Comparing Convex Region-of-Interest Selection Techniques for Surface Geometry

Bret Jackson
bjackson@macalester.edu
Macalester College
Saint Paul, Minnesota

Kayla Beckham
Macalester College
Saint Paul, Minnesota

Anael Kuperwajs Cohen
Macalester College
Saint Paul, Minnesota

Brianna C. Heggeseth
Macalester College
Saint Paul, Minnesota

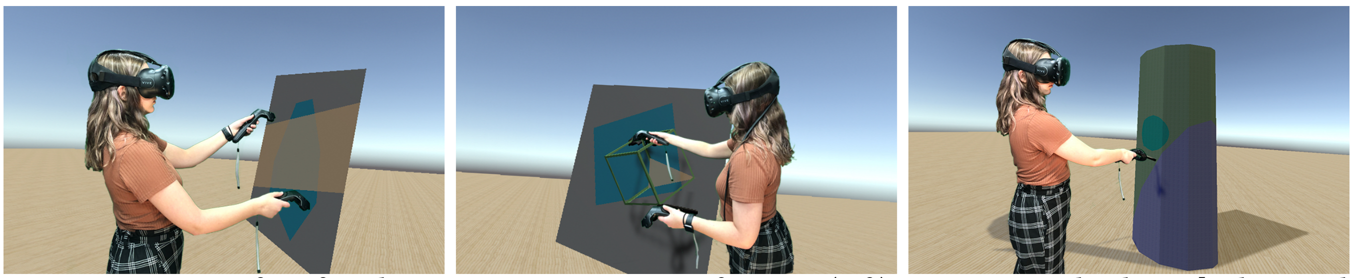

Figure 1: 3D user interfaces for selecting geometry in a region-of-interest: (Left) Yea Big, Yea High Selection[Jackson et al. 2018]; (Center) Two-Handed Volume Cube [Ulinski et al. 2007]; (Right) Slice-n-Swipe [Bacim et al. 2014]
.

ABSTRACT

Selecting 3D regions-of-interest (ROI) in surface geometry is essential for 3D modeling, but few 3D user interfaces using fully manual input for ROI selection exist. Furthermore, their relative performance is not well studied. We present an evaluation comparing three ROI techniques: Volume Cube [Ulinski et al. 2007], Slice-n-Swipe [Bacim et al. 2014], and Yea Big Yea High Selection [Jackson et al. 2018]. Results show that Yea Big Yea High is best for tasks requiring high accuracy and speed, but modifications may be needed for use in dense geometry or with non-convex ROI.

CCS CONCEPTS

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

KEYWORDS

3D User Interfaces, Selection Interfaces, 3D Selection

ACM Reference Format:
Bret Jackson, Kayla Beckham, Anael Kuperwajs Cohen, and Brianna C. Heggeseth. 2019. Comparing Convex Region-of-Interest Selection Techniques for Surface Geometry. In *25th ACM Symposium on Virtual Reality Software and Technology (VRST '19), November 12–15, 2019, Parramatta, NSW, Australia.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3359996.3364258

VRST '19, November 12–15, 2019, Parramatta, NSW, Australia

ACM ISBN 978-1-4503-7001-1/19/11...$15.00
https://doi.org/10.1145/3359996.3364258

1 INTRODUCTION

Driven by the now wide-spread availability of consumer-grade virtual reality (VR) displays, more and more applications that traditionally would have been performed on a desktop display are migrating to VR. For example, artists are increasingly using VR for 3D modeling and other creative pursuits (e.g. with Google Tiltbrush [Inc. 2017], Lift-off [Jackson and Keefe 2016], Cave Painting [Keefe et al. 2001]). In this domain, selection is an essential interaction task for users to be able to quickly and accurately indicate the portions of a 3D model that are the focus for future modeling interactions.

Traditional 3D user interfaces (3DUI) for selecting single or multiple discrete objects, such as the raycast metaphor [Mine 1995] or a virtual hand technique [Mine et al. 1997; Poupyrev et al. 1996], become tedious when selecting hundreds or thousands of triangles in a model. Instead, 3D user interfaces that select a region-of-interest (ROI) must be used. To our knowledge, only four manually controlled — as opposed to structure-aware, which are not generalizable to every model — 3DUI interfaces for ROI selection exist: Tangible Brush [Besançon et al. 2019], Slice-n-Swipe [Bacim et al. 2014], Volume Cube [Ulinski et al. 2007], and Yea Big Yea High Selection [Jackson et al. 2018]. Furthermore, their relative advantages and disadvantages are not yet well explored.

In this paper, we present a controlled study evaluating Slice-n-Swipe, Volume Cube, and Yea Big Yea High — the three manual ROI selection techniques that can be implemented using only readily available consumer-grade VR head-mounted displays and their bundled controllers.

2 RELATED WORK

Although many evaluations exist comparing selection techniques for single objects (e.g. [Bowman et al. 2001; Grossman and Balakrishnan 2006; Lubos et al. 2014]) and multiple objects [Lucas 2005], evaluations comparing ROI selection techniques in a systematic way are lacking. Predominately, evaluations of these techniques only compare to one other baseline. For example, Yu et al. [Yu et al. 2012] compared CloudLasso to Cylinder Selection [Lucas 2005], both representing structural-aware techniques based on a 2D lasso. The most comprehensive study compares the three CAST techniques [Yu et al. 2016] to CloudLasso [Yu et al. 2012] and Cylinder Selection [Lucas 2005]. The techniques were evaluated based on selection speed and the accuracy of the selected region. They found that SpaceCAST was approximately three times as fast as Cylinder Selection and approximately twice as fast as CloudLasso [Yu et al. 2016], while providing similar accuracy.

Besançon et al. [Besançon et al. 2019] present the only study comparing a manual technique, Tangible Brush, to a structural-aware lasso technique, SpaceCAST [Yu et al. 2016]. SpaceCAST was faster to use, but surprisingly less accurate than Tangible Brush.

To our knowledge, no evaluation compares multiple ROI selection techniques for surface geometry based only on fully manual input. In this paper, we propose to fill this gap through a controlled experiment.

3 USER STUDY

3.1 ROI Selection Techniques

Below, we discuss specific implementation details for the three studied selection techniques, and refer the reader to the original papers for a more detailed description of the interface mechanics.

Yea Big, Yea High Selection. [Jackson et al. 2018], shown in Figure 1 Left, uses the metaphor of an infinite cutting plane attached to each hand. At the press of a button, the region between the cutting planes defines the ROI for the selection. The triangles in the selection mesh are then classified as inside or outside the ROI. Triangles intersecting one of the cutting planes are subdivided, re-meshed, and assigned to the appropriate classification. This process can be repeated with the planes in different orientations to progressively refine the volume.

Volume Cube. [Ulinski et al. 2007] uses the metaphor of a six-sided box to define the selection ROI. Ulinski et al. presented three different techniques for controlling the position and orientation of the box. Based on their evaluations in the original paper and follow up studies [Ulinski et al. 2009] that found the "Two-Corner" control method to have the highest selection accuracy, we based our implementation on the "Two-Corner" method. This method attaches opposite corners of the box to the user's hands (See Figure 1 center).

To make for a fairer comparison, we also modified the technique to allow for progressive refinement of the ROI. After an initial selection, subsequent selections operate on only the prior selection region — functioning like a boolean intersection — allowing for selections of arbitrarily shaped ROI beyond a box.

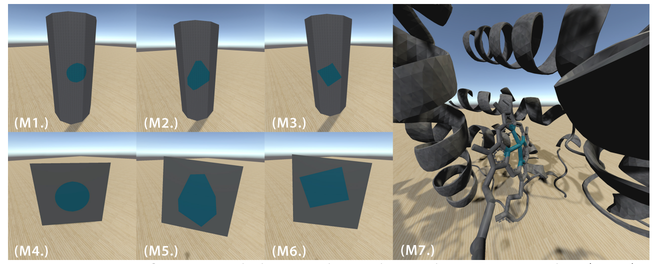

Figure 2: Surface models used in the selection study. (M1.)–(M6.) represent synthetic examples. (M7.) represents a more realistic selection from a hemoglobin molecule.

Slice-n-Swipe. [Bacim et al. 2014] uses a chef's knife metaphor. Originally tracked using a Leap Motion [Ltd. 2019], a slicing movement of a finger creates an infinite cutting plane. Subsequently, a movement perpendicular to the cutting plane is registered as a swipe. Geometry in the swipe direction is removed from the selection. Repeated slice and swiping actions define and filter the geometry to a ROI.

We adapt Slice-n-Swipe to use an HTC Vive controller. A sword line projects one meter along the controller's pointing direction. Moving the controller more rapidly than a threshold (0.03m per frame) indicates a slice plane. A movement roughly perpendicular to the slice plane, while clicking the trigger on the controller held in the dominant hand, creates a swipe.

3.2 Methodology

3.2.1 Participants. We recruited 19 right-handed, undergraduate students. One was unable to complete the study due to technical hardware issues, leaving 18 participants (5 male, 13 female) for analysis. Ages ranged from 18–23 ($M = 20.1$, $SD = 1.1$). Seven participants reported never using VR previously. Nine reported using VR 1–5 times, and two reported more than 20 prior uses. 15 reported prior video game use, with four reporting often or regular use.

3.2.2 Apparatus. The experiment was performed using an HTC Vive head-mounted display with two standard HTC Vive controllers. A 4.4GHz PC running Windows with dual NVidia GTX1070 GPUs drove the display in Unity3D. Participants were allowed to walk in a 2.43m × 2.43m square area free of any obstacles with a virtual floor calibrated to the physical floor.

3.2.3 Task and Procedures. The study used seven surface models (Figure 2). Each model contained a parent object (shown in grey) and a subset selection target (shown in blue). Participants were asked to select the blue target. Before the experiment, participants trained with three additional models (similar in style to M1–M7), and they needed to reach 80% precision and recall on each to advance. The average training time was four minutes for Yea Big Yea High, and six minutes for Slice-n-Swipe and Volume Cube.

3.2.4 Design and Analysis. The study followed a repeated-measures design. The within-subjects independent variables were interaction technique and model. The order of the techniques were fully counter-balanced to avoid learning and order effects. The synthetic

Table 1: Mean accuracy scores and completion times. Standard deviation shown in parentheses.

	Model	F1	MCC	Completion Time
		Slice-n-Swipe		
1	M1: Cylinder/circle	0.94 (0.02)	0.94 (0.02)	74.7s (29.1)
2	M2: Cylinder/irregular	0.95 (0.03)	0.95 (0.03)	63s (31.2)
3	M3: Cylinder/square	0.95 (0.02)	0.95 (0.02)	46.1s (24.8)
4	M4: Plane/circle	0.96 (0.01)	0.95 (0.02)	65.3s (26.6)
5	M5: Plane/irregular	0.96 (0.02)	0.95 (0.02)	61.7s (28.9)
6	M6: Plane/square	0.95 (0.02)	0.94 (0.03)	50.5s (28.6)
7	M7: Molecule	0.88 (0.12)	0.88 (0.11)	101.2s (32.7)
		Volume Cube		
8	M1: Cylinder/circle	0.95 (0.03)	0.95 (0.03)	56.9s (35.8)
9	M2: Cylinder/irregular	0.96 (0.03)	0.96 (0.03)	61.5s (24.8)
10	M3: Cylinder/square	0.96 (0.02)	0.96 (0.02)	51.3s (29.9)
11	M4: Plane/circle	0.95 (0.03)	0.94 (0.04)	62s (36.3)
12	M5: Plane/irregular	0.96 (0.04)	0.94 (0.05)	105.1s (57.7)
13	M6: Plane/square	0.95 (0.05)	0.93 (0.06)	57.1s (34.5)
14	M7: Molecule	0.91 (0.07)	0.91 (0.06)	94.3s (61.1)
		Yea Big Yea High		
15	M1: Cylinder/circle	0.97 (0.01)	0.97 (0.01)	55.8s (25.4)
16	M2: Cylinder/irregular	0.97 (0.02)	0.97 (0.02)	44.8s (19.2)
17	M3: Cylinder/square	0.97 (0.01)	0.97 (0.01)	35.1s (15.8)
18	M4: Plane/circle	0.97 (0.01)	0.97 (0.01)	59.5s (24.6)
19	M5: Plane/irregular	0.98 (0.02)	0.98 (0.02)	61.6s (35.9)
20	M6: Plane/square	0.98 (0.01)	0.98 (0.01)	32s (21.5)
21	M7: Molecule	0.85 (0.21)	0.87 (0.18)	104s (55.6)

models (M1–M6) were presented in random order, while the real-world molecule (M7) was presented last to each participant.

The dependent measures were selection completion time (measured in seconds) and accuracy. To support comparison with Yu et al. [Yu et al. 2016] and Besançon et al. [Besançon et al. 2019], we use two different metrics for accuracy. The first is the F1 score, combining precision (P) and recall (R) based on area of the selection. The second is the Matthew's Correlation Coefficient, MCC. In addition to the true positive (TP), false positive (FP), and false negative (FN) selection areas, the MCC also takes into account true negatives (TN) [Boughorbel et al. 2017].

$$F1=2\cdot\frac{P\cdot R}{P+R} \tag{1}$$

$$MCC=\frac{TP\cdot TN-FP\cdot FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}} \tag{2}$$

Workload was assessed for each interaction technique using NASA's Task Load Index (TLX) [nas 1988]. Perceived exertion was indicated on the Borg CR10 [Borg 1998] scale as a measure of fatigue. This scale is commonly used for measuring exertion during physical activity [Dawes et al. 2005].

Visualizations and numerical summaries (mean and standard deviation) were generated for each outcome by interaction technique and model. The mean outcome differences between interaction techniques were estimated using marginal linear models assuming exchangeable correlation [Liang and Zeger 1986], one fit for each dependent measures. These linear models were fit to the synthetic and real surface models separately due the differences in complexity and existence of outlying values for the real world geometry. This statistical model accounts for different individual ability levels. For each measure, the model included the interaction technique, the selection target shape, surface type, and the number of previous trials with that technique to account for learning. Variables that had no significant effect, based on robust sandwich standard errors [Liang and Zeger 1986], were not included in the final statistical model.

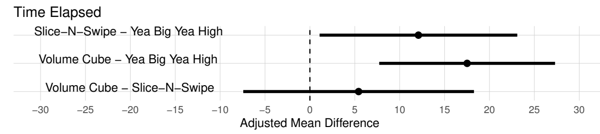

Figure 3: Differences in mean completion time (sec) on synthetic data and 95% confidence intervals estimated from a marginal model adjusted for the surface and target shape.

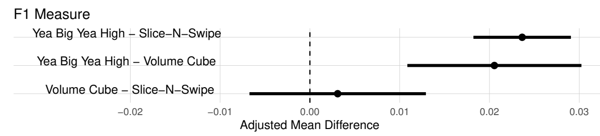

Figure 4: Differences in mean F1 accuracy on synthetic data and 95% confidence intervals estimated from a marginal model that adjusted for the surface shape and experience with the interface.

The NASA TLX and Borg CR10 results were visually and numerically summarized (mean and standard deviation) by interaction technique. We visually explored the NASA TLX subscores by technique to better understand the workload using boxplots; the box represents the values of middle 50% of the scores, the dark middle line represents the median score, and the whiskers represent the extreme non-outlying scores.

3.3 Results

3.3.1 Completion Time. Table 1 includes the mean completion time among the 18 participants for each interaction technique and model. Average completion time is longer for the molecule model as compared to the synthetic scenarios (about 100s as compared to 50-60s). Across the six synthetic models, average completion times are 48s, 60s, and 65s for Yea Big Yea High, Slice-n-Swipe, and Volume Cube, respectively. Figure 3 shows the model-adjusted mean completion time differences between each pair of interaction techniques with corresponding 95% confidence intervals. Completion time is significantly faster for Yea Big Yea High compared to both Slice-n-Swipe and Volume Cube, but the other two techniques are not significantly different in terms of time, keeping surface and target shape fixed.

In terms of the surface model, within an interaction technique, selecting an irregular shape was significantly faster with a cylindrical surface as compared to a flat surface. For both types of surfaces, square targets are significantly faster to select than circles and irregular targets. Irregular targets are faster to select than circles on cylinders but slower to select on flat surfaces.

3.3.2 Accuracy. The molecule model had lower accuracy measures, on average (Table 1). Since F1 and MCC are highly correlated ($r = 0.985$), we present the technique comparison for F1 only. Figure 4 shows the model-adjusted mean F1 differences between each pair of interaction techniques with corresponding 95% confidence intervals. Selection with Yea Big Yea High is significantly more accurate than Slice-n-Swipe and Volume Cube (F1 = 0.975 vs. 0.955, on average), keeping surface and previous trials fixed.

Table 2: Mean and standard deviation for NASA TLX and Borg CR10.

	Selection Technique	NASA TLX Score	Borg CR10
1	Yea Big Yea High	40.21 (14.95)	1.89 (1.24)
2	Slice-n-Swipe	47.72 (15.56)	2.21 (1.44)
3	Volume Cube	60.25 (16.51)	3.69 (2.28)

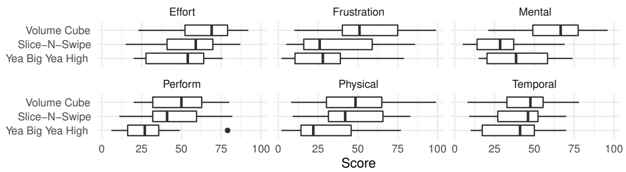

Figure 5: Boxplots (box represents middle 50% of observations and middle line shows median value) of NASA TLX subscores by interaction technique.

3.3.3 Workload and Perceived Exertion. NASA TLX total workload scores are shown in Table 2. Boxplots showing the raw subscores are presented in Figure 5. The data suggest that, on average, Yea Big Yea High requires the lowest total workload (mean = 40.21), then Slice-n-Swipe (mean = 47.72), followed by Volume Cube (mean = 60.25). The subscores generally follow this pattern, except Slice-n-Swipe and Yea Big Yea High were rated much lower on the mental and frustration scales and there are not clear differences in the temporal scale for the three techniques. Yea Big Yea High also had the lowest perceived exertion measured by the Borg CR10, followed by Slice-n-Swipe. Volume Cube exhibited significantly higher mean perceived exertion corresponding with "moderate" exercise.

3.3.4 Qualitative Preferences. 14 participants chose Yea Big Yea High as their preferred technique and 4 chose Slice-n-Swipe. Participants felt that Yea Big Yea High was the most accurate (14 votes, 3 for Slice-n-Swipe, and 1 for Volume Cube) and the most comfortable (12 votes, 6 for Slice-n-Swipe).

4 DISCUSSION

4.1 Study Findings

The results show that Yea Big Yea High outperformed the other two techniques for both completion time and accuracy. Volume Cube and Slice-n-Swipe performed roughly equivalently. Volume Cube was slightly more accurate and slower, but we do not have the statistical power to detect this small effect with the small sample size. Below, we report on the study's findings in more detail.

Yea Big Yea High is best for high accuracy and speed. The lower mean task completion time and higher accuracy for Yea Big Yea High compared to Slice-n-Swipe could be explained in part by the participants' strategies. We observed 11 of the 18 participants using Yea Big Yea High select opposite edges in parallel for the square and circular targets, using fewer operations. In contrast with Slice-n-Swipe, users roughly defined the ROI before refining it with additional slices. Several participants increased accuracy by pre-planning a slice action, slowly moving along the target edge before speeding up to the activation threshold. Even with these accommodations, participants would frequently undo the cut and slice again until the cut plane matched their desired result.

Surprisingly, despite Slice-n-Swipe having slower completion times, 12 participants perceived it as the fastest technique in a post-experiment questionnaire (Yea Big Yea High had six votes, and Volume Cube had zero). We believe this perceptual mismatch can be explained by the game-like feel of the interface. A few participants mentioned that Slice-n-Swipe was the most fun.

Yea Big Yea High's increased accuracy might be explained by the more limited degrees-of-freedom (DOF). It has 6DOF, but this is split to 3DOF per hand, possibly making the selection easier to control. We would expect similar results for Volume Cube, but perhaps this was overwhelmed by the higher mental workload.

Increasing cutting planes leads to higher mental load. Volume Cube had the most NASA TLX mental workload, followed by Yea Big Yea High and then Slice-n-Swipe. This correlates to the increased focus needed for multiple cutting planes. For Yea Big Yea High, participants needed to maintain both hands wider than the target area to keep it within the ROI. Volume Cube further increases the mental requirements because users must focus on all six sides of the box at once. Volume Cube also has the highest perception of exertion on the Borg CR10 [Borg 1998]. Yea Big Yea High has less perceived exertion than Slice-n-Swipe. Despite Slice-n-Swipe allowing a user to rest their hands between slices, these results can be explained by the inaccuracy of the slicing gesture sometimes requiring multiple slices and the greater movement needed to activate it.

Constrained cutting planes are best for dense models. We attribute the slower completion times and less accuracy for the molecule model (M7) to the increased complexity of the dense real-world example, causing participants to spend more time exploring. Completion times also increased due to more physical navigation to filter the selection along the depth axis in front and behind the target for Slice-n-Swipe and Yea Big Yea High. These depth selections were not always optimal, and we observed several participants end the M7 trial without noticing that parts of the molecule were selected outside the target area, leading to lower accuracy. These observations were most common for Yea Big Yea High and Slice-n-Swipe, which use infinite cutting planes. For dense models, Volume Cube may provide a more accurate and intuitive approach since it automatically constrains the selection by the front and back planes of the selection box. Further study is needed.

Performance for ROI selection may depend on visual feedback. A virtual sword was used for visual feedback in the Slice-n-Swipe chef's knife metaphor. While this sword worked particularly well when roughing out the ROI, it was less successful for making precise slices. We observed users focusing on the end of the sword rather than their hand. Occasionally this led to difficulty activating the slice gesture which depended on hand movement. Instead users would keep the position of their hand constant but just rotate the sword to indicate a slice. Future Slice-n-Swipe interfaces should consider different visual feedback or enable the slice gesture based on the motion at the end of the sword rather than the hand.

Users favor selections with more false positives than false negatives. The majority of participants preferred selecting a ROI that was larger than the target area rather than excluding part of it. Undo operations were performed most often when users selected a ROI edge inside the target region. The design of future ROI selection techniques should take this user preference into account.

5 LIMITATIONS AND CONCLUSION

We have presented a controlled evaluation comparing three manual ROI selection techniques. Yea Big Yea High Selection was found to be the fastest and most accurate technique, although some participants preferred Slice-n-Swipe, which is able to select larger than body-scale ROI. One limitation of all the studied techniques is that the ROI is limited to convex shapes. In future work, we plan to develop techniques applicable to non-convex ROI. In conclusion, our findings indicate that future designers of manual ROI selection techniques should: (1) consider limiting the user's focus to one or two regions at once, (2) provide appropriate visual feedback to improve control, and (3) consider alternative approaches to select the ROI's bounds along the depth axis, particularly for dense models.

ACKNOWLEDGMENTS

This research was supported in part through grants from the Clare Booth Luce Foundation.

REFERENCES

1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Human Mental Workload*, Peter A. Hancock and Najmedin Meshkati (Eds.). Advances in Psychology, Vol. 52. North-Holland, 139–83.

F. Bacim, M. Nabiyouni, and D. A. Bowman. 2014. Slice-n-Swipe: A free-hand gesture user interface for 3D point cloud annotation. In *IEEE Symposium on 3D User Interfaces (3DUI)*. 185–186.

Lonni Besançon, Mickael Sereno, Mehdi Ammi, Lingyun Yu, and Tobias Isenberg. 2019. Hybrid Touch/Tangible Spatial 3D Data Selection. *Computer Graphics Forum* 38, 3 (June 2019), 553–567. https://doi.org/10.1111/cgf.13710

Gunnar Borg. 1998. *Borgś perceived exertion and pain scales.* Human kinetics.

Sabri Boughorbel, Fethi Jarray, and Mohammed El-Anbari. 2017. Optimal classifier for imbalanced data using Matthews Correlation Coefficient metric. *PLOS ONE* 12, 6 (06 2017), 1–17. https://doi.org/10.1371/journal.pone.0177678

Doug A Bowman, Donald B Johnson, and Larry F Hodges. 2001. Testbed evaluation of virtual environment interaction techniques. *Presence: Teleoperators & Virtual Environments* 10, 1 (2001), 75–95.

Helen N Dawes, Karen L Barker, Janet Cockburn, Neil Roach, Oona Scott, and Derick Wade. 2005. BorgâĂŹs rating of perceived exertion scales: do the verbal anchors mean the same for different clinical groups? *Archives of physical medicine and rehabilitation* 86, 5 (2005), 912–916.

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*. ACM, 3–12.

Google Inc. 2017. Tilt Brush. https://www.tiltbrush.com/

Bret Jackson, Brighten Jelke, and Gabriel Brown. 2018. Yea Big, Yea High: A 3D User Interface for Surface Selection by Progressive Refinement in Virtual Environments. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 320–326.

Bret Jackson and Daniel F. Keefe. 2016. Lift-Off: Using Reference Imagery and Freehand Sketching to Create 3D Models in VR. *IEEE Transactions on Visualization and Computer Graphics* 22, 4 (2016), 1442–1451.

Daniel F. Keefe, Daniel Acevedo Feliz, Tomer Moscovich, David H. Laidlaw, and Joseph J. LaViola Jr. 2001. CavePainting: A Fully Immersive 3D Artistic Medium and Interactive Experience. In *Proceedings of I3D 2001*. 85–93.

Kung-Yee Liang and Scott L. Zeger. 1986. Longitudinal data analysis using generalized linear models. *Biometrika* 73, 1 (1986), 13–22. https://doi.org/10.1093/biomet/73.1.13

Ultrahaptics Ltd. 2019. Leap Motion. https://www.leapmotion.com/

Paul Lubos, Gerd Bruder, and Frank Steinicke. 2014. Analysis of direct selection in head-mounted display environments. In *2014 IEEE Symposium on 3D User Interfaces (3DUI)*. IEEE, 11–18.

John Finley Lucas. 2005. *Design and Evaluation of 3D Multiple Object Selection Techniques.* Master's thesis. Virginia Polytechnic Institute and State University.

Mark R. Mine. 1995. *Virtual Environment Interaction Techniques.* Technical Report. Chapel Hill, NC, USA.

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*. ACM, 19–26. https://doi.org/10.1145/258734.258747

Ivan Poupyrev, Mark Billinghurst, Suzanne Weghorst, and Tadao Ichikawa. 1996. The Go-go Interaction Technique: Non-linear Mapping for Direct Manipulation in VR. In *Proceedings of the 9th Annual ACM Symposium on User Interface Software and Technology (UIST '96)*. ACM, New York, NY, USA, 79–80.

A. Ulinski, C. Zanbaka, Z. Wartell, P. Goolkasian, and L. F. Hodges. 2007. Two Handed Selection Techniques for Volumetric Data. In *IEEE Symposium on 3D User Interfaces*. 107–114.

Amy C Ulinski, Zachary Wartell, Paula Goolkasian, Evan A Suma, and Larry F Hodges. 2009. Selection performance based on classes of bimanual actions. In *IEEE Symposium on 3D User Interfaces (3DUI)*. IEEE, 51–58.

Lingyun Yu, Konstantinos Efstathiou, Petra Isenberg, and Tobias Isenberg. 2012. Efficient Structure-Aware Selection Techniques for 3D Point Cloud Visualizations with 2DOF Input. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (Dec. 2012), 2245–2254. https://doi.org/10.1109/TVCG.2012.217

Lingyun Yu, Konstantinos Efstathiou, Petra Isenberg, and Tobias Isenberg. 2016. CAST: Effective and efficient user interaction for context-aware selection in 3D particle clouds. *IEEE transactions on visualization and computer graphics* 22, 1 (2016), 886–895.