

UNDERGRADUATE REPORT

Applying VR Technology to Training in Manufacturing Processes

by Kathryn Ledoux and Nikita Deshpande

Advisor:

UG 2006-12



ISR develops, applies and teaches advanced methodologies of design and analysis to solve complex, hierarchical, heterogeneous and dynamic problems of engineering technology and systems for industry and government.

ISR is a permanent institute of the University of Maryland, within the Glenn L. Martin Institute of Technology/A. James Clark School of Engineering. It is a National Science Foundation Engineering Research Center.

Web site <http://www.isr.umd.edu>

Introduction

The use of Virtual Reality to enhance existing training methods is growing increasingly popular in fields such as medicine, emergency services, and flight instruction. It offers an effective training solution in terms of cost, time expenditure, quality of training offered. The Virtual Training Studio was created to offer tutorials on industrial and procedural assembly and disassembly operations, serving the manufacturing industry, whose turnover rate demands constant training for new employees. It currently supports a tutorial instructing users how to assemble a naval rocket. Users are placed in a virtual environment and their position is tracked by four cameras, as well as two gyroscopes, mounted on their HMD and wand controller. The parts of the rocket are positioned in the environment, and step-by-step text instructions are displayed on a projector screen in the room. Users can click buttons next to the projector to select different instruction options. Auto allows them to view an animated automatic assembly, Manual allows them to assemble the object independently using the instructions given on the screen, and Video allows them to view step-by-step video instructions.

New Tutorials

The parts for the naval rocket are simple, with minimal surface features, making their assembly fairly straightforward. In order to continue the project's progress, the system demands more tutorials that require more complex part manipulation

One of the goals for the project, therefore, was to implement two new tutorials, in addition to the assembly of the naval rocket. This involved modeling both the parts of the object to be assembled and any tools necessary for the assembly, as well as modeling an environment for the user to be immersed in when completing the tutorial. We were asked to come up with two objects, one that could be assembled and one to be disassembled. Factors that were taken into

account included the surface features of the object, the number of parts it had, the accessibility of the parts, their complexity (whether they could be modeled relatively easily), and the usefulness and real-world application of having a tutorial on how to assemble or disassemble that part. We chose to create a car, with simplified, to-scale parts under the hood, for a tutorial on how to perform simple maintenance procedures (e.g. change the spark plugs), and a model airplane engine for the assembly tutorial.

Both objects were modeled entirely in Pro/Engineer, an upper level parametric feature-based 3D solid modeling CAD program. By the time all the parts were completed, we both felt that we had gained a much greater understanding of when and how to use most of the features Pro/E supports.

We began the process with modeling the car. Using Kathryn's 1999 Mazda Protegé as our reference, our first step was to take several measurements of the car body itself, as well as the parts under the hood. We created sketches of the individual parts and of the car body, and then used these to model the parts to scale in Pro/E. As we encountered features on parts that we did not know how to create, we had to learn new operations in Pro/E that would allow us to create these features. The final car assembly consisted of the car body, complete with windows, side-view mirrors, doors, trunk, hood area, tires, shock towers, fuse box, air box, air intake, power steering fluid, washer fluid, engine coolant, radiator, battery, engine block, spark plugs, and valve cover. We did not add a cover for the hood since, if it were propped open, it would obstruct the user's view in the virtual world while he was performing the operation.

The body of the car proved to be the most difficult to model. After it was created the first time, the software ran into problems when we tried to make modifications to it. We ended up creating a new model, and were able to modify it without any trouble. We used digital

photographs of a tire and the windshield on Kathryn's car to create and apply texture maps to our model, making it look much more realistic. We had a difficult time finding any instructions on how to apply textures using Pro/E. The two tutorial books we had did not mention it, and none of the online forums we looked through seemed to offer any advice. Finally, by playing around with the Color and Appearance menu, we were able to figure it out.



Figure 1

The engine block, with four spark plugs inserted into it and the valve cover on top, was created as a subassembly and later inserted into the engine bay of the car. We applied texture to the valve cover, giving it a rusted, worn metal appearance. We used the texture mapping method on the tires and car battery, as well. This allowed us to make these parts appear very detailed without spending extra time on modeling. For other parts, such as the fuse box and the oil cap, we modeled fine surface features rather than placing textures on their surfaces.

After the car and its engine were modeled, we next modeled the model airplane engine. This part, too, was modeled using Pro/Engineer. The model airplane engine was chosen for the assembly tutorial because it consisted of a greater number of parts, many of which were more complex and had more detailed surface features, than the naval rocket in the original assembly tutorial. The assembly of the airplane engine also required more steps to assemble, which tended to be more complicated than the simple insertion steps of the original tutorial. Finally, the model airplane engine was chosen because the assembly process taught by the tutorial can be translated

to other more complex machines and parts that have more complicated assemblies, giving the tutorial direct applications to the real world, as organizations, such as the military, can take advantage of this technology.

In order to model this part, an actual model airplane engine was obtained. The parts were disassembled and measurements were taken on all of the parts. This model consisted of about seven main parts, including the muffler, front plate, top plate, piston, crankshaft, and needle valve. The wooden block in the figure below is included in the model as a placeholder for the propeller.




Figure 2




Figure 3

In addition to the parts themselves, we also modeled tools in Pro/E that will eventually be added to the tutorials. In the current tutorial, tools are not included; hand rotation and general alignment govern the assembly procedure. The tools modeled, including a screwdriver, Alan wrench, and ratchet, will eventually be included to add more detail to the tutorial and to provide the user with a more realistic sense of the assembly. Currently, the difficulty that exists with incorporating tools into the tutorial is in the type of movement associated with the tool's own movement within a specific task, such as screwing in a screw or loosening a nut.

In addition to modeling the parts for the objects used in the tutorial, we were asked to model two environments in which our objects, and any future objects incorporated into a tutorial, would be placed. We created two rooms using Lightwave, a 3-D modeling software most commonly used for creating characters and scenes for animation. The software is not very intuitive, so we had trouble getting used to the features. We ordered a book that was very helpful, and we also took advantage of online help. After having modeled several practice objects in Lightwave, we found out that Lightwave files were not easily converted to Virtual Reality Modeling Language (VRML) format (.wrl), the file type necessary to place objects in the virtual world. We wrote to the author of the book, and he recommended a software to us called Deep Exploration that would convert the files.

As we created the garage for the car and the generic room for the airplane engine, we gradually learned more about Lightwave, particularly its limitations for the type of modeling we were trying to do. While the software was not ideal, we were able to create environments fairly close to what we originally envisioned. One of Lightwave's unique and helpful features is its ability to create different layers for each object. This made texturing and coloring of different surfaces much easier. Lightwave also had two different formats, Modeler and Layout. Layout was helpful for visualizing what the completed room would look like with objects added in (e.g. doors). We experienced rendering problems with the Layout module, but they did not cause setbacks.

We were extremely satisfied with all of our final models, although we ran into a time constraint towards the end of the process. One part on the model airplane engine will need to be refined, and the extra space in the hood of the car needs to be filled, either by altering the dimensions of the existing parts, or adding a texture map of extra wires and tubes, or a

combination of both. The dimensions of the rooms may need to be altered as well. Whether this correction needs to be made will not be determined until the room is converted to VRML format and can be tested using the HMD. Each step in the virtual world is equal to two real steps, so the room may need to be enlarged.

Further Exploration of the Wand Interface

Abstract

Currently, the interface of the Virtual Reality Training Studio involves an HMD and a control wand, both with gyroscopes mounted on them. Since precise and timely manipulation of parts' positions is essential to performing assembly and disassembly operations, we decided to conduct a test with the wand interface to determine if it was sufficiently effective and allowed for accurate and quick orienting and positioning of objects. More specifically, we wanted to examine the efficacy of the rotation disk, which allows the user to rotate the part left or right and move the part farther or closer. Our hypothesis states that users will perform the tasks assigned to them more accurately and efficiently when they use the disk as compared to when they only use natural hand rotation.

Experimental Setup

Our experiment was performed on the Virtual Training Studio software. Users wore the Head-Mounted display and controlled part movement within the virtual world with a hand-held wand. Four cameras around the room tracked the user's location within the world and gyroscopes attached to both the HMD and the wand follow the user's exact movements. Our test group consisted of eight college students, two females and six males, none of whom had had prior experience on the virtual system.

The experiment itself was broken into two methods, each consisting of three separate tests, which were repeated for both methods. One method measured the results when the user could only use natural hand rotation of the wand to align the part, and the other method measured the results when the user could use both natural hand rotation and the disk on the wand to rotate the object. The test group was divided into two sub-groups: one group of four subjects used the hands only method for the first set of tests, while the second group of four people used the hybrid method of hand and disk rotation for their first set of tests.

A test included three distinct parts: a cylinder, an equilateral prism, and a scalene prism. The user needed to rotate the selected part with the wand to align its end to a base with the matching shape located near the end of the table. The part that was to be rotated would be highlighted in red, as can be seen in figure 4 below. The same original orientation was used for each part in each test and for every user.

The first two tests measured accuracy under time constraints. For the first timed test, the user was given 20 seconds to rotate and align each part as accurately as he could. For the second test, the user was given 16 seconds per part to complete the same tasks. In order for the user to keep track of how much time he had to finish aligning that part, timers were placed both on the projector screen and also above the base at the end of the table. The third test was an unlimited time test that simply recorded how quickly the user could align the object with the base.



Figure 4



Figure 5

The computer program was set up such that, as long as the part was within 15 degrees of rotation on either side of the axis, the part would turn green to alert the user that they were close to true alignment and would receive a score for that part. But, as long as the user still had time left for that part, he could continue to rotate it to achieve the best alignment as he could judge from his view through the HMD. For the first two tests, scores were recorded by the program as percent accuracy scores; the closer the part was to the true alignment, the higher the percent would be. Three percents were given: two for orientation alignment and one for position, or how close the part's face was to the base. All measurements were referenced from the center of the part and the center of the shape on the base. For the third test, the program simply recorded the amount of time the user took to rotate each part. This test disregarded how accurately the user aligned the object; as long as the object was aligned within the designated region and turned green, the program would allow the user to move on to the next object.

Results

Overall, the results of our experiment were generally inconclusive as to which method, the hands or the hybrid, is better to use in the virtual world. When looking strictly at the overall average accuracy scores of the two tests, the graphs of both the 20 second time test and the 16 second time test are inconclusive regarding which method helps the user to perform the best. The data is quite varied, with the hybrid method sometimes having higher accuracy scores and the hands method having higher scores other times. The only trend that can be observed that seems to hold for most subjects was that, generally, the method that the subject used first tended to result in the lower accuracy score between the two methods. This trend is logical, as the user would have gained more experience with the system by the time they used the second method. In

the overall average score graph for 16 seconds, the accuracy scores were generally higher with the Hands method. This may have been due to the fact that there were more failures during the 16 seconds test with the Hybrid method than with the Hands method.

When comparing the accuracy scores for each individual part, some general observations can be made for some of the graphs. On both the 20 seconds and the 16 seconds graphs for the easy part, the accuracy scores for both methods are very close. This is most likely due to the fact that this part required minimal rotation, so the additional option to use the disk had minimal effect on the accuracy scores, as few people actually chose to use it during the hybrid test. For the medium part, there is also no discernable pattern in which method was better to use. First, a glitch developed in the software during the testing that caused the equilateral prism to turn green when inverted, rather than when aligned correctly. This affected the medium part's scores, as subjects were not able to truly judge how closely aligned their part was. Also, like the easy part, the medium shape required minimal rotation, making the use of the disk nearly obsolete, as most subjects chose not to use it. For the hard part, there was once again no obvious pattern among the data. The only significant observation is that, of those who failed to align the part within the designated range, the majority occurred during the hybrid tests. This could possibly be because the subjects attempted to use the disk, but had trouble actually making the object rotate it in the way they wanted it to, causing them to run out of time in the end.

When comparing the graphs of the 20 second time test to the 16 second time test for each method, it is observed that subjects consistently produced higher average scores in their 20 second trial than their 16 second trial, with one exception. The exception was subject 5, who was observed having trouble using the disc (pushing the button/making it work) during his 20 second trial. For his subsequent 16 second trial, he did not use the disk, and achieved a higher accuracy

score. Subjects 2, 4, and 6 had lower 20 second accuracy scores than 16 second accuracy scores. This may be because they used the Hands method first, and therefore the 20 second trial was their first test. However, this trend was not observed when comparing Hybrid 20 second with the Hybrid 16 second trials.

The graphs comparing using the method first with using the same method second for the 20 second test for both the hands and hybrid methods show interesting results. The average of overall average accuracy scores for subjects who used the hybrid method first (53.96) is lower compared with the average of overall average accuracy scores for subjects who used the hybrid method second (71.03). This demonstrates learning among the users. Those who used a method second produced better scores than their colleagues who used that method first. The average of overall average accuracy scores for subjects who used the Hands Only method first (65.74) is lower compared with the average of overall average accuracy scores for subjects who used the Hands Only method second (69.83). This demonstrates learning among the users. Those who used a method second produced better scores than their colleagues who used that method first. There is a smaller gap between these numbers and the ones previously mentioned for the hybrid method. This may be because hand rotation is more innate, and therefore the learning factor is not as influential on the subject's performance.

When comparing the scores for the easy, medium, and hard parts within both the hands and hybrid methods, there is a fairly noticeable, yet predictable, trend. Subjects tended to align the easy part the most accurately, followed by the medium and then the hard part. Sometimes though, the hard part was aligned more accurately, which may have been due to the computer glitch with the medium part. Finally, the data from the unlimited time tests also shows some slight trends. When observing the data for all three parts for both the hands and hybrid methods,

it must be noted that during the experiment, Subject 1's medium part time was affected by the software glitch, as that person had received no warning beforehand that this problem could possibly arise and became very confused when it occurred. If this data point is disregarded, the graph demonstrates that the times for the easy part and the times for the hard part are pretty consistently equal. The medium part times do not demonstrate a pattern quite as well, probably due to the glitch. When the total times for the hands and hybrid methods are compared, the times for both methods are all generally within the same range, once Subject 1's flawed data point is accounted for.

Even though no real patterns or trends could be established to determine which method was the most efficient one to use in the virtual world, some observations could be made about the possible sources of discrepancy in the data. Where the user grabbed the object affected their ability to use the disc effectively. Some were able to learn how to plan ahead on where to pick up the object to rotate it in the direction they needed to, which helped them improve their scores, but many were not able to do this. The software glitch that arose unexpectedly during some subjects' testing sessions affected their performance on the medium part, causing the scores to show greater variance than it may have otherwise. Another possible source could have been lack of familiarity with the system and the controller. This lack of familiarity affected the subject's depth perception, ability to use the buttons on the wand without being able to see it, and the ability to move objects on a screen using a controller. It was observed that previous experience with gaming systems may have helped user performance in the end, an observation which could be investigated further in future studies.

Conclusion

Virtual reality has the potential to be applied to many real world applications, especially in manufacturing industries. Because of its capability to transfer learning from the virtual world to the real world and to carefully model real world parts and machines using 3D modeling software, virtual reality can be used to train users to perform specific tasks, such as assembly processes. But the technology available to train users is still limited in its ability to provide a true sense of presence to the user. As evidenced by our own experiment on the wand interface, users still have difficulty maneuvering in the virtual world and using the hardware of the system. While a learning trend among the users was established, as evidenced by improved scores and comfort with the system in later tests, a truly accurate representation of motion in the real world was not realized. Therefore, our hypothesis could only be realized if the user realized where the object needed to be picked up to achieve the correct rotation. Further research in VR should focus on developing better hardware that will further increase the user's sense of presence in the system, while also testing the system setups, including both the program created and the hardware used, which best facilitate learning and allow for close and accurate knowledge transfer to the real world.

Appendix A- Experiment Data

