual Reality.. In *ICAT-EGVE*. 83–86.

[2] Douglas Bates, Martin Maechler, Ben Bolker, Steven Walker, Rune Haubo Bojesen Christensen, Henrik Singmann, Bin Dai, and Gabor Grothendieck. 2015. Package 'lme4'. *Convergence* 12, 1 (2015).

[3] Michael Burch, Weidong Huang, Mathew Wakefield, Helen C Purchase, Daniel Weiskopf, and Jie Hua. 2020. The state of the art in empirical user evaluation of graph visualizations. *IEEE Access* 9 (2020), 4173–4198.

[4] William Buxton. 1995. Chunking and phrasing and the design of human-computer dialogues. In *Readings in human–computer interaction*. Elsevier, 494–499.

[5] Nikolas Chaconas and Tobias Höllerer. 2018. An evaluation of bimanual gestures on the microsoft hololens. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 1–8.

[6] Debaleena Chattopadhyay and Davide Bolchini. 2014. Understanding visual feedback in large-display touchless interactions: an exploratory study. (2014).

[7] Weiwei Cui, Hong Zhou, Huamin Qu, Pak Chung Wong, and Xiaoming Li. 2008. Geometry-based edge clustering for graph visualization. *IEEE transactions on visualization and computer graphics* 14, 6 (2008), 1277–1284.

[8] Ivan Golod, Felix Heidrich, Christian Möllering, and Martina Ziefle. 2013. Design principles of hand gesture interfaces for microinteractions. In *Proceedings of the 6th International Conference on Designing Pleasurable Products and Interfaces*. 11–20.

[9] Jerônimo Gustavo Grandi, Henrique Galvan Debarba, and Anderson Maciel. 2019. Characterizing asymmetric collaborative interactions in virtual and augmented realities. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 127–135.

[10] Yves Guiard. 1987. Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. *Journal of motor behavior* 19, 4 (1987), 486–517.

[11] Pantea Habibi and Debaleena Chattopadhyay. 2024. Bimanual, mid-air multi-selection techniques in virtual reality. In *2024 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. IEEE, 561–564.

[12] Ivan Herman, Guy Melançon, and M Scott Marshall. 2002. Graph visualization and navigation in information visualization: A survey. *IEEE Transactions on visualization and computer graphics* 6, 1 (2002), 24–43.

[13] Ken Hinckley, Francois Guimbretiere, Georg Apitz, Nicholas Chen, and Maneesh Agrawala. [n. d.]. Phrasing of Multi-Stroke Gestures in Scriboli. *parameters* 1, 1 ([n. d.]), 0.

[14] Ken Hinckley, Randy Pausch, Dennis Proffitt, and Neal F Kassell. 1998. Two-handed virtual manipulation. *ACM Transactions on Computer-Human Interaction (TOCHI)* 5, 3 (1998), 260–302.

[15] Paul Kabbash, William Buxton, and Abigail Sellen. 1994. Two-handed input in a compound task. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 417–423.

[16] Jinwook Kim, Sangmin Park, Qiushi Zhou, Mar Gonzalez-Franco, Jeongmi Lee, and Ken Pfeuffer. 2025. PinchCatcher: Enabling Multi-selection for Gaze+ Pinch. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–16.

[17] Daniël Lakens, Anne M Scheel, and Peder M Isager. 2018. Equivalence testing for psychological research: A tutorial. *Advances in methods and practices in psychological science* 1, 2 (2018), 259–269.

[18] Hyeongmook Lee, Seung-Tak Noh, and Woontack Woo. 2016. Tunnelslice: Free-hand subspace acquisition using an egocentric tunnel for wearable augmented reality. *IEEE Transactions on Human-Machine Systems* 47, 1 (2016), 128–139.

[19] Ning Liu, Dong-sheng Li, Yi-ming Zhang, and Xiong-lve Li. 2020. Large-scale graph processing systems: a survey. *Frontiers of Information Technology & Electronic Engineering* 21, 3 (2020), 384–404.

[20] Richen Liu, Min Gao, Lijun Wang, Xiaohan Wang, Yuzhe Xiang, Aolin Zhang, Jiazhi Xia, Yi Chen, and Siming Chen. 2022. Interactive extended reality techniques in information visualization. *IEEE Transactions on Human-Machine Systems* 52, 6 (2022), 1338–1351.

[21] Roberto A Montano-Murillo, Cuong Nguyen, Rubaiat Habib Kazi, Sriram Subramanian, Stephen DiVerdi, and Diego Martinez-Plasencia. 2020. Slicing-volume: Hybrid 3d/2d multi-target selection technique for dense virtual environments. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 53–62.

[22] Thomas G Plante, Cara Cage, Sarah Clements, and Allison Stover. 2006. Psychological benefits of exercise paired with virtual reality: Outdoor exercise energizes whereas indoor virtual exercise relaxes. *International Journal of Stress Management* 13, 1 (2006), 108.

[23] Siddhartha Sahu, Amine Mhedhbi, Semih Salihoglu, Jimmy Lin, and M Tamer Özsu. 2017. The ubiquity of large graphs and surprising challenges of graph processing. *Proceedings of the VLDB Endowment* 11, 4 (2017), 420–431.

[24] Rasmus Stenholt. 2012. Efficient selection of multiple objects on a large scale. In *Proceedings of the 18th ACM symposium on Virtual reality software and technology*. 105–112.

[25] Radu-Daniel Vatavu and Bogdan-Florin Gheran. 2025. Intermanual Deictics: Uncovering Users' Gesture Preferences for Opposite-Arm Referential Input, from Fingers to Shoulder. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–16.

[26] Andrew Martin Webb. 2017. *Phrasing Bimanual Interaction for Visual Design*. Ph. D. Dissertation.

[27] Andrew M Webb, Michel Pahud, Ken Hinckley, and Bill Buxton. 2016. Wearables as context for guiard-abiding bimanual touch. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. 287–300.

[28] Qian Zhu, Tao Lu, Shunan Guo, Xiaojuan Ma, and Yalong Yang. 2024. CompositingVis: Exploring interactions for creating composite visualizations in immersive environments. *IEEE Transactions on Visualization and Computer Graphics* (2024).