Augmenting Real Environments

Enylton Machado Coelho
Blair MacIntyre

GVU Center, College of Computing
Georgia Institute of Technology
Atlanta GA
{*machado, blair*}*@cc.gatech.edu*

Abstract

Augmented Reality technology has promised to change the way information is conveyed, creating an impact in activities as diverse as medical procedures, maintenance and entertainment. So far, this promise has failed to materialize. We claim that such failure is caused by the pursuit of perfection instead of usefulness, and propose the creation of systems that function in imperfect environments. In this paper we present the reasons for advocating focusing on usefulness, briefly revisit previous work [3] that laid the first step in this direction, and identify the issues that still need to be addressed. An specific example regarding the application of AR to the industry is used to illustrate these points.

1 Introduction

Augmented Reality (AR) superimposes computer generated graphics on top of physical objects. AR researchers have predicted that this technology will be used to improve medical procedures by displaying Magnetic Resonance Imaging and Computerized Tomography images over the patient, help maintenance personnel do their job by transferring hundreds of manual pages to the computer, support city planning and construction by allowing engineers and workers to *see* the infrastructure that exists below the surface before they dig, etc. But so far AR has been confined to research laboratories and only used in small experiments and demonstrations.

It is largely believed that Augmented Reality is not widely used because the technology is not mature enough: Head Mounted Displays are heavy and cumbersome, wearable computers are bulky and their batteries don't last long enough, trackers are inaccurate, creating a misalignment between the computer generated graphics and the physical world, etc.

It is interesting to note that other technologies that were not mature were deployed before they were completely developed. A prominent example is Microsoft's Windows interface system. Microsoft's Windows system was notorious for "crashing" unexpectedly and was much less reliable than its direct competitors. Nevertheless, users were inclined to accept Window's limitations because of its ease of use.

We argue that the reason technologies are adopted, even before they are fully developed, is because they add value. In the case of Augmented Reality, we claim that it will only be widely adopted by industry when it adds value – be it reduced cost, improved efficiency, etc. – not when this technology works flawlessly.

This paper is structured as follows: the next section presents registration error and argues that contrary to the belief of most researchers, there will never be perfect registration and what is necessary is an infrastructure that accounts for it. Section 1.2 presents an example to substantiate this claim. Section 2 introduces our solution: *Adaptive AR*. We then present how this solution can impact the way information is presented and dealt with. Specifically, we analyze how this could be employed in the industry. Sections 3.1 through 3.5 describe the major hurdles to be overcome to make this possible.

1.1 Registration Error

Most of the research carried out nowadays in Augmented Reality focuses on how to reduce the registration error. Specifically, since Holloway published his PhD thesis [4], where he identified and analyzed the sources of registration error in Augmented Reality and determined that tracker accuracy was one of the main contributors to registration error, research groups have actively been pursuing the Holy Grail of AR: perfect registration. At the 2004 IEEE and ACM International Symposium on Mixed and Augmented Reality, out of the eight sessions, three were on *Tracking and Registration* and one was on *Calibration*.

Registration errors can have a profound impact on the effectiveness of an augmented reality system. The purpose of many AR systems is to provide information to the user about objects by aligning graphics with those objects in the physical world. However, no AR system is perfect. Tracking systems cannot measure the pose of their sensors exactly. Internal system calibration parameters cannot be known perfectly and the world cannot be modeled precisely. As a result, the graphics will not align perfectly with the objects in the physical world. In some situations these errors can be little more than an annoyance. However, in other situations the annotations could be ambiguously placed (it is not clear what object they refer to) or appear to be placed on the wrong object altogether.

The conventional approach to registration errors is to consider them as a type of tracking problem. Apart from a few notable exceptions, none of them recent (e.g., [1], [10]), they are rarely addressed directly. The prevailing assumption seems to be that, given better tracking and faster computers, the major causes of registration errors will be overcome. However, we do not believe that this is the case, especially when one considers mobile augmented reality systems where one cannot rely on accurate, fixed infrastructure in carefully controlled settings.

We believe that a better approach is to assume that registration errors will be inevitable, and provide application developers with tools to help them understand and deal with these errors. In particular, we believe that any AR toolkit should also help developers choose and display annotations in such a way that the effects of registration errors are minimized.

1.2 Motivating Example

One interesting experiment at deploying computation systems to help maintenance tasks was the recent evaluation (by Honda and Microvision[1]) of a wearable maintenance system that uses a non-tracked see-through heads-up display to present *in-situ* automotive maintenance informations to trained technicians. This system was demonstrably useful (resulting in a quoted 38% improvement over the non-wearable version) despite the fact that the graphics were not registered with the physical world. This system raises some interesting questions for the AR community: would an AR version of the system be even more effective? Would precise registration be necessary, or would a coarsely registered version using moderately accurate tracking be as effective (e.g., by allowing the current system's graphics to be generated from the technician's approximate viewpoint, even if they aren't registered)? Would some tasks benefit greatly from precise registration, while others would not? When one considers practical issues of creating, deploying and maintaining such a system, such as cost and robustness, these questions become critical.

To continue with the above example, perhaps some repair shops would have good tracking, and others would not. Perhaps tracking quality would vary with the model of the car (e.g., new cars might have "embedded trackers", old ones would not), or the location in the shop or parking lot. Perhaps trackers break occasionally. For whatever reason, a commercially viable repair system such as this would need to function in a variety of situations, and ideally adapt automatically as the situation changes (e.g., as the technician walks around the shop). As an application designer, it is easy to imagine many different display modes for such a system, based on different amounts of registration error. What is hard is actually computing reasonable estimates of registration error for the different graphical objects in the system, and creating a graphical display system that uses these estimates to select the appropriate display modes. OSGAR [3] was designed to address this problem and it builds on our previous works [7, 6]. This library allows the development of *spatially adaptive AR* and we believe that this is the foundation for AR applications that can be deployed in real world situations.

2 Spatially Adaptive AR

Traditionally, augmented reality applications are developed so that they take into account the characteristics of the tracker and models used. This approach is feasible when one intends to create applications that will work only in predetermined environments. If one does not have complete control over the environment, and *a priori* knowledge of all the trackers and models, such arrangement may fail. Moreover, the causes of registration error are complex, combining the effects of tracking, modeling, and calibration inaccuracies. Requiring an application developer to reason over these combined contributions is impractical, especially when she really only cares about the end result: the registration error.

Our approach is to allow the application developer to focus on the end result, instead of concentrating on the tracker characteristics, i.e., allow the application designer to focus on how to best convey the intent of the augmentation under various levels of registration error, enabling the creation of a new class of Augmented Reality applications that function under a wide range of conditions, namely: *Adaptive Augmented Reality* applications. Every application that communicates information to the user can benefit from this novel approach, be it a maintenance application with images, animations and text or an outdoor tour guide that labels build-

[1] For more information, see `http://www.microvision.com/nomadexpert/field.html`

ings.

In the next two sections, we analyze problems with Augmented Reality applications as a user interface from both the application developer point of view and the user's view point. We then proceed to present our approach: *Spatially Adaptive Augmentations* (section 2.3), and comment on how it addresses both issues.

2.1 Application Developer's Perspective

From the application developer's point of view, Augmented Reality applications are fine tuned to make the best use of the specifications and limitations of the tracking technology available at the development phase. This is very similar to what happened in the first days of computer graphics, when one would develop applications for specific graphics hardware.

Nowadays, the use of an abstraction layer (provided by libraries such as OpenGL) allows one to develop device independent applications, decoupling the application from the underlying hardware infrastructure. Beyond simply providing device independence, such libraries allow the programmer to query the capabilities of the hardware and adapt to them. Similar kinds of abstractions are needed in Augmented Reality if this technology will ever leave the research laboratories and be put to use in real life situations.

Spatially Adaptive Augmented Reality applications allow the application developer to focus on the application itself instead of on the underlying tracking technology.

2.2 User's Perspective

From the point of view of the user, the main concern is whether an Augmented Reality application conveys the information it is supposed to provide. If the registration error gets in the way of the user understanding the intent of the augmentation, the system becomes annoying at best and confusing at worst.

In Holloway's PhD thesis [4], the main sources of registration error he identified were system latency, tracker inaccuracy, optical distortions [5]. Despite the fact that the sources of registration error have been known for a decade, the community is still unable to create an AR system where the amount of registration error, among other problems (such as low quality of the HMDs, bulkiness of the hardware, etc.), does not prevent it from being deployed in real situations. Although the causes of the registration error are well known, we are still a long way from creating a real-time mobile AR system.

For the last decade, the main approach to deal with registration error has been to try to eliminate the sources that cause it by employing faster computers and more accurate trackers. The key observation on which this work is based is that, for many AR applications, perfect registration is not necessary. From the user's point of view, what matters is if the system can convey the proper information, even if severe registration error exists.

Spatially Adaptive Augmented Reality applications concentrate on the quality of the information presented to allow proper information to be communicated under varying circumstances.

2.3 Spatially Adaptive Augmentations

Spatially Adaptive augmentations address both problems. Developers can create device independent systems that automatically select and display the best possible augmentation to communicate the intent established by the designer, without confusing or misleading the user.

Depending on the task to be implemented, different requirements exist. For instance, an application to support a medical procedure requires much higher accuracy than a computerized tour guide. The common approach has been to determine these requirements beforehand and pick the appropriate technology. The problem with this approach is that any change in the tracker characteristics, due for instance to interference, may render the application unusable. To prevent this, it is necessary to support the creation of Augmented Reality applications that are not intimately tied to the tracking system being used [2].

To address the issue presented at section 2.2 (User's Perspective), applications should display the most appropriate augmentation that conveys the intent of the designer but causes as little inconvenience to the user as possible [9]. This approach is preferable to displaying on the wrong location graphics that were supposed to register perfectly with the physical object.

It is our belief that for AR applications to be deployed in situations outside the laboratories, they need to be adaptive. The next section describes the motivation and illustrates the dividends and challenges associated to using this technology in one of its possible application: Industry.

3 AR in the Industry

Historically, in large industrial installations, such as power plants, offshore oil rigs, manufacturing facilities, oil refineries, etc., information and control is centralized. There are usually control rooms where the monitoring of the installation takes place and where decisions are made about how to run the factory, its production lines and/or equipment. The decision process is based on a hierarchical control structure, reminiscent of the administration models developed after World War II.

In the 80's, new management approaches started to modify the decision process a little by starting to decentralize it;

for instance, instead of using quality control at the very end of the production process to reject products that were considered defective, preventive measures were taken and employees were expected to fix the problems as they occured and even stop the production line as necessary. This was a step in the direction of decentralizing decision making, but the then current technology still required information to be concentrated in control rooms.

One of the drivers in today's manufacturing environment is the short life cycle of high volume products. To respond to this trend, companies have invested in flexible plants. Processes and production units are more and more automated. Computer systems are largely employed to monitor processes, schedule work loads, an so on.

The next step in decentralizing information and control will come when information is pervasive, rendering control rooms obsolete and allowing decisions to be made at any physical location where they are needed.

Augmented Reality allows information to be decentralized and presented in the context where it is better suited for decision making. Also, augmented reality is the appropriate interface for human interaction with the robots and the automated production line. It can also support communication between a group of people involved in managing the facility.

In order to support the distribution of information, there is a need to improve both the infrastructure and the software support for it. In the next sections, we examine the challenges that need to be overcome to make this vision come true.

3.1 Ubiquitous Tracking

Trackers need to be ubiquitous [8]. Industrial complexes will have tracking devices of various levels of accuracy throughout their physical installations. Not only people, but also assets need to be tracked. Productions lines will have information about each component and product. Appropriate software infrastructure will be responsible for collecting and distributing this information. Data from multiple trackers will be properly fused to provide a better estimate than each individual tracker could alone.

Trackers with different characteristics (accuracy, latency, etc.) will be installed through out the facility [2]. Depending on the accuracy required, the environment on which a device is installed, the technology used, cost, etc., trackers will have different characteristics and will report the accuracy of their estimates so that the application knows how to best adapt to the current condition.

3.2 Device Independence

Applications will be required to support and interact with this multitude of devices. The most appropriate augmentation will be displayed based on the available information. To function properly, applications need to be independent of the underlying infrastructure. It will interact with the appropriate level of abstraction that allows it to function in the presence of different hardware. For example, when one orders a document to be printed, the specifics of the printer does not need to be specified, it is abstracted from the application and taken care of by a lower level layer. The same is true in the case of AR: one does not want to deal with the tracker's specifics. All an application wants is to know the pose of the tracked object.

3.3 Proper Augmentations

There needs to be research to determine which is the best augmentation that conveys the intent of the application developer in the face of different amounts of error. The best way to augment a physical object is application specific and depends on the intent of the augmentation, but it should be possible to develop some general guidelines and rules. In order to create these guidelines for AR applications, there is a need to run various user studies where different augmentations are tested for various intents and applications.

3.4 User Input

It is not only the display of information that needs to be taken into account. Users will need to interact with the system, for instance to program the machines and direct the robots. In the presence of this multitude of tracker devices, the infrastructure used for information output needs to be leveraged to also take into account user input.

There is a need to advance research in a wide range of input techniques, be it manipulating a tracked device such as a stylus, one or two handed interaction, gesture recognition, multi-modal, etc. All these techniques are susceptible to error and thus suitable to be integrated to the framework that takes into account uncertainty.

3.5 Content Creation

Content creation tools need to facilitate the process of creating the augmentations. Most authors will not have a background in AR/VR or even computer graphics. Content creation tools need to be developed to be used by people who have a diverse set of skills. The creation tasks will be done by teams which will probably involve technical writers, animators, graphics developers, actors, script writers, and so on, and will be as complex as video game development. It will not always be possible for the author to have access to the conditions where the application will be run, so these tools need to allow the simulations of such environments for testing. Also, it needs to expedite the creation of multiple representations, possible by automating parts of the process.

4 Conclusion

The development of Augmented Reality applications that can be actually deployed require the use of adaptable AR, i.e. applications that choose the most appropriate augmentation depending on the expected registration error. In the case of industrial applications of AR, the amount of information available in industrial installations and the need to use flexible plants to respond to the demand of manufacturing short-life-cycle, high-volume products will require the decentralization of the decision making process. Spatially adaptive augmented reality is one technology necessary for this to happen. It will be necessary for trackers to report their characteristics, models to store their accuracy and content creation tools to make it easy for developers to create the appropriate augmentations.

References

[1] R. Azuma and G. Bishop. Improving static and dynamic registration in an optical see-through HMD. In *Computer Graphics (Proc. ACM SIGGRAPH '94)*, Annual Conference Series, pages 197–204, Aug. 1994.

[2] E. M. Coelho and B. MacIntyre. High-level tracker abstractions for augmented reality system design. In *The Int'l Workshop on Software Technology for AR Systems (STARS 2003)*, October 2003.

[3] E. M. Coelho and B. MacIntyre. osgAR: A scenegraph with uncertain transformations. In *International Symposium on Mixed and Augmented Reality (ISMAR 2004)*, pages 6–15, November 2-5 2004.

[4] R. Holloway. *Registration errors in augmented reality systems*. PhD thesis, University of North Carolina at Chapel Hill, 1995.

[5] R. L. Holloway. Registration Error Analysis for Augmented Reality. *Presence: Teleoperators and Virtual Environments*, 6(4):413–432, August 1997.

[6] B. MacIntyre and E. M. Coelho. Adapting to dynamic registration errors using level of error (loe) filtering. In *International Symposium on Augmented Reality (ISAR 2000)*, Oct 5-6 2000.

[7] B. MacIntyre, E. M. Coelho, and S. Julier. Estimating and adapting to registration errors in augmented reality systems. In *IEEE Virtual Reality (VR 2002)*, March 2002.

[8] J. Newman, M. Wagner, M. Bauer, A. MacWilliams, T. Pintaric, D. Beyer, D. Pustka, F. Strasser, D. Schmalstieg, and G. Klinker. Ubiquitous tracking for augmented reality. In *International Symposium on Mixed and Augmented Reality (ISMAR 2004)*, pages 192–201, November 2-5 2004.

[9] C. Robertson and B. MacIntyre. *Adapting to Registration Error in an Intent-Based Augmentation System*. Springer Verlag, 2003.

[10] A. State, M. A. Livingston, G. Hirota, W. F. Garrett, M. C. Whitton, H. Fuchs, and E. D. Pisano. Technologies for augmented-reality systems: realizing ultrasound-guided needle biopsies. In *Proceedings of SIGGRAPH 96*, pages 439–446, New Orleans, LA, August 4–9 1996.