

# Augmented Reality in the Psychomotor Phase of a Procedural Task

Steven J. Henderson\*

Steven K. Feiner\*

Columbia University



Figure 1: AR assistance for psychomotor phase of a procedural task. (Left) Dynamic 3D arrows and labels overlaid on aircraft engine components (center) respond to ongoing user activity to provide assistance in achieving (right) alignment of components. (Captured by a video camera looking through optical-see-through display used in study. A post-render filter was applied to remove camera distortion and vignetting.)

## ABSTRACT

Procedural tasks are common to many domains, ranging from maintenance and repair, to medicine, to the arts. We describe and evaluate a prototype augmented reality (AR) user interface designed to assist users in the relatively under-explored psychomotor phase of procedural tasks. In this phase, the user begins physical manipulations, and thus alters aspects of the underlying task environment. Our prototype tracks the user and multiple components in a typical maintenance assembly task, and provides dynamic, prescriptive, overlaid instructions on a see-through head-worn display in response to the user's ongoing activity. A user study shows participants were able to complete psychomotor aspects of the assembly task significantly faster and with significantly greater accuracy than when using 3D-graphics-based assistance presented on a stationary LCD. Qualitative questionnaire results indicate that participants overwhelmingly preferred the AR condition, and ranked it as more intuitive than the LCD condition.

**Keywords:** augmented reality, maintenance, repair, workpiece

**Index Terms:** H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—Artificial, augmented, and virtual realities; H.5.2 [Information Interfaces and Presentation]: User Interfaces—Training, help, and documentation; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual reality

## 1 INTRODUCTION

Procedural tasks involve people performing established sequences of activities, while interacting with objects in the physical environment, to accomplish specific goals [18]. These tasks occur in many domains and are often accompanied by various forms of assistance designed to help a person correctly understand and perform the procedure. This is particularly true in manufacturing and maintenance, where best practices or regulations often call for the use of manuals, schematics, and other aids for even the most routine tasks. In these

domains, many approaches have been employed to deliver procedural task assistance in practice or in research, including paper manuals, computer-based documentation, job cards, and augmented reality. *Augmented reality* (AR) refers to combining interactive media, such as graphics, with our perception of the real world, so that the added media are geometrically registered in 3D with the real world, interactively and in real time [6].

We focus on applying AR to what Neumann and Majoros [30] refer to as the *workpiece* or *psychomotor* phase of a manufacturing and maintenance task. In the psychomotor phase, the user performs physical manipulations, including comparing, aligning, and adjusting configurations of components. The psychomotor phase is complemented by the *informational* or *cognitive* phase, in which the user directs their attention (localizes), comprehends instructions, and transposes information from the instructions to the actual task environment. This bipartite view of procedural tasks is corroborated by Richardson and colleagues [35], who identify a similar division of activities for assembly tasks.

In earlier work [23, 24], we developed an AR prototype to assist with maintenance activities inside an armored personnel turret, and found that AR assistance was beneficial in the informational phase, allowing mechanics to localize tasks more quickly than computer-based documentation on a stationary display. However, we were unable to show any advantage for AR in the psychomotor phase of our selected tasks, which involved simple actions such as flipping a switch or removing a bolt. We have since applied an improved version of our prototype to a related domain, and are now able to report performance gains afforded by AR in the psychomotor phase of more complex procedural tasks.

In this paper, we present a research prototype that (1) tracks multiple independent physical domain objects and the user's head relative to the world, (2) uses this information to provide an AR user interface that offers dynamic, prescriptive feedback and instructions that guide the user to accomplish an interesting procedural task, and (3) has been validated with a formal user study showing significant improvements in overall time and accuracy in comparison with non-AR 3D-graphics-based electronic documentation on a stationary display. We are not aware of any previous work that meets all three of these criteria, and believe that this combination articulates important opportunities for AR to yield speed and accuracy benefits during psychomotor activities.

\*e-mail: {henderson,feiner}@cs.columbia.edu

We begin with a review in Section 2 of related research on supporting procedural tasks with AR and other computer-based documentation. Then, Section 3 describes the design of our AR prototype (Figure 1) for providing assistance during procedural tasks. This system uses a tracked, see-through, head-worn display to combine 2D and 3D graphics with a user's view of the task environment. Although our prototype is designed to assist the user throughout a procedural task, we focus our discussion on AR techniques specifically designed for psychomotor activities. Next, in Section 4, we present a within-subject user study in which participants completed a series of assembly tasks in a realistic maintenance and repair domain using the prototype AR system and an enhanced form of current computer-based documentation. The assembly tasks involved multiple iterations of challenging component alignments. Our analysis in Section 5 shows that participants were able to complete the alignment and assembly more quickly and more accurately in the AR condition. In Section 6, we review the results of a second, pilot study, which found no statistical differences in performance between the AR condition and an idealized, but often impractical, condition featuring physical labels directly affixed to the components being assembled. Finally, we discuss our results and future work in Section 7 and present our conclusions in Section 8.

## 2 RELATED WORK

The use of AR to assist in the completion of procedural tasks overlaps with a larger body of work involving the use of computer graphics to either directly assist workers or enhance the design of documentation. Booher [9] implemented a programmable task simulator for exploring information presentation techniques, and applied it to procedural tasks. Systems developed by Feiner [14], and Seligmann and Feiner [40], demonstrated techniques and design heuristics for the automated generation of explanatory illustrations. Tversky, Morrison, and Bétrancourt [44] reviewed the efficacy of animated graphics in promoting task cognition, including the comprehension of procedural tasks. Agrawala, Heiser, and colleagues [2, 22] examined cognitive design principles for assembly instructions and developed an algorithm for automatically generating procedural steps and illustrations. Li and colleagues built on this work to support automated and interactive exploded [27, 26] and cutaway [28] views.

Focusing on AR for procedural maintenance tasks, much effort has been devoted to assembly, beginning with the aircraft wire bundle assembly domain explored by Caudell and Mizell [11] and field tested by Curtis and colleagues [12]. Reiner and colleagues [34] demonstrated AR documentation for car door assembly. Separate work by Tang and colleagues [42] and Robertson, MacIntyre and Walker [36], both addressing toy block assembly, demonstrated quantitative performance advantages for AR over traditional media and non-registered overlaid graphics. Zauner, Haller, and Brandl [47] presented, but did not evaluate, step-by-step AR instructions to help in assembly of furniture, while Salonen and Sääski [38] found that users gave high qualitative ratings to their AR system for 3D puzzle assembly. Echlter and colleagues [13] developed a hand-held AR system for assisting welders that was later used in a production environment.

AR systems for repair have addressed domains ranging from end-user servicing of a laser printer [15] to automotive repair [25]. Several industrial and academic consortia have been formed to address AR for maintenance, including ARVIKA [16], Service Training through Augmented Reality (STAR) [33], ARTESAS [5] and AVILUS [1]. Other important work includes a mobile AR prototype designed by Platonov and colleagues [31] that uses markerless tracking to present instructions to automobile mechanics.

There is much notable work on AR for procedural tasks in domains other than maintenance and repair. In medicine, State and colleagues [41] developed an AR system for assisting in breast

biopsies. Follow-on work by Rosenthal and colleagues [37] showed that a skilled surgeon could complete biopsies more accurately with the AR system than with standard imaging. Wacker and colleagues [46] demonstrated and evaluated a needle biopsy system using a swine model. Blum, Sielhorst, and Navab [8] used AR to visualize the path through which a trainee manipulates an obstetric forceps relative to the trajectory followed by an expert. Quarles and colleagues [32] demonstrated how AR could be used to teach the principles of operation of an anesthesia machine. In the arts, systems have been developed to guide users in acting out film scripts while interacting with virtual characters [19], and to assist with playing music [10, 29, 7]. In logistics an AR system for warehouse order picking was shown to enable users to perform more quickly than a conventional paper-list-based approach [39].

All the previously cited AR systems for procedural task assistance track the position and orientation of the user's head relative to the task environment. However, relatively few of these systems simultaneously track one or more movable physical domain objects relative to the environment, in addition to the user's head, as we do. Here, we compare our work with these systems. Several assembly [47, 38] and maintenance [15] applications track task objects, but beyond overlaying models of the tracked objects, they use this information solely at the start and end of psychomotor activities to detect task transitions and verify correct alignment. In contrast, we provide continuous prescriptive feedback for alignment tasks. The needle biopsy systems [41, 37, 46] display virtual representations of tracked biopsy needles within the patient's body. These systems differ from ours in that they intentionally do not provide explicit prescriptive instructions to the user, but instead rely on the skilled surgeon to make her own decisions based on the AR visualization. Blum, Sielhorst, and Navab [8] use AR to depict an expert obstetrician's complex prerecorded actions to a trainee attempting to emulate the expert's behavior, but we are not aware of any quantitative user studies of this work.

## 3 PROTOTYPE

We created a prototype to study the application of AR to procedural tasks encountered in a realistic maintenance domain. This domain is structured around a Rolls-Royce Dart 510 turboprop engine (Figure 2a) located in our lab.

We selected several procedures involving the disassembly, re-assembly, installation and removal of the Dart engine's combustion chambers (Figure 2b). The seven combustion chambers are mounted on the aft section of the engine, as shown in Figure 2(a). Although they all appear similar at first glance, each chamber is different, and contains a unique arrangement of intake and exhaust ports, in addition to minor ports to power ancillary systems on the engine (e.g., deicing). Some of these features are depicted in Figure 2(b-d). Assembling an individual chamber requires a precise alignment between the chamber's truncated conical upper section (shown in Figure 2c, which we will refer to as a "cone") and its mostly cylindrical lower section (shown in Figure 2d, hereafter referred to as a "can"). Each cone and can has a flange circled by a set of 20 evenly spaced holes, which can be seen in the figure. Installing an assembled chamber on the engine requires precise positioning and alignment relative to the other chambers to form an interconnected system of seven chambers. We identified the procedural tasks of assembling and installing these chambers as good opportunities to apply AR for assisting with psychomotor activities.

### 3.1 Software

We developed an AR user interface for exploring these procedural tasks using Goblin XNA, a managed, DirectX-based framework for constructing AR applications [21]. The user wears a tracked, stereo, optical-see-through head-worn display (HWD), described in Section 3.2. The position and orientation of the HWD are used

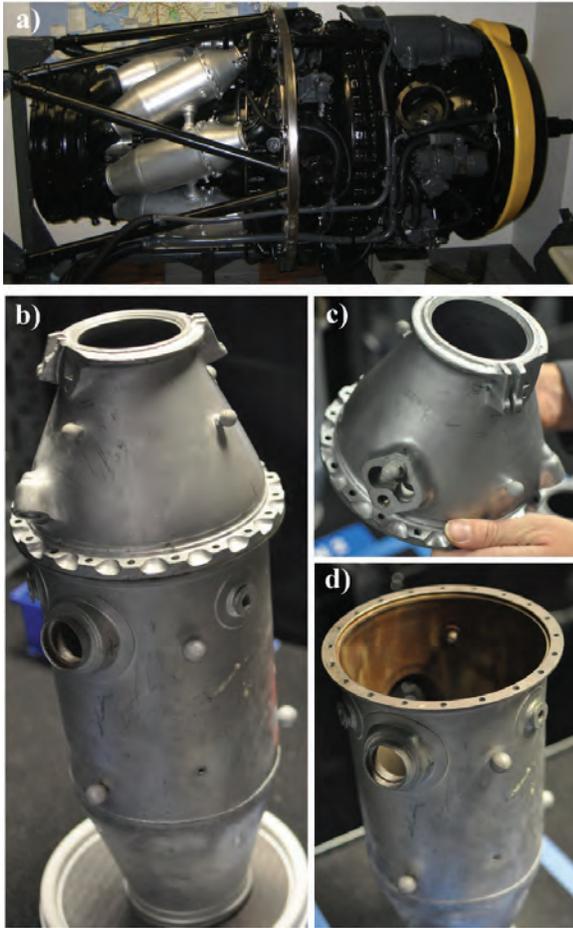


Figure 2: Procedural task domain. (a) Rolls-Royce Dart 510 engine. (b) Combustion chamber. (c) Upper combustion chamber “cone.” (d) Lower combustion chamber “can.” (Small spheres bolted to components are used for tracking.)

to update the location of a virtual stereo camera model, which is calibrated as described in Section 3.2 and used to generate views of a 2D head-up display (HUD) layer and 3D scene graph maintained by the system. The views through each half of the virtual stereo pair are rendered side-by-side against a black backbuffer (for optical-see-through display) within a single  $2560 \times 1024$  viewport at a combined framerate of 60fps. The interface uses a finite state machine to manage visibility of content in the HUD and scene graph, where each state represents a single step in the procedure. State transitions are either manually cued by using an input device, or triggered automatically using a proactive computing model [43] similar to the non-AR prototype demonstrated by Antifakos and colleagues [4]. That is, in certain steps where the worker must reposition an object in order to complete the task, our system uses tracking data to monitor this activity and either automatically transitions to the next step or displays an error message.

### 3.1.1 Informational Phase Assistance

Our prototype delivers several forms of assistance throughout both the informational and psychomotor phases of procedural tasks. Assistance designed primarily for the informational phase includes 2D text drawn in a screen-fixed HUD at the top of the display. This text communicates the currently desired action, and any salient objects and physical locations involved with the action. We also implemented a localization approach similar to our earlier work [23, 24]. A screen-fixed 2D arrow first orients the user in the general direc-

tion of the target task (within  $\pm 90^\circ$  azimuth) and is replaced by a 3D arrow extending from the near clipping plane to the target location. As the user centers the task location in their field of view, the arrow fades to full transparency and a brief highlighting effect is displayed as the arrow disengages. Our informational phase assistance also uses 3D “billboarded” (rotating dynamically to face the user) labels that are positioned on static components of interest to help the user distinguish objects in the task environment. These labels are sometimes accompanied or replaced by static or animated models that are used to help demonstrate desired end states and generalized movements. Our prototype also includes static motion paths that are rendered as Bézier curves depicting an idealized path between an object’s current location and its prescribed destination. These techniques rely on a fixed notion of the task environment, and are not altered in response to ongoing user activity.

### 3.1.2 Psychomotor Phase Assistance

We experimented with the following techniques for assisting the worker in the psychomotor phase of a task:

- *Dynamic 3D arrows.* We designed and tested several kinds of tracked 3D animated arrows that are rendered over or near a movable object to suggest motion or provide feedback about the current orientation compared to a desired end state. As shown in Figures 1 and 3, these arrows alter their size, color, animated direction, or visibility in response to user activity.
- *Dynamic 3D highlights for connection points.* We implemented a series of color-coded, semitransparent highlight effects, shown in Figure 4, which are designed to help in connecting and aligning two rigid bodies, such as the cone and can. As the cone and can are brought into alignment, the matching color-coded highlight on the receiving component (can) fades out until the two connection highlights appear as a single entity.
- *Dynamic billboarded labels.* We extended the static billboarded labels used in informational activities to respond to motion of the user and the labeled object. This includes updating a dynamic visibility occlusion model of the movable objects. Examples are shown in Figures 1 and 3.

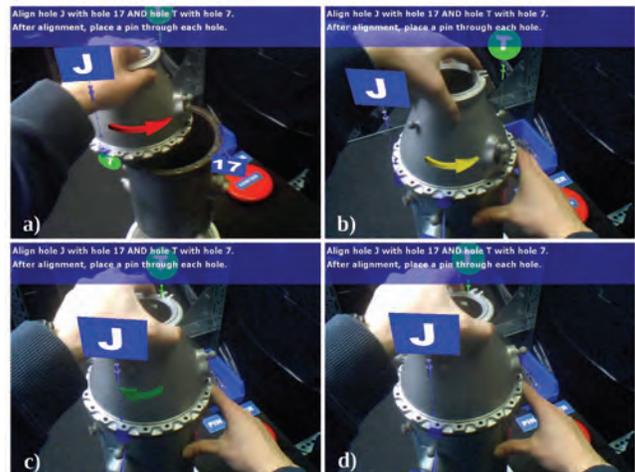


Figure 3: Dynamic 3D arrow. (a) Large, red, dynamic arrow indicates direction and magnitude of motion needed to align can and cone. (b) As can and cone approach alignment, arrow reduces size and changes color through yellow to green and, if necessary, (c) alters direction to specify shortest rotational direction to alignment. (d) When alignment is achieved, arrow fades away. (Images show view through a video-see-through display.)



Figure 4: Dynamic highlights. (Left) Prior to alignment, distinct highlights are rendered on top and bottom holes. (Right) When holes are aligned, bottom highlight slowly fades out, presenting the appearance of a single hole. (Images show cropped enlargement of view through a video-see-through display.)

### 3.2 Hardware

Our prototype uses an NVIS nVisor ST60 color, stereo, optical-see-through HWD (Figure 5). The display has a  $60^\circ$  diagonal field of view per eye, 40% optical transmissivity (of light from the physical task environment), and provides a  $1280 \times 1024$  resolution image to each eye. The display is driven by an ATI Radeon HD5770 Eyefinity graphics card in a quad-core, 3.4GHz AMD Phenom II 965 powered desktop computer with 8GB RAM, running a 64-bit version of Windows 7. The side-by-side stereo pair rendered by our software interface (described in Section 3.1) is mapped by the graphics card display manager to the left and right DVI channels of the HWD. We also support video-see-through HWDs (Vuzix Wrap 920AR and VR920 with CamAR), which we have used to make some of the figures in this paper (as indicated in their captions), to avoid the need to photograph through the nVisor optics.

We support two kinds of optical tracking. The 3D position and orientation of the HWD, cones, and cans are optically tracked at 100Hz using a cluster of 11 NaturalPoint OptiTrack FLEX:V100R2 and FLEX:V100 infrared cameras, mounted around the work area. Three retroreflective sphere markers are fixed to the back of the HWD to create a single rigid body to be tracked by the NaturalPoint TrackingTools application. Each cone and can is outfitted with four markers, so that it is continually tracked within its range of motion. While we acknowledge that it will typically not be possible to add such markers to objects in a production environment, we note that our prototype is a laboratory proof-of-concept system, designed to explore potential future applications that could use alternative tracking technology, such as model-based vision. Moreover, many training applications involving procedural tasks performed with dedicated equipment would permit the addition of markers.

We also experimented with the VTT ALVAR optical tracking library [3], which uses printed fiducial markers that we attached to each combustion chamber. A  $640 \times 480$  resolution Point Grey Firefly MV IEEE 1394a camera that we mounted on a fabricated ledge and attached to the nVisor ST60 (Figure 5) provides input to the ALVAR library. This worked well for tracking large objects held in the user's hands.

We calibrate the AR application by first performing intrinsic and extrinsic calibration of the OptiTrack. Next, we apply the HWD calibration technique of Fuhrmann, Splechtna, and Prikryl [17]. In our implementation, the user holds a  $10\text{cm} \times 10\text{cm}$  patterned target, outfitted so that it can be tracked by the OptiTrack. The user is requested to align the real target with a series of virtual targets, projected at eight 3D locations per eye, first within the left, and then within the right, view frustum of the display. These correspondence points are then used to solve the extrinsic calibration between each eye and the tracked HWD. We have tested this technique with dozens of users and found it to provide excellent results over our tracking volume. We use a modified version of this calibration technique to calibrate the camera used for fiducial marker tracking.



Figure 5: Study participant wearing the HWD. (Photo used with permission of participant.)

### 3.3 AR Assistance During Assembly Task

We applied our AR prototype to the combustion chamber assembly procedural task, which we then evaluated with the user study described in Section 4. A typical combustion chamber assembly using our AR prototype proceeds as follows. The prototype displays text instructions in the 2D HUD of the HWD, instructing the user to locate and pick up a specified can. The localization technique described in Section 3.1.1 guides the user to the can's current location. Virtual labels are provided to help the user identify the can and other objects in the task environment. When our prototype detects the user has begun to move the can, new text is displayed in the HUD, instructing the user to reposition the can to a specified assembly area on a workbench. The prototype also shifts localization cues to this assembly area and presents the virtual motion path leading to this target location. If the user repositions the wrong object, the system is able to detect the error from tracking data and will display an error message asking the user to return the misplaced object to its correct position on the workbench.

After the user successfully places the can in the work area, the process is repeated to locate and move the appropriate cone for assembly. Once the user places the cone within 30cm of the top of the can, the prototype begins to display a dynamic 3D curved arrow, centered about the upright ( $Y$ ) axis of the cone, representing the optimal direction of rotation to bring the cone and can into the desired alignment. As shown in Figures 1 and 3, the size and color of the arrow change dynamically to reflect the magnitude of the motion required to achieve the desired alignment (e.g., large red arrows and small green arrows indicate large and small corrections, respectively). Virtual labels, redundantly encoded using color and shape, identify the assembly's connection points. As explained in Section 3.1.2, dynamic highlights are provided to help the worker identify these connection points and determine when they are aligned. The worker then secures the alignment with a set of pins. This align and pin task requires an  $18^\circ$  level of precision and is traditionally achieved by counting holes from a single, factory-printed reference mark on the edge of each cone and can.

## 4 USER STUDY

We designed a user study to compare the performance and general acceptance of our AR prototype described in Section 3.3 (the AR condition) with that of 3D-graphics-based documentation. The primary procedural task in our study involves assembling a combustion chamber by aligning a can with a cone. We integrated this task into a workbench setting consisting of three cones and three cans, each positioned in one of six bins arrayed on the workbench, as shown in Figure 6. A portion of the workbench was set aside as the work area, where the participant is instructed to assemble a combustion chamber while standing in front of it. A small mechanical turntable was placed in this area to receive the can and facilitate rotation during assembly. We also placed a container of fastening

pins, which are used in our selected task, in the corner of the workbench for easy access by the participant.

We selected the combustion chamber assembly task described in section 3.3 for both the AR and LCD conditions. We chose this task because it could offer a large sample of independent, homogeneous tasks that were resistant to experiential effects. We achieved this by modifying the actual combustion chamber assembly procedure to allow arbitrary pairings and alignments of cans and cones across all combustion chambers. This allowed us to present our users with a large number of possible unique tasks, preventing memorization of distinct arrangements. (Note that the single “correct” assignment of properly aligned cones and cans is not visually obvious, would not be known by anyone who was not trained in the maintenance of this specific engine, and would only be important if the combustion chambers were to be installed in the engine.) As described above, to minimize time spent on the task and facilitate disassembling cone-can pairs for subsequent trials, we further modified the task to use two pins instead of the 20 machine bolts normally used to secure the assembly. We believe that these are valid modifications, since our participants had no prior knowledge of the engine, and our prototype was not intended to provide improved documentation for a set of 20 identical bolt fastening tasks.

#### 4.1 Control Condition

We created a control condition (the LCD condition) that presents 3D-graphics-based documentation corresponding to material in our AR prototype on a 22" diagonal 1920×1080 LCD screen, mounted near the repair area. Because the documentation available for the Dart 510 (which has been out of production for decades) is limited to printed materials, we developed a significantly enhanced version of the computer-based documentation currently employed by many professional mechanics. This computer-based documentation is exemplified by US Department of Defense Interactive Electronic Technical Manuals (IETMs) [45]. These computer-based manuals allow mechanics to use portable computers to browse hypermedia documents containing text, drawings, and historical data. We created an interactive 3D documentation system that employs previously proposed design principles for assembly instructions [22]. This baseline is similar in spirit to that used in our earlier prototype [23, 24]. It incorporates many aspects of the documentation used in the AR condition (e.g., text, instructions, labels, and motion paths), but presents them to the user on a fixed display without using head tracking. Like the AR condition, but unlike current IETMs, the baseline automatically advances the documentation as tracked components are moved.



Figure 6: Study task environment. LCD screen at upper right is used as primary display in control condition and turned off in AR condition.

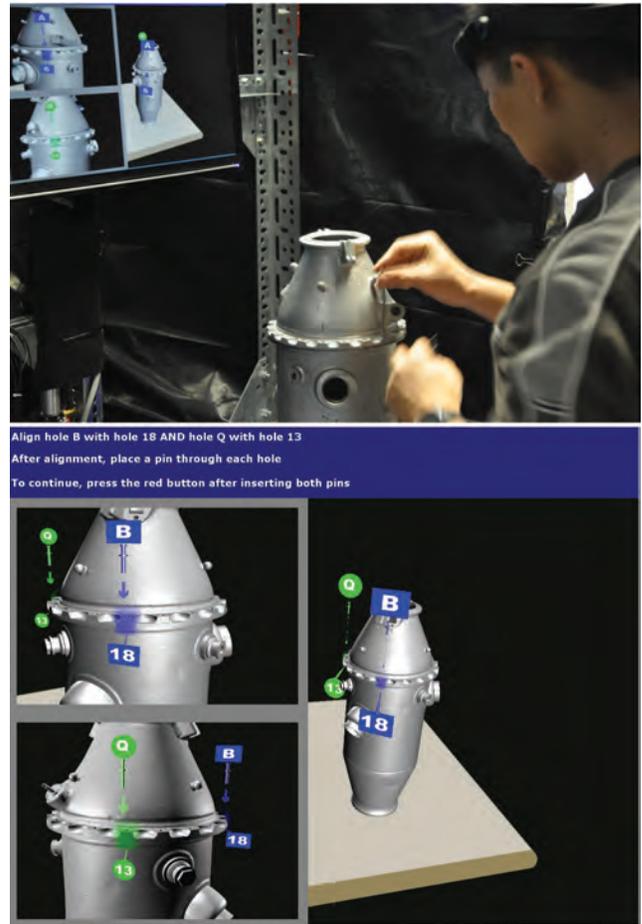


Figure 7: LCD condition. (Top) User performing pinning. (Bottom) Screen captured display from a similar task. (Top photo used with permission of participant.)

Since this documentation is presented on an opaque display, and not overlaid on the task domain, we render it using 3D virtual models corresponding to salient physical objects in the task environment that provide important cues when participants view them directly in the real world in the AR condition. These include a model of the workbench and detailed models of the cans and cones, created from 3D laser scans of the actual components. During localization and movement, this content is used to generate static perspective views rendered from camera poses corresponding to the user’s location at each stage of the task. That is, the 3D models of the can and cone presented on the LCD are rendered in a perspective that reflects the actual position and orientation of the participant’s head at the beginning of the psychomotor step. For the alignment and pinning task, we generate additional close-up views designed to assist the user in identifying the prescribed attachment and pinning points. Figure 7 shows a participant performing pinning in the LCD condition (top) and a screen-captured example display from a similar task (bottom).

#### 4.2 Experiment Design

We designed a within-subject, repeated measures experiment consisting of two conditions (AR, LCD) and randomized iterations of the combustion chamber assembly task. The experiment was blocked by condition, with a five-minute break between blocks. Each block consisted of a set of trials whose number was established through experimentation during a pilot study, as detailed in Section 4.3. Block order was counterbalanced across both con-

ditions to mitigate experiential effects. Each trial was defined as the ordered execution of six steps required for successful assembly of a combustion chamber. We categorized these steps, which are listed in Table 1, according to the predominant human activity employed during that step. This categorization, which we created by adapting the taxonomy of major activity types proposed by Gilbreth and Gilbreth [20] to our task domain, was useful in isolating the steps within the larger procedural task involving psychomotor human abilities. As shown in Table 1, each trial consisted of assembling one of three cones with one of three cans, aligning the cone and can correctly, and inserting two pins through matched pairs of holes on the components. The assignments of cans to cones and the hole pairings were fixed across conditions and participants (i.e., all participants experienced a fixed set of cone to can pairings under both conditions), and were generated pseudorandomly prior to the experiment. Learning effects were controlled by counterbalancing the starting condition.

Participants were screened for stereopsis using the Stereo Optical Co. Stereo Fly test. This test was always administered right before the AR block. If a participant were to fail the stereopsis test, the corresponding data would be excluded from the statistical analysis, but they would still be allowed to complete the study.

Prior to starting the experiment, each participant was asked to sign a consent form and then watch an instructional video corresponding to the participant’s assigned starting condition. The instructional video asked participants to “work as quickly and accurately as possible.” Following the video, participants starting with the AR condition were assisted with donning the nVisor ST60 HWD. Participants starting with the LCD condition were asked to wear a small, lightweight crown affixed with OptiTrack markers, which was used to collect head movement data.

Participants were then given a short practice period involving five trials to become comfortable with the first condition. The participant then began the timed portion of the block, which started with the participant pressing a large red button on the workbench near the turntable (Figure 6). The participant then performed each of the steps listed in Table 1. The completion times for steps 1–4 and 6 were logged automatically, based on state-machine transitions triggered in response to user activity, as described in Section 3.1. The completion time for step 5 was measured when the participant pressed the button near the turntable to confirm completion of the assembly. Prior to transitioning to step 6, the system also calculated the mean alignment error between the can and cone by sampling the rotational difference (yaw) between the can and cone, and comparing this sample mean to the angular difference specified by an optimal alignment.

The block proceeded to the next trial after the participant placed the completed assembly in a designated bin at the conclusion of step 6. Because we used only three combustion chambers in the experiment design, some cone and can recycling was required to support multiple trials. This was accomplished by inserting a disassembly task at certain points in the block, which involved the participant locating one of the completed combustion chambers, moving it to the turntable, removing the pins, and then placing the can and cone back into designated bins.

Following all trials in the first block, the participant removed any headgear, was afforded a five-minute break, and watched the instructional video for the alternate condition. The participant then donned either the HWD (AR) or the tracking crown (LCD), was given a five-trial practice period for the alternate condition, and then proceeded to complete all assembly trials under the alternate condition.

### 4.3 Pilot Testing

Six participants (1 female; ages 18–40,  $\bar{X} = 24.8$ ) were first recruited for a pilot study designed to test our experiment, elicit feed-

Step	Description	Activity Type
1	Locate Can X in Bin W	Locate
2	Move Can X to Turntable	Position
3	Locate Cone Y in Bin V	Locate
4	Place Cone Y on Can X	Position
5	Align Cone Y with Can X; Insert pins	Align
6	Move assembly XY to Bin Z	Position

Table 1: Steps of combustion chamber assembly task.

back about our conditions, and form hypotheses. The participants we recruited by mass email to Computer Science students at our university and by flyers distributed throughout campus, and were paid \$15 each. All participants in the pilot study used a computer multiple times per day. One participant reported having some experience repairing mechanical systems, three reported having a basic level of exposure, and two reported having no experience. Three participants identified themselves as requiring corrective lenses, and wore them in both conditions, since the nVisor provided enough eye relief to accommodate glasses. All six were determined to correctly perceive stereo stimuli.

In this pilot study, each block contained 18 trials. We applied a 2 (Display Condition)  $\times$  3 (Activity Type) repeated measure ANOVA on the mean completion time for each of the steps in Table 1. Using an  $\alpha$  of 0.05, we found a significant main effect of display condition on completion time ( $F_{(1,5)} = 17.14, p = 0.009$ ). A similar analysis found a significant main effect of display condition on accuracy of alignment between the can and cone ( $F_{(1,5)} = 11.51, p = 0.019$ ). An analysis of errors accumulated during localization and positioning activities showed participants made very few errors during these activities in either condition. The pilot study also revealed strong user preference for AR compared to LCD, with all six participants ranking AR as more preferred, as well as the more intuitive of the two systems.

Finally, our pilot study was also helpful in setting the number of trials in each experiment block. We originally intended for participants to perform 18 trials per condition, which would have allowed two pairings of each can with each cone. However, one participant during our pilot trials took almost 75 minutes to complete the entire study and we found no statistical evidence to suggest that cone and can combinations were an effect on our dependent variables. Therefore, we settled on 14 trials per condition for the participants in our formal study. This created a more manageable experiment totaling, on average, 5 min administration + 5 min instructional video (AR) + 1 min stereo vision test + 3 min practice + (14 trials  $\times$  0.5 min assembly under AR) + 5 min break + 5 min instructional video (LCD) + 3 min practice + (14 trials  $\times$  1 min assembly under LCD)  $\approx$  50 minutes per participant. (Note: starting condition was counterbalanced across participants.)

### 4.4 Hypotheses and Testing

Based on our experience with the pilot study, and prior to our experiment, we formulated four hypotheses:

(H1) AR would be the fastest condition during the psychomotor phase of the task (step 5 in Table 1).

(H2) AR would be the most accurate condition during the psychomotor phase of the task.

(H3) AR would be the most preferred condition.

(H4) Participants would rank the AR condition as most intuitive.

These hypotheses were based on our belief that the dynamically tracked arrow, labels, and highlights would simplify the task of identifying and matching the attachment points on each cone and can. We expected participants to find the LCD condition more onerous because identifying the attachment points relies on visually matching salient features on the components.

To help test these hypotheses, we recruited 22 new participants (6 female; ages 18–44,  $\bar{X} = 26.3$ ) from the same target population

Activity Type	Mean Completion Time (s)		Comparison
	AR	LCD	
Locate	2.66	2.39	LCD 0.27 s faster than AR $t_{(21)} = 1.60, p = 0.124$
Position	1.15	1.15	LCD within 0.01 s of AR $t_{(21)} = 0.12, p = 0.905$
Align	24.23	45.55	AR 21.31 s faster than LCD $t_{(21)} = 6.27, p < 0.001$

Table 2: Pairwise comparisons of mean completion time by activity type.

as that of our pilot study. (None of the pilot study participants took part in the new study.) All but two participants used a computer multiple times per day. Two participants reported having some experience repairing mechanical systems, 11 reported having a basic level of exposure, and 9 reported having no experience. Six participants identified themselves as requiring corrective lenses and wore them in both conditions.. All 22 correctly perceived stereo stimuli. Each participant experienced the experiment design described in Section 4.2, with 14 trials for each condition, as finalized in our pilot study (Section 4.3).

## 5 RESULTS

We began our analysis by looking for potential outliers in our completion-time data. We identified suspicious values by examining Tukey plots of the completion times for each major activity type (Locate, Position, Align) across all trials and participants and cross-checked outlying values against videotaped footage of the participants performing the trials. This led us to establish the following ranges for valid completion times in each major activity:  $[0.25, 10]$  s for Locate activities (both conditions),  $[0.25, 10]$  s for Position activities (both conditions),  $[0.25, 60]$  s for AR Align activities, and  $[0.25, 120]$  s for LCD Align activities. We discovered that the predominant source of outliers was tracking errors produced when cones or cans were placed in particular arrangements relative to each other that occasionally formed a spurious constellation of retroreflective markers that the OptiTrack software mistakenly identified as a known rigid body. For example, when cans 1 and 2 are placed close to each other, and both are in a particular orientation, their combined visible configuration of markers may appear as an alias of can 3. This either caused the system to cue the next state prematurely or resulted in the participant stopping until the problem was corrected. We preprocessed our data by removing the outliers we identified, which accounted for 1.49% (1.62% AR; 1.36% LCD) of all calculated completion times across 3 activity types  $\times$  14 trials  $\times$  2 conditions.

### 5.1 Completion Time Analysis

We performed a 2 (Display Condition)  $\times$  3 (Activity Type) repeated measures ANOVA on mean completion times, with our participants as the random variable. Display condition exhibited a significant main effect on completion time ( $F_{(1,21)} = 37.09, p < 0.001$ ). The global mean completion time across all six individual activities was 9.38 s for AR and 16.36 s for LCD. A post-hoc comparison with Bonferroni correction ( $\alpha = 0.0125$ ) revealed that AR was 7.01 s faster than LCD, which was significant ( $p < 0.001$ ). Pairwise comparisons of mean completion times for each activity type are summarized in Table 2 and are depicted in the Tukey plots in Figure 8. As expected, only the Align activity exhibited a statistically significant difference in means, where the mean completion time for AR was 21.31 s (46.79%) faster than that of LCD. We did not expect a localization advantage for AR in this particular experiment because (1) we assigned each can and cone to a specific labeled bin on the workbench, facilitating memorization and (2) we used a small number of cans and cones.

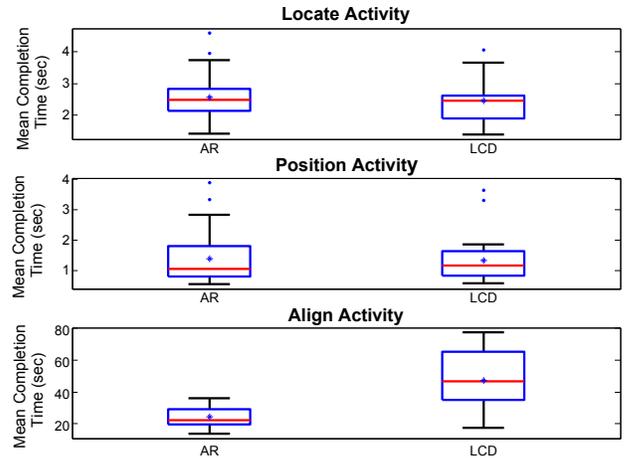


Figure 8: Activity completion times (s) for AR and LCD. An asterisk marks mean task completion for each condition.

### 5.2 Accuracy Analysis

We performed a two (Display Condition) repeated measure ANOVA on mean alignment error measured during the Align activity, as defined in Section 4.2, with our participants as the random variable. Display condition exhibited a significant main effect on alignment error ( $F_{(1,21)} = 48.75, p < 0.001$ ). The mean difference between the optimal orientation and that achieved by the user was 0.08 radians for AR (0.25 inter-hole widths) and 0.36 radians (1.15 inter-hole widths) for LCD and is depicted in the Tukey plot in Figure 9. A post-hoc comparison with Bonferroni correction ( $\alpha = 0.0125$ ) revealed AR was 0.28 radians more accurate than LCD, which was significant ( $p < 0.001$ ).

We also performed a binary accuracy check by counting the number of correct alignments between the can and cone at completion of the task. We defined a correct alignment as a displacement within 0.16 radians (0.5 inter-hole widths) as measured when the participant pressed the confirm button. A McNemar's chi-squared test with Bonferroni correction ( $\alpha = 0.0125$ ) revealed that there was a significant difference ( $\chi^2_{(1, N=303)} = 87.94, p < 0.001$ ) between the binary accuracy rates achieved under the AR and LCD conditions. The mean accuracy rate under the AR condition was 95.3%, compared to 61.7% under the LCD condition.

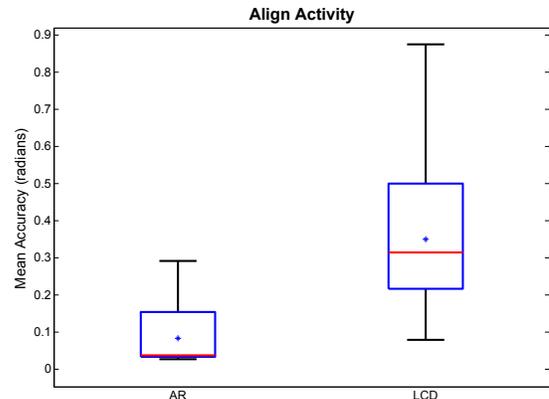


Figure 9: Mean alignment error (radians) for alignment activity in AR and LCD. An asterisk marks mean alignment error for each condition.

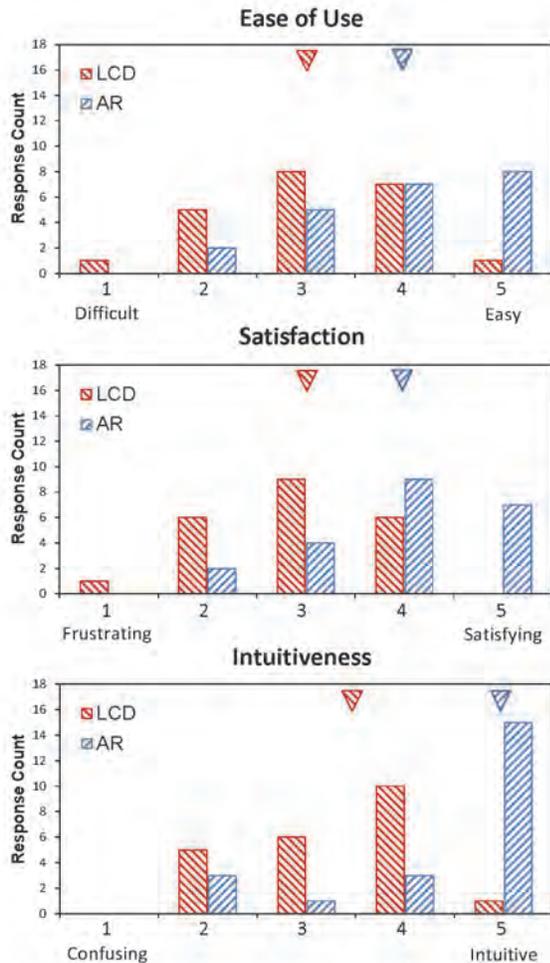


Figure 10: Survey response histograms by condition for ease of use (top), satisfaction (middle), and intuitiveness (bottom). Median values for each condition are displayed as triangles.

### 5.3 Questionnaire Analysis

Participants filled out a post-hoc questionnaire following their experience with both conditions. The questionnaire consisted of five-point Likert-scale questions (1 = most negative, 5 = most positive) to evaluate ease of use, satisfaction level, and intuitiveness for form of assistance. The results for the 22 participants in the formal study who experienced AR and LCD are summarized in Figure 10. Friedman tests revealed significant differences for ease of use ( $\chi^2_{(22,1)} = 8.00, p = 0.005$ ), satisfaction ( $\chi^2_{(22,1)} = 11.84, p = 0.001$ ), and intuitiveness ( $\chi^2_{(22,1)} = 9.80, p = 0.002$ ). Subsequent pairwise Wilcoxon comparisons of AR and LCD revealed AR was rated significantly better than LCD in terms of ease of use ( $p=0.007$ ), satisfaction ( $p=0.005$ ), and intuitiveness ( $p=0.012$ ).

When asked to rank the two forms of documentation based on preference for use, 20 of 22 participants ranked AR first. A Friedman test indicated this was significant ( $\chi^2_{(22,1)} = 11.64, p = 0.001$ ). When asked which form of documentation was the most intuitive, 19 of 22 participants ranked AR first. A Friedman test indicated this was also significant ( $\chi^2_{(22,1)} = 14.73, p < 0.001$ ).

## 6 COMPARISON TO PHYSICAL LABELS

Following up on our experiment, we were interested in how AR would compare to documentation that was physically embedded in the task; for example, by physically labeling or otherwise modify-

ing all components to clearly disambiguate them from each other and clearly distinguish the different ways that the components might be configured. This is a stated goal (although one that is often not achieved) for many manufacturers that develop products, such as furniture and toys, designed to be assembled by consumers. We considered implementing this as our original study baseline, but felt it was not ecologically valid for the combustion chambers and many other objects that may already have been designed and cannot be modified, that may have shapes and surface treatments dictated by other concerns, or that may be routinely subjected to extreme conditions that would damage or obscure superficial documentation. However, we decided that a pilot study comparing AR to such a baseline would be useful for situating our results relative to a broader range of task domains.

To accomplish this, we created a modified version of the LCD condition in which we printed and glued small physical labels to all possible connection points on each can and cone (the PRINTED condition). We also added virtual versions of these printed labels to the virtual models displayed on the LCD. Figure 11 depicts these modifications to the actual cone and can and to the displayed graphics (inset).

We recruited eight additional participants from the same target population, who experienced the same experiment design described in Section 4.2, with PRINTED substituted for LCD, and with the printed labels covered in AR. One participant failed to perceive stereo in the Stereo Fly test, and we excluded his data from our analysis. We also excluded a second participant's data after we noticed he failed to achieve the optimal alignment in all 14 trials of PRINTED. Observing this participant's video suggests he was simply placing the cone on the can, and then inserting pins based solely on the cone's labels (ignoring the can's labels). For the remaining six participants (all male, ages 19–27,  $\bar{X} = 23.5$ ), we performed a 2 (Display Condition)  $\times$  3 (Activity Type) repeated measures ANOVA on mean completion times, with the participants as the random variable. Display condition failed to show evidence of a significant main effect on completion time ( $F_{(1,5)} = 0.67, p = 0.451$ ). The mean completion time for all activities was 7.87 s for AR and 7.29 s for PRINTED. The mean completion times for the Align activity (step 5 in Table 1) were 20.73 s for AR and 19.42 s for PRINTED, and a difference in these means was not significant at the  $\alpha = 0.0125$  level. Our analysis also revealed that display condition failed to show evidence of a significant main effect on accuracy ( $F_{(1,5)} = 1.28, p = 0.31$ ). The mean angular error was 0.034 radians for AR and 0.065 radians for PRINTED, and a difference in these means was also not significant at the  $\alpha = 0.0125$  level. A binary accuracy check applying the same standards as defined in Section 5 revealed an accuracy rate of 100% for AR and 96.4% for PRINTED. A McNemar's chi-squared test with Bonferroni correction ( $\alpha = 0.0125$ ) did not detect a significant difference between rates ( $\chi^2_{(1,84)} = 1.33, p < 0.248$ ).



Figure 11: Components labeled for PRINTED condition. Inset shows graphics displayed on screen.

The six participants experiencing the PRINTED and AR conditions filled out the same post-hoc questionnaire. However, Friedman tests did not reveal significant rating differences between PRINTED and AR in the case of ease of use ( $\chi^2_{(6,1)} = 0.33, p = 0.564$ ), satisfaction ( $\chi^2_{(6,1)} = 0.11, p = 0.956$ ), or intuitiveness ( $\chi^2_{(6,1)} = 2.00, p = 0.157$ ). When asked to rank the two forms of documentation based on preference for use, 5 of 6 participants ranked AR first, and when asked which form of documentation was the most intuitive, 5 of 6 participants ranked AR first. However, a Friedman test indicated that these ranking results were not statistically significant ( $\chi^2_{(6,1)} = 2.667, p = 0.102$ ).

## 7 DISCUSSION AND FUTURE WORK

Our prototype was able to help participants perform the alignment, or psychomotor portion, of the combustion chamber assembly task more quickly in AR than when using the LCD-based documentation. Furthermore, a statistically significant number of participants selected the AR condition as the most preferred and most intuitive condition. That participants overwhelmingly preferred AR was especially encouraging, given that the HWD used in the AR condition (Figure 5) weighs 1.3 kg (not including the added tracking hardware), and was relatively bulky, in comparison with the 205 g tracking crown used in the LCD condition. Notably, our prototype performed better than an enhanced baseline which, in turn, is better than industry-standard IETMs, because our enhanced baseline tracks the user's activity and automatically advances to the next step in the procedure.

We were also pleased that our AR prototype allowed participants to complete the alignment step more accurately than when using the LCD condition. We were somewhat surprised by the much lower accuracy rate of participants in the LCD condition (61.7%). We suspect our instructional video, which asked participants to “work as quickly and accurately as possible”, may have biased participants toward speed. However, we believe any stronger emphasis on accuracy would have resulted in even slower completion times in the LCD condition.

While it is difficult to generalize the findings from our follow-up comparison of AR to the idealized PRINTED condition, given our small sample size, we were encouraged by the lack of statistical evidence to suggest a large difference in performance.

Our prototype was designed to facilitate initial exploration into how AR might improve performance of psychomotor tasks. Having established evidence for such improvement, we see opportunities to extend and generalize our results. First, there is a need to formally explore the specific types of assistance provided by AR during psychomotor activities. The selection of techniques featured in our prototype and described in Section 3.1.2 (e.g., dynamic 3D arrows) was guided by established design heuristics, pilot testing, and intuition. However, we envision a separate set of studies that explore these techniques alone and in combination, enumerating their parameters, and trying to determine the influence that each has on task performance. The end result of these efforts would produce a parameterized catalog of validated AR techniques for all major activities in a procedural task taxonomy.

Second, we are interested in addressing other realistic maintenance and repair tasks. We chose a single parameterized task involving assembly of turboprop engine combustion chambers for our study, and believe that this is representative of a wide range of real tasks that require assembling two rigid bodies that must be precisely rotated about a single axis. We are currently exploring more complex psychomotor alignments involving multiple axes of rotation, such as the one depicted in Figure 12. This task uses an animated arrow and oriented labels to assist the user in proper installation of an assembled combustion chamber on the Dart engine. A successful installation requires careful 3DOF alignment between



Figure 12: A different alignment task with AR assistance. (Captured by a video camera looking through the optical-see-through display. A post-render filter was applied to remove camera distortion and vignetting.)

the combustion chamber and multiple, adjoining components of the engine.

## 8 CONCLUSIONS

We presented an AR prototype for providing assistance during procedural tasks, with an emphasis on applying AR to support the psychomotor phases of these tasks. We applied our prototype to a realistic assembly task encountered in a manufacturing and maintenance domain, and ran a counterbalanced, within-subject, user study comparing the AR prototype with 3D-graphics-based documentation presented on a stationary display. The results of the experiment confirmed that AR was faster and more accurate for psychomotor phase activities, was overwhelmingly preferred by participants, and was considered to be more intuitive, despite the relatively bulky HWD that we used. A small, follow-on pilot study comparing our prototype to an idealized, but often impractical, form of documentation featuring physical labels affixed to components, revealed no statistically significant differences in speed or accuracy. There is still much to be done to address how AR might be used in the wide range of psychomotor activities encountered in our lives. We hope that our work will inspire future research on AR for task assistance, and future commercial implementations, as hardware becomes less bulky and more affordable.

## ACKNOWLEDGMENTS

This research was funded in part by NSF Grant IIS-0905569, ONR Grant N00014-04-1-0005, and generous gifts from VTT and Vuzix.

## REFERENCES

- [1] AVILUS. <http://www.avilus.de>, Sept. 2011.
- [2] M. Agrawala, D. Phan, J. Heiser, J. Haymaker, J. Klingner, P. Hanrahan, and B. Tversky. Designing effective step-by-step assembly instructions. *ACM Trans. on Graphics*, 22(3):828–837, 2003.
- [3] ALVAR. <http://virtual.vtt.fi/virtual/proj2/multimedia/alvar.html>, Jan. 2011.
- [4] S. Antifakos, F. Michahelles, and B. Schiele. Proactive instructions for furniture assembly. In *Proc. Int. Conf. Ubiquitous Computing*, pages 351–360, 2002.
- [5] ARTESAS. ARTESAS-Advanced augmented reality technologies for industrial service applications. <http://www.artesas.de>, May 2011.

- [6] R. Azuma. A survey of augmented reality. *Presence-Teleoperators and Virtual Environments*, 6(4):355–385, 1997.
- [7] I. Barakonyi and D. Schmalstieg. Augmented reality agents in the development pipeline of computer entertainment. In *Proc. Int. Conf. on Entertainment Computing (ICEC '05)*, pages 345–356, 2005.
- [8] T. Blum, T. Sielhorst, and N. Navab. *Advanced Augmented Reality Feedback for Teaching 3D Tool Manipulation*, chapter 25, pages 223–236. Lupiensis Biomedical Publications, 2007.
- [9] H. R. Booher. Relative comprehensibility of pictorial information and printed words in proceduralized instructions. *Human Factors*, 17:266–277, 1975.
- [10] O. Cakmakci, F. Bèrard, and J. Coutaz. An augmented reality based learning assistant for electric bass guitar. In *Proc. Int. Conf. on Human-Computer Interaction*, 2003.
- [11] T. P. Caudell and D. W. Mizell. Augmented reality: An application of Heads-up display technology to manual manufacturing processes. In *Int. Conf. on System Sciences*, pages 73–80, 1992.
- [12] D. Curtis, D. Mizell, P. Gruenbaum, and A. Janin. Several devils in the details: Making an AR application work in the airplane factory. In *Proc. Int. Workshop on Augmented Reality (IWAR '98)*, pages 47–60, 1999.
- [13] F. Echter, F. Sturm, K. Kindermann, G. Klinker, J. Stilla, J. Trilk, and H. Najafi. The intelligent welding gun: Augmented reality for experimental vehicle construction, 2003.
- [14] S. Feiner. APEX: An experiment in the automated creation of pictorial explanations. *IEEE Computer Graphics and Applications*, 5(11):29–37, 1985.
- [15] S. Feiner, B. MacIntyre, and D. Seligmann. Knowledge-based augmented reality. *Communications of the ACM*, 36(7):53–62, 1993.
- [16] W. Friedrich. ARVIKA-Augmented reality for development, production and service. In *Proc. Int. Symp. on Mixed and Augmented Reality (ISMAR '02)*, pages 3–4, 2002.
- [17] A. Fuhrmann, R. Splechtna, and J. Prikryl. Comprehensive calibration and registration procedures for augmented reality. In *Proc. Eurographics Workshop on Virtual Environments*, pages 219–228, 2001.
- [18] R. M. Gagné and W. D. Rohwer. Instructional psychology. *Annual Review of Psychology*, 20(1):381–418, 1969.
- [19] M. Gandy, B. MacIntyre, P. Presti, S. Dow, J. Bolter, B. Yarbrough, and N. O'Rear. AR Karaoke: Acting in your favorite scenes. In *Proc. Int. Symp. on Mixed and Augmented Reality (ISMAR '05)*, pages 114–117, 2005.
- [20] F. B. Gilbreth and L. M. Gilbreth. Classifying the elements of work. *Management and Administration*, 8(2):151–154, 1924.
- [21] GoblinXNA. <http://www.codeplex.com/goblinxna>, May 2011.
- [22] J. Heiser, D. Phan, M. Agrawala, B. Tversky, and P. Hanrahan. Identification and validation of cognitive design principles for automated generation of assembly instructions. In *Proc. Working Conf. on Advanced Visual Interfaces*, pages 311–319, 2004.
- [23] S. Henderson and S. Feiner. Evaluating the benefits of augmented reality for task localization in maintenance of an armored personnel carrier turret. In *Proc. Int. Symp. on Mixed and Augmented Reality (ISMAR '09)*, pages 135–144, Oct. 2009.
- [24] S. Henderson and S. Feiner. Exploring the benefits of augmented reality documentation for maintenance and repair. *IEEE Transactions on Visualization and Computer Graphics*, 17(10):1355–1368, Oct. 2011.
- [25] C. Knöpfle, J. Weidenhausen, L. Chauvigné, and I. Stock. Template based authoring for AR based service scenarios. In *Proc. IEEE Virtual Reality (VR '05)*, pages 249–252, 2005.
- [26] W. Li, M. Agrawala, B. Curless, and D. Salesin. Automated generation of interactive 3d exploded view diagrams. *ACM Trans. Graph.*, 27:101:1–101:7, August 2008.
- [27] W. Li, M. Agrawala, and D. Salesin. Interactive image-based exploded view diagrams. In *Proc. Conf. on Graphics Interface*, pages 203–212, 2004.
- [28] W. Li, L. Ritter, M. Agrawala, B. Curless, and D. Salesin. Interactive cutaway illustrations of complex 3d models. *ACM Trans. Graph.*, 26, July 2007.
- [29] Y. Motokawa and H. Saito. Support system for guitar playing using augmented reality display. In *Proc. Int. Symp. on Mixed and Augmented Reality (ISMAR '06)*, pages 243–244, 2006.
- [30] U. Neumann and A. Majoros. Cognitive, performance, and systems issues for augmented reality applications in manufacturing and maintenance. In *Proc. IEEE Virtual Reality (VR '98)*, pages 4–11, 1998.
- [31] J. Platonov, H. Heibel, P. Meier, and B. Grollmann. A mobile markerless AR system for maintenance and repair. In *Proc. Int. Symp. on Mixed and Augmented Reality (ISMAR '06)*, pages 105–108, 2006.
- [32] J. Quarles, S. Lampotang, I. Fischler, P. Fishwick, and B. Lok. Tangible user interfaces compensate for low spatial cognition. In *Proc. 3DUI*, pages 11–18, 2008.
- [33] A. Raczynski and P. Gussmann. Services and training through augmented reality. In *Proc. European Conf. on Visual Media Production (CVMP '04)*, pages 263–271, 2004.
- [34] D. Reiners, D. Stricker, G. Klinker, and S. Müller. Augmented reality for construction tasks: Doorlock assembly. In *Proc. Int. Workshop on Augmented Reality (IWAR '98)*, pages 31–46, 1999.
- [35] M. Richardson, G. Jones, and M. Torrance. Identifying the task variables that influence perceived object assembly complexity. *Ergonomics*, 47(9):945–964, 2004.
- [36] C. M. Robertson, B. MacIntyre, and B. N. Walker. An evaluation of graphical context when the graphics are outside of the task area. In *Proc. Int. Symp. on Mixed and Augmented Reality (ISMAR '08)*, pages 73–76, 2008.
- [37] M. Rosenthal, A. State, J. Lee, G. Hirota, J. Ackerman, K. Keller, E. D. Pisano, M. Jiroutek, K. Muller, and H. Fuchs. Augmented reality guidance for needle biopsies: An initial randomized, controlled trial in phantoms. *Medical Image Analysis*, 6(3):313–320, 2002.
- [38] T. Salonen and J. Sääski. Dynamic and visual assembly instruction for configurable products using augmented reality techniques. In X.-T. Yan, C. Jiang, and B. Eynard, editors, *Advanced Design and Manufacture to Gain a Competitive Edge*, pages 23–32. Springer London, 2008.
- [39] B. Schwedtfeger, R. Reif, W. A. Gunthner, G. Klinker, D. Hamacher, L. Schega, I. Bockelmann, F. Doil, and J. Tumler. Pick-by-Vision: A first stress test. In *Proc. Int. Symp. on Mixed and Augmented Reality (ISMAR '09)*, pages 115–124, 2009.
- [40] D. D. Seligmann and S. Feiner. Automated generation of intent-based 3D illustrations. In *Proc. Conf. on Computer Graphics and Interactive Techniques*, pages 123–132, 1991.
- [41] A. State, M. A. Livingston, W. F. Garrett, G. Hirota, M. C. Whitton, E. D. Pisano, and H. Fuchs. Technologies for augmented reality systems: Realizing ultrasound-guided needle biopsies. In *Proc. Conf. on Computer Graphics and Interactive Techniques*, pages 439–446, 1996.
- [42] A. Tang, C. Owen, F. Biocca, and W. Mou. Comparative effectiveness of augmented reality in object assembly. In *Proc. SIGCHI Conf. on Human Factors in Computing Systems (CHI '03)*, pages 73–80, 2003.
- [43] D. Tennenhouse. Proactive computing. *Commun. ACM*, 43:43–50, May 2000.
- [44] B. Tversky, J. B. Morrison, and M. Betrancourt. Animation: Can it facilitate? *International Journal Of Human Computer Studies*, 57(4):247–262, 2002.
- [45] US Department of Defense. *Interactive Electronic Technical Manuals-General Content, Style, Format and User-Interaction Requirements (MIL-DTL-87268C)*. 2007.
- [46] F. K. Wacker, S. Vogt, A. Khamene, J. A. Jesberger, S. G. Nour, D. R. Elgort, F. Sauer, J. L. Duerk, and J. S. Lewin. An Augmented Reality System for MR Imageguided Needle Biopsy: Initial Results in a Swine Model1. *Radiology*, 238(2):497–504, Feb 2006.
- [47] J. Zauner, M. Haller, A. Brandl, and W. Hartman. Authoring of a mixed reality assembly instructor for hierarchical structures. In *Proc. Int. Symp. on Mixed and Augmented Reality (ISMAR '03)*, pages 237–246, 2003.