

Interactive Visualization and Augmentation of Mechanical Assembly Sequences

Rajeev Sharma Jose Molineros

The Beckman Institute, University of Illinois at Urbana-Champaign
405 N. Mathews Avenue, Urbana, IL 61801

Abstract

A display system which augments a user's view of the surrounding scene with annotations based on the scene contents, has many potential applications. In order to exploit this potential, systematic schemes need to be developed for presenting the augmentation stimuli. We consider the problem of scene augmentation in the context of a human engaged in assembling a mechanical object from its components. Concepts from robot assembly planning are used to develop a systematic framework for presenting augmentation stimuli for this assembly domain. An interactive simulation tool called AUDIT is then described, which can be used for developing and visualizing multi-modal augmentation schemes for assembly. The interactive nature of the system is useful for evaluating an augmentation scheme and studying its impact on the assembly/disassembly guidance.

Keywords: *augmented reality, assembly, disassembly, interactive system, visualization.*

1 Introduction

Augmented reality (AR) has the goal of enhancing a person's perception of the surrounding world [2, 3, 4, 5, 8, 9], unlike Virtual Reality (VR) that aims at replacing the perception of the world with an artificial one. Augmented reality is also known as augmented virtual reality, enhanced reality [3], mixed reality, embodied virtuality [11], or computer-augmented environment, and is related to ubiquitous computing [12]. An AR system could, for example, consist of a head-mounted, see-through screen on which 3-D graphics is rendered, such that the graphics is superimposed on the view of the surrounding scene [5, 8]. Alternatively, the augmentation could be achieved using a 3-D projection screen, or a fixed see-through screen, or a small hand-held monitor. Another form of augmentation involves superimposing information on the live video image of the real scene, which could be located remotely, as in tele-operation. In addition to the 3-D graphics, a nearby computer monitor could also be used to present additional information that augments the view of the real scene. An essential element of an AR system is that it must possess one or more sensors to monitor the state of the world surrounding the user. The type of sensor would depend on the extent of "scene understanding" required for displaying the augmentation stimuli for a particular task. The applications of AR include: aiding in the repair or maintenance [5], manufacturing [1], education, training, medicine [2], battlefield, etc. Most of these potential applications stem from the fact that the interface can display relevant information, at the appropriate time and location.

The technological challenge of building an AR system can be characterized in terms of the ability of the system to answer the following three questions in real-time:

- **Where to display the augmentation stimuli ?**
This requires the AR system to have information about the "real" objects in the user's vicinity, since the display is generally juxtaposed or situated with the real objects. We use the term *situated graphics* to characterize the specific class of graphics that is used for augmenting a scene; the graphics take on additional meaning when located next to a real object.
- **When to display the augmentation stimuli ?**
This again requires an understanding of the content of the surrounding scene, especially the changes that occur over time. The changes could be, for example, caused by the user interacting with the environment. It is also connected with the current user goal, which could be conveyed to the system by the user with the help of some input device.
- **What augmentation stimuli to display ?** This depends on the current state of the environment, and on the current goal of the user. The content of the augmentation display also depends on human factors that determine the optimal cognitive load on the user for a given configuration of display devices and the task.

Sensing seems to be a big challenge for the practical realization of an AR system and would be important in answering the *when*, *where*, and *what* questions, as outlined above. Another challenge, and the one which we are concerned with in this paper, is developing a systematic scheme for presentation of the augmentation stimuli. The design of a presentation scheme for possibly multi-modal, multi-dimensional, distributed information is a hard problem [5, 11, 12]. The difficulty of designing an efficient and effective information presentation scheme increases with information variety, the number of information sources, and information complexity. Such a scheme must also answer the questions of *when*, *where* and *what* information to present to the user so that it is most effective. This, in turn, needs some way to characterize the state of the world. Unfortunately, there does not seem to be a systematic way of doing it, and hence of presenting the augmentation stimuli. It is hoped that the context of a particular domain will be helpful.

The most notable attempt toward a general scheme for AR display was the use of a *knowledge-based* approach by Feiner [5]. In [5] Feiner described the AR system, KARMA,






Figure 1: Augmented reality for guiding assembly.

that uses a knowledge-based scheme for helping with the repair of a copying machine. The graphic annotations, or illustrations, are iteratively designed by IBIS, a rule-based graphics generator. Other presentation schemes for AR cater to a particular application [1].

In this paper, we use the domain of *assembly* to propose a systematic development of the augmentation scheme. This helps in defining a class of problems characterized in terms of an "assembly world." Figure 1 conceptually illustrates an AR system that will be considered for our discussion. The goal of this system is to guide a person in assembling (or disassembling) a mechanical object from its components. The guidance is provided with the help of see-through goggles that superimpose 3-D graphics over the assembly scene, and with the help of an earphone. There could also be additional information displayed on a nearby computer monitor. The feedback about the current content of the scene is provided to the AR system with the help of a set of video cameras, together with the necessary computer vision algorithms [9].

The task of assembling components to build a mechanical product has been studied in assembly automation for the manufacturing industry. For example, the research area of *robot assembly planning* is concerned with devising ways to program a robot system to carry out the assembly automatically. Some of the ideas developed under the robot assembly planning context are applicable to the assembly task by humans. For example, the notion of an assembly graph from the robotic domain can help in the discrete partitioning of a continuum of possible states of the world. This partitioning helps in managing the presentation of multi-modal augmentation stimuli, including error monitoring and recovery guidance. We investigate this, and show how robotic assembly concepts permit one to develop a systematic AR presentation scheme for assembly.

We present an interactive tool, called AUDIT (which is an acronym for AUGmentation Development Interactive Tool) for studying the design of augmentation schemes for guiding a human operator in an assembly task. The interactive tool helps in finding the best manner of presenting the augmentation stimuli. The AUDIT system is also useful for interactively visualizing the different stages of as-

sembly. The functionality of AUDIT is demonstrated with the help of specific assembly scenarios using three simple mechanical designs. AUDIT provides the following main features:

- an ability to visualize different stages of an assembly; this helps in learning about different geometrically feasible mechanical mating operations at the subsequent assembly steps.
- a way to evaluate different augmentation schemes, by trying out different augmentation displays at different locations for the various stages of the assembly.
- an ability to visualize a high-level decomposition of the assembly process using assembly graphs; this helps in designing a better control sequence for the augmentation display in its varying role in assembly guidance, error monitoring, and error recovery.
- a means of manipulating assembly parts to simulate the assembly operation; the current implementation uses a mouse but this can be changed to other input devices that help in "playing" with different ways in which an object is assembled, while visualizing the role of augmentation.

2 The Assembly Domain

In this section we introduce the theory and techniques from robot assembly planning that are relevant to AR. In particular, this provides the skeleton for developing augmentation schemes for assembly sequences. The schemes will ultimately be used for providing the augmentation stimuli for guiding assembly operations performed by a human.

The assembly involves starting from a set of parts and putting them together in a proper sequence for creating an assembled object. Figure 2, for example, shows different stages in the assembly of a CD from its component parts.




Figure 2: Different stages in assembling a CD.

The issues involved in monitoring and devising a scheme for augmentation require considering all possible assembly states and transitions between them. This is closely related to assembly planning stage in robotics for devising automatic plans for assembling an industrial part from its components [7]. Robot assembly planning involves additional problems such as placement of fixtures, workspace layout design, gross-motion planning, and fine-motion planning. See [6] for a survey of robot assembly and task planning literature. Fortunately, for guiding a human assembly operation, the robotic issues such as fixturing and fine-motion planning become irrelevant, since high level assembly directives are adequate. However, the same formal representation scheme used for robot assembly planning can be



3.2 Augmentation using situated graphics

A unique aspect of the graphics for augmenting a real scene, like the assembly world, is that it is to be positioned or situated next to a real object. We use the term *situated graphics* to denote this form of graphical display. The fact that the user of the AR system would also be viewing the real scene, can be exploited in determining the form and content of situated graphics. The most dominant characteristic of situated graphics would be the simplicity of its contents. Simple line diagrams when superimposed on a real object can convey much of the same information that a more complex graphic display in a purely virtual world would. Similarly, short text labels are much more helpful than extensive text, when co-located with the related real objects. Sometimes situated graphics will involve a detailed rendering of a virtual object, for example, when visualizing a completed assembly by juxtaposing a partial subassembly with the virtual view of the rest of the assembly. At other times, more detailed situated graphics could actually interfere with the user's perception of the surrounding world.

Thus, situated graphics is composed and interpreted with respect to one or more objects in the real world. In the assembly domain, the composition is defined with respect to the individual parts and subassemblies. Since all possible combinations of the parts and subassemblies are known, the pointers to relevant situated graphics can be "stored" in the nodes/arcs of the assembly graph. Note that one or more of the objects, involved in the definition of situated graphics, could be moving. If the user is manipulating certain subassembly, some annotation may be required to be dynamically located next to that subassembly.

3.3 Augmentation through other modalities

In addition to situated graphics, the augmented display could contain other non-situated "modes" of presenting useful information. For these modes, the scene content would determine *when* to display information rather than *where*. These display modalities could include sound, animation, and video clips that aim to help with the assembly goals. Specific sounds can indicate to the user that a specific task has been completed, or notify the user of an error. A sampled voice could mention strategic words like "insert operation" to briefly convey the idea of what to do next. Cases where more than one or two sentences need to be spoken seem unlikely in the assembly domain. Similarly, short videos in a window are an efficient way of visualizing a complex task.

The objective of a well-designed augmentation scheme should be to provide just the necessary information in the most effective way possible, without confusing the operator. Additional stimuli like audio, video or even text will provide extra guidance information, but have to be part of a carefully designed augmentation scheme oriented towards maximizing and optimizing the information presented to the human. Thus, a "beep" could be more effective than a narrator explaining that the user has made an error. For explaining an assembly operation, a simple graphical primitive or animation can convey essentially the same information that a video clip could.

3.4 Augmentation for the next operation

At a given assembly state, the key purpose of the augmentation is to display the choices of the next step, or simply prompt the user to the next assembly operation. Each "assembly step" involves two parts or subassemblies




Figure 6: Displaying a choice in the assembly operation.

and a particular assembly operation. Each choice can be presented by displaying a rubber-band line or arrow between the two parts/subassemblies (Figure 6). This is an example of situated graphics; a simple line carries useful information because of its situatedness with respect to the two objects in the real world. In some applications this may be the most significant role of AR; for example, if the objective is just to correctly assemble an object.

3.5 Augmentation for the mating operation

The final stage of an assembly operation involves "mating" two subassemblies, one of which could be an isolated component. In robotics literature, determination of the robot motion required to bring about this mating is referred to as the fine-motion planning problem. Fortunately, for the domain that we consider, a detailed consideration of fine-motion is not necessary since humans are fairly adept at mating parts using high level "directives." Thus properly positioned graphical icons may provide sufficient information for establishing a liaison between the assembly parts. We analyzed some of the unit operations involved in assem-




Figure 7: Graphical primitives of assembly operations.

bly, and designed a corresponding set of graphical icons or *primitives*. Figure 7 shows some of a basic set of graphical primitives that we designed for guiding assembly. Note that the form of these assembly primitives is simple since



they are designed to be "situated" relative to the assembly, in the sense described earlier. A video clip or graphical animation demonstrating the mating operation can also be very effective, especially for a naive user.

3.6 Augmentation for sensor calibration

Augmentation in a system such as the one illustrated in Figure 1, deals with at least three major subsystem involving the camera, the assembly world, and the graphics system. In order for the subsystems to work together correctly, e.g., for the graphics to be situated at the right position, and for the vision system to estimate the assembly state correctly, the various coordinate systems need to be calibrated. The process of *calibration* can be tedious if it is carried out manually, and difficult if carried out autonomously. A good alternative is to execute the calibration as a "semi-autonomous" or "interactive" process, making use of situated graphics. For example, annotations can help with calibrating the graphics system with the real world. This can be done by placing a cubical object in front of the camera, and letting the user manipulate a graphically generated reference system in space with a mouse or a similar input device. The virtual reference system is moved until its origin and axes are aligned with respect to a corner and the edges of the cubical object.

3.7 The "flashlight technique" for sensing

Even with more than one camera to monitor the assembly, there will be times when a crucial part of the assembly scene is occluded. In general, dealing with occlusions is a hard problem in computer vision. However, we can exploit the interactive nature of the AR system, to overcome these and other shortcomings of the sensing system. In case a




Figure 8: Virtual "flashlight" to help in the inspection.

part is occluded, or when the vision system fails to verify a liaison, we could use the concept of "virtual flashlight," which is a colored cone or frame graphically displayed as emanating from the camera, as shown in Figure 8. This would be followed by a prompt (either visual or audio) to position the subassembly in the "light" cone. This helps the camera in forming an unobstructed view of the subassembly and hence in the identification of a part, or in the verification of a liaison.

3.8 The "pigeon-hole technique" for query

To facilitate interaction between the user and the AR system, another potentially useful technique is the *pigeon-hole technique*. This involves displaying a set of virtual volumes, e.g., boxes, where the user places a part or subassembly in

order to invoke certain kind of information from the system. There could be separate pigeon-holes for answering different types of queries by the user. Figure 9 shows an




Figure 9: Use of pigeon-hole to answer queries.

example in which the CD assembly is inserted in pigeon-hole; in response to this, the system could perhaps graphically depict the stages of assembly process that are still unfinished. The pigeon-hole technique could also be used to position a real object correctly with respect to a more elaborate model of the virtual object. The juxtaposition of the real and the virtual objects could, for example, help in visualizing the functioning of the assembly.

4 The AUDIT System

This section describes an interactive simulation system called AUDIT, that we have implemented. AUDIT is a tool for developing, visualizing, and evaluating AR presentation schemes for the assembly domain. Its ability to help in visualizing the geometric shapes of the subassemblies can also be used in other applications such as robot assembly planning, and in product design. Figure 10 gives an overview of the design of AUDIT. The main functional units are shown which participate in an interactive loop with the human operator. Before the system enters into

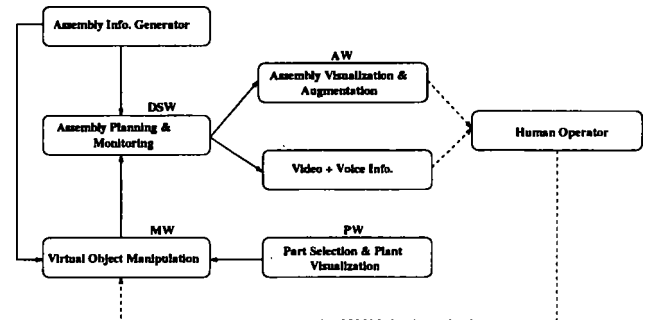


Figure 10: An overview of the AUDIT system.

the interactive loop, the Assembly Information Generator is used to build the necessary representation of the given assembly. The user provides the geometric models of the parts, the liaison graph, and additional information involving assembly feasibility. The assembly graph is then automatically generated and displayed from this information.

The four main functional components involved in the interactive loop are: (1) Assembly planning and monitoring using the assembly graph, (2) Assembly visualization and augmentation, (3) Virtual objects manipulation, and (4) Interactive part selection and plant visualization. For the sake of modularity, each of these functions are imple-



mented with the help of an independent window. The assembly graph is displayed in the Decision Space Window (DSW). The graphics augmentation information is shown in the Augmentation Window (AW). The solid part models are manipulated in the Manipulation Window (MW), and the components are graphically displayed in the Plant Window (PW) for the purpose of choosing the next addition to the assembly.

The four windows are displayed simultaneously but one can use the AUDIT system with only a subset of these windows. A typical session uses all windows. The user selects the parts in the PW, and manipulates them in the MW. The augmentation displayed in the MW explains which parts can be connected and how. If the user wants additional information about the assembly, and how to perform the next operation, the DSW is used. The DSW allows the user to gain a high-level understanding of the assembly process, including the assembly path being followed, and the future options. To help with the visualization, a computer animation showing the assembly operation is displayed in the AW whenever the user selects a node in the assembly graph. Other components of the system are video and voice annotations, which are also activated through the assembly graph. The video clips show a human operator performing different assembly operations, aided by the voice annotations, which provide additional useful information. The user then returns to the MW and completes the assembly step. We next describe in detail each of the main components of AUDIT.

Decision Space Window, DSW The DSW provides visual information about the assembly sequence through the assembly graph. It is also a secondary means of user interaction. The user can connect two subassemblies using the graph as the user interface by selecting one of the nodes in the graph using a mouse. The node selected must be a child of the previous state and it must not be an error node. The user can also obtain visual information about the assembly steps necessary to complete subsequent stages of the assembly by interacting with this window.

The operator starts the system by clicking on the root node and completes the assembly when he clicks on the bottom node. Each node clicked upon pops out, marking the current state. Also, the path leading to the current state is highlighted. Weights can be optionally assigned to each branch in the graph, so that the total cost of following an assembly path can be measured and the utility of different paths compared. There are two possible modes of interaction that the user can choose from, *connect* and *request*. In the connect mode, the new liaison, i.e., the liaison which is present in the current state, but absent in the parent state, is "executed" and the connection between corresponding parts is recorded. The parts of the completed subsequence are displayed in the AW, and the children of the current node (the next allowed states) are highlighted. The next action performed, whether it is in the DSW or in the MW, will lead to one of the children of the current state.

In the request mode, on-line augmentation information is displayed in the AW. This information is presented in several forms: 2-D and 3-D situated graphics, animations, sampled sounds and voice, video, and text. To request

information on a certain state the user clicks on the corresponding node, and the information is displayed in the AW, provided that the node belongs to a valid assembly sequence and is a child of the current state. The most important augmentation is an animation of the parts being fitted together to complete a liaison. An arrow will point to the direction of movement of the parts. A human voice provides the information about the type of operation being performed, and any other information that may be helpful. Also, text and labels are used when convenient and displayed in the AW. In this case, the assembly operation type is displayed in the AW, together with other relevant information. Video clips are optimally displayed using a different pop-up window called the video window (VW).

Augmentation Window, AW The AW is a *passive* window; it is controlled by the actions performed within other windows. The augmentation information includes visual and sound cues such as graphics, animations, voice and plain text, to display relevant information on how to perform a mating operation between two subassemblies. When the user completes an operation within the DSW or the MW, the AW shows the completed subassembly. When the user requests information on the next step to be taken (through the DSW) the AW will show a computer animation of how to perform the assembly operation, along with information about the type of link between the subassemblies being manipulated. A sampled human voice advises the operator how to connect the subassemblies by explaining the type of operation to be performed e.g., insert, snap, screw, push & twist, etc. Graphic primitives or icons are displayed in this window. In addition, the subwindow VW might be spawned off at various times to show a video clip.

Manipulation Window, MW The MW is an active window, like the DSW. The user manipulates virtual models of the real assembly parts, and connects them together using the mouse. The virtual parts are stored as geometric models of the assembly parts. The user can arbitrarily rotate and translate the objects by using the mouse. A graphical menu pops up when the right mouse button is pressed. The options include: clear screen, load objects, select object, rotate object, translate object. To select a part to manipulate, the user chooses an object from the plant model. Collisions among objects are detected efficiently without interfering with real-time manipulation. A liaison is "executed" whenever the 2 subassemblies being connected together are in a specific range of positions and orientations with respect to each other, which depends on the motion freedom that the connection allows to the objects.

The MW provides the main medium for user interaction. The current assembly state is highlighted in the DSW. The augmentation is displayed in AW, but the most relevant augmentation is provided in the MW on-line in the form of situated graphics and sound; these directly help with the assembly manipulation. The situated graphics include lines that connect the selected object with the related objects in the liaison graph, arrows that show where to move the selected object to, and information about the specific assembly operation to be performed. Graphical icons are attached to the arrows, thus providing the necessary information in a compact manner. In addition to the graphics, sound cues are also useful. They are used to



indicate a collision, a successful connection and to provide information about the type of operation.

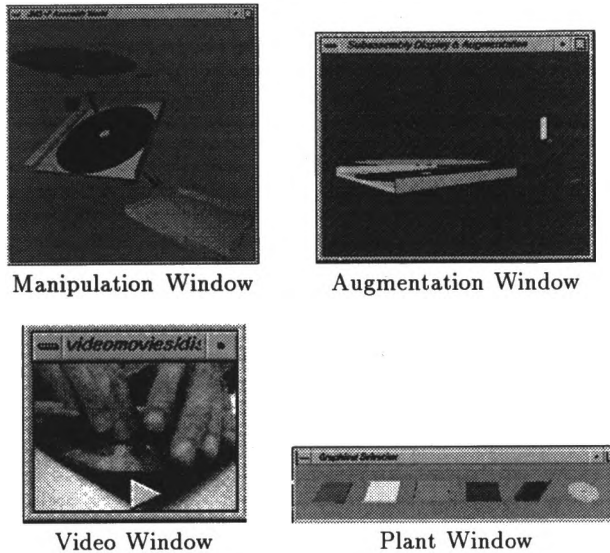


Figure 11: Some of the AUDIT windows.

Plant Window, PW An important aspect of the graphical simulation of assembly, is that the manipulation has to be visual and interactive. The selection of parts, then, also has to be graphical. It is awkward to select the objects directly in the MW as the objects might be occluding the desired part or subassembly, so a separate window, PW, is provided for selecting a part. For selecting a part, the user clicks over the view of the part in the PW using the mouse, and then continues with the assembly in the MW. This window could be generalized to other "plant" situations for part selection.

5 Specific Assembly Scenarios

We present several examples of how AUDIT can be useful in designing augmentation schemes for guiding assembly. A CD is a good example of an assembly, as it is familiar and contains several types of liaisons. The role of the augmentation at all times is to guide the user through the



Figure 12: Some steps in completing the assembly.

assembly. A more advanced operator will want simple and fast on-line information, whereas the unexperienced operator will need a more complete scheme of augmentation. Simple augmentation aids can help the operator identify the correct parts to pick and whether two given parts are related. They are also useful for preventing errors, hinting what type of assembly operation to perform, attracting attention towards the important features, etc.

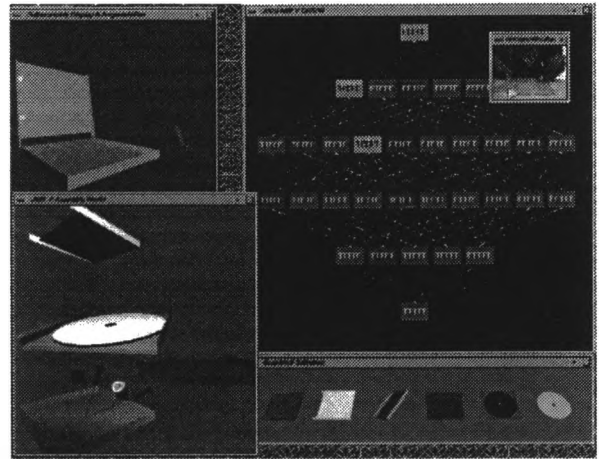


Figure 13: System windows working concurrently.




Figure 14: Dynamically situated augmentation cues.

Figure 12 shows the MW during various stages in an assembly sequence, that ends with the complete CD assembly. Figure 13 gives a snapshot of the AUDIT interface with all the windows running concurrently during the assembly sequence. The most relevant augmentation is presented on-line in the MW. Some examples are: lines joining active parts that are element of a subassembly, arrows pointing towards the direction of contact (activated when the objects are close), graphic primitives for the type of assembly operation to be performed. There are arcs and arrows linking parts that have to be connected together. The direction of the arrows indicate which part is being manipulated. They appear only if the two objects are close. The X next to one arrow in Figure 14 indicates that the liaison between the base and the frame cannot be established because the leaflet has not been connected to the frame yet, and if the operator connects the frame and base, backtracking will be required. Sampled sounds are used for indicating a collision between two objects and also to indicate when a valid connection has been established.

More complex information is provided via other visual and audio stimuli. The DSW, AW and VW present a more complete information that might be used for guiding an unexperienced user. The path followed, and an alternative way to visualize the choices for the next "step" is obtained through the DSW, which is always visible to the user. It is also a way of answering user queries. The children of the current state are highlighted, to present the set of next possible operations. If an assembly step is completed, the AW will show the current state by visualizing the con-



nected virtual parts in place. If the user inquires about a highlighted child in the DSW, the augmentation window will show a 3-D animation of the specific liaison being executed in the AW. The AW also indicates the type of assembly operation required, and a sampled human voice will read the type of operation. In addition, a VW can also be activated. The video will show an operator physically completing the connection. The AW and VW can be seen in Figure 11. Figure 12 shows the MW as the user completes the assembly.

6 Discussion

We have thus shown how the context of the "assembly world" helps in many different ways in developing a systematic AR presentation scheme. The context helps in defining an assembly state which is used to determine *what* to display. The assembly graph helps in building a control mechanism which, together with an appropriate sensing mechanism, determines *when* to display the augmentation. The use of the geometric models, together with an estimate of the current physical locations of the parts, determines *where* to display the augmentation.

The simulation tool, AUDIT, helps in developing the general AR concepts, as well as in visualizing the assembly at different stages. However, a given augmentation scheme will prove its usefulness only in a complete AR setting. One motivation for building AUDIT was to design the augmentation scheme, so that it can be transferred directly to an actual AR system. In fact, the authors have built a computer vision-based AR system called VEGAS that makes use of AUDIT as the augmentation engine [10]. In the implementation, the DSW in connect mode, is used to directly assimilate information about the assembly state changes communicated from the computer vision module. Moreover, AUDIT is incorporated directly into the architecture of the AR system by substituting the assembly simulation (the manipulation window) with assembly manipulation in the real world, and with scene understanding through an appropriate sensing module.

Another potential application of the AUDIT system is as a tool for product design, where the visualization of different aspects of the assembly process can be exploited. The design process can be viewed as a search through a multidimensional space in which the axes are different criteria like testability, durability, ease of use, etc. The ease of assembly/disassembly effects criteria like manufacturability and serviceability of the product. Visualizing assembly sequences, parts, subassemblies, and assembly representations would help in evaluating the ease of assembly of a particular design. The AUDIT system can be easily modified to support this design application.

An augmented reality system has to be ultimately evaluated by the user on the basis of subjective criteria like the ease of use, increase in efficiency, etc. Such evaluation would also involve cognitive issues in the display of complex, distributed information. For example, to judge whether the augmentation stimuli does not "clutter" the user's perception of the real world. The simulation tool, AUDIT, can be used for carrying out a preliminary cognitive evaluation of an augmentation scheme, even before the AR system is built.

7 Conclusions

For realizing the full potential of augmented reality, a critical issue is the development of a systematic augmentation scheme. We have developed such a systematic framework for the presentation of augmentation stimuli within the assembly domain by using concepts from robot assembly planning literature. The interactive system AUDIT was presented to further illustrate and utilize the AR concepts introduced within the context of the assembly domain. The system can be used in several ways. It can be used to design and evaluate the AR display for complete assemblies. The geometric rendering of the subassemblies helps in visualizing the different stages of the assembly. The intuition thus gained can be used to interactively update the assembly graph for subsequent use in AR. The field of augmented reality is as yet in its infancy; the analysis of the AR for a particular domain, like the assembly, should help toward developing better ways to exploit the novel human-computer interface.

References

- [1] J. A. Adam. Virtual reality. *IEEE Spectrum*, 30(10):22-29, 1993.
- [2] M. Bajura, H. Fuchs, and R. Ohbuchi. Merging virtual objects with the real world: Seeing ultrasound imagery within the patient. *Computer Graphics*, 26(2):203-210, 1992.
- [3] J. Bowskill and J. Downie. Extending the capabilities of the human visual system: An introduction to enhanced reality. *Computer Graphics*, 29:61-65, 1995.
- [4] D. Drascic. Stereoscopic video and augmented reality. *Scientific Computing and Automation*, 9(7):31-34, 1993.
- [5] S. Feiner, B. MacIntyre, and D. Seligmann. Knowledge-based augmented reality. *Communications of the ACM*, 36(7):52-63, 1993.
- [6] S. Gottschlich, C. Ramos, and D. Lyons. Assembly and task planning: A taxonomy. *IEEE Robotics and Automation Magazine*, 1(3):4-12, 1994.
- [7] L. S. Homem de Mello and A. C. Sanderson. Representations of mechanical assembly sequences. *IEEE Journal of Robotics and Automation*, 7(2):211-227, 1991.
- [8] M. W. Krueger. Environmental technology: Making the real world virtual. *Communications of the ACM*, 36(7), 1993.
- [9] R. Sharma and J. Molineros. The role of computer vision in augmented virtual reality. In *SPIE Conference on The Engineering Reality of Virtual Reality*, February 1995.
- [10] R. Sharma and J. Molineros. Computer vision-based augmented reality for guiding manual assembly. Technical Report UIUC-BI-AI-RCV-95-03, University of Illinois at Urbana-Champaign, January 1996.
- [11] M. Weiser. The computer for the 21st century. *Scientific American*, 265(3):94-104, 1991.
- [12] P. Wellner. Interacting with paper on the DigitalDesk. *Communications of the ACM*, 36(7):87-96, 1993.

