

Supporting the Design, Selection, and Evaluation of Accessible Interaction Techniques for Virtual Reality

Rachel Franz

A dissertation

submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2024

Reading Committee:

Jacob O. Wobbrock, Chair

Leah Findlater

Anind Dey

Hrvoje Benko

Program Authorized to Offer Degree:

Information School

©Copyright 2024

Rachel Franz

University of Washington

Abstract

Supporting the Design, Selection, and Evaluation of Accessible Interaction Techniques
for Virtual Reality

Rachel Franz

Chair of the Supervisory Committee:

Jacob O. Wobbrock

The Information School

Virtual Reality (VR) is rapidly gaining popularity as a consumer technology, yet accessibility has not been prioritized. VR enables users to interact with immersive virtual worlds, but effective interaction requires the ability to perceive and navigate the environment. Fortunately, over the past thirty years, hundreds of scene-viewing and locomotion techniques that facilitate viewing and navigating in VR environments have been developed that could potentially be accessible to individuals with various physical impairments. However, it remains unclear which techniques are optimal for specific impairments, and the process of designing accessible techniques is largely unexplored. The goal of my dissertation is to provide researchers and designers with guidance on selecting or designing accessible scene-viewing and locomotion techniques to enhance VR accessibility for individuals with upper-body impairments.

To achieve this goal, I developed a taxonomy and design framework to help designers select existing scene-viewing techniques that require minimal head movement as well as create new scene-viewing techniques based on user and virtual environment (VE) characteristics.

Additionally, I created an extensible testbed as a step towards standardizing the evaluation of VR interaction techniques. Using this testbed, I evaluated six locomotion techniques with participants both with and without upper-body impairments and found that Teleport was the most accessible technique. I also found that people with and without impairments performed similarly with Sliding Looking. By analyzing low-level headset and controller data collected during the study, I examined differences in how individuals with and without impairments moved and interacted with the controllers to gain insight into how locomotion technique design impacted performance for both groups. I found that headset metrics differentiated groups across all six techniques and that movement-, button-, and target-related metrics could explain performance differences between groups. Finally, I developed a user model capable of predicting the most suitable locomotion techniques for individual users based on their abilities. The model could predict a user's fastest locomotion technique with 85% accuracy and predicted times were within 16% of actual times it took participants to use the fastest technique.

The thesis this work will demonstrate is:

The accessibility of virtual reality (VR) can be improved by the selection, design, and evaluation of scene-viewing and locomotion techniques using design frameworks, testbed evaluations, and movement analysis and modeling.

TABLE OF CONTENTS

1	List of Figures	xvii
2	List of Tables	xxxiii
3	Acknowledgments.....	xxxviii
4	Introduction.....	1
4.1	Virtual Reality Interaction	3
4.2	Research Approach	5
4.3	Dissertation Structure.....	6
5	Related Work	8
5.1	Ability-Based Design.....	8
5.2	Accessibility in Virtual Reality.....	8
5.2.1	Vision Impairments.....	9
5.2.2	Hearing Impairments	10
5.2.3	Physical Impairments.....	10
5.2.4	Physical Impairments: Identifying Accessibility Problems.....	11
5.2.5	Physical Impairments: Improving Accessibility for People with Ambulatory Impairments	12
5.2.6	Physical Impairments: Improving Accessibility for People with Upper-Body Impairments	13
5.2.7	Physical Impairments: Improving Accessibility for People Experiencing Situational Impairments	13
5.3	Virtual Reality Design Frameworks and Taxonomies.....	14

5.3.1	Using Virtual Reality Design Frameworks and Taxonomies to Improve Accessibility	15
5.4	An Overview of Scene-Viewing Techniques	16
5.4.1	Camera-Control Techniques	16
5.4.2	Scene Modification Techniques.....	17
5.4.3	Social Cue Augmentation	17
5.4.4	Multiple View Techniques.....	17
5.4.5	Field-of-view (FOV) Extension Techniques	18
5.4.6	Multiscale Techniques	18
5.4.7	Projection Distortion Techniques	18
5.4.8	Cue-Based POI Techniques	19
5.5	An Overview of Locomotion Techniques.....	20
5.5.1	Walking-in-place Techniques	21
5.5.2	Steering-Based Techniques.....	21
5.5.3	Selection-Based Techniques	22
5.5.4	Manipulation-Based Techniques	22
5.6	Evaluating Virtual Reality Interaction Techniques.....	23
5.6.1	User Experience and Accessibility of Locomotion Techniques	24
5.6.2	A Lack of a Standardized Approach to Virtual Reality Interaction Technique Evaluation	24
5.6.3	Virtual Reality Testbeds to Support Interaction Technique Evaluation	25
5.7	Measuring the Movement Characteristics of People with Physical Impairments	26
5.8	Modelling User Behavior to Personalize Virtual Reality Interaction.....	28

5.8.1	Performance Modelling for Interface Optimization	28
5.8.2	Performance Modelling in Virtual Reality	29
5.8.3	Authentication in Virtual Reality.....	29
5.9	Summary of Related Work	30
6	A Virtual Reality Scene Taxonomy: Identifying and Designing Accessible Scene-Viewing Techniques	31
6.1	Devising the Scene Taxonomy	32
6.1.1	Theory: Affordances	33
6.1.2	Literature Review Approach.....	34
6.2	The Scene Taxonomy	35
6.2.1	Dimension and Property Definitions	36
6.2.2	Dimension 1: Visual Properties	36
6.2.3	Dimension 2: Task Types	39
6.3	Scene Understanding	41
6.3.1	Scene Definitions	41
6.3.2	Scene Descriptions.....	42
6.4	Scene Properties.....	45
6.4.1	Dimension 1: Visual Properties	45
6.4.2	Dimension 2: Tasks.....	47
6.5	Summary	47
6.6	Application of the Scene Taxonomy.....	48
6.6.1	Common Scene Properties and Tasks in Research Virtual Environments	50
6.7	Using the Scene Taxonomy as an Ideation Tool	51

6.8	Using the Taxonomy to Inform the Design of Scene-Viewing Techniques	53
6.8.1	Implementation Details	53
6.8.2	Design Considerations: Overview	54
6.8.3	Design Considerations: Accessibility	54
6.8.4	Design Considerations: Usability	55
6.8.5	Design Considerations: Realism	55
6.8.6	Design Considerations: Spatial Awareness	55
6.8.7	Design Considerations: User Comfort	56
6.8.8	Design Considerations: User Familiarity with Interaction	56
6.8.9	Design Considerations: Task Objective.....	56
6.8.10	Design Considerations: Summary.....	57
6.9	Technique 1: Object of Interest.....	57
6.10	Technique 2: Proxemics Snapping.....	59
6.11	Technique 3: Rearview Mirror.....	62
6.12	User Study.....	63
6.12.1	Participants.....	63
6.12.2	Apparatus	64
6.12.3	Procedure	64
6.12.4	Design and Analysis	66
6.13	Results.....	67
6.13.1	Technique 1: Object of Interest.....	67
6.13.2	Technique 2: Proxemics Snapping.....	69
6.13.3	Technique 3: Rearview Mirror.....	71

6.14	Discussion	72
6.14.1	Tradeoffs to Consider	73
6.14.2	Reflecting on the Scene Taxonomy: Opportunities and Limitations.....	75
6.14.3	Accessibility for Individuals with Motor Impairments.....	77
6.14.4	Limitations	77
6.15	Summary	78
7	Nearmi: A Framework for Designing Point of Interest Techniques for Virtual Reality Users with Limited Mobility.....	79
7.1	Nearmi.....	80
7.1.1	Representation.....	81
7.1.2	Display	81
7.1.3	Selection.....	81
7.1.4	Transition	82
7.2	Benefits of Nearmi as a Design Framework	82
7.2.1	Foregrounds Accessibility	83
7.2.2	Structures the Design Process	83
7.2.3	Facilitates Evaluation.....	83
7.2.4	Affords Customizability.....	83
7.3	Design Considerations for Nearmi Implementations	84
7.3.1	Accessibility.....	84
7.3.2	Usability	84
7.3.3	Realism	85
7.3.4	Spatial Awareness.....	85

7.3.5	User Comfort	85
7.3.6	User Familiarity with Interaction.....	86
7.3.7	Task Objective	86
7.3.8	Summary	86
7.4	Implementation	86
7.4.1	Representation.....	87
7.4.2	Display	87
7.4.3	Selection.....	88
7.4.4	Transition	90
7.4.5	Summary	91
7.5	High-level Implementation Overview.....	92
7.6	User Interviews	93
7.6.1	Participants.....	94
7.6.2	Apparatus	95
7.6.3	Procedure	96
7.6.4	Analysis.....	98
7.7	Results.....	98
7.7.1	Impact of Mobility Limitations on Virtual Reality Use.....	98
7.7.2	Impact of Limited Head and Body Rotation on Virtual Reality Use.....	99
7.7.3	Design Considerations	101
7.7.4	Design Considerations: Representation	101
7.7.5	Design Considerations: Display.....	102
7.7.6	Design Considerations: Selection	102

7.7.7	Design Considerations: Transition.....	104
7.7.8	Design Considerations: Display × Selection Combination.....	105
7.7.9	Design Considerations: Accessibility Via Customizability.....	105
7.7.10	Applicability of Nearmi to Real Scenarios.....	106
7.7.11	Summary	108
7.8	Discussion	108
7.8.1	Design Considerations Identified by Participants.....	108
7.8.2	Need for Customizability in the Design of Virtual Reality Techniques	110
7.8.3	Customization and Personalization vs. Standardization in Virtual Reality.....	110
7.8.4	How Do We Evaluate Accessible Virtual Reality Techniques?	112
7.8.5	Limitations and Future Work.....	113
7.9	Summary	114
8	A Pattern Overlay Shader for Virtual Reality Users with Color Vision Deficiencies	115
8.1	Theory	116
8.2	Implementation	118
8.3	Limitations and Future Work.....	123
8.4	Summary	124
9	Experimentvr: An Extensible Testbed for Designing and Automating Virtual Reality Studies.....	125
9.1	Design Considerations	126
9.2	Testbed Design and Architecture.....	127
9.3	Experiment Setup: An Example.....	129
9.4	Tasks	130

9.4.1	Tasks: Timed.....	131
9.4.2	Tasks: Locomotion.....	131
9.4.3	Tasks: Selection	133
9.5	Data Logs	134
9.6	Running the Experiment: An Example	135
9.7	Implementation	135
9.8	Experimentvr Pilot Test	136
9.9	Limitations and Future Work.....	136
9.10	Summary	136
10	Comparing Locomotion Techniques With People with Upper-Body Impairments	138
10.1	Study Method.....	139
10.1.1	Participants.....	139
10.1.2	Custom Testbed	140
10.1.3	Locomotion Techniques: Overview.....	141
10.1.4	Locomotion Techniques: Astral Body	143
10.1.5	Locomotion Techniques: Chicken Acceleration.....	143
10.1.6	Locomotion Techniques: Grab and Pull	143
10.1.7	Locomotion Techniques: Sliding Looking	144
10.1.8	Locomotion Techniques: Teleport.....	144
10.1.9	Locomotion Techniques: Throw Teleport	144
10.1.10	Procedure	145
10.1.11	Design and Analysis: Overview.....	147
10.1.12	Design and Analysis: Task Performance Measures.....	147

10.1.13	Design and Analysis: Post-Task Questionnaire Measures.....	148
10.1.14	Design and Analysis: Interview Data	149
10.2	Results.....	149
10.3	Performance Measures.....	149
10.3.1	Trial Time	149
10.3.2	Hit Rate	150
10.3.3	Obstacles Hit.....	151
10.4	Questionnaire Responses	152
10.4.1	Presence	152
10.4.2	Simulator Sickness.....	153
10.4.3	Mental Demand.....	153
10.4.4	Physical Demand	154
10.4.5	Performance	154
10.4.6	Effort.....	155
10.4.7	Frustration	155
10.5	Interview Responses	156
10.5.1	Comparison of Locomotion Techniques.....	156
10.5.2	Accessibility and Interaction Design	157
10.5.3	Perceived Effect of Impairment.....	159
10.5.4	Overall User Experience	160
10.5.5	Enjoyment and Challenge	161
10.6	Summary	162
10.6.1	Astral, Sliding, and Teleport are Best in Terms of Performance and Workload	162

10.6.2	Chicken, Grab, and Throw are Worst in Terms of Performance and Workload	163
10.6.3	Most Comfortable, Fastest and Least Demanding Techniques are not Necessarily Preferred	164
10.6.4	Design Implications	165
11	Comparing Locomotion techniques Among People with and without Upper-Body Impairments	168
11.1	Study Method.....	168
11.1.1	Participants.....	168
11.1.2	Apparatus	169
11.1.3	Procedure	170
11.1.4	Design and Analysis	170
11.1.5	Outcome Measures.....	171
11.2	Results: Performance	172
11.2.1	Trial Time	172
11.2.2	Hit Rate	175
11.2.3	Obstacles Hit.....	177
11.3	Results: Post-Task Questionnaire	177
11.3.1	Presence	178
11.3.2	Sickness.....	178
11.3.3	Mental Demand.....	179
11.3.4	Physical Demand	180
11.3.5	Effort	182
11.3.6	Frustration.....	183

11.3.7	Performance	184
11.4	Summary	185
11.4.1	Sliding Looking: A Universally Optimal Technique?	185
11.4.2	Teleport: Impairment Level Analysis	186
11.4.3	Relative Importance of Performance Measures	187
11.4.4	Importance of Group-Wise vs. Impairment Level-Wise Comparisons	187
11.4.5	Reflecting on the Post-Task Questionnaire.....	188
11.4.6	Summary	188
12	Metrics That Encode Differences in Virtual Reality Interaction between People with and without Upper-Body Impairments	190
12.1	Motivation.....	190
12.2	Study Method.....	191
12.2.1	Raw Data.....	192
12.2.2	Metrics: Overview	192
12.2.3	Movement-Related Metrics	193
12.2.4	Button-Related Metrics	195
12.2.5	Target-Related Metrics	196
12.2.6	Design and Analysis	197
12.3	Results.....	198
12.3.1	Movement-Related Metrics: Device Variability.....	199
12.3.2	Movement-Related Metrics: Device Extent.....	202
12.3.3	Movement-Related Metrics: Device Angle Variability	205
12.3.4	Movement-Related Metrics: Device-Pair Variability.....	206

12.3.5	Movement-Related Metrics: Device Velocity and Acceleration	207
12.3.6	Movement-Related Metrics: Device Angular Velocity and Acceleration	211
12.3.7	Movement-Related Metrics: Submovement Count	214
12.3.8	Button-Related Metrics: Thumbstick Variability and Extent	215
12.3.9	Button-Related Metrics: Device Pressure	217
12.3.10	Target-Related Metrics: Movement Variability, Target Re-Entry Count, and Axis Crossed Count	218
12.4	Summary	220
12.4.1	Head-Related Metrics Reveal Group Differences	220
12.4.2	Unintentional Movements Differentiate Groups	221
12.4.3	Explaining Performance Differences: Astral	223
12.4.4	Explaining Performance Differences: Chicken	223
12.4.5	Explaining Performance Differences: Grab	223
12.4.6	Explaining Performance Differences: Sliding	225
12.4.7	Explaining Performance Differences: Teleport	225
12.4.8	Explaining Performance Differences: Throw	225
12.4.9	Limitations and Future Work	226
12.4.10	Summary	226
13	Predicting Performance with Locomotion techniques based on questionnaire and Calibration Task input	228
13.1	Questionnaire Scenario: Locomotion Technique Recommendation Based on Questionnaire Responses	229
13.2	Questionnaire Scenario: Study Method	230

13.2.1	Model Features.....	230
13.2.2	Performance Measures.....	230
13.2.3	Dataset.....	231
13.2.4	Modelling Method	231
13.2.5	Feature Selection.....	232
13.2.6	Model Evaluation.....	232
13.2.7	Ranking Techniques.....	233
13.3	Results: Questionnaire Scenario	233
13.3.1	Trial Time: Overview	233
13.3.2	Trial Time: Feature Selection	234
13.3.3	Trial Time: Training a Model	236
13.3.4	Trial Time: Ranking Techniques	238
13.3.5	Hit Rate: Overview	241
13.3.6	Hit Rate: Optimal Features	241
13.3.7	Hit Rate: Training a Model	243
13.3.8	Hit Rate: Technique Recommendations	246
13.3.9	Movement Variability: Overview	246
13.3.10	Movement Variability: Optimal Features	247
13.3.11	Questionnaire Scenario: Summary	248
13.4	Calibration Scenario: Using a Locomotion Technique to Calibrate the Recommender System.....	249
13.5	Study Method: Calibration Scenario.....	249
13.5.1	Outcome Predictions.....	250

13.5.2	Dataset.....	250
13.5.3	Modelling Method	251
13.5.4	Feature Selection.....	251
13.5.5	Model Evaluation.....	251
13.5.6	Ranking Techniques.....	252
13.6	Results: Calibration Scenario.....	253
13.6.1	Trial Time: Optimal Features.....	253
13.6.2	Trial Time: Training Models	255
13.6.3	Trial Time: Ranking Techniques	257
13.6.4	Calibration Scenario: Summary	260
13.7	Summary	260
13.7.1	Questionnaire Scenario: Selected Features	260
13.7.2	Calibration Scenario: Selected Features	261
13.7.3	Outcome Measures.....	261
13.7.4	Model Performance and Fit: Scenario 2	262
13.7.5	Comparing Recommendations from the Questionnaire and Calibration Scenarios	262
13.7.6	Recommendations.....	265
13.7.7	Limitations and Future Work.....	266
14	Conclusion	268
14.1	Selection of Existing Techniques for Accessibility	269
14.2	Designing Novel Techniques Based on User and Virtual Environment Characteristics	270

14.3	Identifying Trade-offs in Design Considerations to Personalize Virtual Reality	
	Interaction	271
14.4	Predicting Performance for a New Technique.....	273
14.5	Discussion of the Locomotion Techniques Evaluated.....	274
14.6	Discussion of Participant-Related Observations.....	275
14.7	Contributions	277
14.7.1	Empirical Contributions.....	277
14.7.2	Artifact Contributions	278
14.7.3	Methodological Contributions	278
14.7.4	Theoretical Contributions	278
14.8	Limitations and Future Work.....	279
14.9	Reflections and Insights.....	282
14.9.1	Compensating for Diverse and Small Samples with Large Amounts of Data....	282
14.9.2	Cutting My Losses and Taking Risks	283
14.10	Final Remarks	284
15	References.....	286
A.1	Classification of Scene-Viewing Techniques Using the Taxonomy	303
A.2	Classification of Virtual Reality Applications Using the Scene Taxonomy.....	311
A.3	QuickDASH Questionnaire	313
A.4	Technology-Related (Physical) Impairments Questionnaire (TRIQ).....	314
A.5	Group Comparison Study: Non-Significant Outcome Bar Plots	316
A.5.1	Obstacles Hit by Impairment Level	316
A.5.2	Presence by Group	316

A.5.3	Presence by Impairment Level.....	317
A.5.4	Sickness by Group	317
A.5.5	Sickness by Impairment Level.....	318
A.5.6	Mental Demand by Group	318
A.5.7	Effort by Group.....	319
A.5.8	Performance by Group.....	319
A.6	Metrics Study: Remaining Metrics	320
A.6.1	Movement-Related Metrics: Device Velocity and Acceleration.....	320
A.6.2	Movement-Related Metrics: Device Angular Velocity and Acceleration.....	321
A.6.3	Button-Related Metrics: Thumbstick Variability and Extent	324
A.6.4	Movement-Related Metrics: Device-Pair Variability and Extent.....	325
A.6.5	Movement-Related Metrics: Device Submovement Count	326
A.7	Questionnaire Scenario: Data Table Snippet	327
A.8	Questionnaire Scenario: Data Table Used for Ranking Locomotion Techniques	328
A.9	Calibration Scenario: Data Table Snippet.....	333
A.10	Calibration Scenario: Tuning Parameters	335
A.11	Calibration Scenario: Error	336
A.12	Calibration Scenario: Data Table Used for Ranking Locomotion Techniques	337

1 LIST OF FIGURES

Figure 1. Openness: (from left to right) outdoor, indoor, abstract.....	37
Figure 2. Scale: (from left to right) smaller, larger.....	37
Figure 3. Area: (from left to right), large, small.	38
Figure 4. Scene transitions: (from left to right) frequent, infrequent.....	39
Figure 5. World in miniature (WIM) prototype (Stoakley et al., 1995). Left: A miniature model of the VE. Right: the miniature model of the VE with the human-scale VE behind it. I classified this VE with the scene taxonomy.....	49
Figure 6. A scene from the VR app The Room. This scene would be categorized with the scene taxonomy in the same manner as WIM.	50
Figure 7. Object of Interest: a. The user is facing the octopus in the scene (birds-eye view), b. She selects the turtle icon with her laser pointer, c. Her camera rotates to face the turtle in the scene (birds-eye view), d. She is facing the turtle in the scene.....	58
Figure 8. Proxemics Snapping: a. The user's view is zoomed out so that the environment is in the frame, b. She picks up and moves his dark blue avatar, c. She places his avatar in the light blue snap drop zone where his avatar orients to a socially appropriate distance and orientation relative to another avatar, d. She deactivates proxemics snapping and returns to a first-person view facing the red avatar.	60
Figure 9. Rearview Mirror: a. The user is facing an empty room (birds-eye view), b. She sees an avatar in her rearview mirror, c. She presses a controller button to flip 180° (birds-eye view), d. She faces the blue avatar.	62
Figure 10. Participants' head movements were limited with a chinrest during the study.	64

Figure 11. (a) A user in a virtual environment surrounded by points of interest (POIs). (b) The user pulls up a display of out-of-view POIs and selects a POI to orient to. (c) The user's first-person camera rotates to the selected POI. (d) The user faces the desired POI.	80
Figure 12. The two Representations implemented with Nearmi. Portal view icon of the POI in real time (left). 3D miniature mesh icon of the POI mesh (right).	87
Figure 13. Nearmi's Display and Selection combinations. Columns from left to right show Selection instances: tap, ray cast, wand, console. Rows from top to bottom show Display instances: attached to object, floating.	88
Figure 14. Wand: After tapping an icon as shown in Figure 13, the user brings an icon towards his/her headset. A Selection is made when the headset and icon collide.	89
Figure 15. The three camera Transitions implemented in Nearmi. The first-person camera rotates smoothly in the continuous Transition (left), jumps by 30-degree increments in the discrete Transition (middle), and instantly cuts from a starting orientation to face the POI in the instant Transition (right).	91
Figure 16. A high-level overview of the Nearmi framework manager, the mechanism base classes, and their relationships.	92
Figure 17. The alchemy lab VE used to demo Nearmi. Environment by Unity Technologies. ...	95
Figure 18. Researchers examined how the orientation of hatching lines can be used to differentiate colors. Figure from Sajadi et al. (2013).	116
Figure 19. Different color vision deficiencies have different copunctal points, where the confusion lines converge in xyY space. Left: Deuteranopia (red-green CVD) has a copunctal	

point outside of the color space. Middle: Tritanopia (yellow-blue CVD) and, right: Protanopia (red-green CVD) have copunctal points on the edges of the color space.	117
Figure 20. On a confusion line, accessibility problems occur when there are two different hues, but not when there are different shades of the same hue.	117
Figure 21. Left: The white point in the xyY color space (labelled E) marks the transition between warm and cool hues. Right: The separation line is a circle with the center at the copunctal point and a radius the distance between the center and the white point. In this example, I applied a pattern to the left region of the separation line to make reds distinguishable from greens for people with protanopia.....	118
Figure 22. Three patterns being applied over different hues for people with left: deuteranopia (red-green CVD), middle: protanopia (red-green CVD), right: tritanopia (yellow-blue CVD).	119
Figure 23. The user could edit properties of the pattern including left: which side of the color space it covered and, right: how much of the color space was covered, or in other words, the “sensitivity” of the shader.	119
Figure 24. Top: A person with protanopia’s view of the sports teams. The darker jersey is red, and the lighter jersey is green. Bottom: A person with protanopia’s (red-green CVD) view of the sports teams with a pattern overlaid on the red jerseys.....	121
Figure 25. A scenario in which a user with red-green CVD is looking for apples. Top: It is difficult for the user to immediately identify the red apples. Bottom: With the shader applied to reds, it is easier to identify the apples.	122

Figure 26. Left: An AR view of a red skirt and green dress from the perspective of someone with protanopia (red-green CVD). Right: The user’s view of the same garments with a pattern overlaid on the red skirt (the white dots are the part of the AR surface-tracking capability). 123

Figure 27. The architecture for Experimentvr. In this example, the study is a 2x2 design with one factor being Interaction Technique and the other being Target Distance. There are two scenes, one per technique with two conditions in each scene. Each condition has four trials in which the participant hits a target with one of the interaction techniques. The four trials are repeated three times as a trial block, for a total of 24 trials per technique. Low-level controller and headset data as well as performance data are recorded. 129

Figure 28. The researcher’s view of the testbed environment while the participant is completing a timed task called “Dual wrists flex down hold.” The trial instructions are to “flex both wrists downward and hold” and the participant has 13 seconds before the trial countdown elapses. Instructions only visible to the researcher are displayed in the top left and right corners. Information about advancing through the study as well as detailed information about the study progress is displayed in the top right corner. 131

Figure 29. Left: A view of the participant’s controllers (floating black objects) in the locomotion task environment. The participant is currently using the Teleport technique. The locomotion targets are the blue and green cylinders in the background. The green target is the current target that the participant must navigate to. Right: The participant’s view while they are inside the current target. The blue “7” is the trial countdown and the pink “2” is the dwell countdown. 132

Figure 30. A participant using the ray cast technique to select a target in a Selection task. The dwell countdown shows that the participant must keep the ray on the target for one more second before the trial is considered a “hit”.....	133
Figure 31. (A) A bird’s-eye view of the test environment consisting of six blue targets and paths on a plane. The participant started the task in the center of the plane. The current target would turn from blue to green. (B) The participant’s view when they moved inside the current target, which was green. The blue countdown shows the time remaining to complete the trial and the pink countdown shows the seconds the participant must remain inside the target for the trial to be considered a “hit.” (This is analogous to dwell-based selection on a touch screen.).....	141
Figure 32. (A) The participant’s view of the test environment and 30-second trial countdown. The target that the participant is navigating to is green and other targets in the environment are blue. The participant must also maneuver around the work benches, lighting rigs, and paint cans. (B) A participant performing the task using one hand to stabilize the controller and the other to press the buttons.....	146
Figure 33. (A) Bar plot of Trial Time by Technique. Lower is better. (B) Bar plot of target Hit Rate by Technique. Higher is better. Means are above the bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*), $p<.1$ (.).	150
Figure 34. (A) Bar plot of Obstacles Hit per trial by Technique. Lower is better. Means are above the bars. (B) Boxplot of Likert responses for Presence by Technique. Medians are above and inside the boxes. Higher is greater experienced Presence. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*), $p<.1$ (.).	152

Figure 35. (A) Boxplot of responses for Sickness by Technique. Lower is less experienced Sickness. (B) Boxplot of Likert responses for Mental Demand by Technique. Lower is less perceived Mental Demand. Medians are above and inside the boxes. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*), $p < .1$ (.). 153

Figure 36. (A) Boxplot of Likert responses for Physical Demand by Technique. Lower is less perceived Physical Demand. (B) Boxplot of Likert responses for Performance by Technique. Higher is better perceived Performance. Medians are above and inside the boxes. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*), $p < .1$ (.). 154

Figure 37. (A) Boxplot of Likert responses for Effort by Technique. Lower is less perceived Effort. (B) Boxplot of Likert responses for Frustration by Technique. Lower is less perceived Frustration. Medians are above and inside the boxes. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*), $p < .1$ (.). 155

Figure 38. Bar plot of mean Trial Time by Technique and Group. Lower is better. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*). 173

Figure 39. Bar plot of mean Trial Time by Technique and Impairment Level. Lower is better. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*). . 174

Figure 40. Bar plot of mean Hit Rate by Technique and Group. Higher is better. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*). 175

Figure 41. Bar plot of mean Hit Rate by Technique and Impairment level. Higher is better. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*). 176

Figure 42. Bar plot of mean Obstacles Hit by Technique and Group. Lower is better. Means are at the bottom of error bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*). 177

Figure 43. Bar plot of mean Mental Demand rating (1-7) by Technique and Impairment Level. Lower is better. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).	179
Figure 44. Bar plot of mean Physical Demand rating (1-7) by Technique and Group. Lower is better. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).	180
Figure 45. Bar plot of mean Physical Demand rating (1-7) by Technique and Impairment level. Lower is better. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).	181
Figure 46. Bar plot of mean Effort rating (1-7) by Technique and Impairment Level. Lower is better. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).	182
Figure 47. Bar plot of mean Frustration rating (1-7) by Technique and Group. Lower is better. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*). ..	183
Figure 48. Bar plot of mean Frustration rating (1-7) by Technique and Impairment Level. Lower is better. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).	184
Figure 49. Bar plot of mean Performance rating (1-7) by Technique and Impairment Level. Higher is better. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).	185
Figure 50. Submovements are movements within a rapid aimed movement with dips in velocity (pauses) before and after their occurrence.	195

Figure 51. In this image, the user's path is the blue line, and the optimal path is the horizontal line. Movement Variability is a measure of how much the user's path differs from the optimal path. The figure also depicts a user first entering, then exiting, then re-entering a target. In this case, the Target Re-entry Count would be 1. Additionally, the figure depicts the user crossing the optimal path to the target. In this case the Axis Crossed Count would be 2..... 197

Figure 52. Participants With Impairments had significantly higher Head Variability for all techniques except Chicken. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*). 199

Figure 53. Participants With Impairments had significantly higher Left Controller Variability for Astral, Sliding, and Teleport. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*). 200

Figure 54. Participants With Impairments had significantly higher Right Controller Variability for all techniques except Grab and Throw. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*). 201

Figure 55. Participants With Impairments had significantly higher Head Extent values. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*). 202

Figure 56. Participants With Impairments had significantly higher Left Extent for all techniques except Grab and Throw. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*). 203

Figure 57. Participants With Impairments had significantly higher Right Extents for all techniques except Grab and Throw. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*). 204

Figure 58. Participants With Impairments had significantly higher Head Angle Variability for Teleport. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).....	205
Figure 59. Participants With Impairments had significantly higher Right Hand to Head Variability for Chicken, Grab, and Throw. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).....	206
Figure 60. Participants With Impairments had significantly lower Head Velocity for Chicken. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*)..	207
Figure 61. Participants With Impairments had significantly lower Head Acceleration for all Techniques. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).....	208
Figure 62. Participants With Impairments had a significantly lower Right Velocity for Chicken and Throw. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).....	209
Figure 63. Participants With Impairments had a significantly lower Right Acceleration for Throw. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).....	210
Figure 64. Participants With Impairments had significantly lower Angular Head Velocity for Chicken and Sliding. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).....	211

Figure 65. Participants With Impairments had significantly lower Angular Head Acceleration for all Techniques. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).	212
Figure 66. Participants With Impairments had significantly lower Angular Right Velocity for Grab and Throw. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).	213
Figure 67. Participants With Impairments had significantly Higher Right Submovement Counts for Grab and Throw. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).	214
Figure 68. Participants With Impairments had significantly higher Right Thumbstick Variability for Teleport. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).	215
Figure 69. Participants With Impairments had significantly higher Right Thumbstick Extent for Teleport. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).	216
Figure 70. Participants With Impairments had a significantly higher Right Grip Pressure for Astral, Chicken, and Sliding. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).	217
Figure 71. Participants With Impairments had significantly higher Movement Variability for Astral. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).	218

Figure 72. Participants with impairments had significantly higher Axis Crossed Count for Grab.

Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*). . 219

Figure 73. Participants With Impairments had significantly higher Target Re-entry Count for

Astral and Grab. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$

(**), $p < .05$ (*). 220

Figure 74. Green lines represent head positions of people Without Impairments and blue lines

represent those of people With Impairments. Left: Head positions of all 40 participants while

using Chicken. There was not a significant difference in Head Variability for Chicken. Right:

Head positions of 40 participants while using Teleport. People Without Impairments had a

significantly lower Head Variability, which is reflected in the smaller spread of green

compared to blue head positions for Teleport. 221

Figure 75. A plot of the 20 features selected with recursive feature elimination ordered by their

importance ranking. Technique was highly ranked compared to all other features..... 235

Figure 76. A density plot of R^2 values from cross-validation. There is high variability, but the

distribution is skewed toward higher R^2 variables suggesting a bias towards a subset of

participants. 236

Figure 77. A learning curve showing how performance increases with additional participants. R^2

values increase by 0.10 from 15 to 40 participants..... 237

Figure 78. A plot of actual vs. predicted values, where the line represents perfect prediction. . 238

Figure 79. Slope graph of actual and predicted techniques by rank from fastest to slowest (rank

1-6) for the questionnaire scenario. The y-axis is alphabetically ordered. 239

Figure 80. Left: Heat map of overall ranking accuracy for all participants, middle: participants with impairments, right: and participants without impairments. The model accurately predicted the top three techniques for participants with impairments, but only the top technique for participants without impairments.....	240
Figure 81. A plot of the top 28 features ordered by their importance ranking found during feature selection. Recursive feature elimination resulted in a final set of the top 16 features.	242
Figure 82. A plot of actual vs. predicted values, where the line represents perfect prediction. Values are clustered towards high Hit Rates indicating a ceiling effect.....	244
Figure 83. A density plot of Hit Rate models' R^2 values from cross validation. There is high variability, and the distribution is skewed toward lower R^2 variables suggesting the model underfit the data.	245
Figure 84. A learning curve showing that Hit Rate prediction performance increases with additional participants, but plateaus around an R^2 of 0.20 after 20 participants.....	246
Figure 85. A plot of the top 13 features ordered by their importance ranking found during feature selection. Recursive feature elimination resulted in a final set of the top eight features.	248
Figure 86. The top five features that were selected with RFE when modelling techniques' Trial Times using all other techniques as calibration techniques. Each type of feature was represented across the six feature sets: Button-, movement-, and target-related features from the metrics exploration study, as well as QuickDASH-, TRIQ-,performance, and post-task questionnaire-related features.	254

Figure 87. Learning curves for: Top left: Astral, top middle: Chicken, top right: Grab, bottom left: Sliding, bottom middle: Teleport, bottom right: Throw. Models trained on five-technique datasets.	257
Figure 88. Slope graph of actual and predicted techniques by rank from fastest to slowest (rank 1-6). The y-axis is ordered alphabetically.....	258
Figure 89. Left: Heat map of overall ranking accuracy for all participants, middle: participants with impairments, right: and participants without impairments. The model accurately predicted the top three techniques for ranks 1-3 and 6 for both groups.....	259
Figure 90. Top: Heat maps showing accuracy of the predicted technique rank from the questionnaire scenario. Bottom: Heat maps showing accuracy of the predicted technique rank from the calibration scenario. Left: Accuracy for all participants, middle: participants with impairments, and right: participants without impairments.	263
Figure 91. The QuickDASH questionnaire used in the locomotion technique study to measure upper-body impairment. QuickDASH is a validated questionnaire (Gummesson et al., 2006).	313
Figure 92. The Technology-Related Physical Impairment Questionnaire (TRIQ) questionnaire used in the locomotion studies. It is not standardized but based on literature in accessible computing (Findlater et al., 2010; Mott et al., 2020).	314
Figure 93. Bar plot of mean Obstacles Hit by Technique and Impairment Level. Lower is better. Means are at the bottom of error bars. No significant differences were found.....	316
Figure 94. Bar plot of mean Presence rating (1-6) by Technique and Group. Higher is better. Means are at the bottom of bars. No significant differences were found.	316

Figure 95. Bar plot of mean Presence rating (1-6) by Technique and Impairment Level. Higher is better. Means are at the bottom of bars. No significant differences were found.	317
Figure 96. Bar plot of mean Sickness rating (1-3) by Technique and Group. Lower is better. Means are at the bottom of bars. No significant differences were found.	317
Figure 97. Bar plot of mean simulator Sickness rating (1-3) by Technique and Impairment Level. Lower is better. Means are at the bottom of bars. No significant differences were found. ...	318
Figure 98. Bar plot of mean Mental Demand rating (1-7) by technique and group. Lower is better. Means are at the bottom of bars. No significant differences were found.	318
Figure 99. Bar plot of mean Effort rating (1-7) by Technique and Group. Lower is better. Means are at the bottom of bars. No significant differences were found.	319
Figure 100. Bar plot of mean Performance rating (1-7) by Technique and Group. Higher is better. Means are at the bottom of bars. No significant differences were found.	319
Figure 101. Participants With Impairments had a significantly lower Left Velocity for Teleport.	320
Figure 102. There were no significant differences in Left Acceleration between Groups for all Techniques.	320
Figure 103. Participants With Impairments had a significantly lower Left Angular Velocity for Grab, Throw, and Teleport.....	321
Figure 104. Participants With Impairments had a significantly lower Left Angular Acceleration for Teleport.	321

Figure 105. There were no significant differences in Right Angular Acceleration between Groups for all Techniques.....	322
Figure 106. Participants With Impairments had a significantly higher Left Trigger Pressure for Sliding and Throw and significantly lower Left Trigger Pressure for Astral and Chicken. ..	322
Figure 107. Participants With Impairments had a significantly higher Right Trigger Pressure for Astral, Chicken, Sliding, and Throw and a significantly lower Right Trigger Pressure for Grab.....	323
Figure 108. Participants With Impairments had a significantly higher Left Grip Pressure for Astral, Chicken, and Sliding, and a significantly lower Left Grip Pressure for Throw.....	323
Figure 109. Participants With Impairments had a significantly lower Left Thumbstick Variability for Astral because of an error with Astral’s setup for the other Group.....	324
Figure 110. Participants With Impairments had a significantly lower Left Thumbstick Variability for Astral because of an error with Astral’s setup for the other Group.....	324
Figure 111. Participants With Impairments had significantly higher Left Hand to Head Variability for Grab and Throw.	325
Figure 112. Participants With Impairments had significantly higher Right Hand to Left Hand Variability for Grab and Sliding.	325
Figure 113. Participants With Impairments had significantly higher Head Submovements for Grab.....	326
Figure 114. Participants With Impairments had significantly higher Left Submovements for Astral, Grab, and Teleport.....	326

2 LIST OF TABLES

Table 1. The visual properties of the scene taxonomy.....	36
Table 2. The task type dimension of the scene taxonomy breaks down the scene based on the task the user is performing in it.	40
Table 3. This heatmap shows an overview of the scene-viewing techniques in research and is a summary of Appendix A.1. It reveals popular tasks and visual properties as well as gaps in research. I categorized the virtual environments and tasks that were used to evaluate the scene-viewing technique in the paper and then counted the papers in each category to produce the heatmap. The darker the cell, the more papers on scene-viewing techniques in that category.	51
Table 4. Tasks participants completed with panning and the new scene-viewing techniques.	65
Table 5. This table presents the participant counts of their preferred technique and the technique they perceived as easier. It also presents NASA-TLX median scores and IQR for the new technique and panning. Higher NASA-TLX scores mean higher perceived workload. The data indicated that for some participants, their preferred technique was not the same as the technique they thought was easier.	67
Table 6. Participant demographics and preferred Nearmi implementations.....	96
Table 7. A summary of the videos participants watched and the outcome of each phase.	97
Table 8. Participants' self-reported upper-body motor impairment conditions.	140
Table 9. The six locomotion techniques tested in the study, instructions for using the technique, the control mechanism for forward movement and direction, and the control mechanism type. The table also presents the body part, effort level according to Di Luca et al. (2021), and	

maximum number of hands required to use the technique. Finally, Di Luca et al.'s categorization of the technique in their taxonomy is included.	142
Table 10. Participant responses to final interview questions comparing the locomotion techniques. “Top 1” refers to the technique they would use most, “top 2” is the second-most, and “top 3” is the third-most.	156
Table 11. The table presents the proposed metric, a description of the metric, its type, and the reason the metric might capture differences between groups. Note: “device” refers to the headset and right or left controller. Metrics are technique-agnostic.	193
Table 12. A table of the tuning parameters for the elastic net (alpha and lambda), R^2 values, Bayesian Information Criterion (BIC), root mean squared error (RMSE), and mean absolute error (MAE) for models predicting Trial Time with all questionnaire features, only QuickDASH features, only TRIQ features, and the selected features from feature selection.	234
Table 13. Table of ranking accuracy and mean absolute percent error (MAPE) for participants with impairments, participants without impairments, and all participants. The overall ranking accuracy and MAPE of the six techniques was calculated, as well as the accuracy and MAPE for each rank individually.	241
Table 14. A table of the tuning parameters for regularized random forest: Number of randomly selected predictors (numPred), regularization value (coefReg), and importance coefficient (coefImp), R^2 values, root mean squared error (RMSE), and mean absolute error (MAE) for models predicting Hit Rate with all questionnaire features, only QuickDASH features, only TRIQ features, and selected features from the feature selection.	241

Table 15. A table of the tuning parameters for the elastic net (alpha and lambda), R^2 values, root mean squared error (RMSE), and mean absolute error (MAE) for models predicting the Movement Variability with all questionnaire features, only QuickDASH features, only TRIQ features, and the selected features from feature selection.....	247
Table 16. R^2 values for models tested on single and all calibration techniques to predict Trial Time for the outcome technique. Using Sliding as calibration a technique resulted in the most accurate and best fitting models for most techniques' Trial Times.	256
Table 17. Table showing ranking accuracy, or in other words, the number of techniques that were predicted correctly for a rank out of the total techniques that were actually that rank based on calibration with Sliding. Mean absolute percent error (MAPE) is also reported. Accuracy and MAPE are reported for participants with impairments, participants without impairments, and all participants. The overall ranking accuracy and MAPE of the six techniques were calculated, as well as the accuracy for each rank individually.	259
Table 18. Overall ranking accuracy and within-rank accuracy for the questionnaire scenario (QS) and the calibration scenario (CS). Results are presented for all participants, participants with impairments, and participants without impairments.	264
Table 19. Overall mean absolute percent error (MAPE) and within-rank MAPE for the questionnaire scenario (QS) and the calibration scenario (CS). Results are presented for all participants, participants with impairments, and participants without impairments.....	264
Table 20. Scene-viewing technique papers classified based on the visual properties and tasks that were used to evaluate the scene-viewing technique. Classification is based on the inspection	

of images in the paper and supplemental videos. Papers are grouped and labeled by the type of scene-viewing technique (e.g., “Scene Modification”).....	303
Table 21. Twenty-nine of the most popular virtual reality (VR) applications in 2019 and 2020 classified based on the main visual properties and tasks of the virtual environment (VE). ..	311
Table 22. A sample of the dataset used to train the model in the questionnaire scenario. Data are from two participants, a participant with impairments (PWI) and a participant without impairments (PWO). Column names are provided in the column name key.....	327
Table 23. The data table used to rank locomotion techniques for each participant. Times are predicted with the questionnaire scenario model. Columns represent the Participant ID, Technique, the average Actual Trial Time in seconds, the average Predicted Trial Time in seconds and the Difference between the Predicted and Actual Trial Times in seconds. Participant IDs leading with “P” are participants with impairments and IDs leading with “PC” are participants without impairments.	328
Table 24. A sample of the dataset used to train the model in the calibration scenario. Data are from two participants, a participant with impairments (PWI) and a participant without impairments (PWO). Column names are provided in the column name key.....	333
Table 25. The tuning parameters used for the final models in the calibration scenario.	335
Table 26. The root mean squared error and mean absolute error of the final models in the calibration scenario.	336
Table 27. The data table used to rank locomotion techniques for each participant. Times are predicted with the calibration scenario model. Columns represent the Participant ID, Technique, the average Actual Trial Time in seconds, the average Predicted Trial Time in	

seconds and the Difference between the Predicted and Actual Trial Times in seconds.

Participant IDs leading with “P” are participants with impairments and IDs leading with “PC”

are participants without impairments. 337

3 ACKNOWLEDGMENTS

I extend my gratitude to my family and friends for their support throughout this journey. Special thanks to Jake for his encouragement during both the good times and the challenging ones. Thank you to my mentors during my internships, Martez Mott (Microsoft), José Echevarria (Adobe), and Cole Gleason (Apple) for their helpful research and career advice. I am very appreciative of my committee members for their invaluable insights, which significantly strengthened the final chapters of this dissertation. I also want to thank the Here and Now Project for their assistance in advertising my study; their help was indispensable. This work was generously supported by the Apple AI/ML Fellowship and the Meta Reality Labs Research Fund. I am grateful to these organizations for their support.

To my family

4 INTRODUCTION

Mobility, defined as the freedom to move within the physical environment, is crucial for maintaining a high quality of life (Shafrin et al., 2017). It not only enables access to essential services such as healthcare, banking, and shopping, but also helps reduce social isolation (Zhao et al., 2023). A lack of mobility can diminish autonomy and independence, as individuals may depend on others for daily tasks, thereby decreasing their quality of life (Shafrin et al., 2017).

People with physical impairments, such as those affecting their ability to walk, reach, or grasp, often experience limited mobility, which restricts both their ability to travel and interact physically with their environment. Given the significant impact of mobility on quality of life, it is essential that individuals with physical impairments have the means to overcome these mobility challenges.

Technology has significantly enhanced the lives of individuals with physical impairments by helping them circumvent various barriers. For instance, smart home technology now allows people to control temperature, lighting, doors, and security systems through a smartphone, thus eliminating the need for physical movement to adjust these settings. Additionally, ridesharing apps have improved transportation for wheelchair users by enabling them to request accessible vehicles, thereby expanding their travel options. Customized keyboards, mouse alternatives (like trackballs or joystick mice), and touch screen devices allow individuals with limited fine motor skills to use computers effectively. Voice recognition software can also enable those who find manual input challenging to control their computers through speech. Furthermore, the COVID-19 pandemic accelerated the development and adoption of remote working tools, making it unnecessary for people with limited mobility to commute to schools or workplaces.

Virtual reality (VR) has opened new avenues for individuals with limited mobility as a result of a physical impairment, allowing them to engage in recreational activities like skiing and rock

climbing that would typically be inaccessible. VR not only enables users to virtually "travel" and connect with others in ways previously unattainable with traditional technology, but it also provides a more immersive social experience. This can be particularly transformative for those with physical impairments who may face geographical constraints, helping to mitigate social isolation and foster deeper connections with others.

VR is rapidly becoming a mainstream technology in the consumer market due to its engaging experiences, reduced hardware size, and enhanced affordability. Moreover, VR now appeals to a broader and more diverse user base, as its applications have extended beyond entertainment to encompass social connectivity, education, and professional uses.

However, VR interaction techniques differ significantly from those used on desktop and mobile devices and are still in the early stages of research and development. As is often the case with emerging technologies, the accessibility of VR has not been a primary focus. Studies indicate that VR in its current form is not fully accessible to individuals with disabilities, ranging from physical to vision impairments (Gerling & Spiel, 2021; Mott et al., 2020; Wong et al., 2018). The accessibility of VR is especially crucial for those with physical impairments, as the emphasis on movement and physical interaction can pose significant challenges for this group.

Research demonstrates that people with impairments often prefer to use mainstream technology instead of assistive devices. Reasons for this preference include that assistive devices are expensive, difficult to configure, and burdensome to maintain (Dawe, 2004). In addition, assistive devices can be stigmatizing as they can draw unwanted attention to a user's disability (Shinohara & Wobbrock, 2011). For these reasons, it is important to investigate whether off-the-shelf VR devices can be adapted to people's physical abilities.

4.1 Virtual Reality Interaction

The ultimate goal of VR is to convince users that they are physically inhabiting a virtual world. As a result, VR environments are designed to resemble physical environments and virtual interactions mirror how people interact with the physical world. If an individual has difficulty interacting in the physical world, those same challenges will exist in VR because of the emphasis on immersion and metaphor in virtual interactions.

Before interacting with objects in physical and virtual reality, one must be able to first view or become aware of objects and then travel to them. Viewing and locomotion are therefore the most fundamental ways of interacting with the world for many people, and precede other interactions (e.g., object manipulation, social interaction, etc.). Making these interactions accessible in physical reality has been an area of much research and development and warrants similar attention in VR to enable other types of interactions.

In VR, interactions for viewing and navigating virtual environments (VEs) mirror those in the physical world. Users can view the VE by simply turning their heads or bodies, a process facilitated by head tracking, which involves the headset sensing the user's head orientation and then displaying the corresponding view of the VE. Head tracking represents one form of a scene-viewing technique, which is an interaction technique that facilitates viewing a VE. However, head tracking is just one of many scene-viewing techniques designed for VE.

As for locomotion, the most basic way of travelling in VR is similar to that in physical reality. The user can walk or use a wheelchair in physical space to move to another location in virtual space. The rendered scene is then updated to reflect the corresponding perspective of their new position. This technique is called “real walking” and like head tracking, it is only one example of a locomotion technique, or interaction techniques that enable users to travel in a VE.

Because VR transcends the limitations of the physical world, offering uniquely immersive experiences, VEs often do not match the structure of the physical world. For instance, if a user is navigating in a virtual space that is larger than their physical environment, the user will eventually run into physical a wall, which is a significant safety concern. Furthermore, many users prefer low-exertion interaction techniques to avoid fatigue, which often means staying stationary.

Due to these constraints, many scene-viewing and locomotion techniques have been developed that minimize body movement and rely on hand-held controller input, which enables 3D gestures and button input. For example, a common alternative to head tracking is camera panning, in which the user pushes a thumbstick to rotate the camera view while their head remains forward. A common alternative to room-scale walking is teleportation, in which the user aims their controller to select a location in the environment to jump to. In fact, there are many interaction techniques that facilitate viewing and travelling in VR, and novel techniques continue to be developed to this day.

While stationary, controller-based interaction techniques can improve accessibility for people with lower-body impairments, offloading interactions onto hand-held controllers can introduce new accessibility issues for people with upper-body impairments. Teleportation or camera panning, which primarily use button and thumbstick input, can be difficult for someone with limited dexterity or difficulty gripping a controller. Fortunately, different interaction techniques can be controlled using different parts of the body, enabling users to choose techniques that are minimally affected by their particular impairments.

While empirical research is still needed to fully evaluate the accessibility of scene-viewing and locomotion techniques, designers and researchers can use an informal approach by matching a person's specific impairments with the control mechanisms of various techniques. Ideally, VR

applications could introduce users to a range of scene-viewing and locomotion techniques during the onboarding process, rather than confining them to default interaction techniques. This approach could enhance accessibility and user experience of VR by tailoring interactions to individual needs.

4.2 Research Approach

Currently, there are two ways of identifying scene-viewing or locomotion techniques that are accessible. The first is to select a technique from a range of existing techniques, and the second is to design a novel technique. To ensure that techniques are accessible, designers and researchers must evaluate them. At this point in time there is little support for selecting, designing, or evaluating techniques that are accessible for people with physical impairments. Without tools or frameworks, researchers and designers will not be able to leverage the extensive prior work in VR to inform implementation and could end up redoing past research in a new context. Furthermore, tools and frameworks could guide the design and implementation of VR applications by researchers or designers who do not have an accessibility background, thereby increasing the possibility that mainstream applications are accessible to people with upper-body impairments.

The goal of my work is to facilitate the process of selecting, designing, and evaluating scene-viewing and locomotion techniques to improve the accessibility of VR. To aid in selecting techniques, I created a scene taxonomy that facilitates the process of selecting scene-viewing techniques that are accessible to people with limited head movement and optimal for the characteristics of a VE. I also evaluated six locomotion techniques with people with and without upper-body impairments to identify the most accessible technique for most users. Based on this study, I proposed a predictive model that can be used to recommend techniques to users based on information about their impairments. I also analyzed differences in how people with and without

impairments moved while using locomotion techniques with a set of metrics, which were then used as features in the predictive model.

To help design techniques, I showed how the scene-viewing taxonomy can be used to create novel scene-viewing techniques. I also devised a design framework to support the creation of accessible scene-viewing techniques that enable users to view points of interest in the environment. I then demonstrated how scene-viewing techniques can be designed to make VEs accessible to people with vision impairments, specifically for people with color vision deficiencies.

To help evaluate techniques, I built an extensible testbed that enables researchers to create and automate multi-factor and full-factorial VR experiment designs. The testbed also captures high-level performance data (e.g. time to complete a trial) as well as low-level device data (e.g. controller and headset position) at a user-defined sampling rate.

The thesis this work will demonstrate is:

The accessibility of virtual reality (VR) can be improved by the selection, design, and evaluation of scene-viewing and locomotion techniques using design frameworks, testbed evaluations, movement analysis, and predictive models.

4.3Dissertation Structure

The remainder of this dissertation is structured as follows:

Related Work
Chapter 5

Contextualizes this work in the literature on accessible VR, VR design frameworks, interaction techniques, technique evaluation, and user modelling.

Scene-Viewing Technique
Selection and Design
Chapters 6, 7, and 8

Introduces a taxonomy and design framework to support the selection and design of scene-viewing techniques. Proposes a design for a scene-viewing technique for people with vision impairments.

Supporting VR Interaction
Technique Evaluation
Chapter 9

Discusses the design and implementation of an extensible and customizable testbed for evaluating VR interactions.

Evaluating and Predicting
Performance with Locomotion
Techniques
Chapters 10, 11, 12, 13

Discusses the results of evaluating and predicting users' performance with locomotion techniques.

Conclusion
Chapter 14

Discusses overarching themes from this work and presents reflections on the research process.

5 RELATED WORK

This work is theoretically based in ability-based design and is situated amongst literature related to accessibility in virtual reality (VR). It is inspired by prior work on VR design frameworks, interaction technique evaluation, and user modelling.

5.1 Ability-Based Design

Ability-based design (ABD) is an approach to accessible technology design. Unlike other approaches such as assistive technology design, ABD argues that designers should leverage users' abilities rather than focusing on their disabilities, which can be achieved by designing systems that accommodate users (Wobbrock et al., 2011). As such, ABD systems would not require assistive add-ons or peripheral devices to make technology accessible. This shift of focus also unburdens the user from needing the knowledge, finances, or time to acquire and learn how to use a specialized technology to make another technology accessible (Wobbrock et al., 2011). Instead, this burden is shifted to the technology. Unlike inclusive design, ABD does not require designers to identify all the ways in which users with disabilities might be excluded from the design. Instead, with ABD, the technology design leverages individual abilities. There is evidence to suggest that unburdening the user from technology acquisition and configuration could result in lower abandonment rates by people who use assistive technologies (Bates & Istance, 2002; Goette, 1998). VR systems could benefit from ABD by adapting to users' abilities, making off-the-shelf devices accessible.

5.2 Accessibility in Virtual Reality

As VR becomes a mainstream technology, researchers have focused on improving its accessibility for people with disabilities, including people with vision, hearing, and mobility

impairments. Some communities of practice, such as XR Access¹, are working diligently toward identifying and addressing accessibility barriers endemic to existing VR systems. However, despite progress, VR hardware and interaction design assume a typical body that is able to see, hear, stand, walk, reach, and crouch, making it inaccessible to people whose bodies do not conform to this standard.

5.2.1 Vision Impairments

Researchers have explored solutions that improve the accessibility of existing systems for people who have vision impairments. Zhao et al. (2018) proposed a haptic device, called the Canetroller, that can simulate white cane use in virtual environments (VEs). The Canetroller enabled participants to create an accurate mental model of an indoor space and locate virtual objects within the space.

In addition to haptic information, novel devices have been built to provide auditory information (Gonzalez Penuela et al., 2022). For example, researchers developed a glove to help users request auditory information about objects in VEs using hand gestures. They plan to investigate the types of gestures users prefer and whether the device helps users create a mental map of the environment. Other researchers also investigated sound and proposed SoundBlender, a framework that enables users to differentiate sounds in the environment by manipulating features of the audio stream (Chang et al., 2023)

People who have low vision also struggle to view VEs. Zhao et al. (2019) created a toolkit with 14 visual enhancements, including brightness, magnification, and bifocal and contrast lenses, that made VEs more accessible to people with low vision. Participants were able to complete tasks

¹ <https://xraccess.org/>

faster and with greater enjoyment when using the enhancements, demonstrating the prototype's feasibility as an assistive technology.

Social VR is another context in which accessibility improvements are desirable. Researchers observed how people with vision impairments interacted with a virtual guide when navigating through an environment. The design was inspired by interactions with an actual guide, where a person would hold onto the guide while navigating and pose questions about their surroundings. Participants found the guide to be useful and felt comfortable receiving guidance (Collins et al., 2023).

5.2.2 Hearing Impairments

Researchers have also created VR and augmented reality (AR) enhancements to improve accessibility for deaf and hard of hearing (DHH) individuals. Jain et al. (2021) created a taxonomy of sounds and discussed how the taxonomy dimensions (i.e., sound source and intent) and sound characteristics (i.e., identity, spatial location, loudness, pitch, persistence, and priority) can help translate audio cues into visual and haptic cues. Mathew et al. (2022) found that Access on Demand, an AR application that provides real-time captioning and interprets American Sign Language, enables users to look at the speaker while reading the captions. However, there were several usability issues with the app, such as text being positioned over the speaker (Mathew et al., 2022).

5.2.3 Physical Impairments

VR shares many visual and auditory aspects with other technology paradigms, which hypothetically makes it straightforward to port assistive solutions that improve visual and auditory accessibility from other technology paradigms to VR. However, in an attempt to convince the user that they are physically existing in a virtual space, VR uses a 3D interaction paradigm that does

not exist for touch-based or desktop computing paradigms. As a result, initial exploratory research has been critical for understanding physical accessibility challenges for people who use wheelchairs and people who have upper-body impairments.

5.2.4 Physical Impairments: Identifying Accessibility Problems

People with limited mobility identified several accessibility challenges using a commercial VR system including putting on and taking off the headset, using the buttons, and maintaining a view of their virtual controllers while wearing the headset (Mott et al. 2020). Baker et al. (2020) performed an exploratory study of VR accessibility, which focused on the potential for VR to be used in long-term care facilities. Because all participants in the study used a wheelchair, a significant challenge was reaching virtual objects around the wheelchair without bumping into the armrests.

A survey by the Disability Visibility Project (Wong et al., 2018) found that individuals with self-reported disabilities who had used VR experienced several challenges. One challenge was that the motion tracking system assumed certain abilities (e.g., the ability to stand or use two hands). Another issue was that some respondents were unable to use motion as input because of differences in dexterity, reflexes, and range of motion. Respondents who used wheelchairs had difficulty moving or rotating their wheelchairs to the position expected by the VR system. To overcome these challenges, some respondents tried to use other hardware, such as game pads, however, most VR systems were “locked” and did not allow the use of other input devices.

Reinforcing these findings, Gerling and Spiel (2021) identified ableist assumptions in the design of VR systems and concluded that the design does not accommodate bodies that do not adhere to an ideal standard. They suggested that designers should consider the “actions lent”, or the actions

a user must employ, to use a VR app's interaction paradigm and hardware. In this way, VR designers and researchers can identify physically inaccessible actions (Gerling & Spiel, 2021).

5.2.5 Physical Impairments: Improving Accessibility for People with Ambulatory Impairments

As a result of the above findings, researchers have worked to create systems and techniques that improve the accessibility of VR for people with wheelchairs. WalkinVR² is a commercially available game driver that allows users to remap the relationship between their physical and virtual controllers so that they are able to reach and crouch comfortably.

Gerling et al. (2020) explored the accessibility of VR for people in wheelchairs by building three prototypes of full-body VR games. They highlighted the importance of customizability to accommodate the diversity of types and severities of impairments. Sassi et al. (2023) experimented with redirected wheelchair movement techniques and found that participants in wheelchairs did not notice large redirection gains, suggesting redirected movement has potential as an accessible locomotion technique.

A common accessibility problem for wheelchair users is that their virtual height can hinder their ability to perform tasks that require a higher vantage point. Taheri et al. (2024) investigated the experience of a lifelong wheelchair user when she was virtually walking in VR and noted the benefit of a higher viewpoint. To address this issue, Qorbani et al. (2022) developed a feature that allowed wheelchair users to adjust their elevation in virtual space and found that they were able to complete tasks, such as locating objects, with comparable success to non-wheelchair users. Ganapathi et al. (2022) created a similar feature, which they called 'Floating Locomotion'. Results

² <https://www.walkinvrdriver.com/>

revealed that seated users whose virtual camera elevation was at standing height were able to pick up objects with similar success as participants who were physically standing.

Researchers have also investigated accessibility features for people who have lower body impairments and do not use wheelchairs. For example, Mahmud et al. (2022) demonstrated that participants with multiple sclerosis were able to walk faster with spatial audio feedback when performing real locomotion in VR than they were without audio.

5.2.6 Physical Impairments: Improving Accessibility for People with Upper-Body Impairments

In addition to designing accessible solutions for wheelchair users and people with lower-body impairments, researchers have worked to improve VR accessibility for people with upper-body impairments. Yamagami et al. (2022) developed a design space that helps designers translate bimanual into unimanual interactions for people who only have use of one hand. I (2024) devised a taxonomy to help designers and researchers choose scene-viewing techniques that do not require head movement based on the characteristics of their VE (Chapter 3). I (2021) also developed a framework to guide the creation of accessible point-of-interest techniques, with which users can change their view to points-of-interest in the VE using various controller interactions (Chapter 4).

5.2.7 Physical Impairments: Improving Accessibility for People Experiencing Situational Impairments

Accessible techniques for VR not only benefit people with disabilities but also people experiencing situational impairments. Situational impairments are functional limitations caused by the environment (e.g., low lighting, loud noises, etc.) or an individual's mental, emotional, or physical state (e.g., divided attention, fear, inebriation, etc.) (Sears et al., 2003; Wobbrock, 2019). Since situations like these can hinder technology use, researchers have developed techniques to eliminate or reduce their impact. Most work in VR that addresses physical situational impairments

has focused on locomotion for users who are seated (Funk et al., 2019; Kim et al., 2008; Kitson et al., 2017; Yan et al., 2016) or have limited tracking space (Iwata et al., 2005, 2006; Steinicke et al., 2010; Templeman et al., 1999). Williamson et al. (2019) also investigated the effect of social situational impairments caused by using VR on a plane, finding that reduced movement would make VR use more acceptable in this context.

5.3 Virtual Reality Design Frameworks and Taxonomies

Despite some progress in VR accessibility, VR is still inaccessible for many people. An approach to selecting and designing accessible interaction techniques is through the use of design frameworks. Design frameworks and taxonomies are useful tools within HCI for designing new interaction methods (Surale et al., 2019), categorizing existing and emerging technologies (Mackinlay et al., 1990), and synthesizing knowledge across research areas (Brudy et al., 2019). Many taxonomies that are focused on VR and 3D applications have examined locomotion (Al Zayer et al., 2020), object manipulation (Jerald, 2018), and scene-viewing techniques (Elmqvist & Tsigas, 2008), and are organized by feature and functionality.

In addition to categorizing existing interaction techniques, researchers have also used VR-focused design frameworks to invent novel techniques. Surale et al. (2019) investigated how tablets can be used to perform complex tasks in VR by constructing a design space with dimensions that explored tradeoffs between mid-air gestures and tablet-centric input. Hirzle et al. (2019) proposed a design framework for gaze-based interactions with headsets and proposed two applications to support eye muscle training.

5.3.1 Using Virtual Reality Design Frameworks and Taxonomies to Improve Accessibility

Organizing existing techniques using a taxonomy can guide the researcher when applying them to a new use case, such as selecting a technique to create an accessible experience. Although not originally intended for VR, Elmqvist and Tsigas (2008) developed a taxonomy to organize scene-viewing techniques designed for 3D environments to help users see around occluding objects, which can also be used to address the accessibility problem of identifying scene-viewing techniques that do not require head or body movement. The authors identified five overarching clusters of scene-viewing techniques including techniques that provide multiple views of objects, tour planners that show the highlights of an environment, virtual x-ray techniques that make parts of the scene transparent, projection distorters that adjust the camera's perspective, and volumetric probes, which show exploded views of objects.

VR-focused taxonomies can also be used to select techniques that could be used by someone with a physical impairment. For example, several locomotion taxonomies exist that group techniques using various attributes including control mechanism, realism, and the metaphor used (Al Zayer et al., 2020; Boletsis & Cedergren, 2019; McMahan et al., 2014; Slater & Usoh, 1994; Ware & Osborne, 1990). Some of these have aggregated locomotion techniques by the level of physical exertion and part of the body used, which can point to a technique's accessibility (Slater & Usoh, 1994; Wendt, 2010). Nilsson et al. (2016) grouped techniques in two categories, stationary and mobile. Although some of the stationary techniques they identified still required physical locomotion (e.g., omni-directional treadmills), others could be performed while seated or standing (e.g., hand-based manipulation of the virtual world).

Di Luca et al. (2021) devised a taxonomy that grouped techniques by the motor ability required to use the technique. They found that movement-based and room-scale techniques were generally

less accessible and more realistic than teleportation or techniques that relied on controller input. However, with greater accessibility also came a trade-off in comfort; more accessible techniques usually induced greater simulator sickness because of the disparity between physical and virtual movement.

Design spaces and taxonomies can help researchers select and invent interaction techniques for novel contexts. In the next two sections, I provide an overview of the scene-viewing and locomotion interaction technique design spaces to illuminate the design patterns employed by existing techniques. These techniques are probably accessible for people with mobility impairments because of their design attributes but have not yet been evaluated in this context.

5.4 An Overview of Scene-Viewing Techniques

Typically, scene-viewing techniques are associated with camera-control techniques, where the user orients the camera with their head or controllers. However, there are many other types of techniques that facilitate viewing scenes in VR, which I describe below.

5.4.1 Camera-Control Techniques

Camera-control techniques manipulate the orientation and position of a virtual camera. They can be grouped based on the level of user control they afford. Automatic camera control is achieved through planned camera paths (Andujar et al., 2004; Chittaro et al., 2003; Desroche et al., 2007; Elmqvist et al., 2007; Sokolov et al., 2006) or individually generated viewpoints (Drucker, 1994; Mühler et al., 2007) in a VE. There are also techniques that allow partial user control. For example, users can define viewing constraints for generating viewpoints and animations of environment highlights (Bell et al., 2001; Burtnyk et al., 2002; Gleicher & Witkin, 1992a).

5.4.2 Scene Modification Techniques

Scene modification techniques introduce objects or alter object properties to expose, highlight, or reveal semantically meaningful information in a scene (Lange et al., 2020; Sukan et al., 2014). Researchers have investigated placing text labels in the environment, (Preim et al., 1995, 1997, 1998), enlarging landmarks to help individuals orient themselves in large VEs (2004) exploding object views (Sonnet et al., 2004), distorting terrain to make routes visible (Takahashi et al., 2006), deforming layers in volumetric objects to expose interior parts (McGuffin et al., 2003), changing the transparency of objects, and cutting away parts of objects to reveal occluded components (Feiner & Seligmann, 1992).

5.4.3 Social Cue Augmentation

Social VR experiences are growing in popularity (Tanenbaum et al., 2020) and with this trend comes the challenge of perceiving and understanding non-verbal communication from other avatars, which can be facilitated by social cue augmentation techniques. Researchers identified ways in which avatars are designed to communicate with facial expressions, gestures, posture, and poses to make social cues more understandable and appropriate for other users (Mayer et al., 2020). Others have investigated techniques that facilitate collaboration in VEs (Chenechal et al., 2016; Fleury et al., 2010; Langbehn et al., 2016; Piumsomboon et al., 2019; Zhang & Furnas, 2005).

5.4.4 Multiple View Techniques

Additional views are typically overlaid on a primary view of a VE with multiple view techniques. For example, camera views overlaid on the primary camera view enables exploration of a virtual world from multiple viewpoints (Bares & Lester, 1998; Stoev & Schmalstieg, 2002; Veas et al., 2012). Magic lenses, which are superimposed views that employ a lens metaphor, have been applied to virtual objects to reveal occluded interior components, such as the skeleton of a

hand (Kister, 2018; Kister et al., 2014, 2015; Looser et al., 2004; Luo et al., 2021; Viega et al., 1996).

5.4.5 Field-of-view (FOV) Extension Techniques

Field-of-view (FOV) extension techniques alter the physical hardware of the headset to simulate peripheral vision (Gruenefeld, Stratmann, et al., 2018). Xiao and Benko (2016) used arrays of LEDs inside a headset to replicate the VE colors outside of the user's FOV. Instead of LEDs, Rakkolainen et al. (2017) used smartphones on either side of the user's face to show more of the VE. In addition to adding devices and lights to headsets, camera hardware modifications, such as the use of Fresnel lenses and 360° cameras, have been explored as FOV extension techniques (Ardouin et al., 2012; Yamada & Manabe, 2016).

5.4.6 Multiscale Techniques

Multiscale techniques enable a user to view an environment at different scales. These techniques are specifically made for VEs that offer a range of detail. Challenges involving multiscale techniques are camera speed and control during navigation and zooming (Argelaguet & Maignant, 2016; Fukatsu, 1998; Kopper et al., 2006; Menzner et al., 2020; Piumsomboon et al., 2018; Trindade & Raposo, 2011). Without adjusting camera properties during zooming and navigation, the user could experience double vision and decreased depth perception in VR (Argelaguet & Maignant, 2016). For example, GiANT (Argelaguet & Maignant, 2016), a technique for viewing environments with stereoscopic immersive displays, adjusts the scale factor and speed of the camera based on the user's perceived navigation speed.

5.4.7 Projection Distortion Techniques

Projection distortion techniques integrate multiple camera views to support discovery and artistic expression in VEs (Elmqvist & Tsigas, 2007). For example, Singh (2002) proposed a new

interactive camera model that allows users to create nonlinear perspectives of virtual objects as a form of artistic exploration. Researchers have also investigated projection distortion as a means of overcoming occlusions (Carpendale et al., 1996; Cui et al., 2010). For example, Cui et al. (2010) proposed a technique that bends the camera ray around occluding objects, so that target objects are in the line of sight. Although projection distortion techniques allow users to see more of a target and overcome occlusions, the resulting view can be confusing because it deviates from how people usually perceive 3D objects.

5.4.8 Cue-Based POI Techniques

Cue-based point-of-interest (POI) techniques provide cues about the location and direction of out-of-view POIs. Initially developed for desktop and mobile platforms (Baudisch & Rosenboltz, 2003; Gustafson et al., 2008), cue-based techniques have also been successfully implemented in CAVE systems, AR, and VR (Gruenefeld et al., 2018; Gruenefeld et al., 2017; Petford et al., 2019). For example, classic techniques such as Halo (Baudisch & Rosenboltz, 2003) and Wedge (Gustafson et al., 2008) have been used to support discovery in augmented and VEs (Gruenefeld et al., 2018; Gruenefeld et al., 2017).

The POI techniques discussed above convey distance and direction information, which helps in spatial orientation tasks. However, few POI techniques convey the identity of off-screen POIs, which can aid in exploring POIs. Prior work has suggested that users prefer knowing the identity of POIs when deciding how to orient the camera (Gruenefeld, von Barga, et al., 2018; Lin et al., 2017; Siu & Herskovic, 2013). A novel cue-based technique, Outside-In (Lin et al., 2017), uses picture-in-picture displays to show the identities out-of-view POIs in 360° videos. The orientations and positions of the superimposed displays provide distance and direction cues for the POIs.

Although the POI techniques described above exist to help visualize out-of-view objects, these techniques do not help the user select and orient to a POI. Once a user is aware of an out-of-view POI, the next step is to orient to it so that it is within view. Mobile and PC-based selection and orientation techniques, such as Hopping (Irani et al., 2006) and EdgeSplit (Hossain et al., 2012), allow users to select POI proxies to orient to. Unlike these techniques, existing AR and VR techniques lack control or affordances and instead rely on head movements (Gruenefeld, El Ali, et al., 2018; Gruenefeld, Ennenga, et al., 2017; Lin et al., 2017; Sargunam et al., 2017) or button presses (Pavel et al., 2017) to orient the user.

Within the accessibility community, POI based techniques have focused on identifying sounds of interest for deaf and hard of hearing (DHH) individuals (Ho-Ching et al., 2003; Kaneko et al., 2013; Z. Li et al., 2022, 2023; Matthews et al., 2005). Research suggests that DHH users want to know the source and location of peripheral sounds in their environments (Jain et al., 2019; Matthews et al., 2006). For instance, AR users want to know the sound source and type (e.g., speech vs. non-speech sounds) while having a flexible system that accommodates their preferences for directional granularity and screen layout (Jain et al., 2015).

Prior work has shown that visualizing and orienting to out-of-view POIs are important for VR users. From an accessibility perspective, it is also important to provide users with affordances and controls over the experience (Mott et al., 2020). Nearmi identifies the basic components of cue-based POI techniques in VR and suggests how these components can be used to make the techniques accessible to people with limited head or trunk movement (Franz et al., 2021).

5.5 An Overview of Locomotion Techniques

Like the scene-viewing technique design space, the design space of locomotion techniques is also rich. The simplest and most intuitive method to move in a VE is room-scale walking, in which

the user walks around the physical space to move in a virtual space (Usoh et al., 1999). However, physical constraints such as the size of the room have resulted in researchers and designers inventing locomotion techniques that do not require a one-to-one mapping of physical to virtual locomotion (Rietzler et al., 2020). In addition, walking techniques can require high physical exertion to perform. Designers have focused on controller-based locomotion techniques as well as movement-based techniques that are stationary to address these problems. These techniques also have the potential to be used as accessible alternatives to walking-based techniques. Below I provide an overview of the locomotion technique design space (excluding full-gait walking techniques) based on Al Zayer et al.'s (2020) groupings.

5.5.1 Walking-in-place Techniques

Walking-in-place techniques use the repetitive motion of different parts of the body to control virtual translation while the user stays in place. Researchers have examined use of the lower body to control virtual movement by wiping (Nilsson et al., 2013), stepping (Guy et al., 2015), and tapping with the feet (Nilsson et al., 2013; Zielasko et al., 2016), bending the knees (Guy et al., 2015), and raising and lowering the thighs (Zielasko et al., 2016).

Use of the upper body has also been investigated including finger walking (Kim et al., 2008, 2009), swinging the arms (McCullough et al., 2015), alternately pressing controller buttons using two hands (Sarupuri et al., 2017), and moving the head (Terziman et al., 2010).

5.5.2 Steering-Based Techniques

Steering-based techniques use parts of the body or physical props to control the direction of virtual movement (Brogan et al., 1998; Brooks, 1999). Gaze-directed techniques use the rotation of the head to determine the direction of movement such that the user moves in the direction they are looking (Mine, 1995; Stellmach & Dachsel, 2012). Hand-directed steering techniques are

similar but use the direction of the hands to determine the direction of movement. Other techniques use metaphors such as flying (Boudoin et al., 2008; Bowman et al., 2001; Fairchild et al., 1993; Robinett & Holloway, 1992), skiing (Haydar et al., 2009), or driving a tank (Haydar et al., 2009), while others use novel hand gestures (Cabral et al., 2005; Cardoso, 2016; Cirio et al., 2009, 2012; Haydar et al., 2009; Papadopoulos et al., 2012; Zhang et al., 2017). Lean-directed steering techniques use the position of the user's head as they are leaning to determine direction (Fairchild et al., 1993; Guy et al., 2015; LaViola Jr. et al., 2001). These techniques have been implemented using a variety of devices in addition to mobile and VR setups including a Wii (de Haan et al., 2008; Valkov et al., 2010; Wang & Lindeman, 2011), a touchpad (Wang & Lindeman, 2012), a pressure mat (Carrozzino et al., 2014), and chairs (Beckhaus et al., 2007; Kitson et al., 2015).

5.5.3 Selection-Based Techniques

Selection-based techniques allow the user to select a target, which can either be a position (Bowman et al., 1997, 2001) or object (Brooks, 1999) in the environment, to travel to. One of the most commonly implemented selection-based techniques is Teleportation (Bhandari et al., 2018; Bozgeyikli et al., 2016; Mackinlay et al., 1990; Schroeder et al., 2001). Selection is usually achieved by pointing a ray from a handheld device at the desired position or object and pressing a button to confirm the selection. Other selection-based techniques enable the user to select their desired position on a map of the environment (Bowman et al., 2001) or jump in place while looking at a target to select it (Bolte et al., 2011).

5.5.4 Manipulation-Based Techniques

Manipulation-based techniques use gestures to control the user's viewpoint or environment. These techniques include grabbing and pulling the earth (Coomer et al., 2018), climbing (Slater et al., 1995), hand-walking (Uthe, 2014), and world-in-miniature (WIM), in which the user

manipulates their position in a small model of the environment (Ropinski et al., 2005; Stoakley et al., 1995; Valkov et al., 2010; Wingrave et al., 2006). Researchers have compared different metaphors for manipulating a user's viewpoint (Ware & Osborne, 1990). With eyeball-in-hand, the user positions a virtual camera using a handheld controller. With scene-in-hand, they position the scene relative to the camera. The authors found that the two techniques are optimal for different environments. Scene-in-hand was better for viewing a cube floating in space than navigating through a maze, while participants were divided on which environments eyeball-in-hand was better suited for.

The work reviewed in this section offers some guidance for choosing locomotion techniques based on users' desired levels of physical exertion, and whether the techniques are seated or stationary, which can all be indicators of accessibility. Yet, this guidance is derived mostly from anecdotal evidence, requiring formal empirical investigation into the accessibility of locomotion techniques. Despite the lack of research investigating accessibility, established approaches exist for evaluating VR interaction techniques on performance and user experience factors.

5.6 Evaluating Virtual Reality Interaction Techniques

A variety of factors contribute to the user experience of VR locomotion techniques. Such techniques are often evaluated not only in terms of performance (i.e., speed and accuracy) (Bowman et al., 1997; Christou et al., 2016; Christou & Aristidou, 2017; Ruddle & Lessels, 2009; Suma et al., 2010; Suma et al., 2010) but also by the level of presence (Bhandari et al., 2017; Muhammad et al., 2017; Nilsson et al., 2013; Pfeiffer et al., 2016; Terziman et al., 2010; Usoh et al., 1999; Zambaka et al., 2004), which refers to the extent a technique convinces the user that they are physically existing in the VE (Slater et al., 1994), and level of simulator sickness (Suma et al., 2010; Suma et al., 2009, 2010) they induce. Other important factors include naturalness

(Muhammad et al., 2016; Nilsson, et al., 2013; Slater et al., 1995; Tregillus & Folmer, 2016; Usoh et al., 1999), which refers to a technique's resemblance of real-world interactions, and spatial awareness (Bowman et al., 1998; Christou & Aristidou, 2017; Kim et al., 2008; McCullough et al., 2015), which refers to the extent a technique enables a user to keep track of their surroundings.

5.6.1 User Experience and Accessibility of Locomotion Techniques

Some studies indicate an effect of accessibility on the user experience of locomotion techniques. Individuals with ambulatory impairments as a result of multiple sclerosis (MS) felt greater presence in a VE when walking compared to people without MS, suggesting that an impairment might have an effect on presence (Guo & Quarles, 2012). Techniques that preserve spatial awareness help users maintain an understanding of the spatial relationships of objects in the environment relative to themselves. Although they were rated as more accessible compared to other techniques, teleportation techniques scored low for spatial awareness because the user jumped to a new location rather than moving smoothly through the space, causing disorientation (Di Luca et al., 2021). Preliminary findings such as these demonstrate that user experience factors could be affected by accessibility design considerations. Because of the multi-faceted nature of VR interaction technique evaluation, there is a lack of a standard approach to designing studies.

5.6.2 A Lack of a Standardized Approach to Virtual Reality Interaction Technique Evaluation

Researchers (Bergström et al., 2021; Bowman et al., 1999; Lampton et al., 2014) have noted the absence of a standardized evaluation approach for interaction techniques in VR. Bergström et al. (2021) conducted a survey of 39 papers to discover patterns in evaluation methods for object selection and manipulation techniques. They found that one or a combination of three main independent variables were used in most studies: Interaction Technique, Target, and Task. Most

studies only compared two interaction techniques, while some were more complex and varied target shape, size, distance, and arrangements in addition to the interaction technique. Most studies measured Trial Completion Time and Accuracy (in terms of error rate) as the dependent variables. A common feature in these studies was the use of questionnaires to measure user experience including Satisfaction, Usability, Simulator Sickness, Presence, Workload, Fatigue, and Immersion. They proposed a set of guidelines for performing a study to evaluate object selection and manipulation techniques as a step towards standardizing VR evaluation methods.

5.6.3 Virtual Reality Testbeds to Support Interaction Technique Evaluation

Bowman et al. (1999) also noted the need for a standardized evaluation approach for object selection and manipulation techniques in VR. To address this gap, they built a testbed that comprised two phases, a selection phase in which the user selected the correct object, and a manipulation phase in which the user matched the position and orientation of an object to a target object. They measured Speed, Accuracy, and User Comfort while participants used nine different object selection and manipulation techniques. The researchers discussed several advantages of using a testbed, including reusability in other studies, as well as the ability to produce rich datasets and predictive performance models. Most importantly, they found that there was not a single optimal technique out of the nine they tested, but rather there were several optimal techniques depending on the task, environment, system, and user characteristics.

Poupyrev et al. (1997) created a testbed named VRMAT that enabled researchers to evaluate object manipulation techniques through selection, positioning, and orienting tasks. They also created a novel distance measure called the virtual cubit, which was equal to the length of a user's maximum reach.

Cannavò et al. (2021) developed a testbed and a set of task-based metrics to rank locomotion techniques based on performance and user experience factors. The final ranking was scored using a weighted sum model: Users assigned weights to outcome measures (e.g., Speed, Accuracy, Mental Effort, and Presence, among others) to indicate their relative importance. This approach took into consideration that an optimal technique depended on user preferences in addition to performance. The environments and tasks in the testbed were predefined but enabled researchers to port their own locomotion techniques. Tasks were derived from the literature on locomotion technique evaluation, and the most common tasks were implemented (e.g., traveling in a straight line and avoiding moving objects, among others).

Because studies suggest optimal locomotion and object-selection or manipulation techniques might not exist (Bowman et al., 2001; Cannavò et al., 2021), personalization could be a better approach to VR interaction design. To achieve personalization though, particularly for improving performance with interaction techniques, the first step is to differentiate users based on their movement characteristics.

5.7 Measuring the Movement Characteristics of People with Physical Impairments

Understanding how people with motor impairments move in 2D and 3D space has been a goal for many researchers. Characterizing movement is achieved through the design of metrics that capture differences between individuals with different impairments and people without impairments.

Mackenzie et al. (2001) proposed a set of seven accuracy measures to explain performance differences in terms of Speed and Accuracy of input devices for people without impairments. They found that two measures, ‘movement offset’ (how far the user’s path deviated from the optimal path) and ‘target re-entry’ (how many times a participant exited then entered a target) were

correlated with Speed and Accuracy. Based on this study, Keates et al. (2002) proposed a set of measures to characterize the movement patterns of people with motor impairments. They found that path metrics that measure the distance travelled relative to the cursor displacement and the distribution of distance travelled for a range of cursor speeds could reveal differences between participants with and without impairments as well as differences between participants with more and less impairment severity (Keates et al., 2002).

Hurst et al. (2008) predicted whether people with motor impairments would need an accessibility adaptation when using a mouse. They built machine learning models that achieved 83% accuracy when classifying whether a user could benefit from Steady Clicks, an accessibility feature that kept the cursor stationary while the user was clicking, preventing it from slipping off the target. They used feature selection to identify the best features for building a model and found that click-related features and task completion time were most predictive. In another study in which they built models to distinguish people with and without motor impairments, they found that the task completion time and number of clicks, as well as target re-entry were the most important features when predicting if a user had an impairment.

Instead of overall movement characteristics, Hwang et al. (2004) investigated how submovements, or movements between velocity dips, could reveal differences between people with and without motor impairments. They proposed a set of accuracy measures including ‘number of pauses’ and ‘peak velocity of submovements’, which provided insight into differences based on impairment type as well as differences between people with and without impairments.

Aside from cursor-related measures, touch-related measures have also been proposed. A study by Kong et al. (2022) proposed 15 touch metrics to characterize touch presses. They found that five metrics differentiated people with and without impairments including ‘touch variability’, or

the total distance the finger travelled during a touch, and ‘touch drift’, or the Euclidean distance between the touch press and touch lift, among other measures (Kong et al., 2022).

While the above work compares movement in two dimensions, some researchers have also investigated movement characteristics in three dimensions. A study by Viau et al. (2004) found that people with motor impairments were significantly slower when placing a ball on a virtual target compared to people without motor impairments. Also, the path their hand took was more curved compared to that of people without impairments. However, the authors noted that overall, the movement strategies for both groups were similar (Viau et al., 2004).

5.8 Modelling User Behavior to Personalize Virtual Reality Interaction

Once researchers can measure the movement characteristics of people with impairments, they can use these measures to inform models that predict not only whether a user has an impairment, but also if they can benefit from an adapted interface or interaction technique. Modelling user behavior has made technologies more accessible, usable, and secure.

5.8.1 Performance Modelling for Interface Optimization

User models have been developed to improve touch accuracy and interface layouts on touchscreen devices for people with motor impairments (Montague et al., 2012; Mott et al., 2016; Mott & Wobbrock, 2019; Sarcar et al., 2018). Mott et al. (2016) modeled touch input by participants with motor impairments using template matching to predict participants’ intended touch points, thereby reducing errors. In another study, Mott and Wobbrock (2019) found that touch accuracy could be improved by predicting a touch point using a combination of individual and independent user models. Similarly, Montague et al. (2012) employed shared user models to adjust target sizes and touch duration, finding that these adaptations significantly reduced errors for people with impairments. Gajos et al. (2010) also took the approach of adapting interface

elements to user abilities and found that users with motor impairments were more efficient and accurate when using an interface adapted to their abilities, which typically contained larger targets compared to the default interface.

5.8.2 Performance Modelling in Virtual Reality

Because VR and AR are relatively new technologies, little research has investigated modeling interactions in this context. Cabric et al. (2021) extended the keystroke-level model (KLM) to a mixed reality context and modelled common gestures including controller button clicks, hand raising, mid-air tapping, and head pointing. Erazo and Pino (2015) also used a KLM to predict users' performance time while using touchless hand gestures. These two models predicted performance with 3D gestures but did not factor in the impact of a user's abilities, which could potentially make this technology more accessible.

5.8.3 Authentication in Virtual Reality

Although more work has yet to be done on performance modelling in VR, researchers have measured and characterized movement in VR for the purpose of authenticating users. Studies have shown that authentication can happen in real time (Miller et al., 2019), across different VR systems (e.g., Oculus Quest, HTC Vive, and HTC Vive Cosmos) (Miller et al., 2020), with only head movement (Sivasamy et al., 2020), and by performing different tasks (Kupin et al., 2019; Pfeuffer et al., 2019). Pfeuffer et al. (2019) proposed a set of measures to capture individual differences in movement based on body relation theories. They found that measures of the relationship between device pairs (e.g., distance between headset and right controller) were the most important features in models that identified users. Although not related to detecting movement differences between people with and without impairments, features that predict individual differences could help personalize VR experiences for people with a range of abilities.

5.9 Summary of Related Work

VR accessibility research for people with visual, hearing, and ambulatory impairments has made great advances. However, there is still a gap in the literature making VR accessible for people with upper-body impairments, particularly for those who find it challenging to move their heads, trunks, arms, and hands. Although many scene-viewing and locomotion techniques exist that use controllers instead of body movement, their potential for enhancing accessibility has not been thoroughly explored or evaluated. Part of the reason for this gap is the lack of standardized approaches for evaluating VR interaction techniques for people without impairments, let alone people with impairments. Testbeds and measures that characterize the movement of users with impairments can be useful tools for evaluating accessible interaction techniques and can assist in creating predictive models to adapt VR interaction to users' abilities.

6 A VIRTUAL REALITY SCENE TAXONOMY: IDENTIFYING AND DESIGNING ACCESSIBLE SCENE-VIEWING TECHNIQUES³

Scenes in virtual reality (VR) environments are designed to resemble physical environments, and, as a result, considerable effort goes into making them believable. The richness of scenes and their lack of physical constraints has resulted in the creation of various scene-viewing techniques (i.e., interaction techniques that facilitate the visual perception of a virtual scene) over the past few decades (Elmqvist & Tsigas, 2008; Mayer et al., 2020). These techniques have ranged from camera control techniques to techniques that modify the visual properties of a scene.

However, despite the abundance of scene-viewing technique research, the question remains: how would a designer choose a scene-viewing technique to implement? and more generally, how would a designer reason about scene-viewing techniques given a particular scene? Despite the significant number of available scene-viewing techniques, there is no guidance for determining if a technique is suitable for viewing a scene, particularly when the structure of a scene affects how it can be viewed. For example, a scene that has high object density would lend itself to a variety of techniques, such as techniques that make occluding objects transparent, techniques that distort the camera ray to see around occluding objects, and techniques that enable a user to choose a path through the scene that gives them the best view, but a designer would have to sift through the research to identify and compare these techniques.

To facilitate the selection of appropriate scene-viewing techniques, I approached scenes as spaces with affordances that indicate how they should be viewed. Based on this approach, I devised a scene taxonomy derived from insights in cognitive psychology and computer vision on how people and computers define, view, and describe virtual and physical scenes. In addition, to ground

³ Portions of this chapter were published as Franz, R. L., Junuzovic, S., & Mott, M. (2024). A virtual reality scene taxonomy: Identifying and designing accessible scene-viewing techniques. *ACM Transactions on Computer-Human Interaction*, 31(2), 1–44.

the taxonomy in existing implementations of VR environments, I surveyed 29 commercial VR applications to identify common visual properties and tasks.

Choosing an appropriate scene-viewing technique could be particularly important when designers want to make VR applications accessible. There are many scene-viewing techniques that require little head movement and could be used as alternatives to movement-based techniques, but again, because there are so many techniques, it is difficult to know which to use for a specific context. To address this problem, I applied the scene taxonomy to identify and classify scene-viewing techniques that could be used by people with disabilities or situational impairments that affect their head movement.

In addition to identifying existing scene-viewing techniques that are accessible, I showed how the taxonomy could be used as an ideation tool to design new scene-viewing techniques. I designed and evaluated three scene-viewing techniques with 16 participants temporarily experiencing limited head movement. The findings suggested that although participants thought the new scene-viewing techniques were as easy or easier to use than the default technique implemented in many VR applications, they considered trade-offs related to accessibility, usability, realism, spatial awareness, comfort, familiarity with the interaction, and task objective when determining their preferred technique. As a result of these trade-offs, the technique they thought was easier to use was not always the one they preferred.

6.1 Devising the Scene Taxonomy

To devise the scene taxonomy, I reviewed literature in cognitive psychology and computer vision to understand how people and computers view scenes. First, I present the theoretical underpinning of the taxonomy, which is the concept of affordances. Then I introduce the taxonomy and discuss the literature that the taxonomy is grounded in.

6.1.1 Theory: Affordances

Real-life locations afford different ways of viewing. Humans design houses and buildings based on achieving particular views. For example, tourist destinations are engineered to enable the best lookouts and photo opportunities. Scenes in virtual environments (VEs) also afford a variety of viewing opportunities. I use the concept of affordances to identify how VEs support different views and scene-viewing techniques.

Gibson, a psychologist, proposed that people perceive affordances directly, and that everything in our world is perceived as an affordance (Gibson, 1979). He argued that instead of perceiving qualities or properties of scenes and the objects in them, we immediately perceive the relationship between the observer and the observed (Gibson, 1979). In other words, humans immediately perceive how to interact with an object. For example, the ground is a firm, flat surface that can bear human body weight, so it affords walking. Humans do not need to classify objects based on their characteristics to perceive how they could be interacted with; they do not need to know that the ground can be classified as a hard, flat surface in order to know that it is walkable (Gibson, 1979).

Norman brought the concept of affordances to HCI and defined it in the context of technology design (Norman, 2013). His definition extended Gibson's concept to differentiate between affordances and perceived affordances. He argued that everything affords some interaction. A perceived affordance signals a particular interaction with a particular outcome (Norman, 1999, 2018). So, a user perceives the meaning of a button, which is that it can be pushed to initiate some action. If a perceived affordance is designed well, somebody with little to no experience with the object or interface would know how to interact with it (Norman, 2013). In Norman's words, "When you first see something you have never seen before, how do you know what to do? The answer, I decided, was that the required information was in the world: the appearance of the device could

provide the critical clues required for its proper operation” (Norman, 2013). If a scene in VR is perceived as a space with affordances that suggest how it can be viewed, how would we describe these affordances? I developed a scene taxonomy that breaks down a scene according to its visual properties to address this question.

6.1.2 Literature Review Approach

In this section I discuss how I derived the taxonomy based on literature in two disciplines as well as a review of VR applications.

Using the lens of affordances, scenes afford specific ways of being viewed. In other words, the optimal technique to view a scene might depend on its components and structure. Therefore, I reviewed literature in cognitive psychology and computer vision and popular VR applications to identify ways of describing the structure and main components of a scene. I approached the review wanting to answer two research questions: (1) What is the definition of a scene? And (2) How are scenes described?

The objective of the literature review was to synthesize research across cognitive psychology and computer vision and ground a taxonomy in work that is relevant to scene-viewing techniques. I identified overarching patterns with regards to how scenes are defined and described. The breadth of the review, which spanned two disciplines, and the objective to demonstrate an understanding of relevant work, made a traditional literature review appropriate (Valentine et al., 2009).

Review of cognitive psychology and computer vision literature: I started the literature review with Tversky and Hemenway’s “Categories of Environmental Scenes” (1983), which was the earliest scene taxonomy I could find that was relevant to the research questions. I then searched for more recent scene taxonomies and work related to scene properties in the 442 works that cited

this paper that were published between 1984-2020. I found that cognitive psychology and computer vision papers cited this taxonomy and so I reviewed work in both fields.

Review of VR applications: In July 2019 and August 2020, I reviewed 29 of the most popular VR applications on the HTC Vive App Store⁴ and on the Oculus Rift App Store⁵. I examined the most popular applications so I could explore types of VEs many people are commonly exposed to. If the app was free, I used it myself; otherwise, I viewed YouTube videos of other people using the app. Viewing applications on YouTube⁶ was sometimes preferable to experiencing the app first-hand because parts of the app could only be reached with hours of use. While reviewing VR applications, I focused on the VE's visual properties, and the tasks users were performing in them.

Based on the review of literature in cognitive psychology, computer vision, and the VR applications I examined, I identified properties that describe the visual structure of a scene and the main task.

6.2 The Scene Taxonomy

This taxonomy aims to categorize scene-viewing techniques based on affordances of a virtual scene, namely visual properties and tasks. Appendix A.1 demonstrates how the taxonomy was used to organize scene-viewing techniques. Prior research findings support this approach: researchers found that some scene-viewing techniques are more effective than others depending on the environment structure (Pierce & Pausch, 2004; Ware & Osborne, 1990). Also, according to research in cognitive psychology, people look at scenes in different ways according to the task they are given (Yarbus, 1968). These findings suggest that scene-viewing techniques could be useful for particular environments and tasks. As such, the scene taxonomy has two dimensions:

⁴ <https://www.viveport.com/>

⁵ <https://www.oculus.com/experiences/rift/>

⁶ <https://www.youtube.com/>

visual properties and task types. I present an overview of the taxonomy and discuss the research it is grounded in later.

Table 1. The visual properties of the scene taxonomy.

Visual properties	
Openness	Indoor
	Outdoor
	Abstract
Scale	Smaller
	Human scale
	Larger
Area	Small
	Medium
	Large
Object density	Low
	Moderate
	High
Object tracking	None
	One
	Multiple
Scene transitions	Infrequent
	Frequent
Contains social actors	Yes
	No

6.2.1 Dimension and Property Definitions

I take the approach of breaking down a scene according to its visual properties and tasks. This approach gives designers and researchers a standard language for comparing scenes. Comparing scenes would be challenging if I took other approaches such as describing a scene with a category (e.g., beach), by the names of objects in the scene (e.g., sofa) or by the physical interactions they afford (e.g., walkable) because of the numerous possible descriptors in each category. Also, similarities in the structure of the environment might not be captured with these approaches.

6.2.2 Dimension 1: Visual Properties

In this section I present the visual properties of the scene taxonomy (Table 1).

Openness: The openness of the scene describes whether a scene is enclosed (Figure 1). Indoor scenes are enclosed, such as an office or a cave. Outdoor scenes are open scenes. Examples include a forest or an urban setting. Through the review of VR applications and scene-viewing techniques

I found that many scenes in commercial applications feature a space that is neither indoor nor outdoor. Therefore, I added the value Abstract. Abstract scenes act as a universal space in which 3D objects are designed or viewed (Burtnyk et al., 2002; Grossman et al., 2002).



Figure 1. Openness: (from left to right) outdoor⁷, indoor⁸, abstract⁹.



Figure 2. Scale: (from left to right) smaller¹⁰, larger¹¹.

Scale: The scale of the scene refers to how large or small objects are relative to the size of the user's virtual body (Figure 2). If the scale is human scale, the scene would be at scale relative to the human body. If the user is very large, the objects in the scene would appear small, and the scale of the scene would have the value smaller. Conversely if the user is very small, the scale would have the value larger.

7 Model courtesy of TriForge Assets, <https://assetstore.unity.com/packages/3d/environments/fantasy/fantasy-forest-environment-free-demo-35361>

8 Model courtesy of Omni Studio, <https://assetstore.unity.com/packages/3d/environments/free-medieval-room-131004>

9 Model courtesy of Rhino, <https://www.rhino3d.com/>; Image from: <https://sketchfab.com/blogs/community/preparing-and-uploading-3d-jewelry-designs-in-rhino/>

10 Model courtesy of 255 pixel studios, <https://assetstore.unity.com/packages/3d/environments/urban/city-package-107224>

11 Model courtesy of 255 pixel studios, <https://assetstore.unity.com/packages/3d/environments/urban/city-package-107224>




Figure 3. Area: (from left to right), large¹², small¹³.

Area: The area describes how much area the scene covers (Figure 3). A large scene covers a large area, such as a park. A small scene covers a small area, such as a bathroom. A medium scene is neither large nor small. A scene can be both large and indoor, for example, a warehouse.

Object Density: Object density describes the number of objects in the scene. A forest would be a scene with high object density (Figure 1, left), while an empty warehouse would be a scene with low object density (Figure 3, left). A scene with moderate object density has neither a particularly low nor high number of objects.

Object Tracking: Object tracking describes the quantity of moving objects in the scene. In many VR applications I surveyed, the user must track moving objects (e.g., any app in the offense/defense category). I added object tracking as a property and categorized moving objects in terms of quantity: multiple, one, or none.

¹² Model courtesy of Modular Warehouse, <https://assetstore.unity.com/packages/3d/environments/modular-warehouse-231449>

¹³ Model courtesy of Brick Project Studio, <https://assetstore.unity.com/packages/3d/environments/apartment-kit-124055>



Figure 4. Scene transitions: (from left to right) frequent¹⁴, infrequent¹⁵.

Scene Transitions: Through the review of VR applications, I found that many scenes in commercial applications feature environments where scenes change frequently. Therefore, I added the scene transitions property to refer to how often the users' scenes change (Figure 4). For example, a scene with frequent scene transitions is a large building with multiple hallways and rooms. Users experience a scene transition whenever they enter a new room. An environment with infrequent scene transitions is an underwater landscape where there are no obvious transitions.

Contains Social Actors: This term describes whether the scene contains avatars or characters the user can interact with. Social actors can be avatars controlled by other users or by the app.

6.2.3 Dimension 2: Task Types

Since the way people describe a scene is affected by their task (Yarbus, 1968), I added task type as a second dimension (Table 2) to the scene taxonomy.

¹⁴ Model courtesy of Brick Project Studio, <https://assetstore.unity.com/packages/3d/environments/apartment-kit-124055>

¹⁵ Image courtesy of Llŷr ap Cenydd. Ocean Rift. <https://www.meta.com/experiences/2134272053250863/>

Table 2. The task type dimension of the scene taxonomy breaks down the scene based on the task the user is performing in it.

Tasks
Movement
Socializing and Collaboration
Exploration
Navigation
Offense/Defense
Creativity
Observation
Productivity

Movement: The app's primary goal is to encourage users to move their arms, heads or bodies. For example, the app Dance Central¹⁶ encourages users to try different dance moves.

Socializing and Collaboration: The user's primary goal is to socialize with other users in a virtual space. These VR applications allow users to converse, attend lectures, or play activities together.

Exploration: The primary goal is to actively interact with the scene to discover it without having a specific destination. An example is Ocean Rift¹⁷, in which users explore an underwater sea environment and observe sea animals (Figure 4, right).

Navigation: The primary goal is to travel to a target. An example of navigation in a VR app is Dreadhalls¹⁸, where the objective is to find the exit of a large building with numerous rooms and hallways.

Offense/Defense: These applications' primary goal is to defend against and/or eliminate enemies. An example is Superhot VR¹⁹, where users shoot or punch moving enemies to eliminate them.

¹⁶ <https://www.dancecentral.com/>

¹⁷ <https://www.meta.com/experiences/2134272053250863/>

¹⁸ <http://www.dreadhalls.com/>

¹⁹ <https://superhotgame.com/superhot-vr>

Creativity: The primary purpose of these applications is to enable users to create 3D models, draw, or paint. An example is Kingspray Graffiti²⁰, where the primary purpose is to create murals on virtual walls.

Observation: The primary purpose is to passively view content without interacting with it (e.g., an immersive movie).

Productivity: The objective is to accomplish tasks users would perform on their personal computers. An example is Virtual Desktop²¹, which replaces the 2D monitor with an immersive 3D version of users' desktops.

6.3 Scene Understanding

How we define and describe the environments we occupy is not straightforward, even though we experience scenes every moment of our lives. In this next section, I demonstrate how the scene taxonomy is based on literature in two disciplines: cognitive psychology and computer vision, to investigate how each field defines and describes a scene.

6.3.1 Scene Definitions

Scenes have similar definitions across the two disciplines, which are: object-based, affordance-based, and property-based definitions.

Object-Based definitions: In object-based definitions of scenes, the scene is the background in which objects are located. In other words, "Scenes are a perceptual, spatial generalization of objects; they are the setting or context for objects, the background where objects are figural." (Tversky & Hemenway, 1983) Scenes also contain information about how objects relate to each other and the background, both spatially and semantically (Zelinsky, 2013).

²⁰ <http://infectiousape.com/>

²¹ <https://www.vrdesktop.net/>

An object-based definition of a scene also exists in computer vision literature. Researchers define a scene as a 3D space with 3D objects that have a particular spatial layout and label (Carson et al., 2002; Gupta et al., 2011; Ismail et al., 2018). That is to say, “Virtual scenes are digital 3D representations of natural scenes and contain a set of objects, such as items and avatars that move and interact in 3D space” (Oyekoya et al., 2009).

Affordance-Based Definitions: Scenes can be defined in an affordance-oriented manner as a space that affords human movement, poses, and actions. Xiao et al. (2010) write that a scene is “a place in which a human can act within, or a place to which a human being could navigate”.

Property-Based Definitions: A scene can also be defined based on conceptual and perceptual properties. One definition of a scene, a scene gist, can be thought of as a rough mental sketch of a scene when a person is quickly shown a scene image (Oliva, 2005). It often contains information about the scene’s basic category (e.g., beach), the names and locations of objects, and lower-level features such as color and texture. Humans can understand scene gist within a matter of milliseconds (Thorpe et al., 1996) even when the image is degraded (Schyns & Oliva, 1994).

Drawing from the above definitions, I define a VR scene as a 3D space with 3D objects that a user can interact with (Oyekoya et al., 2009; Tversky & Hemenway, 1983). I add to this definition that a scene is the part of a VE that is immediately visible to the user if they were to view it from any camera angle at their current location. A VE can contain multiple scenes; for example, a VE can contain a house and a forest. The user can be inside the house, which is one scene, or outside in the forest, which is a different scene.

6.3.2 Scene Descriptions

Here I discuss how scenes are described in cognitive psychology and computer vision.

Cognitive Psychology: How people describe scenes reveals information about how scenes are viewed and conceptualized (Li et al., 2007). An individual can look at the same scene in different ways depending on the task they are given (Yarbus, 1968). As a result, a scene can have various descriptions depending on the individual and task. Although viewing patterns can differ significantly, there is evidence to suggest that the human brain processes visual information in a consistent way (Li et al., 2007).

A theory is that two different levels of processing take place during scene perception, they are referred to as perceptual and conceptual processes. Perceptual processes produce perceptual features, which contain purely visual information about a scene such as edges, reflection, texture, etc. (Geisler, 2008; Li et al., 2007). Meanwhile, conceptual processes produce conceptual features, which contain semantic information, such as objects' names and functions (Oliva & Torralba, 2006). Evidence suggests that scene recognition may not be solely dependent on the knowledge of objects in a scene. Instead, individuals can quickly process a scene based on visual features, suggesting that these features alone lead to identification of the scene category and function (Oliva & Torralba, 2001). A theory is that perceptual features are conceptualized based on experience and the probability that they are related to shape and size of common objects in scenes (Geisler, 2008).

Computer Vision: The computer vision research I surveyed focuses on automatically detecting, localizing, recognizing, classifying, and understanding scenes. Scene understanding is a computational approach to the process of perceiving and analyzing a scene. Researchers have used similar mechanisms to those in human scene perception to understand a scene. For example, Oliva and Torralba (2001) showed that perceptual properties can provide a sufficient description of a scene's spatial structure. They call this description the "Spatial Envelope" of a scene and proposed

five quantifiable properties: naturalness, openness, roughness, expansion, and ruggedness (Oliva & Torralba, 2001). They found that by using these features as input, computer vision algorithms could categorize scenes by their perceptual properties rather than knowledge about objects in the scene (Oliva & Torralba, 2001).

Patterson et al. (2014) have also found that adding attributes that people commonly use to describe scenes can improve automatic scene classification. When asked to describe a scene, participants used five attributes: materials (e.g., wood), surface properties (e.g., rusty), affordances (e.g., dining), spatial envelope properties (e.g., openness), and objects (e.g., chair). Some scene attributes are related to their functions. For example, a narrow corridor affords walking through, not sitting, nor laying down. Materials can also suggest related human behaviors; a large body of water might be used for swimming. Finally, embedded objects afford behaviors; chairs and tables can be used for eating. The function of a scene is often present in people's scene descriptions, suggesting that scene function is important to an environment's cognitive schema (a mental framework for organizing knowledge about the world) (Yun et al., 2013)

Another approach to describing a scene is to detect and identify objects in a scene. Satkin et al. (2013) used a database of 3D models of indoor spaces and matched objects in these models to the geometry of objects in a scene image. They used four visual attributes to describe both images and 3D models automatically. Once a 3D model was generated from the 2D image, it was matched with the most similar 3D model in a database of models of indoor spaces (Satkin et al., 2012). Ismail et al. (2018) were able to achieve scene understanding without using a database of models. Their method deconstructed scenes into two parts: an estimation of objects' spatial layout and a model of the spatial relationships between objects.

Functional aspects of environments can be used as descriptors to achieve automatic scene understanding (Hoiem et al., 2005; Lazebnik et al., 2006; Li & Perona, 2005; Oliva & Torralba, 2001; Vogel & Schiele, 2007). In a notable example, Gupta et al. (2011) broadened the concept of scene understanding by estimating the 3D geometry in the space and the potential for human poses, which they call the “workspace” of a scene. They argue that scene understanding means acquiring a human-centric understanding of the space by prioritizing geometry that is important to humans. For example, estimates of the floor would be more important than the estimates of the ceiling because humans normally interact with the floor. They achieved scene understanding by estimating the scene's geometry, modelling common human poses, and using these two sources of information to predict poses supported by the environment. Essentially, Gupta et al.'s (2011) approach describes scenes by the human postures and activities available in them: the scene is sit-able, walk-able, etc.

6.4 Scene Properties

Here I discuss how the scene taxonomy's visual properties are connected to the research discussed above. The task properties were derived from the review of VR applications.

6.4.1 Dimension 1: Visual Properties

Openness: I borrow Oliva & Torralba's (Oliva & Torralba, 2002) spatial envelope term and concept, which refers to how enclosed a space is. However, I do not define openness by the amount of horizon line visible because it can be hard for humans to quantify. This definition is closer to the distinction between “indoor” and “outdoor” observed in psychology experiments. Tversky and Hemenway (1983) devised a taxonomy of environment categories to reflect cognitive schema. Their taxonomy was organized into three levels. The top, most generic level had only “indoor” and “outdoor” as categories. This dichotomous categorization has also been observed by other

researchers in participants' scene descriptions during psychology experiments (Rummukainen et al., 2014; Rummukainen & Mendonça, 2016; Tversky & Hemenway, 1983).

Xiao et al. (2010) proposed a taxonomy of scenes which was used to organize a large database of scene images for scene categorization tasks in computer vision research. Like Tversky and Hemenway's (1983) taxonomy, their scene categories were structured in a three-level taxonomy, with the highest-level categories being "indoor", "outdoor natural", and "outdoor man-made".

Area: Participants used the amount of space covered in their descriptions of scenes in a study investigating how people categorize scenes (Rummukainen et al., 2014). Additionally, the area of scenes communicates the potential for certain human activities (Gupta et al., 2011; J. Xiao et al., 2010), which is relevant to affordance-based scene descriptions.

Scale: Scene understanding research did not surface this property, likely because scenes used in the studies are realistic and human scale. Scale was informed by the VR app review. Some scenes in VR applications varied in scale, such as Google Earth, which had a smaller scale.

Object Density: Object density is related to the spatial envelope property roughness, which describes the number of planes in a scene (Oliva & Torralba, 2001, 2002). Roughness is related to the size and shapes of objects in a scene. A scene with high roughness, such as a forest, will have many small planes and contours, whereas an empty room will have low roughness because it consists of few large planes (i.e., walls). There will be more planes in a scene with a greater number of objects, resulting in a high degree of roughness.

Object Tracking: I included this property based on participants' descriptions of scenes in cognitive psychology research. Rummukainen and Mendonça (2016) found that motion in a scene was one of the first things participants described when asked to recall a scene's properties. To

further highlight the relevance of motion is the existence of dedicated neural processes for perceiving it (Hess et al., 1989).

Social Actors: I included this property based on research in cognitive psychology. One of the most salient features in participants' scene descriptions was the presence of animate objects or objects with faces. In addition, research has confirmed the existence of a dedicated region of the brain for processing face-like arrangements (Grill-Spector et al., 2004). This property is binary (e.g., yes, no) because cognitive psychology research suggests that the mere presence of an animate object is enough to influence participants' scene descriptions (Buswell, 1935; Fletcher-Watson et al., 2008; Yun et al., 2013). Findings consistently show that individuals tend to fixate on faces in a scene (Buswell, 1935; Fletcher-Watson et al., 2008; Yun et al., 2013), suggesting that faces are an important component of scene descriptions.

6.4.2 Dimension 2: Tasks

In the review of popular VR applications (Appendix A.2) and app categories in mainstream VR app stores I found that the most common task types were movement, socializing, exploration, navigation, offense/defense, creativity, observation, and productivity. I further refined the task types and definitions based on the review of scene-viewing techniques in the research literature, discussed in the Related Work chapter.

6.5 Summary

Scenes are defined in various ways within cognitive psychology and computer vision; however, some definitions are similar across these fields. One definition that spans the three fields is that a scene is an immersive, three-dimensional space that contains objects and affords various interactions. Approaches to describing scenes are also similar across fields. A few ways of describing a scene are categorizing it, identifying the objects within it, or listing perceptual and

conceptual features that describe its visual structure. The visual properties of the scene taxonomy are informed by the various approaches to describing scenes. In the next section, I demonstrate how the taxonomy can be used to address an accessibility problem.

6.6 Application of the Scene Taxonomy

I applied the scene taxonomy to an accessibility problem in VR: the most prevalent scene-viewing technique in VR is head tracking, which requires the user to move their head and body to view the scene, yet head tracking might not be accessible to people with limited movement due to disability or situational impairment.

Using the scene taxonomy to organize techniques narrows the scope of suitable scene-viewing techniques to those that are most appropriate for the user's VE and task. Although organizing scene-viewing techniques by scene properties does not guarantee that they are accessible for any type of impairment, it enables designers to choose a technique from a smaller subset of scene-viewing techniques, which facilitates the selection process.

I classified the scene-viewing techniques reviewed in the Related Work chapter based on the types of environments they were designed for and evaluated in (Appendix A.1). I categorized techniques by the visual structure of VEs used in the study apparatus they were evaluated in by inspecting images in the paper and supplemental videos.

The purpose of classifying scene-viewing techniques was to use the scene taxonomy as a designer might when they are examining the design space of scene-viewing techniques to identify some to implement. Classifying scene-viewing techniques can help the designer choose techniques based on the VE they are designing. Classifying techniques can also reveal patterns in the design space such as popular VE types and tasks. Finally, classification can reveal VEs and task types for which scene-viewing techniques do not yet exist.




Figure 5. World in miniature (WIM) prototype (Stoakley et al., 1995). Left: A miniature model of the VE. Right: the miniature model of the VE with the human-scale VE behind it. I classified this VE with the scene taxonomy.

I classified the VE that the Worlds in Miniature (WIM) technique (Stoakley et al., 1995) was evaluated in to demonstrate that scene-viewing techniques are appropriate for specific scene properties (Figure 5). WIM was tested in an indoor space because the VE was the inside of a building. I classified the scale as human- and smaller-scale because the user was inside the building while manipulating a smaller model of the building. The space covered a small to medium area because the user was navigating a room in a building. It had low object density since it was composed of walls, doors, and some shelves. Also, because the environment was a single room, it had infrequent scene changes. Finally, there were no moving objects or social actors because they were not mentioned in the paper nor evident in screenshots of the VE. WIM was evaluated for exploration and navigation tasks (Stoakley et al., 1995). The user manipulated the first-person camera by maneuvering a camera proxy in a miniature 3D model of the enclosed space.

The metaphor of a miniature model was appropriate for the visual structure it was designed for; its boundaries were clearly defined because the VE was enclosed, making manipulating the camera like manually manipulating an object. If the environment was open, large, and had high object density, the miniature model metaphor might not be as effective. Guided camera control techniques

might be more effective for these environments because the user can choose from various paths that present the most important parts of a scene (Elmqvist et al., 2007).



Figure 6. A scene from the VR app The Room²². This scene would be categorized with the scene taxonomy in the same manner as WIM.

WIM could be a suitable technique for a VR app such as The Room (Figure 6) in which the user searches different rooms for clues to unlock the next level. The Room would also be classified as an indoor, human-scale, covering a small to medium area environment with low to moderate object density, infrequent scene changes, without moving objects nor social actors. The user could use WIM to position the camera proxy in the miniature model to wherever they wanted to be in the room.

The above example demonstrates how designers can use the scene taxonomy to identify appropriate scene-viewing techniques for a particular VE. Designers would first classify their VE using the taxonomy and then identify a scene-viewing technique that has been evaluated in a similar environment.

6.6.1 Common Scene Properties and Tasks in Research Virtual Environments


In the classification of scene-viewing techniques (Appendix A.1), I found that the most common scene properties were outdoor, human scale, large spaces, infrequent scene changes, no social

²² <https://www.thevrguid.com/the-room-vr-a-dark-matter/>

actors, and no moving objects to track (Table 3). The most common task was navigation (Table 3). Other tasks, such as offense/defense or socializing, can be considered to have navigation or exploration as a base task. For example, social applications require the user to navigate to an avatar to interact. Because exploration and navigation are the building blocks for other tasks, it is reasonable for researchers to focus on scene-viewing techniques for navigation tasks.

Table 3. This heatmap shows an overview of the scene-viewing techniques in research and is a summary of Appendix A.1. It reveals popular tasks and visual properties as well as gaps in research. I categorized the virtual environments and tasks that were used to evaluate the scene-viewing technique in the paper and then counted the papers in each category to produce the heatmap. The darker the cell, the more papers on scene-viewing techniques in that category.

Legend: Number of papers



		Tasks							
Visual Properties		Movement	Socializing, Collaboration	Offense, Defense	Exploration	Navigation	Creativity	Observation	Productivity
Openness	Indoor								
	Outdoor								
	Abstract								
Scale	Larger								
	Human								
	Smaller								
Area	Large								
	Medium								
	Small								
Object Density	High								
	Moderate								
	Low								
Object Tracking	Multiple								
	One								
	None								
Scene Changes	Frequent								
	Infrequent								
Contains Social Actors	Yes								
	No								

6.7 Using the Scene Taxonomy as an Ideation Tool

The empty cells in Table 3 revealed underexplored areas in the scene-viewing technique design space. In the classification I performed, visual properties that were underexplored were one or multiple objects to track, frequent scene changes and the presence of social actors. Designers and researchers could design novel techniques for these properties by considering the usability or

accessibility issues that panning would introduce in these environments. A designer could also identify a commercial app that has a VE with a particular property and then experience panning without moving their heads or bodies. In this way, a designer could understand what the user experience is when they are limited to a single scene-viewing technique. For example, the designer might find that panning would require multiple button-presses to orient the camera to discover a moving object. A designer or researcher might also classify their VE and identify scene-viewing techniques that were designed for the same VE types and tasks. Existing techniques built for a similar VE could provide inspiration for a novel technique.

Finally, researchers have demonstrated that the scene-viewing techniques that were classified were successful if used in VEs like the VE used for the study apparatus. However, scene-viewing techniques might also be useful in different types of environments from the ones they were tested in. The researcher or designer could create a new VE with properties that the technique has not been tested with. If the technique would not work in an environment with new properties, the designer could adapt the technique so that it could be used in this new environment, which could result in a novel technique.

Up to this point I introduced the scene taxonomy, reviewed the research it is grounded in, and discussed how the taxonomy can be used to classify scene-viewing techniques that require little to no head movement. I then evaluated the taxonomy and demonstrated its utility as an ideation tool by generating new accessible techniques. I wanted to investigate if the scene taxonomy was useful for generating scene-viewing techniques that leverage scene properties. In the following section I demonstrate how I identified gaps in the taxonomy and designed scene-viewing techniques that address the gaps.

6.8 Using the Taxonomy to Inform the Design of Scene-Viewing Techniques

I designed three scene-viewing techniques to demonstrate how the scene taxonomy can be used as an ideation tool. These three scene-viewing techniques can be used as alternatives to or extensions of panning when the user has limited head or trunk movement. A VE might not be accessible for people who cannot easily rotate their heads or bodies to change the camera orientation because most VR experiences rely on head tracking as the predominant scene-viewing technique. The goal of the techniques I developed was to facilitate the perception, discovery, and understanding of a scene without requiring body movement or repetitive controller manipulation.

6.8.1 Implementation Details

All techniques were implemented in Unity 2019.1.7. I used the Oculus Rift S headset and controllers, which was connected to a Lenovo ThinkStation P330.

In addition to the three scene-viewing techniques, basic controls (panning, teleport, pointing) that could be used in parallel with the scene-viewing techniques were implemented. In the implementation of panning, when the user pushed the left controller's thumbstick left or right, the camera would rotate 30° to the left or right on the y-axis. A thirty-degree rotation angle was chosen based on prior work (Sargunam & Ragan, 2018).

The VRTK²³ 3.3 straight pointer was used, which rendered as a ray emitting from the user's left controller. To activate the pointer, the user touched the top of the left thumbstick. To select with the pointer, the user pulled the left trigger. The VRTK 3.3 Bezier point and teleport technique was also used, which rendered as a dotted curve emitting from the user's right controller. Teleport was used as the locomotion technique because it is implemented in many VR applications (Funk et al.,

²³ <https://vrtoolkit.readme.io/docs/summary>

2019). To activate the teleporter, the user touched the right thumbstick with his or her thumb. To teleport to a location, the user pressed the right trigger.

6.8.2 Design Considerations: Overview

The main objective when designing scene-viewing techniques was to make the techniques accessible for people with limited head or trunk movement. As such, techniques that required no head movement and little body movement (e.g., arms, hands) and fewer interactions with the controller than panning were designed. In addition to accessibility, factors that are relevant to VR were also considered. VR interactions for locomotion and object manipulation are often evaluated based on presence, simulator sickness, and spatial awareness (Al Zayer et al., 2020). Additional design considerations were identified based on early VR research on evaluating interaction techniques (Bowman & Hodges, 1997), and a textbook on best practices for designing VR experiences (Jerald, 2015) to guide the design process.

6.8.3 Design Considerations: Accessibility

The main accessibility issue was the effort needed for a user to view a scene, assuming head or trunk rotation was not an option for the user. Panning requires the user to push their thumbstick either continuously or repeatedly in order to orient the camera. The camera then either moves continuously or at discrete angles until the desired orientation is achieved. If a user is in a VE that requires them to be aware of what is around and behind them, it could be tiring to repetitively pan the camera back and forth. This interaction might be challenging for people with conditions that limit their head movement and even more challenging for people who also have poor hand strength or coordination.

6.8.4 Design Considerations: Usability

People who use their controllers to orient their cameras are also likely to need to use their controllers for locomotion and object manipulation. As a result, many different controls are mapped to one controller, which can be confusing and difficult to remember. The user must be able to easily recall and form a mental model of controls with low cognitive effort for the technique to have good usability (Nielsen, 1994b, 1994a).

With current implementations of panning, the user can only orient the camera laterally, (i.e., rotate on the y-axis). Some VEs might have points of interest above or below the user. But the most common implementation of panning in commercial applications does not enable users to orient their camera towards these objects. The lack of freedom when controlling the camera angle with panning could be detrimental to its usability in environments where points of interest surround the user.

6.8.5 Design Considerations: Realism

Many VR experiences aim to make users feel as if they are physically present in the environment. There is evidence to suggest that a relationship between realism of a VE and presence, the feeling of being physically present in a VE, exists (Hvass et al., 2017). As a result, many VR interactions aim to mimic real world interactions.

6.8.6 Design Considerations: Spatial Awareness

Some tasks in VR require the user to orient themselves relative to other objects in the VE. Therefore, a user must have a sense of where they are located. Users could feel disoriented after camera transitions, depending on how much of the scene is shown during the transition. Enabling the user to receive enough information to remain oriented is a part of this design consideration.

6.8.7 Design Considerations: User Comfort

A main design problem in VR is that interactions could inadvertently cause simulator sickness in individuals. Simulator sickness occurs when there is a discrepancy between movement in the virtual and physical environment (Moss & Muth, 2011). For example, continuous camera movement is usually avoided because there is a significant disconnect between perceived virtual movement and the user's physical movement. Simulator sickness is important to consider, especially for scene-viewing techniques that require little movement from the user.

6.8.8 Design Considerations: User Familiarity with Interaction

Users might not be familiar with interaction techniques in VR because it is still emerging as a consumer technology. To help scaffold the learning process, designers can incorporate interactions that are used in more familiar technologies. For example, a user would likely be familiar with a point and click interaction from desktop computing, or a tap from touchscreen interaction. Using familiar interactions could also reduce the cognitive demand imposed by the controller mappings because there would be fewer novel interactions to learn (Nielsen, 1994a).

6.8.9 Design Considerations: Task Objective

Efficiency might be important for some tasks, however considerations such as presence and simulator sickness are also major parts of designing a good interaction technique in VR. Some VEs are rich in detail and the main task is to observe or explore the environment. In this case, an efficient scene-viewing technique might detract from the experience of the VE or make the task too easy. Therefore, understanding the task objective, and whether efficiency is a priority, can inform the design of the scene-viewing technique.

6.8.10 Design Considerations: Summary

It would be ideal to optimize all design considerations but designing the perfect scene-viewing technique might not be possible. Designing a single “perfect” technique is probably not desirable because motor impairments manifest in diverse ways and having different options would likely improve the accessibility of a VE (Trewin, 2000). In addition, accessibility is only one consideration of many when designing a scene-viewing technique. An individual might want to choose a technique that has trade-offs in other design considerations. For example, if a user is not sensitive to simulator sickness, they could choose a technique that is less comfortable but more realistic. Giving users the freedom to choose techniques based on their individual needs and preferences would enable people to take advantage of the diverse benefits that VR affords. I identified these trade-offs when I designed the scene-viewing techniques.

6.9 Technique 1: Object of Interest

There were few techniques designed to track multiple objects in a VE (Appendix A.1), revealing a gap in the design space. This gap surfaces design question: How would a user control the camera view while simultaneously being aware of multiple moving objects? I designed object of interest for outdoor, large spaces with moderate to high object density, and multiple objects to track. The primary focus was to design a technique that would enable users to be aware of and track multiple moving objects in a scene while exploring or navigating. Objects in this environment might appear below or above the user. Unlike panning, object of interest enables users to orient the camera vertically to look above or below themselves.




Figure 7. Object of Interest: a. The user is facing the octopus in the scene (birds-eye view), b. She selects the turtle icon with her laser pointer, c. Her camera rotates to face the turtle in the scene (birds-eye view), d. She is facing the turtle in the scene.

Object of Interest was designed so that icons are visible in front of the user (Figure 7). An icon indicates that an object of interest is outside the user's field of view and appears to the left or the right side of the user, indicating which side the object is on. When the user selects the icon, their camera view cuts to a view that is centered on the object of interest. Icons are visible so long as the user has them toggled on. The user can choose to toggle icons off.

I implemented Object of Interest in an underwater sea environment with sea animals swimming above, below, and around the user. There were four different icons for each sea creature: orca, turtle, dolphin, and octopus. The user could toggle the icons on or off with the left controller grip button. When the user selected the orca button, for example, the user's view automatically shifted to an orca in the scene.

When designing the Object of Interest technique, the focus was on making a technique that was more accessible, usable, and comfortable than panning. The camera automatically oriented to the user's desired object after the corresponding icon was selected, which demonstrated improved accessibility (5.2.1), because the user does not have to push the thumbstick multiple times to orient the camera to the moving object. I used a UI with icons floating in the virtual world so that functionality was offloaded from the controller. Icons provided affordances that facilitated

recognition over recall, which is a usability heuristic (Nielsen, 1994a). The camera instantly oriented to the object of interest in order to prevent simulator sickness (5.2.5 user comfort).

On the other hand, the nature of the instant camera transition and the presence of a UI decreases realism because there is no real-world equivalent (5.2.3 realism). The immediate camera transition also prevents the user from understanding where the object is relative to other objects in the scene, which could diminish their spatial awareness (5.2.4). The selection mechanism, ray-casting, is unique to VR but the underlying principle of pointing and clicking might be familiar enough to help users learn how to use the technique (5.2.6 user familiarity with interaction). The immediate camera transition enables a more efficient and usable way of viewing objects compared to panning; however, efficiency might not be desirable when exploring an environment because the user would miss contextual information when the camera orients to the object (5.2.7 task objective).

6.10 Technique 2: Proxemics Snapping

Despite the growing interest in social VR applications, few scene-viewing techniques assist social interactions (Appendix A.1). The structure of the space changes when social actors are introduced because the space contains zones of interpersonal distance (Hall, 1966). In the physical world, individuals can easily approach and orient themselves relative to other people while being aware of personal space. However, managing personal space is more complex in VR when using controller-based techniques such as panning and teleport.

To use panning and teleport, the user would have to aim their teleporter to an appropriate location in front of another user, which could take several attempts. Next, they would have to fine-tune the camera's orientation using panning, which could take time. I designed Proxemics Snapping for indoor spaces with social actors where the main task is to socialize. This technique

was designed to help users search for other avatars in an environment and achieve socially acceptable distances with little effort.




Figure 8. Proxemics Snapping: a. The user's view is zoomed out so that the environment is in the frame, b. She picks up and moves his dark blue avatar, c. She places his avatar in the light blue snap drop zone where his avatar orients to a socially appropriate distance and orientation relative to another avatar, d. She deactivates proxemics snapping and returns to a first-person view facing the red avatar.

When the user invokes Proxemics Snapping, the camera switches to a third person, zoomed-out view of the environment (Figure 8). Because the technique is designed for indoor spaces, the building's front wall is toggled off, allowing the user to see into it. The user can see their own avatar as well as the other avatars in the scene. The user can then pick up their avatar and place it in one of several pre-designated locations, or "snap drop zones" near other avatars. When users place their avatar in one of these zones, the avatar snaps to a specified orientation to face the avatar they want to interact with. When users deactivate proxemics snapping and resume their first-person view, they will be in front of the other avatar at a socially acceptable distance and orientation.

In the implementation of this technique, the environment was a two-story building with a staircase in the center. Four avatars representing other users were placed around the building. Four snap drop zones were placed in front of the avatars. Snap drop zones were preconfigured distances and orientations relative to the other avatar represented as a light blue shadow of the user's avatar. When the Proxemics Snapping technique was invoked by pressing the "A" button on the right controller, the user's controllers scaled up by a factor of 40 and the front wall of the structure was toggled off, revealing the inside of the house. The user could then pick up her avatar by reaching

into the house and pressing the right controller grip button to pick up the avatar. When the user hovered their avatar over a snap drop zone in front of another avatar, it would appear as a blue, transparent copy of the user's avatar. The user could then release the right grip to release the avatar into the snap drop zone. Upon releasing the avatar, the user would return to first-person view.

With Proxemics Snapping, the user can see all avatars in the environment and automatically be positioned in a socially acceptable way without performing small adjustments, which would require more controller use (5.2.1 accessibility). Since the main task objective was to interact with another avatar, the ability to execute the interaction efficiently was appropriate for this task (5.2.7 task objective). The outline of the avatar served as an affordance indicating where the user could put their avatar. This feature satisfied the recognition over recall heuristic (5.2.2 usability). The interaction was like putting a doll in a dollhouse so was more realistic than using a user interface or teleporting to perform the interaction (5.2.3 realism). While viewing their avatar from a third-person point of view, the user could see where their avatar is relative to the rest of the environment, which could improve spatial awareness (5.2.4). And when the user performed the technique, there was no continuous camera movement that would cause simulator sickness, because the user immediately switched between stationary views (5.2.5 user comfort).

However, the interaction required the user to reach into the environment and pick up their avatar, which could be inaccessible for people who have limited range of motion, arm strength, or arm stability (5.2.1 accessibility). In order to invoke Proxemics Snapping, the user had to remember to press a particular button on the controller, which violated the principle of recognition over recall (5.2.2 usability). Also, the interaction was unlike any other, so users could be unfamiliar with it and would have to learn how to use it (5.2.6 user familiarity with interaction).

6.11 Technique 3: Rearview Mirror

The third gap I identified was the need for scene-viewing techniques for offense/defense tasks with frequent scene changes (Appendix A.1). I designed Rearview Mirror for indoor spaces, with one or multiple moving objects to track, frequent scene transitions, and no social actors. The objective for this technique was to give users greater awareness of their environment and the ability to change their view quickly.




Figure 9. Rearview Mirror: a. The user is facing an empty room (birds-eye view), b. She sees an avatar in her rearview mirror, c. She presses a controller button to flip 180° (birds-eye view), d. She faces the blue avatar.

For this technique, a rearview mirror appeared in front of the user and displayed what was behind their avatar (Figure 9). The user could also rotate 180° to face the opposite direction by pressing a controller button.

In the implementation, I positioned a second camera in front of the user and pointed it in the opposite direction the user was facing. To create the rearview mirror, the camera viewpoint was displayed on a small rectangle in front of the viewer and was always visible. When the user pressed the “A” button on the right controller, they rotated 180°.

Accessibility was the primary consideration when designing the Rearview Mirror technique. The user could see behind themselves without any interaction by simply looking in the rearview mirror. Also, the user could turn to face the opposite direction with a single button press instead

of multiple thumbstick presses (5.2.1 accessibility). Users would be familiar with looking in a rearview mirror if they knew how to drive (5.2.6 user familiarity with the interaction, 5.2.3 realism). Users might be able to maintain spatial awareness, even when the camera flipped to the opposite direction, because the rearview mirror would remain in the same position and always show the environment behind the user (5.2.4 spatial awareness). Users would not experience discomfort because there was no continuous camera movement (5.2.5 user comfort). Because the task was time sensitive, efficiency is important. Therefore, this technique was appropriate for the task because it reduced the time a user would need to look behind themselves and react to whatever was there (5.2.7 task objective).

6.12 User Study

In the user study, I limited participants' head movements to investigate effectiveness of the three scene-viewing techniques discussed above compared to panning. The main goal was to evaluate the scene taxonomy and demonstrate its utility as an ideation tool by designing three scene-viewing techniques that were as easy or easier to use than panning. Designing new usable techniques would give the user more options for ways of viewing scenes in VR. I was also interested in which of the design considerations were important to users and how they perceived trade-offs in design considerations. I designed the techniques to be first and foremost accessible by not requiring head or trunk movement. I also wanted to ensure that I did not sacrifice usability while improving accessibility (e.g., by making the technique difficult to remember or inefficient). So, I was particularly interested in how users would rate the usability of the techniques.

6.12.1 Participants

Sixteen individuals (women: $n=5$, men: $n=11$) with a mean age of 34.6 ($SD=10.6$) participated in the study. Participants rated their expertise with computer systems (e.g., computers, tablets,

smart phones, etc.) on a scale from 1 (novice) to 5 (expert). The median expertise was 3 (*IQR*=1). Five participants had never owned or used VR while 11 had. Of the 11 participants who had used VR, most reported using it a few times a year (*n*=6), followed by almost never (*n*=4), and one person reported using it every few days. Participants who had used or owned VR rated their expertise with VR on a scale from 1 (novice) to 5 (expert). The median expertise was 4 (*IQR*=1). Only one participant reported experiencing motor limitations, which were poor coordination and rapid fatigue. One participant reported hearing loss.

6.12.2 Apparatus

Participants used the Oculus Rift S headset and controllers. They sat in an office chair and placed their chin on a chinrest, which was attached to a desk (Figure 10). Participants' head movements were limited by the chinrest while completing tasks.



Figure 10. Participants' head movements were limited with a chinrest during the study.

6.12.3 Procedure

The participant first signed a consent form and filled out a demographic questionnaire. They then adjusted the chin rest to a comfortable height.

Next, the researcher explained how to operate the VR controllers to rotate the camera, teleport, and point. These basic controls were available to the participant throughout the study. The participant then completed a training session to practice using the basic controls in a simple environment containing a floor and a keyboard that was suspended in the air. The researcher asked

the participant to pan their view left and right using the thumbstick, teleport around the scene, and enter characters into the keyboard with the pointer.

Table 4. Tasks participants completed with panning and the new scene-viewing techniques.

Technique	Task
Object of Interest	"Look at one of each kind of sea animal (orca, turtle, dolphin, octopus)".
Proxemics Snapping	"Stand in front of each avatar as if you were talking to him or her".
Rearview Mirror	"Eliminate as many spheres as you can before 2 minutes is up".

After the basic controls training, the participant practiced using the first technique in an empty scene. Once the participant felt comfortable using the controls, the researcher loaded the task environment. The participant completed the task with the new scene-viewing technique either enabled or disabled (Table 4). When the technique was disabled, the participant could only use the basic controls to perform the task. When the technique was enabled, the participant could complete the task with the technique in addition to the basic controls.

After completing the task, the user took off the headset and completed a NASA-TLX questionnaire (Hart & Staveland, 1988) to measure workload. I was interested in the physical and mental effort it took to use the techniques to measure the extent that usability in the form of workload was compromised by improved accessibility. Participants then put the headset back on and placed their chins in the chinrest. Once participants were in position, they completed the same task again. If the technique was enabled the first-time participants completed the task, it was disabled the second time and vice versa.

After completing the task again (with the technique enabled or disabled), the user filled out a second NASA-TLX questionnaire. The researcher then interviewed the participant and elicited comparisons between panning and the new scene-viewing technique in terms of preference, ease of use, presence, and simulator sickness. I was mainly interested in the accessibility and usability

of the technique, so I did not administer additional questionnaires, such as one for presence (Slater et al., 1994). However, if participants were unsure of what I meant by “presence” I explained that it was the sense that they physically existed in the virtual space.

The researcher also asked the participant for their feedback on the scene-viewing technique and if the participant would use the technique in VR. This procedure was repeated for each of the three techniques.

6.12.4 Design and Analysis

The study was a within-subjects design with one factor, scene-viewing technique and two levels: panning and new scene-viewing technique. I used a 4x4 latin square to counterbalance the techniques. I also counterbalanced the presentation order of the scene-viewing technique such that half of the participants experienced panning first and half experienced the new scene-viewing technique first.

I compared NASA-TLX scores for panning and the new scene-viewing technique for each task. I adapted the original 21-point scale to a 1 to 7 scale by eliminating the “high”, “medium”, and “low” increments for each point on the scale to improve readability (Hart, 2006). Lower scores indicated lower workload. I then compared the raw (unweighted) NASA-TLX scores for panning and new scene-viewing techniques using a Wilcoxon Signed-Rank test because the data were not normally distributed: A Shapiro-Wilk test revealed that the distribution was significantly different from a normal distribution ($W=0.94, p<.05$).

Participants’ interview responses were audio recorded and transcribed. Responses were then qualitatively analyzed using a thematic analysis by finding patterns across participant utterances for each question and deductively grouping utterances by theme (design considerations) (Braun & Clarke, 2008). I used qualitative data to understand how design considerations contributed

participants' perceptions and preferences for scene-viewing techniques. As such, I did not conduct statistical tests on the count data surfaced during the interview because I was interested in how participants in this sample weighed trade-offs in design considerations. Count data could provide insight into participants' thought processes. Qualitative data were also used to provide insight into NASA-TLX scores.

6.13 Results

I summarize results from the user study below. Although participants were able to use panning in all conditions, I refer to the condition in which participants used the new scene-viewing technique with panning as the [technique name] condition. I refer to the condition in which participants used panning without the new scene-viewing technique as the panning condition.

Table 5. This table presents the participant counts of their preferred technique and the technique they perceived as easier. It also presents NASA-TLX median scores and IQR for the new technique and panning. Higher NASA-TLX scores mean higher perceived workload. The data indicated that for some participants, their preferred technique was not the same as the technique they thought was easier.

Technique	Object of Interest			Proxemics Snapping			Rearview Mirror		
	Preferred (count)	Easier (count)	NASA-TLX (Median, IQR)	Preferred	Easier	NASA-TLX	Preferred	Easier	NASA-TLX
Technique + panning	7	9	1.5, 1.0	8	12	1.0, 1.0	11	9	3.3, 1.8
Panning only	8	1	1.5, 1.3	7	1	2.0, 1.0	4	5	3.0, 1.8
No preference	1	6		1	3		1	2	

6.13.1 Technique 1: Object of Interest

Workload: I did not find a significant difference between NASA-TLX scores for panning vs. the Object of Interest technique ($Z=0.20$, *n.s.*). Interview responses to the question, “which was easier to use for this task, the Object of Interest technique or panning?” did not reflect the questionnaire results: Even though participants rated both techniques the same in terms of workload, nine participants responded that the Object of Interest technique was easier, one responded that panning was easier, and six responded that they were the same (Table 5).

Design Considerations: Participants were divided almost equally on preference for panning vs. object of interest. Most participants liked that Object of Interest enabled them to identify and locate objects faster (5.2.7 task objective). For example, P06 explained that object of interest helped her find objects: *“Maybe I didn't know where a particular animal was, then I could use the object of interest technique.”* (P06) (4.2.4 spatial awareness).

Some participants also pointed out that object of interest allowed them to locate moving objects more easily than panning: *“Since the object keeps moving, I think it's easier to have [the object of interest technique] to easily locate it”* (P04) (5.2.2 usability). Some participants found that it was challenging to use panning when they were trying to track a moving object. For example, P11 said: *“[The Object of Interest technique] would make it easy to find things in a super crowded place, especially if you have limited mobility. I didn't have to move much or work very hard to find it. Just, oh, it's there”* (P11) (5.2.1 accessibility).

Icons also provided information about the identity of objects in the scene. P12 explained: *“I liked that I didn't have to guess about whether or not I was seeing something. You go and look and there's something that looks like a sea turtle. Okay well that's a sea turtle because you told me there's a sea turtle.”* (P12). The icons enabled participants to recognize the object that they were looking at, so they knew what types of objects were in the environment (5.2.4 spatial awareness).

However, half of the participants reported that they disliked feeling disoriented after using Object of Interest. P05 explained, *“It was quicker to get there, but I wasn't as oriented as a result. So, it got me there, but I was not 100% about my bearings and I had no idea where any of the other things were.”* (P05). The immediate camera transition to the object of interest caused confusion about where in the environment they were currently looking (5.2.4 spatial awareness).

None of the participants felt simulator sickness using object of interest (5.2.5 user comfort). However, a majority reported feeling less present in the environment with Object of Interest compared to panning. P03 said, *“Just I found, I was focusing more on the buttons than the actual environment itself, I think that's fun to have control like that, but for this scenario, I feel it takes away from the wonderment of just looking around and enjoying the environment”* (P03). P03 raised the point that having a UI in the virtual space detracts from the realism of the environment which reduced his sense of presence (5.2.3 realism).

Usefulness: Despite half of participants feeling disoriented after using the technique, most participants (n=14) said they would want to use object of interest for finding objects if efficiency was important for the task (5.2.7 task objective).

6.13.2 Technique 2: Proxemics Snapping

Workload: I did not find a significant difference between NASA-TLX scores for panning vs. Proxemics Snapping ($Z=1.11$, *n.s.*). Yet, most participants responded that Proxemics Snapping was easier to use (Table 5).

Design Considerations: Although questionnaire and interview results indicated that most participants perceived Proxemics Snapping to be easier, participants' preferences were not always based on ease of use (5.2.2 usability). Participants were almost equally divided in terms of their preferences for Proxemics Snapping vs. panning. Six participants liked Proxemics Snapping because they did not have to correct their position and orientation to achieve an appropriate personal space relative to the other avatars (5.2.1 accessibility, 5.2.2 usability). For example, P05 said, *“You get a perspective on everything in the space and then you can decide and just go to where you want basically in one move, which I think was more efficient”* (P05).

When using panning, participants were often too close to the other avatars, too far, or at an angle that was not optimal for conversing. P12 explained that he settled for a suboptimal orientation relative to the other avatar when using panning because it was difficult to achieve the right positioning (5.2.2 accessibility): “[Panning] did make it more likely that I would keep something I thought was slightly off. So, if I was a little bit too close, then it's a lot of work to turn around and try to teleport sideways and then turn back around so I'd probably just leave it” (P12).

Although some participants reported that they did not dislike anything about Proxemics Snapping, a few reported that interacting with their avatar from a third-person view decreased their presence because it took them out of the first-person view (5.2.3 realism). P14 explained: “You couldn't really explore with [Proxemics Snapping]. I just like being able to walk around the house, not really walk, but just teleport around and look at stuff from different angles” (P14). P15 also thought that using the third-person perspective negatively impacted her sense of presence: “With the [Proxemics Snapping technique], I just dropped the avatar and it kind of oriented itself, so it was a little impersonal. So it was like a dollhouse, it wasn't an experience” (P15).

One participant reported feeling simulator sickness with Proxemics Snapping when she was in the third-person view, but this was not an issue for most participants (5.2.5 user comfort). Most participants mentioned that they experienced less presence with Proxemics Snapping compared to panning (5.2.3 realism).

Usefulness: Despite some participants feeling less present when using the Proxemics Snapping, 13 participants reported that they would want to use the Proxemics Snapping in VR. Participants would use it for locating and moving to other users' avatars in social environments and/or if the environment was large and complex (e.g., with multiple stories and staircases).

6.13.3 Technique 3: Rearview Mirror

Workload: I did not find a significant difference between NASA-TLX scores for panning vs. Rearview Mirror ($Z=0.42$, n.s.). A majority responded that Rearview Mirror was easier to use than panning but their NASA-TLX scores were similar (Table 5).

Design considerations: Most participants preferred Rearview Mirror over panning because they liked being able to see behind themselves, which increased their awareness of the environment (5.2.4 spatial awareness). P14 suggested that Rearview Mirror compensated for not having peripheral vision in VR: *“[I like that there was] just more visibility because you don't have as much peripheral vision. So, it gives you more visibility behind you. I mean you still don't have the peripheral but more visibility”* (P14). Participants also appreciated the fact that Rearview Mirror required only a single button press to turn 180° as opposed to panning which required multiple (5.2.1 accessibility, 5.2.2 usability). P01 explained: *“[I liked] the manner or ability of switching sides very easily. So even if I didn't end up like using that to kill the [enemy] I ended up using it someplace else just to go around faster”* (P01). P09 also mentioned the challenge of repetitively pushing the controller to use panning (5.2.1 accessibility). He said, *“When I got up the stairs it'd flick around and then have a better view instead of constantly having to turn. You have to flick quite a few times”* (P09).

Although participants appreciated the accessibility of Rearview Mirror, some participants felt that it increased cognitive load (5.2.2 usability). P11 explained: *“I felt like there were too many inputs into my brain and too many outputs in my hand that I had to remember. It was the figuring out how to get from seeing behind me to actually utilizing that information that was hard”* (P11). The 180° turn feature violated the heuristic of recognition over recall, which could explain why participants struggled to use it (5.2.2 usability). Like P11, a few other participants reported

forgetting or underutilizing the 180° turn feature despite liking that they could see more with Rearview Mirror (5.2.2 usability). P16 verbalized this issue by saying, “*I'm not really making use of [the turn feature] because... My impression is that there are too many options in my hands so that may make it little difficult for me*” (P16). Despite these comments, the NASA-TLX scores indicated that the workload was not significantly greater with Rearview Mirror compared to panning.

Only one person reported feeling simulator sickness with Rearview Mirror compared to two participants with panning. The fact that two people felt simulator sickness with panning might be because of how frequently they had to pan the camera to see what was behind themselves (5.2.5 user comfort). Most participants reported feeling present with Rearview Mirror and a few even reported feeling more present compared to panning. P05 explained that being able to see more of the environment increased his sense of presence (5.2.3 realism): “*I did feel present a little more, because I could see more, what was going on around me*” (P05). On the contrary, a couple of participants reported feeling less present with Rearview Mirror because they were less focused on the environment (5.2.4 spatial awareness).

Usefulness: Most participants responded that they would use Rearview Mirror in VR if it was available (n=12). Some participants said they would use it for a first-person shooter game to see what was behind themselves.

6.14 Discussion

In this section, I reflect on the results and process of designing new scene-viewing techniques with the scene taxonomy. I also reflect on the opportunities and limitations that the taxonomy affords. Finally, I discuss how the application of the scene taxonomy could inform the design of scene-viewing techniques for individuals with physical impairments.

6.14.1 Tradeoffs to Consider

Most participants reported that the techniques I designed were easier to use than panning, even though there was not a significant difference in their NASA-TLX scores for the new techniques (Table 5). Each technique design centered around specific scene properties, which were identified as gaps in the scene-viewing technique design space. The findings suggested that by leveraging these affordances in the scene, it is possible to generate different options for viewing a scene with little head or trunk movement. I demonstrated the scene taxonomy's usefulness for generating scene-viewing techniques that are just as easy if not easier to use than the default panning technique.

However, participants did not always prefer the technique they considered easier or more efficient. This finding reveals that it is important to consider multiple design considerations when building accessible techniques.

Object of interest: Even though panning and Object of Interest were similar in terms of workload according to the NASA-TLX scores, most participants verbally reported that they thought Object of Interest was easier to use than panning (n=9). Participants indicated trade-offs in spatial awareness, usability, accessibility, and realism: Object of Interest helped users be aware of their surroundings (5.2.4 spatial awareness), helped them identify objects they were viewing (5.2.2 usability), and could switch their view to one of the objects with less controller interaction than panning (5.2.1 accessibility). However, the instant camera transition to the object was disorienting (5.2.4 spatial awareness) and the presence of the UI in the scene detracted from the realism of the environment (5.2.3 realism). These trade-offs appeared to be based on individual preferences since participants were almost equally divided for their preferred technique (Object of Interest: n=7, panning: n=8).

Proxemics snapping: Like Object of Interest, there was not a significant difference in NASA-TLX scores, but most participants reported that Proxemics Snapping was easier to use (n=12). Participants thought Proxemics Snapping enabled them to achieve acceptable personal distances relative to other avatars with less controller manipulation than panning (5.2.1 accessibility, 5.2.2 usability). However, most participants felt less present with Proxemics Snapping when the camera switched from first-person to third-person view, revealing a potential trade-off between usability, accessibility and realism (5.2.3). Also, like Object of Interest, these trade-offs appeared to be based on individual preferences since participants were almost equally divided on their preferred technique (Proxemics Snapping: n=8, panning: n=7).

Rearview Mirror: There was also no significant difference in NASA-TLX scores for panning vs. Rearview Mirror even though most participants reported that they thought Rearview Mirror was easier to use (n=9). Participants appreciated that the Rearview Mirror technique required less physical exertion (5.2.1 accessibility). Even though some participants felt slightly more mentally burdened when using the turn feature (5.2.2 usability), participants felt the same or a greater level of presence when using Rearview Mirror compared to panning (5.2.4 spatial awareness, 5.2.3 realism). Most participants preferred Rearview Mirror to panning (Rearview Mirror: n=11, panning: n=4) possibly because not only did it enhance accessibility, but for some participants it also enhanced their sense of presence.

Summary: The results suggest that NASA-TLX might not have captured participants' concept of workload because their scores did not always reflect the technique they verbally reported to be easier to use. There might be different factors related to workload in VR that the NASA-TLX did not capture. Also, although many participants agreed that the scene-viewing techniques designed with the scene taxonomy were easier to use than panning, this was not reflected in their preferred

technique. This finding demonstrates that accessibility and usability are not the only design considerations that determine a person's preference for a scene-viewing technique.

The results suggested that participants weighed multiple factors when choosing their preferred technique. For Object of Interest and Proxemics Snapping the weighting of factors appeared to be based on individual preference. However, most participants preferred Rearview Mirror over panning which could indicate that they weighed the trade-off between accessibility, usability, and realism in a similar way. Future work could explore what appears to be a cost-benefit analysis for a preferred scene-viewing technique. There might be universally desirable or undesirable trade-offs that could help designers predict if their technique will be preferred over other scene-viewing techniques.

Because trade-offs exist for scene-viewing techniques and panning, designers should provide the option of using alternative scene-viewing techniques in addition to panning. Users could then choose a technique depending on their preferences for accessibility, usability, realism, spatial awareness, comfort, and familiarity with the interaction, as well as the goal of the task they are performing. Greater flexibility would provide users with more options for overcoming limited head movement caused by a disability or situational impairments.

6.14.2 Reflecting on the Scene Taxonomy: Opportunities and Limitations

The scene taxonomy aims to describe the context of a VR experience in a manner that is consistent with research on how people and computers conceptualize scenes. It could have been designed based on users' motor abilities so that designers could select a scene-viewing technique based on the type of impairment an individual has. The problem with this approach is that it is difficult to predict the VR interactions people will struggle with based on their condition and current context. Two people with the same condition could have very different ways of moving.

As a result, it would be necessary to study how people with various motor impairments use the scene-viewing techniques to understand how impairment type and severity relate to performance. Because of this issue, scene-viewing techniques were organized based on the types of environments they were designed for and evaluated in, since the success of the scene-viewing techniques in particular environments is already known. Then, given a VE, a designer could narrow down the subset of scene-viewing techniques that are applicable for their application and evaluate this subset with people with motor impairments to understand their performance.

I demonstrated the use of the scene taxonomy to reason about an accessibility problem in VR, but the taxonomy could be applied to organize the scene-viewing technique design space for a range of purposes. For example, it could be used to organize the design space of movement-based scene-viewing techniques as well. Also, because I constructed the taxonomy in a way that is independent of the properties of any particular set of techniques, it can also be used for categorizing other types of VR interaction such as locomotion techniques, object manipulation techniques, input devices, etc., to reveal patterns and gaps in these design spaces. In general terms, the taxonomy could enable designers to reason about the design space of various types of VR interaction relative to VR scenes and tasks.

As for a limitation of the scene taxonomy, I found that participants considered multiple trade-offs when deciding whether they would use the new techniques. Looking back, it would not have been possible to use the taxonomy to predict the particular trade-offs associated with a technique. These trade-offs would need to be surfaced empirically by evaluating techniques with users. Therefore, while the scene taxonomy can identify patterns in the scene-viewing technique design space, it cannot be used to predict how users will evaluate a given technique.

6.14.3 Accessibility for Individuals with Motor Impairments

The scene-viewing techniques designed with the taxonomy could be useful for individuals with limited head or body movement due to ALS, cerebral palsy, or paralysis. However, controller interactions would also need to be made accessible. For example, if an individual only has use of one hand, interactions for teleport, object manipulation, panning, and scene viewing would need to be usable with one controller. Also, the accessibility and usefulness of some VR interaction techniques could depend on the input devices being used (e.g., multiple switches, game console, foot controller, eye trackers, etc.). Automatically mapping inputs to interaction techniques could be a promising area of future work. Ultimately, the scene taxonomy could contribute to a recommender system for suggesting and mapping VR interaction techniques based on scene properties, task types, hardware affordances, users' abilities, as well as preferences for the design considerations that were identified in this study.

6.14.4 Limitations

I classified scene-viewing techniques based on the environments in which they were evaluated. However, it is possible—and likely—that the techniques could be classified under different visual properties and tasks. Therefore, the taxonomy is only a starting point for classifying scene-viewing techniques, and I acknowledge that any single technique could be classified under different properties based on empirical evidence.

Furthermore, having more participants might have revealed significant differences between NASA-TLX scores. There was larger difference in NASA-TLX scores for Proxemics Snapping and panning compared to Object of Interest and Rearview Mirror, indicating that more participants in this condition might have revealed a significant difference in scores.

6.15 Summary

I devised a scene taxonomy based on the visual properties and tasks associated with VEs from a review of cognitive psychology and computer vision research, as well as a survey of 29 popular VR applications. I then demonstrated how the taxonomy can be used to organize the large body of literature on scene-viewing techniques to address the problem of situational and permanent impairments that affect users' abilities to view scenes. I applied the taxonomy to identify accessible scene-viewing techniques that could be used for different environments and tasks. I also used the taxonomy to identify gaps in the design space that could suggest scene-viewing techniques for people with limited mobility. Based on the gaps I identified, I prototyped three scene-viewing techniques, which I evaluated with participants experiencing limited head movement. I found that most participants thought the new techniques were easier to use compared to panning, however, participants based their preferences on trade-offs in accessibility usability, realism, spatial awareness, comfort, familiarity with the interaction, and the task objective. The taxonomy could potentially be used to reason about various problems where the scene and task affect how users interact with VR.

7 NEARMI: A FRAMEWORK FOR DESIGNING POINT OF INTEREST TECHNIQUES FOR VIRTUAL REALITY USERS WITH LIMITED MOBILITY²⁴

Many applications contain more virtual content than what can be displayed by a device at a given time, requiring users to scroll, pan, or zoom to find out-of-view information. For instance, users might be required to pan or zoom a virtual map to see the location of a restaurant, or to scroll a document to view a specific section. Over the years, researchers have developed numerous approaches to improve the experience of identifying and navigating to out-of-view points of interest (POIs) for various computing platforms, including virtual reality (VR) (Matthews et al., 2006; Pavel et al., 2017).

With VR, users see the virtual environment (VE) through the viewport of their headset, which shows a limited region of the VE. VEs can contain objects, spaces, and avatars that are central to the VR app's objective. However, these POIs might be out-of-view at a given time. To address this problem, POI awareness techniques have been developed to help VR users become aware of out-of-view POIs using visualizations. The user is then expected to use head tracking, where the virtual camera is controlled by the headset's position and orientation, to orient to the POI until it is within view (Osmers & Prilla, 2020).

Although head tracking is a convenient way for many people to explore VEs, it can be inaccessible to people with limited or restricted movement of their head or bodies. People with neck injuries or movement disorders such as cerebral palsy or muscular dystrophy might find it challenging or impossible to rotate their heads or bodies in a steady, fluid manner. In addition, situational impairments, which are functional limitations caused by the environment or an

²⁴ Portions of this chapter were published as Franz, R. L., Junuzovic, S., & Mott, M. (2021). Nearmi: A framework for designing point of interest techniques for VR users with limited mobility. *Proceedings of the 23rd International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '21)*, 1–14.

individual's physical state (e.g., sitting down) (Wobbrock, 2019), can make head or body movement difficult.

Because POI techniques assume that users can easily control their virtual cameras to orient to out-of-view POIs, we need to imagine more accessible methods for POI interaction. But how should designers approach the process of creating accessible POI techniques for people with limited mobility? To address this question, I propose Nearmi, a framework that enables designers to create and customize accessible POI techniques for diverse user abilities and VR environments.

7.1 Nearmi

POI techniques designed with Nearmi would enable users to see representations of out-of-view POIs on a display, select the POI they want to view, and ultimately transition towards the POI's direction (Figure 11).




Figure 11. (a) A user in a virtual environment surrounded by points of interest (POIs). (b) The user pulls up a display of out-of-view POIs and selects a POI to orient to. (c) The user's first-person camera rotates to the selected POI. (d) The user faces the desired POI.

The Nearmi framework outlines a design space that can be explored by implementing instances of Nearmi's four components, each of which I discuss below:

- Representation
- Display
- Selection

- Transition

7.1.1 Representation

POIs are represented as cues to signal their existence in a VE. Representations can be visual, auditory, or haptic. Without a Representation, users would be unaware of nearby POIs. Some Representations also convey information about properties of POIs, such as their distance and direction (Baudisch & Rosenboltz, 2003; Gustafson et al., 2008) relative to the user's position. Representations in existing POI techniques include circles (halos) (Baudisch & Rosenboltz, 2003), triangles (wedges) (Gustafson et al., 2008), arrows (Gruenefeld, El Ali, et al., 2017), and picture-in-picture views (Lin et al., 2017). The Representations designers choose for POIs can depend on numerous factors, including the style of the VE and the tasks available in the environment.

7.1.2 Display

Representations of POIs are presented on a Display. Displays act as a canvas for the arrangement of Representations. In prior work, Representations were Displayed by being directly superimposed over the environment (Gruenefeld, El Ali, et al., 2018) or on an object in the environment (Lin et al., 2017). Displays, like Representations, can also encode information about distance and direction of a POI (Lin et al., 2017). Many POI techniques only implement Representation and Display components and assume that the user can orient the environment or themselves towards POIs independently. Nearmi includes Selection and Transition components to automatically reorient users towards POIs of their choosing.

7.1.3 Selection

POI techniques that require little movement (of head or body) must allow users to indicate to the system which POI they would like to orient towards, which is achieved through the Selection

component. A user can select a POI in many ways, through gesture, sound, eye movements, or input devices, to name a few.

7.1.4 Transition

A Transition is needed to orient the user toward a POI after it has been selected, either by repositioning the user or the environment. Transition serves as an automatic reorientation that removes the burden on users to rotate their heads or bodies. Some prior work has explored Selection and Transition in POI techniques. For example, in work by Pavel et al. (2017), users pushed a controller button, and the system automatically oriented their view to the subject of the 3D video. This technique worked without Representation and Display components because it was designed to be used with 3D videos, which usually only have one or two foci in a scene.

Users engage with POI techniques on two fundamental levels: presentation and interaction. All POI techniques must implement both Representation and Display components—the presentation. Most prior techniques presume users can control the interaction by orienting the environment or themselves toward POIs independently. As a result, the degree to which people with limited mobility can interact with POI techniques will be determined by the general accessibility of VR systems. Although it will take time to envision and build accessible VR systems, frameworks like Nearmi can help accelerate the design of accessible techniques for VR applications.

7.2 Benefits of Nearmi as a Design Framework

I identified four benefits of the Nearmi framework for VR users and designers: (1) it foregrounds accessibility for POI techniques, (2) it structures the design process, (3) facilitates evaluation, and (4) affords customizability of POI techniques. I discuss each benefit below.

7.2.1 Foregrounds Accessibility

POI techniques in previous work typically only included Representation and Display components, as these techniques presume users can orient themselves towards POIs. However, the process of orienting to a POI can be challenging or impossible for VR users with limited mobility who cannot easily move their heads or bodies. Nearmi's Transition component ensures POI techniques designed using the framework will be accessible to VR users who cannot easily move themselves to view POIs. However, implementation of all Nearmi components does not necessarily guarantee the POI technique will be accessible. Designers must consider how specific component implementations can affect accessibility as well.

7.2.2 Structures the Design Process

The Nearmi framework provides designers with a blueprint for building accessible POI techniques. Designers can then explore the design space by mixing and matching implementations of Nearmi components while retaining the basic functionality of a POI technique.

7.2.3 Facilitates Evaluation

When evaluating POI techniques, Nearmi can help designers identify the components that might lead to accessibility or usability breakdowns. Once the component (or combination of components) that is causing the issue is identified, designers can swap in an alternative implementation of the component without entirely redesigning the POI technique or altering its core functionality. Nearmi also enables direct comparison of different implementations in an evaluation.

7.2.4 Affords Customizability

Nearmi could be directly integrated as a setting in VR applications. This setting could enable users to try different implementations of Nearmi and choose a combination that best suits their

abilities and personal preferences. Customizability is known to enhance the accessibility of interaction techniques (Anthony et al., 2013; Oh & Findlater, 2013).

7.3 Design Considerations for Nearmi Implementations

As mentioned above, the Nearmi framework alone cannot guarantee that a POI technique is accessible. VR techniques are often built and evaluated based on various design considerations that affect the overall experience of being in VR. I was interested in exploring the relationships between different design considerations and their manifestation as tradeoffs.

I identified design considerations from prior research and best practices for VR design. For example, Wong et al. (2018) found that 25-29% of survey respondents with self-reported disabilities had difficulty raising or moving arms and rotating their upper bodies in VR. They also reported “spatial awareness” as a challenge experienced by their respondents (Wong et al., 2018), and research by Bowman and Hodges (1997) demonstrated that different “task objectives” could significantly impact users’ performance with various interaction techniques. I identified the remaining design considerations including usability, user comfort, realism, and user familiarity with the interaction from a textbook on best practices for designing VR experiences (Jerald, 2015).

7.3.1 Accessibility

Accessibility could potentially impact three components of Nearmi. For example, a Representation must be clearly visible, audible, etc. Some Displays and Selection implementations might not be accessible for people with limited reach, control or stability.

7.3.2 Usability

Usability was a design consideration for every component of Nearmi. Usable designs enable users to form mental models of technology without excessive cognitive load.

7.3.3 Realism

Realism could affect all four components. The Representation could be a duplicate of the POI (more realistic) or a symbol (less realistic). A Display can be the environment itself (more realism), an object in the environment (more realism), or a UI (less realism). Selection can range from touch (more realistic) to ray cast Selection (less realistic). And the Transition to orient the user POI can be smooth, as if a user was turning, or instant as if they immediately reoriented to the POI (less realistic).

7.3.4 Spatial Awareness

It can be important for the user to understand where they are relative to other objects in the VE depending on the task. The Transition component might affect a user's spatial awareness depending on how much of the environment is presented and the duration of the Transition. Representations can be arranged on a Display to convey spatial information (e.g., direction), and POI Representations can include contextual information that might provide awareness of other items or locations.

7.3.5 User Comfort

User comfort is important for most VR techniques. Simulator sickness is well documented when the user moves in the virtual world without a corresponding movement happening in the physical world. Transition implementations that automatically rotate the virtual camera would be more susceptible to inducing simulator sickness, but designers can utilize approaches that reduce simulator sickness during VR camera movements (Sargunam & Ragan, 2018). Selection also impacts user comfort, and the Selection method should not be cumbersome to use.

7.3.6 User Familiarity with Interaction

Many users might be unfamiliar with VR interactions, regardless of their functional abilities, because VR is still emerging as a consumer technology. Therefore, designers can leverage common interaction patterns from other devices, (e.g., joystick panning on a gaming console), to scaffold the process of learning how to interact in VR. Selection would be affected by this design consideration depending on whether the designer imports interaction techniques from other computing paradigms.

7.3.7 Task Objective

There can be VR experiences where efficiency matters more than realism. Designing for efficiency, however, might take away from the richness of the experience or challenge of the task. Therefore, it is important to consider the context in which POIs are used and whether the technique could interfere with the purpose or objective of the experience.

7.3.8 Summary

Most of these design considerations are relevant for VR techniques in general. However, I was interested in how basic design considerations relate to the accessibility of a POI technique. For example, does increased accessibility mean decreased realism? How do some design considerations impact others in implementations of Nearmi? In the next section I discuss how the design considerations above guided the design process.

7.4 Implementation

I used the Nearmi framework to implement an initial set of component instances that exemplified the design considerations above. Although I could have designed any number of Nearmi component instances, I focused on this subset to surface tradeoffs in design considerations.

In this section, I describe the implementation of eleven Nearmi components and how components work together to help users identify and orient to out-of-view POIs.

7.4.1 Representation

I implemented two Representations: a mesh and a real-time view of the POI that I call portal views (Figure 12). Meshes are static, texture-less, colorless, miniaturized versions of the POI mesh that do not provide details about the POI or its surrounding environment. Portal views are live previews that show the POI and some of its surroundings.




Figure 12. The two Representations implemented with Nearmi. Portal view icon of the POI in real time (left). 3D miniature mesh icon of the POI mesh (right).

Mesh and portal-view Representations exemplify the tradeoffs between accessibility and spatial awareness. The mesh Representation could make it easier to identify the POI because it stands out against the background. However, the mesh Representation lacks detail and information about objects around it. Lack of context might make it more difficult to understand where the POI is in relation to other parts of the VE. On the other hand, portal view might not stand out against the background, making it difficult to see, but could provide important contextual information to support users' spatial awareness.

7.4.2 Display

I implemented two Displays: attached-to-object and floating (Figure 13). Attached-to-object Displays, as the name suggests, can be attached to virtual objects or parts of a user's avatar.

Representations float in the scene on a floating Display. The implementation of the floating Display tracks the user's camera to ensure that icons on the Display are always visible.




Figure 13. Nearmi's Display and Selection combinations. Columns from left to right show Selection instances: tap, ray cast, wand, console. Rows from top to bottom show Display instances: attached to object, floating.

These Displays embody a tradeoff in accessibility and realism. The attached-to-object Display simulates interacting with a Display in the physical world, which could enhance realism. However, the attached-to-object Display might be inaccessible if users cannot see the Display in the scene, for example because they cannot hold their arm in view if the Display is attached to their hand. Therefore, I implemented the floating Display so that it was always visible. Constant visibility, however, might reduce realism because it does not have an equivalent in the physical world.

7.4.3 Selection

I implemented four Selection instances: (1) tap, (2) ray cast, (3) wand, and (4) console (Figure 13 and Figure 14). Tap borrows from the touchscreen interaction paradigm by requiring users to “touch” or intersect their controller and a Representation to select it.




Figure 14. Wand: After tapping an icon as shown in Figure 13, the user brings an icon towards his/her headset. A Selection is made when the headset and icon collide.

To operate ray cast, the user points the ray emitting from one of their controllers and intersects it with a Representation. A button press on the controller, typically the trigger, initiates the Selection. The wand Selection (Figure 14) requires users to tap a Representation. The Representation is then transferred to the controller that performed the tap. Users move the Representation toward their headsets and the Selection is made when the Representation and headset intersect. Console Selection enables users to press the controller's joystick left or right to scroll through the Representations. As the user scrolls through the icons, they grow larger when in focus. Users can either scroll to the next Representation or select the in-focus Representation. The user selects a Representation by pressing the trigger button.

I implemented these Selection techniques to illustrate the relationships between multiple design considerations including accessibility, usability, realism, and familiarity with interaction. Tap enhances realism because it has a physical analogue. In addition, users will likely already be familiar with this interaction paradigm because of the ubiquity of touchscreens. However, using tap might be difficult for users who have difficulty reaching or stabilizing their arms.

VR users will likely be familiar with ray cast because it is a common pointing technique used in many VR applications (D. A. Bowman & Hodges, 1997). Another benefit of ray cast is that it enables users to select objects without reaching. However, ray cast might be inaccessible to people with difficulty pointing. With console Selection, users can quickly and accurately select a

Representation, demonstrating usability. Also, because the thumb (or any finger) is the only body part needed to invoke Selection, console can be accessible for users with limited mobility.

However, console lacks affordances, requiring users to recall how to perform the interaction. Lack of affordances is often categorized as a usability issue (Nielsen, 1994a).

Wand enables users to bring Representations near their faces to examine them up close, enhancing visual accessibility. This technique is especially useful for examining mesh icons: individuals can examine meshes from multiple angles. However, wand requires users to reach and move their arms, which can make it inaccessible.

7.4.4 Transition

There are many ways to orient the first-person camera, each of which has tradeoffs. I implemented three Transitions: (1) continuous, (2) discrete, and (3) instant (Figure 15). The continuous camera Transition smoothly rotates the scene around the user simulating what the user would see if they were to turn their head toward the POI. I chose a rotation speed of 30 degrees per second because there is evidence to support that this is a comfortable speed (Sargunam & Ragan, 2018); Still, the speed can be adjusted in the implementation. The discrete camera Transition rotates the camera using configurable angle increments towards the POI. I used 30-degree increments because prior work has shown that it is a comfortable angle to rotate the camera (Sargunam & Ragan, 2018). The instant Transition immediately orients the camera toward the selected POI. The Transition is visualized as a camera cut.




Figure 15. The three camera Transitions implemented in Nearmi. The first-person camera rotates smoothly in the continuous Transition (left), jumps by 30-degree increments in the discrete Transition (middle), and instantly cuts from a starting orientation to face the POI in the instant Transition (right).

I implemented continuous, discrete, and instant Transitions to surface tradeoffs in spatial awareness, user comfort, and realism. Continuous camera rotations can induce simulator sickness (Ragan et al., 2012) but provide information about the environment during the Transition, which can enhance spatial awareness. Discrete Transitions can induce less simulator sickness than continuous Transitions (Sargunam & Ragan, 2018), however they could also be perceived as less realistic because there is no physical equivalent. The instant Transition is the least realistic of the Transitions I implemented, however, it is unlikely to induce simulator sickness (Ragan et al., 2017). Another tradeoff with the instant Transition is that users might not be able to maintain spatial awareness because they do not see the environment during the orientation process, which is possible with continuous and discrete Transitions.

7.4.5 Summary

Nearmi is a framework with four components—Representation, Display, Selection, and Transition—that supports the design of accessible POI techniques for people with limited mobility. I implemented 11 initial instances of these components that reflect tradeoffs in accessibility, usability, realism, spatial awareness, user comfort, interaction familiarity, and task objective. I was interested in exploring the Nearmi design space to understand how different implementations of Nearmi could speak to the design considerations I identified in prior work and best practices for

VR design. With just these 11 instances, users can already create 48 unique Nearmi implementations ($2 \text{ Representation} \times 2 \text{ Display} \times 4 \text{ Selection} \times 3 \text{ Transition}$) demonstrating customizability and extensibility. It would be impossible to build all possible implementations of Nearmi, but I chose to implement this initial set to demonstrate part of the design space and to surface tradeoffs in the design considerations I identified.

7.5 High-level Implementation Overview

To demonstrate customizability and extensibility, I implemented Nearmi components as a set of Unity scripts written in C#. The basic building block of a Nearmi technique is the Display mechanism, that is, an interface that shows POI icons. It is implemented in the `POIDisplay` base class. To render icon content, each Display contains a Representation mechanism implemented in the `POIIconManager` base class (Figure 16).




Figure 16. A high-level overview of the Nearmi framework manager, the mechanism base classes, and their relationships.

To support interaction with the Display content, each Display is coupled with a Selection mechanism, which is implemented in the `POIInputHandler` and `POIInputGenerator` base classes. The generator provides raw inputs, and the handler preprocesses them before publishing them to other mechanisms. The information exchanged between the Display and the input handler is limited to what is required for providing interaction feedback to the user. For example, the Display contains a supporting mechanism for icon Selection visualization, which is implemented in the `POISelectionVisualizer` base class.

The `POIManager` manages the Display itself, including showing or hiding it and providing it with POIs. The manager is also responsible for discovering POIs in the environment and triggering environment updates once the user selects a POI. In particular, it uses events from the Selection mechanism to trigger a Transition mechanism, which is implemented by the `POICameraTransition` base class. The manager can also illuminate the selected POI using the object highlighter implemented in the `POIObjectHighlighter` base class.

The developer is responsible for the following steps to connect an app to Nearmi: (1) calling functions in the `POIManager` to toggle Display visibility, (2) marking POIs in the scene by attaching the `PointOfInterest` script provided by the framework, which enables Nearmi to discover the POIs, and (3) mapping system inputs to the input generator.

7.6 User Interviews

The goal of this study was to identify design considerations and tradeoffs that were important to people with limited mobility when choosing their preferred implementation of Nearmi. I wanted to understand how they thought about tradeoffs and why they were willing to make specific tradeoffs. I conducted a video elicitation study (Glaw et al., 2017) with seventeen individuals with limited mobility. Video elicitation, which is an interview technique derived from photo elicitation, is used to uncover layers of meaning through images, videos, and other visual media. Images elicit emotions, memories, and information in a more detailed and grounded way than conventional interviews (Glaw et al., 2017). Interviews are appropriate for investigating how and why people perceive phenomena (Weiss, 1994). While participants' perceptions are reflective of how they evaluate VR based on video alone, previous research has demonstrated that people with functional impairments often watch other people use the technology, through online videos or friends, to

decide whether it is worth investing in (Wong et al., 2018). Therefore, the results of this study indicate how they might evaluate POI techniques at the pre-adoption stage.

7.6.1 Participants

I recruited seventeen participants through email invitation (average age of 31.0, $SD=9.0$). Four identified as women and 13 identified as men. I compensated participants with a \$50 Amazon gift card. Twelve participants had prior VR experience. Their VR usage ranged from every few days to a few times a year. Overall, they rated themselves an average of 2.1 ($SD=1.1$) in terms of VR expertise with 1 being novice and 5 being expert. By contrast, all participants rated themselves an average of 4.5 ($SD=0.7$) in terms of computer expertise on the same scale. Participants reported a range of conditions that impacted their abilities to interact with computers (Table 6). Their most common challenges were poor coordination ($n=12$), low strength ($n=11$), and difficulty in gripping ($n=10$). They also reported motor difficulties that affected their neck mobility, such as stiffness ($n=11$), low strength ($n=7$), and pain ($n=6$).



Figure 17. The alchemy lab VE used to demo Nearmi. Environment by Unity Technologies²⁵.

7.6.2 Apparatus

I recorded videos of 48 Nearmi implementations ($\{\text{portal view, mesh}\} \times \{\text{attached to object, floating}\} \times \{\text{tap, ray cast, wand, console}\} \times \{\text{continuous, discrete, instant}\}$). Videos averaged 25.7 seconds ($SD=4.3$).

Videos captured a user visualizing and orienting to POIs in an alchemist's lab using different Nearmi implementations (Figure 17). The user oriented towards the POIs in the same order in each video. I added Nearmi implementations to the game using the steps outlined in the implementation section. In this implementation of the attached-to-object Display, I presented icons on a magic amulet because it reflects the game's occult theme. Videos were presented via screenshare on Zoom²⁶, a video conferencing application.

²⁵ <https://github.com/Unity-Technologies/VRAchemyLab>

²⁶ <https://zoom.us/>

Table 6. Participant demographics and preferred Nearmi implementations.

Identifier	Gender	Age	Condition	Preferred Nearmi implementation
P1	M	26	Spinal muscular atrophy	Portal+attached+console+continuous*
P2	M	26	Muscular dystrophy	Portal+attached+console+instant
P3	M	26	Arthrogryposis	Mesh+floating+console+discrete
P4	M	20	Muscular dystrophy	Mesh+attached+console+continuous
P5	W	32	Muscular dystrophy	Mesh+floating+console+continuous
P6	M	24	Cerebral palsy	Portal+attached+ray+continuous
P7	W	36	Muscular dystrophy	Mesh+attached+console+discrete
P8	M	37	Quadriplegia	Mesh+attached+tap+continuous
P9	M	29	Hemiparesis	Portal+attached+console+continuous*
P10	M	24	Muscular dystrophy	Mesh+attached+tap+instant
P11	W	30	Cerebral palsy	Portal+floating+ray+continuous
P12	M	26	Cerebral Palsy	Portal+attached+ray+continuous
P13	M	61	Difficulty moving arms	Portal+floating+tap+continuous
P14	M	26	Cerebral palsy	Mesh+floating+ray+continuous
P15	W	33	Underdeveloped left hand	Portal+attached+console+continuous*
P16	M	33	Cerebral palsy	Portal+attached+tap+continuous
P17	M	34	Spinal cord injury	Portal+attached+console+continuous*

* signifies the single non-unique preference.

7.6.3 Procedure

Participants completed informed consent and demographic questionnaires prior to the study. At the start of a session, the researcher and participant joined a Zoom meeting room with both parties' video and audio enabled. With participants' consent, the researcher turned on Zoom's record feature.

After describing the study, the researcher asked questions regarding participants' use of VR and whether limited head or body mobility impacted their use. The researcher then described the purpose of Nearmi, its components, and how to interact with the implementations. With the webcam, the researcher demonstrated physical interactions for Selection methods with Oculus Rift VR controllers. While describing Nearmi, the researcher supported the explanation with visuals on a slide deck. The researcher then walked through an example video of a person using Nearmi and explained what was happening as the video was playing.

The remainder of the study consisted of four similar phases. In each phase, the participant watched videos depicting different Nearmi implementations. After watching the videos, the

researcher asked the participant about their preferred component instances, as well as questions related to the learnability, perceived ease of use, and perceived usefulness of their preferred combination. If the participant thought the combination was useful, they were asked which scenarios they would want to use it in. It is important to understand perceived ease of use and usefulness, as they can indicate future adoption (Davis Jr., 1986; Venkatesh et al., 2003).

In the first phase, the videos depicted the different combinations of Representations and Transitions (2 Representation $\times$ 3 Transition implementations) while the Display and Selection components stayed the same (attached to object Display with console Selection). In the second phase, the researcher showed the participant eight videos in which the Display and Selection components (2 Display $\times$ 4 Selection implementations) changed while the Representation and Transition instances remained the same. The Representation and Transition implementations were those the participant preferred in the first phase. Then, the researcher showed the participant a video of the participant's preferred Nearmi implementation built using component instances they identified in the first two phases. The researcher asked if the Nearmi implementation was still the participant's preferred combination and how it could be improved. Participants confirmed their preferences after they had seen all the Nearmi component instances in case they wanted to change them based on information participants learned later in the interview. The three phases are summarized in Table 7.

Table 7. A summary of the videos participants watched and the outcome of each phase.

	Representation	Display	Selection	Transition	Videos	Outcome
Phase 1	All	Attached-to-object	Console	All	6	Participant's preferred Representation $\times$ Transition
Phase 2	Preferred Representation	All	All	Preferred Transition	8	Participant's preferred Display $\times$ Selection
Phase 3	Preferred Representation	Preferred Display	Preferred Selection	Preferred Transition	1	Confirmation of preferences

7.6.4 Analysis

The recorded calls were transcribed, and a thematic analysis of the transcripts was performed (Patton, 2001). The design considerations were used as deductive themes and patterns in the transcripts that summarize participants' thoughts on Nearmi's implementations were found inductively. I coded the 162 data extracts from all interview transcripts and organized codes into themes. Then, another researcher and I discussed representative data extracts for each theme to ensure that deductive constructs were interpreted correctly and that data extracts reflected inductive themes.

7.7 Results

The findings can be categorized into three groups. First, I demonstrate the need for accessible POI techniques in VR. Next, I provide user feedback on individual Nearmi component implementations and their combinations. I highlight design considerations that participants discussed for each component. Finally, I present how participants envisioned using Nearmi in VR. When I use the word "immersion", it is because the participant volunteered the term. I could not evaluate immersion because of the nature of the study procedure.

7.7.1 Impact of Mobility Limitations on Virtual Reality Use

I begin with challenges that participants experienced turning their heads or bodies in VR and how these challenges impacted their ability to use VR.

Challenges Related to Head or Body Rotation. Fifteen of 17 participants reported difficulties turning their heads or bodies. Of the 12 participants who had used VR at least once, ten found it difficult to move their bodies or heads while using VR. Six of the 17 participants reported that they did not have difficulty moving their heads but found it challenging to look behind themselves

because it required them to turn their wheelchairs. For P1, this challenge was exacerbated by the fact that the headset obscured his vision of the physical world. P1 said,

“When I’m sitting in a wheelchair and I’m wearing a headset, I have to be careful that I don’t hit any furniture or injure my legs or something. So, I would kind of take the headset, push it up a little, and try to see underneath it before I turn” (P1).

P4 and P6 reported that their wheelchair headrests made it difficult to turn their heads because the VR headset would bump up against the headrest. As P4 said,

“I have a headrest on my wheelchair and so that kind of gets in the way. And then moving my head around with that (headset) and getting like stuck with the headrest on my chair” (P4).

Four participants felt that the VR headset was heavy, which limited their ability to turn their heads. Seven participants had limited neck mobility, neck spasms, neck pain or stiffness that made it challenging to use the head tracking feature. P11 describes this issue and its effect on his experience of immersion:

“I don’t necessarily like have a limited range of motion, but I just have like a lot of stiff neck issues. Like I remember when you try to move this way (turning head to side), it, it just didn’t—it like didn’t feel natural so that takes you out of the immersion.” (P11)

7.7.2 Impact of Limited Head and Body Rotation on Virtual Reality Use

Challenges related to head and body rotation in VR impacted participants’ VR use. In the best case, participants were able to employ workarounds. For instance, they used joystick panning (P3) and slight head turns (P6) or played games where most of the content appeared in front of them.

Since these workarounds are not always possible or desirable, participants sometimes researched games before buying them to “*make sure someone’s (the user’s) head isn’t like jumping around everywhere*” (P3). None of the participants owned a VR system but had used VR through social ties, at arcades, conferences, and in research studies.

Unfortunately, the challenges related to head turning led two participants (P5 and P13) to abandon their VR systems. P5 explained that she was excited to try VR and bought a PlayStation VR system when it was first released. However, many of the games she wanted to play required body or head movement. She eventually abandoned her VR system, saying,

“There wasn’t like a difficulty setting so I couldn’t, like, just say, oh instead I’m gonna do it on a lower difficulty, or you know, just be worse at it, you know. So, if I could have done that, I probably would have changed it and maybe kept playing” (P5).

Participants have found some ways to overcome limited head and body mobility in VR, yet as the data and other work (Mott et al., 2019, 2020) reveal there is still a need for more accessible solutions for this user group.

VR often requires two hands to view, interact with, and navigate virtual worlds. P15 expressed that the challenge for her was not so much moving her head but rather controlling the virtual camera and performing all other interactions with one hand. She explained,

“I mean the biggest issue for me with (controlling) camera angles is not being able to turn my head or stand for the most part, it’s that usually what’s required with (controlling) camera angles is most games have two joysticks and so to be able to do that with one hand, I just can’t [...] So in this scenario what if I did need to look at the pig’s head and then do something with the staff at the same time?” (P15)

7.7.3 Design Considerations

I now present the next level of findings, which are participants' perceptions of Nearmi and its implementations.

7.7.4 Design Considerations: Representation

The ideal choice of POI Representation type appeared to be driven primarily by a tradeoff between the level of visual information (3.6.4 spatial awareness) and challenge of discovering the object (3.6.7 task objective). Nevertheless, participants were almost evenly split between portal view and mesh for their Representation preference.

Ten participants liked the portal view Representation. Their reasons were that they knew exactly what they were looking for in the scene (3.6.7 task objective), they could see what was around the object (3.6.4 spatial awareness), and portal view aligned better with the game aesthetic (new design consideration: VE aesthetic). By contrast, three participants disliked the portal view because by clearly showing POIs, it took away the challenge of discovery (3.6.7 task objective), which they enjoyed. P5 expressed:

“Sometimes I don’t want all the answers given to me. So, there was an element of like a kind of a hide-and-seek thing to that (mesh) as well.” (P5).

Seven participants preferred the mesh Representation. Their reasons were that it maintained the challenge of discovery to find the object in the scene (3.6.7 task objective) and it looked like a miniature object rather than an image, underscoring the 3D nature of VR (3.6.3 realism). Reasons for disliking mesh were the same reasons that some participants liked the mesh: it had an additional layer of abstraction, making it difficult to recognize the object in the scene (3.6.7 task objective).

7.7.5 Design Considerations: Display

Visual accessibility and reachability (3.6.1 accessibility) of the Display were key factors for participants' preferences, along with realism (3.6.3).

Eleven participants preferred the attached-to-object Display because it was easier to see the Representations (3.6.1 accessibility), it felt more "natural" (3.6.3 realism), easier to reach (3.6.1 accessibility), and it integrated with the aesthetic of the environment (VE aesthetic). P3 disliked the attached-to-object Display because the icons were close together, and he predicted it would be difficult to accurately select (unless the console Selection was being used) (3.6.1 accessibility).

On the other hand, six participants preferred the floating Display because they thought it would be more "immersive" (new design consideration: uniqueness to VR), and there was more space between Representations (3.6.1 accessibility). The main reason for disliking the floating Display was that it made it difficult to distinguish the icons from the background (3.6.1 accessibility).

7.7.6 Design Considerations: Selection

Participants' choices for Selection types were motivated by tradeoffs in accessibility (3.6.1) and realism (3.6.3). Overall, they preferred Selection techniques that required the least arm movement (console or ray cast) (3.6.1 accessibility), yet some participants appreciated mechanisms that increased realism (tap or wand) (3.6.3 realism).

Eight participants preferred the console Selection technique. Their reasons were that it required the least amount of arm movement (3.6.1 accessibility), users could rest their arms while selecting (3.6.1 accessibility), it was one-handed (3.6.1 accessibility), and the interaction was familiar because it was like using game controllers (3.6.6 user familiarity with interaction). Five people disliked the console Selection technique. For these participants, being able to move and interact

with their virtual hands was a major reason for using VR over other gaming systems (uniqueness to VR). To reinforce this point, P15 said:

“If you have the controller and it can sense motion, why wouldn’t you want to use it?

If I’m doing VR, I want what makes VR special, I can use a joystick on any other gaming system” (P15).

P8 and P13 evaluated Selection techniques based on the visual animation of the virtual controllers making a selection, assuming the physical selection could be made using gaze input, which was their primary input method. The reason they disliked the console was because the virtual hands did not interact with anything, taking away from the VR experience.

Four participants preferred ray cast. Reasons were that it required less arm movement and reaching (3.6.1 accessibility). They also liked that this interaction technique is commonly used in VR for performing other actions (e.g., selecting objects, locomotion) (3.6.6 user familiarity with interaction). On the other hand, some participants disliked ray cast because it still required some arm movement and coordination (3.6.1 accessibility). Two other participants reported difficulty with aiming in VR. Finally, P8 disliked ray cast because he felt that it did not fit with the theme of the VE and was better suited to a sci-fi environment (VE aesthetic).

Four participants preferred tap. P8 and P13 enjoyed watching the virtual hands interact with objects. P8 said:

“Even though I can’t like physically make that tap motion, I would enjoy having that animation. I just like the aesthetic of the tap because it looks like you’re reaching out to the menu to select the thing” (P8).

In other words, P8 and P13 wanted to watch their virtual bodies to move in ways that their physical bodies could not (uniqueness to VR). Also, P1 liked that the tap interaction was like touchscreen interaction (3.6.6 user familiarity with interaction). However, many participants disliked tap because it required reaching out, which was difficult to do (3.6.1 accessibility).

Only one participant, P16, preferred wand. Most participants disliked wand because it required the most arm movement (3.6.1 accessibility). However, P8 also liked wand because using gaze input, he did not have to factor in arm movement, and it leveraged the essence of VR that makes it different from other technology (uniqueness to VR). He said,

“It kind of seems like another cool opportunity to just like hammer in the fact that you’re playing VR and it’s awesome. Just having another object move towards you”
(P8).

7.7.7 Design Considerations: Transition

The participants traded off spatial awareness (3.6.4), realism (3.6.3), and comfort (3.6.5) when selecting their preferred camera Transition technique.

Twelve participants preferred continuous Transition because it gave them a better awareness of their environment (3.6.4. spatial awareness) and felt more “immersive” and “natural” compared to the other Transitions (3.6.3 realism).

Two participants preferred the discrete Transition because it was familiar (3.6.6 user familiarity with interaction) and seemed faster (3.6.2 usability) than the continuous Transition. Most participants disliked it because they felt it was “jerky”, “robotic”, and “unnatural” and reminded them of video lag (3.6.3 realism).

A couple of participants preferred the instant Transition because it was the most direct Transition and was faster than the other Transitions (3.6.2 usability). On the other hand, some

participants disliked the instant Transition because they felt that they were missing out on important spatial information (3.6.4 spatial awareness), and it was less “immersive” than the other Transitions (uniqueness to VR).

7.7.8 Design Considerations: Display × Selection Combination

A few participants pointed out that an interaction between the Display and Selection components affected the overall accessibility of the Nearmi implementation. In particular, the floating Display with the tap or wand interaction would have been inaccessible to people who find it challenging to perform reaching movements (3.6.1 accessibility). P3 explained that he would have chosen the attached Display because it was more accessible than the floating Display if he was limited to using tap or wand Selection. He said,

“The floating versus attached thing became simply a matter of like personal visual preference and contrast, whereas if I was considering like a tap and wand then I would have had a total different reason for preferring attached which is I simply couldn’t reach them (the floating icons)” (P3).

On the other hand, P17 pointed out that the attached to object Display had less space between icons, making it difficult to accurately select an icon (3.6.1 accessibility). In addition, P8 expressed that he would use the floating Display and ray cast Selection combination because it would work well with gaze input (new consideration: input device). Therefore, the combination of Display and Selection implementations can have unique accessibility implications that are not apparent based on their individual components.

7.7.9 Design Considerations: Accessibility Via Customizability

While participants considered similar tradeoffs in accessibility (3.6.1) and other factors, they nevertheless created 13 unique preferred Nearmi variants, highlighting the need for

customizability. In fact, five participants explicitly mentioned this customizability is important for accessible design. For instance, P5 said,

“When you have a physical disability, it’s like, you just have to have options. Having this kind of flexibility would definitely help in all kinds of situations in VR” (P5).

P2 also illustrated the need for alternative input methods when a user cannot move in VR traditionally:

“It (Nearmi) makes it a lot more accessible for people that can’t even move around traditionally. So, it’s not going to be like you’re stuck to have to move your arms and things. You’re able to just use a joystick” (P2).

Moreover, P4 highlighted the need for options because different people have different types of mobility:

“I know people who have strong arm movement, but not so much with dexterity. They might do just fine with a tap. Even though we both have limited mobility, people have different types of mobility” (P4).

P11 and P16 had a combination of perceptual and motor impairments reinforcing the need to design for diverse abilities. They assessed Nearmi components based on whether they would provide a frame of reference to support their spatial understanding.

7.7.10 Applicability of Nearmi to Real Scenarios

To investigate participants’ potential adoption of Nearmi implementations, I asked for feedback regarding their expected learnability, perceived ease of use, and usefulness (Davis Jr., 1986; Venkatesh et al., 2003). All participants believed Nearmi would be easy to learn to use. All but three participants expressed it would be easy to use because the interaction felt “natural” and

“intuitive” (3.6.3 realism), and they were already familiar with console interaction (3.6.6 user familiarity with interaction). The remaining participants expressed that they thought the Nearmi implementations would be easy to use if their preferred input method was compatible (input device), the VR controllers were easy to use (3.6.1 accessibility, 3.6.2 usability), and if it integrated well with other interaction techniques for locomotion and object manipulation.

All participants said they would use Nearmi implementations in VR except P15, who said it would depend on the task (3.6.7 task objective). Participants envisioned using Nearmi implementations in a variety of scenarios, including environment exploration, social VR, games, and 360° videos. P9 thought Nearmi would work well in an escape room:

“(Nearmi would be useful) like where you need to solve puzzles or whatever, that would be something where this would be—being able to continuously move the camera and being able to see the whole room rather than just teleporting from one angle to the next, that (Nearmi) would be very helpful. I mean, personally, if there was a game that came out, if there was like an escape room with this (Nearmi), like, I would probably go buy a VR right now, so.” (P9)

P1 explained that Nearmi would be especially useful during game startup:

“Nearmi would be useful in like your (virtual) home kind of setting, and you just want to equip stuff before you head out into the world. And this (Nearmi) would just make it kind of convenient because it’s a known location. There’s no point in game mechanics and trying to make you look around to find stuff” (P1).

While some participants felt Nearmi would be most suitable for slower paced, exploratory VR experiences, A handful of participants could see Nearmi in faster-paced, first-person shooting games where it could be used to auto target:

“It (Nearmi) would be super helpful in a shooting game. I can see it being integrated into that really nicely where it kinda just like auto targets for you if that’s the difficulty level you want it to be at” (P5).

7.7.11 Summary

The findings highlight the need for accessible viewing techniques for exploring VR environments for people with limited mobility. As a first step, I gained insight into participants’ decision-making process when evaluating Nearmi and its component implementations. Although participants had similar justifications for liking or disliking implementations, they created 13 unique Nearmi variants (Table 6), highlighting the need for customizability.

7.8 Discussion

I discuss design considerations that were important for participants when they identified their preferred Nearmi implementation, the need for customizability, and the tension between standardization and customization in the design of VR techniques. I also discuss the need to establish criteria for evaluating accessible VR techniques.

7.8.1 Design Considerations Identified by Participants

Participants discussed most of the design considerations in the form of tradeoffs. One of the main tradeoffs they considered was between accessibility and realism. With a few of the Nearmi implementations, particularly implementations of the Display and Selection components, participants expected that realism would be compromised by increased accessibility. For example, tap was perceived as more realistic than console or ray cast, but it also required that the user reach

and stabilize their arm to perform the Selection. Although accessibility and realism were identified as a tradeoff for some individuals, this tradeoff did not exist for others. For instance, P8, who assumed that he could use gaze input, preferred tap because it was realistic and accessible if he could trigger the tap animation with his gaze. Therefore, tradeoffs were made based on participants' individual abilities and preferences.

Although participants discussed most of the design considerations that I identified before designing the Nearmi implementations, they also identified new ones that I did not consider, including (1) input device, (2) VE aesthetics, and (3) uniqueness to VR. P8 and P13 evaluated techniques in terms of their input device of choice, an eye tracking device. The Nearmi framework does not assume a particular input device and so its implementations should factor in the use of diverse input methods. Participants in Mott et al.'s (2020) study reported being open to using voice and gaze input as alternatives to controller and motion input in VR, indicating the potential for alternative input devices to be desirable. With regards to the VE aesthetic, some participants commented that the aesthetic of some component implementations was not consistent with the VE theme that I used in the videos. For example, even though one participant felt that the ray cast Selection would be accessible, he thought the sci-fi aesthetic of the technique clashed with the occult theme of the VE.

Some participants preferred components that had qualities that were unique to VR. One appeal of VR is that you can do things that you cannot do in real life and experience it in a similar way. P8, who had quadriplegia, wanted to watch an animation of his virtual hands not because it emulated reality, but because it allowed him to experience an alternative reality, which would not be possible with other technologies.

7.8.2 Need for Customizability in the Design of Virtual Reality Techniques

The findings highlight the need to consider the diverse abilities possessed by people with limited mobility. For example, I could have limited the Selection techniques to those that required the least arm movement, but as P4 explained, there are individuals with strong arm movement but limited dexterity, making interacting with joystick or controller buttons challenging. In addition to diversity within a user group, individuals can experience different mobility throughout the day and over the long term. As a result, customizability is not only important for accommodating different individuals, but a single individual at different times.

Although the Selection techniques relied on a VR controller as an input device, it is important to note that alternative input methods such as voice and gaze could also be implemented as either input to trigger controller animations (as P8 and P13 assumed) or as new Selection implementations. For example, a gaze-based Selection implementation could require users to dwell on POI icons to trigger the Transition. Developers could create Nearmi implementations that take advantage of eye tracking, voice recognition, and other input methods once they become commonplace on commodity VR devices.

The findings illustrate that accessible VR hardware is not enough, software should also be accessible, which can be achieved through customizable design. In P15's words:

“I would hope that tech companies not only develop controllers and adaptive parts and stuff, but also think about if we are making a game, is it inclusive?” (P15).

7.8.3 Customization and Personalization vs. Standardization in Virtual Reality

VR is in an early phase as a mainstream consumer technology. As a result, hardware differs across VR systems, and interaction techniques differ across apps. The lack of standardization across hardware and software is likely related to the fact that VR experiences can be evaluated

based on so many criteria that an optimal interaction technique depends on contextual factors (e.g., the VE, the goal of the app, etc.). In VR, design considerations like comfort and immersion are as important, if not more, than speed and accuracy.

Although standardization was a goal for technologies in the past, personalization now makes it possible to dynamically select interaction techniques based on users' characteristics. Numerous VR interaction techniques have been designed in the past 30 years and the number of available techniques will continue to expand as researchers continue to develop new approaches (Al Zayer et al., 2020). VR could benefit from the diversity and quantity of interaction techniques afforded by customization and personalization, rather than focusing on standardization. Accessibility could then be another design consideration alongside realism and comfort, for example, that would factor into selecting an appropriate technique. As a result, VR techniques might not need to be modified for the sole purpose of accommodating users' functional impairments as assistive technologies have done in the past (Wobbrock et al., 2011). Instead, there could be numerous reasons for a user to adopt an interaction technique, one of them potentially being accessibility. This paper advances the potential of customizability to guide future VR interaction design.

A design challenge that the use of personalization presents is introducing customizable techniques, like Nearmi, in a way that users can create their preferred implementation without going through multiple permutations. Options to address this issue include gameplay-based tutorials that ask users to perform specific actions and then recommend options based on their performance, previews that provide users glimpses of how interactions would work without having to try options themselves, and control presets that users can save and share with others (for example, the gaming platform Steam²⁷ allows users to save and share controller configurations for

²⁷ <https://store.steampowered.com/>

different games). There are ongoing efforts in the gaming community to enable customization, such as the more than 60 accessibility options offered by *The Last of Us Part II*²⁸ (i.e., for players with low vision, deaf and hard of hearing players, and players with impaired fine motor control).

7.8.4 How Do We Evaluate Accessible Virtual Reality Techniques?

If personalization and customization are goals for future VR systems, designers must establish a set of criteria for evaluating interaction techniques. Doing so would make it possible to match a user's preferences and abilities with the appropriate interaction techniques. For example, if a user was not sensitive to simulator sickness, they could use techniques that were highly immersive compared to more comfortable techniques that provided less immersion. Although established guidelines for evaluating accessible VR technologies do not yet exist, the design considerations that were surfaced in this study could be applied to any VR interaction technique (e.g., locomotion, object manipulation, etc.), not just POI techniques. This study is a first step towards identifying factors that are important to users when evaluating accessible VR techniques.

This study reveals that not only do we need to understand individual design considerations, we also must understand the relationships between considerations in the form of tradeoffs. Al Zayer et al. (2020), identified design considerations that were consistently related for different types of locomotion techniques. However, accessibility was not one of the factors considered. Therefore, the tradeoffs identified in that paper might not hold for techniques that are designed to be accessible, which highlights the need to establish a set of tradeoffs when accessibility is thrown into the mix. I hope that designers will not only continue to invent accessible VR interactions, but also discuss the rationale for the criteria they choose to design and evaluate their technique with.

²⁸ <https://www.playstation.com/en-us/games/the-last-of-us-part-ii-remastered/>

In this way, our community can move towards a better understanding of how to make VR accessible and engaging at the same time.

7.8.5 Limitations and Future Work

I designed the study so that it could be conducted remotely due to the COVID-19 pandemic, even though Nearmi implementations were ready to be used by participants. While the results are still valuable for understanding how people with limited mobility would evaluate accessible techniques before trying them in VR, a user study in which participants complete tasks in VR would reveal in-depth information about the tradeoffs I identified in this study. People with limited mobility are aware of their physical abilities, so I was pleased to elicit helpful feedback on what interactions might be inaccessible for various reasons (such as difficulty performing mid-air pointing). However, some characteristics of VR use, such as the size and weight of VR controllers, the strength needed to push buttons, and the effects of different Transition instances, cannot be adequately conveyed in videos alone. However, because people with mobility impairments watch other people use VR before investing in it (Wong et al., 2018), it is still critical to understand how they evaluate VR techniques in the pre-adoption stage.

Although I explored how Nearmi can improve accessibility for people with limited mobility, this framework does not exclude its potential applicability to POI techniques for other populations. Future researchers might apply Nearmi to different accessibility contexts using audio or haptic instances of Representation and spatial audio or tactile layouts for Display instances.

Finally, engaging with practitioners is important to bridge the “research-practice gap” (Norman, 2010). Future work can facilitate developers’ use of the framework by creating a toolkit with POI prefabs that developers can place in the VE. Conversations with developers can indicate the practicality and potential usefulness of this approach.

7.9 Summary

I presented Nearmi, a design framework for creating accessible POI techniques for people with limited mobility. I made two additional contributions (1) a set of design considerations identified by participants when choosing their preferred Nearmi implementations (2) participant feedback on the subset of Nearmi components that were implemented. My findings highlight the tradeoffs that participants made when choosing a preferred implementation of Nearmi, and they suggest the need for customizability to accommodate diverse abilities and preferences. This study contributes to a growing research community that aims to make VR usable people with disabilities. As adoption of VR increases, our community not only needs to develop accessible interaction techniques, but also establish the key criteria for evaluating them.

8 A PATTERN OVERLAY SHADER FOR VIRTUAL REALITY USERS WITH COLOR VISION DEFICIENCIES

Scene-viewing techniques can improve the accessibility of virtual reality (VR) for people with color vision deficiency (CVD) in addition to limited head or body movement. CVD affects 4% of the world's population and is a condition in which people cannot differentiate particular colors (Color Blindness, 2023). There are three main types of CVD: protanopia and deuteranopia affect the perception of red and green, and tritanopia affects the perception of blue and yellow.

CVD can affect a person's ability to use technology. For example, many user interface (UI) elements are colored red and green to signal negative and positive actions, respectively (e.g., an "answer call" button is usually green and a "decline" call button is usually red). Researchers have addressed the issue of colors being indistinguishable on UIs by automatically recoloring inaccessible colors (Flatla, 2012). However, recoloring can make other colors look strange and introduce new accessibility problems by making other colors indistinguishable.

Researchers have also explored applying patterns to static images. Sajadi et al. (2013) investigated how the orientation of hatching lines can be used to differentiate hues (Figure 18). Applying a pattern to an object can be a general solution because it does not alter underlying image and will not create new color accessibility issues.




Figure 18. Researchers examined how the orientation of hatching lines can be used to differentiate colors.
Figure from Sajadi et al. (2013).

While applying patterns could be a useful solution, it has not been explored in the context of immersive environments, such as those in VR and augmented reality (AR). Pattern application could be a better solution than recoloring because realism and immersion are important to the experience of virtual environments (VEs) (Di Luca et al., 2021). Recoloring could make the environment look unnatural and unrealistic. Also, recoloring VR or AR objects could introduce more work for a designer because they would have to double check and correct the accessibility of recolored objects. Pattern overlays could be a more general solution than recoloring.

8.1 Theory

To explore using pattern overlays to make colors accessible in VR, I used knowledge of color theory and CVDs to prototype a graphics shader that applies patterns over indistinguishable colors. The xyY color space is a mathematical representation of colors used in the field of color science and is particularly useful for the study of human color perception (Figure 19) (Kerr, 2010).




Figure 19. Different color vision deficiencies have different copunctal points, where the confusion lines converge in xyY space. Left: Deuteranopia (red-green CVD) has a copunctal point outside of the color space. Middle: Tritanopia (yellow-blue CVD) and, right: Protanopia (red-green CVD) have copunctal points on the edges of the color space.

Protanopia, deuteranopia, and tritanopia can be modelled in the xyY color space by drawing lines, called confusion lines, over the color space. The colors under the confusion lines are indistinguishable to people with CVD. The position of the copunctal point, which is the point at which confusion lines converge, varies depending on the type of CVD (Figure 19). Accessibility problems occur when colors are different hues on a confusion line, but not different shades of the same hue. For example, light blue and dark blue can be distinguished correctly by someone with protanopia, but not magenta and light blue (Figure 20).




Figure 20. On a confusion line, accessibility problems occur when there are two different hues, but not when there are different shades of the same hue.

The line that separates indistinguishable colors can be modelled as a circle with its center at the copunctal point and its radius at the white point, which is at the center of the xyY color space

(Figure 21, left). A pattern applied to colors on either side of the separation line indicates that they are different hues (Figure 21, right).




Figure 21. Left: The white point in the xyY color space (labelled E) marks the transition between warm and cool hues. Right: The separation line is a circle with the center at the copunctal point and a radius the distance between the center and the white point. In this example, I applied a pattern to the left region of the separation line to make reds distinguishable from greens for people with protanopia.

8.2 Implementation

To implement the pattern overlay effect, I prototyped a post-processing shader that is applied at the very end of the rendering pipeline. The shader takes the final rendered image and converts the value of each pixel to xyY from RGB. It then checks which side of the separation line the pixel is on. If the pixel is to the left of the separation line, no change is made to the pixel. If it is to the right, it is multiplied by a pattern pixel at the same coordinates to produce a pattern overlay effect.

While creating the shader, some design questions surfaced such as: Is one pattern sufficient? How much of the color space should be covered? How would the user change visual properties of the pattern? I implemented an interface to explore some of these questions. The interface enabled the user to indicate their CVD type, the patterns they wished to use, and how many to use to differentiate multiple hues (Figure 22).




Figure 22. Three patterns being applied over different hues for people with left: deuteranopia (red-green CVD), middle: protanopia (red-green CVD), right: tritanopia (yellow-blue CVD).

It also enabled the user to change the visual properties of the pattern including the scale, opacity, and brightness. Finally, the interface allowed the user to choose the side of the color space to apply patterns over (to the left or right separation line) (Figure 23, left) as well as adjust sensitivity of the shader by changing radius of the circle to cover more or less of the color space (Figure 23, right).




Figure 23. The user could edit properties of the pattern including left: which side of the color space it covered and, right: how much of the color space was covered, or in other words, the “sensitivity” of the shader.

1.1.1.1 Use Cases

I explored a few use cases in which somebody with CVD might experience accessibility issues in VR or AR to determine if pattern overlays could be a feasible solution.

In the first scenario, the user is watching a sports game in which the two teams have red and green jerseys, making it difficult to distinguish the teams (Figure 24, top). The user could overlay a pattern onto the red jerseys to make it clear which players belonged to their respective teams (Figure 24, bottom).



Figure 24. Top: A person with protanopia's view of the sports teams. The darker jersey is red, and the lighter jersey is green. Bottom: A person with protanopia's (red-green CVD) view of the sports teams with a pattern overlaid on the red jerseys.





Figure 25. A scenario in which a user with red-green CVD is looking for apples. Top: It is difficult for the user to immediately identify the red apples. Bottom: With the shader applied to reds, it is easier to identify the apples.

In another scenario, a user is shopping for produce at a grocery store. It could be difficult for them to differentiate fruit (Figure 25, top). It could also be difficult to detect the ripeness of a fruit based on its color, since some types of fruit are green when they are not ripe. I overlaid the pattern onto reds to help the user differentiate pears, which are green, from apples, which are red in the virtual scene (Figure 25, bottom).

I also demonstrated how the shader could be used in AR. In this scenario, the user is shopping, and they want to find a red skirt or dress. Laying the clothes side by side, they look like different shades of green (Figure 26, left), but when they use the AR app, they can see that the skirt is red because the user configured the app to put patterns on reds (Figure 26, right).



Figure 26. Left: An AR view of a red skirt and green dress from the perspective of someone with protanopia (red-green CVD). Right: The user's view of the same garments with a pattern overlaid on the red skirt (the white dots are the part of the AR surface-tracking capability).

8.3 Limitations and Future Work

These are a few of many examples of how the shader could be used to make both virtual and physical reality more accessible for people with CVD. However, a limitation of the shader is that the pattern appears to move as the user's device moves. In other words, it does not "stick" to the object it is applied to. This challenge could affect a user's sense of presence in VR, because objects' surfaces appear to be moving when they should not, which could detract from the realism of the scene. Future work could examine the effect of the pattern overlay compared to the effect of recoloring objects on a user's sense of presence to investigate if the pattern overlay effect lessens the illusion of being immersed in a virtual space.

8.4 Summary

It can be difficult for people with color vision deficiencies (CVDs) to detect colors in physical reality, which can affect their technology use. Researchers have explored changing the colors of UI elements and applying patterns to static images, but applying patterns to immersive environments has not been investigated. I created a graphics shader that overlays patterns onto hues that are difficult for people with different types of CVDs to distinguish and enabled users to change visual properties of the shader. I also explored scenarios in VR and AR in which the shader might be useful, including differentiating team sports jerseys, easily identifying produce in a grocery store, and shopping for clothing.

9 EXPERIMENTVR: AN EXTENSIBLE TESTBED FOR DESIGNING AND AUTOMATING VIRTUAL REALITY STUDIES

As virtual reality (VR) is becoming a more popular consumer technology and new interaction techniques are being developed, there is a need to evaluate interaction techniques to establish if they improve upon the current state-of-the-art. Typically, researchers build their own testbeds to compare users' performance with various interaction techniques. Because researchers evaluate VR interaction techniques using similar tasks and measures (Bergström et al., 2021), a pre-built testbed that can be configured to support various experimental designs and tasks would eliminate the effort needed to build testbeds from scratch.

Past efforts building reusable testbeds to help researchers evaluate VR techniques were focused on a single type of interaction technique (e.g., object manipulation and locomotion) and did not allow for different types of tasks (Bowman et al., 1999; Cannavò et al., 2021; Poupyrev et al., 1997). Cannavò et al.'s (2021) locomotion testbed used pre-defined tasks, such as avoiding moving obstacles or travelling in a straight line, and corresponding measures, such as the number of obstacles hit and linearity of the user's path, were recorded. Because of this design, researchers could not define their own tasks, environments, or measures.

There is a need to standardize the evaluation of interaction techniques in VR to move the field forward, which can be achieved through testbeds (Bowman et al., 1999). However, testbeds must enable the evaluation of different types of VR interaction techniques and allow researchers to customize the task, environment, and data collected. I sought to address this need by developing an extensible, customizable testbed, called Experimentvr, that allows VR researchers to design studies with multiple independent variables, in which the main goals are to compare interaction techniques and log high-level performance data (e.g., trial completion time) and low-level controller and headset data (e.g., device position).

9.1 Design Considerations

Experimentvr is based on three basic uses that are common to many VR evaluation studies in HCI (Bergström et al., 2021). The testbed should support VR researchers in: (1) Designing and automating a between- or within-subjects factorial study design, (2) comparing different interaction techniques, and (3) recording performance data and low-level data from the VR controllers and headset while users perform tasks. I designed Experimentvr to incorporate three design considerations: (1) customizability, (2) extensibility, and (3) ease of use. These three design considerations are intended to guide the design of the testbed to enable researchers to implement their own study designs while leveraging a pre-built tool. Below I specify the attributes Experimentvr should have to adhere to each of the design considerations, which are based on knowledge of user studies in HCI (MacKenzie, 2013b, 2013a) and a survey of studies evaluating VR object manipulation techniques over the last 20 years (Bergström et al., 2021).

Extensibility. A testbed that demonstrates this design consideration should enable the user to design their own studies without having to modify the core code. To enable extensibility:

1. A researcher should be able to define a study with any number of factors and levels per factor.
2. A researcher should also be able to define the number of trials and trial block repetitions.

Customizability. A testbed that demonstrates this design consideration should enable a researcher to modify it to meet the requirements for the study. To enable customizability:

3. The testbed should support within- and between-subjects designs.

4. The researcher should be able to configure the testbed to present different types of tasks that are common in VR studies, including timed tasks, locomotion tasks, and selection tasks.
5. A researcher should be able to easily port their own interaction techniques into the testbed.
6. The test environment should be configurable such that a researcher can define countdown timers and instructions, as well as target objects, sizes, order, and positions.
7. A researcher should be able to define the log file format and sampling rate for recording data.

Ease of use. A testbed that demonstrates this design consideration should make it simple and straightforward to design and configure a study. To improve ease of use:

8. The testbed should eliminate the need for a researcher to code the testbed behavior by providing modular building blocks to implement a study design.
9. The testbed should automate the study procedure including training, presentation of instructions, countdowns to trial start and end, randomization of conditions, and repetition of trials.

9.2 Testbed Design and Architecture

To address these design considerations, I developed a VR testbed composed of several key components: scenes, conditions, trial blocks, trials, and targets. Researchers can utilize these elements to construct a comprehensive study design.

- **Scenes:** Each scene represents a specific interaction technique. Researchers can create a scene for each interaction technique they wish to test.

- **Conditions:** Within a scene, conditions can be defined to test the interaction technique under different scenarios.
- **Trial Blocks and Trials:** Each condition includes a user-defined number of trials, which can be organized into trial blocks for repeated testing.
- **Targets:** Within each trial, researchers can specify a target object.

Besides these main components, there are also ancillary components such as the `ConditionManager`, which keeps track of and randomizes conditions, the `ExperimentManager`, which tracks the experiment state and passes data to the `Logger`, and the `Logger`, which records event data from the scenes and pulls values from the controllers and headset at a user-defined sampling rate. (Figure 27).

The nested design of `Experimentvr` allows for the randomization of techniques and conditions while enabling users to extend the number of study components. Each study component tracks the component within it, facilitating extensibility. This architecture supports the creation of experiments with both multi-factorial and full factorial designs. Additionally, it automates the study procedure by presenting scenes and conditions, thereby streamlining the experimental process.




Figure 27. The architecture for Experimentvtr. In this example, the study is a 2x2 design with one factor being Interaction Technique and the other being Target Distance. There are two scenes, one per technique with two conditions in each scene. Each condition has four trials in which the participant hits a target with one of the interaction techniques. The four trials are repeated three times as a trial block, for a total of 24 trials per technique. Low-level controller and headset data as well as performance data are recorded.

9.3 Experiment Setup: An Example

To illustrate the capabilities of the Experimentvtr, I describe a hypothetical experiment. In the experiment, a researcher conducts a 2x2 within-subjects study, where one of the factors is Interaction Technique with two levels: Technique A and Technique B, and the other is Target Distance, also with two levels: Near and Far (Figure 27).

To implement the study in Experimentvtr, the researcher creates three scenes, which comprise the study environment. The first scene is the starting point of the experiment and is where the `ExperimentManager` and `Logger` objects are initialized. The starting scene also acts as a training environment for the participant, in which they can practice using the controllers.

The remaining two scenes are the testing environments for the interaction techniques, one scene for Technique A and one for Technique B. Within a scene, the researcher creates a `ConditionManager` object, which randomizes and keeps track of the study conditions. The `Target Distance` factor has two levels, so the user adds two `Condition` objects to the `ConditionManager`, one for the Near condition and one for the Far condition. Each `Condition` contains a reference to a `TrialBlock` object, which repeats the trials by a user-defined number, which is three repetitions in this scenario. In this example, the participant experiences four trials, so the researcher adds four `Trial` objects to the `TrialBlock` object (Figure 27). `Trials` can then be configured to enable different types of tasks, including timed, locomotion and selection tasks, which are discussed next.

9.4 Tasks

A `Trial` consists of a task that the participant performs in VR. The testbed supports three different types of tasks: timed, locomotion, and selection. To implement each type of task, the user would attach the `TimedTrial.cs`, `LocomotionTrial.cs`, or `PointingTrial.cs` scripts to game objects in a scene.




Figure 28. The researcher’s view of the testbed environment while the participant is completing a timed task called “Dual wrists flex down hold.” The trial instructions are to “flex both wrists downward and hold” and the participant has 13 seconds before the trial countdown elapses. Instructions only visible to the researcher are displayed in the top left and right corners. Information about advancing through the study as well as detailed information about the study progress is displayed in the top right corner.

9.4.1 Tasks: Timed

During a timed task, the user is instructed to perform an action (e.g., flex both wrists downward and hold) for a certain amount of time as low-level data from controllers and headset are logged. The researcher can also define the duration of the timed task and instructions presented to the participant (Figure 28).

9.4.2 Tasks: Locomotion

In a locomotion task, the participant navigates to a target in an environment using a locomotion technique (Figure 29, left). The testbed currently implements nine locomotion techniques that can be used by a researcher including, Teleport, Grab and Pull, Hand Rockets, Chicken Acceleration, Sliding Looking, Arm Swinging, Throw Teleport, Astral Body, and Cloudstep.




Figure 29. Left: A view of the participant’s controllers (floating black objects) in the locomotion task environment. The participant is currently using the Teleport technique. The locomotion targets are the blue and green cylinders in the background. The green target is the current target that the participant must navigate to. Right: The participant’s view while they are inside the current target. The blue “7” is the trial countdown and the pink “2” is the dwell countdown.

A single locomotion trial involves the participant using a locomotion technique to navigate to a target in the environment. The researcher can define the target object, size, order, and position by placing their desired target objects in the environment (one per `Trial` object) and attaching the `LocomotionTarget.cs` script to the game objects. A configurable trial countdown appears in front of the participant as they navigate to the target, which indicates the time remaining in the trial.

When the participant’s camera collides with the target, a configurable dwell countdown appears in front of the user and starts counting down (Figure 29, right). If the participant exits the target before the dwell countdown elapses, the dwell countdown stops and disappears from view, but the trial countdown continues. If the participant reenters the target, the dwell countdown reappears and restarts.

The target is considered “hit” if the user stays inside the target until the dwell countdown elapses. Otherwise, if the trial countdown elapses before the target is hit, the trial is counted as a “miss.” The current implementation of the locomotion trial also eliminates the need for the participant to visually search for the next target by rotating the participant’s camera to face the next target as soon as a trial is completed.

The locomotion task also enables the researcher to place obstacles in the environment and log obstacle collision events by attaching a script called `ObstacleHit.cs` to the obstacle. `ObstacleHit.cs` creates an event that is bubbled up to the `ExperimentManager` when the XR camera collides with an obstacle.




Figure 30. A participant using the ray cast technique to select a target in a Selection task. The dwell countdown shows that the participant must keep the ray on the target for one more second before the trial is considered a “hit”.

9.4.3 Tasks: Selection

In the current implementation of the selection task, only one technique is implemented, which is ray cast, where a user points at and selects a target using a laser-type line emitting from their controller (Bowman & Hodges, 1997). A single trial in a selection task involves the participant using ray cast to select a target in the environment. In the implementation of the selection task, the target changes color and a configurable dwell countdown appears when the user points at a target (Figure 30).

The behavior of the dwell and trial countdowns are the same as those in the locomotion task when the participant’s ray exits and reenters the target. The target is considered “hit” if the ray remains on the target until the dwell countdown elapses and is considered “missed” when the trial countdown elapses before the dwell countdown expires. Like in the locomotion task, the user can

define the target object, size, and position in the environment by attaching a `PointingTarget.cs` script to target objects.

9.5 Data Logs

Data are logged for each trial by the `Logger` object. The researcher can select either CSV or JSON format for the log file, as well as define the sampling rate in milliseconds. For all task types, basic information about the study is logged including the timestamp, condition name, trial block repetition, trial number, and the trial instructions. The `Logger` also records the start and end time of the trial block, trial, and dwell countdowns. In terms of the device data, the position (x,y,z), rotation (x,y,z), velocity (x,y,z), angular velocity (x,y,z), acceleration (x,y,z), and angular acceleration (x,y,z) are recorded. Controller button information including whether a button is being pressed (bool), the thumbstick position (x,y), as well as grip and trigger pressure (float) are also recorded. Finally, the start and end positions (x,y,z) of invisible rays emanating forward from the center of the headset and controllers are also recorded.

Data described above are recorded for all task types. In addition to these data, task-specific data are also recorded.

- **Locomotion:** These data include the current target name (string), the time at which the participant enters or exits the target (string), whether the target was hit (bool), and times obstacles were hit (string).
- **Selection:** Data logged for the selection task are the same as those for the locomotion task with the exception of obstacle-related data.

The `Logger` object begins recording at the start of a trial block repetition and stops recording when a trial block repetition is complete. One log file is created per trial block repetition, so in the example study above, 12 log files would be created (3 trial block repetitions $\times$ 4 conditions).

9.6 Running the Experiment: An Example

To run the example experiment described above, the researcher would run the program in Unity. The experiment would start with the training scene and display training instructions. In this scenario, the user would learn how to snap turn by pressing the thumbstick.

The experimenter would then advance the study to the next scene by pressing the space bar. The next scene and conditions would be randomized, so the testbed would present Technique B first. The experimenter would then progress the experiment to a practice environment for Technique B in which the user could practice the using controls for the technique. Once the participant finishes practicing with the technique, the researcher would advance the procedure to the first condition, which contains target distances that are Near the user. A configurable countdown would be displayed in front of the user in the virtual environment (VE) and would indicate how much time is left until the first trial in the trial block would begin.

Once the first trial ends, the next trial would automatically start. After all trials are completed, the researcher could then progress to the next trial block repetition, condition, or scene by pressing space again until all repetitions, conditions, and scenes are completed.

9.7 Implementation

I implemented the testbed as a set of C# scripts in Unity because Unity is commonly used by researchers to implement VR techniques and environments. To implement a study design, the researcher would create their desired number of scenes and attach the testbed scripts (e.g., `ConditionManager.cs`, `TrialBlock.cs`, etc.) to empty game objects within the scenes. If the researcher wants to evaluate a novel locomotion technique, they would attach their own script to the camera object to implement that technique. The script would define how the user's

camera performs a virtual translation by mapping VR device inputs to Unity's XR interaction system.

9.8 Experimentvr Pilot Test

I used the testbed to conduct a study in which I evaluated six locomotion techniques with 40 participants. The testbed facilitated the process of designing and automating a study that took two to two-and-a-half hours to conduct, which I present in Chapters 10 through 13. The data logged were of high quality, enabling me to analyze high-level performance metrics (e.g., trial time, hit rate, and obstacles hit) as well as low level device data (e.g. position, rotation, velocity, etc.). A total of 240 log files (40 participants $\times$ 6 techniques $\times$ 2 trial blocks) were created.

9.9 Limitations and Future Work

There are several limitations in the current design of Experimentvr that highlight areas for improvement in future work. Experimentvr supports different VR platforms (e.g., HTC Vive, Oculus Rift, etc.) but does not support different types of input devices, such as eye trackers. Future work could modify Experimentvr to support various input devices. The usability of the testbed could be improved by enabling the researcher to go back to different points in the study, instead of just being able to redo a single trial, which is a limitation of its current design. Currently, only ray casting is implemented as a selection technique. Future work could enable researchers to port other selection techniques into the testbed, however it does not currently support this process. Finally, future work could evaluate the usability and usefulness of the testbed with other VR researchers to investigate if it sufficiently supports different types of VR experiment designs.

9.10 Summary

Experimentvr is an extensible and customizable testbed that is designed based on typical user and VR studies in human-computer interaction research. It enables researchers to design and

implement multi-factor and full factorial experiments as well as automate the study procedure. Experimentvr also logs low-level device data (e.g., position, acceleration) during the experiment to CSV or JSON files at a user-defined sampling rate. It supports three different tasks: Timed, Locomotion, and Selection, which the experimenter can choose depending on the interaction technique or user behavior they wish to measure. In a pilot test, I successfully used Experimentvr to conduct an experiment with 40 participants who used six different locomotion techniques. Future work can investigate how to improve the usability of the testbed by studying its use “in the wild” with other VR researchers.

10 COMPARING LOCOMOTION TECHNIQUES WITH PEOPLE WITH UPPER-BODY IMPAIRMENTS²⁹

One of the most important ways of interacting in virtual reality (VR) is navigating virtual environments (VEs) using locomotion techniques. Since VR interaction started becoming a research focus 30 years ago, researchers and designers have developed numerous locomotion techniques. These techniques are varied in terms of their control mechanisms (body movement-based, controller-based, etc.), the parts of the body used to control the techniques (legs, arms, head, etc.), and the level of physical exertion required to navigate a VE. Although a substantial amount of research has focused on designing locomotion techniques that emulate real-world movement, commercial implementations of locomotion techniques have enabled users to move around virtual spaces while seated because of safety concerns and space constraints. This emphasis on seated techniques also benefits people with limited mobility. As a result of designers looking for ways to overcome real-world constraints, over one hundred distinct locomotion techniques exist to this day (Di Luca et al., 2021), giving users the ability to choose a technique that suits their abilities and preferences.

Despite the large number of locomotion techniques and the distinct opportunity that VR presents to people with physical impairments, data do not exist regarding which techniques are most accessible, particularly for people with upper-body impairments, making it difficult to know which techniques work best for which users and why. Seated techniques can be accessible for people with lower-body impairments who otherwise have typical mobility in their upper body; however, seated techniques are not necessarily accessible for people with upper-body motor impairments.

²⁹ Portions of this chapter were published as Franz, R. L., Yu, J., & Wobbrock, J. O. (2023). Comparing locomotion techniques in virtual reality for people with upper-body motor impairments. *Proceedings of the Conference on Computers and Accessibility (ASSETS'23)*, 1–15.

This study is a first step in identifying the most accessible seated locomotion techniques for people with upper-body motor impairments.

10.1 Study Method

The purpose of the study, which will be referred to as the technique comparison study, was to compare six VR locomotion techniques across several different measures to understand their accessibility and participants' experiences of the techniques, including perceived workload, presence, and simulator sickness.

10.1.1 Participants

Twenty participants with self-reported upper-body motor impairments participated in the study. Eligibility criteria included being fluent in English, 18 years or older, able to provide informed consent, having a permanent or temporary motor impairment that affects the upper body, being able to hold at least one Meta Quest 2 controller, and being able to wear a Quest 2 headset. Four non-binary people, six women, and 10 men participated in the study. The mean age was 39.4 years ($SD=16.6$) and self-reported conditions that caused impairments are listed in Table 8, below. One participant had a high school degree, three completed some college, three had a two-year degree, nine had a bachelor's degree, and four had a master's degree. On a four-point scale from "not proficient" to "very proficient" at using technology, one participant rated themselves as "not proficient," one as "somewhat proficient," six as "average," and 11 as "very proficient." One participant did not respond.

Ten participants had never used VR, six participants had used VR with a headset, and four used VR with a headset and a phone. The first time they had ever used VR ranged from 2016 to 2022. Three participants reported using VR only once before, six reported using it rarely, and one reported using it weekly.

Table 8. Participants' self-reported upper-body motor impairment conditions.

Participant	Condition
P00	Left hand amputee
P01	Spinal stenosis at 3C level
P02	Ehler's Danlos Syndrome, Beals Syndrome (FBN2 gene mutation) causing progressive muscle weakness
P03	C-5 quadriplegia
P04	Peripheral neuropathy
P05	Osteoarthritis, nerve damage
P06	Hand tremor and weakness
P07	Limb Girdle Muscular Dystrophy Type 2A
P08	Cerebral palsy
P09	Chronic joint pain, ongoing carpal tunnel syndrome
P10	Paralysis, quadriplegia
P11	Tetraplegia
P12	Muscular dystrophy
P13	Muscular dystrophy
P14	C5-C6 incomplete spinal cord injury with functional quadriplegia
P15	Nerve damage, severe muscle spasms, arthritis in neck, sacroiliac joint pain
P16	Arthritis in both wrists, elbows, and shoulders from overuse, right arm fatigues easily
P17	Peripheral neuropathy in all extremities
P18	Transverse myelitis
P19	Ehlers-Danlos Syndrome

10.1.2 Custom Testbed

I used a commercial Meta Quest 2 headset and controllers for the VR system, which was connected to an Alienware m15 P79F laptop computer. I developed a custom testbed in C# in Unity 2019.4.17 (see Chapter 9). The testbed environment consisted of a horizontal plane suspended in an empty surrounding world. On the plane were six blue targets, arranged in a circle (Figure 31A).




Figure 31. (A) A bird's-eye view of the test environment consisting of six blue targets and paths on a plane. The participant started the task in the center of the plane. The current target would turn from blue to green. (B) The participant's view when they moved inside the current target, which was green. The blue countdown shows the time remaining to complete the trial and the pink countdown shows the seconds the participant must remain inside the target for the trial to be considered a "hit." (This is analogous to dwell-based selection on a touch screen.)

10.1.3 Locomotion Techniques: Overview

I tested six locomotion techniques selected such that the techniques used different parts of the body (e.g., head, torso, arms, fingers), different control mechanisms (e.g., pointing, controller button manipulation, repetitive movement), and different levels of effort (i.e., low to high). See Table 9 for a breakdown. I used Di Luca et al.'s (2021) Locomotion Vault to identify the six locomotion techniques across a range of categories defined in their paper. I ensured that all techniques could be performed while seated and that any two-handed techniques could be performed with one hand as well. None of the techniques required lower-body movement. Participants could perform a stationary turn using the thumbstick for all locomotion techniques except Astral Body, because Astral Body used a third-person perspective. The stationary turn rotated at discrete 45-degree increments. I added the stationary turn mechanism because only two of the first-person techniques, Chicken Acceleration and Sliding Looking, enabled directional control. Participants only had to learn the stationary turn once because it was consistent across all first-person techniques.

Table 9. The six locomotion techniques tested in the study, instructions for using the technique, the control mechanism for forward movement and direction, and the control mechanism type. The table also presents the body part, effort level according to Di Luca et al. (2021), and maximum number of hands required to use the technique. Finally, Di Luca et al.'s categorization of the technique in their taxonomy is included.

Locomotion technique	Instructions	Forward movement	Direction	Control mechanism	Body part	Effort level	Hands	Di Luca et al. (2021) category
Astral Body	Use the thumbstick to move the avatar	Thumbstick	Thumbstick	Controller input	Thumb	low	1	Controller
Chicken Acceleration	Lean forward and look in the direction you want to move	Lean forward	Rotate head	Movement	Head and torso	medium	0	Relative position
Grab and Pull	Reach forward, press and hold the trigger, pull the controller back toward you, then release the trigger	Reach and retract	Always forward	Movement, controller input	Arm(s), fingers	high	2	Grab
Sliding Looking	Press and hold the "A" or "X" button and look in the direction you want to move	Press the "A" or "X" button	Rotate head	Movement, controller input	Fingers, head	low	1	Controller
Teleport	Aim the blue circle and press trigger to jump to the circle	Aim and press trigger	Always forward	Aiming, controller input	Arm, finger	low	1	Teleport
Throw Teleport	Press and hold the grip to show a ball. Throw the ball and release the grip to release ball. You will move to where the ball lands.	Aim and throw ball	Always forward	Movement, controller input	Arm(s), fingers	high	2	Teleport

10.1.4 Locomotion Techniques: Astral Body

With the Astral Body technique, a participant views their avatar from a third-person perspective looking down on the environment. They control the avatar using the thumbstick by pushing it in the direction that they want the avatar to travel. In the original implementation,³⁰ the user could switch between third- and first-person perspectives, but I disabled this feature because the first-person locomotion could be considered a different technique and would not make trials comparable among participants. I refer to this technique as Astral.

10.1.5 Locomotion Techniques: Chicken Acceleration

With the Chicken Acceleration technique, a participant leans forward to move in the direction they are facing, with increasing speed the more they lean. The participant leans back to slow down and stop. The forward direction is controlled by steering with the head (i.e., by looking in the direction the participant wants to travel). This technique does not require any controller input. I refer to this technique as Chicken.

10.1.6 Locomotion Techniques: Grab and Pull

With the Grab and Pull technique, a participant reaches out with one controller in front of their body, holds down the trigger button, pulls the controller back towards their body, and releases the trigger to grab and pull the ground towards themselves. They can then repeat this movement with one arm or alternate arms, similar to pulling a rope, thereby moving forward continuously. The participant can also grab and pull laterally to the side of their body to move sideways. The participant remains facing forward with this technique unless they change directions using a stationary turn with the thumbstick. I refer to this technique as Grab.

³⁰ <https://vrgamecritic.com/game/alice-mystery-garden>

10.1.7 Locomotion Techniques: Sliding Looking

With the Sliding Looking technique, a participant presses and holds the “A” button (right controller) or, equivalently, the “X” button (left controller) to move forward continuously. They change their forward direction by looking in the direction that they want to move. I refer to this technique as Sliding.

10.1.8 Locomotion Techniques: Teleport

With the Teleport technique, a participant aims the teleporter, which is an arc emanating from the left or right controller. A blue circle appears at the end of the arc where it intersects with the ground. When the participant presses the trigger, they jump to the location of the blue circle. The participant increases or decreases the distance to the blue circle by tilting the controller up or down, thereby adjusting the height of the arc. Once the participant presses the trigger, they experience an instant change to the new viewpoint. The participant remains facing forward with this technique unless they change directions using a stationary turn with the thumbstick. I refer to this technique as Teleport.

10.1.9 Locomotion Techniques: Throw Teleport

With the Throw Teleport technique, a participant presses and holds the grip button to activate a ball that appears on top of their controller. The user then tosses the ball with an underhand or overhand motion. They release the grip button when they want to release the ball and are teleported to where the ball lands. Like with the teleport technique, the participant experiences an instant change in viewpoint once they teleport. The participant remains facing forward with this technique unless they change directions using a stationary turn with the thumbstick. I refer to this technique as Throw.

10.1.10 Procedure

After signing the consent form, a participant filled out a demographic questionnaire. The researcher and participant then reviewed the controller buttons required for the task. After the participant felt comfortable with the controllers, the researcher helped the participant put on and adjust the headset. The researcher loaded the testbed and guided the participant to practice using the thumbsticks to perform a stationary turn. After the participant had practiced turning, the researcher launched an empty environment in which the participant could practice using the locomotion technique. The test phase began when the participant felt comfortable using the locomotion technique.

The participant started the test phase in the center of the circle of targets with an initial five second countdown appearing in front of them. Once the countdown finished, they were rotated toward the first target, which turned green, and a 30-second countdown appeared in front of them. They were instructed to move towards the target using the technique they had just practiced, and the countdown, which moved with the participant, indicated how much time they had to move to the target. Once inside the target, a five second countdown appeared indicating how much time was required to stay inside the target before the trial would be considered a “hit” (Figure 31B). If the participant exited the target before the five-second countdown expired, the countdown disappeared. If they re-entered the target, the countdown restarted at five seconds. The trial was considered a “miss” if the 30-second countdown ended before the participant had stayed inside the target for five seconds. The 30-second countdown was chosen for practical reasons to ensure that the study stayed within a 2.5-hour window.

If the participant failed to move into the target before the 30-second countdown ended, the participant was automatically moved inside the target and rotated towards the next target. I displayed the countdown so that participants were not taken by surprise when they were moved to

the next target. If the participant hit the target before the 30-second countdown ended, they were only rotated towards the next target. Participants were rotated toward targets to eliminate the need for visual search. This process repeated for all six targets in the environment and each path between targets was considered one “trial.” Six trials were considered one “trial block.” For two of the six paths between targets, obstacles appeared, and the participant was instructed to maneuver around them. I introduced obstacles so that participants would have to control their direction, instead of just moving in a straight line. There was no visual, auditory, or haptic feedback if the participant hit an obstacle; they passed through the obstacle and the event was logged. There were 40 obstacles on path 3 and 80 obstacles on path 5 (Figure 31A). The obstacles were construction materials consisting of workbenches, lighting rigs, and paint cans (Figure 32A). Obstacles were a range of heights and sizes (e.g., paint cans were small and sat low to the ground while workbenches were wide and at waist height) to reflect a range of obstacles that might appear in a VE.




Figure 32. (A) The participant’s view of the test environment and 30-second trial countdown. The target that the participant is navigating to is green and other targets in the environment are blue. The participant must also maneuver around the work benches, lighting rigs, and paint cans. (B) A participant performing the task using one hand to stabilize the controller and the other to press the buttons.

After completing one trial block, the participant completed a duplicate trial block. The second trial block was the same as the first, with participants navigating to targets in the same order and

obstacles appearing on the same paths as in the first trial block. The trial block was repeated twice to get a more accurate estimate of performance measures by averaging out any learning effect.

Next, the participant took off the headset and completed a post-task questionnaire on a laptop. Afterwards, the researcher asked the participant brief questions about their experience using the just-completed technique. Questions included what the participant liked and disliked about the technique, how their impairment affected how they used the technique, how difficult it was to use the technique, and whether they struggled to use the technique at a particular point during the task. After the interview, the user put the headset back on and repeated the process for the next technique. Technique order was counterbalanced via random presentation.

After completing this procedure for all six locomotion techniques, the researcher conducted a final interview in which she asked the participant to compare all six techniques. The participant was asked to choose the most and least comfortable technique, their favorite and least favorite, and the three techniques that they would use in VR consistently. An entire study session took between 2 and 2½ hours.

10.1.11 Design and Analysis: Overview

The study employed a single factor within-subjects design. The factor, locomotion Technique, had six levels: Astral, Chicken, Grab, Sliding, Teleport, and Throw. The data set and analysis approaches are explained in the subsections below.

10.1.12 Design and Analysis: Task Performance Measures

There were three outcome measures per locomotion trial. The first measure was the Trial Time, in seconds. The second measure was a dichotomous outcome for whether the target was Hit or not. The third measure was a count of the Obstacles Hit. There were 3987 total data points in this study.

Concerning data analysis, mixed models were used to account for repeated measures (Littell et al., 1996, 1998). Specifically, Trial Time was analyzed using the nonparametric aligned rank transform procedure (Elkin et al., 2021; Higgins & Tashtoush, 1994; Wobbrock, Findlater, et al., 2011) with a linear mixed model analysis of variance (Frederick, 1999) after verifying the ANOVA assumptions were met. Hits (versus misses) were analyzed using mixed logistic regression owing to it having a dichotomous response (Gilmour et al., 1985; Stiratelli et al., 1984). Obstacles Hit were analyzed using mixed negative binomial regression owing to it having a discrete-count response (Hilbe, 2011).

10.1.13 Design and Analysis: Post-Task Questionnaire Measures

The questionnaire consisted of one presence question, which was question #1 from Slater and Steed's (2000) presence questionnaire. I only included question #1 because researchers found that this question elicited the most direct response for presence and had high discriminating power (Slater & Steed, 2000). Responses ranged from 1 to 7, with 1 being "not at all" to 7 being "very much" in terms of feeling present. I also included the first question from the Simulator Sickness Questionnaire (Kennedy et al., 1993) to measure general discomfort as a result of simulator sickness, and included examples of discomfort (e.g., nausea, dizziness, eyestrain, or vertigo). Responses ranged from 0 being "no discomfort" to 3 being "severe discomfort." Finally, I added five questions from the NASA-TLX workload questionnaire (Hart & Staveland, 1988). I excluded the temporal demand question because I did not want the user to focus on the time it was taking for them to complete the task, which might have induced stress. I adapted the original 21-point scale to a 1 to 7 scale by eliminating the "high," "medium," and "low" increments for each point on the scale to improve readability (Hart, 2006). There were 114 data points for all questions except

Presence, which had 112 responses. I analyzed the questionnaire data, all of which were Likert-type scales, using mixed ordinal logistic regression (Hedeker & Gibbons, 1994).

10.1.14 Design and Analysis: Interview Data

I conducted post-task and post-study semi-structured interviews with each participant. I transcribed the audio files for the post-task interviews and the final interview for each participant. P09's audio files were unfortunately lost so the results exclude her interview responses. I then performed an inductive thematic analysis by identifying patterns across participant utterances and organizing them by theme (Braun & Clarke, 2008). I created eight themes and 51 sub-themes. A researcher on the team performed an interrater reliability assessment and achieved a Cohen's Kappa of $\kappa=0.87$, indicating strong agreement (Landis & Koch, 1977).

10.2 Results

I report the performance measures, questionnaire responses, and interview responses in this section.

10.3 Performance Measures

I report the results of statistical tests comparing locomotion techniques with respect to Trial Time, Hit Rate, and Obstacles Hit (a count measure).

10.3.1 Trial Time

To test the normality assumption for a repeated measures ANOVA, an Anderson-Darling test (Anderson & Darling, 1954) was run on the residuals of a repeated measures full-factorial ANOVA model. The test was statistically significant ($A=5.62$, $p<.0001$), indicating that the residuals did not conform to a normal distribution. Therefore, I used the nonparametric aligned rank transform procedure (Higgins & Tashtoush, 1994; Wobbrock, Findlater, et al., 2011) with a linear mixed model (Frederick, 1999) for the Trial Time analyses.

Figure 33A shows a bar plot for Trial Times by Technique. The fastest technique was Teleport. There was a statistically significant effect of Technique on Trial Time ($F(5,1347.4)=184.31$, $p<.0001$). *Post hoc* pairwise comparisons conducted with the ART-C procedure (Elkin et al., 2021), and corrected with Holm's sequential Bonferroni procedure (Holm, 1979), indicated that all pairwise comparisons were significantly different except for Grab vs. Throw ($t(1347.1)=-0.31$, $n.s.$).




Figure 33. (A) Bar plot of Trial Time by Technique. Lower is better. (B) Bar plot of target Hit Rate by Technique. Higher is better. Means are above the bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*), $p<.1$ (.).

10.3.2 Hit Rate

Figure 33B shows bar plots for Hit Rate by Technique. The technique with the highest target Hit Rate was Sliding. Because Hit Rate is a dichotomous dependent variable (i.e., it has a value of 0 or 1), I conducted an analysis using a mixed logistic regression model (Gilmour et al., 1985; Stiratelli et al., 1984). This test indicated a statistically significant effect of Technique on Hit Rate

($\chi^2(5, N=1401)=121.66, p<.0001$). Pairwise comparisons, corrected with Holm's sequential Bonferroni procedure (Holm, 1979), indicated some significant differences, as shown in the figure.

10.3.3 Obstacles Hit

Figure 34A shows bar plots of Obstacles Hit by Technique. Fewer obstacles were hit with Teleport compared to all other techniques. The data were a count response and followed a negative binomial distribution, so I used mixed negative binomial regression, as is typical for count responses (Hilbe, 2011). An omnibus test indicated a statistically significant effect of Technique on Obstacles Hit ($\chi^2(1, N=1401)=104.00, p<.0001$). Pairwise comparisons, corrected with Holm's sequential Bonferroni procedure (Holm, 1979), indicated some significant differences, as shown in the figure.




Figure 34. (A) Bar plot of Obstacles Hit per trial by Technique. Lower is better. Means are above the bars. (B) Boxplot of Likert responses for Presence by Technique. Medians are above and inside the boxes. Higher is greater experienced Presence. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*), $p < .1$ (.).

10.4 Questionnaire Responses

In this subsection, I report the results of statistical tests comparing the locomotion techniques by Presence, Simulator Sickness, Mental Demand, Physical Demand, Performance, Effort, and Frustration. All omnibus tests were conducted using an analysis of variance based on mixed ordinal logistic regression (Hedeker & Gibbons, 1994) and all pairwise comparisons were performed using Z-tests, corrected with Holm's sequential Bonferroni procedure (Holm, 1979).

10.4.1 Presence

Figure 34B shows the distribution of Likert responses (1-7) for Presence by Technique. Chicken, Sliding, and Throw all had similarly high Presence. There was a statistically significant effect of Technique on Presence ($\chi^2(5, N=20)=16.17$, $p < .01$). Pairwise comparisons indicated that Likert scores for some comparisons were significantly different, as shown in the figure.

10.4.2 Simulator Sickness

Figure 35A shows the distribution of responses (0-3) for Simulator Sickness by Technique. Grab induced the greatest Simulator Sickness of all techniques. There was a statistically significant effect of Technique on Simulator Sickness ($\chi^2(5, N=20)=24.74, p<.0001$). Pairwise comparisons indicated that Likert scores for some comparisons involving Teleport were significantly different, as shown in the figure.




Figure 35. (A) Boxplot of responses for Sickness by Technique. Lower is less experienced Sickness. (B) Boxplot of Likert responses for Mental Demand by Technique. Lower is less perceived Mental Demand. Medians are above and inside the boxes. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*), $p<.1$ (.).

10.4.3 Mental Demand

Figure 35B shows the distribution of Likert responses (1-7) for Mental Demand by Technique. Throw had the highest perceived Mental Demand. There was a statistically significant effect of Technique on Mental Demand ($\chi^2(5, N=20)=31.63, p<.0001$). Pairwise comparisons indicated that Likert scores for some comparisons were significantly different, as shown in the figure.

10.4.4 Physical Demand

Figure 36A shows the distribution of Likert responses (1-7) for Physical Demand by Technique. Grab and Throw had the highest perceived Physical Demand. There was a statistically significant effect of Technique on Physical Demand ($\chi^2(5, N=20)=68.04$, $p<.0001$). Pairwise comparisons indicated that Likert scores for some comparisons were significantly different, as shown in the figure.




Figure 36. (A) Boxplot of Likert responses for Physical Demand by Technique. Lower is less perceived Physical Demand. (B) Boxplot of Likert responses for Performance by Technique. Higher is better perceived Performance. Medians are above and inside the boxes. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*), $p<.1$ (.).

10.4.5 Performance

Figure 36B shows the distribution of Likert responses (1-7) for Performance by Technique. Teleport offered the highest perceived Performance. There was a statistically significant effect of Technique on Performance ($\chi^2(5, N=20)=43.01$, $p<.0001$). Pairwise comparisons indicated that Likert scores for some comparisons were significantly different, as shown in the figure.

10.4.6 Effort

Figure 37A shows the distribution of Likert responses (1-7) for Effort by Technique. Chicken, Grab, and Throw seemed to require the greatest Effort. There was a statistically significant effect of Technique on Effort ($\chi^2(5, N=20)=53.74, p<.0001$). Pairwise comparisons indicated that Likert scores for some comparisons were significantly different, as shown in the figure.




Figure 37. (A) Boxplot of Likert responses for Effort by Technique. Lower is less perceived Effort. (B) Boxplot of Likert responses for Frustration by Technique. Lower is less perceived Frustration. Medians are above and inside the boxes. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*), $p<.1$ (.).

10.4.7 Frustration

Figure 37B shows the distribution of Likert responses (1-7) for Frustration by Technique. Chicken, Grab, and Throw were perceived as causing the highest Frustration. There was a statistically significant effect of Technique on Frustration ($\chi^2(5, N=20)=36.97, p<.0001$). Pairwise comparisons indicated that Likert scores for some comparisons were significantly different, as shown in the figure.

10.5 Interview Responses

I report interview responses by discussing how participants compared all six techniques, their thoughts on technique accessibility and interaction design, the effect of their impairments on their use of the techniques, their thoughts on overall user experience, and their enjoyment and challenges with the techniques.

10.5.1 Comparison of Locomotion Techniques

Table 3 shows, for each participant, the technique they each felt was most and least comfortable, most and least preferred, and their top three preferred techniques for actual use in VR.

Table 10. Participant responses to final interview questions comparing the locomotion techniques. “Top 1” refers to the technique they would use most, “top 2” is the second-most, and “top 3” is the third-most.

ID	Most comfortable	Least comfortable	Most preferred	Least preferred	Top 1	Top 2	Top 3
P00	Chicken	Grab	Chicken	Throw	Chicken	Astral	Teleport
P01	Chicken	Throw	Chicken	Throw	Chicken	Sliding	Grab
P02	Astral	Chicken	Teleport	Chicken	Teleport	Siding	Throw
P03	Chicken	Throw	Chicken	Throw	Chicken	Teleport	Sliding
P04	Teleport	Grab	Teleport	Throw	Teleport	Sliding	Chicken
P05	Astral	Throw	Astral	Grab	Astral	Sliding	Throw
P06	Chicken	Throw	Chicken	Throw	Teleport	Chicken	Astral
P07	Astral	Grab	Sliding	Grab	Sliding	Astral	Throw
P08	Sliding	Throw	Sliding	Chicken	Sliding	Grab	Astral
P10	Teleport	Grab	Chicken	Astral	Chicken	Teleport	Sliding
P11	Astral	Sliding	Teleport	Chicken	Teleport	Astral	Throw
P12	Astral	Throw	Astral	Throw	Astral	Teleport	Sliding
P13	Sliding	Astral	Sliding	Grab	Sliding	Astral	Throw
P14	Teleport	Sliding	Teleport	Throw	Teleport	Astral	None
P15	Astral	Sliding	Teleport	Chicken	Teleport	Astral	Throw
P16	Teleport	Grab	Teleport	Grab	Teleport	Astral	Sliding
P17	Astral	Throw	Grab	Chicken	Astral	Grab	Teleport
P18	Astral	Grab	Chicken	Grab	Chicken	Sliding	Throw
P19	Teleport	Grab	Teleport	Chicken	Teleport	Throw	Astral

Astral was considered the most physically comfortable technique by most participants ($n=8$), followed by Teleport ($n=5$), Chicken ($n=4$), and Sliding ($n=2$). Throw was considered the least physically comfortable technique by the most participants ($n=8$), followed by Grab ($n=7$), Sliding ($n=3$), and Chicken ($n=1$). Most participants preferred Teleport ($n=7$) or Chicken ($n=6$), followed

by Sliding ($n=3$), then Astral ($n=2$). Most participants liked Throw the least ($n=7$), followed by Chicken ($n=6$), Grab ($n=5$), and Astral ($n=1$).

Most participants said that they would use Teleport if they were going to use VR ($n=8$). A smaller group would use Chicken ($n=5$). Three participants would use Astral and another three would use Sliding. In terms of the technique that they would use second-most, most participants would use Astral ($n=7$). Other participants would use Sliding ($n=5$), Teleport ($n=3$), or Grab ($n=2$). Finally, most participants would use Throw ($n=7$) third-most. Other participants would use Sliding ($n=4$), Astral ($n=3$), or Teleport ($n=2$). P14 did not like any of the other techniques and would not use a third technique.

10.5.2 Accessibility and Interaction Design

Participants expressed that the controls were easy to use for Astral, Chicken, and Teleport. Astral and Teleport required one button press and Chicken did not require any. Participants also expressed that Astral, Grab, and Sliding gave them a heightened sense of control over navigation. Regarding Sliding, P04 said, *“everything felt a little bit more controlled, and it still felt pretty easy because I was just pressing one button and just moving my head”* (P04). Participants felt that these three techniques were precise, responsive, and made it easy to maneuver around obstacles. Teleport allowed them to navigate the environment particularly quickly and they liked that the teleporter provided visual feedback about where they were going to land after the jump. P12 felt the visual feedback gave her confidence, saying, *“I like the line, like the projection line. That also does help visually, I think. Makes you more confident in utilizing it. Rather than just an arrow showing you where you end up”* (P12).

Some participants expressed that the movements for Chicken and Grab felt natural and intuitive. For Chicken, they liked using their bodies to move forward and control their speed, rather than

using their hands. P00 said, *“the more I leaned forward, the faster it went. If I wanted to go to a different direction, whatever direction I looked and leaned, it went to. So, it's so easy”* (P00). Meanwhile other participants expressed that moving forward with their hands felt natural with Grab. For example, P03 said, *“it feels good to virtually move myself without using my wheelchair, but as if I were dragging myself with my arms”* (P03).

On the other hand, participants felt that they had poor control over their navigation with Chicken and Throw. They found it difficult to slow down, stop, and steer with Chicken, while some found it challenging to aim, avoid obstacles, and time the button release with Throw. Some participants also mentioned that it was challenging to turn with their heads using Sliding and Chicken. These participants expressed that they could not rotate their virtual body far or fast enough, and that it was difficult to coordinate the head rotation with the thumbstick turn. P07 struggled using these two inputs to change direction with Sliding and said, *“the [head] movements were just very subtle with direction, so I was still very much relying on the thumbstick to change direction”* (P07). Leaning forward to move forward with Chicken felt unnatural for some participants because it did not have a real-life equivalent. Other participants also felt it was unnatural to navigate by throwing a ball with Throw. For example, P17 said, *“it didn't seem fluid. It seemed choppy”* (P17).

The stationary turn with the thumbstick was implemented for all techniques except for Astral because Astral was controlled from a third-person perspective. Many participants struggled to control the stationary turn for a few reasons. First, several participants found the discrete turn disorienting because it instantly changed their viewpoint. Second, other participants wanted finer control over the turn by rotating in smaller increments or continuously. Third, P15 related that since she could only push the thumbstick to the right, she had to push more times and got disoriented during the process. She said, *“if you are impaired going one way, then you have to*

keep flipping around. And it seems like if you have to move a quarter of a turn to the right and you move the joystick to the right, it seems like it flips you three quarters or one quarter to the left instead of the right. And you have to backtrack. It's kind of weird" (P15).

10.5.3 Perceived Effect of Impairment

Participants reported that their impairments had an effect on their ability to use all techniques, but particularly on their ability to use Chicken, Grab, and Throw. They discussed either not being able to perform the technique to its maximum capability, or not being able to use the technique how it was intended. For example, P11 used the tilt feature on his wheelchair to mimic leaning back and forth with Chicken. Participants reported that their impairments had little to no effect mainly when using Astral, Sliding, and Teleport.

Some participants could only use one controller because they could not grip it correctly. Instead, they used one hand to stabilize the controller and the other hand to press the buttons (Figure 32B). As a result, these participants reported difficulty positioning the controller, finding and pressing buttons, as well as preventing their fingers from slipping off buttons. A number of participants noted that the controllers were too heavy, leading to hand-shaking and difficulty performing tasks for extended periods of time. Some participants reported that the headset was heavy and caused discomfort, particularly when using Astral because they had to look down on the environment. P17 reflected, *"the actual exercise of having to look down like that for a long time, especially without a kind of counterweight or just understanding how to have to put it so that it doesn't affect me as much; I still felt out of balance and blurry having to look down"* (P17). Participants felt discomfort in their wrists, arms, shoulders, and necks when using Grab and Throw.

Limited range of motion in the neck affected some participants' ability to use Chicken and Sliding, while limited arm reach and range of motion in the shoulders affected some people's

ability to use Grab and Throw. Muscle weakness, especially in the core and arms, also affected some participants' ability to use Grab and Throw. P10 reflected on how core weakness affected her sense of security: *"I don't have a lot of core control, and so the throwing, I think I could have thrown it farther if I felt more secure, like moving more"* (P10).

10.5.4 Overall User Experience

Participants liked that Astral, Chicken, Sliding, and Teleport were not physically demanding. They particularly liked that they had to put in less effort, did not experience pain or discomfort, and only had to use one finger for Astral. P08 liked Astral *"because it was less stressful, less physically demanding. You're just using the thumbstick"* (P08). P04 liked that Chicken was not physically demanding, saying, *"I didn't need to throw my entire body into it. I didn't need to be exerting out of ten in order to get the same response. And so, I like that about it"* (P04).

Participants reported that Astral and Teleport were easy to learn and not mentally demanding to use because they only required one button press to move forward. P12 liked that Teleport required minimal mental effort to plan her route. She said, *"I just know where my destination is after, like, a shorter time, I feel like. So instead of using energy to look for something and decide while I'm moving, I could just decide and then move. [I] feel like it's just faster that way"* (P12).

Some participants verbalized feeling more immersed with Chicken, Grab and Sliding compared to other techniques, even though significant differences did not exist in questionnaire responses. P00 felt more present with Sliding because of the change in visual feedback he experienced when steering with his head. He said, *"I think that the environment is much of a trigger of the visual, and by steering with where you look, it makes you feel like you're in the environment"* (P00). These three techniques were also the only ones that used first-person continuous locomotion in the VE

instead of discrete jumps or a third-person perspective. Some participants felt more oriented in space with Astral because they could see all of the environment.

Participants disliked that Grab and Throw were physically demanding, and thought these techniques required too much movement and that the movement was repetitive. Some participants also mentioned feeling like they had not traveled far enough for the level of effort that they put into the technique when using Grab. For example, P15 said, *“It's harder to move. It feels like you should move more per pull”* (P15).

Participants expressed that Chicken, Grab, Sliding, and Throw were difficult to learn, and that they would need time to be able to use the techniques with ease. P00 reflected on Chicken's learnability saying, *“I think there is a little bit of a learning curve. You got to know how to lean and how far to lean. It takes a little bit of adjustment on that, but it's easy once you got the hang of that”* (P00).

A few participants reported feeling less present with Astral and Teleport and likened the experience to playing a traditional video game. With regards to spatial orientation, some participants felt disoriented when using Chicken and Teleport, particularly when they were getting close to the target. For example, about Chicken, P04 said, *“it was sometimes hard for me to tell whether I was successfully in the target, and sometimes I felt like I was sliding out of the target”* (P04). Only three participants discussed simulator sickness, which they felt when using Sliding and Grab.

10.5.5 Enjoyment and Challenge

Some participants enjoyed using Chicken, Grab, and Throw, and said the techniques were entertaining to use. They liked Throw in particular because using it felt like a game and was reminiscent of shooting baskets. For example, P03 enjoyed the physical and mental challenge of

Throw, he said, *“the physical activity of it I liked and the challenge of seeing how far I could throw the ball and learning what works better, like a higher arc when throwing the ball. So, figuring out different ways to throw it to get where I wanted to go”* (P03). Participants also liked that these three techniques provided exercise and could be used for physical therapy. For example, P08 liked that he exercised an underused part of his body with Grab: *“It gave me the opportunity to use my left hand more, and I got a good arm workout, obviously, but yeah, it was really good to really manipulate my arms and stuff”* (P08).

Other participants felt exactly the opposite about these same techniques. They said that Chicken, Grab, and Throw were frustrating to use because they moved too slowly, were not mentally stimulating, did not feel a sense of satisfaction, and did not enjoy the repetitiveness of Grab and Throw. A couple of participants mentioned that Teleport was too easy to use, which made them feel like they were cheating. P17 said, *“well, it seems like cheating, honestly, because I think the idea is we want our brain to become more keen about visual distance and calculating what we have to do”* (P17).

10.6 Summary

Performance and questionnaire data suggest that the six locomotion techniques can be divided into two groups and that techniques within these groups are similar to each other. I discuss these two groups first, then examine how participants ranked techniques, and finally discuss design implications.

10.6.1 Astral, Sliding, and Teleport are Best in Terms of Performance and Workload

In terms of Trial Time and Hit Rate, Astral, Sliding, and Teleport were the top three techniques. In terms of Obstacles Hit, Teleport outperformed all other techniques with the fewest obstacles hit.

The questionnaire and interview data can shed light on why these techniques had the best performance measures. Based on self-reported Performance and Effort responses, participants felt like they needed to work less to be successful using Astral and Teleport. Some participants also explained that the visual feedback of the teleporter helped them avoid obstacles, which is probably why they hit the fewest obstacles with Teleport compared to the other techniques.

Astral, Teleport, and Sliding had significantly lower Physical Demand scores than Throw. Participants mentioned that they liked that these techniques required little movement and also caused little to no pain or discomfort. Participants explained that low Mental Demand scores for Teleport and Astral were a result of only needing to think about a single button input. In addition, the visual feedback of the teleporter helped them aim and plan their route.

Overall, Teleport had the best performance, lowest workload, and was liked for how easy the controls were to use; participants appreciated how fast it was, how easy it was to aim, and how little their impairments impacted their performance. It was also the most popular technique in terms of physical comfort, preference, and the technique that participants would use most in VR. These findings taken together suggest that Teleport was the best technique for this group of participants.

10.6.2 Chicken, Grab, and Throw are Worst in Terms of Performance and Workload

Chicken, Grab and Throw had the worst performance for Trial Time and target Hit Rate. They were similarly poor in terms of Obstacles Hit. Questionnaire and interview data provide insight into why Chicken, Grab, and Throw were the worst in terms of performance. Participants indicated these techniques had significantly greater Physical Demand than Astral, Sliding, and Teleport. They also reported that their impairments had the greatest effect when they were using these three techniques.

The overall higher workload scores for Grab and Throw might be explained by the fact that participants thought these two techniques were strenuous, slow, and uncomfortable to use. Participants mentioned that Chicken, Grab, and Throw had a learning curve, which was reflected in Mental Demand ratings, particularly for Throw. Participants struggled with Throw on many levels, such as coordinating the button release, aiming the ball, and calibrating the distance to throw. All of these challenges could explain the higher Mental Demand scores for Throw.

10.6.3 Most Comfortable, Fastest and Least Demanding Techniques are not Necessarily Preferred

Participants' responses to the most comfortable technique mostly agreed with the techniques that had the lowest workload. Teleport was most popular as the most comfortable and preferable as well as the technique most participants would use in VR. Grab and Throw were considered the least comfortable and least preferred by most participants, which also is consistent with performance results and subjective responses.

Despite its poor performance and high workload, however, Chicken was the second most popular technique: it was the most comfortable, preferred, and most-used technique. This finding might be explained by the interview results. Some participants felt that Chicken was easy to use and control because it did not require any controller input or arm movement.

Even though most participants said they would use Teleport most frequently and Astral second-most frequently in VR, which is consistent with performance and questionnaire results, most participants said they would use Throw third-most, even though it was one of the worst techniques in terms of performance and workload. Entertainment might explain why Throw was the most popular choice for this category, since some participants said it felt like a game.

The majority of participants rated Astral as most comfortable but only two chose it as their preferred technique. This result might be explained by the fact that participants felt that the controls

were easy to use, but participants also felt less present when using Astral compared to other techniques. Thus, it seems Presence affected the preference for a VR locomotion technique.

Seven participants' preferred technique was different than the technique that they thought was most comfortable. Moreover, 10 participants' least preferred technique was different from the technique they thought was least comfortable. This finding suggests that participants weigh trade-offs in accessibility, user experience, and enjoyment when determining their preference for a locomotion technique. While feeling comfortable and performing well are important factors when ranking techniques, other factors can weigh heavily, depending on the participants' individual preferences.

10.6.4 Design Implications

An implication for locomotion technique design is that using controller button inputs rather than upper limb and torso movement to control the technique would likely lead to better accessibility. Although some participants verbalized difficulty reaching and holding controller buttons, a finding also identified in prior work (Mott et al., 2020), this challenge did not severely affect performance with controller button-based techniques such as Teleport, Astral, and Sliding. This finding is likely a result of many participants holding one controller with two hands, which highlights the need to enable mapping of bimanual to unimanual interactions (Yamagami et al., 2022).

Also, although point and select techniques such as Teleport use some upper limb movement, they are likely to be accessible because, depending on the type of impairment, the user can stabilize their arms while using the techniques. In addition, users can travel further with less input using Teleport compared to other techniques. Some participants liked Chicken because they could control the speed of their virtual movement. Other participants mentioned disliking Grab because they felt like they did not move far with each pull. Designers could add a mechanism to locomotion

techniques to control speed or friction so that users feel like they can travel further with less physical exertion, which could improve accessibility. As found in a previous study in which participants with motor impairments struggled to articulate head gestures (Vatavu & Ungurean, 2022), participants with limited range of motion in their necks struggled to control their direction with Chicken and Sliding. In addition, some participants felt that they were not able to turn quickly enough using their heads. As with speed, designers could enable users to adjust the gain so that a smaller physical head rotation produces a larger virtual rotation.

Another insight is that even though an intended benefit of body movement-based techniques is an improvement in a sense of Presence, I did not see this effect. Rather, findings suggest that using third-person perspective negatively impacts Presence.

The findings taken together suggest that VR designers should provide an array of accessible techniques for users with upper-body impairments instead of defaulting to the most accessible or best performing techniques. Techniques available to a user should represent a range of mental and physical challenges, presence, and entertainment. Participants' preferences for techniques were based on the task objective and whether efficiency or enjoyment were primary concerns, which has been found in prior work (Franz et al., 2021; Yamagami et al., 2022). Identifying where locomotion techniques fall on a range of factors including accessibility requires that researchers and designers evaluate techniques not only on traditional performance, user experience, and workload measures, but also on engagement and entertainment measures. The fact that users evaluate locomotion techniques on multiple dimensions points to an opportunity for personalization. In the future, a system could recommend accessible locomotion techniques to users with impairments based on their abilities and individual preferences, a finding consistent

with and endorsed by the concept of ability-based design (Wobbrock et al., 2018; Wobbrock et al., 2011).

11 COMPARING LOCOMOTION TECHNIQUES AMONG PEOPLE WITH AND WITHOUT UPPER-BODY IMPAIRMENTS

This study, which will be referred to as the group comparison study, builds upon the technique comparison study (Chapter 10) by investigating how performance and workload measures compare between people with and without impairments when using locomotion techniques.

The optimal technique for participants with impairments was Teleport in terms of performance and workload measures in the technique comparison study. Participants also reported in the interview that their impairments had little effect on their ability to use Teleport. Although participants with impairments performed the best with Astral, Sliding, and Teleport, there is no empirical data to compare how much performance and workload with these techniques differ from people without impairments. Comparing these groups could potentially reveal the existence of a universally optimal technique that makes a virtual reality (VR) experience accessible by default. Also, quantifying the difference in performance and workload would establish a baseline upon which designers and researchers can improve to make locomotion techniques more accessible.

11.1 Study Method

The study method was the same as that used in the technique comparison study (Section 10.1), except that it was also conducted with people without impairments.

11.1.1 Participants

Participants included those from the technique comparison study in Section 10.1 as well as participants with no self-reported upper-body impairments. Eligibility criteria included being fluent in English, 18 years or older, able to provide informed consent, and not having a permanent or temporary motor impairment that affects the upper body. Within the group without impairments, ten women and ten men participated in the study. The mean age was 23.2 years ($SD=5.0$). Ten

participants had a high school degree, two completed some college, four had a bachelor's degree, and four had a master's degree. On a four-point scale from "not proficient" to "very proficient" at using technology, one participant rated themselves as "somewhat proficient," 11 as "average," and eight as "very proficient."

Nine participants had never used VR, four participants had used VR with a headset, three had used VR with a smartphone, and four used VR with a headset and a phone. The first time they had ever used VR ranged from 2016 to 2022. Three participants reported using VR only once before, five reported using it rarely, one used it every few months, one used it one or more times a month, and one reported using it weekly.

11.1.2 Apparatus

The same study apparatus was used as that in the technique comparison study. Two questionnaires were also given to all 40 participants, both of which measured upper-body impairment level. The QuickDASH (Gummeson et al., 2006) was the first questionnaire participants responded to. It is validated and is a shortened version of the 30-item Disabilities of the Arm, Shoulder, and Hand (DASH) Questionnaire that measures physical function in the arm, shoulder, and hand (Hudak et al., 1996). This questionnaire was selected because it measures function in multiple areas of the upper body, is agnostic to medical diagnosis, and is short (only 11 items) (Smith et al., 2012). Participants rated their responses on a Likert-type scale (1-5) ordered from low to high difficulty. An example item included, "Please rate your ability to do the following activities in the last week by circling the number next to the appropriate response: Use a knife to cut food." The full questionnaire is available in Appendix A.3. An overall impairment score was calculated to represent the impact a range of musculoskeletal diseases has on upper-

body function (Hudak et al., 1996). The score ranges between 0 to 100, where lower scores represent less functional impairment.

The second questionnaire was a custom questionnaire based on literature in accessible computing. Specifically, it was a combination of participant characteristics reported by Mott et al. (2020) and Findlater et al. (2010) but adapted to questionnaire format. These impairments are likely to affect technology use, so it is referred hereon as the “**Technology-Related physical Impairment Questionnaire**” (TRIQ). Participants indicated on 19 items whether they had the specific impairment, so the response was dichotomous. Examples of responses to the statement, “Check all the impairments you experience” were “slow movements,” “tremor,” and “poor coordination.” The full questionnaire is in Appendix A.4.

11.1.3 Procedure

The same study procedure reported in the technique comparison study was used. People with and without impairments filled out the QuickDASH and TRIQ before starting the study tasks. The QuickDASH took an average of 5.3 minutes ($SD=2.9$) and the TRIQ took an average of 6.9 minutes ($SD=5.2$) to complete.

11.1.4 Design and Analysis

The study employed two, two-factor designs, which I refer to as the between-group comparison and the between-level comparison. The first factor for both designs was Locomotion Technique. It was a within-subjects factor and had six levels: Astral, Chicken, Grab, Sliding, Teleport, and Throw.

Between-Group Comparison: The second factor was Group in this study design. It was a between-subjects factor with two levels: With Impairment and Without Impairment. Participants who had reported having an upper-body impairment were in the With Impairment group and were

the same participants from the technique comparison study reported in Section 10.1. The participants who identified as not having an upper-body impairment were categorized as Without Impairment.

Between-Level Comparison: The second factor was Level in this study design. The QuickDASH score was calculated for all 40 participants. Scores were then converted into four impairment levels to facilitate the interpretability of the data analysis. The four levels and corresponding QuickDASH scores were: Low (0-25), Medium-Low (26-50), Medium-High (51-75), and High (76-100). Due to the presence of a single participant in the High group, I merged this group with the Medium-High group to achieve a more balanced distribution of participants across groups. There were 21 participants in Low, six in Medium-Low, and ten in Medium-High groups. The mean QuickDASH score for participants in the Low group was 5.63 ($SD=7.21$), 41.48 ($SD=7.66$) for Medium-Low, and 61.18 ($SD=10.54$) for Medium-High.

11.1.5 Outcome Measures

The three performance-related outcome measures were the same as those from the technique comparison study reported in Section 10.1.12.: Trial Time, Hit Rate and Obstacles Hit. There were 8199 total data points in this analysis.

The responses to the same post-task questionnaire used in the technique comparison study (Section 10.1.13) were also analyzed as outcome measures. These included questions about Presence and Simulator Sickness, as well as five items from the NASA-TLX workload questionnaire: Mental Demand, Physical Demand, Performance, Effort, and Frustration (Hart & Staveland, 1988).

There were 240 data points for all outcome measures except Presence, which had 238 responses due to two participants not responding.

Concerning data analysis, the same tests were used as those described in the technique comparison study in Section 10.1.11: the aligned rank transform procedure for Trial Time, mixed logistic regression for Hit Rate, mixed negative binomial regression for Obstacles Hit, and ordinal logistical regression for post-task questionnaire outcomes.

11.2 Results: Performance

There were some unexpected differences in performance between groups in terms of Trial Time, Hit Rate, and Obstacles Hit.

11.2.1 Trial Time

An Anderson-Darling test (Anderson & Darling, 1954) was run on the residuals of a repeated measures full-factorial ANOVA models for Group $\times$ Technique and Impairment Level $\times$ Technique. The tests were statistically significant (Group $\times$ Technique: $A=26.59$, $p<.0001$. Impairment Level $\times$ Technique: $A=27.37$, $p<.0001$), indicating that the residuals did not conform to a normal distribution. A nonparametric aligned rank transform procedure was used (Higgins & Tashtoush, 1994; Wobbrock, Findlater, et al., 2011) with a linear mixed model (Frederick, 1999) for the analysis.

There was a statistically significant effect of Technique and Group on Trial Time ($F(5,63.7)=161.58, p<.001$). *Post hoc* pairwise comparisons conducted with the ART-C procedure (Elkin et al., 2021), indicated that all pairwise comparisons were significantly different except for Sliding ($t(62.9)=1.73, n.s.$) (Figure 38).




Figure 38. Bar plot of mean Trial Time by Technique and Group. Lower is better. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

There was a statistically significant effect of Technique and Impairment Level on Trial Time ($F(10,2679.9)=39.56, p<.001$). Participants in the Low group had significantly faster Trial Times compared to those in the Medium-High group for all techniques except Sliding (Figure 39). There were no significant differences between the Medium-Low and Medium-High groups, except with Teleport: Medium-Low participants' Trial Times were significantly higher than participants with Low and Medium-High levels of impairment (Figure 39).




Figure 39. Bar plot of mean Trial Time by Technique and Impairment Level. Lower is better. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

11.2.2 Hit Rate

Because targets Hit (versus missed) was a dichotomous dependent variable (i.e., it has a value of 0 or 1), I conducted an analysis using a mixed logistic regression model (Gilmour et al., 1985; Stiratelli et al., 1984). There was a statistically significant effect of Technique and Group on Hit Rate ($\chi^2(5, N=2733)=26.66, p<.001$). Pairwise comparisons indicated significant differences between groups for all techniques except Sliding and Teleport (Figure 40). Participants with impairments had significantly lower Hit Rates than participants without impairments.




Figure 40. Bar plot of mean Hit Rate by Technique and Group. Higher is better. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

When grouping participants by impairment level, there was also a significant effect of Technique and Impairment Level on Hit Rate ($\chi^2(10, N=2733)=58.55, p<.001$). Pairwise comparisons, corrected with Holm's sequential Bonferroni procedure (Holm, 1979), indicated that there were significant differences between all three levels of impairment for Grab (Figure 41). There were no significant differences between impairment levels for Sliding and Teleport. Participants with a Medium-High Impairment Level hit significantly fewer targets than people

with a Low Impairment Level with Astral, Grab, and Throw. Participants in the Medium-Low group hit significantly fewer targets with Grab and Throw compared to participants in the Low group.




Figure 41. Bar plot of mean Hit Rate by Technique and Impairment level. Higher is better. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).

11.2.3 Obstacles Hit

An omnibus test indicated a statistically significant effect of Technique and Group on Obstacles Hit ($\chi^2(5, N=2733)=22.73, p<.001$). Pairwise comparisons indicated that the only significant difference in number of Obstacles Hit between groups was for Throw, where people Without Impairments hit significantly more obstacles compared to people With Impairments (Figure 42). An omnibus test did not indicate a significant effect of Technique and Impairment Level on Obstacles Hit ($\chi^2(10, N=2805)=12.24, p<.001$) so pairwise comparisons were not analyzed (Appendix A.5.1).




Figure 42. Bar plot of mean Obstacles Hit by Technique and Group. Lower is better. Means are at the bottom of error bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

11.3 Results: Post-Task Questionnaire

In this section, I report the results of statistical tests comparing the Group and Impairment Level responses to the post-task questionnaire for each locomotion technique. All omnibus tests were conducted using an analysis of variance based on mixed ordinal logistic regression (Hedeker &

Gibbons, 1994) and pairwise comparisons for Impairment Level were performed using Z-tests, corrected with Holm's sequential Bonferroni procedure (Holm, 1979).

11.3.1 Presence

There was neither a significant effect of Technique and Group ($\chi^2(5, N=40)=2.69$, *n.s.*) nor Technique and Impairment Level ($\chi^2(10, N=40)=12.12$, *n.s.*) on Presence, so pairwise comparisons were not investigated (Appendices A.5.2 and A.5.3).

11.3.2 Sickness

There was neither a significant effect of Technique and Group ($\chi^2(5, N=237)=3.00$, *n.s.*) nor Technique and Impairment Level ($\chi^2(10, N=237)=10.36$, *n.s.*) on Sickness (Appendices A.5.4 and A.5.5).

11.3.3 Mental Demand

There was no significant effect of Technique and Group on Mental Demand ($\chi^2(5, N=237)=6.54$, *n.s.*) (Appendix A.5.6). However, there was a statistically significant effect of Technique and Impairment Level on Mental Demand ($\chi^2(5, N=237)=29.02$, $p<.01$). *Post hoc* comparisons indicated that Medium-Low participants rated Teleport as being significantly more mentally demanding compared to Low and Medium-High participants (Figure 43).




Figure 43. Bar plot of mean Mental Demand rating (1-7) by Technique and Impairment Level. Lower is better. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

11.3.4 Physical Demand

There was a trending significant effect of Technique and Group on Physical Demand ($\chi^2(5, N=237)=10.88, p=.053$). *Post hoc* comparisons indicated that participants With an Impairment rated Astral, Grab, Teleport and Throw as being significantly more physically demanding than participants Without Impairments (Figure 44).




Figure 44. Bar plot of mean Physical Demand rating (1-7) by Technique and Group. Lower is better. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).

There was a statistically significant effect of Technique and Impairment Level on Physical Demand ($\chi^2(10, N=237)=25.12, p<.01$). Pairwise comparisons indicated that Medium-High participants rated Astral, Grab, and Throw as significantly more physically demanding compared to Low participants. Medium-High participants also rated Grab as significantly more physically demanding than Medium-Low participants. For Teleport, only Medium-Low participants rated it as significantly more physically demanding compared to Low participants (Figure 45).




Figure 45. Bar plot of mean Physical Demand rating (1-7) by Technique and Impairment level. Lower is better. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

11.3.5 Effort

There was no significant effect of Technique and Group on Effort ($\chi^2(5, N=237)=8.26$, *n.s.*) (Appendix A.5.7). There was a statistically significant effect of Technique and Impairment Level on Effort ($\chi^2(10, N=237)=19.3$, $p<.05$). *Post hoc* comparisons indicated people with Medium-High impairment levels rated Astral and Grab as requiring significantly more Effort compared to participants with Low impairment levels. Medium-Low participants rated Teleport as requiring significantly more Effort compared to Low participants (Figure 46).




Figure 46. Bar plot of mean Effort rating (1-7) by Technique and Impairment Level. Lower is better. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

11.3.6 Frustration

There was a statistically significant effect of Technique and Impairment Level on Frustration ($\chi^2(5, N=237)=16.17, p<.01$). *Post hoc* comparisons indicated that people With Impairments found Grab, Teleport, and Throw significantly more frustrating to use compared to people Without Impairments (Figure 47).




Figure 47. Bar plot of mean Frustration rating (1-7) by Technique and Group. Lower is better. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

There was also a statistically significant effect of Technique and Impairment Level on Frustration ($\chi^2(10, N=237)=23.34, p<.01$). Comparisons indicated that people with Medium-High levels of impairment thought Grab was significantly more frustrating to use compared to people with Low levels of impairment (Figure 48).




Figure 48. Bar plot of mean Frustration rating (1-7) by Technique and Impairment Level. Lower is better. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

11.3.7 Performance

There was not a significant effect of Technique and Group on Performance ($\chi^2(5, N=237)=10.43, n.s.$) (Appendix A.5.8). However, there was a statistically significant effect of Technique and Impairment Level on Performance ($\chi^2(10, N=237)=25.64, p<.01$). *Post hoc* comparisons indicated that people with Low levels of impairment thought they performed significantly better with Astral, Grab, Sliding, and Throw compared to people with Medium-High levels of impairment. People with Low levels of impairment also thought they performed better with Grab, Sliding, Teleport, and Throw compared to people with Medium-Low levels of

impairment. People with Medium-Low impairment levels thought they performed significantly better with Grab compared to people with Medium-High impairment (Figure 49).




Figure 49. Bar plot of mean Performance rating (1-7) by Technique and Impairment Level. Higher is better. Means are at the bottom of bars. Significance codes: $p < .001$ (***), $p < .01$ (**), $p < .05$ (*).

11.4 Summary

I discuss the implications of the findings by answering the research question, “is there a universally optimal locomotion technique?” and reflecting on how VR locomotion studies can be improved.

11.4.1 Sliding Looking: A Universally Optimal Technique?

The consistent Trial Times, Hit Rates, and Obstacles Hit across Groups and Impairment Levels when participants used Sliding indicated no distinguishable performance differences. This finding suggests that implementing Sliding as a default locomotion technique in applications could democratize the use of VR locomotion, leveling the playing field for all users regardless of their physical abilities.

Participants also rated Sliding as having similar workload irrespective of their impairment severity. However, participants with Medium-Low and Medium-High levels of impairment thought they performed worse than people with Low levels of impairment when using Sliding, indicating that there was a mismatch between real and perceived performance for these groups. This finding also indicates that perceived performance is not a good proxy for real performance for Sliding.

11.4.2 Teleport: Impairment Level Analysis

An interesting finding surfaced when analyzing performance and workload of Teleport by Impairment Level. Participants with impairments' performance was worse compared to that of people without impairments, as expected. Participants with impairments also reported experiencing a higher workload compared to the other group, which was also expected. However, when grouping participants by Impairment Level, the differences seen in the between-group comparisons applied to Medium-Low and Low participants, while participants with Medium-High levels of impairment did not significantly differ from Low participants on most measures of performance and workload.

Re-watching screen recordings of P01 and P08 in the Medium-Low group who had the highest Trial Times when using Teleport, revealed that both became disoriented after accidentally pressing the thumbstick which caused a snap turn. Also, P01 struggled to isolate his index finger to press the trigger while P08 experienced difficulty aiming towards his lefthand side. These two participants also rated Mental Demand, Physical Demand, and Effort as higher compared to other participants in the Medium-Low group. Therefore, although analyzing outcome measures using Impairment Level resulted in intuitive findings for the most part, it did not capture specific challenges such as finger isolation or aiming that might not have been related to a high functional

impairment score. As a result, the QuickDASH score might be a good overall indicator of impairment severity but might not reflect the severity of particular impairments that affect VR interaction.

11.4.3 Relative Importance of Performance Measures

Trial Time was the most important measure of performance for this experimental design. Trial Time had a binomial distribution because of the 30 second time limit. So, it can be seen as a finer-grained measure of Hit Rate. In other words, not only did Trial Time capture Hit Rate, but it also encoded how “close” the user was to a miss.

Obstacles Hit was not a good performance measure for this study. There were no differences between Groups or Impairment Levels, except for a group-wise difference for Throw. People With Impairments hit significantly more obstacles than people Without Impairments. Also, when looking at Obstacles Hit by Impairment Level for Throw, people with Medium-High levels of impairment did not hit any obstacles. This result is likely a consequence of people with Low levels of impairment being able to travel further and incidentally hitting more obstacles, not because people with higher levels of impairment were more precise when throwing the ball. Obstacles Hit could be measuring how much a participant travelled during the trial, particularly with a physically challenging technique, rather than their precision.

11.4.4 Importance of Group-Wise vs. Impairment Level-Wise Comparisons

Performance and workload differences often depended on the Impairment Level, rather than the existence of an impairment. For example, there was not a significant difference in Mental Demand when looking at the Group comparison for Teleport. But when looking at the Impairment Level comparison, there was a significant difference between Medium-Low and the other groups. This

finding underscores the importance of treating functional impairment as a spectrum rather than a binary attribute in studies with people with disabilities.

11.4.5 Reflecting on the Post-Task Questionnaire

It might be difficult for participants to quantify their perceived performance because Medium-High and Medium-Low participants rated their Performance as significantly worse than Low participants for Sliding, despite there not being a difference in terms of objective performance measures. Otherwise, the responses to individual NASA-TLX items provided some insight into why participants with impairments performed worse with some techniques. For example, Astral, Grab, Teleport, and Throw had a significantly higher Physical Demand ratings for participants With Impairments, but not significantly higher Mental Demand ratings, suggesting that although these techniques were not accessible, it was not challenging for participants With Impairments to learn how to use the techniques.

There was no effect of impairment on user experience-related measures of Presence or Sickness. Even though this is an intuitive finding, it actually contradicts previous work, which reported that Presence might be influenced by having an impairment. People with Multiple Sclerosis (MS) experienced more Presence when walking in a virtual environment (VE) than did people without MS (Guo & Quarles, 2012). This contradiction could be a result of studying a more specific population (i.e., upper-body impairment vs. MS) or the locomotion technique used (i.e., seated locomotion techniques vs. real walking). Future work can explore how different types of impairment and locomotion impact Presence.

11.4.6 Summary

People With Impairments had lower performance and greater workload compared to people Without Impairments when using all techniques except Sliding, where the only difference found

between Impairment Levels was perceived performance. This finding suggests that Sliding could be used as a default technique in VR applications. Reflecting on the experiment design and apparatus provided insight into best practices for conducting VR locomotion studies with people with and without impairments, including measuring impairment on a spectrum, choosing informative performance measures, and being aware of how having an impairment can affect perceptions of performance.

12 METRICS THAT ENCODE DIFFERENCES IN VIRTUAL REALITY INTERACTION BETWEEN PEOPLE WITH AND WITHOUT UPPER-BODY IMPAIRMENTS

This study, which will be referred to as the metrics exploration study, builds on the group comparison study by quantifying the differences in how people with and without impairments interacted with locomotion techniques in an attempt to explain performance differences.

12.1 Motivation

The group comparison study revealed that people with impairments performed worse and reported a higher workload than people without impairments when using all locomotion techniques except Sliding. Yet, this study did not provide insight into why people with impairments performed worse. Interview data from the technique comparison study indicated how participants perceived their impairments affected their interaction with the technique, but there is no empirical data to corroborate their verbal reports. Interviews and think-alouds are known to be limited when applied in user studies because people are often not able to verbalize their rationale for their actions (Nielsen et al., 2002). This limitation could also have impacted participants' abilities in the technique comparison study to identify how their interaction and performance with locomotion techniques was affected by their impairments.

Analyzing low-level controller and headset data while people with impairments were using locomotion techniques can provide insight into how the techniques' control mechanisms were affected by participants' impairments. For example, participants with impairments might have rotated their headsets less when using Chicken or Sliding, indicating that they had less range of motion in their necks compared to participants without impairments, which could have affected their performance with these techniques. Identifying differences in how users with and without impairments interacted with a technique could indicate how the technique could be adapted to

improve its accessibility. For example, the head rotation gain could be adjusted for Chicken and Sliding so that users with limited range of motion in their necks could rotate in the virtual world as much as users with typical range of motion.

There was a high amount of variability in types and severities of the upper-body impairments participants had. For example, some participants could move their arms easily but had difficulty turning their heads, while the opposite was true for other participants. Despite this variability, there could be ways in which having an upper-body impairment consistently impacted interaction, resulting in lower performance. Measuring how participants with impairments interacted and moved when using locomotion techniques could reveal commonalities that are not specific to particular impairments. Identifying these commonalities could help designers pinpoint specific interaction components within locomotion techniques that are inaccessible for people with a range of upper-body disabilities. Identifying these commonalities could ultimately help designers make locomotion techniques accessible to more people with different types of disabilities.

Metrics that quantify the effect of an impairment at the interaction level could also be used as features for a model that predicts a new user's performance with a locomotion technique. Measures that capture the relationship between interaction, impairment status, and performance could be used to train models that are able to predict performance with multiple locomotion techniques based on an initial calibration task with a locomotion technique. These metrics would be sensitive to the presence of an upper-body impairment and could lead to more accurate predictions of performance.

12.2 Study Method

The participants, study apparatus, and procedure were the same as those in the technique and group comparison studies in Chapters 10 and 11 respectively. This section reports additional data

collected during the study and presents metrics that could capture interaction-level differences between participants with and without impairments.

12.2.1 Raw Data

Low-level device data were collected from the motion controllers and virtual reality (VR) headset. These data included the x, y, and z values for position, rotation, velocity, angular velocity, acceleration, and angular acceleration of all three devices. Button data from the left and right controllers were also recorded and included thumbstick position, trigger and grip pressure, trigger and grip presses, primary and secondary button presses, and primary and secondary button touches. Data recorded also included the position of the controllers and headset within the virtual environment (VE).

12.2.2 Metrics: Overview

Movement-, button-, and target-related metrics were chosen based on previous literature related to measuring movement differences between individuals and between people with and without impairments (Hwang et al., 2004; Kong et al., 2022; Mackenzie et al., 2001; Pfeuffer et al., 2019). Some metrics were also supported by participants' feedback during interviews in the technique comparison study (Section 10.5). All metrics were technique-agnostic to be able to compare metrics between techniques and groups. Also, some metrics were hypothesized to apply to some techniques and not others because of their control mechanism (e.g., head-related metrics for Chicken and Sliding), so differences seen between groups for some techniques and not others would support the idea that differences were related to interaction with the technique. The metrics explored in this study are summarized in Table 11.

Table 11. The table presents the proposed metric, a description of the metric, its type, and the reason the metric might capture differences between groups. Note: “device” refers to the headset and right or left controller. Metrics are technique-agnostic.

Metric	Type	Description	Rationale
Device variability (Kong et al., 2022)	Movement-related	The total distance the device travelled	Could reflect the efficiency of movement
Device extent (Kong et al., 2022)	Movement-related	The Euclidean distance between the two furthest device positions	Could reflect the range of motion and efficiency of movement
Device angle variability (Kong et al., 2022)	Movement-related	The cumulative angular distance the device rotated	Could reflect the device stability
Device angle extent (Kong et al., 2022)	Movement-related	The angular distance between the two furthest angles	Could reflect the range of motion and device stability
Device velocity and acceleration	Movement-related	The average velocity and acceleration of the device	Could reflect the ability to control movement speed
Device angular velocity and acceleration	Movement-related	The average angular velocity and acceleration of the device	Could reflect the ability to control movement rotation speed
Total distance between device pairs (Pfeuffer et al., 2019)	Movement-related	Total Euclidean distance between device pairs (left-right, left-headset, right-headset)	Could reflect the efficiency of movement
Extent between device pairs (Pfeuffer et al., 2019)	Movement-related	Greatest Euclidean distance between device pairs	Could reflect range of motion
Number of submovements (Hwang et al., 2004)	Movement-related	The number of movements between dips in velocity	Could reflect the efficiency of movement while performing the interaction
Thumbstick distance (Kong et al., 2022)	Button-related	The total distance the thumbstick traveled	Could reflect thumb strength and range of motion
Thumbstick extent (Kong et al., 2022)	Button-related	The angular distance between the two furthest angles	Could reflect thumb range of motion
Trigger and grip pressure	Button-related	Average trigger and grip pressure	Could reflect finger strength and control
Button press count	Button-related	Number of times a button was pressed	Could reflect finger dexterity and isolation
Target re-entry (Mackenzie et al., 2001)	Target-related	Number of times the user exited and re-entered the target	Could reflect the ability to control speed, direction, and position
Target axis crossed count (Mackenzie et al., 2001)	Target-related	Number of times the user crossed the optimal path to the target	Could reflect the ability to control speed, direction, and position
Movement variability (Mackenzie et al., 2001)	Target-related	How straight the users' path is compared to the optimal path	Could reflect the ability to control speed, direction, and position

12.2.3 Movement-Related Metrics

The Device (Angular) Variability and Extent are based on Kong et al.'s (Kong et al., 2022) work on quantifying touch characteristics. Device (Angular) Variability could measure the efficiency with which the user performed the movement. While Device (Angular) Extent could indicate both

efficiency of movement and the range of motion of the user's torsos (headset) and arms (controllers). Interview data suggests that these metrics could capture between-group differences because participants with impairments found it difficult to turn their heads when using Chicken and Sliding, which could be reflected in Head Angular Variability, Extent, Velocity, and Acceleration. Participants also discussed having difficulty reaching with Grab and Throw as well as repetitively performing the motions, which could be related to Controller Variability, Extent, Velocity, and Acceleration. The average Device (Angular) Velocity and Acceleration were found by computing the mean of all values within a trial. The remaining metrics were found using the formulae below:

$$Device\ Variability = \sum_{i=1}^{n-1} \sqrt{(x_i - x_{i-1})^2 + (y_i - y_{i-1})^2 + (z_i - z_{i-1})^2} \in [0, \infty)$$

$$Device\ Extent = \max_{i,j \in 0, \dots, n-1} \sqrt{(x_j - x_i)^2 + (y_j - y_i)^2 + (z_j - z_i)^2} \in [0, \infty)$$

Device Angular Variability and Extent were found in a similar way, where R represents the Euler angle:

$$Device\ Angular\ Variability = \sum_{i=1}^{n-1} \sqrt{(R_{ix} - R_{i-1x})^2 + (R_{iy} - R_{i-1y})^2 + (R_{iz} - R_{i-1z})^2} \in [0, \infty)$$

$$Device\ Angular\ Extent = \max_{i,j \in 0, \dots, n-1} \sqrt{(R_{jx} - R_{ix})^2 + (R_{jy} - R_{iy})^2 + (R_{jz} - R_{iz})^2} \in [0, 360]$$

The Device-Pair Variability and Extent are Pfeuffer et al.'s (2019) work in behavioral biometrics in VR. The authors identified movement-related metrics that captured individual differences between how users moved in an effort to use movement as an authentication method in VR. They found that when used in a predictive model that identified the user, device-pair metrics were used in the optimized feature sets after feature selection, suggesting the predictive power of device-pair metrics. Because these metrics identified differences between individual users, they

could also be used to distinguish users with and without impairments. Device-Pair Variability and Extent were found using the following formulae:

$$Device\ Pair\ Variability = \sum_{i=1}^{n-1} \sqrt{(x_{device1_i} - x_{device2_i})^2 + (y_{device1_i} - y_{device2_i})^2 + (z_{device1_i} - z_{device2_i})^2} \in [0, \infty)$$

$$Device\ Pair\ Extent = \max_{i,j \in 0, \dots, n-1} \sqrt{(x_{device1_i} - x_{device2_j})^2 + (y_{device1_i} - y_{device2_j})^2 + (z_{device1_i} - z_{device2_j})^2} \in [0, \infty)$$




Figure 50. Submovements are movements within a rapid aimed movement with dips in velocity (pauses) before and after their occurrence.

The Number of Submovements metric is based on Hwang et al.'s (2004) work on analyzing movements while motion-impaired users interacted with a mouse. They defined a submovement as a movement within a rapid aimed movement where velocity dips before and after the movement (Figure 50). The Number of Submovements is the sum of submovements during a trial.

12.2.4 Button-Related Metrics

The button-related metrics are not based on prior work but could indicate a number of user abilities including finger isolation, control, range of motion, and strength. Thumbstick Variability and Extent could capture all of the described characteristics. In their interviews, some participants with impairments stated that it was difficult to turn their heads so they would have to compensate by pushing the thumbstick to turn, which would be reflected in Thumbstick Variability and Extent. Participants with impairments also mentioned having difficulty with pressing buttons and reported

that their fingers would sometimes slip off buttons. Thumbstick Variability and Extent were computed with the following formulae:

$$\text{Thumbstick Variability} = \sum_{i=1}^{n-1} \sqrt{(x_i - x_{i-1})^2 + (y_i - y_{i-1})^2} \in [0, \infty)$$

$$\text{Thumbstick Extent} = \max_{i,j \in 0, \dots, n-1} \sqrt{(x_j - x_i)^2 + (y_j - y_i)^2} \in [0, 1.20]$$

Average Grip and Trigger Pressure were computed by finding the average of pressure values during a trial. Button Press Counts were the sum of trigger, grip, primary or secondary buttons presses during a trial.

12.2.5 Target-Related Metrics

Target-related metrics are based on MacKenzie et al.'s (2001) work on evaluating the accuracy of pointing devices and are adapted to a 3D environment. These metrics could indicate a participant's ability to control their position, direction, and speed using the locomotion technique by comparing their path to the optimal path, which is a straight line from the start to end targets in the testbed environment (Figure 51). During interviews, some participants with impairments reported that they had poor control over navigation with Chicken and Throw and had difficulty knowing if they were inside the target with Chicken and Teleport, suggesting that target-related metrics could reveal why people with impairments performed differently than people without impairments. The first metric, Movement Variability (Figure 51), represents "the extent to which the sample points lie in a straight line along an axis parallel to the task axis," (Mackenzie et al., 2001) and is computed with the following formula:

$$\text{Movement variability} = \sqrt{\frac{\sum_{i=1}^{n-1} y_i - \bar{y}}{n-1}}$$




Figure 51. In this image, the user's path is the blue line, and the optimal path is the horizontal line. Movement Variability is a measure of how much the user's path differs from the optimal path. The figure also depicts a user first entering, then exiting, then re-entering a target. In this case, the Target Re-entry Count would be 1. Additionally, the figure depicts the user crossing the optimal path to the target. In this case the Axis Crossed Count would be 2.

The second metric, Target Re-entry Count, is the sum of instances the participant exited the target and re-entered it (Figure 51). The third metric, Axis Crossed Count, is the number of times the participant crossed the optimal path to the target (Figure 51).

12.2.6 Design and Analysis

The between-subjects factor in the analysis was the same as that used in the group comparison study (Section 11.1.4). Group had two levels: With Impairment and Without Impairment. The within-subjects factor, Technique, was also the same as that used in the group comparison study. It had six levels: Astral, Chicken, Grab, Sliding, Teleport, and Throw. Eight participants in the With Impairment group who did not use the left controller because they could only grip one controller were excluded from analysis of left-handed and dual-handed metrics. Two participants who were left-handed in the Without Impairment group were excluded from analysis, except for when analyzing headset metrics.

Outcome measures were the metrics discussed above, with one value calculated per metric per trial, resulting in 2,159 data points for the left hand, 2,721 data points for the right hand, and 2,763 data points for the head. An Anderson-Darling test (Anderson & Darling, 1954) was run on the residuals of a repeated measures full-factorial ANOVA models for Group $\times$ Technique. The tests were statistically significant for every metric, indicating that the residuals did not conform to a normal distribution. The nonparametric aligned rank transform procedure was used (Higgins &

Tashtoush, 1994; Wobbrock, Findlater, et al., 2011) with a linear mixed model (Frederick, 1999) for the analysis.

12.3 Results

The analysis revealed some interesting differences in movement-, button-, and target-related metrics between Groups that could explain differences in performance as well as provide researchers with a technique-agnostic way of identifying whether a user has an impairment. Non-significant results and results that were unclear how to interpret are available in Appendix A.6.

12.3.1 Movement-Related Metrics: Device Variability

There was a statistically significant effect of Technique and Group on Head Variability ($F(5,2780.2)=29.99, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly higher Head Variability for all techniques except for Chicken (Figure 52).




Figure 52. Participants With Impairments had significantly higher Head Variability for all techniques except Chicken. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

There was a statistically significant effect of Technique and Group on Left Variability ($F(5,2297.3)=10.18, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly higher left controller variability for Astral, Sliding, and Teleport (Figure 53).




Figure 53. Participants With Impairments had significantly higher Left Controller Variability for Astral, Sliding, and Teleport. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

There was a statistically significant effect of Technique and Group on Right Variability ($F(5,2786.2)=11.87, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly higher Right Controller Variability for all techniques except Grab and Throw (Figure 54).




Figure 54. Participants With Impairments had significantly higher Right Controller Variability for all techniques except Grab and Throw. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

12.3.2 Movement-Related Metrics: Device Extent

There was a statistically significant effect of Technique and Group on Head Extent ($F(5,2778.3)=15.56, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly higher Head Extent for all techniques (Figure 55).




Figure 55. Participants With Impairments had significantly higher Head Extent values. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

There was a statistically significant effect of Technique and Group on Left Extent ($F(5,2296.2)=21.80, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly higher Left Controller Extents for all techniques except Grab and Throw (Figure 56).




Figure 56. Participants With Impairments had significantly higher Left Extent for all techniques except Grab and Throw. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

There was a statistically significant effect of Technique and Group on Right Extent ($F(5,2785.7)=67.79, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly higher Right Controller Extents for all techniques except Grab and Throw (Figure 57).




Figure 57. Participants With Impairments had significantly higher Right Extents for all techniques except Grab and Throw. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

12.3.3 Movement-Related Metrics: Device Angle Variability

There was a statistically significant effect of Technique and Group on Head Angle Variability ($F(5,2780.6)=4.72$, $p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly higher Head Angle Variability for Teleport (Figure 58).




Figure 58. Participants With Impairments had significantly higher Head Angle Variability for Teleport. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

12.3.4 Movement-Related Metrics: Device-Pair Variability

There was a statistically significant effect of Technique and Group on Right Hand to Head Variability ($F(5,2785.9)=40.21, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly higher Right Hand to Head Variability for Chicken, Grab, and Throw (Figure 59).




Figure 59. Participants With Impairments had significantly higher Right Hand to Head Variability for Chicken, Grab, and Throw. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

12.3.5 Movement-Related Metrics: Device Velocity and Acceleration

There was a statistically significant effect of Technique and Group on Head Velocity ($F(5,2779.9)=60.25, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly lower Head Velocity for Chicken (Figure 60).




Figure 60. Participants With Impairments had significantly lower Head Velocity for Chicken. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

There was a statistically significant effect of Technique and Group on Head Acceleration ($F(5,2779.9)=49.44, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly lower Head Acceleration for all techniques (Figure 61).




Figure 61. Participants With Impairments had significantly lower Head Acceleration for all Techniques. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

There was a statistically significant effect of Technique and Group on Right Velocity ($F(5, 2673.1)=101.53, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly lower Right Velocity for Chicken and Throw (Figure 62).




Figure 62. Participants With Impairments had a significantly lower Right Velocity for Chicken and Throw. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

There was a statistically significant effect of Technique and Group on Right Acceleration ($F(5, 2673.1)=140.18, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly lower Right Acceleration for Throw (Figure 63).




Figure 63. Participants With Impairments had a significantly lower Right Acceleration for Throw. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

12.3.6 Movement-Related Metrics: Device Angular Velocity and Acceleration

There was a statistically significant effect of Technique and Group on Angular Head Velocity ($F(5,2782.1)=19.08, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly lower Angular Head Acceleration for Chicken and Sliding (Figure 64).




Figure 64. Participants With Impairments had significantly lower Angular Head Velocity for Chicken and Sliding. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

There was a statistically significant effect of Technique and Group on Angular Head Acceleration ($F(5,2780.1)=32.41$, $p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly lower Angular Head Acceleration for all techniques (Figure 65).




Figure 65. Participants With Impairments had significantly lower Angular Head Acceleration for all Techniques. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

There was a statistically significant effect of Technique and Group on Angular Right Velocity ($F(5,2781.4)=146.99, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly lower Angular Right Velocity for Grab and Throw (Figure 66).




Figure 66. Participants With Impairments had significantly lower Angular Right Velocity for Grab and Throw. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

12.3.7 Movement-Related Metrics: Submovement Count

There was a statistically significant effect of Technique and Group on Right Submovement Count ($F(5,2780.9)=28.04$, $p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly higher Right Submovement Counts for Grab and Throw (Figure 67).




Figure 67. Participants With Impairments had significantly Higher Right Submovement Counts for Grab and Throw. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

12.3.8 Button-Related Metrics: Thumbstick Variability and Extent

There was a statistically significant effect of Technique and Group on Right Thumbstick Variability ($F(5,2787.7)=11.26$, $p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly higher Right Thumbstick Variability for Teleport (Figure 68).




Figure 68. Participants With Impairments had significantly higher Right Thumbstick Variability for Teleport. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

There was a statistically significant effect of Technique and Group on Right Thumbstick Extent ($F(5,2787.2)=11.15, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly higher Right Thumbstick Extents for Teleport (Figure 69).




Figure 69. Participants With Impairments had significantly higher Right Thumbstick Extent for Teleport. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

12.3.9 Button-Related Metrics: Device Pressure

There was a statistically significant effect of Technique and Group on Right Grip Pressure ($F(5, 2679.0)=35.85, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly higher Right Grip Pressure for Astral, Chicken, and Sliding (Figure 70).




Figure 70. Participants With Impairments had a significantly higher Right Grip Pressure for Astral, Chicken, and Sliding. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

12.3.10 Target-Related Metrics: Movement Variability, Target Re-Entry Count, and Axis Crossed Count

There was a statistically significant effect of Technique and Group on Movement Variability ($F(5,2791.1)=4.32, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly higher Movement Variability for Astral (Figure 71).




Figure 71. Participants With Impairments had significantly higher Movement Variability for Astral. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

There was a statistically significant effect of Technique and Group on Axis Crossed Count ($F(5,2791.8)=3.81, p<.01$). *Post hoc* comparisons indicated that people With Impairments had significantly higher Axis Crossed Count for Grab (Figure 72).




Figure 72. Participants with impairments had significantly higher Axis Crossed Count for Grab. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

There was a statistically significant effect of Technique and Group on Target Re-entry Count ($F(5,2791.1)=5.44, p<.001$). *Post hoc* comparisons indicated that people With Impairments had significantly higher Target Re-entry Count for Astral and Grab (Figure 73).




Figure 73. Participants With Impairments had significantly higher Target Re-entry Count for Astral and Grab. Means are at the bottom of bars. Significance codes: $p<.001$ (***), $p<.01$ (**), $p<.05$ (*).

12.4 Summary

Head-related metrics revealed differences between groups regardless of technique. Other metrics can explain differences in performance found in the group comparison study for all techniques.

12.4.1 Head-Related Metrics Reveal Group Differences

There were several head-related metrics that revealed Group differences across Techniques. Head Extent was greater for people With Impairments compared to people Without Impairments for all Techniques, suggesting that people With Impairments moved their headset a greater distance between two points in space than people Without Impairments (Figure 55). People With

Impairments could have been less efficient with their movements, been less balanced, and moved their whole bodies compared to people Without Impairments when using locomotion techniques.

Head Acceleration and Angular Acceleration were greater for people Without Impairments across all techniques, suggesting that this group was able to react more quickly and with greater intensity compared to people With Impairments (Figure 61 and Figure 65). This finding also suggests that people With Impairments could have moved less efficiently because they were not able to change their speed as quickly as people Without impairments. Head Extent, Acceleration and Angular Acceleration could reveal that people With Impairments were less efficient with their movements compared to the other Group, explaining performance differences between the Groups.




Figure 74. Green lines represent head positions of people Without Impairments and blue lines represent those of people With Impairments. Left: Head positions of all 40 participants while using Chicken. There was not a significant difference in Head Variability for Chicken. Right: Head positions of 40 participants while using Teleport. People Without Impairments had a significantly lower Head Variability, which is reflected in the smaller spread of green compared to blue head positions for Teleport.

12.4.2 Unintentional Movements Differentiate Groups

A related finding is that participants With Impairments appeared to make more unintentional movements. Participants With Impairments moved their headset significantly more with all Techniques except for the Technique that was controlled exclusively with head movement, Chicken. Figure 74 depicts all 40 participants' head positions when using the technique. The head position of the two groups looks similar for Chicken (Figure 74, left), however, people Without

Impairments moved their headsets less than people With Impairments when using Teleport (Figure 74, right).

This same pattern was apparent for Left and Right controller Variability and Extent (Figure 53, Figure 54, Figure 56, and Figure 57). People With Impairments moved their controllers a greater distance than people Without Impairments when they were using their controllers to perform techniques that did not require movement (i.e. Astral, Chicken, Sliding, and Teleport). Prior research found that older adults can overcome noise in their movements by applying greater force (Walker et al., 1997). Also, people with intention tremor performed ballistic movements in a similar way to people without tremor but made more errors during the corrective phase (Flowers, 1976). These two findings suggest that noise in movement can be masked by intentional force, which could be an explanation for why intentional movements did not appear different between groups in this study.

The presence of unintentional movements might also explain why participants With Impairments had significantly higher Right Grip Pressure for Astral, Chicken, and Sliding. These three techniques did not require participants to use the grip or trigger buttons (which are close together) and as a result, participants could have unintentionally pressed the grip button more than people Without Impairments when using these locomotion techniques. This explanation is also likely because many participants gripped the controller with both hands and could have pressed the grip button with their left hand unknowingly (the grip button is on the left-hand side of the right controller).

Although unintentional movements do not explain performance differences between groups, they can provide researchers with a way of identifying if a user is likely to have an impairment through movement-related metrics.

12.4.3 Explaining Performance Differences: Astral

Participants With Impairments had significantly greater Movement Variability when using Astral compared to participants Without Impairments. This finding could be a result of participants having less fine motor control and dexterity with their thumbs, leading to a greater overall deviation from the optimal path compared to users Without Impairments. This finding is also supported by interview responses, where participants With Impairments reported having difficulty keeping their fingers on buttons. Participants With Impairments had a significantly greater Target Re-entry Count, which could also be a consequence of having less thumb dexterity, resulting in less control over the avatar's position.

12.4.4 Explaining Performance Differences: Chicken

People Without Impairments had a higher Head Velocity and Angular Head Velocity when using Chicken compared to participants With Impairments (Figure 60 and Figure 64), suggesting that participants Without Impairments performed better with Chicken than participants With Impairments because they were able to move and turn their heads faster, and as a consequence, also turn their virtual speed and direction faster. Although Head Variability was not significantly different for Chicken, Right Hand to Head Variability was, which could suggest that participants With Impairments made more unintentional movements with their right hands when using Chicken, increasing the Total Distance Between the Device Pair.

12.4.5 Explaining Performance Differences: Grab

Several metrics revealed differences between Groups for Grab. There were not any significant differences between Groups for Left and Right Variability, which is consistent with the finding that intentional movements are similar for both Groups. However, participants With Impairments

had significantly higher Head Variability for Grab, which could explain the significantly greater Variability between Left and Right controllers and the Headset for participants With Impairments.

The difference in performance could be explained by participants With Impairments making more unintentional and less precise movements with their heads or torsos relative to their controllers when using Grab. Participants With Impairments made significantly more submovements with their right controllers when using Grab compared to the other group, additional evidence that participants were less fluid with their controller movements. This finding is also supported by prior research which reported that older adults broke up their movement into a greater number of submovements compared to younger adults as a way of compensating for the higher noise in their movements (Walker et al., 1997).

It is possible that the same effect was not visible with the Left Submovement Counts because many participants With Impairments did not use their left hands compared to people Without Impairments, and the ones who did may have been more efficient with their left controllers, which is why they used them in the first place. Participants Without Impairments had greater Right Angular Velocity, which indicates that they were able to re-orient their controller faster when doing the grab and pull motion compared to people With Impairments.

In terms of target-related metrics, participants With Impairments had a greater Axis Crossed Count and Target Re-entry Count. This result suggests that participants were less able to control their position in the VE with the grab and pull movement, likely because most participants only used one arm to grab and pull making their movements unbalanced and requiring more movement from the right arm, which could have caused fatigue.

12.4.6 Explaining Performance Differences: Sliding

The only metric that showed a significant difference between Groups for Sliding was Angular Head Velocity. Participants Without Impairments were able to turn their heads faster when using Sliding, which is also supported by qualitative feedback from participants With Impairments. They reported feeling like they were not able to turn their virtual bodies fast enough with Sliding. However, this difference must not have had a large impact on performance because there were no significant differences on Trial time, Hit Rate, or Obstacles Hit between groups in the group comparison study.

12.4.7 Explaining Performance Differences: Teleport

Participants With Impairments had greater Head Angle Variability as well as Right Thumbstick Variability and Extent than the other Group. When looking back at screen recordings of the participants With Impairments, several participants overshot targets when using Teleport. Interview findings also revealed that participants struggled to see where the target was when they were close to it and had to turn their heads or re-orient their virtual bodies using the thumbstick to face the target. These metrics reveal differences between Groups in the corrective phase of the interaction, which is consistent with prior research that found that people with tremor made more corrections during the corrective phase of movement than people without tremor (Flowers, 1976).

12.4.8 Explaining Performance Differences: Throw

Like Grab, several metrics revealed differences between Groups, which could indicate why participants With Impairments performed worse compared to the other Group when using Throw. Participants With Impairments had a significantly greater Left and Right Hand to Head Variability, likely because they made more inefficient movements with their heads or torsos when using the technique compared to participants Without Impairments. This same pattern existed for Grab

because both Techniques required gross motor movement with the arms. Like with Grab, participants With Impairments had more Right-handed Submovements compared to the other group, again suggesting they were less fluid with gross motor movements. Also, participants Without Impairments had a greater Right Angular Velocity, suggesting that they were able to re-orient their controller faster than people With Impairments during the throw motion, which could have impacted performance.

12.4.9 Limitations and Future Work

Not all techniques were dual handed, which made it difficult to compare left and right metrics across techniques. In addition, many participants with impairments did not use the left controller, which could have made participants With Impairments who used it appear similar to participants Without Impairments, leading to significant differences in the right-handed metrics but not in the left-handed ones. Future work could adapt Grab and Throw to be single-handed techniques, as well as require all participants to use only one controller to eliminate the potential confound of technique handedness.

Future analysis could also investigate how metric values vary across impairment levels based on QuickDASH scores, where Impairment Level is used as an independent variable. The current analysis did not require a finer grained analysis of impairment level to reveal interesting results, but could be relevant to other studies, for example, to investigate how much a metric is able to distinguish levels of impairment.

12.4.10 Summary

Head-related metrics revealed differences between groups across all techniques. Device Variability and Extent consistently revealed that participants with impairments made more unintentional movements compared to people without impairments. However, Device-Pair

Variability captured differences between groups when Device Variability did not, likely because when participants were making intentional movements with one device (e.g., right controller with Grab and Throw), they were also making unintentional movements with another device (e.g., headset with Grab and Throw). Differences in performance for all techniques might be explained by movement-, button-, and target-related metrics, which could serve as features in a model that predicts performance with techniques for new VR users.

13 PREDICTING PERFORMANCE WITH LOCOMOTION TECHNIQUES BASED ON QUESTIONNAIRE AND CALIBRATION TASK INPUT

Over one hundred distinct locomotion techniques exist to this day (Di Luca et al., 2021). If a user is new to virtual reality (VR) and has an impairment, it could be challenging to try multiple locomotion techniques before finding one that is most appropriate for his or her abilities. A recommender system could help new users overcome this barrier by suggesting a technique based on information about their abilities.

Results from the technique comparison study (Chapter 10) indicated that the optimal technique in terms performance varied for individual participants. Teleport was the fastest technique for most participants with impairments ($n=14$). However, the fastest technique was Astral for three participants, Sliding for two participants, and Chicken for one participant. There was less variability in the group of participants without impairments: Teleport was the top technique for the majority of participants, but Grab was the top technique for one participant, and Chicken for another.

In terms of accuracy (Hits), there was less variability between participant groups compared to Trial Time: most participants hit all six targets in a trial. However, many participants with impairments did not achieve perfect accuracy, suggesting there could be enough variability in this group to predict Hits for all participants.

There were no significant differences between groups in terms of Movement Variability, with the exception of Astral. However, a predictive model trained on different types of participant data could pick up on individual differences, making Movement Variability a good performance measure. Movement Variability could reflect a user's ability to control their virtual position and orientation with locomotion techniques.

Given the inconsistency between participants in terms of the optimal technique as well as the potential to model different aspects of performance, there is an opportunity to personalize the VR experience and make it more accessible by recommending the most appropriate technique for a new user based on their abilities. There are two scenarios of interest in which a new VR user could be recommended a locomotion technique.

In one scenario, the questionnaire scenario, the VR user would respond to a questionnaire about their upper-body impairments. Using this data as input, the model would then predict the optimal techniques for different measures of performance. A recommender system could then produce a ranked list of techniques that the user would likely be able to use well according to their abilities.

In the second scenario, the calibration scenario, a new user would respond to a questionnaire about their impairments, similar to the initial step in the questionnaire scenario. Afterwards, the user would perform a calibration task with a locomotion technique and then fill out a second questionnaire about their experience using the technique. Based on this input, a recommender system could produce a prioritized list of techniques, mirroring the format of the questionnaire scenario.

This chapter reports on a study, which will be referred to as the prediction study, that uses machine learning techniques to build predictive models for the scenarios described. First, the process for building the predictive model based on questionnaire data is reported followed by the process for the process for the calibration scenario. Lastly, findings from and a comparison of the two scenarios are discussed.

13.1 Questionnaire Scenario: Locomotion Technique Recommendation Based on Questionnaire Responses

In the questionnaire scenario, a new VR user would fill out a questionnaire about the type of upper-body impairment they experience and the severity of their condition. The recommender

system would then use this data to predict the locomotion techniques that are most appropriate for their abilities.

13.2 Questionnaire Scenario: Study Method

This section discusses the method used to create the predictive models and includes the outcome variables, datasets, feature selection method, machine learning algorithms, and model evaluation approach. The data used for this part were collected in the group comparison study described in Section 11.1.

13.2.1 Model Features

The individual responses to QuickDASH and Technology-Related Physical Impairment Questionnaire (TRIQ) questions as well as the computed QuickDASH score (described in Section 11.1) were used as model features. There was a total of 32 features, including locomotion technique. The full QuickDASH questionnaire is available in Appendix A.3 and the full TRIQ questionnaire is available in Appendix A.4.

13.2.2 Performance Measures

Three outcome predictions that captured different facets of performance were used including Trial Time, Hit Rate, and Movement Variability. The same performance measures were used as those in the technique and group comparison studies with the exception of Obstacles Hit. Obstacles Hit was not useful for quantifying performance because it likely captured how far a participant travelled rather than their ability to navigate with the locomotion technique (see Section 11.4.3 for a detailed explanation). Movement Variability, which was proposed as a target-related metric in Section 12.3.10, measures how straight a user's path is compared to the optimal path and could indicate a user's ability to control the technique. Movement Variability was chosen as the third outcome prediction for this reason.

Outcome predictions were aggregated at the technique level by finding means across all 12 trials (2 trial blocks $\times$ 6 trials) for each participant.

13.2.3 Dataset

- Features: There was a total of 32 features, which included the locomotion technique, responses to individual QuickDASH and TRIQ questions, as well as the overall QuickDASH score.
- Instances: Each instance represented one technique performed by one participant, for a maximum of 240 instances (40 participants $\times$ 6 techniques). However, three participants skipped a questionnaire response, resulting in the final dataset containing 237 instances.

A snippet of the dataset is available in Appendix A.7.

13.2.4 Modelling Method

Linear regression with elastic net regularization was used to model Trial Time and Movement Variability. A regularized random forest was used to model Hit Rate. The simplicity of the algorithms used coupled with regularization techniques provided a safeguard against overfitting the dataset.

Elastic net regularization discourages reliance on any single feature or small sets of highly correlated features by penalizing large coefficients, making the model less sensitive to noise in the training data (Zhang et al., 2017). Regularization when building a Random Forest tends to produce sparsity, so only the most important rules and features are retained (Liu et al., 2014).

I performed a grid search of tuning parameters for each model, resulting in the highest performing models in terms of R^2 . The `train` function in the R `caret`³¹ package was used to train the models.

³¹ <https://cran.r-project.org/web/packages/caret/caret.pdf>

13.2.5 Feature Selection

Feature selection was performed with `caret`'s implementation of recursive feature elimination (RFE) with the same algorithm used to train the model as the base algorithm. RFE, which is a wrapper-based feature selection technique was used because wrapper-based techniques are generally better compared to filter-based techniques at selecting the correct features by removing correlated features (Bee Wah et al., 2018). Because there will likely be correlated features in this dataset, such as the QuickDASH score and question responses, a wrapper-based feature selection technique was preferable.

13.2.6 Model Evaluation

Cross validation was performed with 25 to 27 folds, in which one or two participants were left out of the training set to estimate the performance of the model with unseen data. Folds were created using `caret`'s `groupKFold` function, which created training and testing sets based on participant ID. The average R^2 was computed across folds for an estimate of the model's accuracy. R^2 was selected as the performance metric because it is easily interpretable with values ranging from 0 to 1, where higher values equate to better performance. The R^2 value represents the amount of variability in the data that the model can explain. The Bayesian Information Criterion (BIC) was also used to compare models in terms of model fit and complexity. BIC is effective for comparing the fit of different models, however, the absolute value of the BIC does not have inherent meaning. A lower BIC indicates a better model fit relative to other models. BIC is more conservative measure than Akaike Information Criterion (AIC), another measure of model fit, because it penalizes models with many parameters. Root mean squared error and mean absolute error were also recorded to capture the average magnitude of prediction errors.

13.2.7 Ranking Techniques

To rank techniques for each participant from best to worst, a ten-fold cross validation was conducted, and the outcome measure was predicted for each hold-out set. Leave-one-participant-out (LOPO) cross validation tends to be overly pessimistic with small datasets as test sets are less representative of the data as a whole and do not generalize well (Hurst et al., 2008; James et al., 2021). To reduce variability in training and test sets, I used k-fold cross validation, which can lead to models that have more reliable performance estimates compared to models trained with LOPO cross validation. After extracting predictions from the k-fold cross validation, predicted and actual values were used to rank each technique from 1 to 6, where 1 represented the fastest technique and 6 represented the slowest technique. Rankings for actual Trial Times and predicted Trial Times were recorded in a dataset, which also contained participant ID and technique. Mean absolute error (MAE), root mean squared error (RMSE), and mean absolute percent error (MAPE) were also calculated.

13.3 Results: Questionnaire Scenario

In this section I discuss findings related to building predictive models for Trial Time, Hit Rate, and Movement Variability based on questionnaire input.

13.3.1 Trial Time: Overview

After applying a logarithmic transformation over Trial Time, a Kolmogorov-Smirnov test revealed that the distribution was not significantly different from a normal distribution ($D=0.07$, $p=0.23$). So, the logarithm of Trial Time was used as the outcome predictor in the questionnaire scenario. For the sake of brevity, the outcome predictor will be referred to Trial Time for the remainder of the section.

There was low variability within feature values because they were repeated across techniques for each participant and many participants without impairments had similar feature values, making the model rank deficient when using a linear mixed modelling method. Elastic net regularization proved to be helpful in overcoming this issue by adding a small amount of bias to reduce the variance of estimates.

Table 12. A table of the tuning parameters for the elastic net (alpha and lambda³²), R² values, Bayesian Information Criterion (BIC), root mean squared error (RMSE), and mean absolute error (MAE) for models predicting Trial Time with all questionnaire features, only QuickDASH features, only TRIQ features, and the selected features from feature selection.

	Alpha	Lambda	R ²	BIC	RMSE	MAE
All features	0.01	0.0031	0.52	25243.04	0.19	0.15
QuickDASH features	0.55	0.0312	0.54	10725.92	0.20	0.16
TRIQ features	0.10	0.0027	0.52	13830.63	0.22	0.18
Selected features	0.55	0.0031	0.58	12908.48	0.21	0.17

The four models had similar R² and error values. The model trained on the selected feature set performed the best (R²=0.58). However, the model trained on only QuickDASH features had the lowest BIC, indicating a better fit than the model built with the selected features. It also had a lower RMSE and MAE than the model built with the selected feature set.

13.3.2 Trial Time: Feature Selection

The RFE determined that 20 of the 32 features resulted in the best performing models. The top ten features were all from the QuickDASH questionnaire, with the exception of locomotion technique (Figure 75).

³² The lower value of alpha spreads the penalty more evenly across the coefficients, which can be beneficial when many predictors have a small or moderate effect. A higher value of lambda increases the regularization strength, leading to smaller values of model coefficients. This can help reduce overfitting by making the model less sensitive to noise in the training data.




Figure 75. A plot of the 20 features selected with recursive feature elimination ordered by their importance ranking. Technique was highly ranked compared to all other features.

Technique had an importance ranking that was about three times that of the second most important feature, painQD, which suggests that technique was much more important to the model's performance compared to all other features. This finding could indicate that Trial Time could be predicted with some accuracy by technique alone, without knowing a users' impairment status. Training a model for Trial Time with technique as the only feature resulted in an R^2 of 0.22, indicating that impairment-related features were still critical to the model's performance.

Features with less importance than pain from the TRIQ questionnaire, which was the thirteenth feature, did not contribute much more to the model performance than did previous features (Figure 75), so pain was chosen as the cutoff, resulting in a selected feature set of 13 features.

The fact that painQD was the first QuickDASH feature, and pain was the first TRIQ feature selected by the RFE revealed the significance of pain on the participants' Trial Times. PainQD

referred to the severity arm, hand, or shoulder pain participants experienced (see Appendix A.3 for QuickDASH questionnaire), while pain from the TRIQ questionnaire referred to whether the participant experienced pain in their core, neck, shoulders, arms, or hands (see Appendix A.4 for TRIQ).

13.3.3 Trial Time: Training a Model

Interpreting a density plot of R^2 values, a learning curve, and a plot of actual vs. predicted values revealed more about the fit of the model of Trial Time trained on the QuickDASH feature set.




Figure 76. A density plot of R^2 values from cross-validation. There is high variability, but the distribution is skewed toward higher R^2 variables suggesting a bias towards a subset of participants.

The density plot of R^2 values from the cross validation revealed that model performance varied, but was skewed towards higher R^2 values, indicating that the model could be slightly biased toward a subset of participants, likely the participants without impairments (Figure 76). Adding more participants could lead to less variability in the model's accuracy.




Figure 77. A learning curve showing how performance increases with additional participants. R^2 values increase by 0.10 from 15 to 40 participants.

The learning curve generated by averaging R^2 values across models trained on increasingly larger numbers of participants showed only a slight increase in values from 0.50 to 0.60 as the number of participants in the training set increased (Figure 77). This result could indicate that the performance is unlikely to make large improvements with a few more participants. The model performance could be improved by adding several more participants, particularly people with impairments, to increase the variability in the data. Improvements could also be made by introducing features, such as the metrics discussed in the metric exploration study in Chapter 12.




Figure 78. A plot of actual vs. predicted values, where the line represents perfect prediction.

The plot of actual vs. predicted values suggests that the model underestimates the value of higher trial times, indicating that the performance could improve with more data from people with impairments (Figure 78).

13.3.4 Trial Time: Ranking Techniques

Because the QuickDASH feature set led to a high performing model with a better fit than the model built with the selected features, it was used to rank the techniques. Appendix A.8 presents the actual and predicted Trial Times by participant and technique. The MAE was 2.63 seconds, the RMSE was 3.65 seconds, and the MAPE was 15.0%.




Figure 79. Slope graph of actual and predicted techniques by rank from fastest to slowest (rank 1-6) for the questionnaire scenario. The y-axis is alphabetically ordered.

The model predicted the same technique within ranks 1-3, indicating that it was not able to model subtle differences between participants by their self-reported impairments for these ranks (Figure 79). So, although the model was moderately accurate for predicting the Trial Time, achieving an R^2 of 0.54, it was mostly basing its estimates on technique, which is also what the RFE results showed. The learning curve indicated that performance would not increase much with more participants, which is also likely because the model gave technique a large weight. The fact that technique alone could not explain as much of the variance of Trial Time as the selected features indicates that self-reported impairments did help the model predict Trial Time more accurately, but the model's relatively high performance did not translate to technique ranking.




Figure 80. Left: Heat map of overall ranking accuracy for all participants, middle: participants with impairments, right: and participants without impairments. The model accurately predicted the top three techniques for participants with impairments, but only the top technique for participants without impairments.

Heat maps showed that the model learned mainly from people with impairments for ranks 1-3 and applied the same rule to people without impairments (Figure 80). The model predicted with 70% accuracy that the top three techniques for people with impairments were Teleport, Astral, and Sliding (Table 13). Predicting the same three techniques for people without impairments led to less accuracy than people with impairments, but also lower MAPE (Table 13). This indicates that the model's predictions were close to the actual values on average, but it was making wrong predictions for a portion of the data.

Table 13. Table of ranking accuracy and mean absolute percent error (MAPE) for participants with impairments, participants without impairments, and all participants. The overall ranking accuracy and MAPE of the six techniques was calculated, as well as the accuracy and MAPE for each rank individually.

	Overall Ranking Accuracy (MAPE)	Top 1 Technique Accuracy (MAPE)	Top 2 Technique Accuracy (MAPE)	Top 3 Technique Accuracy (MAPE)	Top 4 Technique Accuracy (MAPE)	Top 5 Technique Accuracy (MAPE)	Top 6 Technique Accuracy (MAPE)
With Impairments	47% (17%)	70% (25%)	70% (14%)	70% (15%)	15% (17%)	25% (12%)	30% (19%)
Without Impairments	39% (13%)	89% (16%)	63% (13%)	32% (10%)	21% (13%)	11% (9%)	16% (16%)
All	43% (15%)	79% (21%)	67% (13%)	51% (12%)	18% (15%)	18% (11%)	23% (17%)

13.3.5 Hit Rate: Overview

Hit Rate did not follow a particular distribution as a continuous variable, so a regularized random forest (RRF) was used to train the model. RRF is suitable algorithm for modelling Hit Rate because it is relatively simple, does not make strict assumptions about the distribution of the outcome predictor, and regularization safeguards against overfitting the data (Liu et al., 2014). As expected, the optimal feature set selected using RFE resulted in the highest performance ($R^2=0.38$) and only had a slightly higher BIC compared to the QuickDASH model (Table 14).

Table 14. A table of the tuning parameters for regularized random forest: Number of randomly selected predictors (numPred)³³, regularization value (coefReg)³⁴, and importance coefficient (coefImp)³⁵, R^2 values, root mean squared error (RMSE), and mean absolute error (MAE) for models predicting Hit Rate with all questionnaire features, only QuickDASH features, only TRIQ features, and selected features from the feature selection.

	numPred	coefReg	coefImp	R^2	BIC	RMSE	MAE
All features	2	0.01	0.0	0.35	20502.95	0.16	0.13
QuickDASH features	2	1.00	0.5	0.36	20439.83	0.15	0.12
TRIQ features	2	0.01	0.0	0.37	20541.12	0.16	0.13
Selected features	2	1.00	0.0	0.38	20441.39	0.15	0.12

13.3.6 Hit Rate: Optimal Features

The RFE indicated that sixteen features resulted in the most parsimonious models of Hit Rate (Figure 81).

³³ The number of randomly selected predictors is the number of features that are considered at each split of the tree. A smaller number prevents overfitting by increasing the randomness of the model.

³⁴ The regularization value indicates how much the model is penalized for the inclusion of additional features or additional splits in the tree.

³⁵ The importance coefficient refers to how much weight is given to features in the model training process.

RecreationalActivitiesQD, or the ability of the respondent to perform recreational activities in which they take some force or impact through their arm, shoulder, or hand (e.g., golf, hammering, tennis, etc.) was the most important feature for predicting Hit Rate. This finding also makes intuitive sense because Hit Rate is a measure of accuracy, which many recreational activities require, such as sports. In terms of the locomotion techniques, Grab and Throw required the participant to be accurate with gross motor movements, while Teleport required the participant to aim their controller accurately.




Figure 81. A plot of the top 28 features ordered by their importance ranking found during feature selection. Recursive feature elimination resulted in a final set of the top 16 features.

Distance was the most important TRIQ feature and referred to whether or not the respondent experienced difficulty controlling the distance of their arm movement. This finding also makes

intuitive sense because Grab, Throw, and Teleport required the ability to reach or aim to hit targets successfully. Isolation, which was the second-most important TRIQ feature related to how well participants could isolate their fingers, which would affect their ability to control the snap turn on all of the techniques with the thumbstick.

Features that were ranked lower than difficultySleepingQD had limited importance to the model. So difficultySleepingQD was chosen as the cutoff point, leaving a total of 13 features. There were three TRIQ features in the final feature set and the remaining features, besides technique, were from the QuickDASH questionnaire.

13.3.7 Hit Rate: Training a Model

The selected feature set was used to train the final model for Hit Rate because it had the highest performance, one of the lowest BIC, and accounted for correlated features with RFE. The plot of predicted vs. actual values indicated that lower Hit Rates were not modelled accurately (Figure 82). Also, the lack of variability in Hit Rate was evident because most of the data points were clustered towards high Hit Rates, indicating a ceiling effect.




Figure 82. A plot of actual vs. predicted values, where the line represents perfect prediction. Values are clustered towards high Hit Rates indicating a ceiling effect.

The density plot showed the high variability in the model's performance (Figure 83). The distribution was skewed towards lower R^2 values, suggesting that the model underfit the data and could not account for more subtle patterns, likely because it assumed a high Hit Rate regardless of features. The low variability in Hit Rate meant that there were fewer instances with low Hit Rates to learn from, which resulted in underfitting.




Figure 83. A density plot of Hit Rate models' R^2 values from cross validation. There is high variability, and the distribution is skewed toward lower R^2 variables suggesting the model underfit the data.

The learning curve suggests that adding more participants will not increase the model's performance significantly because its performance plateaus at 0.20 after 20 participants (Figure 84). This finding is likely a consequence of the ceiling effect due to the study design rather than particular participant characteristics.




Figure 84. A learning curve showing that Hit Rate prediction performance increases with additional participants, but plateaus around an R^2 of 0.20 after 20 participants.

Overall, the model underfit Hit Rate, likely because of the ceiling effect as a result of the generous trial time limit. A time limit of less than 30 seconds to complete a trial could have reduced the ceiling effect.

13.3.8 Hit Rate: Technique Recommendations

The ceiling effect made it impossible to rank techniques. Most participants achieved a Hit Rate of 1 for each technique. There was not enough variability within an individual participant's Hit Rate to be able to recommend a top technique.

13.3.9 Movement Variability: Overview

After applying a logarithmic transformation to Movement Variability, a Kolmogorov-Smirnov test revealed that the distribution was not significantly different from a normal distribution ($D=0.05$, $p=0.60$). Like Trial Time, linear regression with elastic net regularization was also used to model Movement Variability. Comparing the models trained on the four feature sets revealed that the models performed similarly, but the model trained on the selected features had the best model fit (Table 15).

Table 15. A table of the tuning parameters for the elastic net (alpha and lambda), R^2 values, root mean squared error (RMSE), and mean absolute error (MAE) for models predicting the Movement Variability with all questionnaire features, only QuickDASH features, only TRIQ features, and the selected features from feature selection.

	Alpha	Lambda	R^2	BIC	RMSE	MAE
All features	0.55	0.0035	0.42	52933.44	0.34	0.27
QuickDASH	0.10	0.0035	0.40	26670.28	0.34	0.27
TRIQ	0.10	0.0358	0.40	41851.53	0.35	0.27
Selected features	0.10	0.0348	0.41	24699.34	0.32	0.25

13.3.10 Movement Variability: Optimal Features

The RFE resulted in a final feature set of eight features (Figure 85). Technique had an importance ranking that was much greater than any of the questionnaire features for estimating Movement Variability. Besides technique, the optimal feature set only contained QuickDASH features. Training a model with technique as the only feature revealed that the same amount of variance in Movement Variability could be explained by technique alone ($R^2=0.42$), suggesting that the variance in Movement Variability had little to do with participants' impairments. This finding was also in line with the lack of significant between-group differences in Movement Variability found in the metrics exploration study, with the exception of Astral (Section 12.3.10). Therefore, Movement Variability was not a good measure of participants' performance when using locomotion techniques.




Figure 85. A plot of the top 13 features ordered by their importance ranking found during feature selection. Recursive feature elimination resulted in a final set of the top eight features.

13.3.11 Questionnaire Scenario: Summary

Modelling Trial Time with QuickDASH features, resulted in a model that predicted participants' Trial Times moderately well. However, when ranking the predicted Trial Times and comparing them to their actual Trial Times, the model predicted the same technique within each of the top three ranks. If a researcher or designer wanted to predict the Trial Time for a new user, this model could be useful, but it does not capture enough individual differences to work as a recommender system. Questionnaire responses do not appear to encode enough information about participant characteristics to explain the variance in Trial Time.

Because of the ceiling effect that was a result of the study design, Hit Rate did not have enough variability to result in an acceptable model. Lack of variability also prevented the ability to rank

techniques based on Hit Rate. Movement Variability proved to be a poor proxy for success when using locomotion techniques because it appeared to be unrelated to participants' impairments.

13.4 Calibration Scenario: Using a Locomotion Technique to Calibrate the Recommender System.

Questionnaire data might not provide enough information about a person's ability to use VR, including their ability to hold the controller, sit upright, and push controller buttons. In other words, the questionnaires might not capture abilities that are specific to VR use that could affect users' performance with locomotion techniques. For this reason, having a new user perform a short calibration task in which they use one or two locomotion techniques could inform the model and improve its performance.

In the calibration scenario, the user would fill out a questionnaire about their impairment, and then perform a navigation task with one or two locomotion techniques. After performing the task, the user would fill out a questionnaire related to their experience using the calibration technique.

The input data to the model would be impairment-related questionnaire responses, movement-, button-, and target-related metrics logged during the calibration task, measures of success performing the calibration task (e.g. Trial Time), and responses to the post-task questionnaire. The model would then predict the optimal technique for the new user, which could include the techniques they used to calibrate the system. Having a variety of input data types could provide the model with a multi-dimensional snapshot of a user's abilities, enabling the model to capture subtle differences between users.

13.5 Study Method: Calibration Scenario

The model was trained, and its performance was analyzed to gain insight into whether calibration data could improve performance compared to the use of questionnaires alone.

13.5.1 Outcome Predictions

There was only one outcome prediction in this scenario, which was Trial Time. Hit Rate was not a good measure of performance because of the ceiling effect discussed in the questionnaire scenario (Section 13.3.5), so it was not used in the analysis. Movement Variability was not a good measure of performance either because it did not appear to vary with impairment-related measures (Section 13.3.9).

In this scenario it did not make sense to predict the performance for the calibration technique, because the performance would be known. So, Trial Time was broken up into six outcome measures: Astral Trial Time, Chicken Trial Time, Grab Trial Time, Sliding Trial Time, Teleport Trial Time, and Throw Trial Time. The prediction was made at the trial level.

13.5.2 Dataset

Features: There was a total of 68 features, which included:

- QuickDASH questions and score
- TRIQ questions
- Movement-related metrics
- Button-related metrics
- Target-related metrics
- Performance measures (Trial Time, Hit Rate)
- Post-task questionnaire questions (NASA-TLX)

Left-handed metrics were excluded because many participants with impairments did not use the left controller. A finding in the metrics comparison study was that between-group differences for left-handed metrics were not consistent with differences for right-handed metrics, suggesting that participants with impairments who used their left hand were similar to participants without impairments (Section 12.4.5). So left-handed metrics could introduce noise.

Instances: Each row represented one trial, so there was a maximum of 2808 instances (6 techniques $\times$ 2 trial blocks $\times$ 6 trials $\times$ 39 participants). PC01, a participant without impairments

was excluded from analysis because he did not use all six techniques. Some trials were missing due to technical issues during the study and some questions on questionnaires were skipped by participants, so the final dataset contained 2799 instances.

The training dataset included all participants' data from all trials. A snippet of the dataset is available in Appendix A.9.

13.5.3 Modelling Method

Models of Trial Times were trained for each technique using all techniques for calibration besides the outcome technique that was being predicted, which will be referred to as the five-technique set, as well as a single technique as a calibration technique, which will be referred to as the single-technique set. Models were trained using elastic net regularization and linear regression to avoid overfitting the dataset.

13.5.4 Feature Selection

An optimal feature set was selected using recursive feature elimination (RFE) for each outcome technique based on the five-technique test set. The base algorithm was linear regression. The same features were also used when evaluating the model with the single-technique test set.

Four out of six of the final feature sets contained all 80 features. To avoid overfitting, the top 20 were used to train models. Twenty was arbitrarily selected as the number of selected features, however, this number is well below the 1:10 to 1:20 feature to instances ratio that is generally recommended for regression modelling (Harrell, 2015).

13.5.5 Model Evaluation

During the leave-one-participant-out cross validation, the model was trained on 38 participants' data and tested on data from the assumed calibration technique of the excluded participant. So, if Astral Trial Time was being predicted, the dataset included data from 38 participants, and was

evaluated on data logged when the excluded participant was using the other techniques besides Astral.

The model was evaluated with two different types of test sets. One set contained data from the excluded participants' data from all six techniques, minus the technique that Trial Time was being predicted for (the five-technique set). The other test set contained data from only one technique performed by the excluded participant. This test set reflects the use of one calibration technique in this scenario (the single-technique set).

Although it is not realistic to assume that a new VR user would use five locomotion techniques in a calibration task, models tested on five techniques can serve as baselines to compare other models to. Models tested on single techniques should perform similarly to the models tested on five techniques, which could indicate that the model is not overfitting the smaller dataset.

The same performance measures as those in the questionnaire scenario were used: R^2 , Bayesian Information Criterion (BIC) as well as the root mean squared error (RMSE) and mean absolute error (MAE).

13.5.6 Ranking Techniques

To rank the techniques, Trial Times were predicted for the hold out sets during a ten-fold cross validation. A k-fold cross validation was used for the same reason as described in Section 13.2.7. To investigate whether using one calibration technique was feasible for recommending techniques, only the instances that corresponded to the calibration technique were used to rank techniques. So, since Sliding was determined to be the best calibration technique, values captured when participants were using Sliding were used to rank the other techniques based on their predicted Trial Times. Mean absolute error (MAE), root mean squared error (RMSE), and mean absolute percent error (MAPE) were also calculated.

The means of predicted and actual Trial Times were used to rank each technique from 1 to 6, where 1 represented the fastest technique and 6 represented the slowest technique. Rankings for actual Trial Times and predicted Trial Times were recorded in a dataset, which also contained participant ID and technique.

13.6 Results: Calibration Scenario

In this section I discuss findings related to building predictive models for Trial Time based on calibration and questionnaire data.

13.6.1 Trial Time: Optimal Features

Figure 86 presents the top five features for predicting Trial Time for each technique based on the five-technique dataset. The top two most important features were Movement Variability and Sliding Trial Time. This finding is intuitive because a person's performance with calibration techniques is likely to be similar to their performance with the outcome technique. Movement Variability in this model is likely being treated as a more detailed proxy for technique. Technique explained a considerable amount of variance when predicting Movement Variability and technique had a high importance ranking in the Trial Time model in the questionnaire scenario (Section 13.3.10).




Figure 86. The top five features that were selected with RFE when modelling techniques' Trial Times using all other techniques as calibration techniques. Each type of feature was represented across the six feature sets: Button-, movement-, and target-related features from the metrics exploration study, as well as QuickDASH-, TRIQ-, performance, and post-task questionnaire-related features.

QuickDASH features were present in all feature sets, with choresQD appearing in four out of six feature sets. There were more QuickDASH features than any other feature across the six feature sets, indicating their importance in explaining variance in the data. Only two movement-related features appeared in the top five feature sets: Head Extent and Distance between Right Hand and Head. Head Extent is intuitively important for Chicken and Sliding because these techniques are controlled with the head. Also Distance Between Right Hand and Head is an intuitive feature for Throw because the user extends their arm away from their body to throw the ball.

Button-related features were also selected. Secondary Button Press Count was a top feature for Astral, which is likely because a few participants with impairments struggled to keep their fingers

from sliding off buttons (Section 10.5). As a result, the secondary button could have accidentally been pressed by these participants because it is next to the thumbstick, which was used to control Astral. Average Trigger Pressure was a top button-related feature for Grab. This feature is intuitive because the user pressed the trigger to “grab” the ground.

Interestingly, one post-task questionnaire feature, Effort, was a top feature for Sliding. The group comparison study showed that Sliding was the only technique for which there were no between group differences in Trial Time. It is possible that Effort was used as a proxy for performance because objective between-group differences in Trial Time and Hit Rate did not exist for Sliding.

13.6.2 Trial Time: Training Models

R^2 and BIC values for the models built with the five-technique and single-technique datasets are presented in Table 16. Tuning parameters are available in Appendix A.10 and RMSE and MAE are available in Appendix A.11 and. mModels trained on single-technique datasets tended to perform better than those trained on the five-technique datasets and BIC values were much lower for single-technique models, indicating a better fit (Table 16). Sliding was slightly better as a calibration technique for predicting Trial Time compared to the other five techniques, resulting in models with an average R^2 of 0.36 and BIC of 240317.00.

Table 16. R^2 values for models tested on single and all calibration techniques to predict Trial Time for the outcome technique. Using Sliding as calibration a technique resulted in the most accurate and best fitting models for most techniques' Trial Times.

		Calibration Technique						
		Astral	Chicken	Grab	Sliding	Teleport	Throw	Five Techniques
Predicting Trial Time for	Astral (R^2 , BIC)		0.31 246401.30	0.43 221451.30	0.44 223784.90	0.31 270009.10	0.28 236625.10	0.22 1000518.00
	Chicken (R^2 , BIC)	0.35 262749.80		0.30 264781.90	0.29 269215.90	0.29 240469.10	0.23 265384.00	0.19 1277563.00
	Grab (R^2 , BIC)	0.49 223670.70	0.37 261140.90		0.49 268849.00	0.34 250300.20	0.32 231127.80	0.26 1055080.00
	Sliding (R^2 , BIC)	0.45 219699.60	0.32 230668.90	0.42 211246.60		0.32 203788.40	0.27 242211.40	0.24 1237405.00
	Teleport (R^2 , BIC)	0.30 252557.20	0.32 297061.90	0.23 264155.90	0.30 211014.80		0.27 297097.20	0.19 1324103.00
	Throw (R^2 , BIC)	0.33 258282.10	0.27 239502.30	0.36 280115.50	0.29 228720.60	0.30 251189.00		0.24 1161711.00
	Mean (R^2 , BIC)	0.38 243391.90	0.32 254955.10	0.35 248350.20	0.36 240317.00	0.31 243151.20	0.27 254489.10	0.22 1176063.00

Learning curves of the Trial Times for each technique suggest that performance is unlikely to increase much with more participants, because they all plateau at an R^2 value of 0.20 to 0.30 (Figure 87). This phenomenon suggests that the data is good quality and representative of the population. The model could improve with more complexity or more features, which would require more data to avoid overfitting.




Figure 87. Learning curves for: Top left: Astral, top middle: Chicken, top right: Grab, bottom left: Sliding, bottom middle: Teleport, bottom right: Throw. Models trained on five-technique datasets.

13.6.3 Trial Time: Ranking Techniques

Sliding was used as a calibration technique to predict Trial Times and rank techniques. Appendix A.12 presents the actual and predicted Trial Times by participant and technique. The MAE was 1.87 seconds, the RMSE was 2.63 seconds, and the MAPE was 11.3%.

Unlike the model in the questionnaire scenario, the models used to predict Trial Time in this scenario were able to capture subtle differences between participants and predict different techniques within the top three ranks (Figure 88). For example, the models were able to correctly predict that Astral, Sliding, and Chicken were the fastest techniques for a few participants with impairments. The models were also able to differentiate between groups within ranks. For example, for rank 6, models correctly predicted that Chicken was the slowest for a few participants

without impairments, while Grab and Throw were the slowest techniques for participants with impairments. So, although the models were less accurate for predicting raw Trial Time compared to the model from the questionnaire scenario, they were able to pick up on participant characteristics to rank techniques with greater accuracy compared to the questionnaire scenario.




Figure 88. Slope graph of actual and predicted techniques by rank from fastest to slowest (rank 1-6). The y-axis is ordered alphabetically.

Heat maps showed that the models' overall accuracy when ranking six techniques was moderately high (Figure 89). The prediction for the fastest technique was the most accurate for both groups compared to all other ranks.




Figure 89. Left: Heat map of overall ranking accuracy for all participants, middle: participants with impairments, right: and participants without impairments. The model accurately predicted the top three techniques for ranks 1-3 and 6 for both groups.

The models predicted the top technique for people with impairments with 80% accuracy, and that for people without impairments with 89% accuracy (Table 17). However, the MAPE was higher for the top 1 and 2 techniques compared to the other ranks (Table 17). The models were less accurate predicting the top fourth, fifth, and sixth techniques for both groups of participants, and more so for people without impairments compared to people with impairments (Table 17). However, they predicted the worst technique for both groups with greater accuracy than they did the fourth and fifth worst techniques.

Table 17. Table showing ranking accuracy, or in other words, the number of techniques that were predicted correctly for a rank out of the total techniques that were actually that rank based on calibration with Sliding. Mean absolute percent error (MAPE) is also reported. Accuracy and MAPE are reported for participants with impairments, participants without impairments, and all participants. The overall ranking accuracy and MAPE of the six techniques were calculated, as well as the accuracy for each rank individually.

	Calibration Scenario: Accuracy (MAPE)	Top 1 Technique Accuracy (MAPE)	Top 2 Technique Accuracy (MAPE)	Top 3 Technique Accuracy (MAPE)	Top 4 Technique Accuracy (MAPE)	Top 5 Technique Accuracy (MAPE)	Top 6 Technique Accuracy (MAPE)
With Impairments	55% (11%)	80% (15%)	55% (13%)	55% (8%)	45% (11%)	40% (10%)	55% (10%)
Without Impairments	54% (11%)	89% (17%)	68% (14%)	47% (11%)	37% (7%)	37% (9%)	47% (10%)
All	55% (11%)	85% (16%)	62% (13%)	51% (10%)	41% (9%)	38% (10%)	51% (10%)

13.6.4 Calibration Scenario: Summary

Selected features for predicting Trial Time included both feature types: questionnaire and calibration features. A comparison of calibration techniques revealed that Sliding was the best technique for predicting Trial Time for the other techniques. Ranking techniques based on Trial Time revealed that models were able to detect variability within and between participant groups. They were moderately accurate when ranking the six techniques for all participants, but particularly accurate when ranking the top technique for both groups.

13.7 Summary

I discuss insights I gained while the building predictive models based on the questionnaire and calibration scenarios. I then provide a list of recommendations for developing a future ability-based recommender system for locomotion techniques. Lastly, I identify some limitations of the investigation and describe future work inspired by this research.

13.7.1 Questionnaire Scenario: Selected Features

The QuickDASH features had a higher importance ranking than TRIQ features during feature selection for Trial Time, Hit Rate, and Movement Variability. QuickDASH is a standardized questionnaire and questions have been designed to measure different dimensions of upper-body impairments. TRIQ, on the other hand, has not been validated so there are correlated questions that likely do not capture various facets of a person's impairments. Additionally, QuickDASH encodes more information about impairment severity by using a Likert-Type scale, while TRIQ only encodes the presence of an impairment. TRIQ features are sparse because of their dichotomous nature, and many are likely redundant, which can explain why more QuickDash features were selected than TRIQ features.

Even though QuickDASH was had greater importance to the model than TRIQ features, a challenge with the QuickDASH features is that their meaning is not immediately evident. For example, it's not clear how difficulty doing chores relates to Trial Time.

Technique was the most important feature for Trial Time, indicating that much of the variance in performance could be explained by the locomotion technique rather than participants' impairment-related characteristics. A *post hoc* analysis using only technique as a feature revealed that technique only explained half the variance in Trial Time compared to what the selected feature set could explain. Yet, when ranking techniques by predicted Trial Times, the model predicted the same technique within each rank for the top three ranks. This result could indicate that impairment-related features were not given enough weight to capture variability in participants.

13.7.2 Calibration Scenario: Selected Features

The top features for predicting Trial Time were performance related. QuickDASH features were present in almost every model of Trial Time by technique, indicating that they still had predictive power even when calibration-data was introduced. Yet, the fact that the top features selected for each technique covered all the data types (e.g., questionnaire, movement-based, button-based, post-task, etc.) reveals that having different data types could have provided the model with a richer representation of an individual, leading to correct predictions of different techniques within ranks.

13.7.3 Outcome Measures

Hit Rate was not a good measure of performance for this analysis. There was a ceiling effect as a result of the time limit for a trial being too generous. It was too easy for participants to achieve high accuracy, which resulted insufficient variance to build a good predictive model. It was also a dichotomous variable and, as such, did not encode as much information about performance as a continuous variable like Trial Time. Therefore, even if the ceiling effect was eliminated with a

different experimental design, Hit Rate still might not have been very informative, particularly if Trial Time encoded hit or misses, which it did in this study design because of the time limit.

The variance in Movement Variability could be explained entirely by technique. This finding is also in line with the lack of significant between-group differences in Movement Variability in the metrics exploration study, with the exception of Astral (Section 12.3.10). Ultimately, Hit Rate and Movement Variability were poor measures for predicting performance.

13.7.4 Model Performance and Fit: Scenario 2

Sliding was slightly better than other calibration techniques for predicting Trial Times for the other five techniques. This finding is surprising because performance did not significantly vary between groups or even impairment levels in the group comparison study (Section 11.2.1). It is possible that Sliding could magnify small differences in how the participant performed and experienced the technique.

In the trials without obstacles, it was possible to remain completely still and hit the target because the user was re-oriented at the beginning of a trial so that the target was directly in front of the user. As a result, a small deviation in how a participant used Sliding would be magnified, making it easier for the model to learn from the calibration data.

13.7.5 Comparing Recommendations from the Questionnaire and Calibration Scenarios

The model from questionnaire scenario ($R^2=0.54$) performed better than that from the Sliding calibration scenario when predicting Trial Time ($R^2=0.36$). However, when ranking techniques from 1-6, from fastest to slowest, the calibration model was able to capture individual and group differences within each rank, while the questionnaire model predicted the same technique within each of the top three ranks. Looking closer at the predictions from the calibration scenario, the

predicted Trial Times became more accurate when they were aggregated at the technique level. This finding explains why the questionnaire scenario model had a higher R^2 but was less accurate in ranking techniques compared to the calibration scenario model.




Figure 90. Top: Heat maps showing accuracy of the predicted technique rank from the questionnaire scenario. Bottom: Heat maps showing accuracy of the predicted technique rank from the calibration scenario. Left: Accuracy for all participants, middle: participants with impairments, and right: participants without impairments.

The models from the calibration scenario achieved 55% accuracy for all participants across all ranks, while the questionnaire scenario model achieved 42% accuracy (Table 18). The MAPE was also lower for the calibration scenario compared to the questionnaire scenario: 11% versus 15% (Table 19). The Sliding calibration scenario models were particularly more accurate when predicting the top technique for participants with impairments, achieving an accuracy of 80% compared to 70% with the questionnaire model. The calibration scenario also had a lower MAPE

for the top technique (15%) compared to the questionnaire scenario (25%). However, predicting Teleport for every participant as the top technique already resulted in a high accuracy of 79% for all participants. The calibration model was slightly more accurate, with an accuracy of 85%.

Table 18. Overall ranking accuracy and within-rank accuracy for the questionnaire scenario (QS) and the calibration scenario (CS). Results are presented for all participants, participants with impairments, and participants without impairments.

	All Techniques Accuracy		Top 1 Technique Accuracy		Top 2 Technique Accuracy		Top 3 Technique Accuracy		Top 4 Technique Accuracy		Top 5 Technique Accuracy		Top 6 Technique Accuracy	
	QS	CS	QS	CS	QS	CS	QS	CS	QS	CS	QS	CS	QS	CS
With Impairments	49%	55%	70%	80%	70%	55%	70%	55%	20%	45%	40%	40%	25%	55%
Without Impairments	39%	54%	89%	89%	63%	68%	32%	47%	21%	37%	11%	37%	16%	47%
All	42%	55%	79%	85%	67%	62%	51%	51%	18%	41%	21%	38%	23%	51%

Table 19. Overall mean absolute percent error (MAPE) and within-rank MAPE for the questionnaire scenario (QS) and the calibration scenario (CS). Results are presented for all participants, participants with impairments, and participants without impairments.

	All Techniques MAPE		Top 1 Technique MAPE		Top 2 Technique MAPE		Top 3 Technique MAPE		Top 4 Technique MAPE		Top 5 Technique MAPE		Top 6 Technique MAPE	
	QS	CS	QS	CS	QS	CS	QS	CS	QS	CS	QS	CS	QS	CS
With Impairments	17%	11%	25%	15%	14%	13%	15%	8%	17%	11%	12%	10%	19%	10%
Without Impairments	13%	11%	17%	17%	13%	14%	10%	11%	13%	7%	9%	9%	16%	10%
All	15%	11%	21%	16%	13%	13%	12%	10%	15%	9%	11%	10%	17%	10%

The Sliding calibration scenario models were better at ranking the fourth, fifth, and sixth techniques compared to the questionnaire scenario model (Figure 90). There was much more variability in the actual rank of the last three techniques compared to the top three, and it appears that the calibration scenario models were able to capture more of it. This finding suggests that the calibration scenario models identified underlying patterns in participants' device usage and post-task questionnaire data because these types of data were present in the training data for the calibration scenario but not the questionnaire scenario.

13.7.6 Recommendations

I list recommendations for creating a recommender system for locomotion techniques based on the study.

Recommendation 1: Use a calibration technique plus a questionnaire.

The calibration scenario led to higher accuracy in technique ranking compared to the questionnaire scenario and feature selection revealed that calibration-related metrics were important to the model.

Recommendation 2: Use QuickDASH alone or a combination of QuickDASH and TRIQ for the questionnaire.

There were more QuickDASH features selected as top features than TRIQ features, however, two TRIQ features were still present in the top five features for predicting Trial Time in the calibration scenario: isolation and low strength.

Recommendation 3: Use Sliding as a calibration technique but the testbed must re-orient the user to face the target at the beginning of a trial.

Sliding resulted in the highest performing and best fit models on average for predicting Trial Time in the calibration scenario. Sliding is also reasonable as a calibration technique because it had a low perceived workload by participants with impairments. Therefore, the user would not have difficulty doing the calibration task. The predictive power of Sliding might be in part due to the fact that it did not require movement in trials without obstacles because the user was automatically re-oriented to the next target, which could have made the model more sensitive to superfluous movements. Therefore, it was likely a combination of the task design and the technique that made Sliding better as a calibration technique compared to other techniques.

Recommendation 4. Do not use Movement Variability as a performance measure.

An investigation into how much technique influenced the prediction for Trial Time in the questionnaire scenario revealed that the same amount of variance in Trial Time could be explained by technique alone compared to the selected feature set.

Recommendation 5: If using Hit Rate as a performance measure, ensure that the task is difficult to increase variability or just use Trial Time.

The problem I encountered with Hit Rate was that the generous time limit participants were given to hit a target resulted in a ceiling effect, because most participants were able to hit the target. However, making it more difficult to hit a target could discourage participants, particularly participants with impairments who do not use VR often, and they might not try as hard or be turned off of VR as a result. Trial Time actually encoded accuracy because of the time limit and could be seen as a continuous version of Hit Rate because it essentially measured how close a participant was to missing the target. Therefore, Trial Time could be used instead of Hit Rate if it is important that participants not be discouraged.

13.7.7 Limitations and Future Work

Future work could investigate improving the validity of the TRIQ questionnaire. Currently, there is no standardized questionnaire for measuring how motor impairments affect performance with technology. Researchers could redesign the TRIQ questionnaire to use Likert-type scales.

Learning curves for Trial Time from both scenarios indicated only a slight increase in performance with more participants. Future researchers could recruit more participants to avoid overfitting the data when increasing model complexity or adding more features.

Although this study only investigates six techniques, it is likely that the same modelling approach would generalize to other techniques as well. Features used in the questionnaire scenario were not technique-dependent and rather represented participants' upper-body impairment. Also, the top two features for the calibration scenario were Movement Variability and Trial Time, indicating that how participants performed with other techniques is a good predictor of how they will perform with the technique for which Trial Time is being predicted. QuickDASH features

were present in all top five feature sets which further suggests some of the models' predictive abilities come from impairment-related rather than technique-specific measures.

Future research could investigate whether Sliding is still a good calibration technique when used with a different task design where the user isn't automatically re-oriented to face the target. Furthermore, future work could investigate building a predictive model based on participants' performance with a calibration task that is more representative of a real VR app, for example, one in which they search for an object, avoid moving objects, or follow a path.

14 CONCLUSION

In this chapter I discuss overarching themes from the work in this dissertation and list the contributions of this research. I also discuss future work and reflect on my experience doing this research.

Across this work, I found that people with upper-body impairments can use VR interaction techniques developed for commercial VR hardware but that the interaction techniques need to be adapted to the user and virtual environment to be accessible. The hypothesis that different abilities could be better suited to different locomotion techniques proved to be accurate, as there was a range of optimal techniques for people with impairments. This finding supported the need for a predictive model to personalize the experience of VR.

However, a limitation of the predictive model proposed in this work is that it defines the “best” technique as the most accessible technique, which is the technique with which participants could navigate an environment in the least amount of time. In the technique comparison study, it was evident that the fastest or most comfortable techniques were not necessarily preferred because participants weighed other factors such as presence and entertainment, which is a finding supported by prior work (Yamagami et al., 2022). Despite this current limitation, the model was able to predict a user’s top technique with 85% accuracy when they used Sliding to calibrate the model. This finding suggests that a participant can identify their top technique out of a set of six techniques with relatively high accuracy by filling out a questionnaire and performing a simple navigation task. Yet, despite the relatively high accuracy the model achieved, the mean absolute percent error (MAPE) was also relatively high (16%) for the top technique predictions, indicating that there was a small number of data points for which the model made inaccurate predictions. This result is expected because there was a lot of variability in the data. More data, particularly

from people with impairments, could reduce variability, which could lead to more accurate predictions.

14.1 Selection of Existing Techniques for Accessibility

One of the main findings of this dissertation is that it is possible to select existing interaction techniques that can improve the accessibility of virtual reality (VR). The scene taxonomy enabled the identification of existing scene-viewing techniques in the research literature going back thirty years that required little to no head movement (Chapter 6), The Nearmi design framework demonstrated how an existing technique, the point-of-interest (POI) technique, can be discretized and its components adapted to the preferences and abilities of individual users (Chapter 7). Evaluating locomotion techniques revealed that Teleport was the most accessible for individuals with upper-body impairments. Additionally, their performance with Sliding was comparable to that of individuals without impairments (Chapters 10 and 11). The ranking of techniques using a predictive model enabled the best technique to be selected based on an individual's abilities with some accuracy (Chapter 13).

These findings indicate that the existing corpus of scene-viewing and locomotion techniques can be used and adapted to cater to an individual's type and severity of impairment. This result is promising because it means that novel techniques do not need to be designed to make VR accessible, which would take time to design and evaluate. It also means that people with impairments can use the same techniques as people without impairments, which can reduce stigma and enhance self-efficacy (Shinohara & Wobbrock, 2011; Wobbrock, Kane, et al., 2011) Finally, adapting existing techniques to be more accessible can enhance the usability of VR for everyone, as is what often happens when technologies are adapted for people with disabilities (Trewin, 2000).

14.2 Designing Novel Techniques Based on User and Virtual Environment Characteristics

This work also demonstrated that designing interaction techniques can improve the accessibility of VR. Furthermore, it reveals that a specific approach, which involves deconstructing existing designs into their component parts, can facilitate the design process and ground it in an understanding of virtual environments (VE), interaction techniques, and user characteristics.

In Chapter 6 I discussed how segmenting a VE into its visual and task components can facilitate the design of scene-viewing techniques that many participants thought were easier to use than the default technique. In Chapter 7 I illustrated how breaking down the POI technique into its functional components can make it easier to design new adaptations that address a range of design considerations including accessibility, usability, realism, etc. In Chapter 8 I showed how breaking down the xyY color space into different regions can make indistinguishable colors in VEs accessible to people with color vision impairments. In Chapter 12, I identified how different movements, parts of the body, and controller interactions influenced performance with locomotion techniques, which could suggest adaptations to locomotion technique design. For example, I found that people with impairments had significantly lower Angular Head Acceleration and Velocity when using Chicken and Sliding, which suggests that increasing the head rotation gain so that the virtual world rotates around the user faster with less head movement could make the techniques more accessible. The need to adapt head rotation gain is also supported by findings from a previous study in which participants with motor impairments struggled to articulate head gestures (Vatavu & Ungurean, 2022).

In the prediction study (Chapter 13), the TRIQ questionnaire was designed to isolate the components of upper-body impairment to reveal how specific characteristics influenced performance with locomotion techniques. However, this approach of discretizing user

characteristics was less successful compared to approaches discussed above. Fewer TRIQ features were selected for the predictive model compared to QuickDASH features, which assess the impact of upper-body impairments on daily tasks. This finding suggests a future research opportunity to refine and validate the TRIQ questionnaire, with the aim of clarifying the connection between impairment and locomotion technique performance.

14.3 Identifying Trade-offs in Design Considerations to Personalize Virtual Reality Interaction

A finding that is consistent throughout this work is that there is no optimal scene-viewing or locomotion technique. People with physical and situational impairments made trade-offs in design considerations when choosing their preferred techniques, which included accessibility, usability, realism, comfort, familiarity with the interaction, task objective, input device, aesthetic of the implementation, and whether the interaction was unique to VR. These design considerations impacted which scene-viewing participants preferred (Chapter 6), which POI technique components they preferred (Chapter 7), and which locomotion techniques they preferred (Chapter 10).

Although accessibility was an important factor, it was not always the most important factor. When choosing a preferred scene-viewing technique, some participants preferred a more realistic technique even if it was less accessible. Also, even though most participants thought the scene-viewing techniques designed with the taxonomy were easier to use than panning, they did not always prefer the new scene-viewing technique.

In Chapter 7, the 17 participants in the study created 13 unique implementations of Nearmi, meaning that nearly every participant chose a unique combination of Nearmi components. Their rationale for choosing components varied, but a consistent trade-off participants made was, again, between accessibility and realism, which was the most common trade-off across the studies. For

example, participants liked the tap Selection component because it mimicked tapping a touchscreen, however some participants were concerned they would not be able to steady their arm to make the selection. The fact that the trade-off between accessibility and realism was common makes intuitive sense because VR aims to replicate interactions typical of a person without impairments within the physical world. When the physical world is inaccessible, this issue can extend into VR as well. Overcoming physical challenges virtually can disrupt the sense of realism and presence because a real-world equivalent often does not exist. To enhance accessibility without compromising realism, virtual environments themselves can be designed to be more accessible. For example, virtual shelves can be made lower, eliminating the need for users to reach. Another option is to use vehicles for locomotion, such as a virtual wheelchair or car. These vehicles are used for locomotion in the real world, and therefore would be more realistic as locomotion techniques compared to the ones studied in this work.

When choosing their preferred locomotion technique, the technique participants perceived as most physically comfortable was not always consistent with the technique that they liked the most or would use frequently in VR—some participants chose different techniques for each category (Chapter 10). For example, even though most participants thought Throw was the least comfortable and liked it the least, most participants chose Throw as one of the top three techniques that they would use frequently in VR. Reasons for this choice were that it was a form of exercise and rehabilitation therapy, and it also had an entertainment value that other techniques did not have. In addition, participants did not always prefer or want to use the technique they performed the best with. For example, Chicken was the second-most liked, considered comfortable, and technique participants would use, despite being one of the worst in terms of performance and workload measures.

This finding suggests a difference between accessible design and using performance as a proxy for accessibility. Participants liked that Chicken did not require any controller input and could be controlled exclusively with the head, making it an accessible technique by design even though it was one of the least accessible in terms of performance measures. Future work can explore the relationship between these concepts as many studies use performance as a proxy for accessibility. There might be a need to distinguish accessible performance and accessible design.

The fact that participants weighed multiple design considerations when choosing their preferred scene-viewing or locomotion technique and the existence of hundreds of interaction techniques necessitates personalizing VR interaction. Instead of trying various techniques to find their preference, users could rank user experience measures (i.e., presence, comfort, etc.) by importance and calibrate the system to their abilities, the system could then suggest techniques that maximize each criterion. This approach, which is aligned with Ability-Based Design (Wobbrock et al., 2018; Wobbrock, Kane, et al., 2011), could ensure that techniques remain accessible for a range of disabilities without compromising the user experience of VR.

14.4 Predicting Performance for a New Technique

A significant limitation of the predictive modelling approach is that the model would not be able to predict performance for a new technique. If a designer was going to predict the trial time for a new technique, they would have to retrain the questionnaire and calibration scenario models with additional data. Without any examples of how people perform with a new technique to use as training data, it would not be possible to make a prediction with the modelling approach I used. However, it could be possible to estimate how a new user would perform with a new technique by drawing comparisons to the techniques examined in the prediction study. For example, if the new technique uses gross arm movements like in Throw and Grab, a designer could input the new user's

questionnaire and calibration data when they were using Sliding to predict their trial times for Throw or Grab to estimate their time for the new technique.

Adopting an alternative approach to creating a predictive model could be more effective in addressing the limitation of the current models only being able to predict the performance of techniques they have previously encountered. For example, breaking down interactions into their component parts could help create a model that generalizes its predictions to new techniques. Interaction techniques can often be decomposed into discrete steps or "interaction units." For example, selecting a virtual target using a laser pointer in VR involves two primary interaction units: aiming the controller and pressing the trigger button. Each unit demands specific physical capabilities, such as arm strength, stability, grip strength, and finger dexterity. Performance predictions for new techniques can be made by developing a comprehensive vocabulary of interaction units common to several interaction techniques and mapping these units to the physical abilities required to perform them effectively. If associations between a user's abilities, interaction units, and performance with techniques that are composed of interaction units can be estimated, a predictive model could be built by mixing and matching interaction units. Predictions could then be made by plugging in the values for user abilities to produce a performance estimate without requiring training data for a specific technique.

14.5 Discussion of the Locomotion Techniques Evaluated

The six locomotion techniques evaluated in the last four chapters of this dissertation were chosen to cover different aspects of abilities, including gross motor movement of the arms, fine motor control with the fingers, the ability to aim, and the ability to move the head and torso. These techniques were selected to ensure coverage of different parts of the body and combinations of movements. For instance, Sliding required head and finger movement, Throw, Grab, and Teleport

required arm movement, Chicken required head and torso movement, and Astral required only finger control.

The selection of these techniques was guided by questions concerning the presence and absence of specific abilities. For example, what if someone could not hold a controller? What if someone lacked fine motor skills? What if they could not move their head or torso? By considering these scenarios, it was imperative that individuals without certain abilities should still be able to use at least one of the locomotion techniques.

An interesting finding that surfaced when evaluating locomotion techniques was that even though participants with impairments performed as well as people without impairments when using Sliding, they perceived their performance to be worse when responding to the NASA-TLX. It is possible that participants were comparing their performance to an assumed “typical” performance, whereas the people without impairments were not comparing their performance to an imagined standard. Physical and Mental Demand as well as Effort and Frustration were similar across impairment levels for Sliding. So, even though people with impairments did not perceive the workload for Sliding to be greater, they still perceived their performance to be lower compared to people without impairments.

14.6 Discussion of Participant-Related Observations

Some themes related to participants with upper-body impairments became evident through the course of this work, which led to several design recommendations for new interaction techniques. Controller-based scene-viewing and locomotion techniques were generally more accessible than movement-based techniques. For example, when creating their preferred Nearmi implementation, most participants chose the controller Selection component, in which they could scroll through representations using a thumbstick. In the technique comparison study, the top three techniques

mainly relied upon controller button input: Sliding relied on a button press and hold, Astral required thumbstick movement, and Teleport required the user to select using the Trigger button. Therefore, if a designer wants to make a technique accessible for people with upper-body impairments they should employ the following:

Design Recommendation: Techniques should mainly be operated through controller button input.

Another observation about participants with upper-body impairments was that the majority had lower-body impairments and used a wheelchair or other type of mobility aid. Therefore, even though this work focused on enhancing the accessibility of VR for people with upper-body impairments, the techniques needed to be usable while seated to be accessible.

Design Recommendation: Techniques should be able to be performed while seated.

Several participants with impairments used both hands to operate one controller. They typically used one hand to stabilize the device, and the other to push buttons. Therefore, most participants with impairments were limited to single-handed techniques. This observation suggests that the physical hardware still requires people to find workarounds to operate it. One adaptation that could make controllers more accessible is to enable users to attach the controller to their hands or wrists, so users are not required to grip and hold it. On the software side, not requiring two hands to use interaction techniques would improve their accessibility. Frameworks such as the one developed by Yamagami et al. (2022) can be used to adapt bimanual to unimanual techniques.

Design Recommendation: Do not require that a controller be gripped or held; instead, attach it to the wrist or hand.

Design Recommendation: Ensure interaction techniques can be used with one controller.

14.7 Contributions

Taken together, the work I have completed advances knowledge about selecting, designing, and evaluating accessible scene-viewing and locomotion techniques for VR. Below I list each contribution my work makes and group contributions by type (Wobbrock & Kientz, 2016).

14.7.1 Empirical Contributions

1. A demonstration of how a novel scene taxonomy can be employed to classify and generate scene-viewing techniques that require little to no head movement.
2. Quantitative and qualitative data demonstrating participants' perceptions of workload and user experience of three scene-viewing techniques while they experienced limited head mobility.
3. Challenges that people with mobility impairments experienced when turning their heads or bodies in VR and how these challenges impacted their ability to use VR.
4. Design considerations and trade-offs that were important to people with mobility impairments when choosing their preferred implementations of Nearmi.
5. Quantitative and qualitative data from people with upper-body impairments related to their performance, perceived workload, Presence, Simulator Sickness, as well as perceptions of accessibility and user experience when using six locomotion techniques in VR.
6. A comparison of data from Contribution 5 to data collected from people without upper-body impairments.

7. A comparison of movement-, button-, and target-related metrics computed from low-level device data while people with and without impairments used locomotion techniques.
8. The evaluation of a user model that predicts performance with locomotion techniques based on impairment-related questionnaire responses as well as low-level device data.

14.7.2 Artifact Contributions

1. Three scene-viewing techniques that require little to no head movement designed for scenes with uncommon visual properties and tasks.
2. Multiple implementations of Nearmi components that demonstrate a set of accessibility- and user experience-related design considerations.
3. A scene-viewing technique that make colors accessible to people with different types of color vision deficiencies by overlaying patterns on indistinguishable colors.
4. An extensible and customizable testbed that enables researchers to design and conduct multifactor VR studies with timed, locomotion, and pointing tasks.

14.7.3 Methodological Contributions

1. An approach to developing a model that predicts the performance of people with and without upper-body impairments with six locomotion techniques.

14.7.4 Theoretical Contributions

2. A taxonomy for identifying the visual properties of and tasks in a VR scene.
3. A design framework for creating accessible point-of-interest techniques.

14.8 Limitations and Future Work

An interesting finding was that, within the groups with impairments and without impairments, Teleport was the fastest technique for most people. However, Sliding showed no significant differences in Trial Time, Hit Rate, or Obstacles Hit between the groups. People without impairments performed significantly better with Teleport in terms of Trial Time than those with impairments. In the metrics exploration study, it was revealed that some participants with impairments overshot the target and had to correct their position, which could explain why performance with Teleport was not comparable between groups. Therefore, Teleport might be the best technique to implement if a designer is concerned with optimizing individual performance and Sliding might be the best if a designer is interested in ensuring users have similar performance regardless of impairment. While Sliding might be a universally accessible technique in the short term, further research is needed to determine if people with and without impairments can continue to perform similarly with extended use.

One limitation of this study is that only six locomotion techniques were investigated. There might be techniques even more accessible than Teleport, and further studies are needed to determine if the modeling approach remains effective with a broader range of techniques.

This design approach of discretizing VEs and interaction techniques into their component parts can also be useful in training generative AI systems to design interaction techniques or VEs based on individuals' abilities. Breaking down designs into interaction units and isolating their relationship with technology performance could help a model learn the relationship between specific components of the VE or interaction technique and user characteristics and performance directly, rather than through broader concepts such as impairments or interaction techniques. By using instances of VEs or interaction techniques that share components, a model can be trained on a large dataset, learning how different types of techniques and environments that share particular

components contribute to outcomes. The generative model could then use the user characteristics and preferences to mix and match design components to create a personalized interaction technique or VE.

There is a significant opportunity to develop a questionnaire that can predict how accessible VR or any technology will be based on an individual's impairments. Currently, there is no technology-related physical impairment questionnaire. As demonstrated in the prediction study, although QuickDASH was better at capturing the relationship between impairment and VR use than TRIQ, it was not clear how responses (i.e., ability to carry a briefcase) related to someone's ability to use VR. By contrast, the ability to isolate one's fingers, a TRIQ response, is directly related to pushing a button. Future work could use machine learning techniques to find relationships between multiple questionnaire items, low-level interaction data, and performance on technology-related tasks such as moving a mouse, using a touchscreen, or operating a VR controller. This approach could demonstrate the predictive power of a questionnaire in determining how accessible a technology will be for a new user.

There could be several reasons for the lack of predictive power of TRIQ compared to QuickDASH. One reason might be that participants responded to TRIQ questions without any context. It might have been easier to respond when asked about a specific task, for example, how easy or difficult it was for them to type or scroll on a touchscreen. This would enable participants to imagine their impairments relative to technology use, similar to how QuickDASH relates impairments to daily tasks. Another reason for the lack of predictive power is that TRIQ did not encode as much information as QuickDASH because it was dichotomous while QuickDash features were ordinal. Future research could investigate if Likert-type responses could improve TRIQ's predictive power.

In the technique comparison study, participants' preferred techniques were not always the most performant. Participants weighed factors such as presence and entertainment value when choosing their preferred technique. The models created in the prediction study predicted the fastest technique for users, which might not necessarily be their preferred technique. Future work can examine whether the preferred techniques can be predicted using questionnaire and calibration data. The NASA-TLX and user experience questions related to presence and simulator sickness could be given more weight in the new models because they could provide insight into how participants experienced the techniques, which could lead to more accurate predictions of preferred techniques.

A limitation of the current work is that the findings only apply to people whose upper-body impairments do not affect their ability to hold and operate at least one controller. Therefore, people with complete paralysis in their arms and hands would not be able to use the techniques discussed in this work. In the Nearmi study, there were two participants who did not have use of their arms or hands and discussed that they would use Nearmi if it could be controlled using voice commands or eye gaze techniques. Mott et al. (2020) also found that people with impairments would be inclined to use VR if it could be controlled with voice or eye-gaze input.

Another limitation of this work is that it prioritizes user preferences over designer preferences. Some techniques could lend themselves to providing an experience that the designer envisions, while others could detract from the experience.

However, one finding from this work that could lead to a designer choice is whether the designer wants users to perform well or perform similarly with a locomotion technique. In some VR experiences, such as games in which users are playing with other users, it could be important that users perform similarly, thus levelling the playing field. In this case, enabling a user to perform

their personal best with a locomotion technique might not be ideal because it would give some users an advantage that is not necessarily related to gameplay. To ensure users perform similarly, a designer could make Sliding the default locomotion technique. However, if users are not playing with others, or are in environments in which they might want to travel efficiently, Teleport could be a good default.

Based on this work, the optimal use case would be one in which a user can choose among a variety of techniques. However, this use case would limit how much control a designer has over orchestrating a particular experience. A compromise between designer and user preferences could be found by enabling the designer to choose a default technique that is accessible (e.g., Sliding, Teleport, or Astral) while also enabling a user to choose a different technique if they prefer. Also, the designer could allow the user to choose from a range of techniques but indicate the ones they recommend for the experience. With this approach, designers could still guide the user to have the experience they envisioned for the VR app, but ensure that the experience is accessible to users with a range of abilities and preferences.

14.9 Reflections and Insights

In this section I reflect on the process of doing this research.

14.9.1 Compensating for Diverse and Small Samples with Large Amounts of Data

Conducting research with older adults during my master's program and the first three years of my PhD was challenging due participants' diversity in terms of impairment types and severities, as well as attitudes toward and experience with technology. Additionally, finding participants, particularly institutionalized older adults, was difficult. I often used questionnaires to collect quantitative data or counted events during user studies, but the small and diverse sample sizes

frequently led to non-significant results, which was frustrating given the considerable time, energy, and travel involved in conducting these studies.

I learned that to compensate for these sample-related issues, it was essential to collect as much quantitative data as possible during each session. This realization led to the study design for the locomotion study. Recruiting and conducting the study took months, and the sessions were lengthy. However, collecting questionnaire, performance, interview, and low-level device data ultimately resulted in four chapters of this dissertation. I discovered that I could offset the limitations of small and diverse samples by gathering large amounts of data.

14.9.2 Cutting My Losses and Taking Risks

Three years into my PhD, after years of master's and PhD research and multiple publications on aging and technology that could not be used for my dissertation, I started on a completely new topic. Because of the COVID-19 pandemic and not knowing when I would be able to do research with older adults again, I started research on VR and accessibility. I chose to explore a topic that was just being investigated by the research community, with little prior research on VR accessibility for physical impairments to learn from. Although Professor Wobbrock is an expert in accessibility, VR, and the research problems in that space were not his primary areas of expertise. While he could advise me on accessibility, I had to rely on my firsthand experience with VR and the literature to guide my research. Despite these challenges, I chose to cut my losses with aging and technology research and pivot to VR and accessibility. Looking back, I am glad I did, I think my research improved and I was able to make different types of contributions as a result of having to learn a new research area quickly and under challenging circumstances.

14.10 Final Remarks

This dissertation presents my research on improving VR accessibility for people with physical impairments. This research is an attempt to demonstrate the following thesis:

The accessibility of virtual reality (VR) can be improved by the design, selection, and evaluation of scene-viewing and locomotion techniques using design frameworks, testbed evaluations, movement analysis, and predictive models.

The dissertation began with the development of a taxonomy and design framework aimed at assisting designers in selecting and designing scene-viewing techniques that require minimal head or torso movement. I then demonstrated how a scene-viewing technique could potentially enhance VR accessibility for individuals with color vision impairments. To advance the standardization of interaction technique evaluation in VR, I developed an extensible and customizable testbed. Using this testbed, I evaluated six locomotion techniques with participants with and without upper-body impairments. I then examined differences between groups through metrics computed from low-level VR device data. Additionally, I outlined a process for creating a user model that predicts an individual's optimal locomotion techniques.

Throughout this research, participants have expressed preferences for different design considerations such as accessibility, presence, comfort, and spatial orientation, among others, when selecting their preferred scene-viewing and locomotion techniques. The finding that the most accessible technique for individuals with impairments varies depending on their specific impairment underscores the necessity for personalized VR experiences. Personalization can make VR accessible to individuals with physical impairments without requiring specially designed interaction techniques or hardware. Furthermore, users would not be limited to the most accessible technique; they could define a list of accessibility and user experience criteria important to their VR experience, and a system could identify a technique that embodies these criteria. This approach

would allow individuals with impairments to enjoy the unique aspects of VR, such as the sense of physical presence in a virtual space, and the ability to engage in activities not possible in the real world (e.g., flying, inhabiting different bodies), while still being able to view and navigate the VE. Ultimately, this work and the vision of personalized VR paves the way for an inclusive future where this technology transcends accessibility barriers, enabling all individuals to experience the boundlessness of virtual reality.

15 REFERENCES

- [1] Al Zayer, M., MacNeilage, P., & Folmer, E. (2020). Virtual locomotion: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 26(6), 2315–2334. <https://doi.org/10.1109/TVCG.2018.2887379>
- [2] Anderson, T. W., & Darling, D. A. (1954). A test of goodness of fit. *Journal of the American Statistical Association*, 49(268), 765–769. <https://doi.org/10.2307/2281537>
- [3] Andujar, C., Vazquez, P., & Fairen, M. (2004). Way-Finder: Guided tours through complex walkthrough models. *Computer Graphics Forum*, 23(3), 499–508. <https://doi.org/10.1111/j.1467-8659.2004.00781.x>
- [4] Anthony, L., Kim, Y., & Findlater, L. (2013). Analyzing user-generated YouTube videos to understand touchscreen use by people with motor impairments. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, 1223–1232. <https://doi.org/10.1145/2470654.2466158>
- [5] Ardouin, J., Lécuyer, A., Marchal, M., Riant, C., & Marchand, E. (2012). FlyVIZ: A novel display device to provide humans with 360° vision by coupling catadioptric camera with HMD. *Proceedings of the 18th ACM Symposium on Virtual Reality Software and Technology (VRST '12)*, 41–44. <https://doi.org/10.1145/2407336.2407344>
- [6] Argelaguet, F., & Maignant, M. (2016). GiAnt: Stereoscopic-compliant multi-scale navigation in VEs. *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology (VRST '16)*, 269–277. <https://doi.org/10.1145/2993369.2993391>
- [7] Aspin, R., & Hoang Le, K. (2007). Augmenting the CAVE: An initial study into close focused, inward looking, exploration in IPT systems. *Symposium on Distributed Simulation and Real-Time Applications Augmenting*, 217–224. <https://doi.org/10.1109/DS-RT.2007.9>
- [8] Baker, S., Waycott, J., Robertson, E., Carrasco, R., Neves, B. B., Hampson, R., & Vetere, F. (2020). Evaluating the use of interactive virtual reality technology with older adults living in residential aged care. *Information Processing and Management*, 57(3), 1–13. <https://doi.org/10.1016/j.ipm.2019.102105>
- [9] Bares, W. H., & Lester, J. C. (1998). Intelligent multi-shot visualization interfaces for dynamic 3D worlds. *Proceedings of the 4th International Conference on Intelligent User Interfaces (IUI '98)*, 119–126. <https://doi.org/10.1145/291080.291101>
- [10] Bates, R., & Istance, H. O. (2002). Why Are Eye Mice Unpopular? — A Detailed Comparison of Head and Eye Controlled Assistive Technology Pointing Devices. In *Universal Access and Assistive Technology* (Vol. 2, Issue 3, pp. 75–86). Springer London. https://doi.org/10.1007/978-1-4471-3719-1_8
- [11] Baudisch, P., & Rosenbaltz, R. (2003). Halo: A technique for visualizing off-screen locations. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*, 481–488. <https://doi.org/10.1145/642611.642695>
- [12] Beckhaus, S., Blom, K., & Haringer, M. (2007). ChairIO—the Chair-Based Interface. In *Concepts and Technologies for Pervasive Games: A Reader for Pervasive Gaming Research* (Vol. 1, pp. 231–264).
- [13] Bee Wah, Y., Ibrahim, N., Abdul Hamid, H., Abdul-Rahman, S., & Fong, S. (2018). Feature selection methods: Case of filter and wrapper approaches for maximising classification accuracy. *Pertanika Journal of Science and Technology*, 26(1), 329–340. <http://www.pertanika.upm.edu.my/>
- [14] Bell, B., Feiner, S., & Höllerer, T. (2001). View management for virtual and augmented reality. *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology (UIST '01)*, 101–110. <https://doi.org/10.1145/502348.502363>
- [15] Bergström, J., Dalsgaard, T.-S., Alexander, J., & Hornbæk, K. (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 (CHI '21)*, 1–31. <https://doi.org/https://doi.org/10.1145/3411764.3445193>
- [16] Bhandari, J., MacNeilage, P., & Folmer, E. (2018). Teleportation without spatial disorientation using optical flow cues. *Proceedings of the 44th Graphics Interface Conference (GI '18)*, 162–167. <https://doi.org/10.20380/GI2018.22>
- [17] Bhandari, J., Tregillus, S., & Folmer, E. (2017). Legomotion: Scalable walking-based virtual locomotion. *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology (VRST '18)*. <https://doi.org/10.1145/3139131.3139133>

- [18] Boletsis, C., & Cedergrén, J. E. (2019). VR locomotion in the new era of virtual reality: An empirical comparison of prevalent techniques. *Advances in Human-Computer Interaction*, 2019, 1–15. <https://doi.org/10.1155/2019/7420781>
- [19] Bolte, B., Steinicke, F., & Bruder, G. (2011). The jumper metaphor: An effective navigation technique for immersive display setups. *Proceedings of Virtual Reality International Conference (VRIC '11)*, 6–8. <https://api.semanticscholar.org/CorpusID:51960695>
- [20] Boudoin, P., Otmane, S., & Mallem, M. (2008). Fly Over, a 3D interaction technique for navigation in virtual environments independent from tracking devices. *10th International Conference on Virtual Reality (VRIC '08)*, 7–13. <https://hal.science/hal-00339449>
- [21] Bowman, D. A., & Hodges, L. F. (1997). An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments. *Proceedings of the 1997 Symposium on Interactive 3D Graphics (I3D '97)*, 35–38. <https://doi.org/10.1145/253284.253301>
- [22] Bowman, D. A., Johnson, D. B., & Hodges, L. F. (1999). Testbed evaluation of virtual environment interaction techniques. *Virtual Reality Software and Technology (VRST '99)*, 26–33. <https://doi.org/10.1145/323663.323667>
- [23] Bowman, D. A., Johnson, D. B., & Hodges, L. F. (2001). Testbed evaluation of virtual environment interaction techniques. *Presence: Teleoperators and Virtual Environments*, 10(1), 75–95. <https://doi.org/10.1162/105474601750182333>
- [24] Bowman, D. A., Koller, D., & Hodges, L. F. (1997). Travel in immersive virtual environments: An evaluation of viewpoint motion control techniques. *Proceedings of IEEE 1997 Annual International Symposium on Virtual Reality*, 45–52. <https://doi.org/10.1109/VRAIS.1997.583043>
- [25] Bowman, D., Koller, D., & Hodges, L. F. (1998). A methodology for the evaluation of travel techniques for immersive virtual environments. *Virtual Reality*, 3(2), 120–131. <https://doi.org/10.1007/BF01417673>
- [26] Bowman, D., Wingrave, C., Campbell, J., & Ly, V. (2001). Using pinch gloves for both natural and abstract interaction techniques in virtual environments (pp. 1–5). <https://eprints.cs.vt.edu/archive/00000547/>
- [27] Bozgeyikli, E., Raij, A., Katkooi, S., & Dubey, R. (2016). Point & Teleport locomotion technique for virtual reality. *Proceedings of the 2016 Annual Symposium on Computer-Human Interaction in Play (CHI PLAY '16)*, 205–216. <https://doi.org/10.1145/2967934.2968105>
- [28] Braun, V., & Clarke, V. (2008). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- [29] Brogan, D. C., Metoyer, R. A., & Hodgins, J. K. (1998). Dynamically simulated characters in virtual environments. *IEEE Computer Graphics and Applications*, 18(5), 58–69. <https://doi.org/10.1109/38.708561>
- [30] Brooks, F. P. (1999). What's real about virtual reality? *IEEE Computer Graphics and Applications*, 19(6), 16–27. <https://doi.org/10.1109/38.799723>
- [31] Brown, L. D., & Hua, H. (2006). Magic lenses for augmented virtual environments. *IEEE Computer Graphics and Applications*, 26(4), 64–73. <https://doi.org/10.1109/MCG.2006.84>
- [32] Brudy, F., Holz, C., Rädle, R., Wu, C. J., Houben, S., Klokmoose, C. N., & Marquardt, N. (2019). Cross-device taxonomy: Survey, opportunities, and challenges of interactions spanning across multiple devices. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, 1–28. <https://doi.org/10.1145/3290605.3300792>
- [33] Burtnyk, N., Khan, A., Fitzmaurice, G., Balakrishnan, R., & Kurtenbach, G. (2002). StyleCam: Interactive stylized 3D navigation using integrated spatial and temporal controls. *Proceedings of the 15th Annual ACM Symposium on User Interface Software and Technology (UIST '02)*, 101–110. <https://doi.org/10.1145/571985.572000>
- [34] Buswell, G. T. (1935). *How People Look at Pictures: A Study of the Psychology and Perception in Art*. University of Chicago Press.
- [35] Cabral, M. C., Morimoto, C. H., & Zuffo, M. K. (2005). On the usability of gesture interfaces in virtual reality environments. *Proceedings of the 2005 Latin American Conference on Human-Computer Interaction (CLIHIC '05)*, 100–108. <https://doi.org/10.1145/1111360.1111370>
- [36] Cabric, F., Dubois, E., & Serrano, M. (2021). A predictive performance model for immersive interactions in mixed reality. *2021 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 1–10.

- <https://doi.org/10.1109/ISMAR52148.2021.00035i>
- [37] Cannavò, A., Calandra, D., Praticò, F. G., Gatteschi, V., & Lamberti, F. (2021). An evaluation testbed for locomotion in virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 27(3), 1871–1889. <https://doi.org/10.1109/TVCG.2020.3032440>
 - [38] Cardoso, J. C. S. (2016). Comparison of gesture, gamepad, and gaze-based locomotion for VR worlds. *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology (VRST '16)*, 319–320. <https://doi.org/10.1145/2993369.2996327>
 - [39] Carpendale, M. S. T., Cowperthwaite, D. J., & Fracchia, F. D. (1996). Distortion viewing techniques for 3-dimensional data. *Proceedings IEEE Symposium on Information Visualization*, 46–54. <https://doi.org/10.1109/INFVIS.1996.559215>
 - [40] Carrozzino, M., Avveduto, G., Tecchia, F., Gurevich, P., & Cohen, B. (2014). Navigating immersive virtual environments through a foot controller. *Proceedings of the 20th ACM Symposium on Virtual Reality Software and Technology (VRST '14)*, 23–26. <https://doi.org/10.1145/2671015.2671121>
 - [41] Carson, C., Belongie, S., Greenspan, H., & Malik, J. (2002). Blobworld: Image segmentation using expectation-maximization and its application to image querying. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(8), 1026–1038. <https://doi.org/10.1109/TPAMI.2002.1023800>
 - [42] Chang, R. C., Hung, C. S., Jain, D., & Guo, A. (2023). SoundBlender: Manipulating sounds for accessible mixed-reality awareness. *Adjunct Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*, 1–4. <https://doi.org/10.1145/3586182.3615787>
 - [43] Chenechal, M. Le, Lacoche, J., Royan, J., Duval, T., Gouranton, V., & Arnaldi, B. (2016). When the giant meets the ant: An asymmetric approach for collaborative and concurrent object manipulation in a multi-scale environment. *2016 IEEE Third VR International Workshop on Collaborative Virtual Environments*, 18–22. <https://doi.org/10.1109/3DCVE.2016.7563562>
 - [44] Chittaro, L., Ranon, R., & Ieronutti, L. (2003). Guiding visitors of Web3D worlds through automatically generated tours. *International Conference on 3D Web Technology*, 27–38. <https://doi.org/10.1145/636593.636598>
 - [45] Christou, C. G., & Aristidou, P. (2017). Steering versus teleport locomotion for head-mounted displays. In L. De Paolis, P. Bourdot, & A. Mongelli (Eds.), *Augmented Reality, Virtual Reality, and Computer Graphics (AVR '17)* (Vol. 10325, pp. 431–446). Springer International Publishing.
 - [46] Christou, C., Tzanavari, A., Herakleous, K., & Poullis, C. (2016). Navigation in virtual reality: Comparison of gaze-directed and pointing motion control. *18th Mediterranean Electrotechnical Conference (MELECON '16)*, 1–6. <https://doi.org/10.1109/MELCON.2016.7495413>
 - [47] Cirio, G., Marchal, M., Regia-Corte, T., & Lécuyer, A. (2009). The magic barrier tape: A novel metaphor for infinite navigation in virtual worlds with a restricted walking workspace. *Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology (VRST '09)*, 155–162. <https://doi.org/10.1145/1643928.1643965>
 - [48] Cirio, G., Vangorp, P., Chapoulie, E., Marchal, M., Lecuyer, A., & Drettakis, G. (2012). Walking in a Cube: Novel metaphors for safely navigating large virtual environments in restricted real workspaces. *IEEE Transactions on Visualization and Computer Graphics*, 18(4), 546–554. <https://doi.org/10.1109/TVCG.2012.60>
 - [49] Collins, J., Jung, C., Jang, Y., Montour, D., Won, A. S., & Azenkot, S. (2023). “The guide has your back”: Exploring how sighted guides can enhance accessibility in social virtual reality for blind and low vision people. *Proceedings of the 25th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '23)*, 1–14. <https://doi.org/10.1145/3597638.3608386>
 - [50] Color Blindness. (2023, March 17). Cleveland Clinic. <https://my.clevelandclinic.org/health/diseases/11604-color-blindness>
 - [51] Coomer, N., Clinton, W., Bullard, S., & Williams-Sanders, B. (2018). Evaluating the effects of four VR locomotion methods: Joystick, arm-cycling, point-tugging, and teleporting. *Proceedings of the 15th ACM Symposium on Applied Perception (SAP '18)*, 1–8. <https://doi.org/10.1145/3225153.3225175>
 - [52] Cui, J., Rosen, P., Popescu, V., & Hoffmann, C. (2010). A curved ray camera for handling occlusions through continuous multiperspective visualization. *IEEE Transactions on Visualization and Computer Graphics*,

- 16(6), 1235–1242. <https://doi.org/10.1109/TVCG.2010.127>
- [53] Davis Jr., F. D. (1986). A Technology Acceptance Model for Empirically Testing New End-User Information Systems: Theory and Results [Dissertation, Massachusetts Institute of Technology]. <http://hdl.handle.net/1721.1/15192>
 - [54] Dawe, M. (2004). Complexity, cost and customization: Uncovering barriers to adoption of assistive technology. Refereed Poster at the ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '04).
 - [55] de Haan, G., Griffith, E. J., & Post, F. H. (2008). Using the Wii Balance Board™ as a low-cost VR interaction device. Proceedings of the 2008 ACM Symposium on Virtual Reality Software and Technology (VRST '08), 289–290. <https://doi.org/10.1145/1450579.1450657>
 - [56] Desroche, S., Jolivet, V., & Plemenos, D. (2007). Towards a plan-based automatic exploration of virtual worlds. International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision, 25–32.
 - [57] Di Luca, M., Seifi, H., Egan, S., & Gonzalez-Franco, M. (2021). Locomotion Vault: The extra mile in analyzing VR locomotion techniques. Proceedings of the Conference on Human Factors in Computing Systems (CHI '21), 1–10. <https://doi.org/10.1145/3411764.3445319>
 - [58] Drucker, S. (1994). Intelligent Camera Control for Graphical Environments [Dissertation]. Massachusetts Institute of Technology.
 - [59] Elkin, L. A., Kay, M., Higgins, J. J., & Wobbrock, J. O. (2021). An aligned rank transform procedure for multifactor contrast tests. Proceedings of the 34th Annual ACM Symposium on User Interface Software and Technology (UIST '21), 754–768. <https://doi.org/10.1145/3472749.3474784>
 - [60] Elmqvist, N., & Tsigas, P. (2007). View-projection animation for 3D occlusion management. Computers & Graphics, 31, 864–876. <https://doi.org/10.1016/j.cag.2007.09.006>
 - [61] Elmqvist, N., & Tsigas, P. (2008). A taxonomy of 3D occlusion management for visualization. IEEE Transactions on Visualization and Computer Graphics, 14(5), 1095–1109. <https://doi.org/10.1109/TVCG.2008.59>
 - [62] Elmqvist, N., & Tudoreanu, M. E. (2007). Occlusion management in immersive and desktop 3D virtual environments: Theory and evaluation. The International Journal of Virtual Reality, 6(1), 1–13.
 - [63] Elmqvist, N., Tudoreanu, M. E., & Tsigas, P. (2007). Tour generation for exploration of 3D virtual environments. Proceedings of the 2007 ACM Symposium on Virtual Reality Software and Technology (VRST '07), 207–210. <https://doi.org/10.1145/1315184.1315224>
 - [64] Erazo, O., & Pino, J. A. (2015). Predicting task execution time on natural user interfaces based on touchless hand gestures. Proceedings of the 20th International Conference on Intelligent User Interfaces (IUI '15), 97–109. <https://doi.org/10.1145/2678025.2701394>
 - [65] Fairchild, K. M., Lee, B. H., Loo, J., Ng, H., & Serra, L. (1993). The heaven and earth virtual reality: Designing applications for novice users. Proceedings of IEEE Virtual Reality Annual International Symposium, 47–53. <https://doi.org/10.1109/VRAIS.1993.380799>
 - [66] Feiner, S. K., & Seligmann, D. D. (1992). Cutaways and ghosting: satisfying visibility constraints in dynamic 3D illustrations. The Visual Computer, 8(5–6), 292–302. <https://doi.org/10.1007/BF01897116>
 - [67] Findlater, L., Jansen, A., Shinohara, K., Dixon, M., Kamb, P., Rakita, J., & Wobbrock, J. O. (2010). Enhanced area cursors: Reducing fine pointing demands for people with motor impairments. Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology (UIST '10), 153–162.
 - [68] Flatla, D. R. (2012). Accessibility for individuals with color vision deficiency. Proceedings of the 24th Annual ACM Symposium Adjunct on User Interface Software and Technology (UIST '12), 102, 8–12. <https://doi.org/10.1145/2140446.2140448>
 - [69] Fletcher-Watson, S., Findlay, J. M., Leekam, S. R., & Benson, V. (2008). Rapid detection of person information in a naturalistic scene. Perception, 37, 571–583. <https://doi.org/10.1068/p5705>
 - [70] Fleury, C., Chauffaut, A., Duval, T., Gouranton, V., & Arnaldi, B. (2010). A generic model for embedding users' physical workspaces into multi-scale collaborative virtual environments. Proceedings of the 20th International Conference on Artificial Reality and Telexistence (ICAT '10), 1–9. <https://hal.science/inria-00534096v1>

- [71] Flowers, K. A. (1976). Visual “closed-loop” and “open-loop” characteristics of voluntary movement in patients with Parkinsonism and intention tremor. *Brain*, 99(2), 269–310. <https://doi.org/10.1093/brain/99.2.269>
- [72] Franz, R. L., Junuzovic, S., & Mott, M. (2021). Nearmi: A framework for designing point of interest techniques for VR users with limited mobility. *Proceedings of the 23rd International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '21)*, 1–14. <https://doi.org/10.1145/3441852.3471230>
- [73] Franz, R. L., Junuzovic, S., & Mott, M. (2024). A virtual reality scene taxonomy: Identifying and designing accessible scene-viewing techniques. *ACM Transactions on Computer-Human Interaction*, 31(2), 1–44. <https://doi.org/10.1145/3635142>
- [74] Frederick, B. N. (1999). Fixed-, random-, and mixed-effects ANOVA models: A user-friendly guide for increasing the generalizability of ANOVA results. In B. Thompson (Ed.), *Advances in Social Science Methodology* (Vol. 111). JAI Press.
- [75] Fukatsu, S. (1998). Intuitive control of “bird’s eye” overview for navigation in an enormous virtual environment. *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST '98)*, 67–76. <https://doi.org/10.1145/293701.293710>
- [76] Funk, M., Müller, F., Fendrich, M., Shene, M., Kolvenbach, M., Dobbartin, N., Günther, S., & Mühlhäuser, M. (2019). Assessing the accuracy of point & teleport locomotion with orientation indication for virtual reality using curved trajectories. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, 1–12. <https://doi.org/10.1145/3290605.3300377>
- [77] Gajos, K. Z., Weld, D. S., & Wobbrock, J. O. (2010). Automatically generating personalized user interfaces with Supple. *Artificial Intelligence*, 174(12–13), 910–950. <https://doi.org/10.1016/j.artint.2010.05.005>
- [78] Ganapathi, P., Thiel, F. J., Swapp, D., & Steed, A. (2022). Head in the clouds - Floating locomotion in virtual reality. *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW '22)*, 668–669. <https://doi.org/10.1109/VRW55335.2022.00185>
- [79] Geisler, W. S. (2008). Visual perception and the statistical properties of natural scenes. *Annual Review of Psychology*, 59, 167–192. <https://doi.org/10.1146/annurev.psych.58.110405.085632>
- [80] Gerling, K., Mason, L., & Dickinson, P. (2020). Virtual reality games for people using wheelchairs. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*, 1–11. <https://doi.org/10.1145/3313831.3376265>
- [81] Gerling, K., & Spiel, K. (2021). A critical examination of virtual reality technology in the context of the minority body. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*, 1–14. <https://doi.org/10.1145/3411764.3445196>
- [82] Gibson, J. J. (1979). The Theory of Affordances. In *The Ecological Approach to Visual Perception* (pp. 127–137). Lawrence Erlbaum Associates, Inc., Publishers.
- [83] Gilmour, A. R., Anderson, R. D., & Rae, A. L. (1985). The analysis of binomial data by a generalized linear mixed model. *Biometrika*, 72(3), 593–599. <https://doi.org/10.1093/biomet/72.3.593>
- [84] Glaw, X., Inder, K., Kable, A., & Hazelton, M. (2017). Visual methodologies in qualitative research: Autophotography and photo elicitation applied to mental health research. *International Journal of Qualitative Methods*, 16(1), 1–8. <https://doi.org/10.1177/1609406917748215>
- [85] Gleicher, M., & Witkin, A. (1992a). Through-the-lens camera control. *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '92)*, 331–340. <https://doi.org/10.1145/133994.134088>
- [86] Gleicher, M., & Witkin, A. (1992b). Through-the-lens camera control. *ACM SIGGRAPH Computer Graphics*, 26(2), 331–340.
- [87] Goette, T. (1998). Factors leading to the successful use of voice recognition technology. *Proceedings of the Third International ACM Conference on Assistive Technologies (ASSETS '98)*, 189–196. <https://doi.org/10.1145/274497.274532>
- [88] Gonzalez Penuela, R. E., Poremba, W., Trice, C., & Azenkot, S. (2022). Hands-On: Using gestures to control descriptions of a virtual environment for people with visual impairments. *Adjunct Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*, 1–4. <https://doi.org/10.1145/3526114.3558669>

- [89] Grill-Spector, K., Knouf, N., & Kanwisher, N. (2004). The fusiform face area subserves face perception, not generic within-category identification. *Nature Neuroscience*, 7(5), 555–562. <https://doi.org/10.1038/nn1224>
- [90] Grossman, T., Balakrishnan, R., Kurtenbach, G., Fitzmaurice, G., Khan, A., & Buxton, B. (2002). Creating principal 3D curves with digital tape drawing. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '02)*, 121–128. <https://doi.org/10.1145/503376.503398>
- [91] Gruenefeld, U., El Ali, A., Boll, S., & Heuten, W. (2018). Beyond Halo and Wedge: Visualizing out-of-view objects on head-mounted virtual and augmented reality devices. *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '18)*, 1–11. <https://doi.org/10.1145/3229434.3229438>
- [92] Gruenefeld, U., El Ali, A., Heuten, W., & Boll, S. (2017). Visualizing out-of-view objects in head-mounted augmented reality. *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '17)*, 1–7. <https://doi.org/10.1145/3098279.3122124>
- [93] Gruenefeld, U., Ennenga, D., Ali, A. El, Heuten, W., & Boll, S. (2017). EyeSee360: Designing a visualization technique for out-of-view objects in head-mounted augmented reality. *Proceedings of the 5th Symposium on Spatial User Interaction (SUI '17)*, 109–118. <https://doi.org/10.1145/3131277.3132175>
- [94] Gruenefeld, U., Stratmann, T. C., El Ali, A., Boll, S., & Heuten, W. (2018). RadialLight: Exploring radial peripheral LEDs for directional cues in head-mounted displays. *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '18)*, 1–6. <https://doi.org/10.1145/3229434.3229437>
- [95] Gruenefeld, U., von Barga, R., & Heuten, W. (2018). Identification of out-of-view objects in virtual reality. *Proceedings of the 2018 ACM Symposium on Spatial User Interaction (SUI '18)*, 182. <https://doi.org/10.1145/3267782.3274678>
- [96] Gummesson, C., Ward, M. M., & Atroshi, I. (2006). The shortened disabilities of the arm, shoulder and hand questionnaire (QuickDASH): Validity and reliability based on responses within the full-length DASH. *BMC Musculoskeletal Disorders*, 7(1), 44. <https://doi.org/10.1186/1471-2474-7-44>
- [97] Guo, R., & Quarles, J. (2012). Differences in presence between healthy users and users with multiple sclerosis. *IEEE VR Workshop on Perceptual Illusions in Virtual Environments*, 1–6. <https://doi.org/10.1109/PIVE.2012.6229792>
- [98] Gupta, A., Satkin, S., Efros, A. A., & Hebert, M. (2011). From 3D scene geometry to human workspace. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '11)*, 1961–1968. <https://doi.org/10.1109/CVPR.2011.5995448>
- [99] Gustafson, S., Baudisch, P., Gutwin, C., & Irani, P. (2008). Wedge: Clutter-free visualization of off-screen locations. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*, 787–796. <https://doi.org/10.1145/1357054.1357179>
- [100] Guy, E., Punpongson, P., Iwai, D., Sato, K., & Boubekeur, T. (2015). LazyNav: 3D ground navigation with non-critical body parts. *2015 IEEE Symposium on 3D User Interfaces (3DUI '15)*, 43–50. <https://doi.org/10.1109/3DUI.2015.7131725>
- [101] Hall, E. T. (1966). *The Hidden Dimension* (Vol. 609). Garden City, NY: Doubleday.
- [102] Harrell, F. E. (2015). *Regression Modeling Strategies*. Springer.
- [103] Hart, S. G. (2006). Nasa-Task Load Index (NASA-TLX); 20 years later. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 50(9), 904–908. <https://doi.org/10.1177/154193120605000909>
- [104] Hart, S. G., & Staveland, L. E. (1988). Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. *Advances in Psychology*, 52, 139–183. [https://doi.org/10.1016/S0166-4115\(08\)62386-9](https://doi.org/10.1016/S0166-4115(08)62386-9)
- [105] Haydar, M., Maidi, M., Roussel, D., & Mallem, M. (2009). A new navigation method for 3D virtual environment exploration. *2nd Mediterranean Conference on Intelligent Systems and Automation (CISA '09)*, 1107, 190–195. <https://doi.org/10.1063/1.3106472>
- [106] Hedeker, D., & Gibbons, R. D. (1994). A random-effects ordinal regression model for multilevel analysis. *Biometrics*, 50(4), 933–944. <https://doi.org/10.2307/2533433>
- [107] Hess, R. H., Baker, C. L. Jr., & Zihl, J. (1989). The “motion-blind” patient: low-level spatial and temporal filters. *The Journal of Neuroscience*, 9(5), 1628–1640. <https://doi.org/10.1523/JNEUROSCI.09-05->

01628.1989

- [108] Higgins, J. J., & Tashtoush, S. (1994). An aligned rank transform test for interaction. *Nonlinear World*, 1(2), 201–211.
- [109] Hilbe, J. M. (2011). *Negative Binomial Regression* (2nd ed., Vol. 2011). Cambridge University Press.
- [110] Hirzle, T., Gugenheimer, J., Geiselhart, F., Bulling, A., & Rukzio, E. (2019). A design space for gaze interaction on head-mounted displays. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, 1–12. <https://doi.org/10.1145/3290605.3300855>
- [111] Ho-Ching, F. W.-L., Mankoff, J., & Landay, J. A. (2003). Can you see what I hear? The design and evaluation of a peripheral sound display for the deaf. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*, 161–168. <https://doi.org/10.1145/642611.642641>
- [112] Hoiem, D., Efros, A. A., & Hebert, M. (2005). Geometric context from a single image. *Tenth IEEE International Conference on Computer Vision (ICCV '05)*, 1, 654–661. <https://doi.org/10.1109/ICCV.2005.107>
- [113] Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2), 65–70.
- [114] Hossain, Z., Hasan, K., Liang, H., & Irani, P. (2012). EdgeSplit: Facilitating the selection of off-screen objects. *Proceedings of the 14th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '12)*, 79–82. <https://doi.org/10.1145/2371574.2371588>
- [115] Hudak, P. L., Amadio, P. C., & Bombardier, C. (1996). Development of an upper extremity outcome measure: The DASH (disabilities of the arm, shoulder, and head). *American Journal of Industrial Medicine*, 29(6), 602–608. [https://doi.org/10.1002/\(SICI\)1097-0274\(199606\)29:6<602::AID-AJIM4>3.0.CO;2-L](https://doi.org/10.1002/(SICI)1097-0274(199606)29:6<602::AID-AJIM4>3.0.CO;2-L)
- [116] Hurst, A., Hudson, S. E., Mankoff, J., & Trewin, S. (2008). Automatically Detecting Pointing Performance. *Proceedings of the 13th International Conference on Intelligent User Interfaces (IUI '08)*, 3205, 11–19. <https://doi.org/10.1145/1378773.1378776>
- [117] Hvass, J. S., Larsen, O., Vendelbo, K. B., Nilsson, N. C., Nordahl, R., & Serafin, S. (2017). The effect of geometric realism on presence in a virtual reality game. *IEEE Virtual Reality*, 339–340. <https://doi.org/10.1109/VR.2017.7892315>
- [118] Hwang, F., Keates, S., Langdon, P., & Clarkson, J. (2004). Mouse movements of motion-impaired users: A submovement analysis. *Proceedings of the 6th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '04)*, 102–109.
- [119] Irani, P., Gutwin, C., & Yang, X. D. (2006). Improving selection of off-screen targets with Hopping. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*, 299–308. <https://doi.org/10.1145/1124772.1124818>
- [120] Ismail, A. S., Seifelnasr, M. M., & Guo, H. (2018). Understanding indoor scene: Spatial layout estimation, scene classification, and object detection. *Proceedings of the 3rd International Conference on Multimedia Systems and Signal Processing (ICMSSP '18)*, 64–70. <https://doi.org/10.1145/3220162.3220182>
- [121] Iwata, H., Yano, H., Fukushima, H., & Noma, H. (2005). CirculaFloor. *IEEE Computer Graphics and Applications*, 25(1), 64–67. <https://doi.org/10.1109/MCG.2005.5>
- [122] Iwata, H., Yano, H., & Tomioka, H. (2006). Powered Shoes. In *ACM SIGGRAPH Emerging Technologies* (pp. 28-es). <https://doi.org/10.1145/1179133.1179162>
- [123] Jain, D., Findlater, L., Volger, C., Zotkin, D., Duraiswami, R., & Froehlich, J. (2015). Head-mounted display visualizations to support sound awareness for the deaf and hard of hearing. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*, 241–250. <https://doi.org/10.1145/2702123.2702393>
- [124] Jain, D., Junuzovic, S., Ofek, E., Sinclair, M., Porter, J., Yoon, C., MacHanavajhala, S., & Ringel Morris, M. (2021). A taxonomy of sounds in virtual reality. *Proceedings of the 2021 ACM Designing Interactive Systems Conference (DIS '21)*, 160–170. <https://doi.org/10.1145/3461778.3462106>
- [125] Jain, D., Lin, A. C., Amalachandran, M., Zeng, A., Guttman, R., Findlater, L., & Froehlich, J. (2019). Exploring sound awareness in the home for people who are deaf or hard of hearing. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, 1–13. <https://doi.org/10.1145/3290605.3300324>

- [126] James, G., Witten, D., Hastie, T., & Tibshirani, R. (2021). *An Introduction to Statistical Learning*. Springer US. <https://doi.org/10.1007/978-1-0716-1418-1>
- [127] Jerald, J. (2015). *The VR Book: Human-Centered Design for Virtual Reality*. ACM and Morgan & Claypool. <https://doi.org/10.1145/2792790>
- [128] Jerald, J. (2018). A taxonomy of spatial interaction patterns and techniques. *IEEE Computer Graphics and Applications*, 38(1), 11–19. <https://doi.org/10.1109/MCG.2018.011461524>
- [129] Kaneko, Y., Chung, I., & Suzuki, K. (2013). Light-emitting device for supporting auditory awareness of hearing-impaired people during group conversations. *IEEE International Conference on Systems, Man, and Cybernetics*, 3567–3572. <https://doi.org/10.1109/SMC.2013.608>
- [130] Keates, S., Hwang, F., Langdon, P., Clarkson, P. J., & Robinson, P. (2002). Cursor measures for motion-impaired computer users. *Proceedings of the Fifth International ACM Conference on Assistive Technologies (ASSETS '02)*, 135–142. <https://doi.org/10.1145/638249.638274>
- [131] Kennedy, R. S., Lane, N. E., Berbaum, K. S., & Lilienthal, M. G. (1993). Simulator Sickness Questionnaire: An enhanced method for quantifying simulator sickness. *The International Journal of Aviation Psychology*, 3(3), 203–220. https://doi.org/10.1207/s15327108ijap0303_3
- [132] Kerr, D. A. (2010). The CIE XYZ and xyY color spaces. *Colorimetry*, 1(1), 1–16.
- [133] Kim, J. S., Gračanin, D., Matković, K., & Quek, F. (2008). Finger Walking in Place (FWIP): A traveling technique in virtual environments. In A. Butz, B. Fisher, A. Krüger, P. Olivier, & M. Christie (Eds.), *Smart Graphics: Vol. 5166 LNCS* (pp. 58–69). Springer. https://doi.org/10.1007/978-3-540-85412-8_6
- [134] Kim, J. S., Gračanin, D., Matković, K., & Quek, F. (2009). iPhone/iPod touch as input devices for navigation in immersive virtual environments. *IEEE Virtual Reality Conference*, 261–262. <https://doi.org/10.1109/VR.2009.4811045>
- [135] Kister, U. (2018). *Interactive Visualization Lenses: Natural Magic Lens Interaction for Graph Visualization* [Dissertation, Dresden University of Technology]. <http://nbn-resolving.de/urn:nbn:de:bsz:14-qucosa-236782>
- [136] Kister, U., Reipschläger, P., & Dachelt, R. (2014). Multi-touch manipulation of magic lenses for information visualization. *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces (ITS '14)*, 431–434. <https://doi.org/10.1145/2669485.2669528>
- [137] Kister, U., Reipschläger, P., Matulic, F., & Dachelt, R. (2015). BodyLenses: Embodied magic lenses and personal territories for wall displays. *Proceedings of the 2015 International Conference on Interactive Tabletops & Surfaces (ITS '15)*, 117–126. <https://doi.org/10.1145/2817721.2817726>
- [138] Kitson, A., Hashemian, A. M., Stepanova, E. R., Kruijff, E., & Riecke, B. E. (2017). Comparing leaning-based motion cueing interfaces for virtual reality locomotion. *IEEE Symposium on 3D User Interfaces (3DUI)*, 73–82. <https://doi.org/10.1109/3DUI.2017.7893320>
- [139] Kitson, A., Riecke, B. E., Hashemian, A. M., & Neustaedter, C. (2015). NaviChair: Evaluating an embodied interface using a pointing task to navigate virtual reality. *Proceedings of the 3rd ACM Symposium on Spatial User Interaction (SUI '15)*, 123–126. <https://doi.org/10.1145/2788940.2788956>
- [140] Kong, J., Zhong, M., Fogarty, J., & Wobbrock, J. O. (2022). Quantifying touch: New metrics for characterizing what happens during a touch. *Proceedings of the 24th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '22)*, 1–13. <https://doi.org/10.1145/3517428.3544804>
- [141] Kopper, R., Ni, T., Bowman, D. A., & Pinho, M. (2006). Design and evaluation of navigation techniques for multiscale virtual environments. *IEEE Virtual Reality Conference*, 175–182. <https://doi.org/10.1109/VR.2006.47>
- [142] Kupin, A., Moeller, B., Jiang, Y., Banerjee, N. K., & Banerjee, S. (2019). Task-driven biometric authentication of users in virtual reality (VR) environments. *MultiMedia Modeling (MMM '19)*, 55–67. https://doi.org/10.1007/978-3-030-05710-7_5
- [143] Lampton, D. R., Bliss, J. P., & Morris, C. S. (2014). Human performance measurement in virtual environments. *Handbook of Virtual Environments*, 779–810. <https://doi.org/10.1201/b17360-40>
- [144] Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33(1), 159–174. <https://doi.org/10.2307/2529310>
- [145] Langbehn, E., Bruder, G., & Steinicke, F. (2016). Scale matters! Analysis of dominant scale estimation in

- the presence of conflicting cues in multi-scale collaborative virtual environments. *IEEE Symposium on 3D User Interfaces (3DUI)*, 211–220. <https://doi.org/10.1109/3DUI.2016.7460054>
- [146] Lange, D., Stratman, T. C., Gruenfeld, U., & Boll, S. (2020). HiveFive: Immersion preserving attention guidance in virtual reality. *Proceedings of the 2020 SIGCHI Conference on Human Factors in Computing Systems (CHI '20)*, 1–13. <https://doi.org/10.1145/3313831.3376803>
 - [147] LaViola Jr., J. J., Feliz, D. A., Keefe, D. F., & Zeleznik, R. C. (2001). Hands-free multi-scale navigation in virtual environments. *Proceedings of the 2001 Symposium on Interactive 3D Graphics (I3D '01)*, 9–15. <https://doi.org/10.1145/364338.364339>
 - [148] Lazebnik, S., Schmid, C., & Ponce, J. (2006). Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '06)*, 2, 2169–2178. <https://doi.org/10.1109/CVPR.2006.68>
 - [149] Li, F.-F., Iyer, A., Koch, C., & Perona, P. (2007). What do we perceive in a glance of a real-world scene? *Journal of Vision*, 7(1), 10. <https://doi.org/10.1167/7.1.10>
 - [150] Li, F.-F., & Perona, P. (2005). A Bayesian hierarchical model for learning natural scene categories. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR '05)*, 524–531. <https://doi.org/10.1109/CVPR.2005.16>
 - [151] Li, Z., Connell, S., Dannels, W., & Peiris, R. (2022). SoundVizVR: Sound indicators for accessible sounds in virtual reality for deaf or hard-of-hearing users. *Proceedings of the 24th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '22)*, 1–13. <https://doi.org/10.1145/3517428.3544817>
 - [152] Li, Z., Shinohara, K., & Peiris, R. L. (2023). Exploring the use of the SoundVizVR plugin with game developers in the development of sound-accessible virtual reality games. *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, 1–7. <https://doi.org/10.1145/3544549.3585750>
 - [153] Lin, Y. T., Liao, Y. C., Teng, S. Y., Chung, Y. J., Chan, L., & Chen, B. Y. (2017). Outside-in: Visualizing out-of-sight regions-of-interest in a 360° video using spatial picture-in-picture previews. *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology (UIST '17)*, 255–265. <https://doi.org/10.1145/3126594.3126656>
 - [154] Littell, R. C., Henry, P. R., & Ammerman, C. B. (1998). Statistical analysis of repeated measures data using SAS procedures. *Journal of Animal Science*, 76(4), 1216–1231. <https://doi.org/10.2527/1998.7641216x>
 - [155] Littell, R. C., Milliken, G. A., Stroup, W. W., & Wolfinger, R. D. (1996). Analysis of repeated measures data. *SAS System for Mixed Models*, 87–134. <https://doi.org/10.4148/2475-7772.1427>
 - [156] Liu, S., Dissanayake, S., Patel, S., Dang, X., Mlsna, T., Chen, Y., & Wilkins, D. (2014). Learning accurate and interpretable models based on regularized random forests regression. *BMC Systems Biology*, 8(3), 1–9. <https://doi.org/10.1186/1752-0509-8-S3-S5>
 - [157] Looser, J., Billingham, M., & Cockburn, A. (2004). Through the looking glass: The use of lenses as an interface tool for augmented reality interfaces. *International Conference on Computer Graphics and Interactive Techniques in Australasia and Southeast Asia (GRAPHITE '04)*, 204–211. <https://doi.org/10.1145/988834.988870>
 - [158] Luo, W., Goebel, E., Reipschläger, P., Ellenberg, M. O., & Dachsel, R. (2021). Exploring and slicing volumetric medical data in augmented reality using a spatially-aware mobile device. *IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR '21)*, 334–339. <https://doi.org/10.1109/ISMAR-Adjunct54149.2021.00076>
 - [159] MacKenzie, I. S. (2013a). Designing HCI Experiments. In *Human-Computer Interaction: An Empirical Research Perspective* (pp. 157–189). Elsevier Inc.
 - [160] MacKenzie, I. S. (2013b). *Human-Computer Interaction: An Empirical Research Perspective*. Elsevier Inc. <https://doi.org/10.1016/C2012-0-02819-0>
 - [161] Mackenzie, I. S., Kauppinen, T., & Silfverberg, M. (2001). Accuracy measures for evaluating computer pointing devices. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '01)*, 9–16. <https://doi.org/10.1145/365024.365028>
 - [162] Mackinlay, J., Card, S. K., & Robertson, G. G. (1990). A semantic analysis of the design space of input devices. *Human-Computer Interaction*, 5(2–3), 145–190. <https://doi.org/10.1080/07370024.1990.9667153>

- [163] Mackinlay, J. D., Card, S. K., & Robertson, G. G. (1990). Rapid controlled movement through a virtual 3D workspace. *Proceedings of the 17th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '90)*, 171–176. <https://doi.org/10.1145/97879.97898>
- [164] Mahmud, M. R., Stewart, M., Cordova, A., & Quarles, J. (2022). Auditory feedback for standing balance improvement in virtual reality. *IEEE Conference on Virtual Reality and 3D User Interfaces*, 782–791. <https://doi.org/10.1109/VR51125.2022.00100>
- [165] Mathew, R., Mak, B., & Dannels, W. (2022). Access on demand: Real-time, multi-modal accessibility for the deaf and hard-of-hearing based on augmented reality. *Proceedings of the 24th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '22)*, 1–6. <https://doi.org/10.1145/3517428.3551352>
- [166] Matthews, T., Fong, J., Ho-Ching, F. W.-L., & Mankoff, J. (2006). Evaluating non-speech sound visualizations for the deaf. *Behaviour & Information Technology*, 25(4), 333–351. <https://doi.org/10.1080/01449290600636488>
- [167] Matthews, T., Fong, J., & Mankoff, J. (2005). Visualizing non-speech sounds for the deaf. *Proceedings of the 7th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '05)*, 52–59. <https://doi.org/10.1145/1090785.1090797>
- [168] Mayer, S., Reinhardt, J., Schweigert, R., Jelke, B., Schwind, V., Wolf, K., & Henze, N. (2020). Improving humans' ability to interpret deictic gestures in virtual reality. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*, 1–14. <https://doi.org/10.1145/3313831.3376340>
- [169] McCullough, M., Xu, H., Michelson, J., Jackoski, M., Pease, W., Cobb, W., Kalescky, W., Ladd, J., & Williams, B. (2015). Myo arm: Swinging to explore a VE. *Proceedings of the ACM SIGGRAPH Symposium on Applied Perception (SAP '15)*, 107–113. <https://doi.org/10.1145/2804408.2804416>
- [170] McGuffin, M. J., Tancau, L., & Balakrishnan, R. (2003). Using deformations for browsing volumetric data. *IEEE Visualization (VIS '03)*, 401–408. <https://doi.org/10.1109/VISUAL.2003.1250400>
- [171] McMahan, R. P., Kooper, R., & Bowman, D. A. (2014). Principles for designing effective 3D interaction techniques. In K. S. Hale & K. M. Stanney (Eds.), *Handbook of Virtual Environments: Design, Implementation, and Applications* (pp. 285–311). CRC Press.
- [172] Menzner, T., Gesslein, T., Otte, A., & Grubert, J. (2020). Above surface interaction for multiscale navigation in mobile virtual reality. *IEEE Conference on Virtual Reality and 3D User Interfaces*, 372–381. <https://doi.org/10.1109/VR46266.2020.00057>
- [173] Miller, R., Ajit, A., Kholgade Banerjee, N., & Banerjee, S. (2019). Realtime behavior-based continual authentication of users in virtual reality environments. *IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR '19)*, 253–254. <https://doi.org/10.1109/AIVR46125.2019.00058>
- [174] Miller, R., Banerjee, N. K., & Banerjee, S. (2020). Within-System and Cross-System Behavior-Based Biometric Authentication in Virtual Reality. *Proceedings - 2020 IEEE Conference on Virtual Reality and 3D User Interfaces, VRW 2020*, 311–316. <https://doi.org/10.1109/VRW50115.2020.00070>
- [175] Mine, M. R. (1995). *Virtual Environment Interaction Techniques* [Dissertation]. UNC Chapel Hill.
- [176] Montague, K., Hanson, V. L., & Cobley, A. (2012). Designing for individuals: Usable touch-screen interaction through shared user models. *Proceedings of the 14th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '12)*, 151–158. <https://doi.org/10.1145/2384916.2384943>
- [177] Moss, J. D., & Muth, E. R. (2011). Characteristics of head-mounted displays and their effects on simulator sickness. *Human Factors*, 53(3), 308–319. <https://doi.org/10.1177/0018720811405196>
- [178] Mott, M., Cutrell, E., Franco, M. G., Holz, C., Ofek, E., Stoakley, R., & Morris, M. R. (2019). Accessible by design: An opportunity for virtual reality. *IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR '19)*, 451–454. <https://doi.org/10.1109/ISMAR-Adjunct.2019.00122>
- [179] Mott, M., Tang, J., Kane, S. K., Cutrell, E., & Morris, M. R. (2020). “I just went into it assuming that I wouldn't be able to have the full experience”: Understanding the accessibility of virtual reality for people with limited mobility. *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '20)*, 1–13. <https://doi.org/10.1145/3373625.3416998>
- [180] Mott, M., Vatavu, R.-D., Kane, S. K., & Wobbrock, J. O. (2016). Smart Touch: Improving touch accuracy for people with motor impairments with template matching. *Proceedings of the 2016 CHI Conference on*

- Human Factors in Computing Systems (CHI '16), 1934–1946. <https://doi.org/10.1145/2858036.2858390>
- [181] Mott, M., & Wobbrock, J. O. (2019). Cluster Touch: Improving touch accuracy on smartphones for people with motor and situational impairments. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, 1–14. <https://doi.org/10.1145/3290605.3300257>
 - [182] Muhammad, A. S., Ahn, S. C., & Hwang, J.-I. (2016). Effect of using walk-in-place interface for panoramic video play in VR. *Proceedings of the 2016 Symposium on Spatial User Interaction (SUI '16)*, 203. <https://doi.org/10.1145/2983310.2989197>
 - [183] Muhammad, A. S., Ahn, S. C., & Hwang, J.-I. (2017). Active panoramic VR video play using low latency step detection on smartphone. *IEEE International Conference on Consumer Electronics (ICCE '17)*, 196–199. <https://doi.org/10.1109/ICCE.2017.7889284>
 - [184] Mühler, K., Neugebauer, M., Tietjen, C., & Preim, B. (2007). Viewpoint selection for intervention planning. *IEEE VGTC Symposium on Visualization*, 1–8. <https://doi.org/10.2312/VisSym/EuroVis07/267-274>
 - [185] Ngoc, L. Le, & Kalawsky, R. S. (2013). Evaluating usability of amplified head rotations on base-to-final turn for flight simulation training devices. *IEEE Virtual Reality*, 51–54. <https://doi.org/10.1109/VR.2013.6549359>
 - [186] Nielsen, J. (1994a). Heuristic evaluation. In *Usability Inspection Methods* (pp. 25–64). John Wiley & Sons.
 - [187] Nielsen, J. (1994b). *Usability Engineering*. Morgan Kaufmann.
 - [188] Nielsen, J., Clemmensen, T., & Yssing, C. (2002). Getting access to what goes on in people's heads?-Reflections on the think-aloud technique.
 - [189] Nilsson, N. C., Serafin, S., Laursen, M. H., Pedersen, K. S., Sikstrom, E., & Nordahl, R. (2013). Tapping-In-Place: Increasing the naturalness of immersive Walking-In-Place locomotion through novel gestural input. *IEEE Symposium on 3D User Interfaces (3DUI '13)*, 31–38. <https://doi.org/10.1109/3DUI.2013.6550193>
 - [190] Nilsson, N. C., Serafin, S., & Nordahl, R. (2013). The perceived naturalness of virtual locomotion methods devoid of explicit leg movements. *Proceedings of Motion on Games (MIG '13)*, 155–164. <https://doi.org/10.1145/2522628.2522655>
 - [191] Nilsson, N. C., Serafin, S., & Nordahl, R. (2016). Walking in place through virtual worlds. In M. Kurosu (Ed.), *Proceedings, Part II, of the 18th International Conference on Human-Computer Interaction (Vol. 9732, pp. 37–48)*. Springer-Verlag. https://doi.org/10.1007/978-3-319-39516-6_4
 - [192] Norman, D. A. (1999). Affordance, conventions, and design. *Interactions*, 6(3), 38–43. <https://doi.org/10.1145/301153.301168>
 - [193] Norman, D. A. (2010). The research-practice gap: The need for translational developers. *Interactions*, 17(4), 9–12. <https://doi.org/10.1145/1806491.1806494>
 - [194] Norman, D. A. (2013). *The Design of Everyday Things: Revised and Expanded Edition*. Basic Books.
 - [195] Norman, D. A. (2018). Affordances and Design. *Jnd.Org*. https://jnd.org/affordances_and_design/
 - [196] Oh, U., & Findlater, L. (2013). The challenges and potential of end-user gesture customization. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, 1129–1138. <https://doi.org/10.1145/2470654.2466145>
 - [197] Oliva, A. (2005). Gist of the scene. In L. Itti, G. Rees, & J. K. Tsotsos (Eds.), *Neurobiology of Attention* (pp. 251–256). Academic Press. <https://doi.org/10.1016/B978-012375731-9/50045-8>
 - [198] Oliva, A., & Torralba, A. (2001). Modeling the shape of the scene: A holistic representation of the spatial envelope. *International Journal of Computer Vision*, 42(3), 145–175. <https://doi.org/10.1023/A:1011139631724>
 - [199] Oliva, A., & Torralba, A. (2002). Scene-centered description from spatial envelope properties. In H. H. Bülthoff, C. Wallraven, S.-W. Lee, & T. A. Poggio (Eds.), *Biologically Motivated Computer Vision. BMCV 2002. Lecture Notes in Computer Science (Vol. 2525, pp. 263–272)*. Springer. https://doi.org/10.1007/3-540-36181-2_26
 - [200] Oliva, A., & Torralba, A. (2006). Building the gist of a scene: The role of global image features in recognition. In S. Martinez-Conde, S. L. Macknik, L. M. Martinez, J.-M. Alonso, & P. U. Tse (Eds.), *Progress in Brain Research (Vol. 155, Issue Part B, pp. 23–36)*. Elsevier. [https://doi.org/10.1016/S0079-6123\(06\)55002-2](https://doi.org/10.1016/S0079-6123(06)55002-2)
 - [201] Osmers, N., & Prilla, M. (2020). Getting out of out of sight: Evaluation of AR mechanisms for awareness

- and orientation support in occluded multi-room settings. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*, 1–11. <https://doi.org/10.1145/3313831.3376742>
- [202] Oyekoya, O., Steptoe, W., & Steed, A. (2009). A saliency-based method of simulating visual attention in virtual scenes. *Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology (VRST '09)*, 199–206. <https://doi.org/10.1145/1643928.1643973>
 - [203] Papadopoulos, C., Sugarman, D., & Kaufman, A. (2012). NuNav3D: A touch-less, body-driven interface for 3D navigation. *IEEE Virtual Reality Workshops (VRW '12)*, 67–68. <https://doi.org/10.1109/VR.2012.6180885>
 - [204] Patterson, G., Xu, C., Su, H., & Hays, J. (2014). The SUN attribute database: Beyond categories for deeper scene understanding. *International Journal of Computer Vision*, 108(1–2), 59–81. <https://doi.org/10.1007/s11263-013-0695-z>
 - [205] Patton, M. Q. (2001). *Qualitative Evaluation and Research Methods* (3rd Edition). SAGE Publications, Inc.
 - [206] Pavel, A., Hartmann, B., & Agrawala, M. (2017). Shot orientation controls for interactive cinematography with 360° video. *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology (UIST '17)*, 289–297. <https://doi.org/10.1145/3126594.3126636>
 - [207] Petford, J., Carson, I., Nacenta, M. A., & Gutwin, C. (2019). A comparison of guiding techniques for out-of-view objects in full-coverage displays. *Proceedings of the 2019 SIGCHI Conference on Human Factors in Computing Systems (CHI '19)*, 1–13. <https://doi.org/10.1145/3290605.3300288>
 - [208] Pfeiffer, T., Schmidt, A., & Renner, P. (2016). Detecting movement patterns from inertial data of a mobile head-mounted-display for navigation via walking-in-place. *IEEE Virtual Reality*, 263–264. <https://doi.org/10.1109/VR.2016.7504754>
 - [209] Pfeuffer, K., Geiger, M. J., Prange, S., Mecke, L., Buschek, D., & Alt, F. (2019). Behavioural biometrics in VR. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, 1–12. <https://doi.org/10.1145/3290605.3300340>
 - [210] Pierce, J. S., & Pausch, R. (2004). Navigation with place representations and visible landmarks. *IEEE Virtual Reality*, 173–288. <https://doi.org/10.1109/VR.2004.1310071>
 - [211] Piumsomboon, T., Lee, G. A., Ens, B., Thomas, B. H., & Billinghurst, M. (2018). Superman vs. giant: A study on spatial perception for a multi-scale mixed reality flying telepresence interface. *IEEE Transactions on Visualization and Computer Graphics*, 24(11), 2974–2982. <https://doi.org/10.1109/TVCG.2018.2868594>
 - [212] Piumsomboon, T., Lee, G. A., Irlitti, A., Ens, B., Thomas, B. H., & Billinghurst, M. (2019). On the shoulder of the giant: A multi-scale mixed reality collaboration with 360 video sharing and tangible interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, 1–17. <https://doi.org/10.1145/3290605.3300458>
 - [213] Poupyrev, I., Weghorst, S., Billinghurst, M., & Ichikawa, T. (1997). A framework and testbed for studying manipulation techniques for immersive VR. *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST '97)*, 21–28. <https://doi.org/10.1145/261135.261141>
 - [214] Preim, B., Michel, R., Hartmann, K., & Strothotte, T. (1998). Figure captions in visual interfaces. *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI '98)*, 235–246. <https://doi.org/10.1145/948496.948528>
 - [215] Preim, B., Raab, A., & Strothotte, T. (1997). Coherent zooming of illustrations with 3D-graphics and text. *Proceedings of the Graphics Interface 1997 Conference*, 105–113. <https://graphicsinterface.org/proceedings/gi1997/gi1997-12/>
 - [216] Preim, B., Ritter, A., Strothotte, T., Pohle, T., Bartram, L., & Forsey, D. A. (1995). Consistency of rendered images and their textual labels. *Proceedings of CompuGraphics*, 201–210.
 - [217] Qorbani, H. S., Abdinejad, M., Arya, A., & Joslin, C. (2022). Improving accessibility of elevation control in an immersive virtual environment. *IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR '22)*, 26–35. <https://doi.org/10.1109/AIVR56993.2022.00012>
 - [218] Ragan, E. D., Scerbo, S., Bacim, F., & Bowman, D. A. (2017). Amplified head rotation in virtual reality and the effects on 3D search, training transfer, and spatial orientation. *IEEE Transactions on Visualization and Computer Graphics*, 23(8), 1880–1895. <https://doi.org/10.1109/TVCG.2016.2601607>
 - [219] Ragan, E. D., Wood, A., McMahan, R. P., & Bowman, D. A. (2012). Trade-offs related to travel techniques

- and level of display fidelity in virtual data-analysis environments. Joint Virtual Reality Conference of ICAT - EGVE - EuroVR, 1–4. <https://doi.org/10.2312/EGVE/JVRC12/081-084>
- [220] Rakkolainen, I., Raisamo, R., Turk, M., & Höllerer, T. (2017). Field-of-view extension for VR viewers. Proceedings of the 21st International Academic Mindtrek Conference (AcademicMindtrek '17), 1–4. <https://doi.org/10.1145/3131085.3131088>
 - [221] Rietzler, M., Deubzer, M., Dreja, T., & Rukzio, E. (2020). Telewalk: Towards free and endless walking in room-scale virtual reality. Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20), 1–9. <https://doi.org/10.1145/3313831.3376821>
 - [222] Robinett, W., & Holloway, R. (1992). Implementation of flying, scaling, and grabbing in virtual worlds. Proceedings of the 1992 Symposium on Interactive 3D Graphics (I3D '92), 189–192. <https://doi.org/10.1145/147156.147201>
 - [223] Ropinski, T., Steinicke, F., & Hinrichs, K. (2005). A constrained road-based VR navigation technique for travelling in 3D city models. Proceedings of the 2005 International Conference on Augmented Tele-Existence (ICAT '05), 228–235. <https://doi.org/10.1145/1152399.1152441>
 - [224] Ruddle, R. A., & Lessels, S. (2009). The benefits of using a walking interface to navigate virtual environments. ACM Transactions on Computer-Human Interaction, 16(1), 1–19. <https://doi.org/10.1145/1502800.1502805>
 - [225] Rummukainen, O., & Mendonça, C. (2016). Reproducing reality: Multimodal contributions in natural scene discrimination. ACM Transactions on Applied Perception, 14(1), 1–10. <https://doi.org/10.1145/2915917>
 - [226] Rummukainen, O., Radun, J., Virtanen, T., & Pulkki, V. (2014). Categorization of natural dynamic audiovisual scenes. PLoS ONE, 9(5), 1–14. <https://doi.org/10.1371/journal.pone.0095848>
 - [227] Sajadi, B., Majumder, A., Oliveira, M. M., Schneider, R. G., & Raskar, R. (2013). Using patterns to encode color information for dichromats. IEEE Transactions on Visualization and Computer Graphics, 19(1), 118–129. <https://doi.org/10.1109/TVCG.2012.93>
 - [228] Sarcar, S., Jokinen, J. P. P., Oulasvirta, A., Wang, Z., Silpasuwanchai, C., & Ren, X. (2018). Ability-based optimization of touchscreen interactions. IEEE Pervasive Computing, 17(1), 15–26. <https://doi.org/10.1109/MPRV.2018.011591058>
 - [229] Sargunam, S. P., Moghadam, K. R., Suhail, M., & Ragan, E. D. (2017). Guided head rotation and amplified head rotation: Evaluating semi-natural travel and viewing techniques in virtual reality. IEEE Virtual Reality, 19–28. <https://doi.org/10.1109/VR.2017.7892227>
 - [230] Sargunam, S. P., & Ragan, E. D. (2018). Evaluating joystick control for view rotation in virtual reality with continuous turning, discrete turning, and field-of-view reduction. Proceedings of the 3rd International Workshop on Interactive and Spatial Computing (IWISC '18), 74–79. <https://doi.org/10.1145/3191801.3191815>
 - [231] Sarupuri, B., Chipana, M. L., & Lindeman, R. W. (2017). Trigger Walking: A low-fatigue travel technique for immersive virtual reality. IEEE Symposium on 3D User Interfaces (3DUI '17), 227–228. <https://doi.org/10.1109/3DUI.2017.7893354>
 - [232] Sassi, V. F. P., Porcino, T., Clua, E. W. G., & Trevisan, D. G. (2023). Redefining redirected movement for wheelchair-based interaction for virtual reality. IEEE 11th International Conference on Serious Games and Applications for Health (SeGAH '23), 1–8. <https://doi.org/10.1109/SeGAH57547.2023.10253796>
 - [233] Satkin, S., & Hebert, M. (2013). 3DNN: Viewpoint invariant 3D geometry matching for scene understanding. IEEE International Conference on Computer Vision, 1873–1880. <https://doi.org/10.1109/ICCV.2013.235>
 - [234] Satkin, S., Lin, J., & Hebert, M. (2012). Data-driven scene understanding from 3D models. Proceedings of British Machine Vision Conference (BMVC '12), 1–11.
 - [235] Schroeder, R., Huxor, A., & Smith, A. (2001). Activeworlds: Geography and social interaction in virtual reality. Futures, 33(7), 569–587. [https://doi.org/10.1016/S0016-3287\(01\)00002-7](https://doi.org/10.1016/S0016-3287(01)00002-7)
 - [236] Schyns, P. G., & Oliva, A. (1994). From blobs to boundary edges: Evidence for time-and spatial-scale-dependent scene recognition. Psychological Science, 5(4), 195–200. <https://doi.org/10.1111/j.1467-9280.1994.tb00500.x>
 - [237] Sears, A., Lin, M., Jacko, J., & Xiao, Y. (2003). When computers fade: Pervasive computing and situationally-induced impairments and disabilities. HCI International, 2(3), 1298–1302.

- [238] Shafrin, J., Sullivan, J., Goldman, D. P., & Gill, T. M. (2017). The association between observed mobility and quality of life in the near elderly. *PLoS ONE*, 12(8). <https://doi.org/10.1371/journal.pone.0182920>
- [239] Shinohara, K., & Wobbrock, J. O. (2011). In the shadow of misperception: Assistive technology use and social interactions. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*, 705–714. <https://doi.org/10.1145/1978942.1979044>
- [240] Singh, K. (2002). A fresh perspective. *Proceedings of the Graphics Interface 2002 Conference*, 17–24. <https://doi.org/10.20380/GI2002.03>
- [241] Siu, T., & Herskovic, V. (2013). SidebARs: Improving awareness of off-screen elements in mobile augmented reality. *Proceedings of the 2013 Chilean Conference on Human-Computer Interaction (ChileCHI '13)*, 36–41. <https://doi.org/10.1145/2535597.2535608>
- [242] Sivasamy, M., Sastry, V. N., & Gopalan, N. P. (2020). VRCAuth: Continuous authentication of users in virtual reality environment using head-movement. *5th International Conference on Communication and Electronics Systems (ICCES '20)*, 518–523. <https://doi.org/10.1109/ICCES48766.2020.9137914>
- [243] Slater, M., & Steed, A. (2000). A virtual presence counter. *Presence: Teleoperators and Virtual Environments*, 9(5), 413–434. <https://doi.org/10.1162/105474600566925>
- [244] Slater, M., & Usoh, M. (1994). Body-centered interaction in immersive virtual environments. In N. Magnenat, N. M. Thalmann, & D. Thalmann (Eds.), *Artificial Life and Virtual Reality* (pp. 125–148). John Wiley & Sons, Inc. <https://api.semanticscholar.org/CorpusID:15339904>
- [245] Slater, M., Usoh, M., & Steed, A. (1994). Depth of presence in virtual environments. *Presence: Teleoperators and Virtual Environments*, 3(2), 130–144. <https://doi.org/10.1162/pres.1994.3.2.130>
- [246] Slater, M., Usoh, M., & Steed, A. (1995). Taking steps: The influence of a walking technique on presence in virtual reality. *ACM Transactions on Computer-Human Interaction*, 2(3), 201–219. <https://doi.org/10.1145/210079.210084>
- [247] Smith, M. V., Calfee, R. P., Baumgarten, K. M., Brophy, R. H., & Wright, R. W. (2012). Upper extremity-specific measures of disability and outcomes in orthopaedic surgery. *Journal of Bone and Joint Surgery*, 94(3), 277–285. <https://doi.org/10.2106/JBJS.J.01744>
- [248] Sokolov, D., & Plemenos, D. (2007). High level methods for scene exploration. *Journal of Virtual Reality and Broadcasting*, 3(12), 1–12. <https://doi.org/10.20385/1860-2037/3.2006.12>
- [249] Sokolov, D., Plemenos, D., & Tamine, K. (2006). Viewpoint quality and global scene exploration strategies. *Proceedings of the First International Conference on Computer Graphics Theory and Applications (VISIGRAPP '06)*, 184–191. <https://doi.org/10.5220/0001350101840191>
- [250] Sonnet, H., Carpendale, S., & Strothotte, T. (2004). Integrating expanding annotations with a 3D explosion probe. *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI '04)*, 63–70. <https://doi.org/10.1145/989863.989871>
- [251] Steinicke, F., Bruder, G., Jerald, J., Frenz, H., & Lappe, M. (2010). Estimation of detection thresholds for redirected walking techniques. *IEEE Transactions on Visualization and Computer Graphics*, 16(1), 17–27. <https://doi.org/10.1109/TVCG.2009.62>
- [252] Stellmach, S., & Dachselt, R. (2012). Designing gaze-based user interfaces for steering in virtual environments. *Proceedings of the Symposium on Eye Tracking Research and Applications (ETRA '12)*, 131–138. <https://doi.org/10.1145/2168556.2168577>
- [253] Stiratelli, R., Laird, N., & Ware, J. H. (1984). Random-effects models for serial observations with binary response. *Biometrics*, 40(4), 961–971.
- [254] Stoakley, R., Conway, M. J., & Pausch, R. (1995). Virtual reality on a WIM: Interactive worlds in miniature. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '95)*, 265–272. <https://doi.org/10.1145/223904.223938>
- [255] Stoev, S. L., & Schmalstieg, D. (2002). Application and taxonomy of through-the-lens techniques. *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST '02)*, 57–64. <https://doi.org/10.1145/585740.585751>
- [256] Sukan, M., Elvezio, C., Oda, O., Feiner, S., & Tversky, B. (2014). ParaFrustum: Visualization techniques for guiding a user to a constrained set of viewing positions and orientations. *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology (UIST '14)*, 331–340.

<https://doi.org/10.1145/2642918.2647417>

- [257] Suma, E. A., Finkelstein, S. L., Clark, S., Goolkasian, P., & Hodges, L. F. (2010). Effects of travel technique and gender on a divided attention task in a virtual environment. *IEEE Symposium on 3D User Interfaces (3DUI '10)*, 27–34. <https://doi.org/10.1109/3DUI.2010.5444726>
- [258] Suma, E. A., Finkelstein, S. L., Reid, M., Ulinski, A., & Hodges, L. F. (2009). Real walking increases simulator sickness in navigationally complex virtual environments. *IEEE Virtual Reality Conference*, 245–246. <https://doi.org/10.1109/VR.2009.4811037>
- [259] Suma, E., Finkelstein, S., Reid, M., Babu, S., Ulinski, A., & Hodges, L. F. (2010). Evaluation of the cognitive effects of travel technique in complex real and virtual environments. *IEEE Transactions on Visualization and Computer Graphics*, 16(4), 690–702. <https://doi.org/10.1109/TVCG.2009.93>
- [260] Surale, H. B., Gupta, A., Hancock, M., & Vogel, D. (2019). TabletInVR: Exploring the design space for using a multi-touch tablet in virtual reality. *Proceedings of the 2019 SIGCHI Conference on Human Factors in Computing Systems (CHI '19)*, 1–13. <https://doi.org/10.1145/3290605.3300243>
- [261] Taheri, A., Caetano, A., & Sra, M. (2024). Virtual steps: The experience of walking for a lifelong wheelchair user in virtual reality. *IEEE Conference Virtual Reality and 3D User Interfaces (VR '24)*, 168–178. <https://doi.org/10.1109/VR58804.2024.00040>
- [262] Takahashi, S., Yoshida, K., Shimada, K., & Nishita, T. (2006). Occlusion-free animation of driving routes for car navigation systems. *IEEE Transactions on Visualization and Computer Graphics*, 12(5), 1141–1148. <https://doi.org/10.1109/TVCG.2006.167>
- [263] Tanenbaum, T. J., Hartoonian, N., & Bryan, J. (2020). "How do I make this thing smile?": An inventory of expressive nonverbal communication in commercial social virtual reality platforms. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*, 1–13. <https://doi.org/10.1145/3313831.3376606>
- [264] Templeman, J. N., Denbrook, P. S., & Sibert, L. E. (1999). Virtual locomotion: Walking in place through virtual environments. *Presence: Teleoperators and Virtual Environments*, 8(6), 598–617. <https://doi.org/10.1162/105474699566512>
- [265] Terziman, L., Marchal, M., Emily, M., Multon, F., Arnaldi, B., & Lécuyer, A. (2010). Shake-Your-Head: Revisiting walking-in-place for desktop virtual reality. *Proceedings of the 17th ACM Symposium on Virtual Reality Software and Technology (VRST '10)*, 27–34. <https://doi.org/10.1145/1889863.1889867>
- [266] Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, 381, 520–522. <https://doi.org/10.1038/381520a0>
- [267] Tregillus, S., & Folmer, E. (2016). VR-STEP: Walking-in-place using inertial sensing for hands-free navigation in mobile VR environments. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*, 1250–1255. <https://doi.org/10.1145/2858036.2858084>
- [268] Trewin, S. (2000). Configuration agents, control and privacy. *Proceedings on the 2000 Conference on Universal Usability (CUU '00)*, 9–16. <https://doi.org/10.1145/355460.355466>
- [269] Trindade, D. R., & Raposo, A. B. (2011). Improving 3D navigation in multiscale environments using cubemap-based techniques. *Proceedings of the 2011 ACM Symposium on Applied Computing (SAC '11)*, 1215–1221. <https://doi.org/10.1145/1982185.1982454>
- [270] Tversky, B., & Hemenway, K. (1983). Categories of Environmental Scenes. *Cognitive Psychology*, 15(1), 121–149. [https://doi.org/10.1016/0010-0285\(83\)90006-3](https://doi.org/10.1016/0010-0285(83)90006-3)
- [271] Usoh, M., Arthur, K., Whitton, M. C., Bastos, R., Steed, A., Slater, M., & Brooks, F. P. (1999). Walking > walking-in-place > flying, in virtual environments. *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99)*, 359–364. <https://doi.org/10.1145/311535.311589>
- [272] Uthe, N. (2014). Hand Walking. <https://www.youtube.com/watch?v=Bu6qUcJ-NL8>
- [273] Valentine, J. C., Hedges, L. V., & Cooper, H. M. (2009). *The Handbook of Research Synthesis and Meta-Analysis*. Russell Sage Foundation.
- [274] Valkov, D., Steinicke, F., Bruder, G., & Hinrichs, K. H. (2010). Traveling in 3D virtual environments with foot gestures and a multi-touch enabled WIM. *Proceedings of Virtual Reality International Conference*, 171–180.

- [275] Vatavu, R.-D., & Ungurean, O.-C. (2022). Understanding gesture input articulation with upper-body wearables for users with upper-body motor impairments. *Conference on Human Factors in Computing Systems (CHI '22)*, 1–16. <https://doi.org/10.1145/3491102.3501964>
- [276] Veas, E., Grasset, R., Kruijff, E., & Schmalstieg, D. (2012). Extended overview techniques for outdoor augmented reality. *IEEE Transactions on Visualization and Computer Graphics*, 18(4), 565–572. <https://doi.org/10.1109/TVCG.2012.44>
- [277] Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478. <https://doi.org/10.2307/30036540>
- [278] Viau, A., Feldman, A. G., McFadyen, B. J., & Levin, M. F. (2004). Reaching in reality and virtual reality: A comparison of movement kinematics in healthy subjects and in adults with hemiparesis. *Journal of NeuroEngineering and Rehabilitation*, 1, 1–7. <https://doi.org/10.1186/1743-0003-1-11>
- [279] Viegas, J., Conway, M. J., Williams, G., & Pausch, R. (1996). 3D magic lenses. *Proceedings of the 9th Annual ACM Symposium on User Interface Software and Technology (UIST '96)*, 51–58. <https://doi.org/10.1145/237091.237098>
- [280] Vogel, J., & Schiele, B. (2007). Semantic modeling of natural scenes for content-based image retrieval. *International Journal of Computer Vision*, 72(2), 133–157. <https://doi.org/10.1007/s11263-006-8614-1>
- [281] Walker, N., Philbin, D. A., & Fisk, A. D. (1997). Age-related differences in movement control: Adjusting submovement structure. *The Journals of Gerontology*, 52B(1), 40–53. <https://doi.org/10.1093/geronb/52B.1.P40>
- [282] Wang, J., & Lindeman, R. (2011). Isometric versus elastic surfboard interfaces for locomotion in virtual reality. *Proceedings of the 6th IASTED International Conference on Human-Computer Interaction, HCI 2011*. <https://doi.org/10.2316/P.2011.747-007>
- [283] Wang, J., & Lindeman, R. (2012). Leaning-based travel interfaces revisited: Frontal versus sidewise stances for flying in 3D virtual spaces. *Proceedings of the 18th ACM Symposium on Virtual Reality Software and Technology (VRST '12)*, 121–128. <https://doi.org/10.1145/2407336.2407360>
- [284] Ware, C., & Osborne, S. (1990). Exploration and virtual camera control in virtual three-dimensional environments. In *Proceedings of the 1990 symposium on Interactive 3D graphics (I3D '90)* (pp. 175–183). <https://doi.org/10.1145/91394.91442>
- [285] Weiss, R. S. (1994). *Learning From Strangers: The Art and Method of Qualitative Interview Studies*. The Free Press.
- [286] Wendt, J. D. (2010). *Real-Walking Models Improve Walking-in-Place Systems*. University of North Carolina at Chapel Hill.
- [287] Williamson, J. R., McGill, M., & Outram, K. (2019). PlaneVR: Social acceptability of virtual reality for aeroplane passengers. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, 1–14. <https://doi.org/10.1145/3290605.3300310>
- [288] Wingrave, C. A., Hacıahmetoglu, Y., & Bowman, D. A. (2006). Overcoming world in miniature limitations by a scaled and scrolling WIM. *IEEE Symposium on 3D User Interfaces (3DUI '06)*, 11–16. <https://doi.org/10.1109/VR.2006.106>
- [289] Wobbrock, J. O. (2019). Situationally aware mobile devices for overcoming situational impairments. *Proceedings of the ACM SIGCHI Symposium on Engineering Interactive Computing Systems (EICS '19)*, 1–18. <https://doi.org/10.1145/3319499.3330292>
- [290] Wobbrock, J. O., Findlater, L., Gergle, D., & Higgins, J. J. (2011). The aligned rank transform for nonparametric factorial analyses using only ANOVA procedures. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*, 143–146. <https://doi.org/10.1145/1978942.1978963>
- [291] Wobbrock, J. O., Gajos, K. Z., Kane, S. K., & Vanderheiden, G. C. (2018). Ability-based design. *Communications of the ACM*, 61(6), 62–71. <https://doi.org/10.1145/3148051>
- [292] Wobbrock, J. O., Kane, S. K., Gajos, K. Z., Harada, S., & Froehlich, J. (2011). Ability-based design: Concept, principles, and examples. *ACM Transactions on Accessible Computing*, 3(3), 1–36. <https://doi.org/10.1145/1952383.1952384>
- [293] Wobbrock, J. O., & Kientz, J. A. (2016). Research contributions in human-computer interaction. *Interactions*, May-June, 38–44. <https://doi.org/10.1145/2907069>

- [294] Wong, A., Gillis, H., & Peck, B. (2018). VR Accessibility Survey for People with Disabilities. <https://drive.google.com/file/d/0B0VwTVwReMqLMFIzdZVVaVdaTFk/view>
- [295] Xiao, J., Hays, J., Ehinger, K. A., Oliva, A., & Torralba, A. (2010). SUN database: Large-scale scene recognition from abbey to zoo. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 3485–3492. <https://doi.org/10.1109/CVPR.2010.5539970>
- [296] Xiao, R., & Benko, H. (2016). Augmenting the field-of-view of head-mounted displays with sparse peripheral displays. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*, 1221–1232. <https://doi.org/10.1145/2858036.2858212>
- [297] Yamada, W., & Manabe, H. (2016). Expanding the field-of-view of head-mounted displays with peripheral blurred images. *Adjunct Proceedings of the 29th Annual ACM Symposium on User Interface Software and Technology (UIST '16)*, 141–142. <https://doi.org/10.1145/2984751.2985735>
- [298] Yamagami, M., Junuzovic, S., Gonzalez-Franco, M., Ofek, E., Cutrell, E., Porter, J. R., Wilson, A. D., & Mott, M. E. (2022). Two-In-One: A design space for mapping unimanual input into bimanual interactions in VR for users with limited movement. *ACM Transactions on Accessible Computing*, 15(3), 1–25. <https://doi.org/10.1145/3510463>
- [299] Yan, Z., Lindeman, R. W., & Dey, A. (2016). Let your fingers do the walking: A unified approach for efficient short-, medium-, and long-distance travel in VR. *IEEE Symposium on 3D User Interfaces (3DUI '16)*, 27–30. <https://doi.org/10.1109/3DUI.2016.7460027>
- [300] Yarbus, A. L. (1968). *Eye Movements and Vision*. In *Neuropsychologia* (Issue 4). Springer New York. <https://doi.org/10.1007/978-1-4899-5379-7>
- [301] Yun, K., Peng, Y., Samaras, D., Zelinsky, G. J., & Berg, T. L. (2013). Exploring the role of gaze behavior and object detection in scene understanding. *Frontiers in Psychology*, 4, 1–14. <https://doi.org/10.3389/fpsyg.2013.00917>
- [302] Zambaka, C., Babu, S., Xiao, D., Ulinski, A., Hodges, L. F., & Lok, B. (2004). Effects of travel technique on cognition in virtual environments. *IEEE Virtual Reality*, 149–286. <https://doi.org/10.1109/VR.2004.1310068>
- [303] Zelinsky, G. J. (2013). Understanding scene understanding. *Frontiers in Psychology*, 4, 1–3.
- [304] Zhang, F., Chu, S., Pan, R., Ji, N., & Xi, L. (2017). Double hand-gesture interaction for walk-through in VR environment. *IEEE/ACIS 16th International Conference on Computer and Information Science (ICIS '17)*, 539–544. <https://doi.org/10.1109/ICIS.2017.7960051>
- [305] Zhang, X., & Furnas, G. W. (2005). MCVEs: Using cross-scale collaboration to support user interaction with multiscale structures. *Presence*, 14(1), 31–46. <https://doi.org/10.1162/1054746053890288>
- [306] Zhang, Z., Lai, Z., Xu, Y., Shao, L., Wu, J., & Xie, G. Sen. (2017). Discriminative elastic-net regularized linear regression. *IEEE Transactions on Image Processing*, 26(3), 1466–1481. <https://doi.org/10.1109/TIP.2017.2651396>
- [307] Zhao, D., Chai, S., Gao, T., Li, J., & Zhou, C. (2023). Physical mobility, social isolation and cognitive function: Are there really gender differences? *The American Journal of Geriatric Psychiatry*, 31(9), 726–736. <https://doi.org/https://doi.org/10.1016/j.jagp.2023.04.002>
- [308] Zhao, Y., Bennett, C. L., Benko, H., Cutrell, E., Holz, C., Morris, M. R., & Sinclair, M. (2018). Enabling people with visual impairments to navigate virtual reality with a haptic and auditory cane simulation. *Proceedings of the 2018 SIGCHI Conference on Human Factors in Computing Systems (CHI '18)*, 1–4. <https://doi.org/10.1145/3170427.3186485>
- [309] Zhao, Y., Morris, M. R., Cutrell, E., Holz, C., & Wilson, A. D. (2019). SeeingVR: A set of tools to make virtual reality more accessible to people with low vision. *Proceedings of the 2019 SIGCHI Conference on Human Factors in Computing Systems (CHI '19)*, 1–14. <https://doi.org/10.1145/3290605.3300341>
- [310] Zielasko, D., Horn, S., Freitag, S., Weyers, B., & Kühlen, T. W. (2016). Evaluation of hands-free HMD-based navigation techniques for immersive data analysis. *IEEE Symposium on 3D User Interfaces (3DUI '16)*, 113–119. <https://doi.org/10.1109/3DUI.2016.7460040>

A.1 CLASSIFICATION OF SCENE-VIEWING TECHNIQUES USING THE TAXONOMY

Table 20. Scene-viewing technique papers classified based on the visual properties and tasks that were used to evaluate the scene-viewing technique. Classification is based on the inspection of images in the paper and supplemental videos. Papers are grouped and labeled by the type of scene-viewing technique (e.g., “Scene Modification”).

Visual Properties		Tasks						
		Movement	Socializing, Collaboration	Offense, Defense	Exploration	Navigation	Creativity	Observation
Openness	Indoor		Social Cues (Mayer et al., 2020; Tanenbaum et al., 2020), Multi-scale (Piumsomboon et al., 2019)		Automatic (Desroche et al., 2007), Direct (Stoakley et al., 1995)	Automatic (Andujar et al., 2004; Chittaro et al., 2003), Scene Modification (Sukan et al., 2014), Amplified (LaViola Jr. et al., 2001; Ragan et al., 2017; Sargunam et al., 2017), Multi-scale (Argelaguet & Maignant, 2016), Direct (Ware & Osborne, 1990) (Eyeball in Hand, Flying vehicle control)		
	Outdoor		Social Cues (Tanenbaum et al., 2020), Multi-scale (Chenechal et al., 2016; X. Zhang & Furnas, 2005)		Multi-view (Looser et al., 2004; Stoev & Schmalstieg, 2002), Projection Distortion (Cui et al., 2010), Guided (Bell et al., 2001), Multi-scale (Langbehn et al., 2016)	Direct (Funk et al., 2019), Scene Modification (Lange et al., 2020; Pierce & Pausch, 2004; Takahashi et al., 2006), Multi-view (Bares & Lester, 1998; Brown & Hua, 2006; Veas et al., 2012), Amplified (Ngoc & Kalawsky, 2013), FOV Extension (Ardouin et al., 2012; R. Xiao & Benko, 2016), Multi-scale (Argelaguet & Maignant, 2016; Fukatsu, 1998; Trindade & Raposo, 2011), Guided (Elmqvist et al., 2007)	Guided (Gleicher & Witkin, 1992b)	FOV extension (Yamada & Manabe, 2016), Cue-based (Lin et al., 2017), Guided (Pavel et al., 2017)
	Abstract				Automatic (Sokolov et al., 2006; Sokolov & Plemenos, 2007), Scene Modification (McGuffin et al., 2003;	Cue-based (Gruenefeld, El Ali, et al., 2017, 2018), Multi-scale (Aspin & Hoang Le, 2007; Kopper et al.,	Direct (Grossman et al., 2002), Guided	

Visual Properties	Tasks							
	Movement	Socializing, Collaboration	Offense, Defense	Exploration	Navigation	Creativity	Observation	Productivity
				Preim et al., 1995, 1997, 1998), Multi-view (Viega et al., 1996), Projection Distortion (Carpendale et al., 1996; Singh, 2002)	2006; Menzner et al., 2020), Projection Distortion (Elmqvist & Tudoreanu, 2007), Direct (Ware & Osborne, 1990) (Scene in Hand), Scene Modification (Feiner & Seligmann, 1992; Sonnet et al., 2004)	(Burtnyk et al., 2002)		
Scale	Larger				Multi-scale (Argelaguet & Maignant, 2016; Kopper et al., 2006)(Aspin & Hoang Le, 2007), Direct (Ware & Osborne, 1990) (Scene in Hand)		Cue-based (Lin et al., 2017)	
	Human	Social Cues (Mayer et al., 2020; Tanenbaum et al., 2020), Multi-scale (Chenechal et al., 2016; Piumsomboon et al., 2019; X. Zhang & Furnas, 2005)		Automatic (Desroche et al., 2007) , Projection Distortion (Cui et al., 2010), Direct (Stoakley et al., 1995), Multi-scale (Langbehn et al., 2016)	Automatic (Andujar et al., 2004; Chittaro et al., 2003), Direct (Funk et al., 2019; Ware & Osborne, 1990) (Eyeball in hand, Flying vehicle control), Scene Modification (Lange et al., 2020; Sukan et al., 2014; Takahashi et al., 2006), Multi-view (Veas et al., 2012), Amplified (LaViola Jr. et al., 2001; Ngoc & Kalawsky, 2013; Ragan et al., 2017; Sargunam et al., 2017), FOV Extension (Ardouin et al., 2012; R. Xiao & Benko, 2016), Multi-scale (Argelaguet & Maignant, 2016; Aspin & Hoang Le, 2007; Fukatsu, 1998; Kopper et al., 2006; Menzner et al., 2020; Trindade & Raposo, 2011), Guided (Elmqvist et al., 2007), Cue-based		FOV extension (Yamada & Manabe, 2016), Guided (Pavel et al., 2017)	

Visual Properties	Tasks							
	Movement	Socializing, Collaboration	Offense, Defense	Exploration	Navigation	Creativity	Observation	Productivity
					(Gruenefeld, El Ali, et al., 2017, 2018)			
	Smaller	Multi-scale (Chenechal et al., 2016; Piumsomboon et al., 2019; X. Zhang & Furnas, 2005)		Automatic (Sokolov et al., 2006; Sokolov & Plemenos, 2007), Scene Modification (McGuffin et al., 2003; Preim et al., 1995, 1997, 1998), Multi-view (Looser et al., 2004; Stoev & Schmalstieg, 2002; Viega et al., 1996), Guided (Bell et al., 2001), Projection Distortion (Carpendale et al., 1996; Cui et al., 2010; Singh, 2002), Direct (Stoakley et al., 1995), Multi-scale (Langbehn et al., 2016)	Scene Modification (Feiner & Seligmann, 1992; Pierce & Pausch, 2004; Sonnet et al., 2004), Multi-view (Bares & Lester, 1998; Brown & Hua, 2006), Amplified (LaViola Jr. et al., 2001), Multi-scale (Argelaguet & Maignant, 2016; Fukatsu, 1998; Menzner et al., 2020; Trindade & Raposo, 2011), Projection Distortion (Elmqvist & Tudoreanu, 2007)	Direct (Grossman et al., 2002), Guided (Burtnyk et al., 2002; Gleicher & Witkin, 1992b)		
Area	Large	Multi-scale (Chenechal et al., 2016; X. Zhang & Furnas, 2005)		Automatic (Sokolov et al., 2006; Sokolov & Plemenos, 2007), Scene Modification (McGuffin et al., 2003; Preim et al., 1995, 1997, 1998), Multi-view (Stoev & Schmalstieg, 2002; Viega et al., 1996), Projection Distortion (Carpendale et al., 1996; Cui et al., 2010; Singh, 2002), Guided (Bell et al., 2001), Multi-scale (Langbehn et al., 2016)	Direct (Funk et al., 2019; Ware & Osborne, 1990) (Scene in Hand), Scene Modification (Lange et al., 2020; Pierce & Pausch, 2004; Sonnet et al., 2004; Takahashi et al., 2006), Multi-view (Bares & Lester, 1998; Brown & Hua, 2006; Veas et al., 2012), Amplified (LaViola Jr. et al., 2001; Ngoc & Kalawsky, 2013; Ragan et al., 2017), FOV Extension (R. Xiao & Benko, 2016), Multi-scale (Argelaguet & Maignant, 2016; Aspin & Hoang Le, 2007; Fukatsu, 1998; Kopper et al., 2006; Menzner et al., 2020; Trindade & Raposo, 2011), Guided (Elmqvist et al.,	Direct (Grossman et al., 2002), Guided (Burtnyk et al., 2002; Gleicher & Witkin, 1992b)	FOV extension (Yamada & Manabe, 2016), Cue-based (Lin et al., 2017)	

Visual Properties		Tasks							
		Movement	Socializing, Collaboration	Offense, Defense	Exploration	Navigation	Creativity	Observation	Productivity
						2007), Projection Distortion (Elmqvist & Tudoreanu, 2007)			
	Medium		Social Cues (Tanenbaum et al., 2020), Multi-scale (Piumsomboon et al., 2019)		Direct (Stoakley et al., 1995)	Automatic (Andujar et al., 2004; Chittaro et al., 2003), Amplified (LaViola Jr. et al., 2001), FOV extension (Ardouin et al., 2012), Multi-scale (Argelaguet & Maignant, 2016), Cue-based (Gruenefeld, El Ali, et al., 2017, 2018)		Guided (Pavel et al., 2017)	
	Small		Social Cues (Chenechal et al., 2016; Mayer et al., 2020)		Automatic (Desroche et al., 2007), Multi-view (Looser et al., 2004)	Scene Modification (Sukan et al., 2014), Amplified (Sargunam et al., 2017), Multi-scale (Argelaguet & Maignant, 2016), Direct (Ware & Osborne, 1990) (Eyeball in hand, Flying vehicle control)			
Object Density	High		Multi-scale (Chenechal et al., 2016)		Multi-view (Stoev & Schmalstieg, 2002), Guided (Bell et al., 2001), Multi-scale (Langbehn et al., 2016), Projection Distortion (Carpendale et al., 1996; Cui et al., 2010)	Scene Modification (Lange et al., 2020) (Feiner & Seligmann, 1992), Multi-view (Bares & Lester, 1998; Brown & Hua, 2006), Amplified (Ragan et al., 2017), FOV Extension (R. Xiao & Benko, 2016), Multi-scale (Argelaguet & Maignant, 2016; Trindade & Raposo, 2011), Guided (Elmqvist et al., 2007), Projection Distortion (Elmqvist & Tudoreanu, 2007)			
	Moderate		Social Cues (Tanenbaum et al., 2020), Multi-scale (Chenechal et al., 2016;		Automatic (Desroche et al., 2007), Direct (Stoakley et al., 1995)	Automatic (Andujar et al., 2004; Chittaro et al., 2003), Direct (Funk et al., 2019), FOV extension (Ardouin et al., 2012; R. Xiao & Benko,		FOV extension (Yamada & Manabe, 2016), Cue-based (Lin et al., 2017),	

Visual Properties	Tasks							
	Movement	Socializing, Collaboration	Offense, Defense	Exploration	Navigation	Creativity	Observation	Productivity
		Piumsomboon et al., 2019; X. Zhang & Furnas, 2005)			2016), Multi-scale (Fukatsu, 1998; Menzner et al., 2020), Multi-view (Veas et al., 2012)		Guided (Pavel et al., 2017)	
	Low	Social Cues (Chenechal et al., 2016; Mayer et al., 2020)		Automatic (Sokolov et al., 2006; Sokolov & Plemenos, 2007), Scene Modification (McGuffin et al., 2003; Preim et al., 1995, 1997, 1998), Direct (Ware & Osborne, 1990) (Flying vehicle control), Multi-view (Looser et al., 2004; Viega et al., 1996), Projection Distortion (Singh, 2002)	Scene Modification (Pierce & Pausch, 2004; Sukan et al., 2014; Takahashi et al., 2006), Amplified (LaViola Jr. et al., 2001; Ngoc & Kalawsky, 2013; Sargunam et al., 2017), Multi-scale (Kopper et al., 2006) (Aspin & Hoang Le, 2007; Trindade & Raposo, 2011), Cue-based (Gruenefeld, El Ali, et al., 2017, 2018), Direct (Ware & Osborne, 1990) (Scene in Hand)	Direct (Grossman et al., 2002), Guided (Burtnyk et al., 2002; Gleicher & Witkin, 1992b)		
Object Tracking	Multiple				Multi-view (Bares & Lester, 1998)	Guided (Gleicher & Witkin, 1992b)	Cue-based (Lin et al., 2017), Guided (Pavel et al., 2017)	
	One	Social Cues (Mayer et al., 2020), Multi-scale (Piumsomboon et al., 2019; X. Zhang & Furnas, 2005)			Automatic (Chittaro et al., 2003), Scene Modification (Lange et al., 2020; Takahashi et al., 2006), FOV extension (R. Xiao & Benko, 2016)			
	None	Social Cues (Tanenbaum et al., 2020), Multi-scale (Chenechal et al., 2016)		Automatic (Desroche et al., 2007; Sokolov et al., 2006; Sokolov & Plemenos, 2007) Scene Modification (McGuffin et al., 2003; Preim et al., 1995, 1997, 1998), Multi-view (Looser et al., 2004; Stoev & Schmalstieg, 2002; Viega et al., 1996), Projection Distortion (Carpendale et al., 1996; Cui et al., 2010;	Automatic (Andujar et al., 2004), Multiview (Brown & Hua, 2006), Direct [18] [53] (Scene in Hand), Scene Modification (Feiner & Seligmann, 1992; Pierce & Pausch, 2004; Sonnet et al., 2004; Sukan et al., 2014), Amplified (LaViola Jr. et al., 2001; Ngoc & Kalawsky, 2013; Ragan et al., 2017; Sargunam et al.,	Direct (Grossman et al., 2002), Guided (Burtnyk et al., 2002)	FOV extension (Yamada & Manabe, 2016), Multi-scale (Menzner et al., 2020)	

		Tasks							
Visual Properties		Movement	Socializing, Collaboration	Offense, Defense	Exploration	Navigation	Creativity	Observation	Productivity
					Singh, 2002), Guided (Bell et al., 2001), Multi-scale (Langbehn et al., 2016), Direct (Sonnet et al., 2004; Stoakley et al., 1995)	2017), FOV extension (Ardouin et al., 2012), Guided (Elmqvist et al., 2007), Multi-scale (Argelaguet & Maignant, 2016; Aspin & Hoang Le, 2007; Fukatsu, 1998; Kopper et al., 2006; Menzner et al., 2020; Trindade & Raposo, 2011), Projection Distortion (Elmqvist & Tudoreanu, 2007), Cue-based (Gruenefeld, El Ali, et al., 2017, 2018)			
Scene Changes	Frequent		Multi-scale (Chenechal et al., 2016)		Direct (Stoakley et al., 1995; Ware & Osborne, 1990) (Flying vehicle control)	Amplified (LaViola Jr. et al., 2001; Ragan et al., 2017; Sargunam et al., 2017), FOV extension (Ardouin et al., 2012), Multi-scale (Argelaguet & Maignant, 2016; Kopper et al., 2006; Trindade & Raposo, 2011)		Cue-based (Lin et al., 2017), Guided (Pavel et al., 2017)	
	Infrequent		Social Cues (Mayer et al., 2020; Tanenbaum et al., 2020), Multi-scale (Piumsomboon et al., 2019; X. Zhang & Furnas, 2005)		Automatic (Desroche et al., 2007; Sokolov et al., 2006; Sokolov & Plemenos, 2007), Scene Modification (McGuffin et al., 2003; Preim et al., 1995, 1997, 1998; Sonnet et al., 2004), Multi-view (Looser et al., 2004; Stoev & Schmalstieg, 2002; Viega et al., 1996), Projection Distortion (Carpendale et al., 1996; Cui et al., 2010; Singh, 2002), Guided (Bell et al., 2001), Multi-scale (Langbehn et al., 2016)	Automatic (Andujar et al., 2004; Chittaro et al., 2003), Direct [18] [7,53] (Scene in hand), Scene Modification (Feiner & Seligmann, 1992; Lange et al., 2020; Pierce & Pausch, 2004; Sonnet et al., 2004; Sukan et al., 2014; Takahashi et al., 2006), Multi-view (Bares & Lester, 1998; Ngoc & Kalawsky, 2013) (Veas et al., 2012), Multi-scale (Fukatsu, 1998; Menzner et al., 2020) (Aspin & Hoang Le, 2007), Guided (Elmqvist et al., 2007),	Direct (Grossman et al., 2002), Guided (Burtnyk et al., 2002; Gleicher & Witkin, 1992b)	FOV extension (Yamada & Manabe, 2016)	

		Tasks							
Visual Properties		Movement	Socializing, Collaboration	Offense, Defense	Exploration	Navigation	Creativity	Observation	Productivity
						Cue-based (Gruenefeld, El Ali, et al., 2017, 2018), Projection Distortion (Elmqvist & Tudoreanu, 2007), FOV extension (R. Xiao & Benko, 2016)			
Contains Social Actors	Yes		Social Cues (Mayer et al., 2020; Tanenbaum et al., 2020), Multi-scale (Chenechal et al., 2016; Piumsomboon et al., 2019; X. Zhang & Furnas, 2005)			Automatic (Chittaro et al., 2003), Multi-view [5] [23], FOV extension (Ardouin et al., 2012)			
	No				Automatic (Desroche et al., 2007; Sokolov et al., 2006; Sokolov & Plemenos, 2007), Direct (Sonnet et al., 2004; Stoakley et al., 1995; Ware & Osborne, 1990) (Flying vehicle control), Scene Modification (McGuffin et al., 2003; Preim et al., 1995, 1997, 1998), Multi-view (Looser et al., 2004; Stoev & Schmalstieg, 2002; Viega et al., 1996), Projection Distortion (Carpendale et al., 1996; Cui et al., 2010; Elmqvist & Tudoreanu, 2007; Singh, 2002), Guided (Bell et al., 2001), Multi-scale (Langbehn et al., 2016)	Automatic (Andujar et al., 2004), Direct (Funk et al., 2019), Scene Modification (Lange et al., 2020; Pierce & Pausch, 2004; Sonnet et al., 2004; Sukan et al., 2014; Takahashi et al., 2006), Amplified (LaViola Jr. et al., 2001; Ngoc & Kalawsky, 2013; Ragan et al., 2017; Sargunam et al., 2017), FOV Extension (R. Xiao & Benko, 2016), Multi-scale (Argelaguet & Maignant, 2016; Aspin & Hoang Le, 2007; Fukatsu, 1998; Menzner et al., 2020; Trindade & Raposo, 2011), Guided (Elmqvist et al., 2007), Cue-based (Gruenefeld, El Ali, et al., 2017, 2018), Projection Distortion (Elmqvist & Tudoreanu, 2007), Direct (Brown & Hua, 2006; Ware	Direct (Grossman et al., 2002), Guided (Burtnyk et al., 2002; Gleicher & Witkin, 1992b)	FOV extension (Yamada & Manabe, 2016), Multi-scale (Menzner et al., 2020), Cue-based (Lin et al., 2017), Guided (Pavel et al., 2017)	

Visual Properties	Tasks							
	Movement	Socializing, Collaboration	Offense, Defense	Exploration	Navigation	Creativity	Observation	Productivity
					& Osborne, 1990) (Scene in Hand), Scene Modification (Feiner & Seligmann, 1992), Multi- view (Veas et al., 2012)			

A.2 CLASSIFICATION OF VIRTUAL REALITY APPLICATIONS USING THE SCENE TAXONOMY

Table 21. Twenty-nine of the most popular virtual reality (VR) applications in 2019 and 2020 classified based on the main visual properties and tasks of the virtual environment (VE).

Visual Properties		Tasks Movement	Socializing, Collaboration	Offense, Defense	Exploration	Navigation	Creativity	Observation	Productivity
Openness	Indoor	Beatsaber, Job Simulator, Dance Central, Vacation Simulator	VRChat, Altspace	Blade and Sorcery, Saints and Sinners, Robo Recall, Arizona Sunshine, Superhot VR, Pistol Whip, Gorn, Onward	Lone Echo, I Expect You to Die, The Room VR	Dreadhalls			
	Outdoor	The Climb			Ocean Rift, Minecraft, Google Earth	MazeVR	Kingspray Graffiti, Tiltbrush	Within	Virtual Desktop
	Abstract						Quill		Supermedium
Scale	Larger						Quill		Virtual Desktop
	Human	Beatsaber, The Climb, Job Simulator, Vacation Simulator, Dance Central	VRChat, Altspace	Blade and Sorcery, Saints and Sinners, Robo Recall, Arizona sunshine, Superhot VR, Onward, Pistol Whip, Gorn	Lone Echo, I Expect You to Die, Ocean Rift, Minecraft, The Room VR	MazeVR, Dreadhalls	Kingspray graffiti, Tiltbrush, Quill	Within	Supermedium
	Smaller				Google Earth				
Area	Large	Beatsaber, Job Simulator, Dance Central, The climb		Blade and Sorcery, Onward, Pistol Whip	Google Earth, Ocean Rift, Minecraft		Tiltbrush, Quill	Within	Virtual Desktop, Supermedium
	Medium	Vacation Simulator	VRChat, Altspace	Saints and Sinners, Robo Recall, Gorn, Superhot VR	I Expect You to Die		Kingspray graffiti		
	Small			Arizona Sunshine	Lone Echo, The Room VR	MazeVR, Dreadhalls			
Object Density	High			Saints and Sinners, Robo Recall, Onward, Pistol Whip	Lone Echo, Google Earth, Ocean Rift, Minecraft				
	Moderate	Vacation Simulator		Superhot VR	I Expect You to Die		Kingspray Graffiti		
	Low	Beatsaber, Job Simulator	VRChat, Altspace	Blade and Sorcery, Arizona Sunshine, Gorn	The Room VR	DreadHalls, MazeVR	Tiltbrush, Quill	Within	Virtual Desktop, Supermedium
Object Tracking	Multiple	Beatsaber		Blade and Sorcery, Superhot VR, Onward, Pistol Whip	Ocean Rift				
	One	Job Simulator, Vacation Simulator, Dance Central		Arizona Sunshine	Lone Echo, Gorn				

	None	The Climb	VRChat, Altspace	Saints and Sinners, Robo Recall	I Expect You to Die, Google Earth, Minecraft, The Room VR	Dreadhalls, MazeVR	Tiltbrush, Quill, Kingspray Graffiti	Within	Virtual Desktop, Supermedium
Scene Changes	Frequent			Saints and Sinners, Arizona Sunshine, Superhot VR,	Lone Echo	Dreadhalls, MazeVR		Within	
	Infrequent	Beatsaber, Job Simulator, Vacation Simulator, Dance Central, The Climb	VRChat, Altspace	Blade and Sorcery, Onward, Pistol Whip, Gorn	I Expect You to Die, Google Earth, Ocean Rift, Minecraft, The Room VR		Kingspray Graffiti, Tiltbrush, Quill		Virtual Desktop, Supermedium
Contains Social Actors	Yes	Job Simulator, Vacation Simulator, Dance Central	VRChat, Altspace	Onward	Lone Echo		Kingspray Graffiti		
	No	Beatsaber, The Climb		Blade and Sorcery, Saints and Sinners, Robo Recall, Arizona Sunshine, Superhot VR, Pistol Whip	I Expect to Die, Google Earth, Ocean Rift, Minecraft, The Room VR, Gorn	Dreadhalls, MazeVR	Tiltbrush, Quill	Within	Virtual Desktop, Supermedium

A.3 QUICKDASH QUESTIONNAIRE

QuickDASH					
Please rate your ability to do the following activities in the last week by circling the number below the appropriate response.					
	NO DIFFICULTY	MILD DIFFICULTY	MODERATE DIFFICULTY	SEVERE DIFFICULTY	UNABLE
1. Open a tight or new jar.	1	2	3	4	5
2. Do heavy household chores (e.g., wash walls, floors).	1	2	3	4	5
3. Carry a shopping bag or briefcase.	1	2	3	4	5
4. Wash your back.	1	2	3	4	5
5. Use a knife to cut food.	1	2	3	4	5
6. Recreational activities in which you take some force or impact through your arm, shoulder or hand (e.g., golf, hammering, tennis, etc.).	1	2	3	4	5

	NOT AT ALL	SLIGHTLY	MODERATELY	QUITE A BIT	EXTREMELY
7. During the past week, to what extent has your arm, shoulder or hand problem interfered with your normal social activities with family, friends, neighbours or groups?	1	2	3	4	5

	NOT LIMITED AT ALL	SLIGHTLY LIMITED	MODERATELY LIMITED	VERY LIMITED	UNABLE
8. During the past week, were you limited in your work or other regular daily activities as a result of your arm, shoulder or hand problem?	1	2	3	4	5

Please rate the severity of the following symptoms in the last week. (circle number)					
	NONE	MILD	MODERATE	SEVERE	EXTREME
9. Arm, shoulder or hand pain.	1	2	3	4	5
10. Tingling (pins and needles) in your arm, shoulder or hand.	1	2	3	4	5

	NO DIFFICULTY	MILD DIFFICULTY	MODERATE DIFFICULTY	SEVERE DIFFICULTY	SO MUCH DIFFICULTY THAT I CAN'T SLEEP
11. During the past week, how much difficulty have you had sleeping because of the pain in your arm, shoulder or hand? (circle number)	1	2	3	4	5

QuickDASH DISABILITY/SYMPTOM SCORE = $\left(\frac{\text{sum of } n \text{ responses}}{n} - 1 \right) \times 25$, where n is equal to the number of completed responses.

A QuickDASH score may not be calculated if there is greater than 1 missing item.

Figure 91. The QuickDASH questionnaire used in the locomotion technique study to measure upper-body impairment. QuickDASH is a validated questionnaire (Gummesson et al., 2006).

A.4 TECHNOLOGY-RELATED (PHYSICAL) IMPAIRMENTS QUESTIONNAIRE (TRIQ)

Check all the impairments that you experience:

Slow movements
Low strength in core, shoulders, neck, arms, hands or fingers
Tremor
Poor coordination
Rapid fatigue when using core, shoulders, neck, arms, hands or fingers
Difficulty gripping objects
Difficulty holding objects
Lack of sensation in core, shoulders, neck, arms, hands or fingers
Difficult controlling direction of hand or arm movement
Difficulty controlling distance of hand or arm movement
Limited range of motion in core, shoulders, neck, arms, hands, wrists or fingers
Pain in core, shoulders, neck, arms, hands, wrists or fingers
Poor precision with finger(s)
Poor finger isolation
Limited wrist extension or flexion
Difficulty moving core, shoulders, neck, arms, hands or fingers quickly
Difficulty moving core, shoulders, neck, arms, hands or fingers at the right time
Difficulty balancing while seated
Limited mobility of lower body (legs or feet)

Figure 92. The **T**echnology-**R**elated **P**hysical **I**mpairment **Q**uestionnaire (TRIQ) questionnaire used in the locomotion studies. It is not standardized but based on literature in accessible computing (Findlater et al., 2010; Mott et al., 2020).

A.5 GROUP COMPARISON STUDY: NON-SIGNIFICANT OUTCOME BAR PLOTS

A.5.1 Obstacles Hit by Impairment Level




Figure 93. Bar plot of mean Obstacles Hit by Technique and Impairment Level. Lower is better. Means are at the bottom of error bars. No significant differences were found.

A.5.2 Presence by Group




Figure 94. Bar plot of mean Presence rating (1-6) by Technique and Group. Higher is better. Means are at the bottom of bars. No significant differences were found.

A.5.3 Presence by Impairment Level




Figure 95. Bar plot of mean Presence rating (1-6) by Technique and Impairment Level. Higher is better. Means are at the bottom of bars. No significant differences were found.

A.5.4 Sickness by Group




Figure 96. Bar plot of mean Sickness rating (1-3) by Technique and Group. Lower is better. Means are at the bottom of bars. No significant differences were found.

A.5.5 Sickness by Impairment Level




Figure 97. Bar plot of mean simulator Sickness rating (1-3) by Technique and Impairment Level. Lower is better. Means are at the bottom of bars. No significant differences were found.

A.5.6 Mental Demand by Group




Figure 98. Bar plot of mean Mental Demand rating (1-7) by technique and group. Lower is better. Means are at the bottom of bars. No significant differences were found.

A.5.7 Effort by Group




Figure 99. Bar plot of mean Effort rating (1-7) by Technique and Group. Lower is better. Means are at the bottom of bars. No significant differences were found.

A.5.8 Performance by Group




Figure 100. Bar plot of mean Performance rating (1-7) by Technique and Group. Higher is better. Means are at the bottom of bars. No significant differences were found.

A.6 METRICS STUDY: REMAINING METRICS

A.6.1 Movement-Related Metrics: Device Velocity and Acceleration




Figure 101. Participants With Impairments had a significantly lower Left Velocity for Teleport.




Figure 102. There were no significant differences in Left Acceleration between Groups for all Techniques.

A.6.2 Movement-Related Metrics: Device Angular Velocity and Acceleration




Figure 103. Participants With Impairments had a significantly lower Left Angular Velocity for Grab, Throw, and Teleport.




Figure 104. Participants With Impairments had a significantly lower Left Angular Acceleration for Teleport.




Figure 105. There were no significant differences in Right Angular Acceleration between Groups for all Techniques.

Button-Related Metrics: Button Pressure




Figure 106. Participants With Impairments had a significantly higher Left Trigger Pressure for Sliding and Throw and significantly lower Left Trigger Pressure for Astral and Chicken.




Figure 107. Participants With Impairments had a significantly higher Right Trigger Pressure for Astral, Chicken, Sliding, and Throw and a significantly lower Right Trigger Pressure for Grab.




Figure 108. Participants With Impairments had a significantly higher Left Grip Pressure for Astral, Chicken, and Sliding, and a significantly lower Left Grip Pressure for Throw.

A.6.3 Button-Related Metrics: Thumbstick Variability and Extent




Figure 109. Participants With Impairments had a significantly lower Left Thumbstick Variability for Astral because of an error with Astral's setup for the other Group.




Figure 110. Participants With Impairments had a significantly lower Left Thumbstick Variability for Astral because of an error with Astral's setup for the other Group.

A.6.4 Movement-Related Metrics: Device-Pair Variability and Extent




Figure 111. Participants With Impairments had significantly higher Left Hand to Head Variability for Grab and Throw.




Figure 112. Participants With Impairments had significantly higher Right Hand to Left Hand Variability for Grab and Sliding.

A.6.5 Movement-Related Metrics: Device Submovement Count



Figure 113. Participants With Impairments had significantly higher Head Submovements for Grab.




Figure 114. Participants With Impairments had significantly higher Left Submovements for Astral, Grab, and Teleport.

A.7 QUESTIONNAIRE SCENARIO: DATA TABLE SNIPPET

Table 22. A sample of the dataset used to train the model in the questionnaire scenario. Data are from two participants, a participant with impairments (PWI) and a participant without impairments (PWO). Column names are provided in the column name key.

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	34	25	26	27	28	29	30	31	32	33	34
PWI	Astral	2.57	2	4	3	5	1	3	3	3	3	4	52.27	3	0	0	1	1	0	1	0	1	0	0	1	1	0	0	1	0	0	1	1
PWI	Chicken	2.93	2	4	3	5	1	3	3	3	3	4	52.27	3	0	0	1	1	0	1	0	1	0	0	1	1	0	0	1	0	0	1	1
PWI	Grab	2.90	2	4	3	5	1	3	3	3	3	4	52.27	3	0	0	1	1	0	1	0	1	0	0	1	1	0	0	1	0	0	1	1
PWI	Sliding	2.78	2	4	3	5	1	3	3	3	3	4	52.27	3	0	0	1	1	0	1	0	1	0	0	1	1	0	0	1	0	0	1	1
PWI	Teleport	2.26	2	4	3	5	1	3	3	3	3	4	52.27	3	0	0	1	1	0	1	0	1	0	0	1	1	0	0	1	0	0	1	1
PWI	Throw	2.85	2	4	3	5	1	3	3	3	3	4	52.27	3	0	0	1	1	0	1	0	1	0	0	1	1	0	0	1	0	0	1	1
PWO	Astral	2.68	2	1	1	1	1	1	1	1	1	1	2.27	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
PWO	Chicken	2.70	2	1	1	1	1	1	1	1	1	1	2.27	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
PWO	Grab	2.87	2	1	1	1	1	1	1	1	1	1	2.27	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
PWO	Sliding	2.72	2	1	1	1	1	1	1	1	1	1	2.27	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
PWO	Teleport	2.27	2	1	1	1	1	1	1	1	1	1	2.27	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
PWO	Throw	2.73	2	1	1	1	1	1	1	1	1	1	2.27	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Column Name Key

1	iD	12	painQD	23	sensation	34	lowerBody
2	technique	13	tinglingQD	24	direction		
3	logTrialTime	14	scoreQD	25	distance		
4	openJarQD	15	difficultySleepingQD	26	rom		
5	choresQD	16	gripping	27	pain		
6	carryBagQD	17	slowMovements	28	precision		
7	washBackQD	18	lowStrength	29	isolation		
8	knifeCutFoodQD	19	tremor	30	wrist		
9	recreationalActivitiesQD	20	coordination	31	quickly		
10	socialActivitiesQD	21	fatigue	32	rightTime		
11	activitiesLimitedQD	22	holding	33	balancing		

A.8 QUESTIONNAIRE SCENARIO: DATA TABLE USED FOR RANKING LOCOMOTION TECHNIQUES

Table 23. The data table used to rank locomotion techniques for each participant. Times are predicted with the questionnaire scenario model. Columns represent the Participant ID, Technique, the average Actual Trial Time in seconds, the average Predicted Trial Time in seconds and the Difference between the Predicted and Actual Trial Times in seconds. Participant IDs leading with “P” are participants with impairments and IDs leading with “PC” are participants without impairments.

Participant	Technique	Actual Trial Time (s)	Predicted Trial Time (s)	Difference (s)
P00	Astral	12.30	14.49	-2.20
P00	Chicken	15.08	18.74	-3.66
P00	Grab	13.67	18.85	-5.18
P00	Sliding	13.72	15.98	-2.26
P00	Teleport	8.77	12.55	-3.79
P00	Throw	16.91	19.00	-2.09
P01	Astral	31.56	17.99	13.57
P01	Chicken	21.26	23.84	-2.57
P01	Grab	30.77	23.41	7.36
P01	Sliding	22.79	20.24	2.56
P01	Teleport	30.22	15.07	15.15
P01	Throw	31.53	24.05	7.47
P02	Astral	13.44	16.00	-2.56
P02	Chicken	25.27	19.94	5.33
P02	Grab	21.06	20.09	0.97
P02	Sliding	17.52	17.18	0.34
P02	Teleport	12.62	13.34	-0.73
P02	Throw	19.90	20.70	-0.80
P03	Astral	14.67	20.05	-5.38
P03	Chicken	19.67	25.14	-5.47
P03	Grab	24.64	24.65	-0.01
P03	Sliding	19.60	21.48	-1.88
P03	Teleport	15.75	16.15	-0.40
P03	Throw	27.48	24.71	2.77
P04	Astral	14.46	13.76	0.70
P04	Chicken	15.41	17.86	-2.45
P04	Grab	22.27	17.37	4.91
P04	Sliding	14.57	15.26	-0.69
P04	Teleport	10.35	11.77	-1.42
P04	Throw	17.48	18.18	-0.70
P05	Astral	14.47	15.52	-1.05
P05	Chicken	25.02	19.29	5.73
P05	Grab	26.26	19.14	7.12
P05	Sliding	16.03	16.99	-0.96
P05	Teleport	14.22	13.00	1.22
P05	Throw	23.70	19.74	3.97
P06	Astral	12.58	13.52	-0.94
P06	Chicken	22.18	16.97	5.21
P06	Grab	22.56	16.96	5.60
P06	Sliding	14.83	14.85	-0.02
P06	Teleport	10.03	11.59	-1.56
P06	Throw	15.87	17.77	-1.90

P07	Astral	14.76	17.61	-2.85
P07	Chicken	18.35	22.73	-4.38
P07	Grab	27.57	21.12	6.44
P07	Sliding	16.80	19.42	-2.62
P07	Teleport	14.24	14.74	-0.50
P07	Throw	25.96	22.24	3.72
P08	Astral	19.24	19.50	-0.27
P08	Chicken	30.39	24.77	5.61
P08	Grab	23.25	25.70	-2.46
P08	Sliding	15.70	23.00	-7.30
P08	Teleport	25.13	15.67	9.46
P08	Throw	31.52	25.34	6.18
P09	Astral	16.21	15.83	0.37
P09	Chicken	18.39	20.34	-1.95
P09	Grab	22.69	20.08	2.61
P09	Sliding	16.69	17.24	-0.55
P09	Teleport	12.62	13.47	-0.85
P09	Throw	23.71	20.40	3.31
P10	Astral	25.84	20.76	5.08
P10	Chicken	22.16	27.51	-5.34
P10	Grab	27.18	26.90	0.28
P10	Sliding	19.49	23.25	-3.76
P10	Teleport	10.71	18.62	-7.91
P10	Throw	29.23	27.36	1.87
P11	Astral	11.76	16.75	-4.99
P11	Chicken	25.81	19.97	5.84
P11	Grab	26.23	19.97	6.26
P11	Sliding	17.31	17.56	-0.25
P11	Teleport	10.66	13.88	-3.22
P11	Throw	26.02	20.44	5.59
P12	Astral	17.03	18.51	-1.49
P12	Chicken	23.95	23.33	0.63
P12	Grab	28.56	22.96	5.60
P12	Sliding	20.18	20.04	0.14
P12	Teleport	13.92	15.73	-1.81
P12	Throw	30.07	23.60	6.48
P13	Astral	30.19	20.63	9.56
P13	Chicken	31.55	28.11	3.44
P13	Grab	29.85	28.63	1.22
P13	Sliding	16.80	26.16	-9.36
P13	Teleport	31.56	17.06	14.51
P13	Throw	31.56	29.16	2.40
P14	Astral	14.94	16.61	-1.67
P14	Chicken	24.21	20.62	3.59
P14	Grab	26.12	20.33	5.80
P14	Sliding	19.93	17.64	2.29
P14	Teleport	13.84	13.72	0.13
P14	Throw	28.16	20.77	7.39
P15	Astral	14.76	18.51	-3.75
P15	Chicken	20.52	23.53	-3.01
P15	Grab	27.36	22.57	4.79
P15	Sliding	16.44	20.20	-3.76
P15	Teleport	15.20	15.24	-0.03

P15	Throw	22.93	23.76	-0.83
P16	Astral	13.11	15.51	-2.40
P16	Chicken	18.69	19.40	-0.71
P16	Grab	18.18	19.42	-1.24
P16	Sliding	16.12	16.65	-0.53
P16	Teleport	9.61	13.42	-3.81
P16	Throw	17.21	19.91	-2.70
P17	Astral	14.55	16.01	-1.46
P17	Chicken	19.21	20.30	-1.09
P17	Grab	16.88	20.36	-3.49
P17	Sliding	15.43	17.52	-2.10
P17	Teleport	13.37	13.35	0.02
P17	Throw	22.33	20.55	1.77
P18	Astral	14.82	17.30	-2.47
P18	Chicken	18.19	22.29	-4.10
P18	Grab	20.43	21.87	-1.44
P18	Sliding	19.05	18.64	0.41
P18	Teleport	17.93	13.93	3.99
P18	Throw	21.53	22.37	-0.84
P19	Astral	13.00	18.45	-5.45
P19	Chicken	17.77	23.41	-5.64
P19	Grab	22.95	22.65	0.31
P19	Sliding	16.44	20.04	-3.60
P19	Teleport	10.25	15.71	-5.46
P19	Throw	23.41	23.22	0.19
PC00	Astral	12.53	12.51	0.03
PC00	Chicken	20.60	15.97	4.63
PC00	Grab	20.46	15.77	4.69
PC00	Sliding	14.05	13.70	0.35
PC00	Teleport	11.92	10.60	1.33
PC00	Throw	17.06	16.44	0.62
PC02	Astral	11.11	12.44	-1.33
PC02	Chicken	11.61	16.04	-4.43
PC02	Grab	10.96	16.10	-5.14
PC02	Sliding	13.44	13.64	-0.20
PC02	Teleport	8.15	10.64	-2.49
PC02	Throw	14.43	16.38	-1.95
PC03	Astral	11.93	12.19	-0.25
PC03	Chicken	19.53	15.43	4.10
PC03	Grab	15.08	15.57	-0.49
PC03	Sliding	13.66	13.35	0.31
PC03	Teleport	10.29	10.28	0.01
PC03	Throw	16.87	15.89	0.97
PC04	Astral	13.08	12.52	0.57
PC04	Chicken	21.24	15.92	5.32
PC04	Grab	14.41	16.05	-1.64
PC04	Sliding	17.21	13.62	3.59
PC04	Teleport	12.40	10.48	1.92
PC04	Throw	18.89	16.38	2.51
PC05	Astral	19.20	11.77	7.43
PC05	Chicken	14.79	15.24	-0.45
PC05	Grab	13.09	15.28	-2.18
PC05	Sliding	14.86	12.96	1.89

PC05	Teleport	12.16	10.06	2.10
PC05	Throw	16.71	15.70	1.01
PC06	Astral	14.54	12.44	2.10
PC06	Chicken	14.85	16.06	-1.21
PC06	Grab	17.61	15.94	1.67
PC06	Sliding	15.21	13.58	1.63
PC06	Teleport	9.71	10.67	-0.96
PC06	Throw	15.28	16.37	-1.09
PC07	Astral	12.58	13.60	-1.02
PC07	Chicken	15.06	17.37	-2.31
PC07	Grab	16.42	17.33	-0.91
PC07	Sliding	14.13	14.82	-0.69
PC07	Teleport	10.56	11.53	-0.98
PC07	Throw	14.27	17.84	-3.57
PC08	Astral	11.01	12.23	-1.22
PC08	Chicken	16.01	15.51	0.50
PC08	Grab	9.18	15.75	-6.57
PC08	Sliding	13.84	13.35	0.49
PC08	Teleport	9.39	10.39	-1.00
PC08	Throw	13.69	16.06	-2.37
PC09	Astral	13.21	12.57	0.64
PC09	Chicken	16.50	16.06	0.44
PC09	Grab	14.19	15.93	-1.74
PC09	Sliding	14.52	13.74	0.78
PC09	Teleport	10.41	10.62	-0.21
PC09	Throw	13.45	16.53	-3.08
PC10	Astral	11.40	12.46	-1.06
PC10	Chicken	10.97	16.22	-5.25
PC10	Grab	14.52	16.01	-1.49
PC10	Sliding	14.08	13.77	0.31
PC10	Teleport	11.26	10.58	0.68
PC10	Throw	15.91	16.36	-0.45
PC11	Astral	12.79	12.18	0.61
PC11	Chicken	23.56	15.34	8.23
PC11	Grab	17.35	15.40	1.96
PC11	Sliding	15.69	13.25	2.44
PC11	Teleport	11.66	10.25	1.40
PC11	Throw	17.56	15.95	1.61
PC12	Astral	13.60	12.12	1.49
PC12	Chicken	17.06	15.48	1.58
PC12	Grab	16.09	15.56	0.54
PC12	Sliding	15.14	13.21	1.93
PC12	Teleport	11.42	10.20	1.22
PC12	Throw	13.69	16.03	-2.34
PC13	Astral	15.25	14.16	1.09
PC13	Chicken	26.17	17.88	8.30
PC13	Grab	17.30	18.47	-1.17
PC13	Sliding	15.92	15.77	0.16
PC13	Teleport	15.25	11.92	3.33
PC13	Throw	19.51	18.68	0.82
PC14	Astral	15.55	13.51	2.04
PC14	Chicken	16.44	17.62	-1.18
PC14	Grab	18.37	17.49	0.88

PC14	Sliding	19.04	14.71	4.33
PC14	Teleport	13.67	11.41	2.26
PC14	Throw	17.71	17.99	-0.28
PC15	Astral	14.82	12.13	2.69
PC15	Chicken	12.96	15.67	-2.71
PC15	Grab	14.69	15.63	-0.94
PC15	Sliding	14.77	13.30	1.47
PC15	Teleport	10.75	10.35	0.39
PC15	Throw	14.84	15.95	-1.11
PC16	Astral	13.02	13.62	-0.60
PC16	Chicken	14.63	17.58	-2.95
PC16	Grab	13.10	17.67	-4.57
PC16	Sliding	17.37	14.87	2.50
PC16	Teleport	10.54	11.64	-1.10
PC16	Throw	14.18	18.08	-3.90
PC17	Astral	15.37	12.04	3.34
PC17	Chicken	13.96	15.74	-1.78
PC17	Grab	14.97	15.43	-0.46
PC17	Sliding	16.25	13.27	2.98
PC17	Teleport	10.06	10.29	-0.23
PC17	Throw	14.83	15.99	-1.16
PC18	Astral	14.63	13.37	1.26
PC18	Chicken	17.87	17.27	0.60
PC18	Grab	15.61	17.05	-1.44
PC18	Sliding	14.74	14.69	0.05
PC18	Teleport	11.74	11.37	0.38
PC18	Throw	16.88	17.63	-0.75
PC19	Astral	11.76	12.21	-0.45
PC19	Chicken	18.86	15.41	3.45
PC19	Grab	12.88	15.64	-2.76
PC19	Sliding	13.39	13.32	0.07
PC19	Teleport	8.39	10.37	-1.98
PC19	Throw	10.53	16.16	-5.63

A.9 CALIBRATION SCENARIO: DATA TABLE SNIPPET

Table 24. A sample of the dataset used to train the model in the calibration scenario. Data are from two participants, a participant with impairments (PWI) and a participant without impairments (PWO). Column names are provided in the column name key.

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29							
PWI	Grab	1	1	12.4	14.0	17.2	17.9	8.4	16.2	2.1	1.4	2.3	0.7	0.7	0.7	695.6	101.9	1614.8	695.6	101.9	1614.8	0	0.0	0	0.0	0	0	0.3							
PWI	Grab	1	2	12.5	16.8	26.6	15.3	17.8	24.2	2.6	1.4	3.8	0.7	0.7	0.7	3016.2	101.9	2186.8	3016.2	101.9	2186.8	0	0.0	0	0.0	0	0	0.3							
PWI	Grab	1	3	14.6	18.8	31.0	31.5	15.3	29.2	2.3	1.6	3.4	0.7	0.7	0.7	2846.0	101.5	1154.8	2846.0	101.5	1154.8	0	0.0	0	0.0	0	0	0.3							
PWI	Grab	1	4	14.3	18.7	27.2	22.2	10.4	30.7	2.3	1.5	3.7	0.7	0.7	0.7	2536.0	101.0	1511.5	2536.0	101.0	1511.5	0	0.0	0	0.0	0	0	0.3							
PWI	Grab	1	5	23.7	31.5	31.5	31.6	13.5	31.5	2.2	1.4	3.4	0.7	0.7	0.7	3264.8	101.2	1826.2	3264.8	101.2	1826.2	0	0.0	0	0.0	0	0	0.4							
PWI	Grab	1	6	12.9	17.9	19.8	13.7	11.7	31.5	2.3	1.4	2.9	0.8	0.7	0.7	3363.4	86.5	2672.6	3363.4	86.5	2672.6	0	0.0	0	0.0	0	0	0.2							
PWO	Throw	2	1	31.6	8.9	8.2	9.6	6.5	10.3	1.3	0.8	2.7	0.6	0.4	0.9	1391.6	410.7	1412.2	1391.6	410.7	1412.2	0	0.6	0	0.3	0	0	0.0							
PWO	Throw	2	2	11.5	11.3	10.7	13.8	7.6	12.6	1.4	0.9	4.1	0.6	0.4	1.0	1694.4	412.1	2302.6	1694.4	412.1	2302.6	0	0.5	0	0.2	0	0	0.0							
PWO	Throw	2	3	11.6	11.7	11.2	13.5	7.8	13.3	1.4	0.9	4.5	0.6	0.4	0.9	1906.4	410.9	2896.9	1906.4	410.9	2896.9	0	0.6	0	0.2	0	0	0.0							
PWO	Throw	2	4	12.2	11.6	11.2	15.1	8.9	14.8	1.4	0.9	4.7	0.6	0.4	1.0	2767.8	412.7	2138.5	2767.8	412.7	2138.5	0	0.0	0	0.0	0	0	0.0							
PWO	Throw	2	5	18.6	13.9	13.2	16.2	10.7	16.4	1.4	0.9	4.0	0.6	0.4	1.2	1658.2	411.0	2351.9	1658.2	411.0	2351.9	0	0.8	0	0.4	0	0	0.0							
PWO	Throw	2	6	10.8	10.2	10.5	12.5	8.2	13.2	1.3	0.8	4.0	0.6	0.4	1.1	3096.8	745.8	1909.3	3096.8	745.8	1909.3	0	0.0	0	0.0	0	0	0.0							
		30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
PWI		0.1	0	0	0	0	0.0	0	0	0	0.0	0	0.0	6.0	10.0	0	154.5	42.6	124.1	0.0	1.0	0.2	0	0	0.0	0.1	0	0.1	0	0	0.0	0.2	0.0	1.5	0
PWI		0.0	0	0	0	0	0.0	0	0	0	0.0	0	0.0	0.0	15.0	0	235.9	59.0	198.5	0.0	6.0	0.3	0	0	0.0	0.1	0	0.2	0	0	0.0	0.2	0.0	1.2	0
PWI		0.0	0	0	0	0	0.0	0	0	0	0.0	0	0.0	0.0	19.0	0	282.6	70.2	240.3	0.0	2.0	1.0	0	0	0.0	0.0	0	0.1	0	0	0.0	0.1	0.0	1.2	0
PWI		0.0	0	0	0	0	0.0	0	0	0	0.0	0	0.0	0.0	17.0	0	246.0	62.2	207.1	0.0	2.0	0.2	0	0	0.0	0.1	0	0.2	0	0	0.0	0.2	0.0	1.5	0
PWI		0.0	0	0	0	0	0.0	0	0	0	0.0	0	0.0	1.0	27.0	0	290.2	66.2	251.9	1.0	2.0	1.2	0	0	0.0	0.1	0	0.2	0	0	0.0	0.2	0.0	1.7	0
PWI		0.0	0	0	0	0	0.0	0	0	0	0.0	0	0.0	0.0	13.0	0	286.3	76.3	243.7	0.0	2.0	0.3	0	0	0.0	0.1	0	0.1	0	0	0.0	0.2	0.0	1.0	0
PWO		0.1	0	0	0	0	0.0	0	0	2.0	0	3.0	2.0	0.0	0	62.5	52.9	32.9	0.0	0.0	0.1	0	0	0.1	0.0	0	0.4	0	0	0.2	0.4	0.1	3.5	0	
PWO		0.2	0	0	0	0	0.0	0	0	7.0	0	3.0	5.0	0.0	0	76.8	53.0	61.7	0.0	4.0	0.5	0	0	0.1	0.0	0	0.5	0	0	0.2	0.5	0.3	4.4	0	
PWO		0.2	0	0	0	0	0.0	0	0	6.0	0	2.0	3.0	0.0	0	79.7	54.4	61.6	0.0	0.0	0.5	0	0	0.1	0.0	0	0.6	0	0	0.2	0.4	0.1	3.0	0	
PWO		0.2	0	0	0	0	0.0	0	0	11.0	0	1.0	7.0	0.0	0	91.4	62.9	76.8	0.0	2.0	0.2	0	0	0.1	0.0	0	0.6	0	0	0.3	0.4	0.2	3.5	0	
PWO		0.1	0	0	0	0	0.0	0	0	9.0	0	3.0	5.0	0.0	0	105.3	92.7	57.5	1.0	2.0	0.2	0	0	0.1	0.0	0	0.5	0	0	0.2	0.5	0.3	3.4	0	
PWO		0.2	0	0	0	0	1.0	0	0	7.0	0	3.0	4.0	0.0	0	95.0	77.7	54.4	0.0	2.0	0.5	0	0	0.1	0.0	0	0.6	0	0	0.2	0.4	0.5	4.1	0	

	64	65	66	67	68
PWI	0	1.0	17.2	1	0
PWI	0	0.0	26.6	1	0
PWI	0	3.0	31.0	1	0
PWI	0	1.0	27.2	1	0
PWI	0	0.0	31.5	0	4
PWI	0	2.0	19.8	1	0
PWO	0	0.0	10.3	1	0
PWO	0	0.0	12.6	1	0
PWO	0	0.0	13.3	1	1
PWO	0	0.0	14.8	1	0
PWO	0	0.0	16.4	1	2
PWO	0	0.0	13.2	1	0

Column Name Key

1 iD	22 rightAngleExtent	43 rightTriggerButtonCount	64 numLeftSubmovements
2 technique	23 leftThumbstickDistance	44 rightMenuButtonCount	65 numRightSubmovements
3 trialBlock	24 rightThumbstickDistance	45 leftHeadDistance	66 trialTime
4 trial	25 leftThumbstickExtent	46 rightHeadDistance	67 hit
5 trialTimeAstral	26 rightThumbstickExtent	47 leftRightDistance	68 numObstaclesHit
6 trialTimeChicken	27 leftAvgTriggerPressure	48 targetEnteredCount	
7 trialTimeGrab	28 leftAvgGripPressure	49 axisCrossedCount	
8 trialTimeSliding	29 rightAvgTriggerPressure	50 movementVariability	
9 trialTimeTeleport	30 rightAvgGripPressure	51 headAvgVelocity	
10 trialTimeThrow	31 leftPrimaryButtonCount	52 leftAvgVelocity	
11 headTotalDistance	32 leftPrimaryButtonCount	53 rightAvgVelocity	
12 leftTotalDistance	33 leftSecondaryButtonCount	54 headAvgAcceleration	
13 rightTotalDistance	34 leftSecondaryTouchCount	55 leftAvgAcceleration	
14 headExtent	35 leftGripButtonCount	56 rightAvgAcceleration	
15 leftExtent	36 leftTriggerButtonCount	57 headAvgAngVelocity	
16 rightExtent	37 leftMenuButtonCount	58 leftAvgAngVelocity	
17 headAngleVariability	38 rightPrimaryButtonCount	59 rightAvgAngVelocity	
18 leftAngleVariability	39 rightPrimaryButtonCount	60 headAvgAngAcceleration	
19 rightAngleVariability	40 rightSecondaryButtonCount	61 leftAvgAngAcceleration	
20 headAngleExtent	41 rightSecondaryButtonCount	62 rightAvgAngAcceleration	
21 leftAngleExtent	42 rightGripButtonCount	63 numHeadSubmovements	

A.10 CALIBRATION SCENARIO: TUNING PARAMETERS

Table 25. The tuning parameters used for the final models in the calibration scenario.

Calibration Technique	Predicting Trial Time for											
	Astral (alpha, lambda)		Chicken (alpha, lambda)		Grab (alpha, lambda)		Sliding (alpha, lambda)		Teleport (alpha, lambda)		Throw (alpha, lambda)	
Astral			1.00	0.72	1.00	0.98	1.00	0.52	1.00	0.71	1.00	0.95
Chicken	0.55	0.05			0.55	0.10	0.55	0.43	0.10	0.61	0.55	0.10
Grab	0.55	0.61	1.00	0.72			1.00	0.48	0.10	0.59	0.10	0.99
Sliding	0.55	0.60	1.00	0.60	0.55	0.98			1.00	0.05	0.55	0.95
Teleport	0.10	0.68	1.00	0.70	0.10	0.98	0.55	0.41			0.55	0.95
Throw	1.00	0.62	1.00	0.72	1.00	1.01	0.10	0.43	0.10	0.73		
Five Techniques	1.00	0.06	0.55	0.63	0.34	1.01	0.10	0.41	0.10	0.59	0.55	0.09

A.11 CALIBRATION SCENARIO: ERROR

Table 26. The root mean squared error and mean absolute error of the final models in the calibration scenario.

Calibration Technique	Predicting Trial Time for													
	Astral (RMSE, MAE)		Chicken (RMSE, MAE)		Grab (RMSE, MAE)		Sliding (RMSE, MAE)		Teleport (RMSE, MAE)		Throw (RMSE, MAE)			
Astral			6.32	5.40	4.60	3.90	3.81	2.75	4.53	3.58	5.09	4.26	4.87	4.00
Chicken	4.91	3.90			4.60	3.73	4.22	3.14	4.75	3.70	5.29	4.26	4.80	3.75
Grab	4.25	3.28	6.40	5.40			4.19	3.02	5.13	4.01	4.85	3.85	5.00	3.91
Sliding	4.29	3.29	6.31	5.32	4.56	3.81			4.88	3.82	5.04	4.14	5.02	4.08
Teleport	4.72	3.58	6.47	5.53	4.86	4.01	4.13	3.07			5.03	4.13	5.04	4.06
Throw	4.43	3.45	6.32	5.41	4.82	4.10	4.06	3.05	4.34	3.44			4.79	3.89
Five Techniques	4.56	3.46	6.43	5.42	4.64	3.78	4.09	3.02	4.68	3.65	4.93	3.91	4.88	3.87

A.12 CALIBRATION SCENARIO: DATA TABLE USED FOR RANKING LOCOMOTION TECHNIQUES

Table 27. The data table used to rank locomotion techniques for each participant. Times are predicted with the calibration scenario model. Columns represent the Participant ID, Technique, the average Actual Trial Time in seconds, the average Predicted Trial Time in seconds and the Difference between the Predicted and Actual Trial Times in seconds. Participant IDs leading with “P” are participants with impairments and IDs leading with “PC” are participants without impairments.

Participant	Technique	Actual Trial Time (s)	Predicted Trial Time (s)	Difference (s)
P00	Astral	12.3	13.87	-1.58
P00	Chicken	15.08	17.07	-1.99
P00	Grab	13.67	20.13	-6.46
P00	Sliding	13.72	13.72	0
P00	Teleport	8.77	11.35	-2.58
P00	Throw	16.91	18.53	-1.62
P01	Astral	31.58	22.98	8.6
P01	Chicken	21.26	21.18	0.09
P01	Grab	30.77	26.83	3.94
P01	Sliding	22.34	22.34	0
P01	Teleport	30.22	25.07	5.15
P01	Throw	31.53	25.05	6.47
P02	Astral	13.44	17.29	-3.84
P02	Chicken	13.44	19.97	-6.53
P02	Grab	21.06	23.58	-2.52
P02	Sliding	17.52	17.52	0
P02	Teleport	12.62	11.84	0.77
P02	Throw	19.9	22.61	-2.71
P03	Astral	14.67	18.58	-3.91
P03	Chicken	19.67	21.76	-2.08
P03	Grab	24.64	25.6	-0.96
P03	Sliding	19.6	19.6	0
P03	Teleport	15.75	18.9	-3.15
P03	Throw	27.48	25.99	1.5
P04	Astral	14.46	17.3	-2.84
P04	Chicken	15.41	17.88	-2.47
P04	Grab	22.27	22.79	-0.52
P04	Sliding	14.57	14.57	0
P04	Teleport	10.35	9.51	0.84
P04	Throw	17.48	17.96	-0.48
P05	Astral	14.47	13.18	1.29
P05	Chicken	25.02	20.07	4.95
P05	Grab	27.03	24.45	2.57
P05	Sliding	16.03	16.03	0
P05	Teleport	14.22	14.01	0.21
P05	Throw	23.7	23.69	0.01
P06	Astral	12.58	13.69	-1.11
P06	Chicken	22.18	17.13	5.05
P06	Grab	22.56	20.32	2.24
P06	Sliding	14.83	14.83	0
P06	Teleport	10.03	9.97	0.07

Participant	Technique	Actual Trial Time (s)	Predicted Trial Time (s)	Difference (s)
P06	Throw	15.87	19.28	-3.41
P07	Astral	14.76	17.78	-3.01
P07	Chicken	18.35	19.91	-1.56
P07	Grab	27.57	26.85	0.72
P07	Sliding	16.8	16.8	0
P07	Teleport	14.24	12.61	1.64
P07	Throw	25.96	26.09	-0.13
P08	Astral	19.24	15.66	3.58
P08	Chicken	30.39	24.8	5.59
P08	Grab	23.25	21.58	1.67
P08	Sliding	15.7	15.7	0
P08	Teleport	25.13	20.58	4.55
P08	Throw	31.52	27.41	4.11
P09	Astral	16.21	12.4	3.81
P09	Chicken	18.39	18.93	-0.54
P09	Grab	22.69	23.79	-1.11
P09	Sliding	16.69	16.69	0
P09	Teleport	12.62	11.65	0.97
P09	Throw	23.71	20.25	3.46
P10	Astral	25.84	23.07	2.77
P10	Chicken	22.16	24.11	-1.95
P10	Grab	28.43	26.11	2.32
P10	Sliding	19.32	19.32	0
P10	Teleport	10.71	17.65	-6.94
P10	Throw	29.23	28.46	0.77
P11	Astral	11.76	13.55	-1.79
P11	Chicken	25.81	22.12	3.69
P11	Grab	29.53	25.05	4.48
P11	Sliding	17.26	17.26	0
P11	Teleport	10.66	9.53	1.13
P11	Throw	26.02	26.39	-0.37
P12	Astral	17.03	15.19	1.84
P12	Chicken	23.95	20.63	3.33
P12	Grab	30.13	28.36	1.77
P12	Sliding	19.69	19.69	0
P12	Teleport	13.92	12.08	1.84
P12	Throw	31.52	28.67	2.86
P13	Astral	30.19	22.74	7.45
P13	Chicken	31.55	25.57	5.98
P13	Grab	31.65	28.54	3.11
P13	Sliding	16.8	16.8	0
P13	Teleport	31.56	22.72	8.85
P13	Throw	31.56	28.81	2.75
P14	Astral	14.94	15.04	-0.1
P14	Chicken	24.21	21.11	3.11
P14	Grab	26.12	27.39	-1.26
P14	Sliding	19.93	19.93	0
P14	Teleport	13.84	10.71	3.13
P14	Throw	28.16	27.14	1.02
P15	Astral	14.76	13.33	1.43
P15	Chicken	20.52	20.87	-0.35

Participant	Technique	Actual Trial Time (s)	Predicted Trial Time (s)	Difference (s)
P15	Grab	27.36	23.17	4.19
P15	Sliding	16.44	16.44	0
P15	Teleport	15.2	16.03	-0.82
P15	Throw	22.93	22.29	0.64
P16	Astral	13.11	13.15	-0.04
P16	Chicken	18.69	19.9	-1.21
P16	Grab	18.18	20.89	-2.71
P16	Sliding	16.12	16.12	0
P16	Teleport	9.61	12.88	-3.28
P16	Throw	17.21	20.27	-3.06
P17	Astral	14.55	11.49	3.06
P17	Chicken	19.21	18.35	0.86
P17	Grab	16.88	21.16	-4.28
P17	Sliding	15.43	15.43	0
P17	Teleport	13.37	14.1	-0.72
P17	Throw	22.33	19.87	2.46
P18	Astral	14.82	16.43	-1.6
P18	Chicken	18.19	20.72	-2.53
P18	Grab	20.43	20.58	-0.15
P18	Sliding	19.05	19.05	0
P18	Teleport	17.93	19.53	-1.6
P18	Throw	21.53	23.87	-2.34
P19	Astral	13	12.51	0.49
P19	Chicken	17.77	21.13	-3.37
P19	Grab	22.95	26.43	-3.47
P19	Sliding	16.44	16.44	0
P19	Teleport	10.25	14.57	-4.31
P19	Throw	23.41	26.02	-2.61
PC00	Astral	12.53	13.69	-1.15
PC00	Chicken	20.6	17.12	3.48
PC00	Grab	20.46	18.48	1.97
PC00	Sliding	14.05	14.05	0
PC00	Teleport	11.92	10.63	1.3
PC00	Throw	17.06	16.58	0.48
PC02	Astral	16.85	13.3	3.55
PC02	Chicken	11.61	17.92	-6.31
PC02	Grab	10.96	15.05	-4.09
PC02	Sliding	13.44	13.44	0
PC02	Teleport	8.15	10.86	-2.72
PC02	Throw	14.43	16.9	-2.47
PC03	Astral	11.93	13.47	-1.54
PC03	Chicken	19.53	17.03	2.5
PC03	Grab	15.08	13.86	1.22
PC03	Sliding	13.66	13.66	0
PC03	Teleport	10.29	9.68	0.61
PC03	Throw	16.87	16.92	-0.05
PC04	Astral	13.08	14.75	-1.67
PC04	Chicken	21.24	18.02	3.22
PC04	Grab	14.41	16.38	-1.97
PC04	Sliding	17.21	17.21	0
PC04	Teleport	12.4	11.44	0.96

Participant	Technique	Actual Trial Time (s)	Predicted Trial Time (s)	Difference (s)
PC04	Throw	18.89	17.93	0.95
PC05	Astral	13.08	13.83	-0.74
PC05	Chicken	14.79	16.85	-2.06
PC05	Grab	13.09	14	-0.91
PC05	Sliding	14.86	14.86	0
PC05	Teleport	12.16	11.78	0.38
PC05	Throw	16.71	15.48	1.22
PC06	Astral	14.54	14.23	0.31
PC06	Chicken	14.85	17.18	-2.33
PC06	Grab	17.61	16.07	1.54
PC06	Sliding	15.21	15.21	0
PC06	Teleport	9.71	10.85	-1.14
PC06	Throw	15.28	16.09	-0.81
PC07	Astral	12.58	14.1	-1.51
PC07	Chicken	15.06	17.41	-2.35
PC07	Grab	16.42	16.22	0.2
PC07	Sliding	14.13	14.13	0
PC07	Teleport	10.56	11.96	-1.4
PC07	Throw	14.27	16.84	-2.57
PC08	Astral	11.01	12.15	-1.14
PC08	Chicken	16.01	17.34	-1.32
PC08	Grab	9.18	12.59	-3.41
PC08	Sliding	13.84	13.84	0
PC08	Teleport	9.39	9.9	-0.51
PC08	Throw	13.69	13.7	-0.01
PC09	Astral	13.21	13.07	0.14
PC09	Chicken	16.5	17.66	-1.16
PC09	Grab	14.19	14.67	-0.48
PC09	Sliding	14.52	14.52	0
PC09	Teleport	10.41	11.19	-0.78
PC09	Throw	13.45	13.84	-0.4
PC10	Astral	11.4	13.59	-2.19
PC10	Chicken	10.97	17.14	-6.17
PC10	Grab	14.52	15.24	-0.72
PC10	Sliding	14.08	14.08	0
PC10	Teleport	11.26	10.67	0.58
PC10	Throw	16.21	16.67	-0.46
PC11	Astral	12.7	14.14	-1.44
PC11	Chicken	25.25	18.54	6.71
PC11	Grab	15.79	14.49	1.3
PC11	Sliding	15.7	15.7	0
PC11	Teleport	11.53	11.93	-0.4
PC11	Throw	19.24	16.27	2.97
PC12	Astral	11.88	13.82	-1.94
PC12	Chicken	17.16	17.25	-0.09
PC12	Grab	16.52	14.28	2.24
PC12	Sliding	14.86	14.86	0
PC12	Teleport	11.03	10.26	0.77
PC12	Throw	14.09	15.72	-1.63
PC13	Astral	15.23	14.4	0.83
PC13	Chicken	27.02	18.97	8.06

Participant	Technique	Actual Trial Time (s)	Predicted Trial Time (s)	Difference (s)
PC13	Grab	16.76	15.69	1.07
PC13	Sliding	15.6	15.6	0
PC13	Teleport	15.01	11.34	3.67
PC13	Throw	19.95	15.72	4.23
PC14	Astral	15.19	15.15	0.04
PC14	Chicken	16.43	18.63	-2.2
PC14	Grab	18.56	16.68	1.88
PC14	Sliding	19.06	19.06	0
PC14	Teleport	13.31	12.3	1.02
PC14	Throw	18.04	15.87	2.16
PC15	Astral	14.76	13.68	1.08
PC15	Chicken	12.84	17.69	-4.84
PC15	Grab	14.77	13.95	0.82
PC15	Sliding	14.85	14.85	0
PC15	Teleport	10.34	11.02	-0.68
PC15	Throw	15.2	15.38	-0.18
PC16	Astral	12.83	15.43	-2.59
PC16	Chicken	14.74	19.57	-4.83
PC16	Grab	13.15	15.62	-2.47
PC16	Sliding	17.45	17.45	0
PC16	Teleport	10.3	13.78	-3.47
PC16	Throw	14.45	18.58	-4.13
PC17	Astral	15.17	13.76	1.41
PC17	Chicken	14.02	17.84	-3.82
PC17	Grab	15.24	14.32	0.93
PC17	Sliding	16.09	16.09	0
PC17	Teleport	9.69	11.33	-1.64
PC17	Throw	15.08	16.46	-1.38
PC18	Astral	14.54	14.52	0.02
PC18	Chicken	18.56	17.24	1.33
PC18	Grab	15.06	15.07	-0.02
PC18	Sliding	15.02	15.02	0
PC18	Teleport	11.18	11.52	-0.34
PC18	Throw	17.49	19.41	-1.92
PC19	Astral	11.27	13.47	-2.2
PC19	Chicken	20.53	16.48	4.05
PC19	Grab	11.22	14.06	-2.85
PC19	Sliding	13.56	13.56	0
PC19	Teleport	8.01	10.55	-2.54
PC19	Throw	10.65	15.89	-5.24