Cross-Reality 3D Sketching: Human-Computer Interfaces to Support Working in Extended Reality

Bret Jackson

Carleton College, May 22, 2025

Macalester College

Carleton Class of 2008

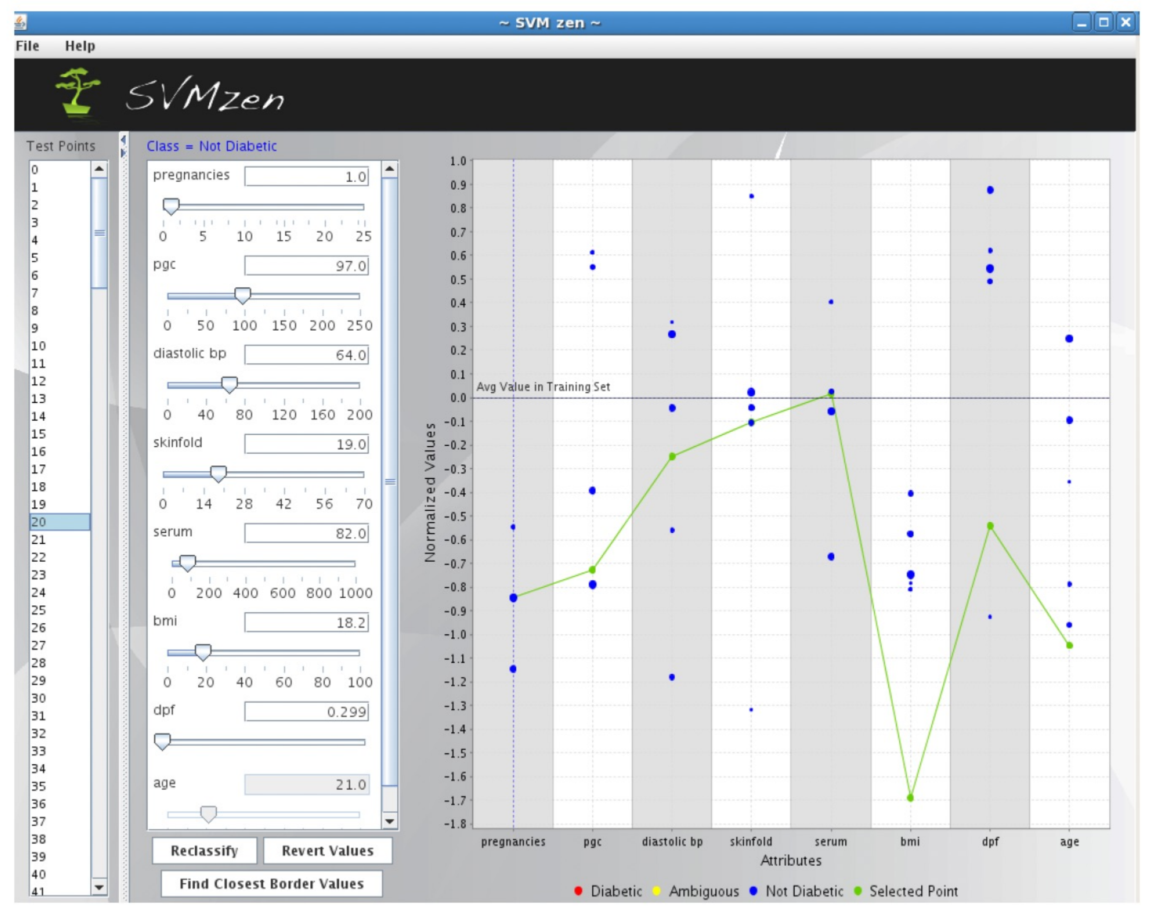

Understanding Support Vector Machine Classifications: A Local Approach

David Barbella
Carleton College
300 North College St.
Northfield, Minnesota
barbelld@carleton.edu

Sami Benzaid
Carleton College
300 North College St.
Northfield, Minnesota
benzaids@carleton.edu

Janara Christensen
Carleton College
300 North College St.
Northfield, Minnesota
christej@carleton.edu

Bret Jackson
Carleton College
300 North College St.
Northfield, Minnesota
jacksonb@carleton.edu

Victor Qin
Carleton College
300 North College St.
Northfield, Minnesota
qinv@carleton.edu

David Musicant
Carleton College
1 North College St.
Northfield, Minnesota
dmusican@carleton.edu

ABSTRACT

Support vector machines are a valuable tool for making classifications, but their black-box nature means that they lack the natural explanatory value that many other classifiers possess. In this paper, we suggest two novel methods for providing insight into local classifications produced by a support vector machine. In the first, we report the support vectors most influential in the final classification for a particular test point. In the second, we determine which features of that test point would need to be altered (and by how much) in order to be placed on the separating surface between the two classifications. We call the latter technique "border classification." In addition to introducing these explanatory techniques, we also present a free-for-download software tool that enables users to visualize these insights graphically.

Categories and Subject Descriptors

I.5.2 [**Pattern Recognition**]: Design Methodology—*Classifier design and evaluation*; H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*Graphical user interfaces, Screen design* design and evaluation]

General Terms

Algorithms, Design, Human Factors

Keywords

Support vector machine, classification, explanation

1. INTRODUCTION

Support vector machines (SVMs) [5, 17] are a well-known supervised learning technique for performing binary classification. SVMs are very accurate and generalize well to a wide range of applications. Because support vector machines are "black-box" classifiers, the decisions they make are not always easily explainable. By this we mean that the model produced does not naturally provide any useful intuitive reasons about why a particular point is classified in one class rather than another. For instance, consider an SVM used by a bank to determine to whom they will loan money. If a customer's loan application is rejected and they would like to know why, then it is not very useful to only be able to say that the algorithm came back with a number lower than some required threshold. It would be much more satisfactory to the customer to be able to tell them that they were denied credit because their income is too low and they have six outstanding loans. Decision trees can provide this level of explainability, but cannot always achieve results as accurate as those produced by an SVM. In addition to applications in the credit industry, understandable classifiers can be used to assist in medical diagnoses.

Recently Fung et al. [6] described the importance of understandable classifications for computer aided diagnosis (CAD). CAD systems provide analysis and interpretation of medical images. Although these systems have been shown to be highly successful in both research labs and clinical settings, the U.S. Government has been reluctant to allow CAD systems to be used outside of experimental settings. Furthermore, doctors are also reluctant to trust a diagnosis without receiving a convincing reason for it. Without understandable reasons for the diagnoses produced by these systems, they face problems when undergoing regulation and acceptance in the medical community.

We propose two techniques to explain classifications provided by support vector machine classifiers for continuous data, without sacrificing accuracy. A key distinction between our approach and other approaches for explain SVM classifiers [6, 12, 14] is that our techniques explain decisions at the *local level*, i.e. for an individual test point. The first technique involves finding the most relevant support vectors for an individual point, that is those that were most influential in determining the class into which the point was placed. The second technique involves determining the minimal change necessary in a test point's features to place that





Ultralingua, inc.

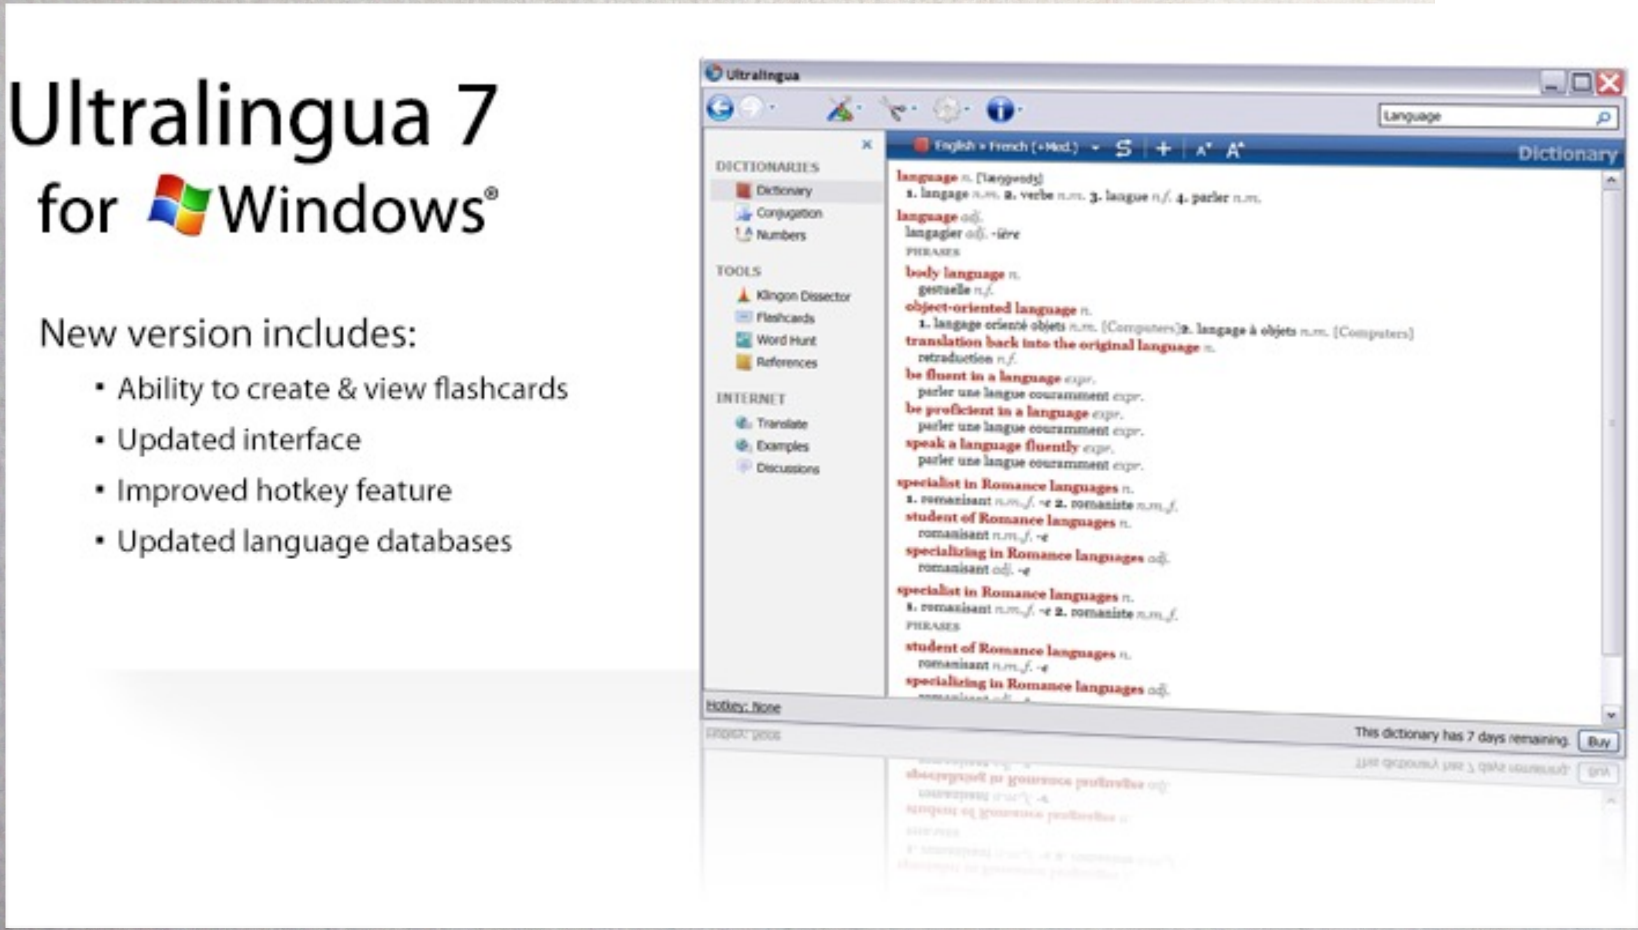

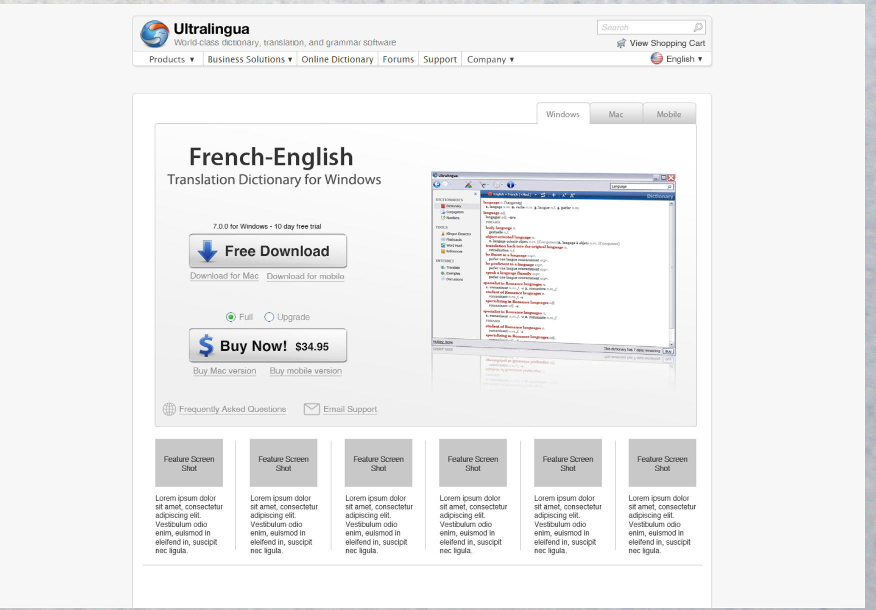

Interactive
Visualization Lab
UMN

ivlab.cs.umn.edu

Macalester







Cross-Reality 3D Sketching:
Human-Computer Interfaces to Support Working in Extended Reality

Bret Jackson


BLUE HAWK
BLUE HAWK

Lift-Off: Using Reference Imagery and Freehand Sketching to Create 3D Models in VR

Two problems limit creativity and applicability of 3D modeling tools to VR:
(1) the lack of control for freehand modeling
(2) the difficulty of starting from scratch

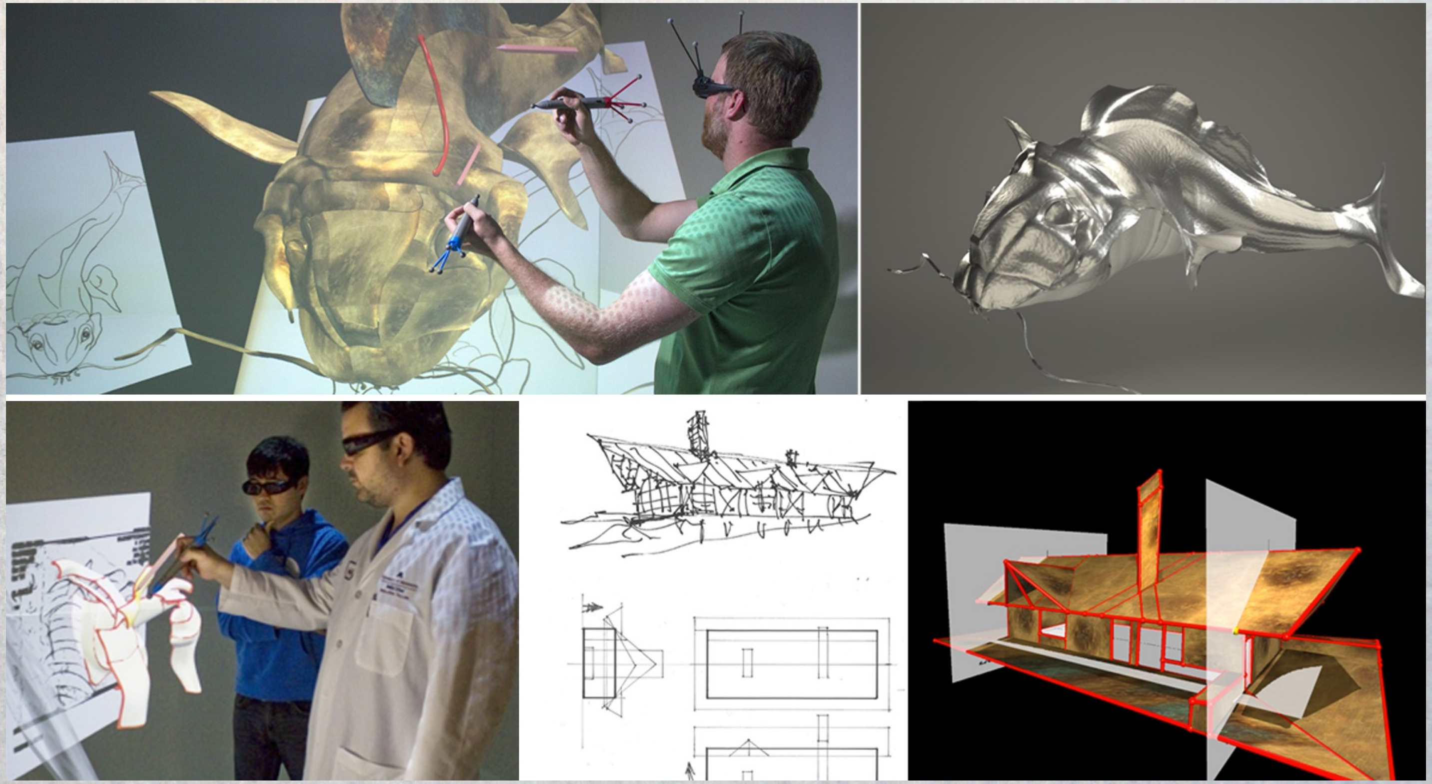

Jackson, B., & Keefe, D. F.. 2016. Lift-Off: Using Reference Imagery and Freehand Sketching to Create 3D Models in VR. *IEEE Transactions on Visualization and Computer Graphics* 22, 4, 1442–1451. Also presented at the Best of IEEE VR session at SIGGRAPH 2016 DOI:https://doi.org/10.1109/TVCG.2016.2518099.


Slide Picker

LIFT-OFF: WORKFLOW

Rotation of styluses maps to curve handles

Guide curve serves as visual feedback and functions like a magnetic rope to move the selection curve to a contour in the sketch

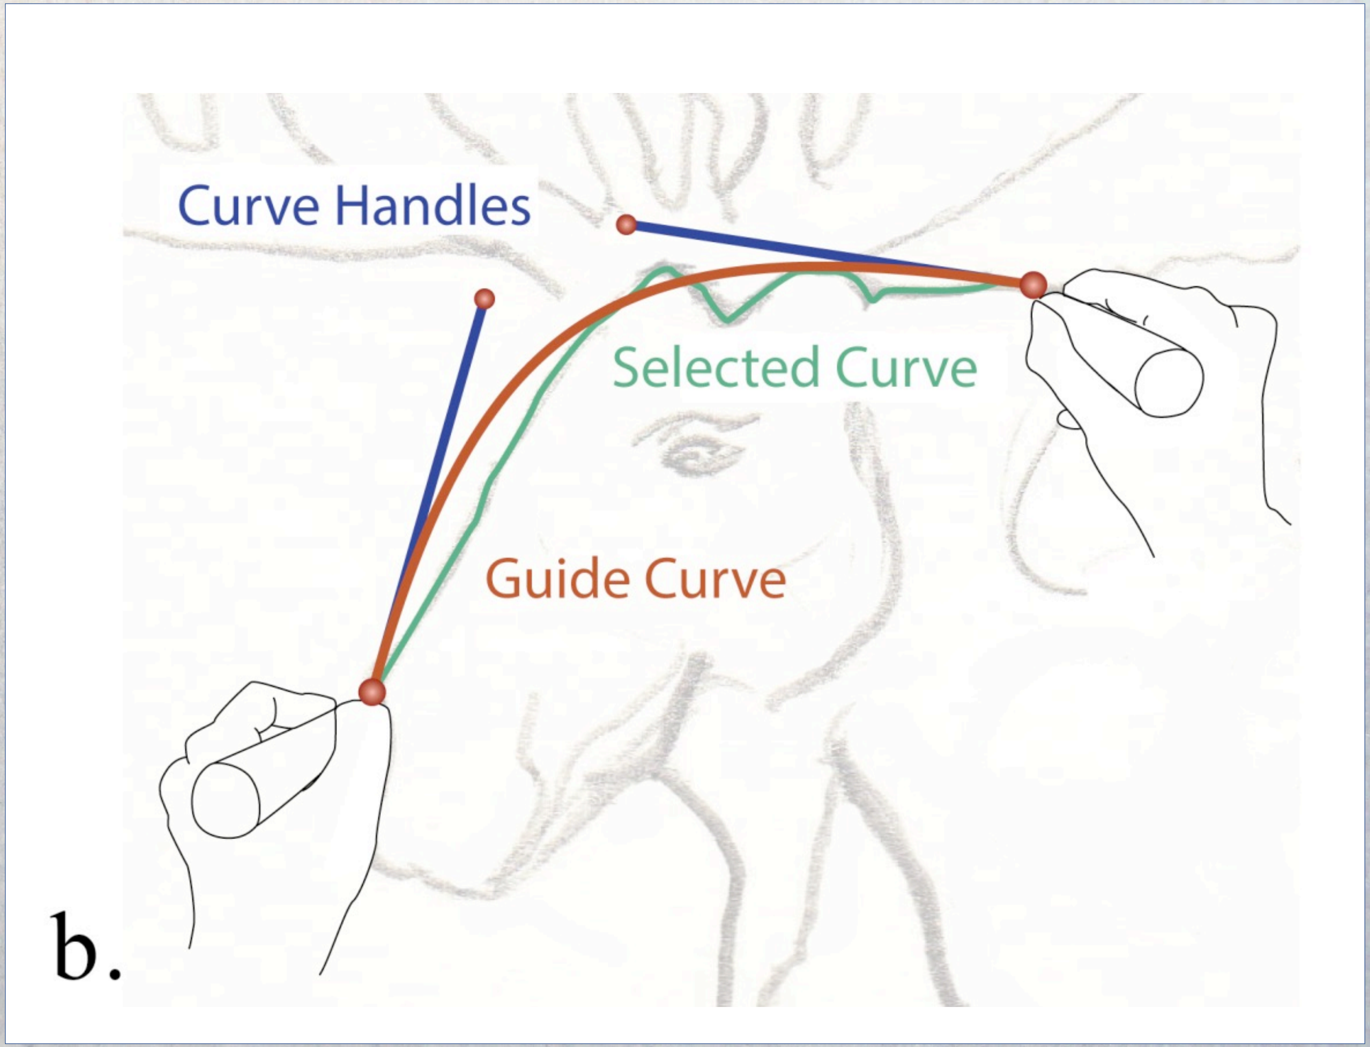

Yea Big, Yea High: A 3D User Interface for Surface Selection by Progressive Refinement in Virtual Environments

Bret Jackson, Brighten Jelke, and Gabriel Brown. "Yea Big Yea High: A 3D User Interface for Surface Selection by Progressive Refinement in Virtual Environments". In: Proceedings of IEEE Virtual Reality. 2018.

Bret Jackson, Kayla Beckham, Anael Kuperwajs Cohen, Brianna C. Heggeseth. "Comparing Convex Region-of-Interest Selection Techniques for Surface Geometry". In: Proceedings of the 25th ACM Symposium on Virtual Reality Software and Technology (VRST). 2019.

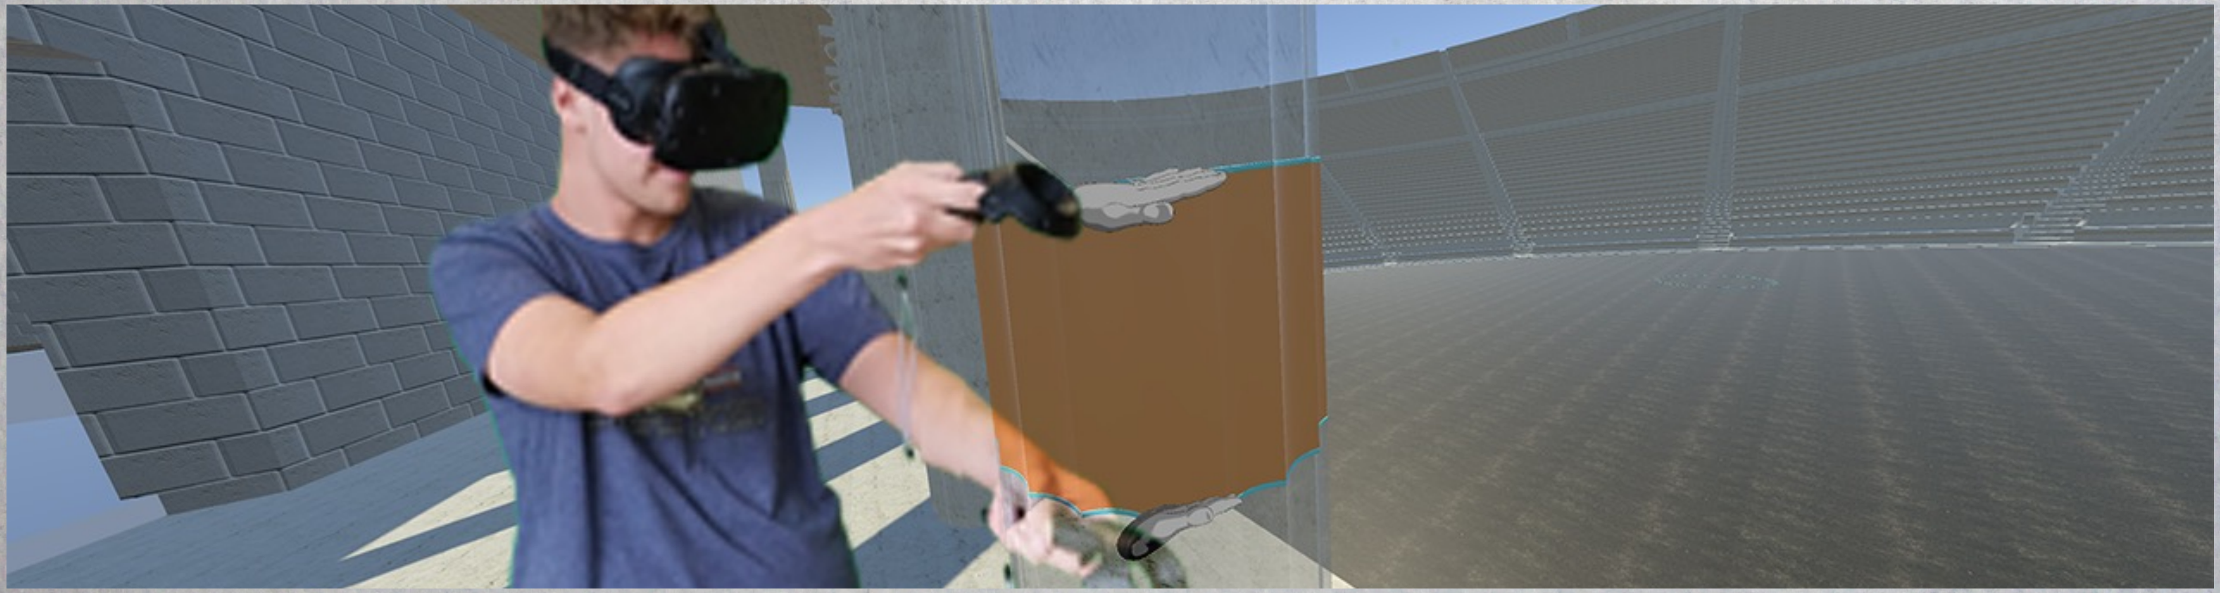

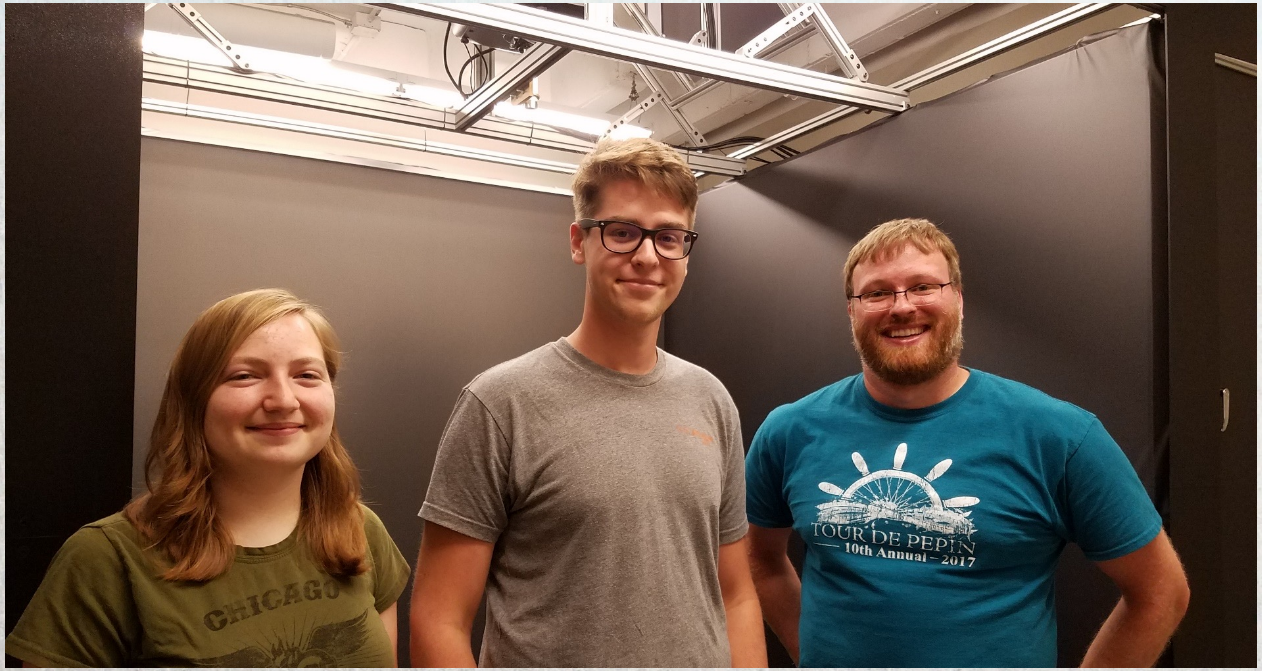

Brighten Jelke
Gabriel Brown

Anael Kuperwajs Cohen
Kalya Beckham

Brianna Heggeseth

Problem: How to efficiently select 3D subspaces?

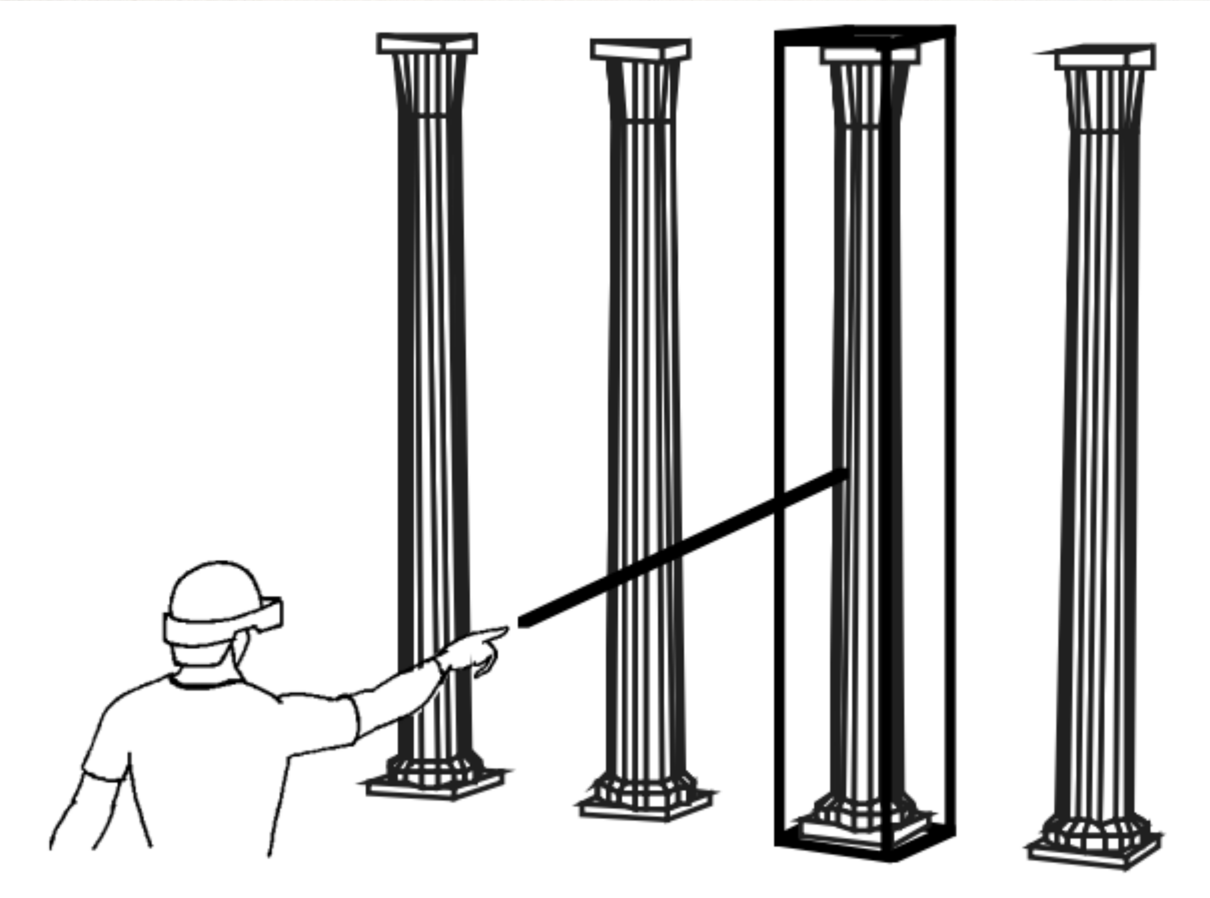

Mine, 1995

Interaction Design

Selection boundaries are determined by two infinite cutting planes attached to the user's hands

- Bi-manual, symmetric-synchronous actions
- 3-DOF per hand

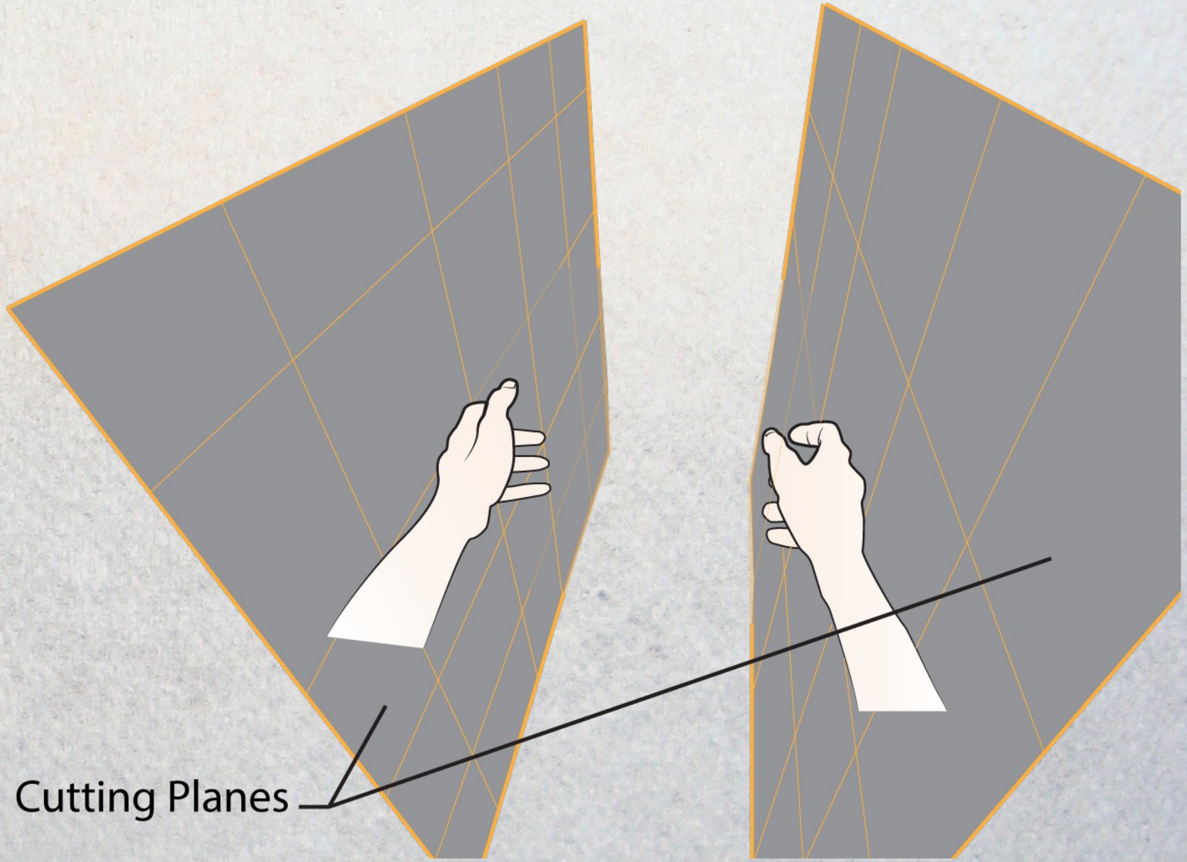

Surface Selection

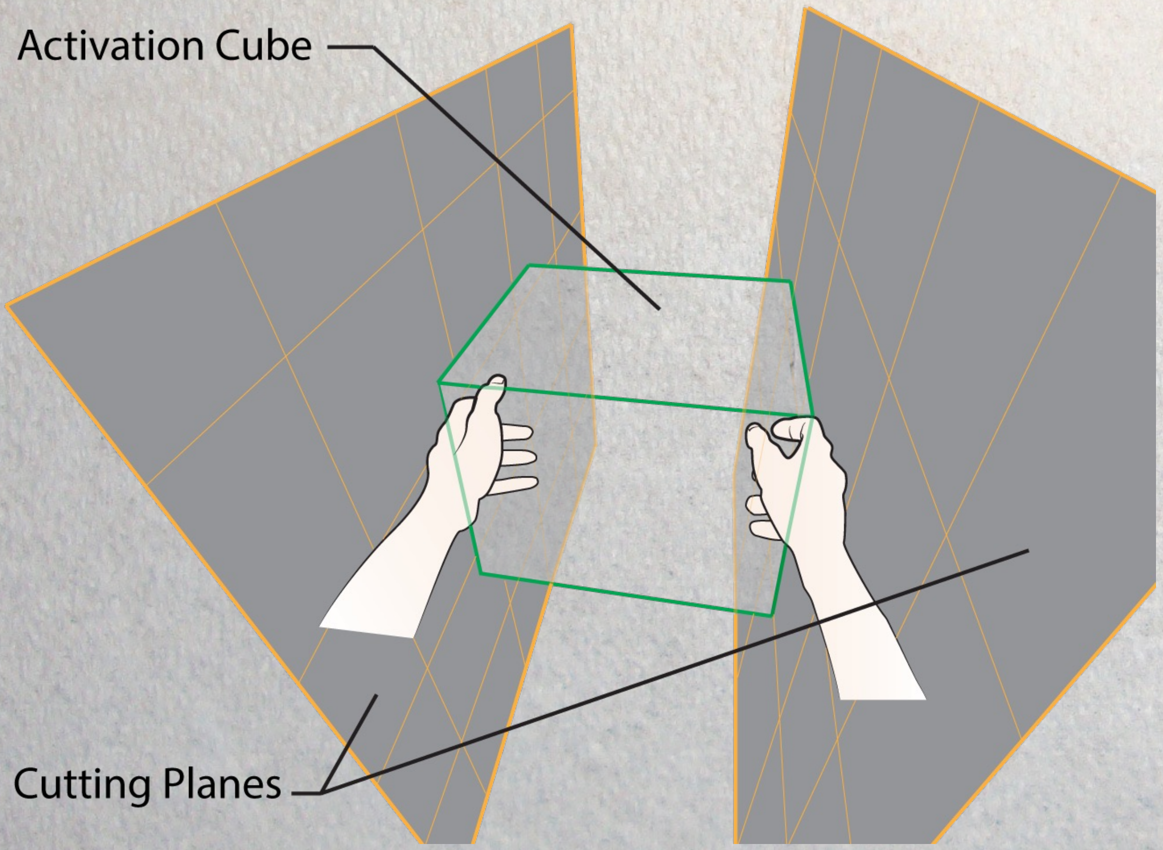

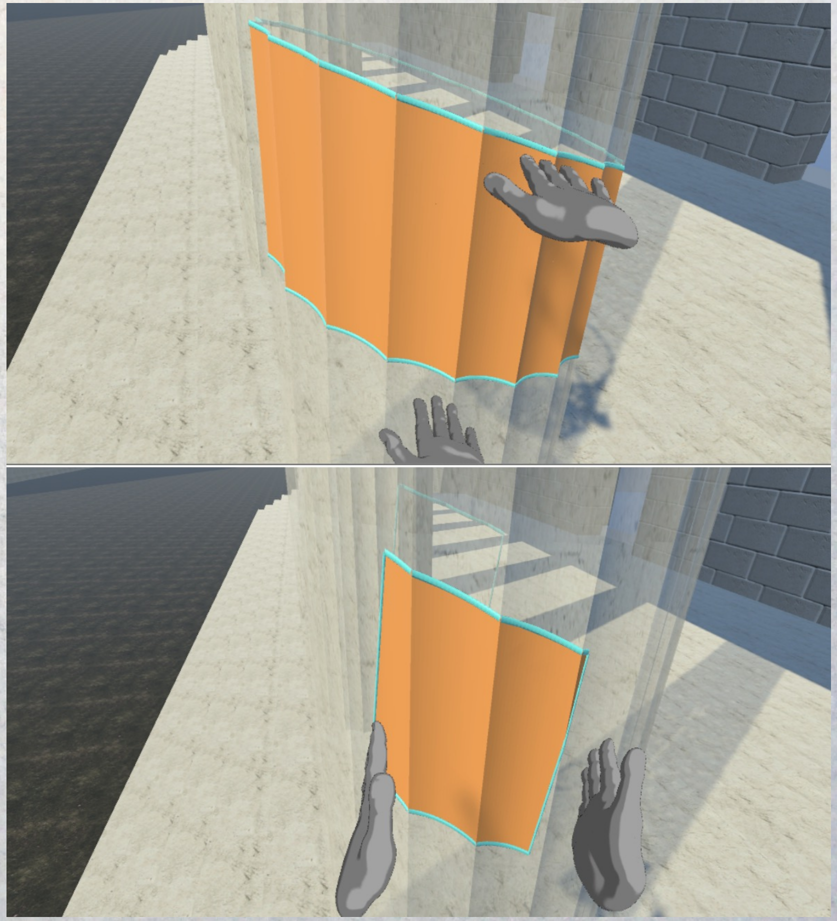


A SELECTION INTERFACE
DESIGNED FOR VIRTUAL ENVIRONMENTS

Progressive refinement

The selection action can be repeated to progressively refine the selection

Enables more complex selection volumes

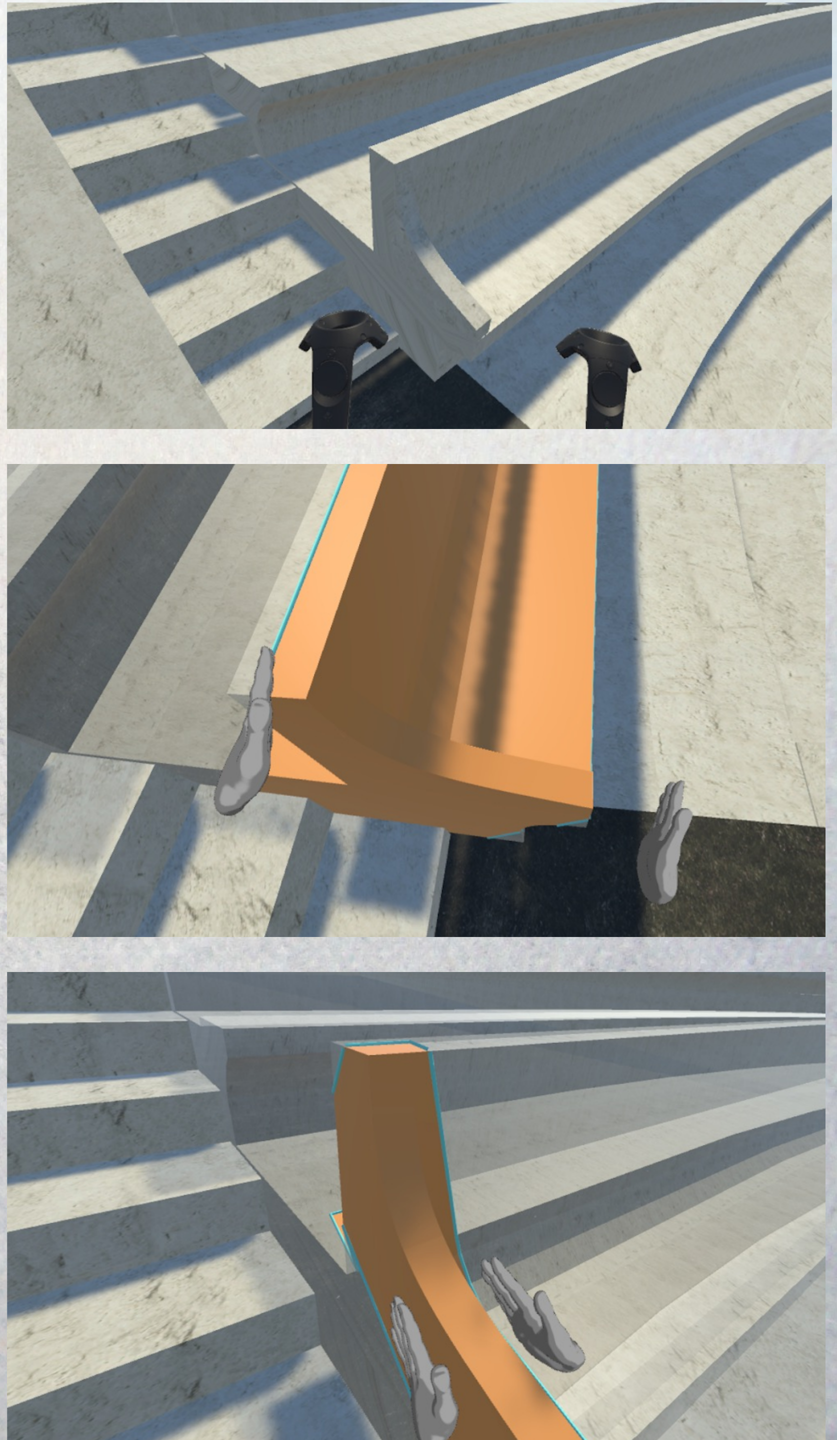

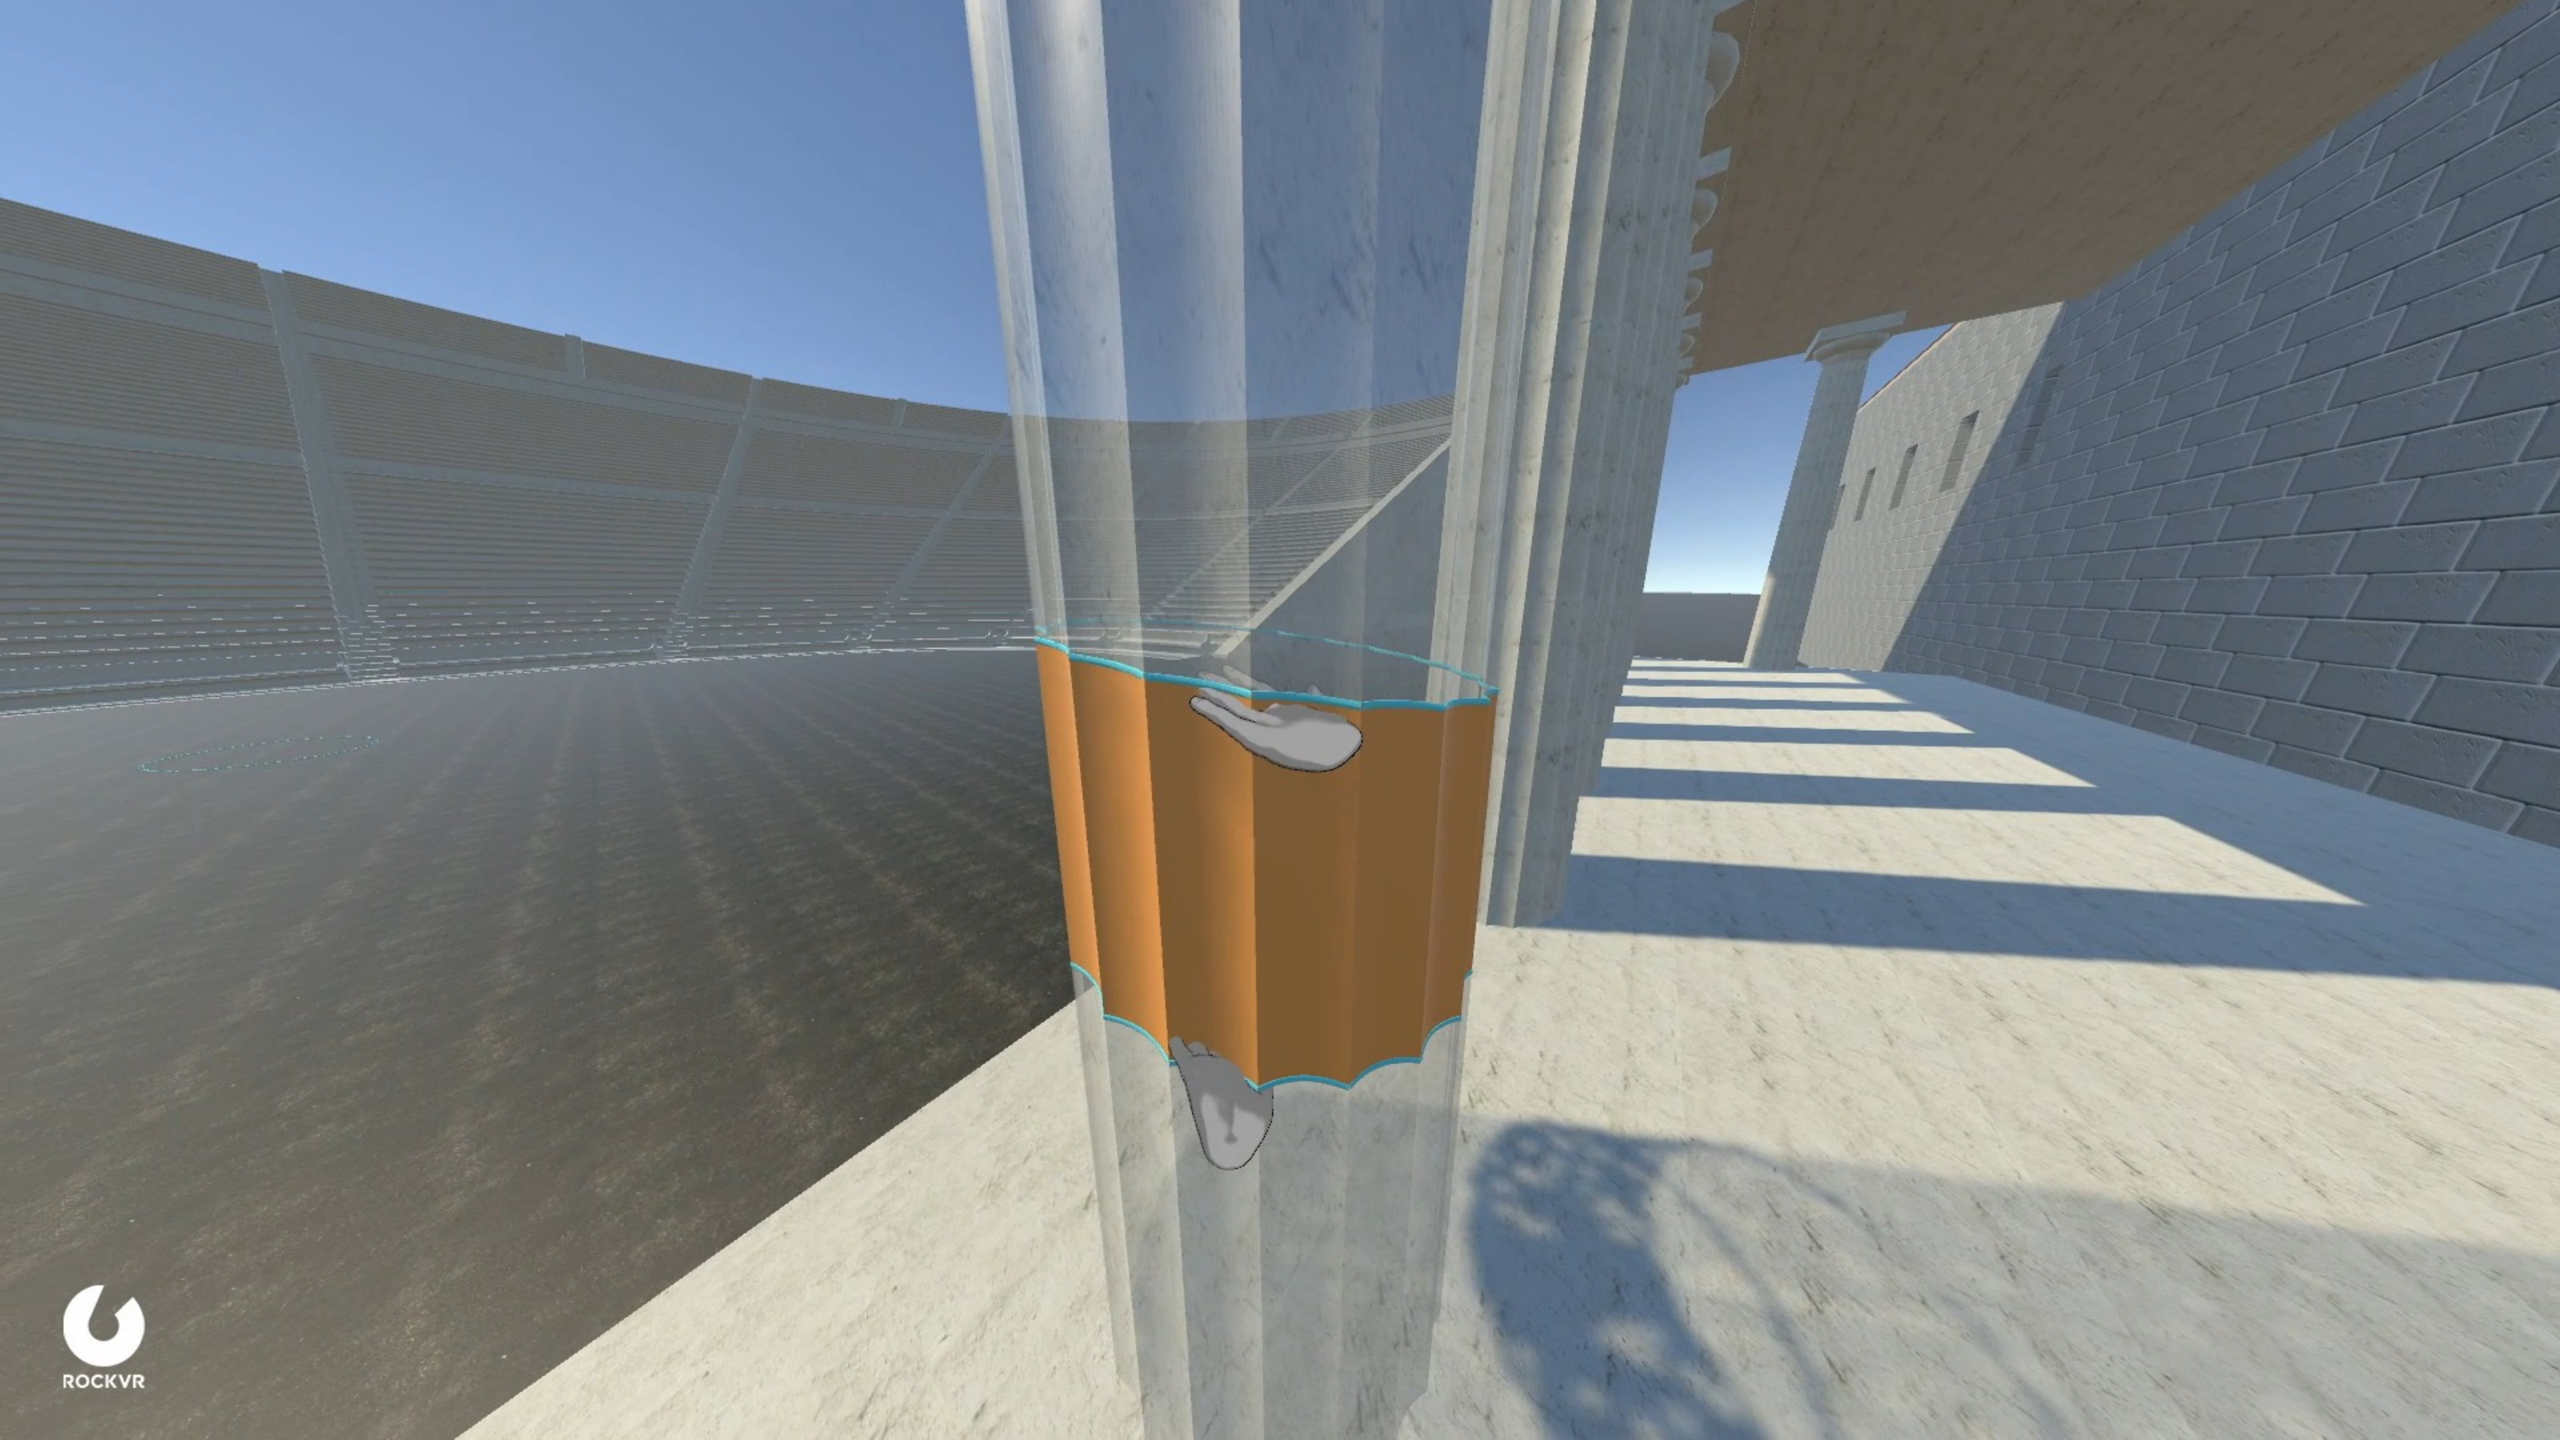
ROCKVR

Implementation

Implemented in Unity3D

Uses a state design pattern with two states:
- Navigation State
- Selection State

In selection state, each mesh intersecting the activation cube is divided into selected and unselected submeshes

Algorithm 1: Selection algorithm implemented using Unity3D

```
for each object colliding with the activation cube do
    for each triangle in the object's mesh do
        if left hand or right hand plane intersects the triangle
         then
            split triangle at intersection and re-tessellate;
            add new triangles to the Selected and Unselected
              submeshes;
        else if the triangle is in the normal half-space of both
         planes then
            add it to Selected submesh;
        else
            add it to the Unselected submesh;
```

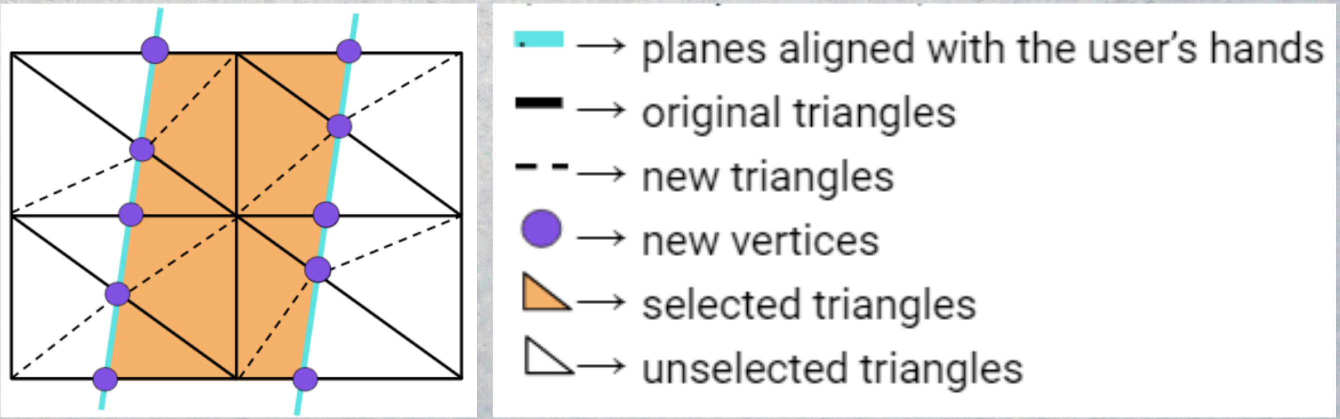

Drawing from Life

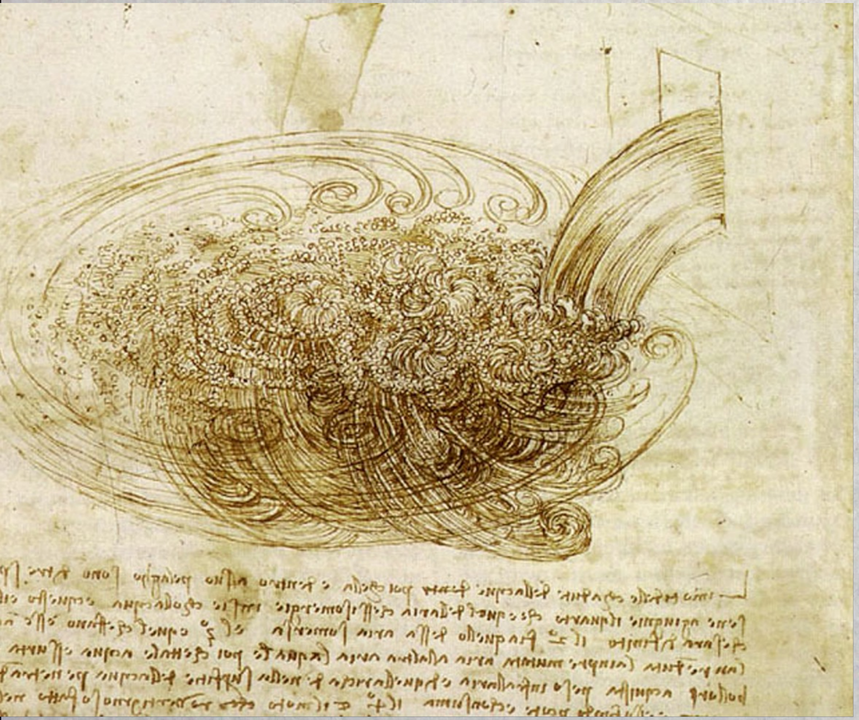

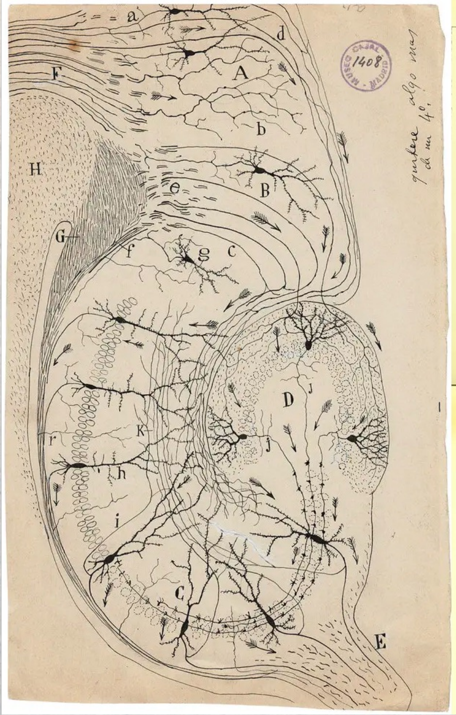

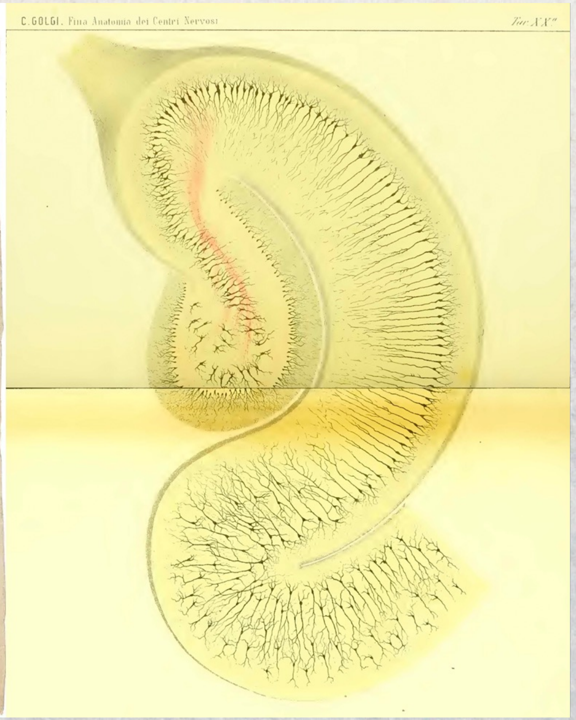

Drawing from Life

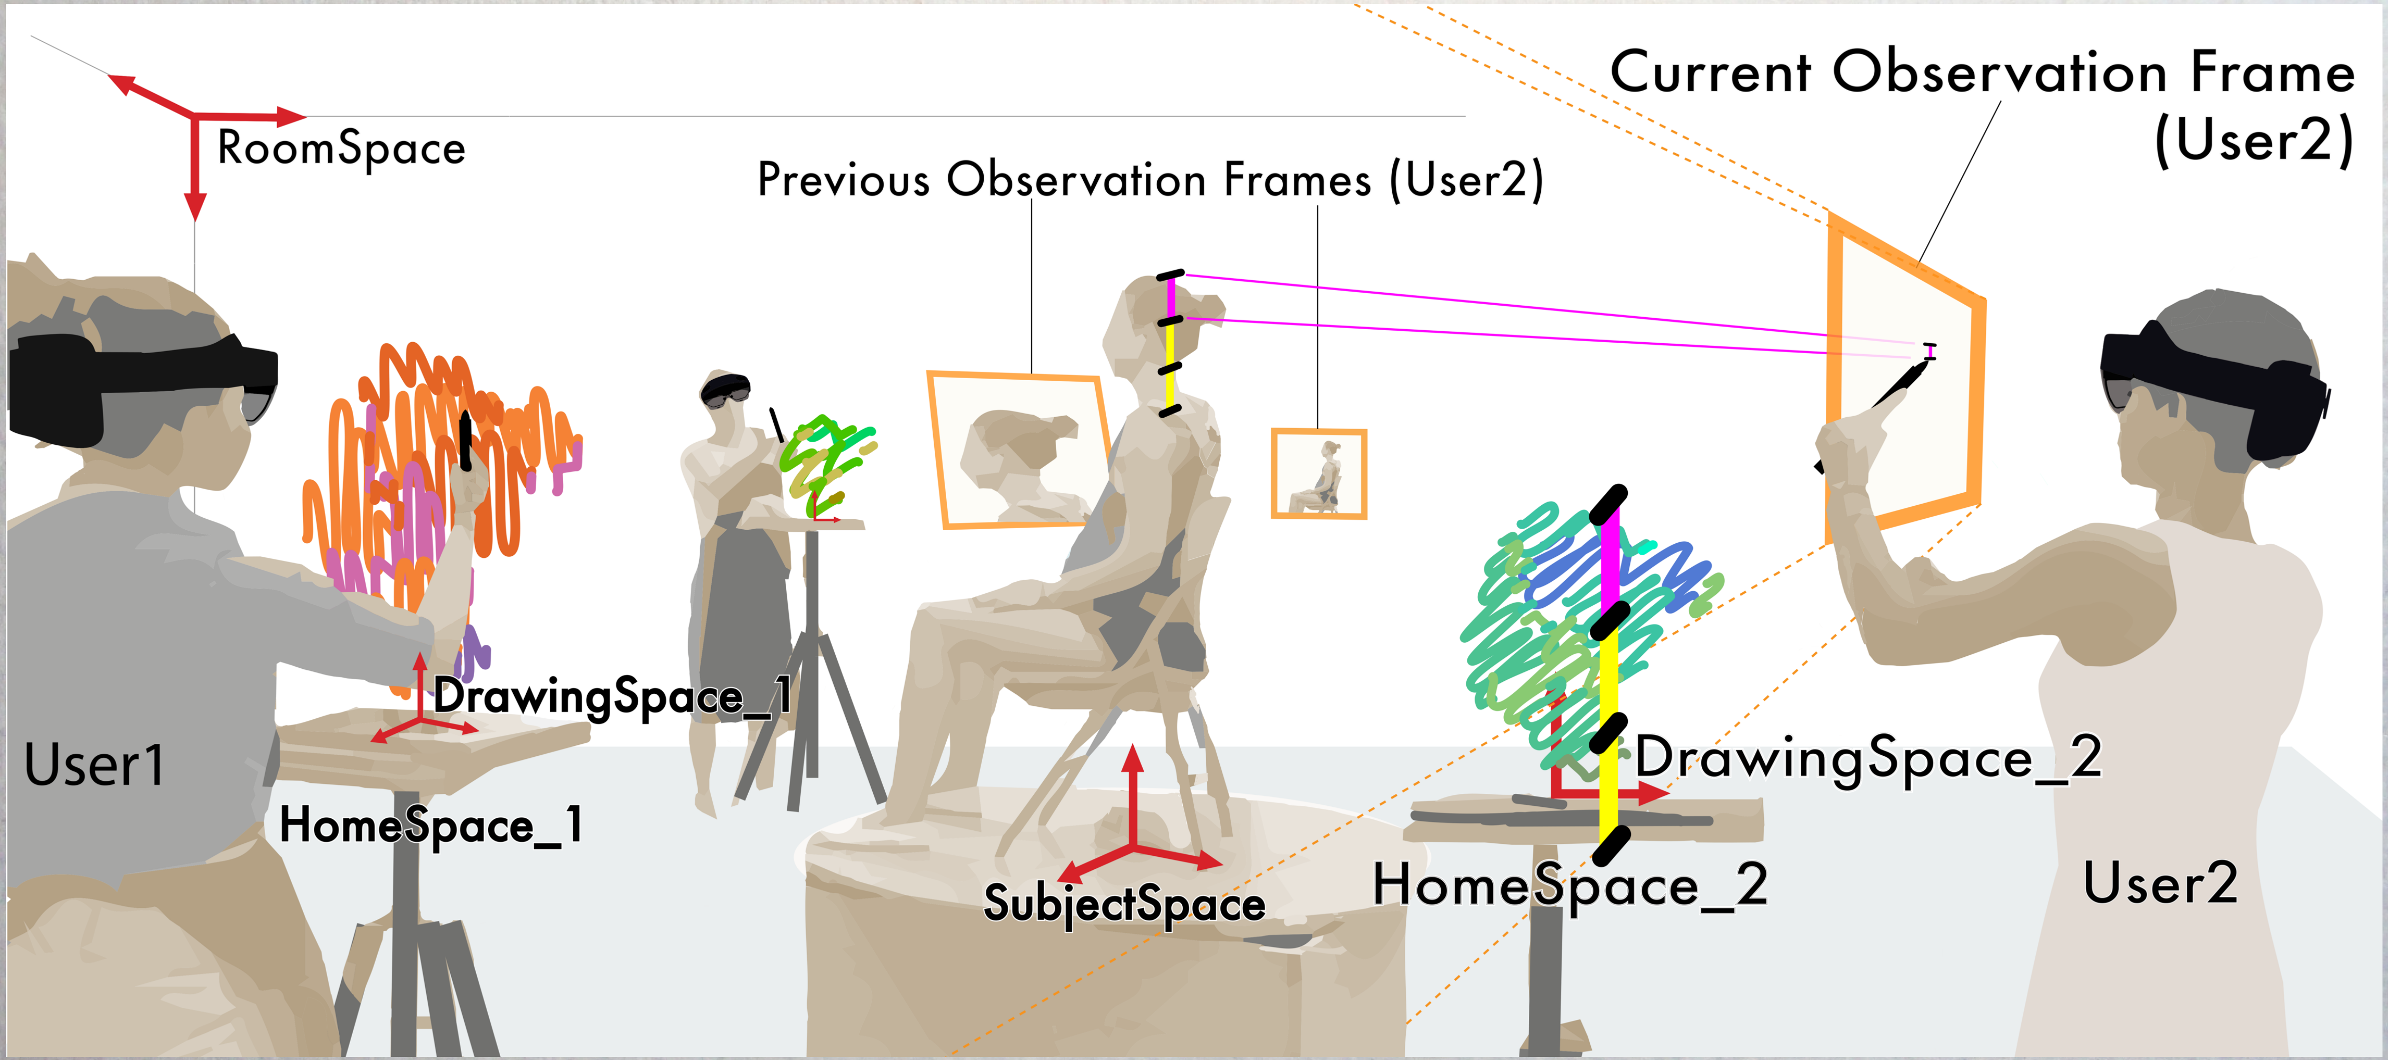

Workspace Guardian:
Investigating Awareness of Personal Workspace Between Co-Located Augmented Reality Users

Bret Jackson, Linda Lor, Brianna C. Heggeseth

B. Jackson, L. Lor and B. C. Heggeseth, "Workspace Guardian: Investigating Awareness of Personal Workspace Between Co-Located Augmented Reality Users," in *IEEE Transactions on Visualization and Computer Graphics*, vol. 30, no. 5, pp. 2724-2733, May 2024, doi: 10.1109/TVCG.2024.3372073.

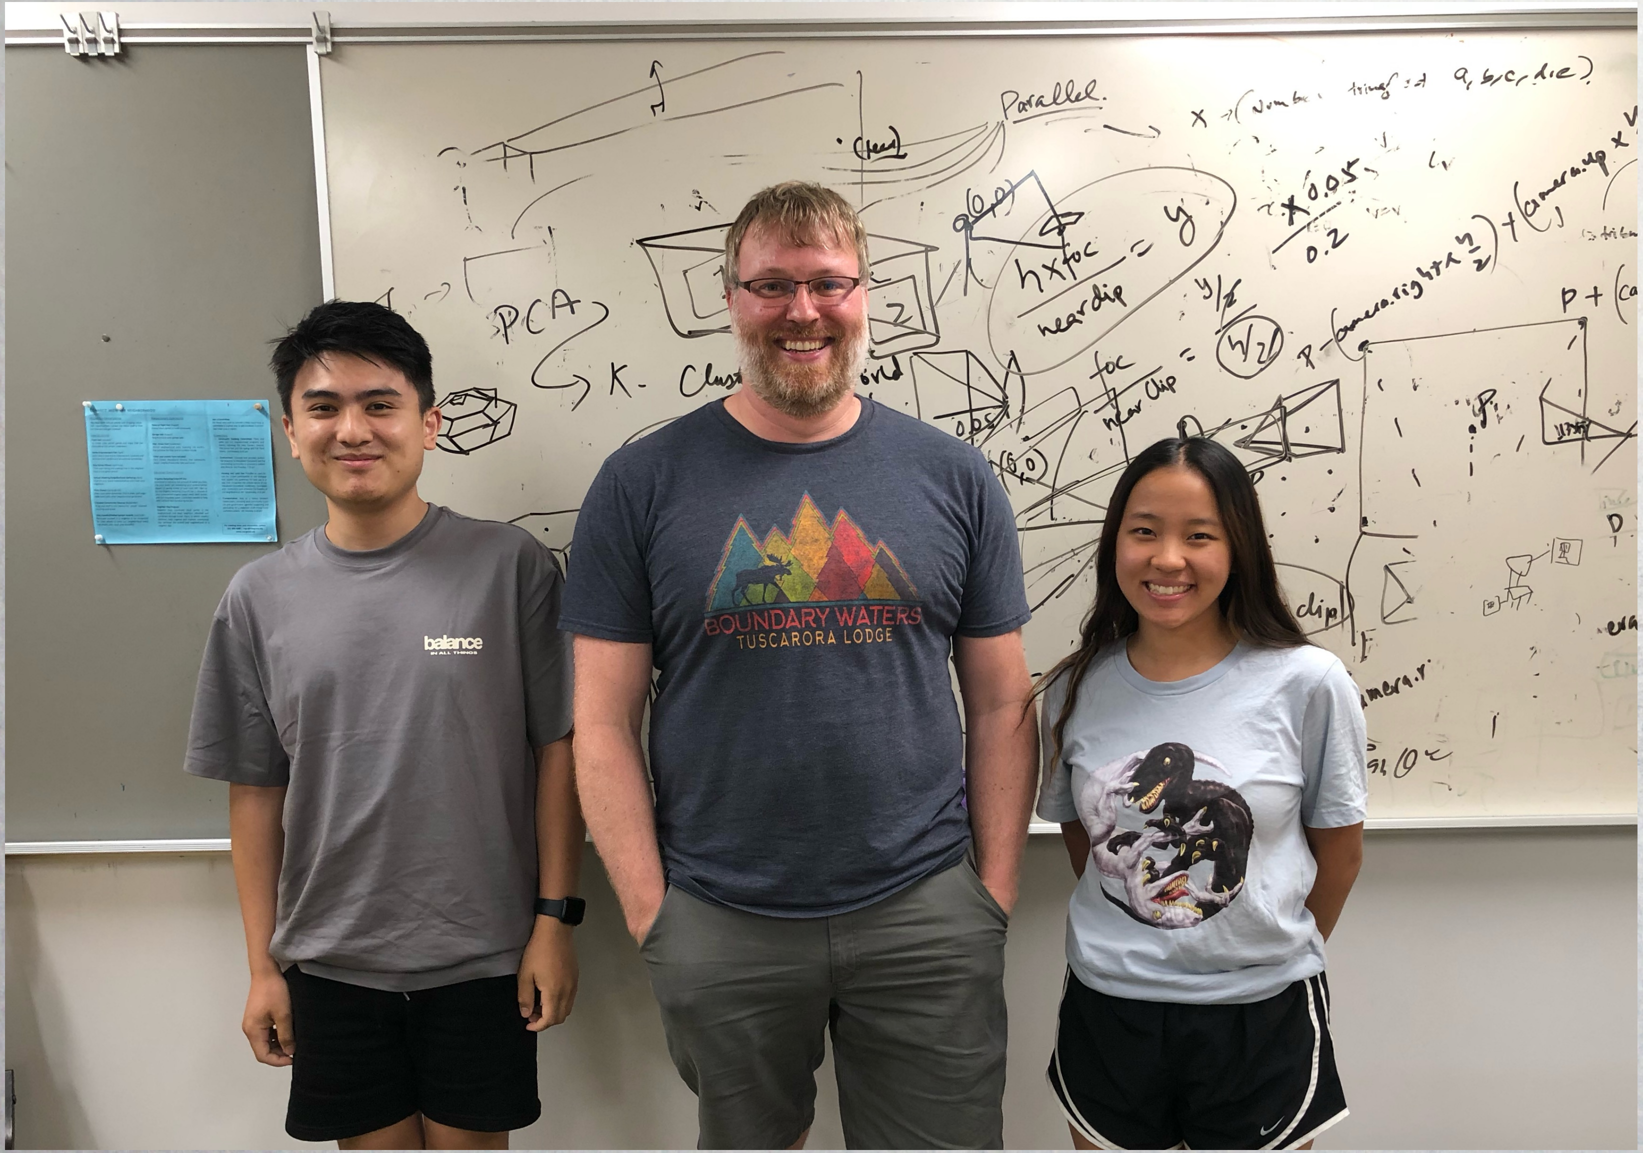
Parallel.
PCA
balance
BOUNDARY WATERS
TUSCARORA LODGE


MAPS
make a U-turn

Drawing from Life

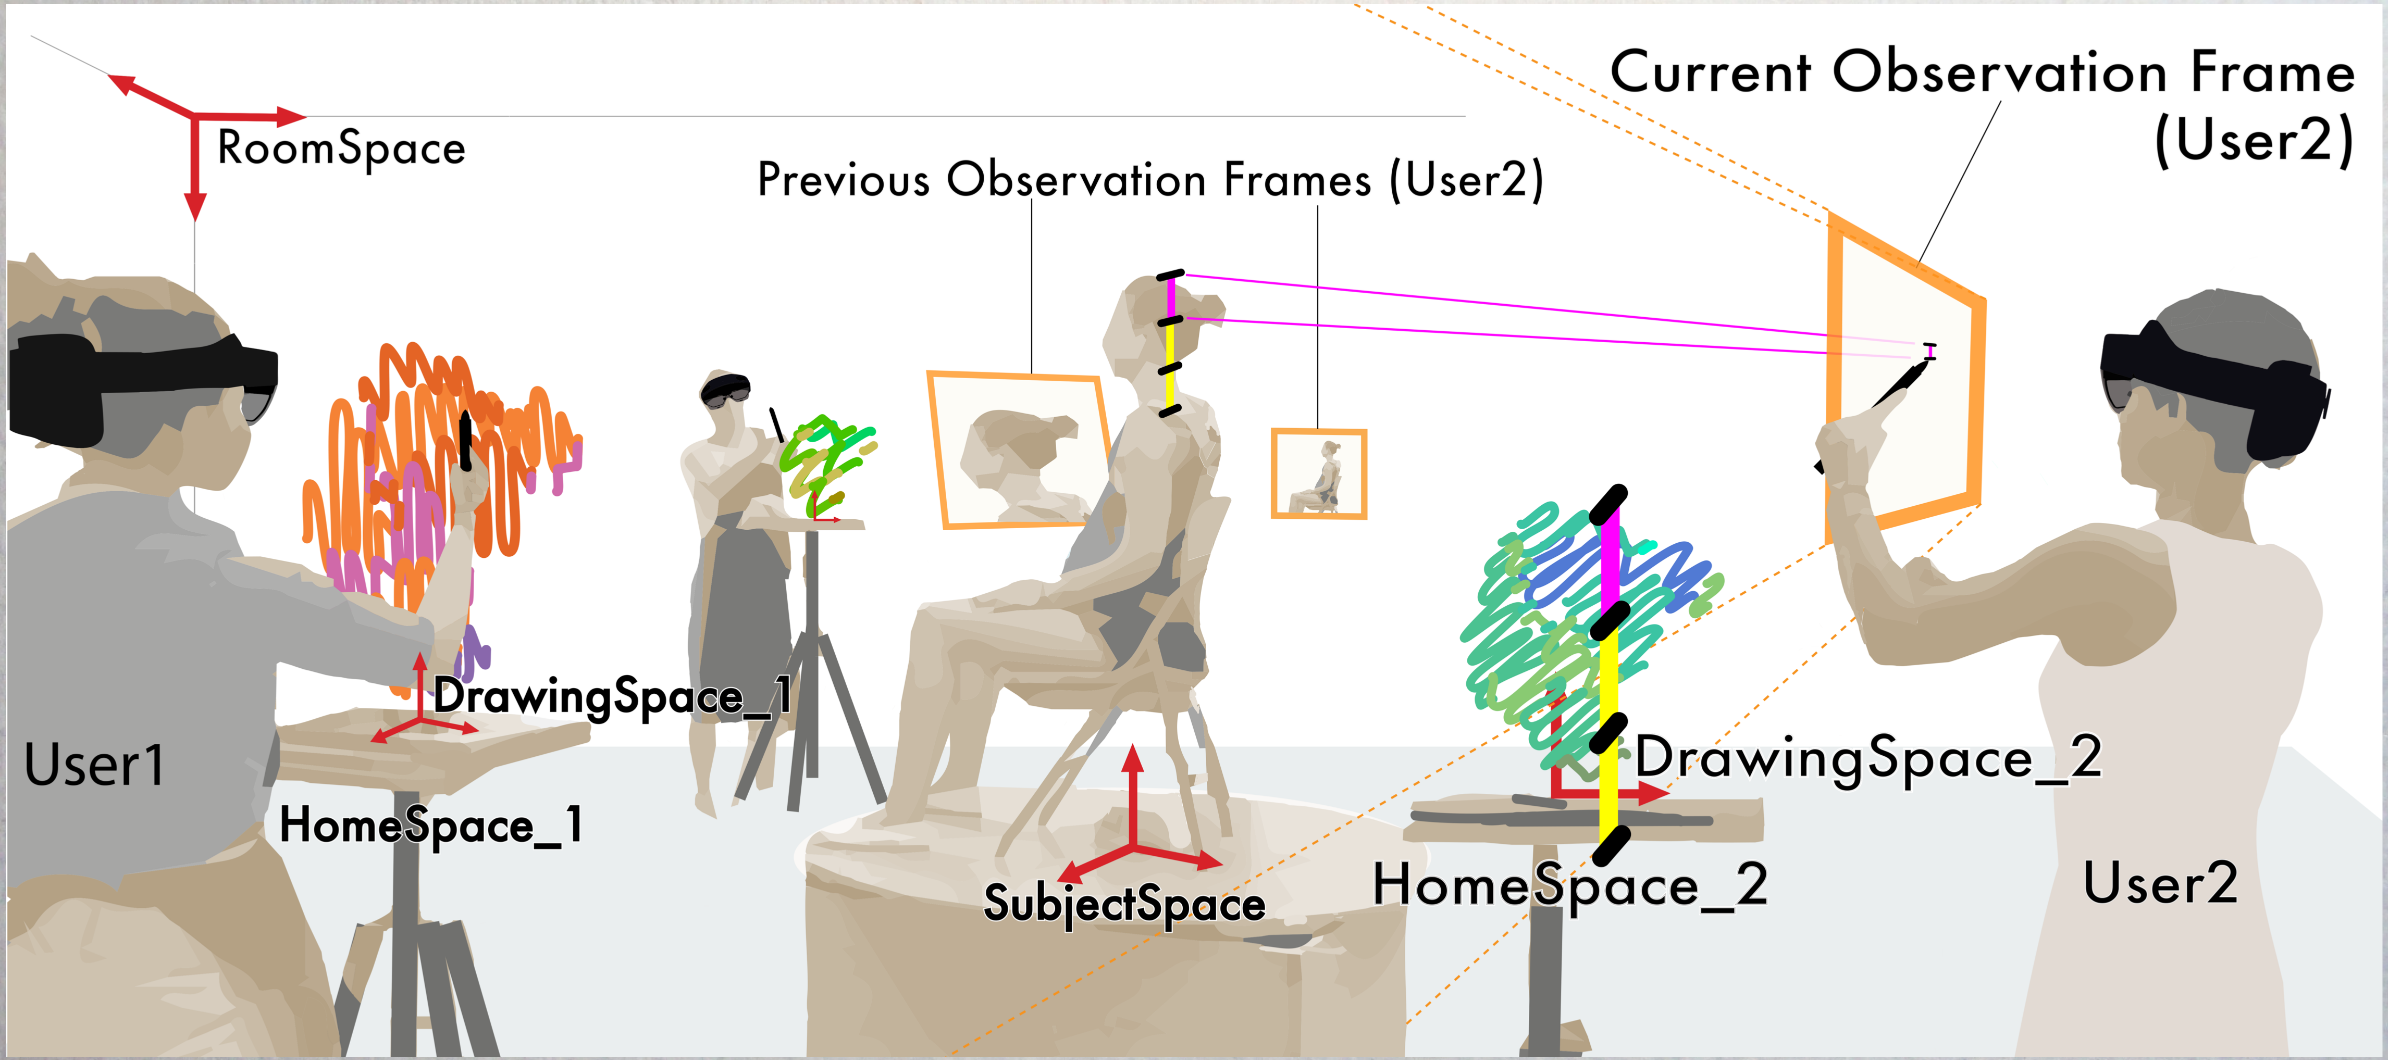

AR Workspace Guardians

One of the largest use-cases for AR in the future is for physically co-located users who are performing independent tasks rather than collaborative ones

Personal workspaces – The area that is directly used for viewing and interacting with virtual content

Workspace Guardians - Visual techniques for communicating personal workspace boundaries between AR users

Research Questions

RO1: Automatic vs. Self-Defined Boundaries

How the methods of defining boundaries and indicating them to other users impacts workspace encroachment

RO2: Privacy Concerns

How privacy considerations influence user preferences of workspace guardian techniques

Full Content View Guardian

Other co-located users can see all the content and interactions, but not directly interact with them

Content Outlines Guardian

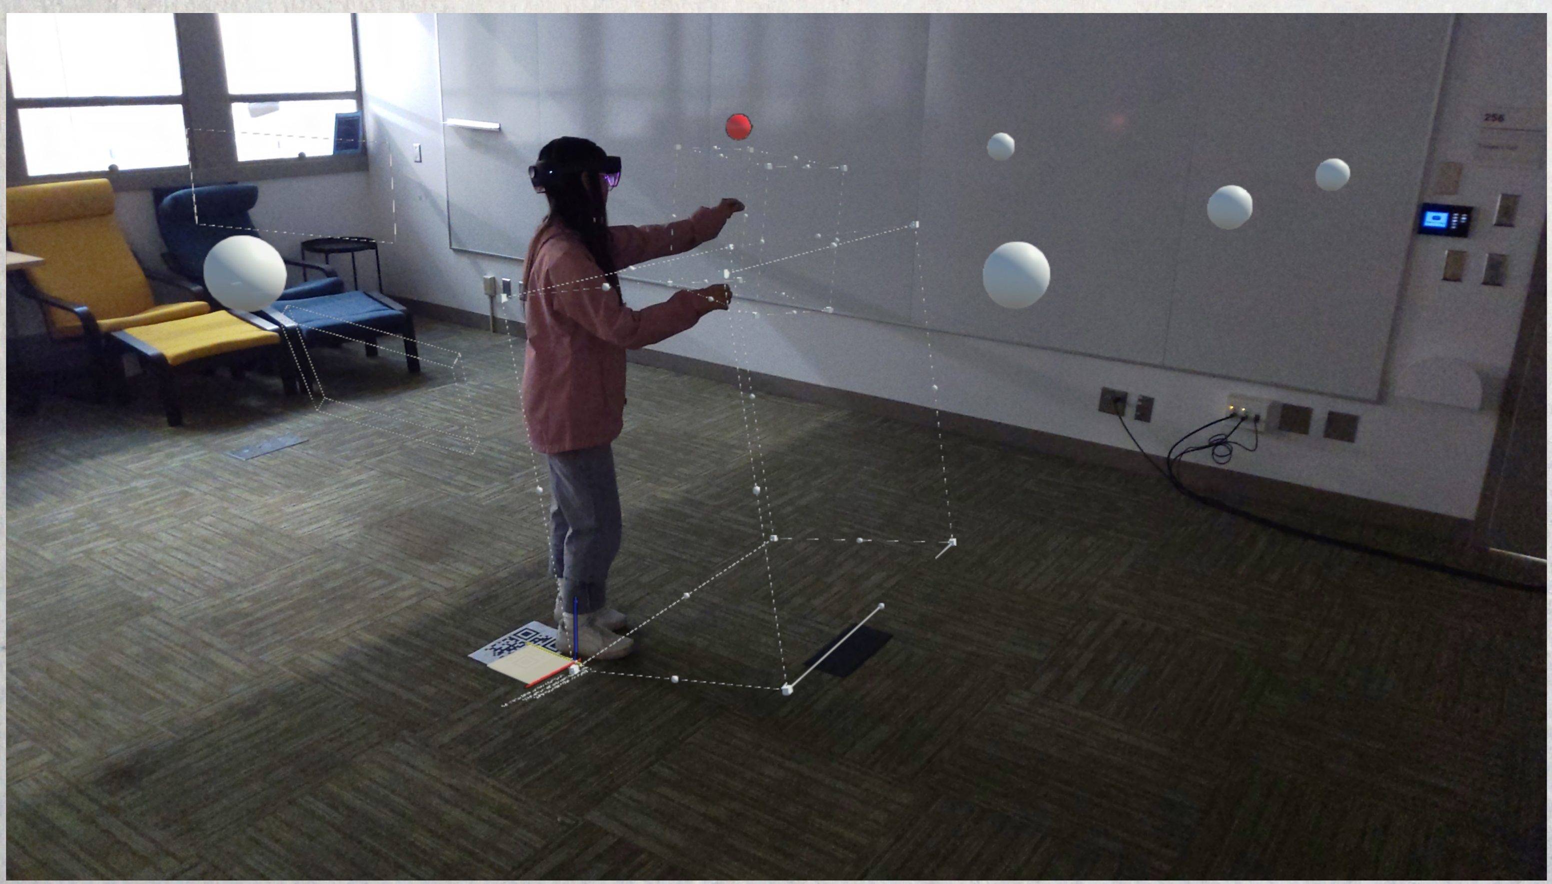

Bounding box outlines of virtual holograms are communicated to other users

Self-Defined Boundary Guardian

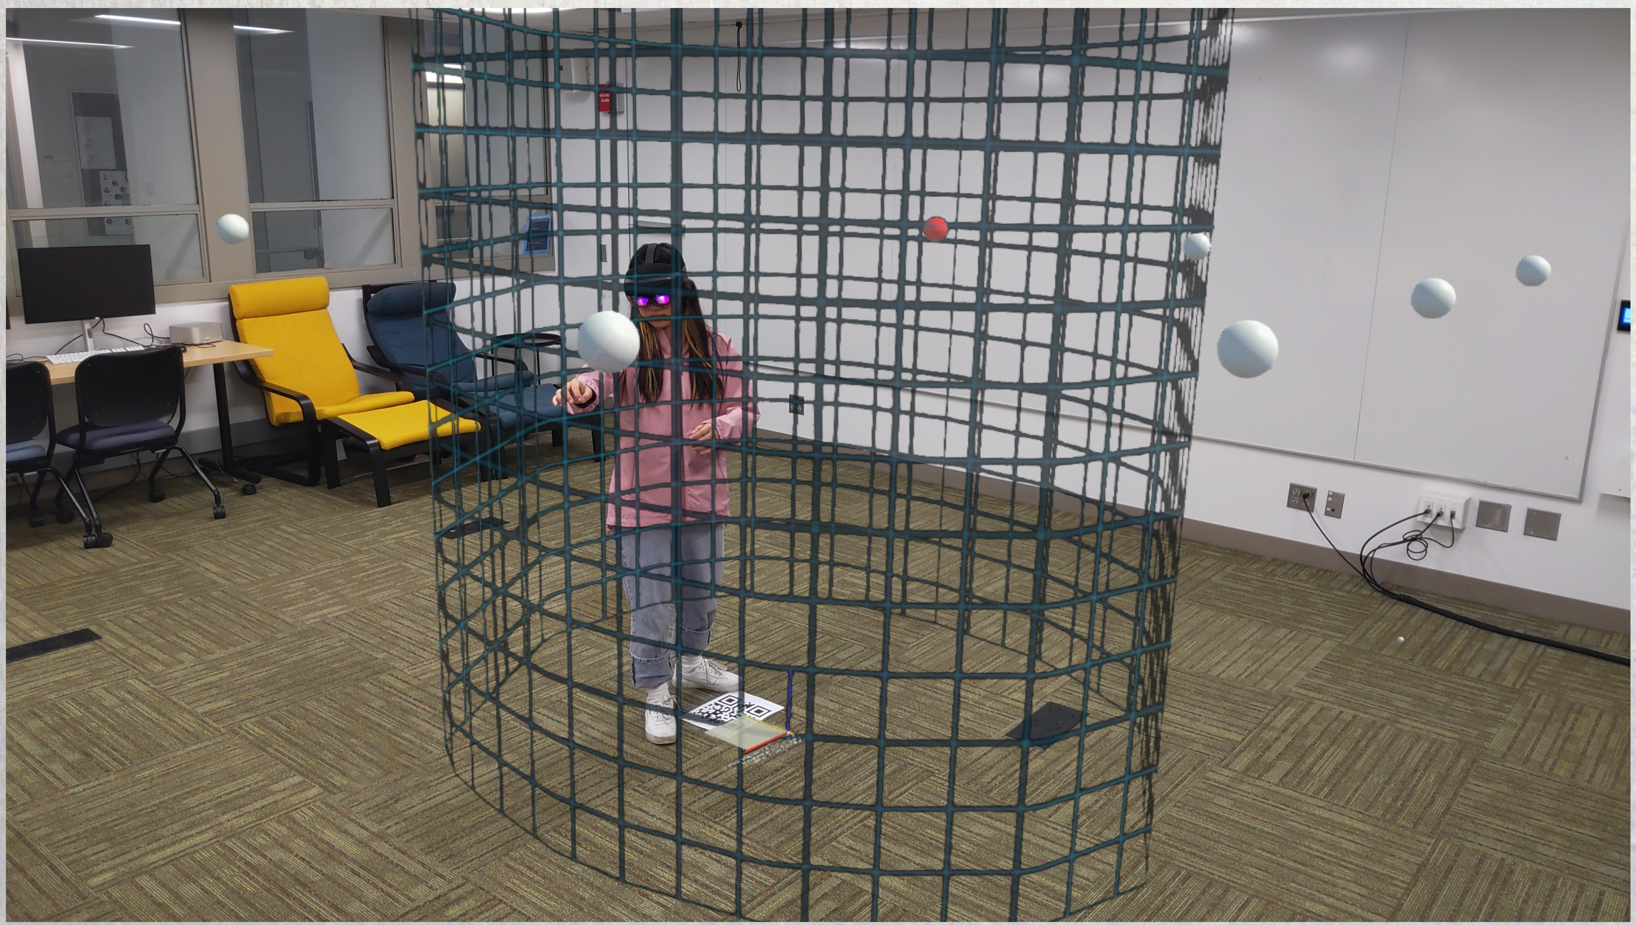

Users self-define their workspace boundary by drawing it on the floor.
The boundary is represented as a grid, similarly to VR guardian systems.

Evaluation Methodology

Participants:

Recruited 36 participants (18 pairs)

Self-identified as 22 male, 10 female, 4 other

Age: 18–35 (M = 21.9, SD = 4.31)

Frequency of Prior XR Use	Number of Participants
No prior use	13
1–5 times	15
5–20 times	2
More than 20 times	6

31 reported prior video game experience (11 monthly, 8 weekly, and 12 daily use)

Task:

Colorer: paint a 3D model

Walker: search and find the red sphere

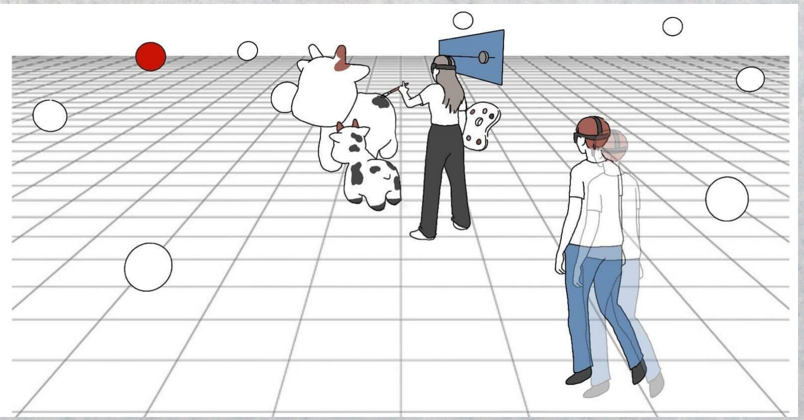

Apparatus:

Two Microsoft Hololens 2 AR HMDs

To facilitate networking, input/output streamed to the HMDs using the Microsoft Holographic Remoting application from two desktop computers

Implemented in Unity3D with MRTK2

4.8768m × 4.8768m square area free of any obstacles

Results: Collisions

Head:

Hands:

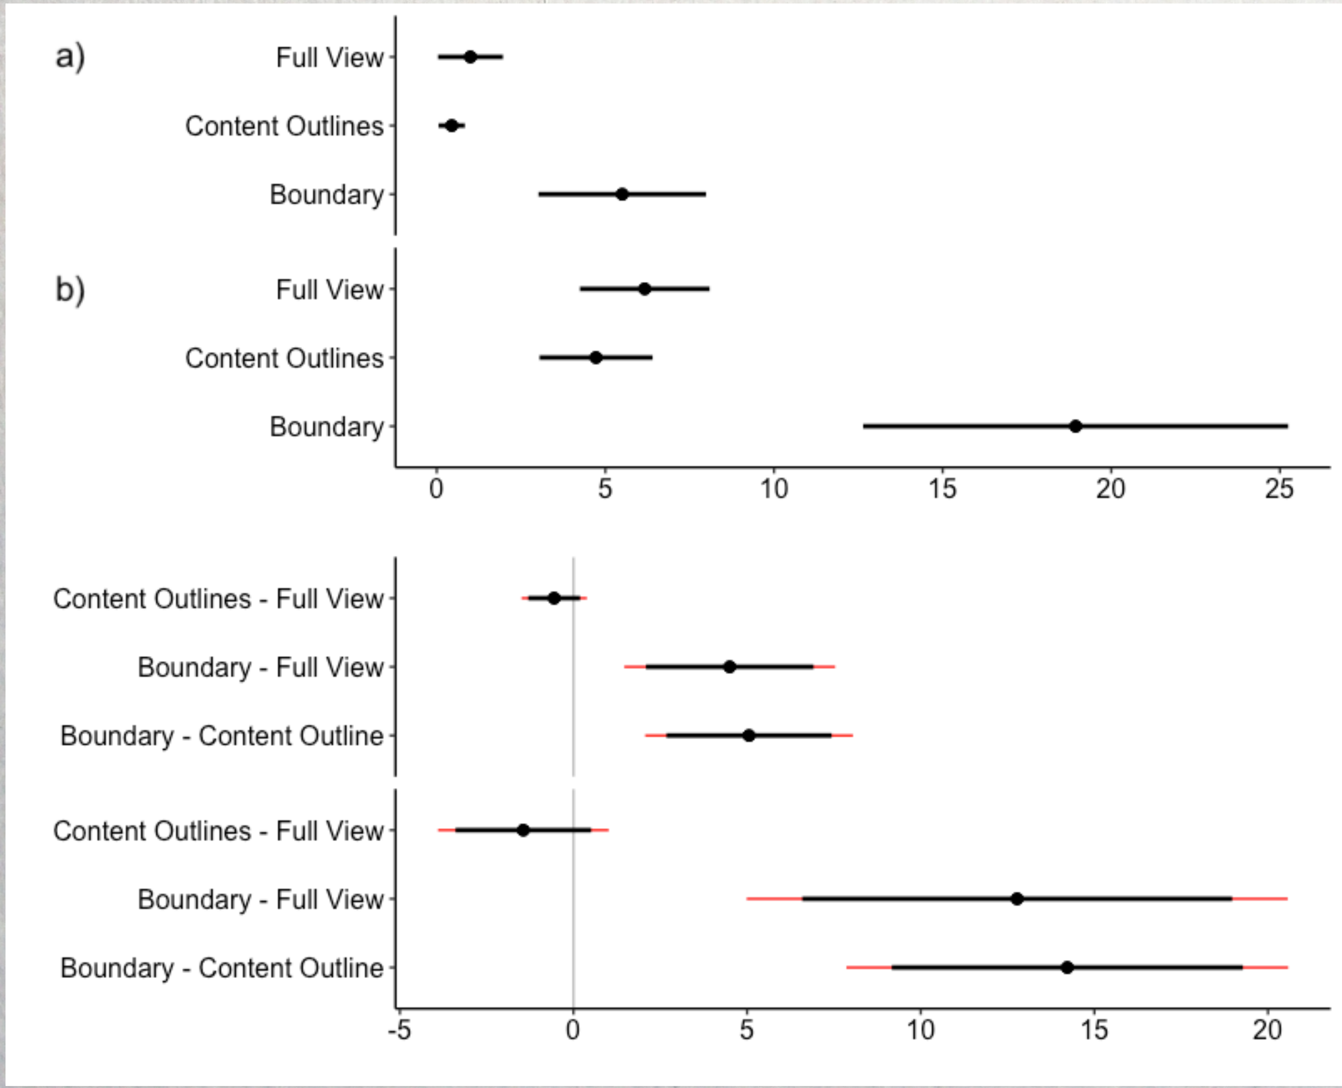

Head:

Hands:

Mean Collisions and 95% CIs

Red bars: CIs for Bonferroni-corrected within subject comparison

Results: Spatial and Temporal Invasion Rate

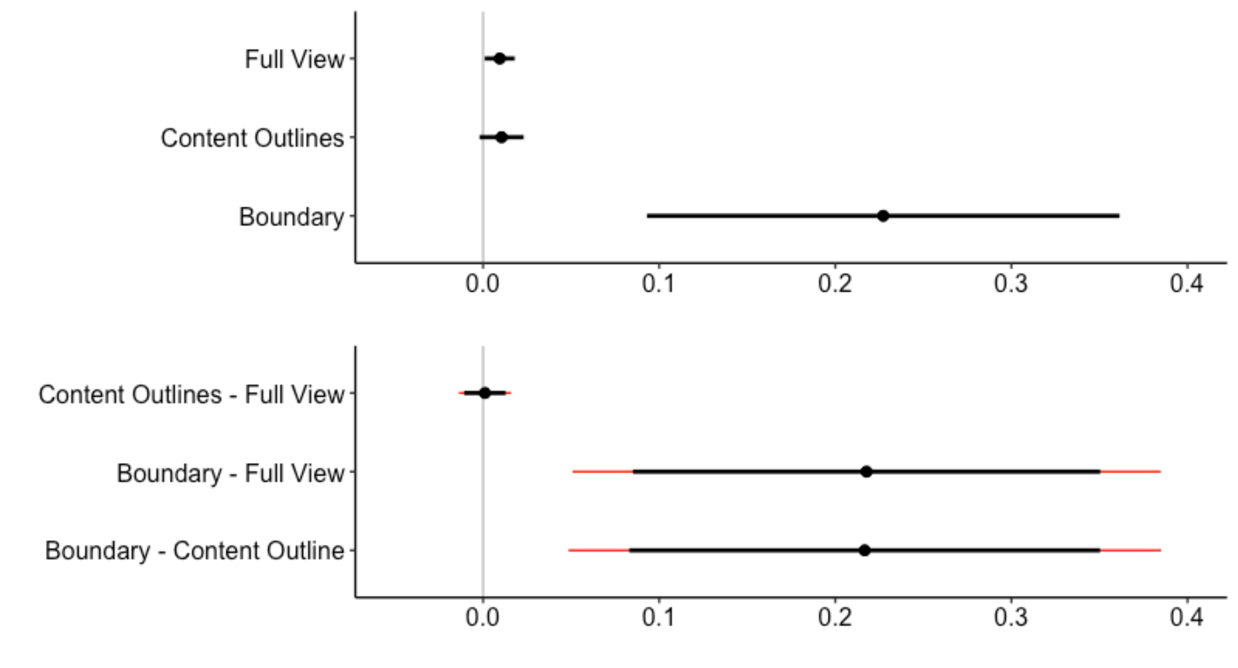

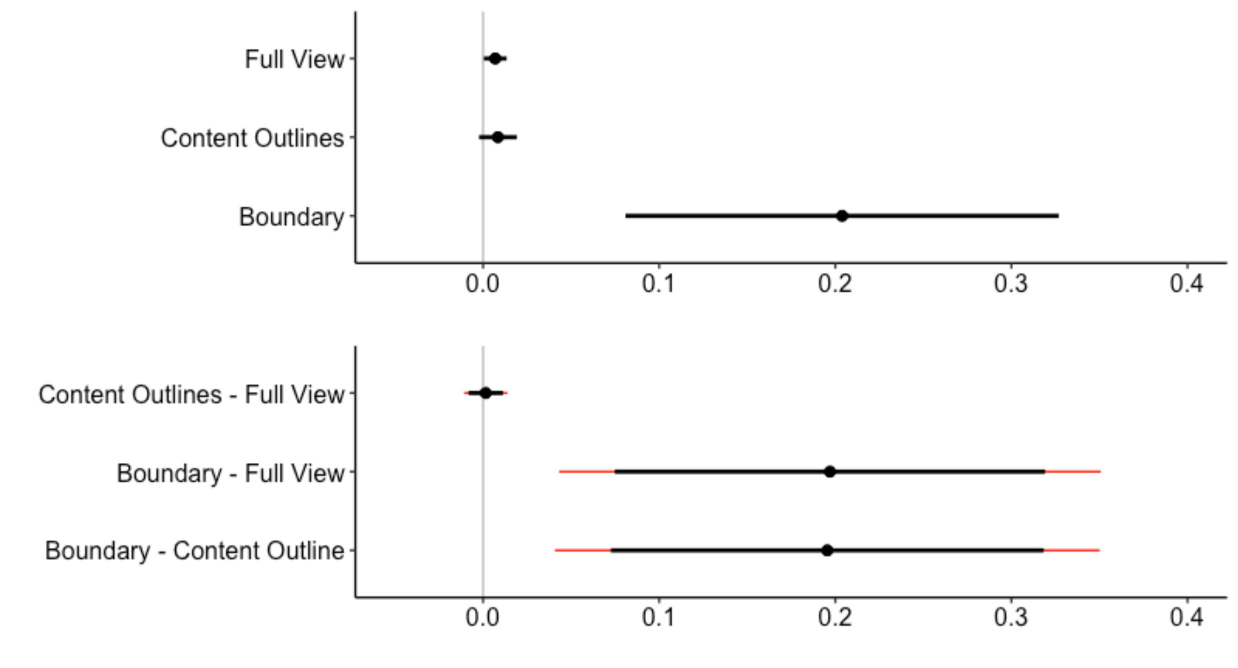

Spatial Invasion Rate - the distance Walkers moved while colliding with the Colorer's virtual content divided by the total distance they traveled in each trial.

Temporal Invasion Rate - the amount of time spent colliding or within the boundary divided by the total task completion time

Significant Qualitative Results between Task Groups

The colorers generally felt more neutral or disagreement that the walker could understand their workspace bounds

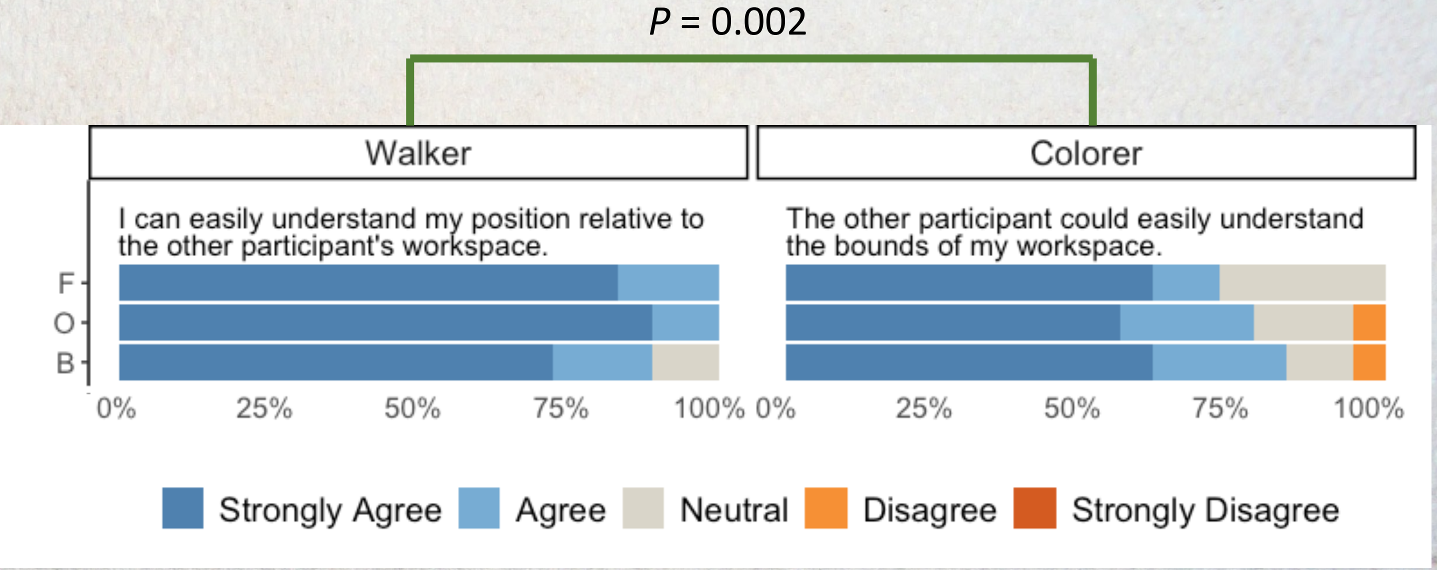

The colorers generally felt more disagreement that their interactions were directed by the other's projections

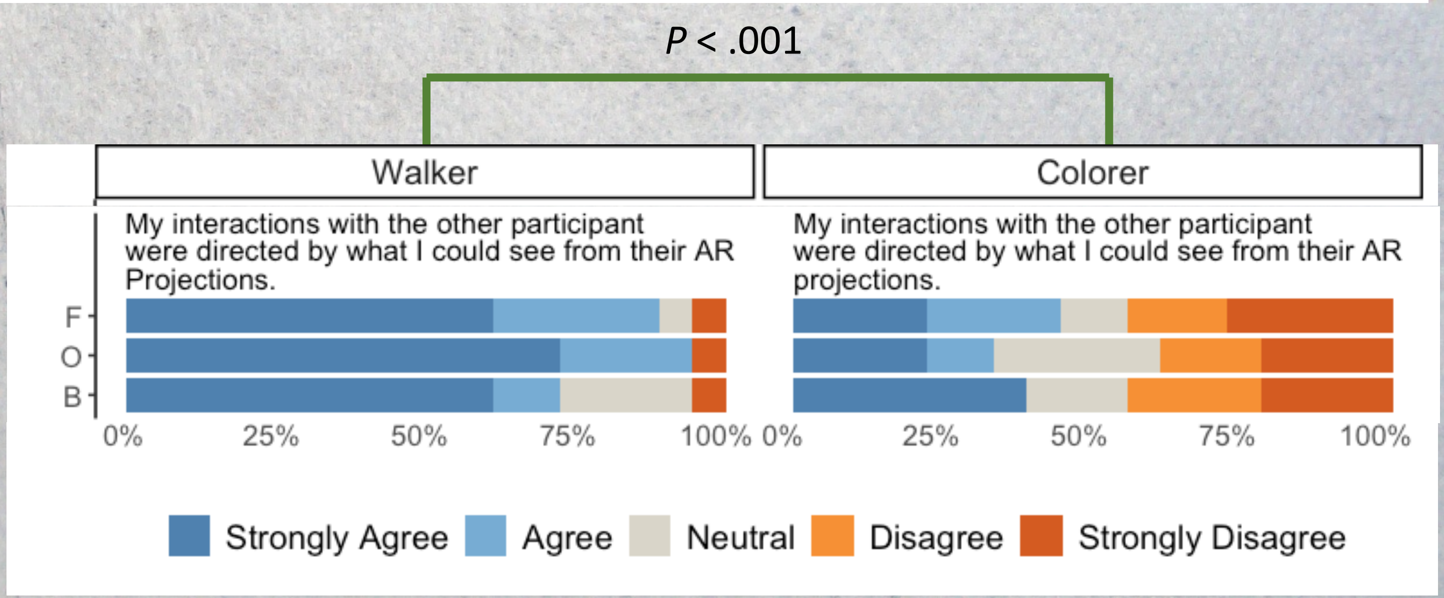

Significant Qualitative Results between Workspace Guardians

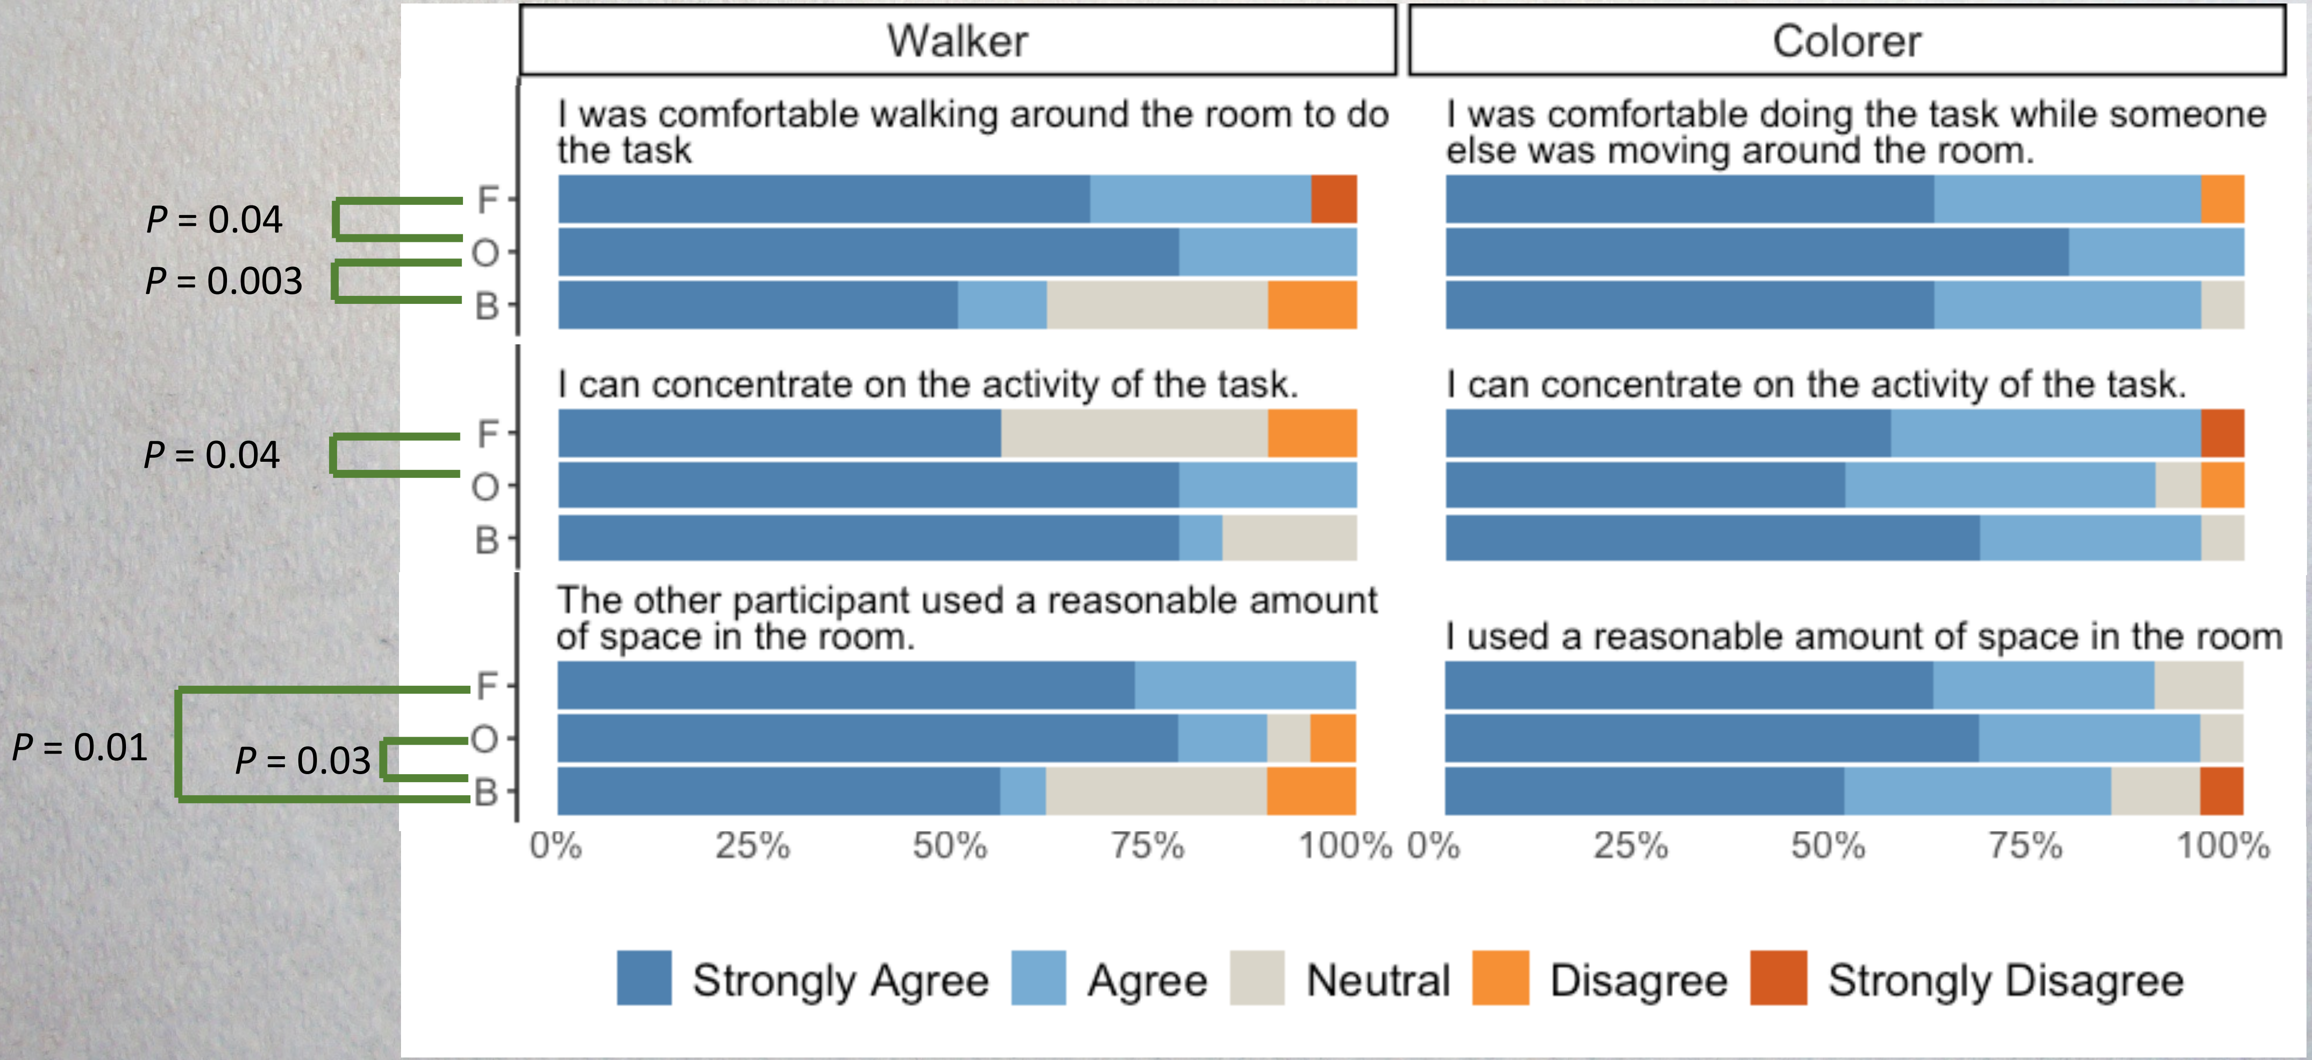

Qualitative Rankings

For the walker:

A clear preference for the Full View followed by the Content Outlines ($\chi 2 = 22$, $P < 0.001$).

For the colorer:

No strong preferences for most preferred guardian

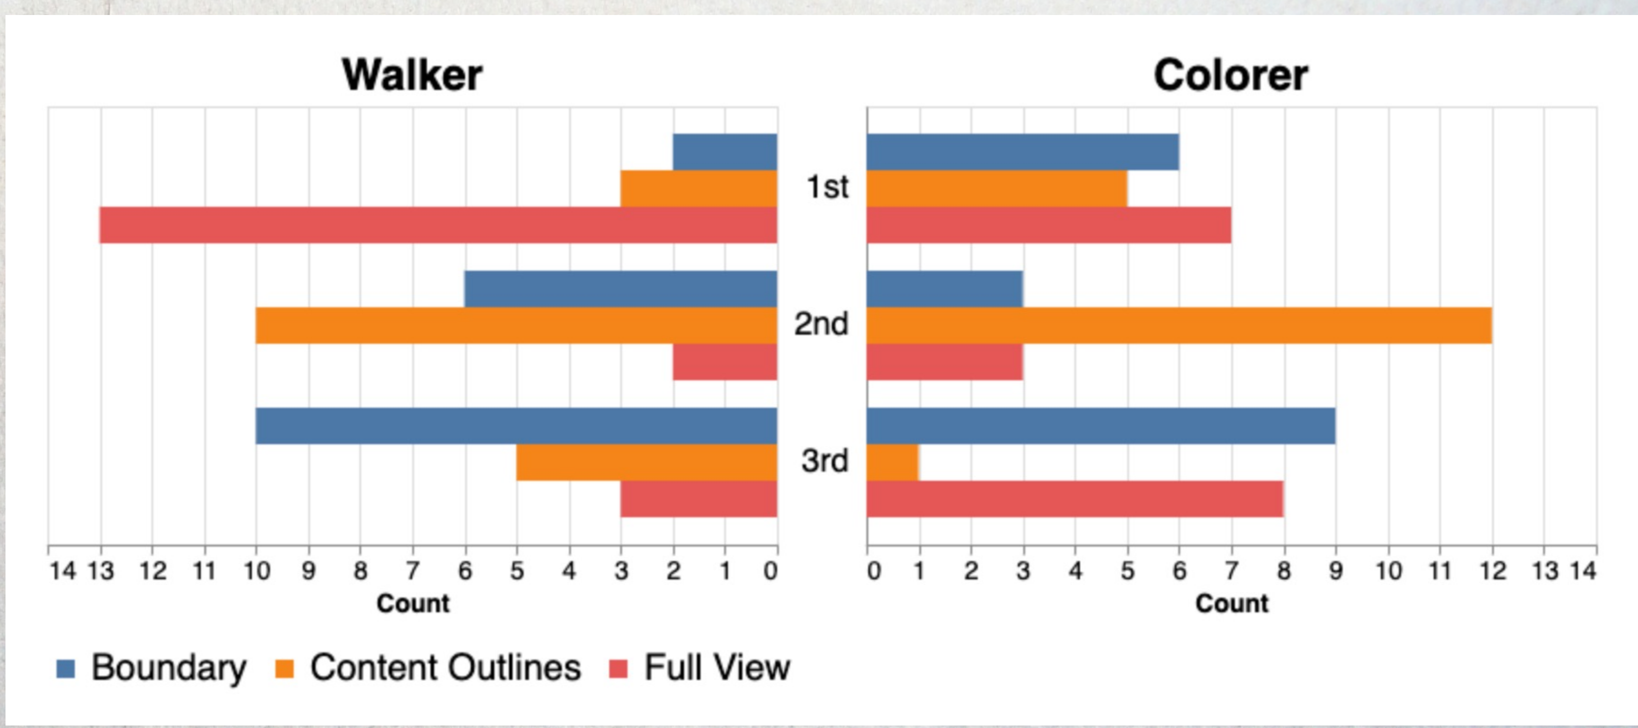

Walker

Aspect	1st	2nd	3rd	4th	5th	Avg Rank
Prevention of physical collisions	11	2	1	2	1	1.72
Being able to view what others are doing in AR	4	4	2	5	3	2.94
Self-defined boundaries	2	4	7	3	3	3.22
Privacy of personal content from neighboring users	1	5	4	3	5	3.33
Prevention of collisions with AR projections	0	3	4	5	6	3.78

Colorer

Aspect	1st	2nd	3rd	4th	5th	Avg Rank
Prevention of physical collisions	13	2	2	1	0	1.50
Privacy of personal content from neighboring users	2	6	3	5	2	2.94
Prevention of collisions with AR projections	1	4	6	7	0	3.06
Self-defined boundaries	2	4	4	5	3	3.17
Being able to view what others are doing in AR	0	2	3	0	13	4.33

The importance rankings for different design aspects of workspace guardians are inverted depending on task

Design Implications

All the workspace guardians influenced behavior

When made aware of workspace boundaries people will try to avoid encroaching on them

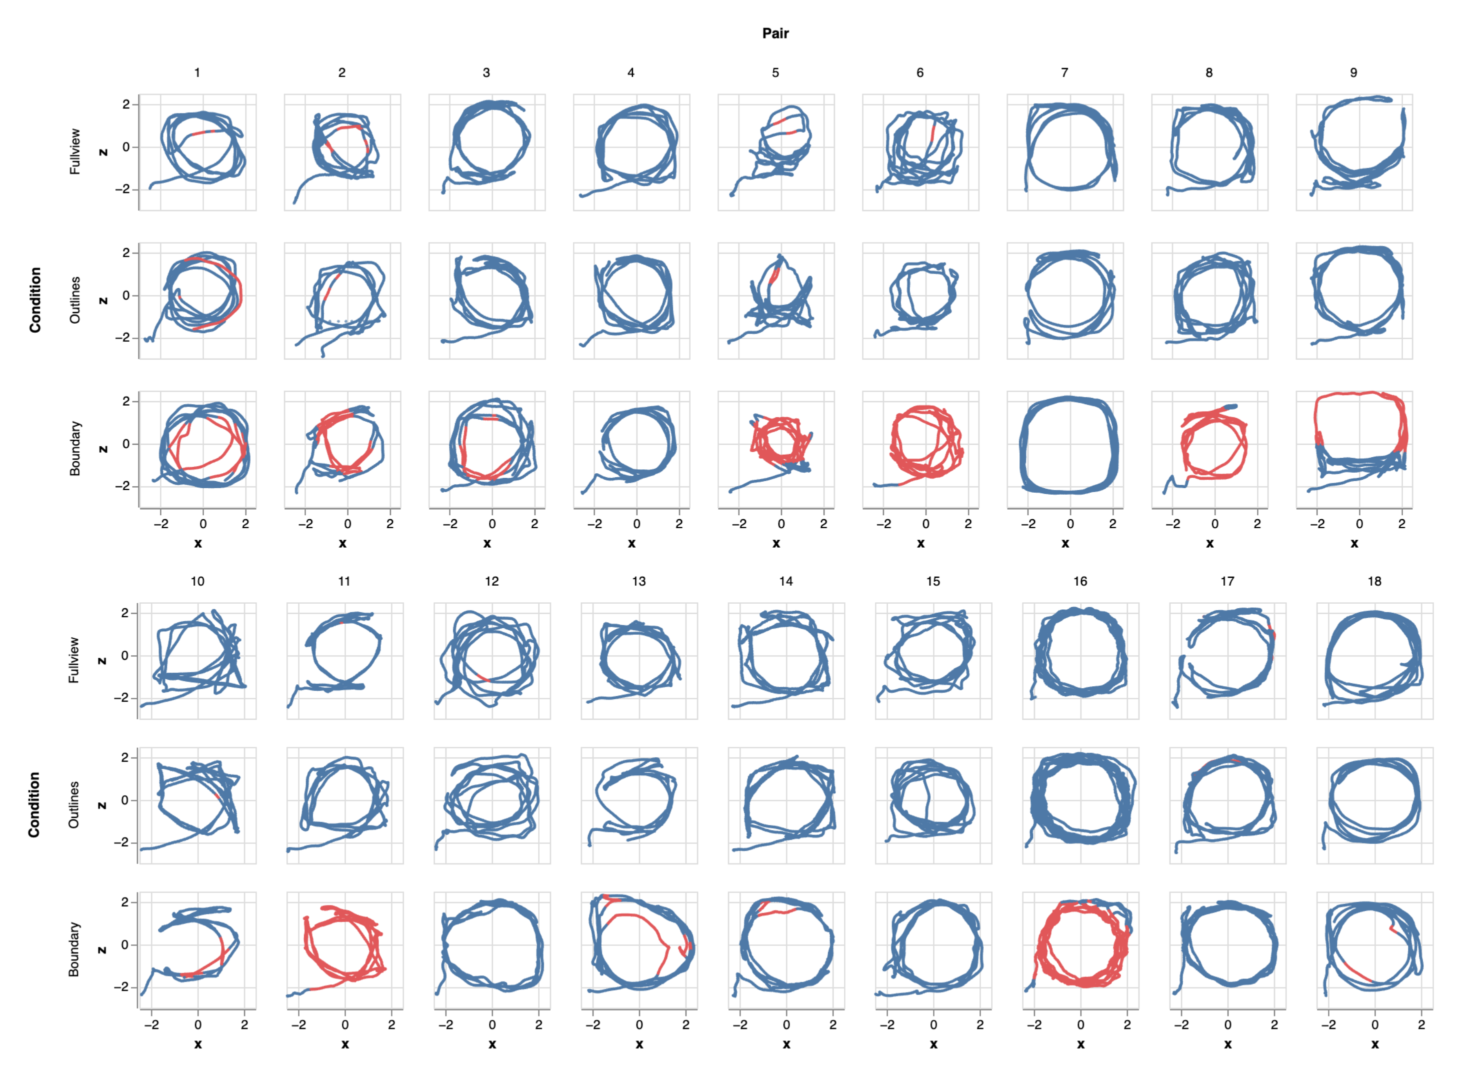

Design Implications

Favor automatic-creation vs. self-defined workspace guardians

Automatically defined guardians had fewer collisions and were generally preferred

Colorers often self-defining workspace boundaries that were larger than needed

> *“The boundary was significantly larger and harder to avoid than simply seeing the full objects which were much smaller than the boundary area and I could walk around easier".*

The visual representations are important and need future work

> *“Boundary view was least preferred because I became more aware of the AR boundary and real life boundary and I didn’t want to disturb the other person’s task."*

> *“I preferred the content outlines option over the boundary option because it felt like there was much more space to move around in even though the boundary option never felt like it took up too much space".*

Design Implications

Tension between privacy and comfort

Participants’ opinions on privacy differed depending on their task

Walkers were curious and wanted to know not just whether they were encroaching on the other’s workspace but whether that would be an issue

> *“With the full view I felt like I understood what the other participant was doing and knew if/why I would be a bother in his space".*

The content outlines guardian balanced privacy and distraction

> *“I found the content outlines to be the happy medium between enough information but not distracting".*

> *“I was able to concentrate more on my task if I couldn’t see what my partner was doing."*

Cross-Reality 3D Sketching

Recenter user-interface development on the task of *observing* and boosting perceptual skills

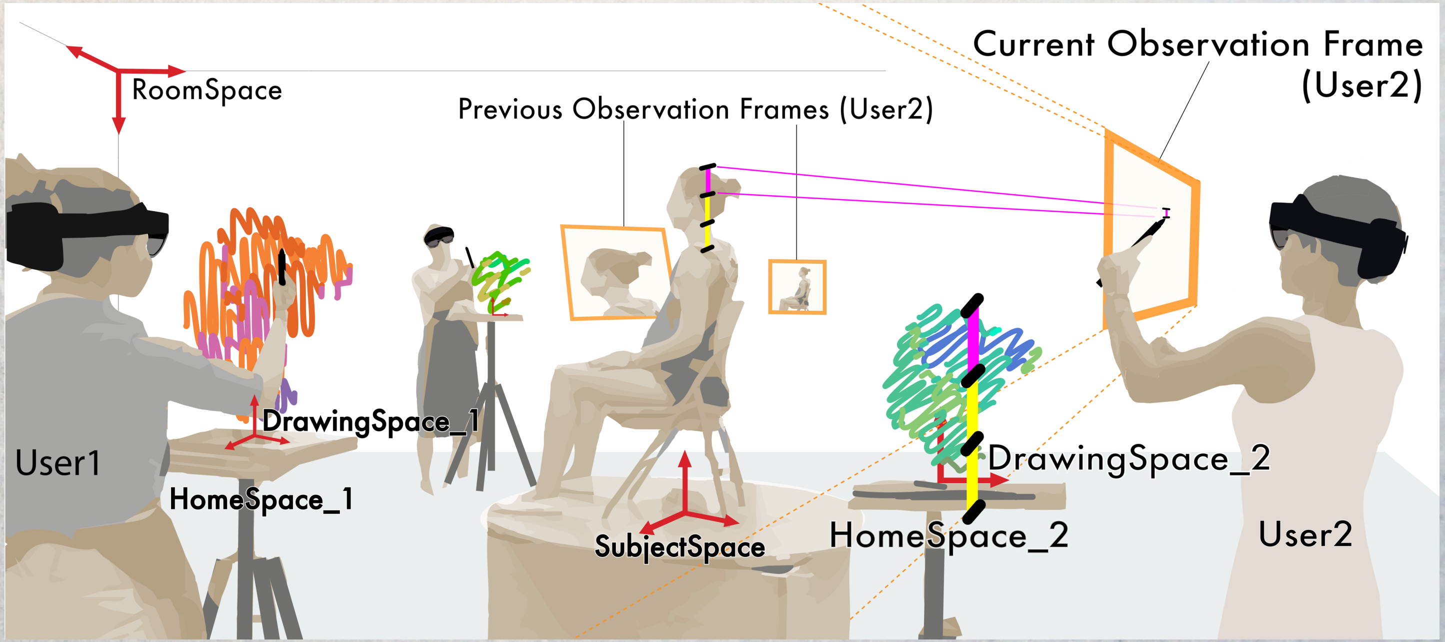

Sighting Points

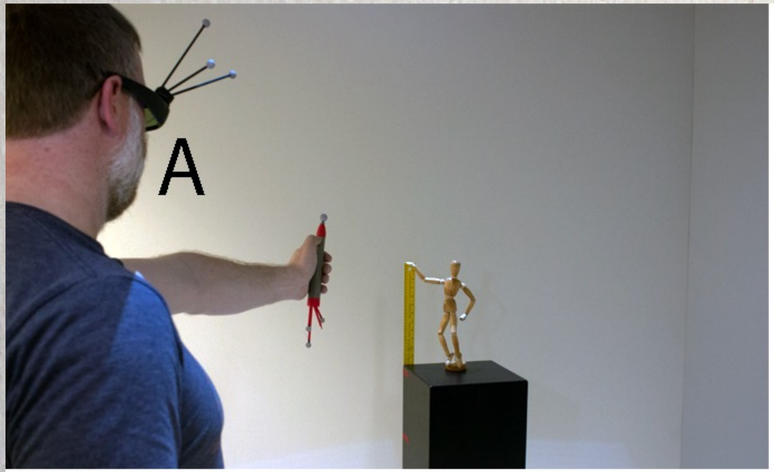

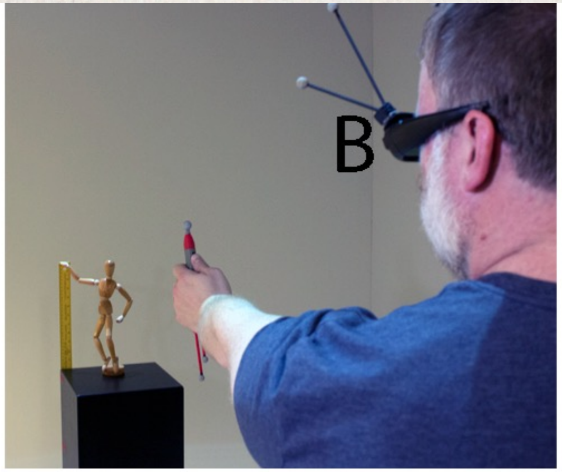

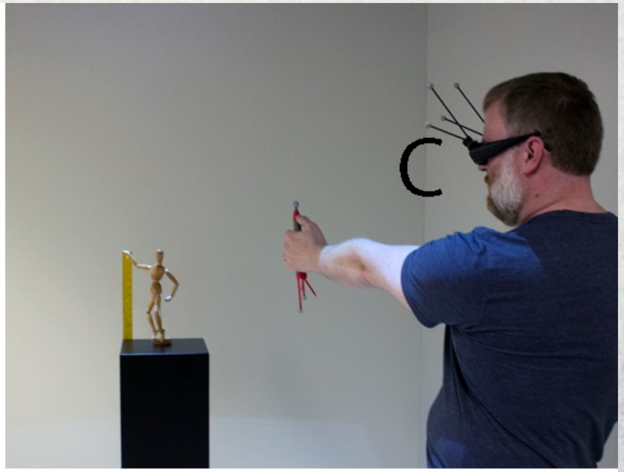

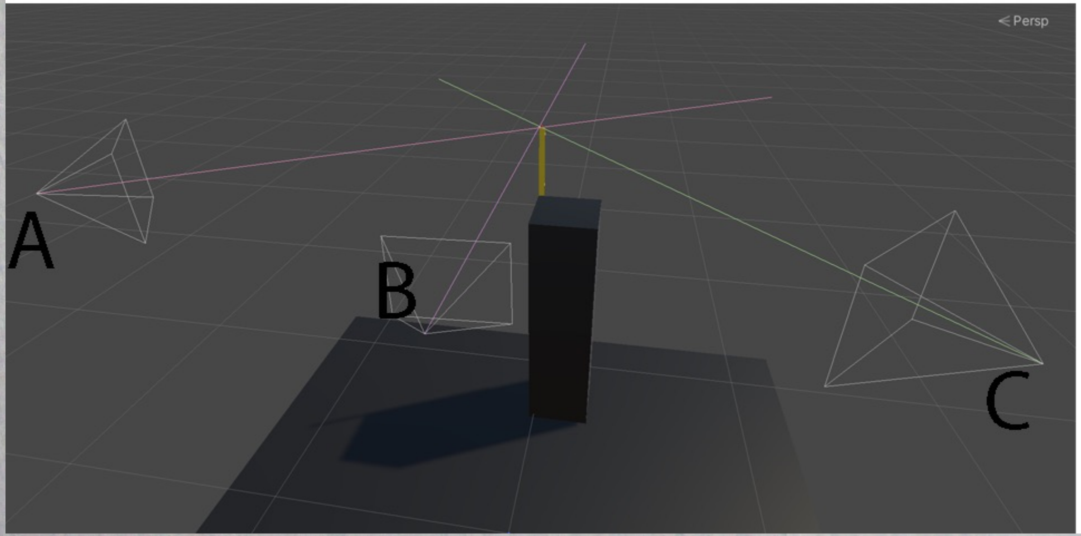

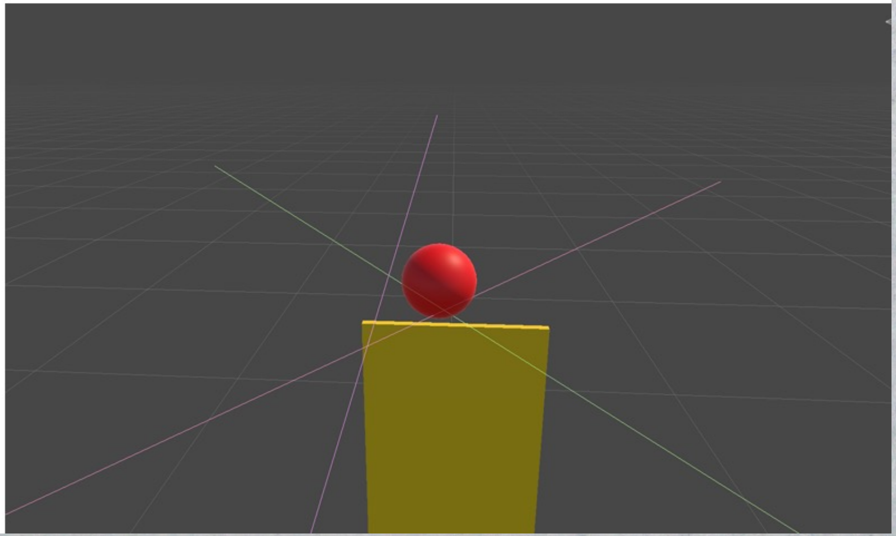

MACALESTER COLLEGE

Thank you!

Bret Jackson (bjackson@macalester.edu)
http://bret-Jackson.com

This work was supported in part through grants from the Clare Booth Luce Foundation, Macalester College, and the National Science Foundation under Grant No.2326999