A Testbed for Exploring Multi-Level Precueing in Augmented Reality

Jen-Shuo Liu*
Department of Computer Science
Columbia University

Barbara Tversky†
Department of Human Development
Teachers College
Columbia University

Steven Feiner‡
Department of Computer Science
Columbia University

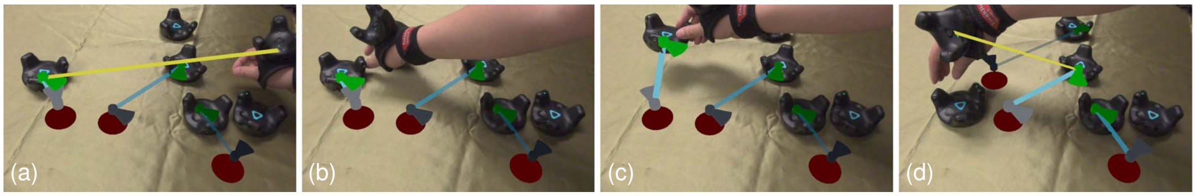

Figure 1: Cue and precue visualizations for a manual object movement and orientation task in AR. Successive subtasks each involve picking up a designated object, moving and rotating it, and placing at a designated target location. The yellow line connects the Vive tracker on the user's dominant hand and the current object (also a Vive tracker) to be moved. Each movement precue is a cyan line connecting the object to its destination target (shown as a red circle). Each rotation precue is a pair of wedges: a green wedge shows the current orientation of the object, while a gray wedge shows the target orientation. Successive phases in a subtask: (a) The user moves their unloaded hand to the object to be manipulated, guided by the yellow line. (b) The user is about to pick up the object and the yellow line has disappeared. (c) The user moves and rotates the object in their loaded hand, transporting it to its destination. (d) The user deposits the object at its destination with the specified orientation, their unloaded hand is now connected to the next object to be manipulated, and the cues and precues are updated.

Abstract

Precueing information about upcoming subtasks prior to performing them has the potential to make an entire task faster and easier to accomplish than cueing only the current subtask. Most AR and VR research on precueing has addressed path-following tasks requiring simple actions at a series of locations, such as pushing a button or just visiting that location. We present a testbed for exploring multi-level precueing in a richer task that requires the user to move their hand between specified locations, transporting an object between some of them, and rotating it to a designated orientation.

Index Terms: Human-centered computing—Human computer interaction (HCI)—Interaction paradigms—Mixed / augmented reality; Computing methodologies—Computer graphics—Graphics systems and interfaces—Perception

1 Introduction

Many maintenance and assembly tasks involve a series of movement and rotation actions, such as putting a nut on a bolt and then using a driver to fasten it, or moving a correctly oriented part into the right slot in a chassis to assemble a computer. With sufficient experience, experts can plan and prepare a future action while performing the current one, increasing efficiency of performance. Precueing may be especially helpful for users who are not familiar with the task because the alternative would be checking instructions for each step of the procedure, a disruptive and time-consuming practice that can lead to error.

To increase efficiency and accuracy, AR and VR can be used to present *cues* for the current step and *precues* for future steps. Previous work on precueing [3, 5, 9] used relatively simple path-following tasks. Yet many tasks involve more than paths; they also require actions at each step of the path. There are numerous action possibilities; users may be required to attach objects in a variety of ways or detach or change settings, increasing the complexity of possible cues and precues for kinds of actions. To address these complexities, we have developed a testbed to study the efficacy of cues and precues for specifying actions as well as paths. We designed a compound task that contains a series of subtasks. In each subtask, the user must move their hand to an object, pick up the object, transport it rotated by a specified angle to a destination target, and deposit it at the destination.

*e-mail: jl5004@columbia.edu
†e-mail: bt2158@tc.columbia.edu
‡e-mail: feiner@cs.columbia.edu

2 Related Work

Hertzum and Hornbæk [3] investigated single-level precueing (showing information about the first upcoming step in addition to the current step) in touchpad/mouse input for a desktop 2D study. They found their participants spent less time moving to a precued target than to non-precued ones. Volmer et al. [9] studied the effect of using single-level precueing in a path-following task performed in a projector-based AR environment. Their study showed that using a single precue can improve the user's performance. More recently, Liu et al. [5] investigated the effect of using multi-level precueing (showing information about multiple upcoming steps in addition to the current step) in a VR path-following task. They showed that, for the visualization styles they studied, participants performed the best when given given two to three precues for styles that connect targets with lines.

3 Task

The task we explore in our testbed consists of a series of subtasks. In each subtask, the user needs to move their hand to a designated object, pick it up, move the object to a specified destination while rotating it by a specified angle, and put it down at the destination. Figure 1 shows an example of one of these subtasks. The objects are placed on a horizontal tabletop. The user needs to move the objects between different locations on the tabletop and rotate each about an axis perpendicular to the tabletop. Based on whether the user's hand is loaded (full) or unloaded (empty), as defined by Gilbreth and Gilbreth [2], each subtask can be divided into the following phases:

Unloaded hand: As shown in Figure 1(a–b), at the beginning of each subtask, the user's hand is empty (unloaded). The user needs

978-1-6654-8402-2/22/$31.00 ©2022 IEEE
DOI 10.1109/VRW55335.2022.00121

2022 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW) | 978-1-6654-8402-2/22/$31.00 ©2022 IEEE | DOI: 10.1109/VRW55335.2022.00121

to move their unloaded hand to the designated object and pick it up. At this point, the subtask enters the loaded-hand phase.

Loaded hand: As shown in Figure 1(c), after the user picks up the object, they must move it to the specified destination, accomplishing the visualized rotation, and put down the object. The current subtask must be completed before proceeding to the next subtask. While movement and rotation of the object can be performed simultaneously, information about where to move and how much to rotate is shown using separate visualizations, as described in Sections 4.1 and 4.2. As shown in Figure 1(d), when the user places the object at the specified destination with the specified orientation, the subtask is complete and the unloaded-hand phase of the next subtask begins.

4 Cueing and Precueing

Our testbed supports cues and precues for movement and cues and precues for rotation.

Cues: Cues guide the user to perform the current subtask by picking up one of the objects, moving it to its destination while rotating it by a specified angle, and putting it down. As the user finishes the kth subtask, the cue visualizations update to guide the user to perform the $k+1$st subtask.

Precues: For a condition with m precued subtasks, when the user is working on subtask k, the precued subtasks are subtasks $k+1$, $k+2$, $k+3$, ..., $k+m$.

Once the user completes the current subtask k, the cued subtask becomes subtask $k+1$ and the precued subtasks become $k+2$, $k+3$, ..., $k+m+1$. This continues until the user finishes the entire task. We use brightness, transparency, and size differences to help users differentiate the cues and precues.

4.1 Movement Cue and Precues

We decided to use lines as movement cues and precues, based on prior work by Volmer et al. [9] and Liu et al. [5] that shows that using a line to connect an origin and destination in a movement cue or precue will yield better performance than visualizations that do not use a line.

A movement cue shows the movement needed to complete the current subtask. A movement cue consists of two lines: One line connects the user's empty hand to the object to be picked up, while another line, which is colored differently, connects the object to its destination. Once the user's hand is sufficiently close to the object, the line connecting their hand to the object will disappear.

Movement precues show information about future subtasks. Our movement precues each have only one line: the one that connects the object to its destination. We excluded the one connecting the user's empty hand to the object based on our pilot studies, which did not show these lines as providing additional benefit. In Figure 1(a), the movement cue consists of the yellow line connecting the user's hand to an object and the cyan line connecting that object to its destination target. The two dimmer cyan lines are the first and second movement precues.

4.2 Rotation Cue and Precues

Our 1DoF rotation cue shows the amount of rotation required in the subtask. It consists of two wedges: a green wedge shows the object's current orientation, while a gray wedge shows the target orientation. We developed the wedge visualization through pilot studies. Each wedge is large enough to be easily observed, even when the user is doing a movement task, and its shape is different enough from the movement visualization to not cause confusion. Figure 1 shows an example of how the rotation cue and precues update during a subtask and after it completes. One wedge of the rotation cue shown in Figure 1 (b–c) is attached to the object being moved, similar to how Sodhi et al. [6] project a movement cue on their user's hand.

5 Implementation

We implemented our testbed using Unity 2019.4.18f1 [7], and the open-source Mercury Messaging framework for Unity [1]. We use a Varjo XR-3 video–see-through AR headset [8] with a 115° horizontal (134° diagonal) field of view and a 90Hz refresh rate. The XR-3 is tracked with four HTC SteamVR 2.0 Base Stations. We use a Vive Tracker 3.0 for hand tracking and five Vive Tracker 2.0 units as the tracked objects for participants to move, all tracked by the SteamVR Base Stations. The headset runs on a computer powered by an Intel® Core™ i9-11900K Processor and an Nvidia GeForce RTX 3090 graphics card.

6 Conclusions and Future Work

We have developed an AR testbed that supplements graphical cues with graphical precues to provide information about future subtasks for a heterogeneous manual task. We plan to test and compare different configurations of the movement and rotation precues to determine how well they help participants perform this task, and how many cues and precues participants can use effectively. We are also interested in extending our testbed to address additional DoF, bimanual interaction, and more complex tasks, as well as to include feedback about accuracy that could help users improve their performance in some tasks (e.g., Jeanne et al. [4]).

Acknowledgments

This research was funded in part by National Science Foundation Grant CMMI-2037101. We thank Portia Wang for her assistance in making the video.

References

[1] C. Elvezio, M. Sukan, and S. Feiner. Mercury: A messaging framework for modular UI components. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, pp. 588:1–588:12. ACM, New York, NY, USA, 2018. doi: 10.1145/3173574.3174162

[2] F. B. Gilbreth and L. M. Gilbreth. *Applied motion study: A collection of papers on the efficient method to industrial preparedness*. Macmillan, 1919.

[3] M. Hertzum and K. Hornbæk. The effect of target precuing on pointing with mouse and touchpad. *International Journal of Human–Computer Interaction*, 29(5):338–350, 2013. doi: 10.1080/10447318.2012.711704

[4] F. Jeanne, I. Thouvenin, and A. Lenglet. A study on improving performance in gesture training through visual guidance based on learners' errors. In *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology*, VRST '17. Association for Computing Machinery, New York, NY, USA, 2017. doi: 10.1145/3139131.3139144

[5] J.-S. Liu, C. Elvezio, B. Tversky, and S. Feiner. Using multi-level precueing to improve performance in path-following tasks in virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 27(11):4311–4320, 2021. doi: 10.1109/TVCG.2021.3106476

[6] R. Sodhi, H. Benko, and A. Wilson. LightGuide: Projected visualizations for hand movement guidance. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, p. 179–188. Association for Computing Machinery, New York, NY, USA, 2012. doi: 10.1145/2207676.2207702

[7] Unity. *Unity*. `https://unity.com/`, last accessed on 12/31/2021.

[8] Varjo. *Varjo XR-3*. `https://varjo.com/products/xr-3/`, last accessed on 12/31/2021.

[9] B. Volmer, J. Baumeister, S. Von Itzstein, I. Bornkessel-Schlesewsky, M. Schlesewsky, M. Billinghurst, and B. H. Thomas. A comparison of predictive spatial augmented reality cues for procedural tasks. *IEEE Transactions on Visualization and Computer Graphics*, 24(11):2846–2856, 2018. doi: 10.1109/TVCG.2018.2868587