A Testbed for Exploring Virtual Reality User Interfaces for Assigning Tasks to Agents at Multiple Sites

Jen-Shuo Liu
Columbia University
New York, NY, USA
jl5004@columbia.edu

Chongyang Wang
Columbia Teachers College
New York, NY, USA
gw2437@tc.columbia.edu

Barbara Tversky
Columbia Teachers College
New York, NY, USA
bt2158@tc.columbia.edu

Steven Feiner
Columbia University
New York, NY, USA
feiner@cs.columbia.edu

Figure 1: Spatial layouts of secondary sites. In each layout, a selected secondary site (highlighted in white) has become the primary site, located at the center, at which the user assigns tasks. (a) 2×4, flat. (b) 2×4, curved. (c) 1×8, flat. (d) 1×8, curved.

ABSTRACT

In virtual reality (VR) teleoperation and remote task guidance, a remote user may need to assign tasks to local technicians or robots at multiple sites. We are interested in scenarios where the user works with one site at a time, but must maintain awareness of the other sites for future intervention. We present an instrumented VR testbed for exploring how different spatial layouts of site representations impact user performance. In addition, we investigate ways of supporting the remote user in handling errors and interruptions from sites other than the one with which they are currently working, and switching between sites. We conducted a pilot study and explored how these factors affect user performance.

ACM Reference Format:

Jen-Shuo Liu, Chongyang Wang, Barbara Tversky, and Steven Feiner. 2023. A Testbed for Exploring Virtual Reality User Interfaces for Assigning Tasks to Agents at Multiple Sites. In *The 2023 ACM Symposium on Spatial User Interaction (SUI '23), October 13–15, 2023, Sydney, NSW, Australia.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3607822.3618004

1 INTRODUCTION

Virtual reality (VR) and augmented reality (AR) have been used in remote task guidance [1, 3, 6, 9, 10] and teleoperation [4, 5] to allow a remote user to guide or assign tasks to agents (e.g, technicians or robots) at physical job sites. While the remote user might collaborate with just a single site, it could be more efficient if they work with multiple sites, even if they do so one at a time, especially since physical presence isn't required. After the user addresses the needs of one site, they could turn their attention to another site.

Previous work by Otsuki et al. [10] presented a desktop AR user interface for a remote expert whose screen includes a camera view of each of up to four technicians that they advise, one at a time. They explored how the expert's performance is affected by the number of technicians they advise. Related research has investigated simultaneously visualizing multiple datasets in VR and AR. Liu et al. [7] examined how the layout of many equal-sized visualizations in VR can affect user performance. Unlike this previous data-visualization work, we are interested in users who primarily operate on one data object (a factory site in our case) at a time, while also needing to maintain peripheral awareness of and ability to switch to and from others. We are developing a VR testbed for exploring tradeoffs between ways to lay out these multiple sites (Figure 1) and, at the same time, to support switching between them, including when errors and interruptions arise (Figure 2).

2 TESTBED

Our testbed is implemented in Unity 2019.4.12f1 and includes a primary site, where the remote user assigns tasks by manipulating virtual replicas of physical task objects, and multiple secondary sites (Figure 1). All sites in our testbed rely on simulations to support user studies in which the remote user is the sole participant.

The participant first selects a secondary site that will become the primary site, causing its secondary site representation to be highlighted in white. Next, they manipulate representations of objects in the primary site to specify tasks to be performed. Our testbed indicates the necessary manipulations, to avoid the need for domain expertise. After the participant has assigned tasks for that site, they can select a different site at which to work, while the agents at the previous site accomplish their tasks.

While proposing tasks for the primary site, the participant must also monitor the secondary sites (e.g., to recognize errors or other requests for assistance). Our testbed highlights in red a site with errors (Figure 2a). When errors appear at a secondary site, the participant must select it, causing it to become the primary site. The participant must then assign a set of specified tasks to correct the errors before returning to the previous primary site.

We are using our testbed to explore how to help the participant understand and recover from these interruptions, which can reduce

SUI '23, October 13–15, 2023, Sydney, NSW, Australia

ACM ISBN 979-8-4007-0281-5/23/10.
https://doi.org/10.1145/3607822.3618004

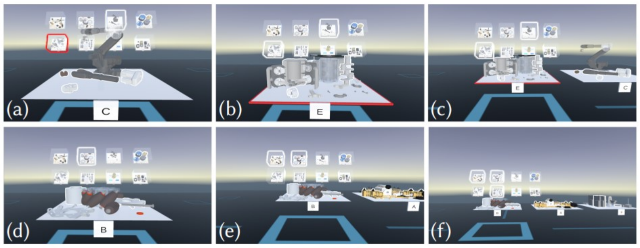

Figure 2: Error handling and upcoming site queue. (a) Error appears in a secondary site, indicated by red highlight. (b–c) User selects site with errors, and it becomes the primary site. Previous site is (b) hidden or (c) pushed onto a stack at the right. (d–f) Upcoming queue with size (d) 0, (e) 1, and (f) 2.

user performance [2]. Approaches we are studying include hiding the site that was replaced (Figure 2b), pushing it onto a stack (Figure 2c), or placing it in some other temporary holding area from which it can be easily retrieved.

If a remote site finishes its tasks, indicated by a green highlight (Figure 1a), the user will need to check and confirm that the tasks succeeded. To enable this, the user can select the site, adding it to an upcoming site queue so that when the user finishes working with the current primary site, a queued site can easily become the next primary site. Depending on the condition, the upcoming site queue will show a maximum of 0, 1, or 2 upcoming sites selected by the user (Figures 2d–f). We are also testing different secondary site layouts (Figure 1), allowing us to compare our results with those of Liu et al. [7].

3 PILOT STUDY

In preparation for a formal study, we ran a small pilot study, approved by the review board at our institution. Participants wore a Varjo XR-3 video–see-through AR headset driven by a computer with an Intel® Core™ i9-11900K processor and an Nvidia GeForce RTX 3090 graphics card and tracked using four HTC SteamVR Base Station 2.0 units. A Vive hand-held controller was used to manipulate objects.

We recruited five right-handed participants (two female) 20–33 years old (average 24.6), all with normal stereo vision and color vision. Each participant received a USD 15 gift card. We tested the four secondary-site layouts shown in Figure 1, the mechanisms for recovering from interruptions (Figures 2b and 2c), and the three upcoming queue sizes, shown in Figures 2 (d–f). Each condition encompassed two trials, each featuring an eight-step sequence with two steps containing errors. We counterbalanced the order of conditions and used trial completion time to evaluate user performance.

3.1 Observations

Our pilot study revealed strong learning effects in the first four to five trials for the participants. To address this, we plan to add more practice trials in our formal study. After excluding trials with prominent learning effects, we observed several trends: Users spent less time finishing trials with errors when the previous site was pushed onto a stack to help recover from the interruption (73.82s) than when it was hidden (68.89s). For upcoming site queue size, results suggest similar performance for queue sizes 0 (72.44s) and 1 (71.80s), but poorer outcomes for queue size 2 (75.12s). Note that the p-values for pairwise t-tests were not significant due to the low number of data points.

For secondary site layouts, we did not see a trend toward different performance between flat and curved layouts, unlike Liu et al. [7]. This could be because we placed secondary sites in the user's field of view, aided by highlights for locating upcoming sites or errors.

4 CONCLUSIONS AND FUTURE WORK

Our testbed is allowing us to explore strategies for presenting and interacting with tasks assigned to sites by a remote user. In an initial small pilot study, we did not see a trend related to secondary site layout. When addressing system-indicated errors, pushing the previous site onto a stack, rather than hiding it, tended to decrease completion time. Further, there was a trend that an appropriately sized upcoming site queue led to a shorter task completion time. To reduce the learning effects we found, our upcoming formal study will include additional practice trials. We are also developing adaptive secondary site layouts, adjusted for user eye gaze, with the goal of significantly improving user performance.

Currently, we present site status designed to support a particular order in which sites should be visited, similar to the approach used by Liu et al. [8]. This is to avoid the confound caused by different participants' preferences for the next site to select. To accommodate user preferences, we plan to conduct studies in which the participant can decide the order in which the tasks are assigned. We are also interested in mitigating interruptions in multi-user systems similar to Otsuki et al. [10].

ACKNOWLEDGMENTS

This research was funded in part by National Science Foundation Grant CMMI-2037101.

REFERENCES

[1] M. Adcock, D. Ranatunga, R. Smith, and B. H. Thomas. 2014. Object-Based Touch Manipulation for Remote Guidance of Physical Tasks. In *Proc. SUI 2014*. 113–122.
[2] B. P. Bailey and S. T. Iqbal. 2008. Understanding Changes in Mental Workload during Execution of Goal-Directed Tasks and Its Application for Interruption Management. *ACM Trans. Comput.-Hum. Interact.* 14, 4, Article 21 (2008).
[3] T. Duval, T. T. H. Nguyen, C. Fleury, A. Chauffaut, G. Dumont, and V. Gouranton. 2014. Improving awareness for 3D virtual collaboration by embedding the features of users' physical environments and by augmenting interaction tools with cognitive feedback cues. *J. Multimodal User Interfaces* 8 (2014), 187–197.
[4] I. Farkhatdinov and J.-H. Ryu. 2008. Teleoperation of multi-robot and multi-property systems. In *Proc. INDIN 2008*. 1453–1458.
[5] A. Kheddar, C. Tzafestas, P. Coiffet, T. Kotoku, and K. Tanie. 1996. Multi-robot teleoperation using direct human hand actions. *Adv. Robot.* 11, 8 (1996), 799–825.
[6] S. Kim, G. Lee, W. Huang, H. Kim, W. Woo, and M. Billinghurst. 2019. Evaluating the Combination of Visual Communication Cues for HMD-Based Mixed Reality Remote Collaboration. In *Proc. CHI 2019*. 1–13.
[7] J. Liu, A. Prouzeau, B. Ens, and T. Dwyer. 2020. Design and Evaluation of Interactive Small Multiples Data Visualisation in Immersive Spaces. In *Proc. IEEE VR 2020*. 588–597.
[8] J.-S. Liu, B. Tversky, and S. Feiner. 2022. Precueing Object Placement and Orientation for Manual Tasks in Augmented Reality. *IEEE TVCG* 28, 11 (2022), 3799–3809.
[9] O. Oda, C. Elvezio, M. Sukan, S. Feiner, and B. Tversky. 2015. Virtual Replicas for Remote Assistance in Virtual and Augmented Reality. In *Proc. UIST 2015*. 405–415.
[10] M. Otsuki, T.-Y. Wang, and H. Kuzuoka. 2022. Assessment of Instructor's Capacity in One-to-Many AR Remote Instruction Giving. In *Proc. VRST 2022*. Article 8.