# Augmented Reality in Computer Science Education: Exploring the Benefits of 6DOF Augmented Reality for Teaching Inheritance in Object Oriented Programming

Mark Steelman*


## Abstract

Object oriented programming is an important set of concepts taught early on in computer science education, and many of the examples used to explain object-oriented programming involve real world examples with tangible objects. These examples are often just explained in writing and quickly become abstracted and shown as lines of code, making it hard to grasp these concepts well. Augmented reality (AR) has the potential to make these examples more intuitive by actually visualizing real world examples with 3D objects as if they exist in a students real world environment. In this study we compare the effectiveness of a 6DOF AR version of a tutorial on the concept of inheritance in object oriented programming with an equivalent desktop computer version of the same tutorial. Our hypothesis was that subjects that go through the AR version of the tutorial would be able to recall the contents of the tutorial better than subjects who do the desktop computer version, and that subjects would enjoy the AR version more. We collected both quantitative and qualitative data on a total of 18 subjects split between two groups and found that on average subjects performed better in recalling content from the tutorial with the AR version over the desktop computer version. Also, all subjects liked the AR version more for a variety of reasons including that it was more interactive, more fun, and could provide benefits that standard learning mediums could not. Thus, the qualitative and quantitative results support our hypothesis, although there are still disadvantages to the AR medium and we discuss how this use of AR can be improved for future use.


## 1 Introduction

Augmented reality is an inherently engaging and interactive medium and has the potential to significantly improve the learning experience for students. Technology is also becoming increasingly important and integrated into people's daily lives, and as a result the relevance of understanding technology and topics like computer science is also increasing. Despite its increasing relevance, computer science is a complicated subject and many people find it too complex and boring to be worth studying. Given the potential of AR as an educational tool, we would like to see if AR might be able to make computer science education more approachable by making it more fun and easier to learn. So far, there are few educational tools that provide an interactive and engaging experience, and most of them are in 2D platforms which inherently limits their potential interactivity and induced focus. For almost every subject in education, teachers teach the same content in the same medium year after year and we believe that it would be more effective to add new mediums such as AR so that students who struggle with a given subjects have more options of how they can learn that subject. There are many unique affordances of AR that provide the option for experiential and more natural learning processes. There are also an increasing number of online learning platforms available today, and AR could potentially be added to these ecosystems and make use of features such as multiplayer capabilities so that you could have many students all present in the same educational experience from anywhere in the world. We believe that gathering evidence for the baseline capabilities of AR for improving education quality is essential to eventually building a robust educational platform that incorporates technology such as AR and VR.

Object oriented programming is a core aspect of many coding languages, and it is one of the first set of concepts taught to students taking introductory level computer science courses. Since object oriented programming is so important and often involves real world examples with tangible objects to explain object oriented programming concepts, we feel that AR fits in well to teach these concepts as it allows for digital 3D objects to be viewed and interacted with in a student's real world environment. Of the core concepts in object oriented programming, the concept of inheritance and the common examples used to explain inheritance with hierarchies of objects was easy to translate into an 3D AR explanation. The purpose of this study is not only to evaluate AR as a general educational tool, but also to evaluate AR as a medium that can be used to teach computer science. As discussed, computer science education is becoming increasingly relevant and important, so the results of this research could have a significant impact on the field of computer science education.

## 2 Literature Review

There have been many research projects done that explore the effects of using Virtual Environments (VE's) as a learning tool. We consider the AR medium to be a form of a virtual environment as it involves 3D virtual objects in a user's real world environment, but can also augment that real world environment to blend the virtual and real worlds. In regards to education, the range of the uses of virtual environments have varied greatly. For example, there have been projects that have explored the effects of interactivity in virtual environments within the education space [5]. With Hussein and Natterdal's study, it relied on qualitative methods to form conclusions. Important findings included realizing that virtual environments allowed participants to remain more focused because they felt immersed and present in the virtual world as compared to the real world (where participants felt that they could be easily distracted by the world around them). However, another factor to consider for the design of our study is the use of text since text is used as an aid to the audio clips (as described in section 4). Hussein and Natterdal concluded that text was harder to read in a virtual application and that a subject with lots of text to read would be most suitable for 2D interfaces, while subjects that require simulations and 3D representations would be a better fit for virtual environments. As our study primarily focuses on learning through audio narration and visual examples, difficulties with text reading should not make a significant impact. There are also other projects that provide interesting insight into techniques that could potentially provide effective 3D learning. Ragan, Sowndararajan, Kopper, and Bowman [9] discuss the values and limitations of supplementing conceptual information with spatial information in educational virtual environments. Participants

*e-mail:mark.steelman@duke.edu

in their study memorized procedures in a virtual environment and then attempted to recall those procedures. Their study was helpful for detailing ways in which learning within virtual environments could potentially be tested. Other studies focused on memorization in the context of spatial awareness within virtual environments, specifically how virtual environments aid spatial memory recollection [1]. The conclusions from this study indicated that spatial and object memory resulting from interactions with the virtual environment were enhanced only for aspects of the environment directly involved in the interaction, specifically, the spatial layout through which participants were required to navigate. An important aspect to consider with memory testing is time allocation for recollection, as it plays a huge factor in the results. In the Brooks, Attree, Rose, Clifford, Leadbetter study [1], participants were given five minutes to draw their recollection of the layout of the virtual reality rooms. This study also indicated that enhanced memory recollection only occurred in the participants that were active and navigated the virtual environment with the joystick, and had negligible effects on passive participants who observed the actions of the active participants.

An essential component to the formation of our research topic is that AR technology has the power to facilitate learning. While the two technologies are not the same, both virtual reality (VR) and AR technologies make use of interaction with digital 3D objects, and for that reason prior research related to VR in education is relevant to our study. Studies have shown that the application of VR technology in education enriches teaching and learning in the current education model. Virtual learning environments (VLEs) improve students abilities to analyze problems and explore new concepts. When immersive, interactive, and imaginational aspects are also integrated, the virtual learning space allows learners to model, act, and express anything they want as long as the environment provides the tools to do so [7]. This study defines four principle components of VLEs: Knowledge Space, Communication Community, Active Action, and Facility Toolkit. These four principles will be helpful in guiding us in the direction of creating a holistic teaching and learning tool that follows the guidelines of virtual learning environments. Furthermore, this study provides insight into interaction enhancing active learning, which is consistent with other studies that have been conducted within the virtual environment space. From these studies, strong conclusions can be drawn that interaction in virtual environments is necessary for learning in VLEs to be successful. Interaction can come in a variety of forms, including more passive interaction such as the inclusion of additional sensory inputs/outputs that have been shown in other studies outside of virtual reality to improve memory as a result of the emotional and other psychological states that other senses trigger. Dinh, Walker, Song, Kobayashi, and Hodges [4] investigated the effects of tactile, olfactory, and audio sensory cues on the sense of presence and memory in a virtual environment, finding a positive correlation between increased sensory modalities and the users sense of presence and memory for objects in the VE. Interestingly, while increasing the fidelity of the three aforementioned senses increased memory and presence, increasing visual detail did not correspond with increased memory or presence. There have been previous predictions of such observations (such as Ivan Sutherlands) about visual realism and that it would not necessarily contribute to increased presence, so we will probably not focus too heavily on the visual quality of our experimental design as long as it is minimally convincing.

Another important aspect crucial to the success of our study is how information within a virtual environment will be presented. As we would like to explore VEs within any educational context, the aim is to be able to use this virtual environment to help facilitate learning of abstract or complex symbolic concepts (potentially related to the space). Bowman, Hodges, Allison, Wineman [2] conducted a study to compare learning between traditional lectures and classroom material enhanced with the use of a virtual environment in the context of a zoo exhibit. The participants were given tests on material relating to specific information about the design of the zoo exhibit habitat. Students who had their learning augmented by the virtual environment performed better on the test, suggesting that students were able to draw on their experiences in VEs in educational settings where virtual environments were used to augment learning. From studies like this one, it is intuitive that virtual environments are better at creating associations between spatial and abstract information, and they add a strong experiential component to educational settings. An important factor in this study was the time after the study in which the test was given - students were unaware that they were participating in an experiment and that they would be tested, and the test was given five days after students were presented the content so that they could internalize the material as they would in a classroom setting, and so that too short of a time post the experiment wouldnt serve as a confounding variable.

As designers of the AR interface that will be used for our study, we recognize that the success of the study also depends on the usability of the AR interface. Wickens [10] argues that the naturalness of the interface must be increased to reduce the cognitive effort required by the user to navigate and interpret the VE. By aspiring to make the interface as natural as possible, we avoid the possibility of participants experiencing the cognitive overload of processing new material while learning how to navigate an augmented environment in unfamiliar ways. As Wickens points out, we want to mitigate any distractions from learning due to usability issues of the interface.

Since part the motivation for our research is to work towards a widely applicable educational tool that uses AR to enhance any and all relevant topics in education, it is important that we consider the attributes of VLEs in relationship to educational theory and pedagogical practice. Bricken [3] discusses several key aspects of VR that have inherent ties to educational theory, including the experiential nature of VR, the natural interaction with information that is afforded, the possibility of shared experiences, and the flexibility of VR to be tailored to individuals. In education there are almost always multiple paths to achieving learning outcomes, but usually certain paths are more efficient and effective than others, and AR may be able to assist in providing access to those better paths when otherwise unavailable. For instance, there are many situations in which abstract methods are used to describe or teach inherently experiential or physically interactive concepts like architectural design where there are usually insufficient resources to learn this through non-abstract methods. Many educational institutions struggle to provide sufficient individual attention to students due to the lack of teachers or resources, but since AR can be used to collect individual data that can be analyzed, the barrier to individual tailorship can be significantly reduced. However, despite the many affordances of AR, there are still various challenges such as information accuracy (sometimes difficult to translate content accurately into AR), fear of overburdening the teachers, lack of historical perspective, and other issues that have yet to even be discovered.

There has also been research showing the affective outcomes of using AR for STEM education. Ibez and Delgado-Kloos [6] found in their meta-analysis of studies related to AR's use in STEM education that some of the most frequent outcomes were motivation, attitude, enjoyment, engagement, and interest. This further supports our intuition that AR is an appropriate medium to improve the quality of computer science education. In order to come up with a tutorial for teaching the concept of inheritance in object oriented programming in a way that is purely conceptual and involves an example with real world objects so that both the AR and desktop computer versions of the tutorial can be created with the same content, we looked for well reviewed conceptual explanations of object oriented programming and inheritance. We found a simple but effective article that gives this explanation and used this as a basis for the creation of our tutorial [8].

## 3 Hypothesis

Our overall hypothesis was two fold: subjects that completed the AR version of the inheritance tutorial would be able to better recall the contents of the tutorial, and at the same time they would find the tutorial more enjoyable and preferable. The first part of this hypothesis was influenced by the fact that the experiential nature of AR and the natural interaction with information that is afforded provide a solid foundation for using AR as an educational tool [3]. The content of the tutorials is also the same, so the difference between the two versions of the tutorial are just the interaction interface, and since the AR interface requires more active and involved interactions, we would assume that the subject is more actively involved with the AR version of the tutorial. The second part of this hypothesis was formed because we knew that prior studies with AR have shown positive sentiments related to the use of AR devices [8]. Our study involves the use of an AR device of much higher quality and fidelity, so we expected these sentiments to be even stronger. Also, since the device that we used in the study is relatively new, we expected that this would be a new experience for the majority of our subjects which itself is likely to contribute to the subjects enjoying and being fascinated by the AR version.

## 4 Experimental Design

In this experiment, there were two groups of participants: the group that went through the AR version of the inheritance tutorial first and the group that went through the desktop computer version first. For the sake of simplicity, we will call these groups the AR group and the desktop group respectively. Participants were randomly placed into one of the two groups. Before beginning the tutorial, subjects from both groups read and signed an informed consent form and also completed a background information questionnaire that assesses their experience with video games and virtual/augmented reality environments. Both versions of the tutorials have the same 12 steps, but the audio clips, visualizations, and interaction methods are modified slightly to adapt appropriately to the device. These modifications do not affect the content of the tutorial, which is the same across both versions.

After completing the consent form and background questionnaire, the subjects from the AR group were given a training phase to get used to the AR device and the controls necessary for completing the AR version of the inheritance tutorial. Once ready, the subjects from the AR group then went through the AR version of the tutorial and upon completion were asked to sit down and answer five interview questions to see how well they understand the parts of the tutorial. Then they were given a subjective feedback questionnaire to complete that asked them about their feelings about the AR tutorial. Next, the subjects from the AR group went through the desktop computer version of the tutorial and upon completion were given another subjective feedback questionnaire to complete that asked them questions about comparing the two versions of the tutorial. After this the AR group was done with the experiment.

For the desktop group, after completing the consent form and background questionnaire, the subjects were not given a training phase as it was assumed that they were familiar with the standard computer device and immediately went through the desktop computer version of the tutorial. Upon completion the subjects from the desktop group were asked to answer five interview questions to see how well they understand the parts of the tutorial. Then they were given a subjective feedback questionnaire to complete that asked them about their feelings about the desktop computer tutorial. Next, the subjects from the desktop group went through the AR version of the tutorial and upon completion were given another subjective feedback questionnaire to complete that asked them questions about comparing the two versions of the tutorial. After this the desktop group was done with the experiment.

### 4.1 Participants

18 voluntary, unpaid participants were recruited in total, out of which 11 were male and 7 were female. The mean age of the participants was 21 years old. From the background survey the participants completed prior to the experiment, 14 of the participants had previous experience in a virtual environment using a head-mounted device. The participants came from a variety of different educational backgrounds and levels ranging from public policy to biology, and all were undergraduate students except for one graduate student. None of the participants had any experience with computer science education or stereoscopic AR glasses.

### 4.2 Experimental Procedure

In this section we will describe in more detail the experimental procedure for both the AR group and the desktop group. For the AR group, first they go through a training phase to get used to the AR device. This training phase involves showing the subjects from the AR group how to put on the headset correctly, how to hold the controller, how to reach out and select/grab an object with the controller using the bumper button, and how to deselect/drop an object with the controller using the bumper button. Then they begin the tutorial by grabbing the digital blue cube on the table (Figure 1) next to the set of 3D printed tool objects and dropping the blue cube onto the pink platform on the ground behind them (Figure 2). This is a simple way to add 3D interaction to the procedure so that they are getting their body and senses involved. This causes the first powerpoint slide (Figure 3) to float upwards from the ground where the blue cube was dropped and the audio clip narration for step 1 begins (see Appendix for all slides and audio clip scripts). At the end of audio clip for step 1, the subject is instructed to pick up the blue cube again from the same spot on the table and drop it on the next pink platform to move on to the next step. This process repeats for the first 5 steps which include the 5 slides to go along with the audio narration. However, the pink platform is in a different place at each of these first 5 steps so that the subject is forced to move around the room a little more, adding to the interaction and engagement. Also, on step 3, the subject is prompted to look at the variables and methods of the tool class floating above the 3D printed tool handle on the table (Figure 4). At the end of step 5, the subject is prompted to move on to the tool example with the set of 3D printed tool objects on the table. Steps 6-12 guide the subject through grabbing the tool objects and creating different child classes of the tool class (see Appendix for pictures of these steps). Then the subject is instructed to return the headset and controller to the experiment supervisor and finishes the remaining parts of the experiment that we previously described (interview, subjective feedback questionnaire 1, desktop computer tutorial, and subjective feedback questionnaire 2).

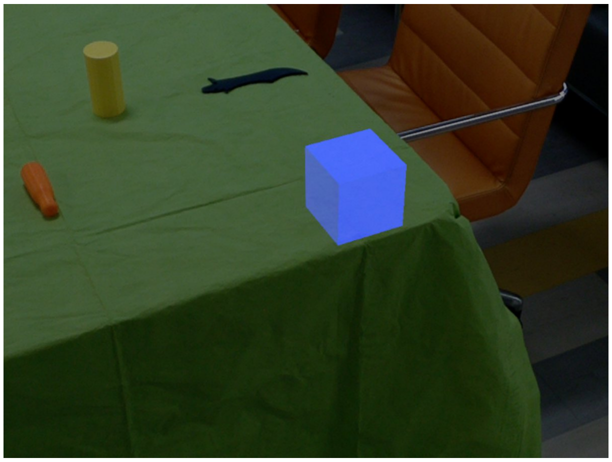

Figure 1: Blue cube that subject grabs to drop on pink platform

Figure 2: Pink platform that subject drops blue cube on to cause the next slide to appear

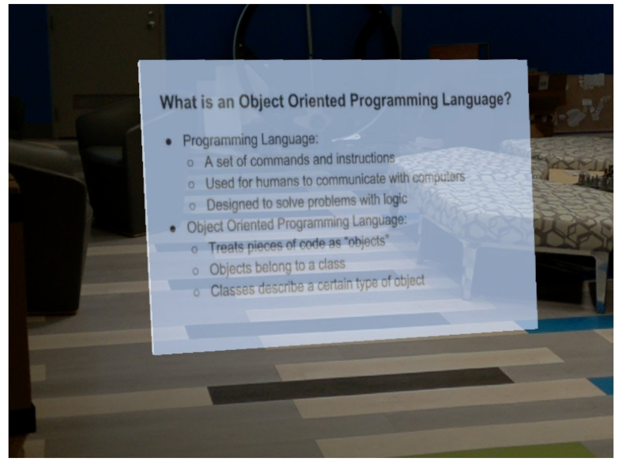



Figure 3: Slide that floats upwards from the ground where the blue cube was dropped

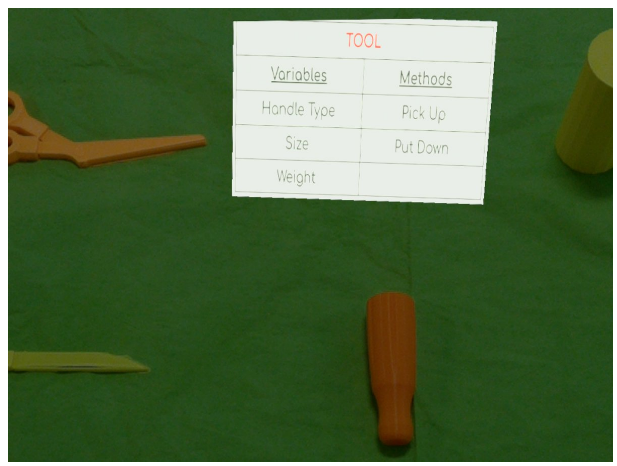



Figure 4: Variables and methods of the tool class floating above the 3D printed tool handle

For the desktop group, they begin the tutorial by clicking the start tutorial button which causes the audio clip for step 1 to begin and the first slide to appear (see Appendix for all slides and audio clip scripts). Just like the AR version of the tutorial, the first 5 steps all have slides to go along with the audio clip. For these 5 steps, there is a common interface (Figure 5) that the subject use to go onto the next steps with a next step button that the subject is prompted to click at the end of the audio clip for each step. For steps 6-12, there is the tool example interface (Figure 6) that the subject uses to go through creating copies of the tool objects and creating different child classes of the tool class (see Appendix for audio clip scripts for these steps that describes what the subject is instructed to do at steps 6-12). Then the subject finishes the remaining parts of the experiment that we previously described (interview, subjective feedback questionnaire 1, AR tutorial, and subjective feedback questionnaire 2).

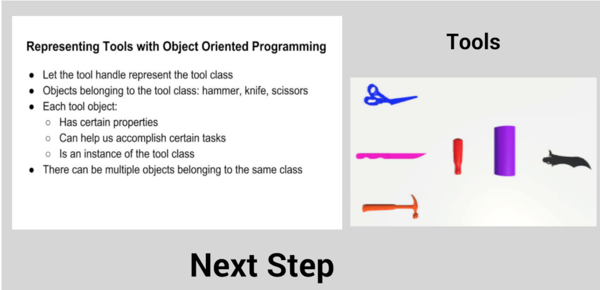


Figure 5: Slide interface for steps 1-5 of the desktop computer tutorial version

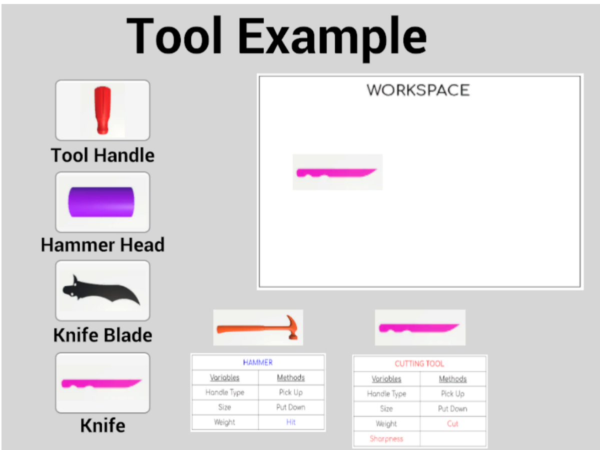


Figure 6: Tool example interface for steps 6-12 of the desktop computer tutorial version

For both the AR group and the desktop group, the interview questions and subjective feedback questionnaires were exactly the same so that we could compare the data between the two groups. There were 5 questions for the interview. Question 1: What is a programming language? Question 2: What is an object oriented programming language? Question 3: What are the features of an object in an object oriented programming language? Question 4: Why is it useful to use the principle of inheritance in object oriented programming? Question 5: How can we use inheritance to create new classes of objects? For all 5 questions, there was a rubric (see Appendix for questions and rubrics) for giving a score between 0-3 that assessed how well the subject was able to answer the question based on the content of the tutorial. The reason we decided to use an interview with open ended questions as our method for quantitatively assessing the participants' ability to recall content from the tutorial was because it allowed us to collect quantitative data while making sure that there was not a ceiling effect that would be possible had we given a multiple choice test since the content was so fresh on their mind. We felt that this would also make participants more likely to give extensive answers than if they were just to fill out short answer questions on paper.

The first subjective feedback questionnaire asks the subject 5 questions using a 1-7 Likert scale to assess how they felt about the first tutorial they completed (see Appendix for subjective feedback questionnaire questions). After those 5 questions there is a space for subjects to leave any comments or questions they had. This is used to see how the subject felt about the tutorial before they have had a chance to go through the other version of the tutorial. The second subjective feedback questionnaire asks the subject 4 open ended questions about both versions of the tutorial since they will have gone through both versions of the tutorial before they answer this second questionnaire (see Appendix for subjective feedback questionnaire questions). This questionnaire is used to see what the subject thinks about the two versions of the tutorial in hindsight of knowing what both are like. Both of these subjective feedback questionnaires are useful for us to understanding the participants' experience of the study and to provide some reasoning for the quantitative results.

### 4.3 Apparatus

In our experiment, for the AR version of the tutorial we used the Magic Leap One augmented reality glasses and the Magic Leap Control (the controller) that provided participants with the ability to see and interact with 3D digital objects in the study environment such as digital replicas of the 3D printed tool objects laid out on the table (Figure 7). The headset has integrated speakers for each ear that allows the user to hear audio input. For the desktop computer version of the tutorial we used a laptop. Both versions of the tutorial were created using Unity, with 3D models downloaded from Thingiverse. Those same 3D models were used to make 3D prints of the tools that were put on the table for the AR tutorial. Due to the irregularity of the tracking of the real world environment on the Magic Leap device, we put a green paper table cloth over the table so that the headset was able to consistently track the table since the table cloth was not reflective and reflection without the table cloth caused the headset to improperly spatially map the table.

Figure 7: Magic Leap One apparatus and table with 3D printed tool objects

## 5 Results

We performed an analysis of qualitative and quantitative results from the interview questions and subjective feedback questionnaires that participants were asked to complete. The following sections summarize the data and information collected from the 18 participants that completed the experiment.

### 5.1 Quantitative Results

As previously mentioned, we had 2 groups of subjects, the AR group and the desktop group, each of which answered the same set of interview questions after completing the AR version of the tutorial and desktop computer version of the tutorial respectively. We calculated the average score for each of the 5 questions (out of 3) as well as the average total score for all 5 questions (Figure 8). These values are the primary performance measure that we used to assess the participants' ability to recall the essential content from the tutorial on object oriented programming and inheritance. Per the first part of our hypothesis, we expected to see higher scores for the AR group than the desktop group.

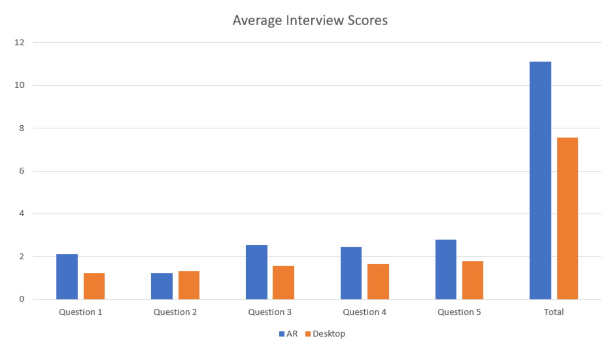


Figure 8: Average group scores for interview questions

Using these average score values for each group, we also calculated the standard deviation for each of the 5 questions to see how much variability there was in the performance of subjects in each group. Given the small range of possible scores, we expected the standard deviations to be small and not provide much insight, but in the case of a higher standard deviation that could tell us what parts of the tutorial gave some subjects more trouble than others.

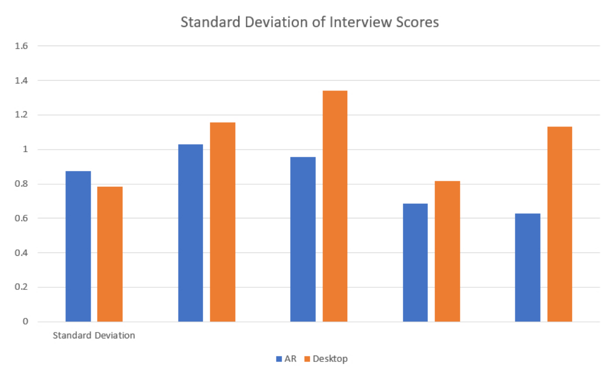


Figure 9: Standard deviation of scores for interview questions per group

Since one goal of our study is to see how the AR version of the tutorial can potentially improve recall ability over the desktop computer version, analyzing the score differences between the different groups is essential to evaluating our hypothesis. We also tracked the length of the interviews to see if there was any difference between the two groups in terms of how long it took for subjects to answer the interview questions, but did not expect a significant difference since the amount of content being asked about was small. We calculated the average interview times for each group. For the AR group, the average interview time was about 109 seconds, and for the desktop group the average interview time was about 117 seconds.

### 5.2 Qualitative Results

The participants were given subjective questionnaires so that the information we received would be as holistic as possible. It was useful because it gave participants to explain their feelings towards the first tutorial they did before trying the other as well as their feelings towards both versions of the tutorial after they tried both. In the first subjective questionnaire after completing the interview, participants were asked to rate their opinions on 5 aspects of the tutorial based on a 7 point Likert scale. We calculated the average of these scores for each group, and per the second part of our hypothesis we expected to see higher scores for the AR group than the desktop group.

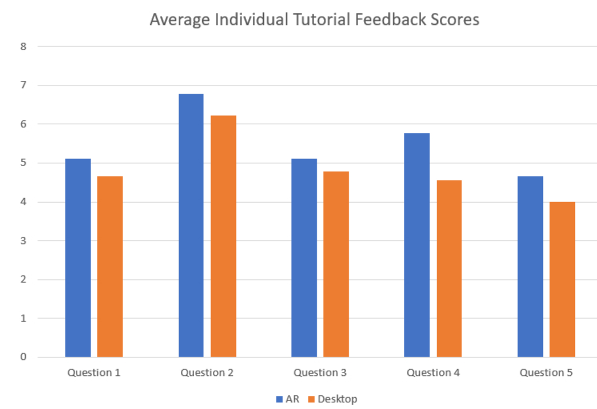


Figure 10: Average group scores for subjective feedback questions after first version of tutorial was completed

As part of the first subjective feedback questionnaire before the subjects tried the other version of the tutorial, we also gave them the opportunity to share any comments or questions they had regarding the tutorial. We will summarize the responses from both the AR group and the desktop group. For the AR group, 3 subjects mentioned that the hands on experience with AR was helpful for learning. The remaining comments were not repeated across multiple subjects: it was difficult to read text in AR, it was easy to be distracted by AR since it is new (but over time this would diminish), sometimes the AR interface was confusing, the headset gave a bit of a headache, and the field of vision was not great.

For the desktop group, 4 subjects said that the tool example was most helpful. 2 subjects said that the interface was clear, user friendly, and easy to use. 2 subjects felt that the slide portion was boring and harder to remember. The remaining comments were not repeated across multiple subjects: combining objects was a little conceptually unclear, and the desktop version was generally boring.

With the second subjective feedback questionnaire after the subjects had tried both versions of the tutorial, we will summarize the responses from both groups together. For the first question about which tutorial the subject liked better and why, all 18 subjects preferred the AR version and most of the commentary was positive with respect to the AR version. 10 subjects said it was more interactive, 2 said it forced them to focus, 2 said it was more engaging and interesting, 2 said it was more fun, and 2 said it gave a tactile sense that you can't get on a computer. The remaining comments with respect to this first question were not repeated across multiple subjects: the interaction made concepts easier to understand, the combination of tactile and auditory learning styles helped, there was a game like feel to AR, it was easier to follow, there was more active learning over passive learning, there was better visualization, and the example was especially done well. While all of these were positive comments with respect to the AR version of the tutorial, there were two comments that were more negative (not repeated): they paid more attention to the content of the slides during the desktop version, and sometimes lost focus.

For the second question about the advantages of each version of the tutorial, we will summarize the advantages mentioned for each version. With the AR version, 6 subjects said it was more engaging, 5 said it was more interactive, 2 said it was more likely to capture attention, 2 said it was more entertaining, and 2 said it was better for hands on learners. The remaining comments with respect to AR in this second question were not repeated across multiple subjects: it has experiential learning, it was more physical, it was more memorable for a visual learner, it was more immersive, it was more gamelike, it was easier to learn the material, there was active learning, the abstract concept was taught in a more representative manner, they could see how AR could visualize things that couldn't be visualized otherwise, it was cool because it was new, and the visualization of merging tools was helpful for explaining inheritance. With the desktop computer version, 5 subjects said it was easier to focus, 4 said it was more direct, 4 said it was more accessible, 3 said it was quicker, 3 said it was easy to navigate, 2 said it was familiar, and 2 said it was less distracting. The remaining comments with respect to the desktop computer version in this second question were not repeated across multiple subjects: it was less blurry, it was less glitchy, it was better for visual learners, it was interactive, it was more user friendly, and it was easier to read.

For the third question about the disadvantages of each version of the tutorial, we will summarize the disadvantages mentioned for each version. With the AR version, 5 subjects said it was distracting, 4 said it requires getting used to the device, 4 said there were more bugs/glitches, 3 said it was not accessible, 2 said it was difficult to read, and 2 said there was a limited field of view. The remaining comments with respect to AR in this third question were not repeated across multiple subjects: it was challenging to navigate, the visuals were less clear, it takes extra time, it was harder to set up, it is confusing to newcomers, it was harder to focus because of the cool factor, it could be more interesting, and it was disorienting. With the desktop computer version, 11 subjects said it was boring, 7 said it was harder to pay attention, 3 said it was less engaging, 2 said it was less interactive, and 2 said it was easier to passively watch. There were only two non-repeated comments with respect to the desktop computer version in this second question: it was harder to remember, and it was not hands on.

For the fourth question about whether the subject would like to use either of the interfaces for future tutorials and why, we will summarize what subjects said about both the AR and desktop computer interfaces. With the AR interface, all 18 subjects said they would like to use the AR interface for future tutorials. As for the repeated reasons why, 4 subjects said it has lots of potential, 4 said it was fun, 3 said it was engaging, 2 said it was novel, 2 said it was interesting, and 2 said it can make people interested in topics that would otherwise be uninteresting. The remaining comments with respect to AR in this fourth question were not repeated across multiple subjects: it was interactive, it was exciting, it was clear, there was a well rounded understanding of material, it was immersive, it was good for complicated topics, and it could bring in more users. With the desktop computer interface, 7 of the 18 subjects said they would also use the desktop computer interface for future tutorials and the following comments were mentioned but not repeated across multiple subjects: it used concrete examples, it is portable, it is interactive, and it is simple but effective.

Since the second goal of our study is to see if AR can make computer science education more enjoyable and preferable, analyzing the subjective feedback from the two groups is also essential to evaluating our hypothesis. In addition, understanding the general trends from this qualitative data can help provide an explanation for the quantitative data gathered.

## 6 Discussion

Both the quantitative and qualitative results of our study are very valuable and in general have both supported the two parts of our hypothesis. For the most part, the results speak for themselves, but nevertheless we discuss our interpretations of the results and their implications for the future of AR in computer science education in the following sections.

### 6.1 Interpreting the Quantitative Results

The quantitative results showed a relatively consistent pattern of higher recall scores for the interview questions with the AR group over the desktop group. This can be seen pretty clearly in Figure 8, as the average score for every question other than question 2 was higher for the AR group than it was for the desktop group. Using the total average score and the total standard deviation for all 5 questions for each group, we calculated a p-value to see if the difference between these total scores is statistically significant. The p-value we calculate was approximately 0.027, which is below the commonly accepted significance threshold of 0.05, so we can accept the first part of our hypothesis as being supported by our quantitative results. As for the difference in the average time it took for subjects to complete the interview, this difference was not close to statistically significant and there is no reason to believe that there is any correlation with our hypothesis.

A number of the qualitative results provide us with insight as to why this may have occurred, but the ones that stand out the most are that the AR tutorial was more interactive, caused the subjects to focus more, was more engaging and fun, and had the ability to make subjects interested in an educational topic that they would otherwise have no interest in. While this finding is certainly a positive indication of AR's capability to be an effective tool in computer science education, it is still important that we acknowledge that this result pertains to a tutorial on a very small set of content and further research must be conducted to see just how broadly AR can be applied in computer science education or other educational topics. The concept of inheritance is often explained in a somewhat abstract way, so it is certainly positive that AR can help explain a somewhat abstract concept, although compared to many other concepts in computer science inheritance is far less abstract. Another factor that is important not to ignore is that none of our subjects had ever experience AR glasses before, so it makes sense that they would be more excited and pay more attention to this novel device. As a result, it is unclear whether the qualitative descriptions about the AR tutorial being more exciting and attention grabbing would hold if the subjects had been used to the device. Despite these reservations, the statistical significance of the quantitative results absolutely warrants further research.

### 6.2 Interpreting the Qualitative Results

The first part of the qualitative results from the first subjective feedback questionnaire interestingly did not show that the AR group gave significantly higher scores on the 7 point Likert scale questions. For all 5 of these 7 point Likert scale questions, the AR group did give a higher rating than the desktop group, but only slightly, and the p-value calculated for this difference was not close to being significant. However, we believe these results may be explained by the content of the tutorial rather than the interface. Question 1 for instance asks about how much the subject feels that they have learned, and since the content being learned is the same for both groups, it does make sense that the two groups feel that they have learned about the same amount. These questions also do not directly ask the subjects to compare the two interfaces. The responses in the question asking the subjects to leave any comments or questions after the Likert scale questions may also explain why these scores were not significantly different. For the AR tutorial, while 3 subjects said that hands on experience with AR was helpful for learning, the rest of the comments were negative and talked about the various issues of using the device including the difficulty to read text, ease of distraction, confusion, headaches, and poor field of view. These negative aspects may explain why subjects in the AR group did not give higher Likert scale scores.

The second part of the qualitative results from the second subjective feedback questionnaire, however, very consistently indicated that the AR tutorial was preferable and more enjoyable, which supports the second part of our hypothesis. In the first question of this second subjective feedback questionnaire, all 18 participants preferred the AR version, and this unanimous result arguably trumps the lack of support from the results of the Likert scale questions in the first subjective feedback questionnaire. The fact of the matter is, once the subjects had seen both versions of the tutorial, there was pretty much no question which one they liked more. Sure, the question that followed about the disadvantages revealed a number of downsides to the AR version, but the benefits of the AR version clearly outweigh these downsides. It still shouldn't be ignored that 7 of the participants also would use the desktop computer interface in the future, and these results by no means suggest that we should just replaced traditional computers with AR devices, but there were many benefits from the AR device that could not be attained with the desktop computer device. Again, we should still be wary of the fact that none of the participants had ever tried AR glasses before, so it's possible that the novelty of the device played an important part in the qualitative results observed. What was especially positive though was that 2 subjects explicitly said that AR could make people interested in topics that would otherwise be uninteresting to them. This speaks to the core goal of the study: addressing the problem that computer science education is increasingly relevant but often unapproachable to students from a non-STEM background. For students that would end up being interested in computer science education on their own, AR could certainly make that education process easier and more enjoyable, but we are not as concerned with students that are already interested in computer science education. By making introductory computer science concepts easier and more enjoyable to learn, we can theoretically bring in many more students into computer science education which as we discuss is increasingly important to address many of the world's problems that involve computer technology.

## 7 Conclusions and Future Work

The use of augmented reality in education has been explored to some extent already, but 6DOF AR devices have only been widely available since the release of the first Hololens device and there have not been many studies investigating the educational potential of 6DOF AR. Beyond that, there has been even less investigation into the use of AR for computer science education and the abstract concepts within it. This study has provided some important findings that 6DOF AR can be used both to improve students' abilities to retain information from introductory computer science curricula and at the same time enjoy the process more. Our study has shown that AR can make students more interested in computer science topics that they would otherwise not be interested in, that it can make certain concepts feel less abstract, and that there are benefits like levels of engagement and interaction that cannot be achieved with other existing educational mediums. Along with these benefits, however, are a number of hurdles that still need to be overcome in order for AR to become widely used in educational settings. AR needs to become more accessible, easier to use, and less glitchy (especially the tracking reliability). Aspects of the optical quality of AR devices need to be improved as well, such as readability and field of view.

While this study is a good start, a lot of future research needs to be done in addition to improvements to the hardware in order for educational institutions to feel confident that adding AR devices to their ecosystem of educational tools is worthwhile. We need to test the application of AR on larger, more complex curricula, both in the field of computer science and otherwise. We need to see if findings like the ones from this study will hold true with participants who are well versed with AR technology and do not have the novelty factor influencing their experience. We also need to explore the implications of multiplayer AR experiences and the benefits, downsides, and complexities that come along with them. Since we have observed a wide variety of subjective responses in this study, running these kinds of experiments on larger and more diverse populations is also essential to understanding the robustness of AR's potential as an educational tool and who it is for. Despite the long road ahead before the potential widespread adoption of AR, these initial findings present an exciting road map for the future of technology in education.

## References


[1] B. BM, A. EA, R. FD, C. BR, and L. AG. The specificity of memory enhancement during interaction with a virtual environment. *US National Library for Medicine*, 7(1):65–78, Jan. 1999.

[2] D. A. Bowman, L. F. Hodges, D. Allison, and J. Wineman. The educational value of an information-rich virtual environment. *Presence*, 8(3):317–331, June 1999.

[3] M. Bricken. Virtual reality learning environments: Potentials and challenges. *SIGGRAPH Comput. Graph.*, 25(3):178–184, July 1991.

[4] H. Q. Dinh, N. Walker, C. Song, A. Kobayashi, and L. F. Hodges. Evaluating the importance of multi-sensory input on memory and the sense of presence in virtual environments. *IEEE*, pages 222–228, 1999.

[5] M. Hussein and C. Natterdal. The benefits of virtual reality in education. *Department of Computer Science and Engineering*, 1(2):1–15, June 2015.

[6] Ibanez and Delgado-Kloos. Augmented reality for stem learning: A systematic review. *Science Direct*, 123(1):109–123, Aug 2018.

[7] Z. Pan, A. D. Cheok, H. Yanga, J. Zhua, and J. Shia. Virtual reality and mixed reality for virtual learning environments. *Elsevier Ltd*, 30(1):20–28, Jan. 2005.

[8] Petkov. How to explain object-oriented programming concepts to a 6-year-old. *freeCodeCamp*, 1(1):1, Jun 2018.

[9] E. D. Ragan, A. Sowndararajan, R. Kopper, and D. A. Bowman. The effects of higher levels of immersion on procedure memorization performance and implications for educational virtual environments. *MIT Press Journals*, 19(6):527–543, Dec. 2010.

[10] C. D. Wickens. Virtual reality and education. In *[Proceedings] 1992 IEEE International Conference on Systems, Man, and Cybernetics*, pages 842–847 vol.1, Oct 1992.

# A Audio Clip Scripts

## A.1 AR Version

(Initialize app with orientation of objects before giving headset to subject)
(Give headset to subject)

(Not pre-recorded) In this tutorial, you can select objects by reaching out to them so your controller is touching the object and then press the bumper button to select. The tutorial will prompt you to select 3D objects, so make sure to only select the digital objects and dont bump the controller into real-world objects.

(Not pre-recorded) Please make sure to listen carefully until each audio clip has fully finished and then follow the instructions at the end of each audio clip. Also, please dont play around or do anything you are not instructed to do. When youre ready, select the blue cube on the table, bring it over the pink cube on the ground, and press the bumper button again to drop the blue cube onto the pink cube. This will start the tutorial.

Step 1: At a very basic level, a programming language is a set of commands and instructions that allow humans to communicate with a computer to create a software program. Just like we use human languages to tell humans to do something, we use programming languages to tell computers to do something. Even though programming languages can seem extremely foreign and too complex to understand, they are designed in a way that allows us to solve problems with the same kind of logic that we use to solve everyday problems. Object oriented programming languages are a type of programming language that treats pieces of code as objects, and these objects belong to what is called a class that describes a certain type of object. Pick up the blue cube from the table again and drop it onto the next pink cube on the floor to go to the next step.

Step 2: One way of thinking about objects and classes is in the case of tools like the ones on this table. Each tool has certain properties and can help us accomplish certain tasks. Lets say we represent a tool class with this tool handle, and the objects that belong to this class are tools like the hammer, knife, and scissors. Each individual tool object is considered an instance of the tool class, and there can be multiple objects that belong to the same class. Drop the blue cube on the next pink cube to continue.

Step 3: In object oriented programming, classes and instances of classes will have properties that we call variables, and they will also have functions that allow them to do something that we call methods. For instance, if you look at the tool handle you can see the tool class has three variables: handle type, size, and weight. It also has has two methods: pick up and put down. There are a number of principles of designing and creating objects in object oriented programming, and one of the core concepts used in creating objects is called inheritance. Drop the blue cube on the next pink cube to continue.

Step 4: When we are designing different objects to create different functionality that we want, often times these objects are very similar. They share common logic. But theyre not entirely the same. When creating a lot of objects, it can be tedious to have to do the same work over and over to create new classes of objects. How do we reuse the common logic and extract the unique logic into a separate class so that we dont have to do redundant work? Inheritance is a concept in object oriented programming that helps us achieve this. Drop the blue cube on the next pink cube to continue.

Step 5: The idea behind inheritance is that you create a child class by deriving from another parent class. This way, we form a hierarchy. The child class reuses all variables and methods of the parent class and can implement its own unique variables and methods. Lets consider the tool example to see how to use the concept of inheritance. Select the tool example cube above the tools on the table to learn how to create new classes of tools.

Step 6: Say we want to create a hammer child class that belongs to the tool parent class. Create a new instance of the tool class by selecting the digital handle representing a tool object above the real tool handle. This will cause you to grab a new copy of a tool object.

Step 7: Now hover your controller over the table to the right of the tool handle and press the bumper button again to drop the tool object copy on the table.

Step 8: Next, grab a copy of the hammerhead part by selecting the digital hammerhead part above the hammerhead. Then bring your controller and the hammerhead copy over to the tool object you just dropped on the table until the hammerhead touches the tool object.

Step 9: Good job, youve created a hammer child class that has the same variables and methods as the tool parent class as well as a new unique method, hit. Now lets create a knife child class. Like you did before, grab another digital copy of the tool object and drop it on the table to the right of the hammer you just created.

Step 10: Now grab a digital copy of the knife blade part above the knife blade and bring the copy over to the tool object you just dropped on the table until the two objects touch.

Step 11: Awesome, now youve created a second child class of the tool parent class, the cutting tool class. This class inherits the variables from its parent class and adds a new method, cut, as well as a new variable, sharpness. With inheritance, you can also create a child class of a child class that becomes part of the class hierarchy. Lets do this by creating a scissors class that is a child of the cutting tool class. To do this, you will see a digital knife copy above the physical knife object on the left. Grab a new copy of the knife object and then drag it over to the knife object you just created until the two knives touch.

Step 12: You got the hang of it. Now youve created the scissors class which inherits everything from the tool and cutting tool classes, and adds a new method, snip.

By using the principle of inheritance, each class adds only what is necessary for it while reusing common logic from the parent classes. Thats the basics! You have now finished the tutorial. Please give the controller and headset to the supervisor now.

## A.2 Desktop Version

(Not pre-recorded) Please make sure to listen carefully until each audio clip has fully finished and then follow the instructions at the end of each audio clip. Also, please dont play around or do anything you are not instructed to do. When youre ready, select the start tutorial button.

Step 1: At a very basic level, a programming language is a set of commands and instructions that allow humans to communicate with a computer to create a software program. Just like we use human languages to tell humans to do something, we use programming languages to tell computers to do something. Even though programming languages can seem extremely foreign and too complex to understand, they are designed in a way that allows us to solve problems with the same kind of logic that we use to solve everyday problems. Object oriented programming languages are a type of programming language that treats pieces of code as objects, and these objects belong to what is called a class that describes a certain type of object. Select the next step button to continue.

Step 2: One way of thinking about objects and classes is in the case of tools like the ones shown. Each tool has certain properties and can help us accomplish certain tasks. Lets say we represent a tool class with the tool handle, and the objects that belong to this class are tools like the hammer, knife, and scissors. Each individual tool object is considered an instance of the tool class, and there can be multiple objects that belong to the same class. Select the next step button to continue.

Step 3: In object oriented programming, classes and instances of classes will have properties that we call variables, and they will also have functions that allow them to do something that we call methods. For instance if you look at the tool handle you can see the tool class has three variables: handle type, size, and weight. It also has has two methods: pick up and put down. There are a number of principles of designing and creating objects in object oriented programming, and one of the core concepts used in creating objects is called inheritance. Select the next step button to continue.

Step 4: When we are designing different objects to create different functionality that we want, often times these objects are very similar. They share common logic. But theyre not entirely the same. When creating a lot of objects, it can be tedious to have to do the same work over and over to create new objects. How do we reuse the common logic and extract the unique logic into a separate class so that we dont have to do redundant work? Inheritance is a concept in object oriented programming that helps us achieve this. Select the next step button to continue.

Step 5: The idea behind inheritance is that you create a child class by deriving from another parent class. This way, we form a hierarchy. The child class reuses all variables and methods of the parent class and can implement its own unique variables and methods. Lets consider the tool example to see how to use the concept of inheritance. Select the next step button to continue.

Step 6: Say we want to create a hammer child class that belongs to the tool parent class. Create a new instance of the tool class by selecting the tool handle. This will cause you to grab a new copy of a tool object.

Step 7: Now hover your mouse over the workspace box to the right of the tool handle and click the mouse button again to drop the tool object copy into the workspace.

Step 8: Next, select the hammerhead part to grab a copy of the hammerhead part. Then drop the hammerhead part in the workspace to add the hammerhead to the tool handle.

Step 9: Good job, youve created a hammer child class that has the same variables and methods as the tool parent class as well as a new unique method, hit. Now lets create a knife child class. Like you did before, grab another copy of the tool object and drop it into the workspace.

Step 10: Now grab a copy of the knife blade part and drop it in the workspace to add the knife blade to the tool handle.

Step 11: Awesome, now youve created a second child class of the tool parent class, the cutting tool class. This class inherits the variables from its parent class and adds a new method, cut, as well as a new variable, sharpness. With inheritance, you can also create a child class of a child class that becomes part of the class hierarchy. Lets do this by creating a scissors class that is a child of the cutting tool class. Grab a copy of the knife object on the left and drop it in the workspace, then grab another copy of the knife object and drop it in the workspace to combine them.

Step 12: You got the hang of it. Now youve created the scissors class which inherits everything from the tool and cutting tool classes, and adds a new method, snip.

By using the principle of inheritance, each class adds only what is necessary for it while reusing common logic from the parent classes. Thats the basics! You have now finished the tutorial.

## B PICTURES OF STUDY

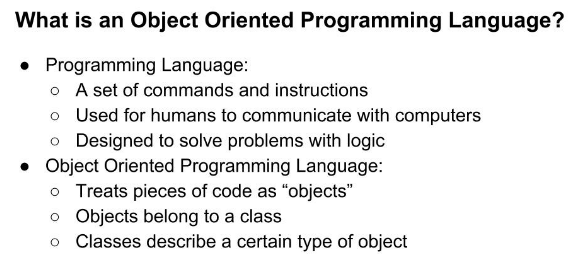


Figure 11: Slide 1 (at Step 1)

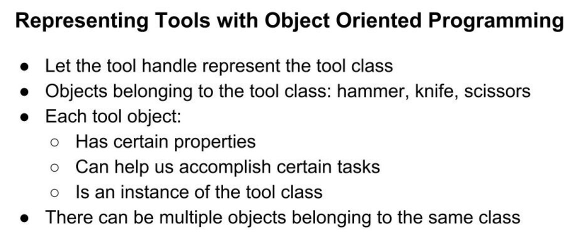


Figure 12: Slide 2 (at Step 2)

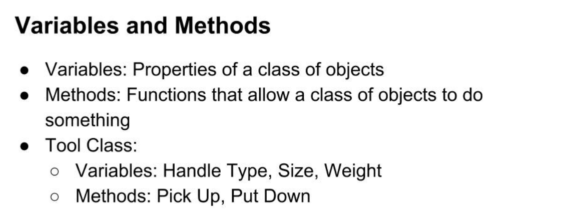


Figure 13: Slide 3 (at Step 3)

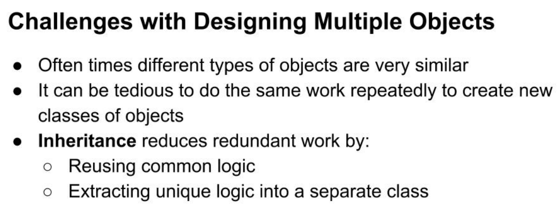


Figure 14: Slide 4 (at Step 4)

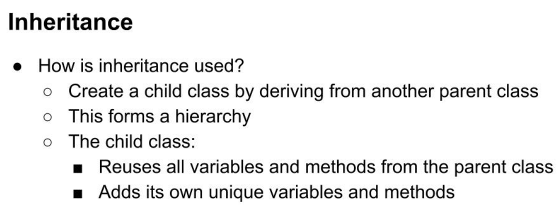


Figure 15: Slide 5 (at Step 5)

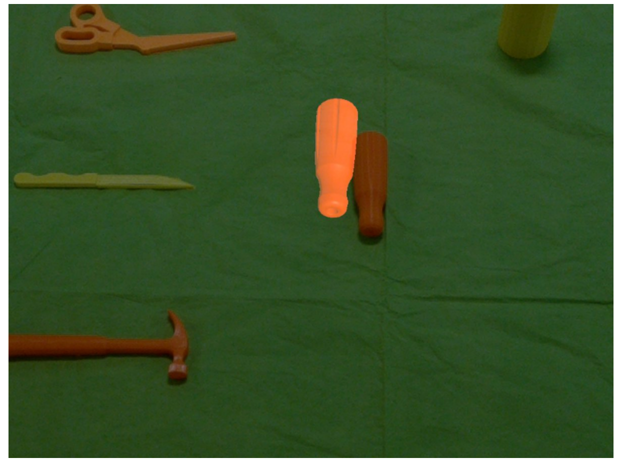

Figure 16: Step 6 (AR Version)

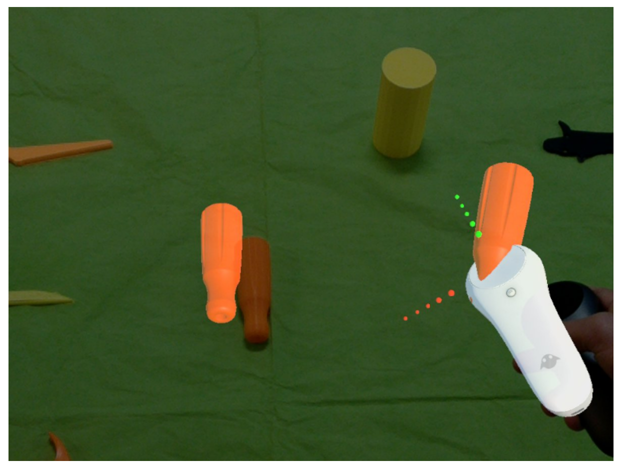

Figure 17: Step 7 (AR Version)

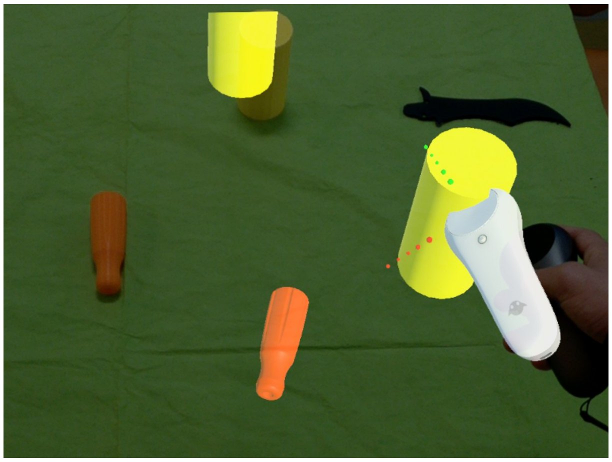

Figure 18: Step 8 (AR Version)

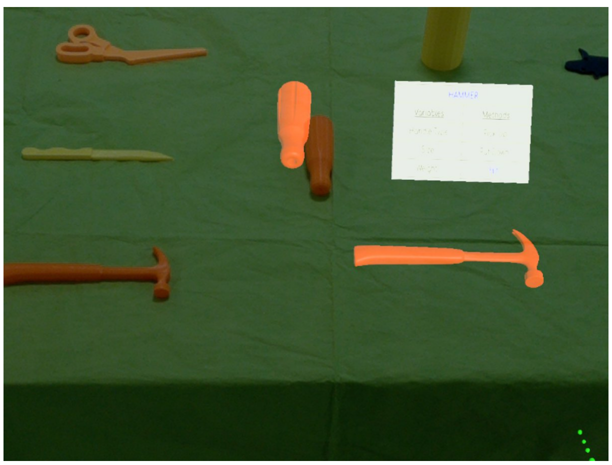

Figure 19: Step 9 (AR Version)

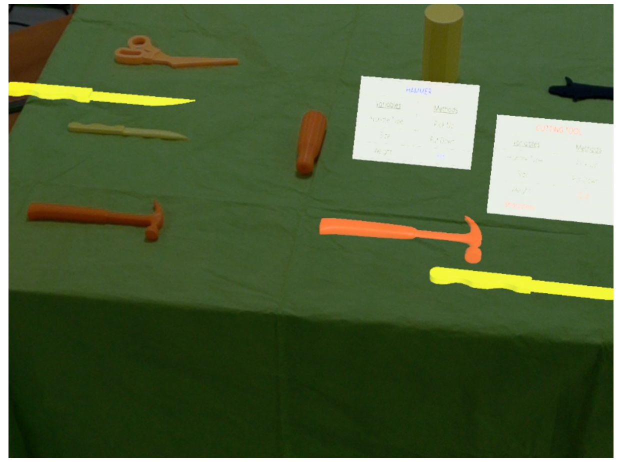

Figure 20: Step 11 (AR Version)

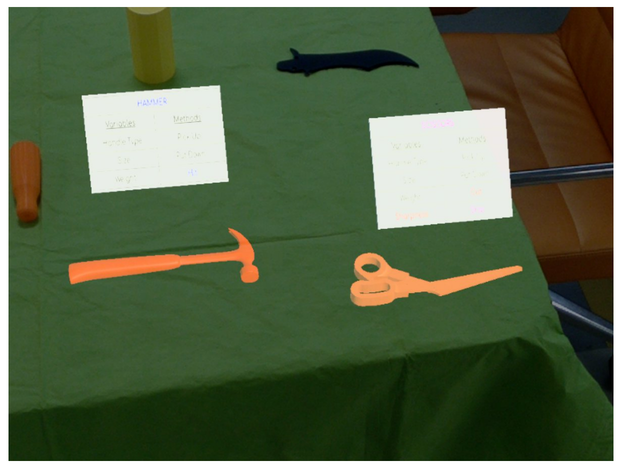

Figure 21: Step 12 (AR Version)

## C QUESTIONNAIRES

## INTERVIEW QUESTIONS TO ASSESS OBJECT INHERITANCE UNDERSTANDING

1) What is a programming language?

Criterion for assessment:
0 – Subject does not know what a programming language is
1 – Subject acknowledges that a programming language is a set of commands and instructions
2 – Subject acknowledges that a programming language is used for humans to communicate with computers
3 – Subject acknowledges that a programming language is a set of commands and instructions used for humans to communicate with computers

2) What is an object oriented programming language?

Criterion for assessment:
0 – Subject does not know what an object oriented programming language is
1 – Subject acknowledges that an object oriented programming language treats pieces of code like objects
2 – Subject acknowledges that classes in an object oriented programming language describe a certain type of object
3 – Subject acknowledges that objects can be used to accomplish certain tasks and that there can be multiple instances of objects from the same class to accomplish similar tasks

3) What are the features of an object in an object oriented programming language?

Criterion for assessment:
0 – Subject does not know what the features of an object are
1 – Subject acknowledges that objects have properties called variables
2 – Subject acknowledges that objects have functions that allow them to do something called methods
3 – Subject acknowledges that objects have both properties called variables and functions called methods

4) Why is it useful to use the principle of inheritance in object oriented programming?

Criterion for assessment:
0 – Subject does not know why it is useful to use the principle of inheritance
1 – Subject acknowledges that there is a need to create new functionality
2 – Subject acknowledges that there is a benefit to reducing redundant work
3 – Subject acknowledges that there is a need to create new functionality while reducing redundant work

5) How can we use inheritance to create new classes of objects?

Criterion for assessment:
0 – Subject does not know how to use inheritance to create new classes of objects
1 – Subject acknowledges that one starts with a parent class of objects and creates a child class that derives from the parent class
2 – Subject acknowledges that a hierarchy is formed and that the child class reuses all variables and methods from the parent class
3 – Subject acknowledges that the child class adds its own unique variables and methods to go along with the preexisting variables and methods inherited from the parent class

Figure 22: Interview Questions and Scoring Rubrics

**SUBJECTIVE FEEDBACK SURVEY ON TUTORIAL EXPERIENCE**

Please answer the following questions to your best ability.

1) How much do you feel that you have learned as a result of this tutorial?
Nothing at all O 1 O 2 O 3 O 4 O 5 O 6 O 7 I learned a lot

2) Was the interface and set of instructions on what to do clear?
Not at all O 1 O 2 O 3 O 4 O 5 O 6 O 7 Extremely clear

3) How much did you enjoy the learning process of this tutorial?
No enjoyment O 1 O 2 O 3 O 4 O 5 O 6 O 7 I enjoyed the process a lot

4) How much did you enjoy using the interface of this tutorial?
No enjoyment O 1 O 2 O 3 O 4 O 5 O 6 O 7 I enjoyed the interface a lot

5) Are you now any more likely to want to learn more about computer science/computer science concepts?
Not at all more likely O 1 O 2 O 3 O 4 O 5 O 6 O 7 Much more likely

6) Please leave any comments or questions about your experience below:

Please answer the following questions after you have completed both versions of the tutorial.

1) Which version of the tutorial did you like better and why?

2) What do you think the advantages of each version of the tutorial are?

3) What do you think the disadvantages of each version of the tutorial are?

4) Would you like to use either of these interfaces for future tutorials? If so, why?

Figure 23: Subjective Questionnaire