Precueing Sequential Rotation Tasks in Augmented Reality

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

Barbara Tversky
Department of Human Development
Teachers College
Columbia University
New York, NY, USA
bt2158@tc.columbia.edu

Steven Feiner
Department of Computer Science
Columbia University
New York, NY, USA
feiner@cs.columbia.edu

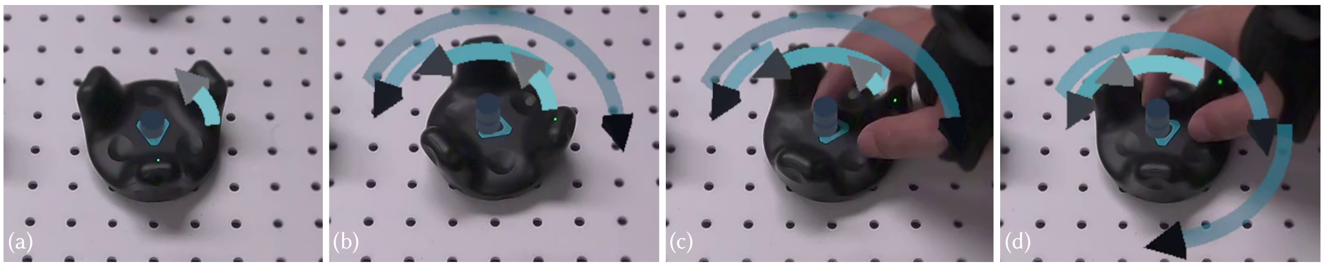

Figure 1: Visualizations guiding a user to rotate a task object (Vive Tracker). (a) Circular arrow cue for current step. The base of the pyramidal arrowhead represents the direction and magnitude of the rotation, while the arrow tail indicates the point on the object to rotate to the arrowhead base. (b) A cue and three precues for future steps. Arrow bodies are centered about the axis of rotation. Later precues are defined by larger circles, on lower planes, and are darker and more transparent. (c) As the user (whose hand is tracked by a wrist-worn tracker) rotates the object, the arrow tail rotates accordingly. (d) When the tail reaches the arrowhead base, the step is complete: The first precue becomes the cue, it gets brighter and more opaque, its radius becomes smaller, and its plane becomes higher; the other precues update accordingly, and a precue for the next step is added.

ABSTRACT

Augmented reality has been used to improve sequential-task performance by *cueing* information about a current task step and *precueing* information about future steps. Existing work has shown the benefits of precueing movement (translation) information. However, rotation is also a major component in many real-life tasks, such as turning knobs to adjust parameters on a console. We developed an AR testbed to investigate whether and how much precued rotation information can improve user performance. We consider two unimanual tasks: one requires a user to make sequential rotations of a single object, and the other requires the user to move their hand between multiple objects to rotate them in sequence.

We conducted a user study to explore these two tasks using circular arrows to communicate rotation. In the single-object task, we examined the impact of number of precues and visualization style on user performance. Results show that precues improved performance and that arrows with highlighted heads and tails, with each destination aligned with the next origin, yielded the shortest completion time on average. In the multiple-object task, we explored whether rotation precues can be helpful in conjunction with movement precues. Here, using a rotation cue without rotation precues in conjunction with a movement cue and movement precues performed the best, implying that rotation precues were not helpful when movement was also required.

CCS CONCEPTS

• **Human-centered computing → Mixed / augmented reality**; **User studies**; **Empirical studies in HCI**; • **Computing methodologies → Perception**.

KEYWORDS

precueing, cueing, object rotation

ACM Reference Format:

Jen-Shuo Liu, Barbara Tversky, and Steven Feiner. 2022. Precueing Sequential Rotation Tasks in Augmented Reality. In *28th ACM Symposium on Virtual Reality Software and Technology (VRST '22), November 29-December 1, 2022, Tsukuba, Japan.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3562939.3565641

1 INTRODUCTION

Many real-life tasks involve rotating objects. For example, a power plant operator adjusts the values of parameters by rotating knobs on a control console. An optical technician rotates components on an optical bench to proper angles. The knobs on a mixer allow a DJ to adjust values such as volume, equalization, and tempo. In these tasks, users often need to perform these rotations one by one in a particular order. This may be difficult for novices who are unfamiliar with the required order, directions, and amounts.

VRST '22, November 29-December 1, 2022, Tsukuba, Japan

ACM ISBN 978-1-4503-9889-3/22/11…$15.00
https://doi.org/10.1145/3562939.3565641

These users might benefit from *cueing* rotation information for the current step and *precueing* rotation information for future steps.

Previous work has shown that cues and precues help improve user performance in tasks that involve only movement (translation) [23, 28, 31, 49, 50]. In contrast, we analyze the utility of cues and precues in tasks that require rotations, performed in sequence, with or without interleaved translations. We consider two unimanual tasks: one requires a user to perform sequential rotations of a single object (Figure 1), and the other requires the user to move their hand between multiple objects to rotate them in sequence. We developed an augmented reality (AR) testbed for exploring the benefits of rotation precues in these sequential rotation tasks and conducted a user study through which we make the following contributions:

- We investigate in the single-object task whether users can benefit from rotation precues, and how the number of rotation precues impacts task completion time.
- We study in the single-object task how the style of rotation precues impacts task completion time.
- We study in the multiple-object task whether using rotation precues interleaved with hand movement precues helps users perform faster.

2 RELATED WORK

2.1 Systems That Provide Rotation Information

Some systems provide visual feedback for rotation, but do not use cues or precues. Berthaut and Jones [12] use AR to overlay current values on the knobs of audio consoles. Similarly, Trujano et al. [44] draw AR arcs to show the values of parameters controlled by a knob on an electronic keyboard. To visualize interactions in a simulated holographic workbench, Underkoffler and Ishii [47] developed a system in which users rotate (and translate) tracked objects representing optical components. This causes light paths annotated with lengths and angles to be projected on the workbench. Suzuki et al. [43] extend this concept by supporting end-user design of the augmented environment. Their users sketch annotations on tracked objects, selectively indicating lines, angles, and parameters that are dynamically visualized in AR.

In contrast, other systems guide users to perform specific tasks by providing a cue with or without precues. In gaming, *Joy-Con Rotation* [4] uses a numeric value to cue the angle by which each player should rotate their controller. In *Rotaeno* [7] and *Chrono Circle* [2], a player needs to rotate and tap a gaming device in response to moving lines and circles that cue and precue required actions. The player must act with the right rhythm rather than as fast as possible. Although some of these systems use rotation precues, we know of no analyses of their benefits.

Researchers have long used circular arrows to cue rotation in AR (e.g., [17, 20, 48]). Liu et al. [27] designed a mobile-phone–based AR system that overlays instructions on physical control panels. Circular arrows indicate the amounts knobs are to be rotated. Their study shows that giving a user feedback with AR by updating the visualization as they perform the task can reduce task completion time. Schönauer et al. [40] developed a system for guiding motion and rotation through haptics. They place vibrotactors on a user's arms and use vibration to cue the user to rotate their arms. Oda et al. [35] use AR to cue participants to move and rotate physical equipment parts in assembly tasks. A virtual replica of a part shows its 6DoF goal pose in place and virtual "rubberband" lines connect locations on the physical part to corresponding locations on the virtual replica to guide the user. Sukan et al. [42] designed alternative visual cues to guide freehand 3DoF rotations and compared their effectiveness. These projects developed and analyzed rotation cues, but do not support precueing. In contrast, we address precueing rotations in sequential tasks, for a single object and for multiple objects.

2.2 Studies of Precueing

Hertzum and Hornbæk [23] studied the effect of using one precue for a movement task with desktop 2D touchpad/mouse input. They showed that participants moved faster to a precued target than to targets that were not precued. Volmer et al. [50] used one movement precue in a projector-based spatial AR system. They found that users performed faster with one precue than with no precue, and that a line precue connecting the current and next destinations improved performance the most. Their follow-on work for the same task [49] showed that, even when sleep-deprived, users still benefited from one movement precue and performed the best with a line precue. Liu et al. [28] investigated the effectiveness of using multiple movement precues in a VR path-following task. They showed that study participants could use two to three precues if the precues were based on lines and only one precue if lines were not used. Later, Liu et al. [31] explored the impacts of using multiple cues for potentially parallel steps in an unordered task. They found that using gaze to determine the steps on which the user was likely to be working and showing just the corresponding cues helped reduce task-performance time.

The systems mentioned above [23, 28, 31, 49, 50] focused on translational movement rather than rotation. Liu et al. [29, 30] presented an AR testbed for exploring cueing and precueing tasks in which a user must move and rotate each of a series of objects, one at a time, placing it at a designated position and orientation. Movement and rotation could be performed simultaneously and rotation direction (clockwise or counterclockwise) was unspecified. This is unlike many rigid-body tasks, such as using a screwdriver, in which the direction of rotation matters. In contrast, we address cueing and precueing tasks in which rotation directions are specified, and a series of rotations are performed on the same object or on a set of objects that the user must move between in sequence.

3 TESTBED AND TASKS

We built a testbed using nine Vive Tracker 2.0 units [14] and a pegboard. The pegboard is placed on a table parallel to the ground. The trackers are the objects to be rotated and are arranged in a 3×3 square grid on the pegboard (Figure 2). Each is mounted on a 1/4"-20 bolt inserted through a pegboard hole, so that the tracker can be rotated about an axis vertical to the face of the pegboard. Tracker centers are 6" apart on each axis in the plane of the pegboard to ensure that the visualizations for one tracker will not overlap with those of adjacent trackers. We developed the user interface and associated application with Unity 2019.4.18f1 [8] and the Mercury Messaging framework [19]. The headset and computer used in our user study are described in Section 4.2.2.

Figure 2: Rotation precueing testbed.

3.1 Tasks

We consider two tasks: a *single-object task* that requires the user to perform sequential rotations of a single object and a *multiple-object task* that requires the user to move their hand to and rotate each of a set of objects in sequence.

Each task consists of multiple steps. In each step, the user needs to rotate a specific object by a designated angle (±40°, ±80°, ±120°, or ±160°) and remain within an angular tolerance of ±15° for 0.15s for the step to be considered complete. At that point, a beep sound is played and the next step begins. We chose the angular tolerance and time threshold based on user feedback in our pilot studies. We picked rotation angles smaller than 180° because work has shown that larger rotations are hard to perform with a single hand motion [16, 45, 51] and therefore harder to mentally simulate [36, 37]. If a real-life task requires a rotation larger than 180°, it might be possible to decompose it into a series of smaller ones.

3.2 Cueing and Precueing

In our system, *cues* display information about the current step, while *precues* show information about future steps. For a condition with m precued steps, when the user is working on step k, the information about step k is cued, while information about steps $k+1$, $k+2$, ..., $k+m$ are precued. After the user completes step k, step $k+1$ becomes the cued step, and steps $k+2$, $k+3$, ..., $k+m+1$ become the precued steps. This continues until the user finishes the entire task.

3.3 Single-Object Task

3.3.1 Rotation Visualization. The rotation cue and precues are circular arrows. For our single-object task, shown in Figure 1(b), the innermost circular arrow is the rotation cue, while the outer arrows are the rotation precues, all centered about the vertical axis of rotation, each in a separate plane. Precues for later steps are successively of larger radius, lower plane, darker, and more transparent. For the rotation cue, the end of the arrow tail follows the object's orientation when the user rotates the object (Figure 1c). The base of the pyramidal arrowhead indicates the goal orientation and does not move when the user rotates the object.

In each step, the user needs to rotate the object to make the tail of the rotation cue align with the arrowhead base within the tolerance range (Section 3.1). If the user rotates too far or in the wrong direction, the cue shows the rotation needed to compensate. Upon completing the step, the first precue becomes the cue: it becomes brighter and more opaque, its radius becomes smaller, and its plane becomes higher. The other precues update similarly and a new precue is added at the end (Figure 1d). This process repeats until the user finishes the task. We picked the colors of the cue and precues through pilot studies to ensure that users could easily distinguish them.

We use 4cm, 5.5cm, 7cm, and 8.5cm radii for the arrow bodies of the cue, first precue, second precue, and third precue, respectively, in the single-object task. Cue and precue arrow bodies are each 1cm wide. Each arrowhead is a pyramid with a 2cm×1cm rectangular base, and 2.5cm height.

3.3.2 Circular Arrow Styles. Our pilot studies revealed that highlighting different parts of the arrows can affect user performance. Therefore, we created the following arrow styles:

Unhighlighted (U). In the Unhighlighted style (Figure 3a), the arrowhead and arrow tail are not highlighted. The step is complete when the tail is within the (unvisualized) angular tolerance for the time threshold.

Head-Highlighted (H). In the Head-Highlighted style (Figure 3b), the arrowhead is highlighted by being replaced with a thick grey segment centered about the base of the head. The grey segment helps highlight the goal orientation. The grey segment occupies 30°, which corresponds to our ±15° angular tolerance. Consequently, the user can finish the step simply by moving the tail into the grey segment for the time threshold. We designed this style to compare whether highlighting the place to which the user needs to rotate the object or the direction information provided by the arrowhead is more helpful.

Tail-Highlighted (T). In the Tail-Highlighted style (Figure 3c), the arrow tail is highlighted by being replaced with a thick green segment centered about the tail. The green segment occupies 10° and highlights the current orientation of the object in the cue, or the origin orientation in a precue. The step is complete when the center of the green segment is within the angular tolerance for the time threshold. This style visualizes a tighter version of the angular tolerance: if the green segment overlaps the arrowhead base, it is within the middle 2/3 of the angular tolerance.

Head-and-Tail-Highlighted (HT). In the Head-and-Tail-Highlighted style (Figure 3d), the head and tail are replaced by a grey segment and a green segment, respectively. The green segment is the same width as that of T (10°), but the grey segment is wider than that of H by half of the width of the green segment on each side (totaling 40° instead of 30°). Consequently, the user can finish the step simply by moving the tail completely into the grey segment for the time threshold. (When this occurs, the tail, which is at the center of the green highlight, will be within the angular tolerance.)

3.3.3 Circular Arrow Alignment Approaches. Our pilot studies found that how we align the cue and precue arrows significantly impacts task performance. Therefore, we designed the following two alignment approaches:

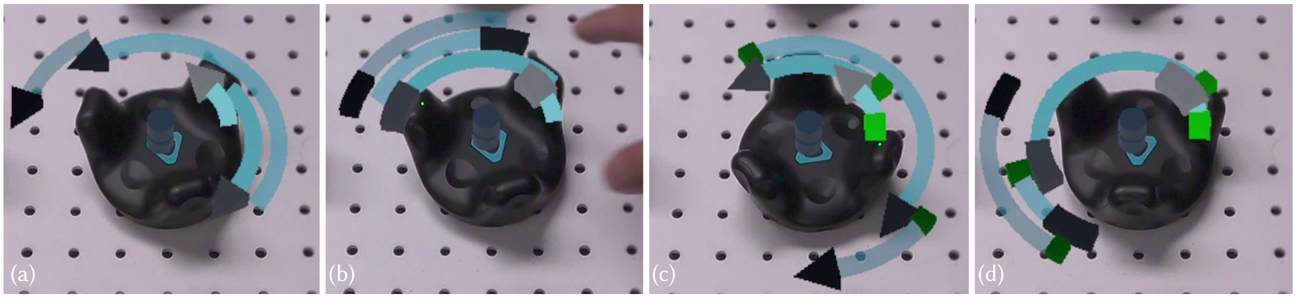

Figure 3: Circular arrow styles tested in our study. (a) Unhighlighted. (b) Head-Highlighted. (c) Tail-Highlighted. (d) Head-and-Tail-Highlighted.

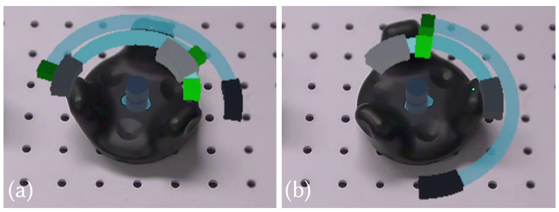

Figure 4: Alignment approaches. (a) Destination–Origin. (b) Origin–Origin.

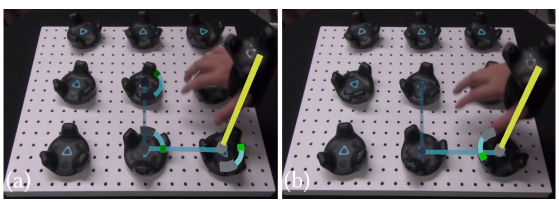

Figure 5: Part 3 visualizations. Both use a movement cue and two movement precues. (a) C+P: Includes a rotation cue and two rotation precues. (b) C-Only: Includes a rotation cue but no rotation precues.

Destination–Origin (DO). In this approach (Figure 4a), we align the destination (head) of a cue or precue with the origin (tail) of the next precue. This is intended to allow the user to more easily follow successive precues.

Origin–Origin (OO). In this approach (Figure 4b), we align the origins of the cue and all precues. This is intended to allow the user to more easily find the origins of the cue and precues, as they are always at the same angular offset around the object.

We also tested another alignment approach in which the middle points of the cue and precues are aligned. However, its performance was sufficiently similar to OO that we decided to exclude it from the formal study.

3.4 Multiple-Object Task

In the multiple-object task, the user needs to move their hand to a new object in each step. We use lines to cue and precue movement (translation) between objects. Previous work on precueing [28, 49, 50] has shown that straight lines (with or without arrowheads) can effectively provide users with translation information. For our line visualizations, we decided not to use arrowheads since they could obscure the rotation cue or precues. As shown in Figure 5, we use a yellow line connecting a wrist-worn tracker to the current object as the movement cue, while cyan lines connect successive objects as movement precues. The yellow cue line is hidden when the wrist-worn tracker is less than 15cm from the current task object.

The user rotates a different object in each step, so each object is augmented with only a single rotation cue or precue whose circular arrow body has a 4cm radius. We designed the multiple-object task to explore whether participants can use rotation precues in conjunction with movement precues. We consider two visualizations:

Rotation-Cue-and-Precues (C+P). This visualization (Figure 5a) has the same number of movement and rotation precues.

Rotation-Cue-Only (C-Only). This visualization (Figure 5b) has a movement cue and precues and a rotation cue, but no rotation precues.

4 USER STUDY

4.1 Hypotheses

Based on our design goals, observations, and results from our pilot studies, we formulated four hypotheses:

H1. *Participants will perform faster with a rotation precue than with no rotation precue in the single-object task.* Previous work [23, 28, 49, 50] has shown that providing one or more precues for future steps in movement tasks can help users prepare for future steps and therefore shorten task completion. Our pilot studies also found this was the case for our rotation task.

H2. *The HT arrow style will reduce step completion time more than the other styles in the single-object task.* More specifically, we expect the HT style will perform better than the other three in Section 3.3.2. Based on our pilot studies, we hypothesized that highlighting the goal and object orientations would help participants perform the task more efficiently.

H3. *The DO alignment approach will reduce step completion time compared to the OO alignment approach in the single-object task.* We hypothesize this based on our pilot studies. Since the participant can more easily find the next origin after a destination with DO alignment, they can perform better than with OO alignment.

H4. *Users will not perform worse in C-Only than in C+P in the multiple-object task.* We hypothesize this based on our pilot studies, in which we found that when movement is required, the user did not perform worse when given only the rotation cue but no rotation precue.

4.2 Method

4.2.1 Study Design. Our study used a single-session, within-subject design. It consisted of three parts. In Part 1, we tested the impact of the visualization styles of Section 3.3.2 on performance in the single-object task. Part 1 uses 16 conditions: (U, H, T, HT) × (0 Precues, 1 Precue, 2 Precues, 3 Precues). We decided to test 0–3 precues since cognitive science research has shown that human ability to track objects visually is limited [15, 32, 38], and performance in our pilot studies was fastest with one or two precues. Since the aim of Part 1 was to find the best styles, we blocked conditions by number of precues rather than styles, counterbalancing block order, and style order within blocks. We did this to avoid any participant encountering a style only at the beginning, middle, or end of Part 1.

In Part 2, we tested the impact of the alignment approaches of Section 3.3.3 on performance in the single-object task, using the HT circular arrow style. Part 2 uses eight conditions: (DO, OO) × (0 Precues, 1 Precue, 2 Precues, 3 Precues). As in Part 1, we blocked conditions by number of precues to avoid any participant encountering an alignment approach only at the beginning, middle, or end of Part 2.

In Part 3, we tested whether rotation precues are helpful when used in conjunction with movement precues in the multiple-object task, using the HT circular arrow style and the visualizations of Section 3.4. Part 3 uses seven conditions: (C+P, C-Only) × (1 Precue, 2 Precues, 3 Precues) and (C+P) × (0 Precues). (C-Only) × (0 Precues) is equivalent to (C+P) × (0 Precues), so we decided not to test it twice. As in Parts 1 and 2, we blocked the conditions by number of precues.

Participants performed three trials for each condition: one practice trial and two timed trials. We created three 16-step task sequences for Part 1 and Part 2. They were used for the first, second, and third trials in each condition, respectively. Each of the eight angles of Section 3.1 was used twice in a sequence. For the multiple-object task, we used the same three 16-step sequences. However, for each step, we picked a different object to rotate. In addition, no object was cued or precued by more than one movement visualization at a time. The sequences were designed to form a closed loop (in terms of object orientation for both tasks and hand position for the multiple-object task). This means we were able to pick a random start point in the sequence while ensuring when a specific step was performed, the user would see the same upcoming rotations.

Since, at the beginning of a trial, the user's hand might be in a random position, to avoid a confound, we inserted one preparation step, which was not included in the 16 steps, to guide the user to the central object in the single-object task or the origin from which the user should move their dominant hand in the first step of the multiple-object task. The visualizations worked the same way in the preparation step as in the 16 regular steps.

To avoid confounding effects related to a participant's anticipation of completing a trial, we presented an entire set of additional steps, past the point of the last step to be completed, so that additional precues would be rendered past the end. Thus, unless a participant was counting steps, which we found to not be the case in our pilot studies, they would not anticipate the end of a trial.

4.2.2 Equipment. Study participants wore a Varjo XR-3 video–see-through AR headset [9] with a 115° horizontal field of view (134° diagonal at 12 mm eye relief) and a 90Hz refresh rate. The XR-3 was 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. Each participant wore a Vive Tracker 3.0 on their dominant hand, attached with a Rebuff Reality TrackStrap. This was the only hand they were allowed to use to interact with the testbed described in Section 3. All trackers viewed the same base stations as the XR-3. The headset, trackers, and table were sanitized using 70% isopropanol before each session. The headset was also sanitized using a Cleanbox CX1 [3] UVC system.

4.2.3 Participants. Our study was approved by our institutional review board. Using convenience sampling through department email lists and posted flyers, we recruited from our institution 17 participants: three female, 19–40 years old (average 26.6), two left-handed. Three participants owned or once owned a VR headset, one had used VR in class projects, five had used AR/VR several times, and eight had no AR/VR experience. Of the recruited participants, two finished only Part 1, and one finished only Part 1 and Part 2. Thus, 17 participants completed Part 1, 15 completed Part 2, and 14 completed Part 3. Each participant received a USD 15 gift card as compensation.

4.2.4 Procedure. Each participant was first welcomed by the study coordinator and presented with an information sheet. After giving informed consent, the participant was then introduced to the study flow and given the Stereo Optical Co. Inc. Stereo Fly Test (SFT) [6], which contains nine questions, to screen for stereo vision and the Ishihara Pseudo-Isochromatic Plate (PIP) test [25] to screen for color deficiencies. One participant answered five of the SFT questions correctly, two participants answered seven correctly, one answered eight correctly, and the rest answered all correctly. For the PIP test, one participant had a color-vision deficiency and could not answer any question in the PIP test correctly, and the rest answered all correctly. Results from the SFT and PIP tests were not used to exclude participants.

The participant was seated in front of the table with the testbed, at a distance from which they could easily reach any of the nine trackers. The study coordinator then put the headset on the participant and attached a Vive Tracker 3.0 to the participant's dominant hand using a Rebuff Reality TrackStrap. The study coordinator started the study application, which recorded the 6DoF pose of each tracker at the end of each Unity frame (ca. 90 fps) and the completion time of each step. The participants were instructed to finish each rotation as fast as possible, and were informed of the angular tolerance and time threshold requirements. We decided not to ask the participants to achieve speed and accuracy at the same time, since participants might have different preferences for one over the other, which could potentially cause a confound.

After finishing all trials, the participant was asked to fill out a questionnaire that included demographic questions, a modified unweighted NASA TLX [33], and a request to rank the rotation precueing styles of Part 1 and alignment approaches of Part 2 based on their effectiveness. The TLX survey was modified to use a 1–7 scale, with 1 as best (as per [1]) rather than the original 0–20 scale. Each participant rated each of the rotation precueing styles for each TLX metric. Images of each rotation precueing style were displayed during the rating process to remind participants of the styles they tried.

4.3 Results

On study completion, we processed completion-time results generated automatically by our system before analyzing them. We used Tukey's outlier filter [46] to label outliers before analysis. Steps that took more than the third quartile plus 1.5 times interquartile range (third quartile minus first quartile) or less than the first quartile minus 1.5 times interquartile range were labeled as outliers. The "outside fence" for each condition and user was computed separately. We expected the conditions would have a significant effect on completion time, and we noticed that some users performed substantially better than others. The percentage of steps labeled as outliers in each condition can be found in the supplementary material.

We evaluated all hypotheses for significance with $\alpha = .05$. For each part of the study, we fit a separate linear mixed-effects model to our data using the MATLAB Statistics and Machine Learning Toolbox [5]. Table 1 lists the parameters used in each model. **M1**–**M3** denote the models used to fit the data from Part 1–3, respectively. Each model used the step completion time as the measurement. We used step (the specific step in the 16-step sequence) and user as the random-effect variables. There is no interaction term between random-effect variables. In contrast, we put interaction terms between the fixed-effect variables in each model. The terms in each model were picked based on comparing the current models and other alternative models using a likelihood ratio test. By comparing different linear mixed-effects models, we found that handedness, AR/VR experience, color-vision deficiency, and stereo vision are not significant factors. These may be because they were absorbed by the user random effect. **M1**–**M3** all have large effect sizes [13].

The rotation angle in a step significantly affected the step completion time. For example, we would expect a participant to take significantly less time to perform a 40° rotation than a 160° rotation. Therefore, instead of directly calculating the standard error of all steps in a condition, which would reflect differences in step difficulty and participant performance rather than instabilities in repeated measurements, we used linear mixed-effects models to calculate the contribution of each factor. For this reason, we do not include error bars in Figures 6–8, since they would be misleading. Please see the linear mixed-effects model results in the supplementary material for the standard error values.

The results of Part 1 are shown in Figure 6. In the rest of this section, we use *Parameter–NumberOfPrecues* to denote the condition in which *Parameter* is *style* in Part 1, *alignment approach* in Part 2, and *visualization* in Part 3. For **M1**, we used HT–0 Precues as the baseline condition. To evaluate **H1**, we checked the p-value of the 1 Precue term, which is .004. This shows that adding in the first precue reduces the step completion time.

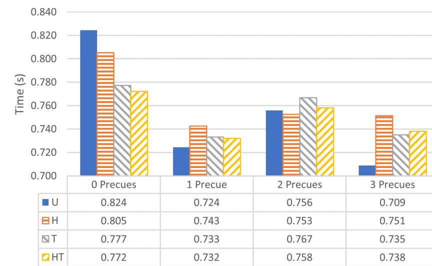

Figure 6: Part 1 Step CompletionTime.

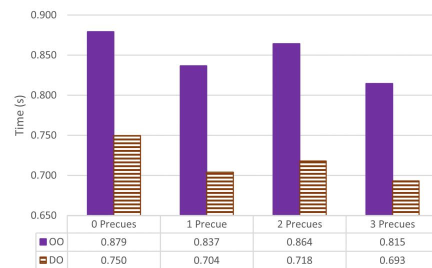

Figure 7: Part 2 Step CompletionTime.

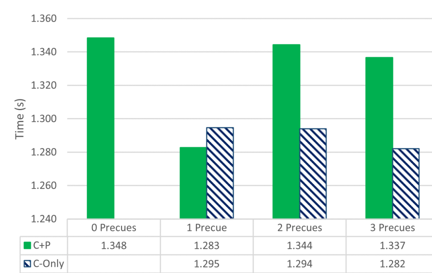

Figure 8: Part 3 Step CompletionTime.

To evaluate **H2**, we checked the p-values of the U, H, and T styles, which are $< .001$, .015, and .613, respectively. This shows that HT performed better than U and H, but not better than T. Therefore, **H2** is partially supported.

Figure 7 shows the results of Part 2. For **M2**, we use DO–0 Precues as the baseline condition. To evaluate **H3**, we checked the p-value of the OO term, which is $< .0001$. This shows that DO performed significantly better than OO, supporting **H3**. The p-values of the 1 Precue, 2 Precues, and 3 Precues terms are $< .001$, .034, and $< .001$, respectively. This shows that, consistent with Part 1, using one or more precues helps reduce step completion time. However, the p-values of the interaction terms are all $> .465$, indicating that the interactions between alignment approach and number of precues are not strong.

Figure 8 shows the results of Part 3. For **M3**, we do not fit the model to the data of C+P–0 Precues as there is no corresponding C-Only–0 Precues condition. We use C+P–1 Precue as the baseline condition. The C-Only term does not have a significant p-value (.530), while 2 Precues and 3 Precues have significant p-values (both

Table 1: Parameters of the linear mixed-effects models.

Model	Data Source	Hypothesis	Fixed-effect variable(s)	Effect Sizes
M1	Part 1	**H1, H2**	Precueing style, Number of precues	$\eta^2 = 0.5769$, Cohen's $d = 0.9992$ (Large)
M2	Part 2	**H3**	Alignment approach, Number of precues	$\eta^2 = 0.5788$, Cohen's $d = 1.0025$ (Large)
M3	Part 3	**H4**	Rotation Cue Only or All, Number of precues	$\eta^2 = 0.6412$, Cohen's $d = 1.1106$ (Large)

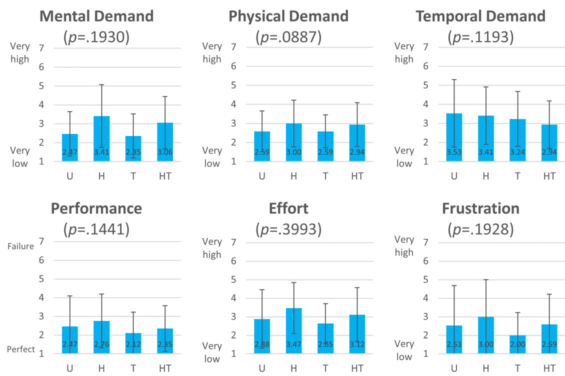

Figure 9: Part 1 NASA TLX results.

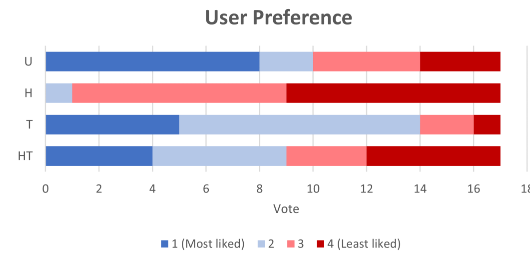

Figure 10: Preferences for circular arrow styles.

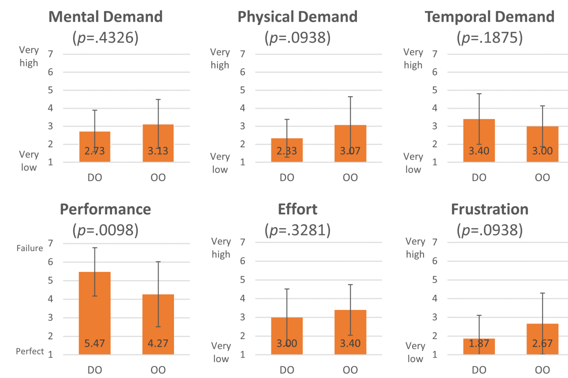

Figure 11: Part 2 NASA TLX results.

.003). The C-Only–2 Precues and C-Only–3 Precues interaction terms have significant p-values of .003 and .010, respectively. To evaluate **H4**, we calculate the contrast between C+P–2 Precues and C-Only–2 Precues, and the contrast between C+P–3 Precues and C-Only–3 Precues. The p-values are .012 and .002, respectively. For C+P–1 Precue and C-Only–1 Precue, we ran a non-inferiority test on them and picked a non-inferiority margin of 0.06s (completion time increased by 0.06s). The p-value of the test is .01, indicating that using C-Only will not degrade the performance under the non-inferiority margin. Therefore, **H4** is supported.

To avoid type-I errors, we used the Holm–Bonferroni method [24]. We checked a total of eight p-values (one for **H1**, three for **H2**, one for **H3**, and three for **H4**) to validate our hypotheses and seven of them are significant before correction. Among these seven p-values, two are $< .001$, and the remaining five are .002, .004, .010, .012, and .015 (lowest-to-highest). With this order, these p-values are smaller than .05/8, .05/7, ..., .05/2, respectively, meaning they survive their corresponding Holm–Bonferroni-corrected α.

4.3.1 User Feedback. Figure 9 shows the unweighted NASA TLX results for Part 1 and the p-values for each metric calculated using Friedman tests. We did not find a significant difference between the circular arrow styles for any of the metrics ($p > .05$).

The participants ranked the circular arrow styles based on their preferences (Figure 10). Participants preferred U the most, followed by T and HT, and finally H. A Friedman test on the preference data of Part 1 yielded $p = .0043$, showing the difference in preference is significant. Participants were also asked how many precues they thought were the most useful: no participants answered zero, seven participants answered one, eight participants answered two, and two participants answered three.

The unweighted NASA TLX results of Part 2 are shown in Figure 11. Since there were only two alignment approaches, we used pairwise t-tests rather than Friedman tests on the Part 2 NASA TLX results. The resulting p-values appear in the corresponding subfigures in Figure 11. Except for the performance metric ($p = .0098$), we did not find a significant difference between the two alignment approaches ($p > .05$). For Part 2 preferences, 12 out of the 15 participants preferred DO alignment over OO alignment. A Wilcoxon signed-rank test on the preference data of Part 2 yielded $p = .0352$, showing the difference in preference is significant.

4.3.2 Error Rate. We examined the steps labeled as outliers by Tukey's method [46] to understand the kinds of errors participants made in the study. Note that if a participant made an error, they still had to complete the current step (rotating in the right direction, within the angular tolerance for the time threshold) before proceeding to the next step. Since errors could have multiple causes, we decided to label them manually. We found that participants made the following major types of errors:

Fast (0.08%). Some steps were finished faster than the Tukey lower fence and labeled as outliers, but there were no obvious anomalies in the tracker trace.

Wrong Direction (0.78%). The participant rotated the tracker in the wrong direction. This typically happened when there was a cued or precued destination about 180° from the origin visualization of the cue. The participants parsed the destination information but somehow took the wrong direction.

Missed Target (1.25%). In these cases, the participant rotated the tracker towards the goal orientation but failed to meet the angular tolerance or exited the angular tolerance without satisfying the time threshold.

Work on Precue (0.81%). The participant tried to work on a precued orientation step rather than the cued goal orientation. This happened when information about multiple steps was shown in the scene (one or more precues were used) and the participant rotated to a precued orientation without completing the current step.

Tracking (0.13%). In rare cases, tracking would fail and cause the users to spend additional time on some steps. This only happened in the multiple-tracker task.

Wrong Object (0.05%). In rare cases in the multiple-tracker task, the participant's hand moved to the wrong tracker.

Others (0.89%). This category includes all errors not covered by the other error types. Most of these involve a combination of the other types. In some other cases, the participants spent a long time planning the step, but no anomalies were found in the hand-motion path or tracker rotations.

The percentages of each participant's error types under each condition can be found in the supplementary material. We used chi-square tests to analyze if the participants made more errors of a specific type under different conditions. All comparisons have either an insignificant p-value ($> .05$) or a small effect size ($\phi < 0.1$), indicating that a participant's performance in each condition was reflected in their completion time rather than their error rate.

4.3.3 Accuracy Analysis. In addition to the error rates, we calculate the average orientation error per step that participants made in each condition. Note that participants were asked to finish each as fast as they could, and the step was considered complete as soon as the tracker's orientation was within the angular tolerance for the time threshold. We calculated the accuracy to explore whether it differed between the conditions tested. The average error ranged from 8° to 10°, with no significant differences between conditions. The error data per condition can be found in the supplementary materials.

5 DISCUSSION

5.1 Comparison Between Circular Arrow Styles

Overall, HT and T performed the best, as shown in Section 4.3. To explore further, we considered the number of precues. For 0 Precues, T and HT performed better than H, and U performed the worst. This shows that highlighting either the head or the tail helped participants perform the task. Regarding why T performed better than H, it is possible that the original arrowhead acts as a highlight, though different from the grey segment in that it does not provide tolerance range information. Thus, for 0 Precues, performance is mostly decided by whether the tail is highlighted.

In 1 Precue and 2 Precues, the results were not different between styles; however, in 3 Precues, U performed better than the others. This may be because the number of objects people can track is limited [15, 32, 38]. U and H have only one highlight or arrowhead in each visualization, while in T and HT there are two: the tail highlight plus the head highlight or the arrowhead. In addition, U may have performed better than H in 3 Precues because the clear direction information provided by the arrowhead is more useful.

One interesting trend is that the average step completion time increased in 2 Precues, but decreased in 3 Precues. This is similar to the trend in work by Liu et al. [28]. One possible explanation is that the way participants segment the visualizations changes when the number of precues increases. When there are few precues, the participants might segment the highlights and the arrow body separately. However, when there are more precues, participants might interpret each precue as a whole.

5.2 Interaction Between Numbers of Movement and Rotation Precues That Are Useful

Our study results in Section 4 show that in the single-object task, participants could use one to three rotation precues. In the multiple-object task, participants were able to use multiple movement precues but not multiple rotation precues. This indicates that when multiple objects were involved, performance was predominantly determined by the number of movement precues rather than the number of rotation precues.

Using precues requires imagining future actions. Novick and Tversky [34] showed that people solving geometric analogies first imagine changes of location and then changes of orientation. People imagine and apply the transformations sequentially; they do not integrate them or imagine them at the same time. In each step of our multiple-object task, the participant moved their hand to an object and then rotated it. As in Novick and Tversky, they appeared to imagine moving before rotating.

5.3 Performance Difference Between Experimenter and Participants

The results of Part 3 show that our study participants could not use rotation precues in the multiple-object task. This is different from our self-pilot study, in which we found that the first author, highly practiced in this sort of task, benefited from one rotation precue. In addition, the first author performed faster than the participants in the formal study and performed best with the HT style and with one to two precues in Part 1 and Part 2. This shows that with training, people might be able to use rotation precues in the multiple-object task, and the best number of precues may be different in the single-object task.

5.4 Design Guidelines for Rotation Precues

Our study results in Section 4 show that participants performed best with the DO alignment approach and the HT circular arrow style. Thus, in our fast-paced task, using shapes to highlight the current and the goal orientations, and aligning the visualizations in a way that allows the participants to smoothly switch between visualizations, can improve task performance. Combining these, we infer that when designing rotation-precueing visualizations, one should properly highlight the current and final orientations of the object to be manipulated. In addition, if multiple rotation precues are displayed on a single object, one should properly align the subsequent visualizations to help participants find the next starting point. This can be applied to rhythm games such as *Rotaeno* [7].

When multiple objects are involved, the movement precues will dominate user performance. In this situation, it would be good if the rotation cue/precues on the upcoming objects were easily seen, so the user can process them before their hand reaches out to the object, as their gaze moves ahead of their hand [11, 52]. Possible applications include guiding a user to cook when multiple burners are used or training a novice power plant worker to adjust values in a system following a procedure. One can use a rotation cue to indicate the required rotation of the current object and a movement cue and precues to indicate future objects to be rotated. Of course, the best visualizations will vary depending on the tasks being addressed, so additional studies will be necessary to fine-tune visualizations for new tasks.

6 LIMITATIONS AND FUTURE WORK

6.1 Participant Demographics

As mentioned in Section 4, we recruited the participants from our institution. They were relatively young, and the proportion of female participants was low. Though we did not find a sex-based performance difference, further study is needed to determine if our results generalize to different participant demographics.

In addition, our study used a single-session design, which made it difficult to measure the performance of trained users. We believe that this underlies the difference between the experimenter and participants noted in Section 5.3. We would like to address this in future work.

6.2 Area Cursors and Circular Arrow Styles

Kabbash and Buxton [26] investigated how the size of the mouse cursor impacts the difficulty of selection, developing the concept of an *area cursor* that is larger than the objects to be selected. Our T style can be considered as an example, with the large tail highlight "selecting" the arrow base to complete the step. Further, the HT style uses a large tail highlight to "select" an even larger head highlight by being fully encompassed by it to complete the step. HT could thus be thought of in terms of a cursor and target, each having a separate width parameter, as contemplated by Kabbash and Buxton. In our case, the difference between the widths embodies the angular tolerance, as described in Section 3.3.2.

6.3 Multi-Tasking and Task Requiring Accuracy

In our study, the participants were asked to perform one step at a time and to finish as fast as possible. However, in some real-life tasks, people may need to perform multiple steps simultaneously with high accuracy. For example, when playing video games, the player may use a keyboard to control the movement of a character while using a mouse to perform actions. Although expert users can exhibit impressive mastery, novices can require much training to learn these skills. While there has been research on people's ability to multi-task [22, 41] and on precueing parallel steps of the same type [31], precueing different types of steps simultaneously with different speed and accuracy constraints will be interesting future work.

6.4 6DoF Cueing and Precueing

Steps in assembly tasks often involve 6DoF operations, such as moving a machine part to a destination with a specified pose. Applications of precues to these tasks might be similar to Part 3 of our study: providing only movement precues without rotation precues may be sufficient for many users. On the other hand, highly trained users may be able to exploit movement precues and rotation precues simultaneously, as suggested in Section 5.3. Wearable tracking devices [39] could make 6DoF body precueing possible. Additional studies will be needed to decide the most effective information to include in cues and precues and to design application-specific visualizations.

6.5 Continuous and Smooth Precueing

Our study used discrete cues and precues. That is, the participants needed to move to and manipulate a specific object and rotate it in a designated direction to a specific orientation. These goals were checked at discrete time points. However, some tasks might require users to move in more complicated paths at changing speeds with specific temporal constraints, all checked continuously. For example, the task explored by Ellis et al. [18] required the participant to move a ring along a complex curved path, with continuous constraints on its angle. In a situation such as this, the way in which cues and precues are formulated may need to be reconsidered, further taking into account the user's changing ability to process them. It is worth investigating whether our results could be transformed into a continuous form, similar to how Fitts' law [21] can be transformed into the steering law [10].

7 CONCLUSIONS

We investigated different approaches to cueing and precueing sequential rotation tasks in AR. In our single-object task, we found participants were able to use rotation precues. In addition, among the circular arrow styles tested, either highlighting the tail or highlighting both the head and tail yielded the best results when averaged over all numbers of precues. In addition, participants performed best when the destination of a cue or precue and the origin of the next precue were aligned. In the multiple-object task, rotation precues were not beneficial to the participants, even though the first author, who has used the system extensively during development, routinely performs better when provided with a single rotation precue in addition to movement precues. Finally, we believe our results are applicable to a variety of real-world tasks such as adjusting parameters by rotating knobs on a control panel.

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] MeasuringU 2022. *10 Things to Know about the NASA TLX*. MeasuringU. Retrieved 07/29/2022 from https://https://measuringu.com/nasa-tlx/
[2] Andamiro Co., LTD. 2022. *Chrono Circle*. Andamiro Co., LTD. Retrieved 07/29/2022 from https://chrono-circle.com/
[3] Cleanbox Technology 2022. *CleanboxCX1*. Cleanbox Technology. Retrieved 07/29/2022 from https://cleanboxtech.com/products/
[4] Nintendo 2022. *Joy-Con Rotation in 1-2-Switch*. Nintendo. Retrieved 07/29/2022 from https://www.nintendo.com/store/products/1-2-switch-switch/
[5] The MathWorks, Inc. 2022. *Matlab Statistics and Machine Learning Toolbox*. The MathWorks, Inc. Retrieved 07/29/2022 from https://www.mathworks.com/help/stats/index.html
[6] Stereo Optical Co. Inc. 2022. *Original Stereo Fly Stereotest*. Stereo Optical Co. Inc. Retrieved 07/29/2022 from https://www.stereooptical.com/products/stereotests-color-tests/original-stereo-fly/
[7] Dream Engine Games 2022. *Rotaeno*. Dream Engine Games. Retrieved 07/29/2022 from https://www.rotaeno.com/en
[8] Unity 2022. *Unity*. Unity. Retrieved 07/29/2022 from https://unity.com/
[9] Varjo 2022. *Varjo XR-3*. Varjo. Retrieved 07/29/2022 from https://varjo.com/products/xr-3/
[10] Johnny Accot and Shumin Zhai. 1997. Beyond Fitts' Law: Models for Trajectory-Based HCI Tasks. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) *(CHI '97)*. Association for Computing Machinery, New York, NY, USA, 295–302. https://doi.org/10.1145/258549.258760
[11] Cristina Becchio, Valeria Manera, Luisa Sartori, Andrea Cavallo, and Umberto Castiello. 2012. Grasping intentions: from thought experiments to empirical evidence. *Frontiers in Human Neuroscience* 6 (2012). https://doi.org/10.3389/fnhum.2012.00117
[12] Florent Berthaut and Alex Jones. 2016. ControllAR: Appropriation of Visual Feedback on Control Surfaces. In *Proceedings of the 2016 ACM International Conference on Interactive Surfaces and Spaces* (Niagara Falls, Ontario, Canada) *(ISS '16)*. Association for Computing Machinery, New York, NY, USA, 271–277. https://doi.org/10.1145/2992154.2992170
[13] Marc Brysbaert and Michaël Stevens. 2018. Power analysis and effect size in mixed effects models: A tutorial. *Journal of cognition* 1, 1 (2018).
[14] HTC Corporation. 2022. *VIVE Tracker 2.0*. Retrieved 07/29/2022 from https://www.vive.com/us/accessory/vive-tracker/
[15] Nelson Cowan, C Morey, and Zhijian Chen. 2007. The legend of the magical number seven. In *Tall tales about the brain: Separating fact from fiction*, Sergio Della Sala (Ed.). Oxford University Press, 45–59.
[16] Andrew Crossan, John Williamson, Stephen Brewster, and Rod Murray-Smith. 2008. Wrist Rotation for Interaction in Mobile Contexts. In *Proceedings of the 10th International Conference on Human Computer Interaction with Mobile Devices and Services* (Amsterdam, The Netherlands) *(MobileHCI '08)*. Association for Computing Machinery, New York, NY, USA, 435–438. https://doi.org/10.1145/1409240.1409307
[17] Francesca De Crescenzio, Massimiliano Fantini, Franco Persiani, Luigi Di Stefano, Pietro Azzari, and Samuele Salti. 2011. Augmented Reality for Aircraft Maintenance Training and Operations Support. *IEEE Computer Graphics and Applications* 31, 1 (2011), 96–101. https://doi.org/10.1109/MCG.2011.4
[18] Stephen R Ellis, Francois Breant, B Manges, Richard Jacoby, and Bernard D Adelstein. 1997. Factors influencing operator interaction with virtual objects viewed via head-mounted see-through displays: Viewing conditions and rendering latency. In *Proceedings of IEEE 1997 Annual International Symposium on Virtual Reality*. IEEE, 138–145.
[19] Carmine Elvezio, Mengu Sukan, and Steven Feiner. 2018. Mercury: A Messaging Framework for Modular UI Components. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. ACM, New York, NY, USA, Article 588, 12 pages. https://doi.org/10.1145/3173574.3174162
[20] Steven Feiner, Blair Macintyre, and Dorée Seligmann. 1993. Knowledge-Based Augmented Reality. *Commun. ACM* 36, 7 (jul 1993), 53–62. https://doi.org/10.1145/159544.159587
[21] Paul M Fitts. 1954. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology* 47, 6 (1954), 381.
[22] Anthony G Greenwald. 2003. On doing two things at once: III. Confirmation of perfect timesharing when simultaneous tasks are ideomotor compatible. *Journal of Experimental Psychology: Human Perception and Performance* 29, 5 (2003), 859.
[23] Morten Hertzum and Kasper Hornbæk. 2013. The Effect of Target Precuing on Pointing With Mouse and Touchpad. *International Journal of Human–Computer Interaction* 29, 5 (2013), 338–350. https://doi.org/10.1080/10447318.2012.711704
[24] Sture Holm. 1979. A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics* 6, 2 (1979), 65–70. http://www.jstor.org/stable/4615733
[25] S. Ishihara. 1972. *Ishihara's Tests for Colour-blindness*. Kanehara Shuppan.
[26] Paul Kabbash and William A. S. Buxton. 1995. The "Prince" Technique: Fitts' Law and Selection Using Area Cursors. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '95)*. ACM Press/Addison-Wesley Publishing Co., USA, 273–279. https://doi.org/10.1145/223904.223939
[27] Can Liu, Stephane Huot, Jonathan Diehl, Wendy Mackay, and Michel Beaudouin-Lafon. 2012. Evaluating the Benefits of Real-Time Feedback in Mobile Augmented Reality with Hand-Held Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) *(CHI '12)*. Association for Computing Machinery, New York, NY, USA, 2973–2976. https://doi.org/10.1145/2207676.2208706
[28] Jen-Shuo Liu, Carmine Elvezio, Barbara Tversky, and Steven Feiner. 2021. Using Multi-Level Precueing to Improve Performance in Path-Following Tasks in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 27, 11 (2021), 4311–4320. https://doi.org/10.1109/TVCG.2021.3106476
[29] Jen-Shuo Liu, Barbara Tversky, and Steven Feiner. 2022. Precueing Object Placement and Orientation for Manual Tasks in Augmented Reality. *IEEE Transactions on Visualization and Computer Graphics* 28, 11 (2022). https://doi.org/10.1109/TVCG.2022.3203111
[30] Jen-Shuo Liu, Barbara Tversky, and Steven Feiner. 2022. A Testbed for Exploring Multi-Level Precueing in Augmented Reality. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. 540–541. https://doi.org/10.1109/VRW55335.2022.00121
[31] Jen-Shuo Liu, Portia Wang, Barbara Tversky, and Steven Feiner. 2022. Adaptive Visual Cues for Guiding a Bimanual Unordered Task in Virtual Reality. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*.
[32] George A Miller. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review* 63, 2 (1956), 81.
[33] NASA. 1986. Task Load Index (TLX) v. 1.0 Manual. *Human Performance Research Group, NASA-Ames Research Center, Moffett Field* (1986).
[34] Laura R Novick and Barbara Tversky. 1987. Cognitive constraints on ordering operations: The case of geometric analogies. *Journal of Experimental Psychology: General* 116, 1 (1987), 50–67.
[35] Ohan Oda, Carmine Elvezio, Mengu Sukan, Steven Feiner, and Barbara Tversky. 2015. Virtual Replicas for Remote Assistance in Virtual and Augmented Reality. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology* (Charlotte, NC, USA) *(UIST '15)*. Association for Computing Machinery, New York, NY, USA, 405–415. https://doi.org/10.1145/2807442.2807497
[36] Lawrence M Parsons. 1987. Imagined spatial transformations of one's hands and feet. *Cognitive Psychology* 19, 2 (1987), 178–241. https://doi.org/10.1016/0010-0285(87)90011-9
[37] Lawrence M Parsons. 1994. Temporal and kinematic properties of motor behavior reflected in mentally simulated action. *Journal of Experimental Psychology: Human Perception and Performance* 20, 4 (1994), 709.
[38] Zenon W Pylyshyn and Ron W Storm. 1988. Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial vision* 3, 3 (1988), 179–197.
[39] Farshid Salemi Parizi, Wolf Kienzle, Eric Whitmire, Aakar Gupta, and Hrvoje Benko. 2021. RotoWrist: Continuous Infrared Wrist Angle Tracking Using a Wristband *(VRST '21)*. Association for Computing Machinery, New York, NY, USA, Article 26, 11 pages. https://doi.org/10.1145/3489849.3489886
[40] Christian Schönauer, Kenichiro Fukushi, Alex Olwal, Hannes Kaufmann, and Ramesh Raskar. 2012. Multimodal Motion Guidance: Techniques for Adaptive and Dynamic Feedback. In *Proceedings of the 14th ACM International Conference on Multimodal Interaction* (Santa Monica, California, USA) *(ICMI '12)*. Association for Computing Machinery, New York, NY, USA, 133–140. https://doi.org/10.1145/2388676.2388706
[41] Mariano Sigman and Stanislas Dehaene. 2008. Brain Mechanisms of Serial and Parallel Processing during Dual-Task Performance. *Journal of Neuroscience* 28, 30 (2008), 7585–7598. https://doi.org/10.1523/JNEUROSCI.0948-08.2008 arXiv:https://www.jneurosci.org/content/28/30/7585.full.pdf
[42] Mengu Sukan, Carmine Elvezio, Steven Feiner, and Barbara Tversky. 2016. Providing Assistance for Orienting 3D Objects Using Monocular Eyewear. In *Proceedings of the 2016 Symposium on Spatial User Interaction* (Tokyo, Japan) *(SUI '16)*. Association for Computing Machinery, New York, NY, USA, 89–98. https://doi.org/10.1145/2983310.2985764
[43] Ryo Suzuki, Rubaiat Habib Kazi, Li-yi Wei, Stephen DiVerdi, Wilmot Li, and Daniel Leithinger. 2020. RealitySketch: Embedding Responsive Graphics and Visualizations in AR through Dynamic Sketching *(UIST '20)*. Association for Computing Machinery, New York, NY, USA, 166–181. https://doi.org/10.1145/3379337.3415892
[44] Fernando Trujano, Mina Khan, and Pattie Maes. 2018. ARPiano Efficient Music Learning Using Augmented Reality. In *Innovative Technologies and Learning*, Ting-Ting Wu, Yueh-Min Huang, Rustam Shadiev, Lin Lin, and Andreja Istenič Starčič (Eds.). Springer International Publishing, Cham, 3–17.
[45] Mike Tucker and Rob Ellis. 1998. On the relations between seen objects and components of potential actions. *Journal of Experimental Psychology: Human*

Perception and Performance 24, 3 (1998), 830.

[46] John W Tukey. 1977. *Exploratory data analysis*. Vol. 2. Addison-Wesley Pub. Co., Reading, Mass.

[47] John Underkoffler and Hiroshi Ishii. 1998. Illuminating Light: An Optical Design Tool with a Luminous-Tangible Interface. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Los Angeles, California, USA) *(CHI '98)*. ACM Press/Addison-Wesley Publishing Co., USA, 542–549. https://doi.org/10.1145/274644.274717

[48] David E. Usevitch, Adam J. Sperry, and Jake J. Abbott. 2020. Translational and Rotational Arrow Cues (TRAC) Navigation Method for Manual Alignment Tasks. *ACM Trans. Appl. Percept.* 17, 1, Article 1 (feb 2020), 19 pages. https://doi.org/10.1145/3375001

[49] Benjamin Volmer, James Baumeister, Raymond Matthews, Linda Grosser, Stewart Von Itzstein, Siobhan Banks, and Bruce H. Thomas. 2022. A Comparison of Spatial Augmented Reality Predictive Cues and their Effects on Sleep Deprived Users. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 589–598. https://doi.org/10.1109/VR51125.2022.00079

[50] Benjamin Volmer, James Baumeister, Stewart Von Itzstein, Ina Bornkessel-Schlesewsky, Matthias Schlesewsky, Mark Billinghurst, and Bruce H. Thomas. 2018. A Comparison of Predictive Spatial Augmented Reality Cues for Procedural Tasks. *IEEE Transactions on Visualization and Computer Graphics* 24, 11 (2018), 2846–2856. https://doi.org/10.1109/TVCG.2018.2868587

[51] Andreas Wohlschläger and Astrid Wohlschläger. 1998. Mental and manual rotation. *Journal of Experimental Psychology: Human perception and performance* 24, 2 (1998), 397.

[52] J. Wolf, Q. Lohmeyer, C. Holz, and M. Meboldt. 2021. Gaze Comes in Handy: Predicting and Preventing Erroneous Hand Actions in AR-Supported Manual Tasks. In *2021 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE Computer Society, Los Alamitos, CA, USA, 166–175. https://doi.org/10.1109/ISMAR52148.2021.00031