

The Importance of Being Mobile: Some Social Consequences of Wearable Augmented Reality Systems

Steven K. Feiner
Department of Computer Science
Columbia University
New York, NY 10027
feiner@cs.columbia.edu
www.cs.columbia.edu/graphics

Abstract

What are the consequences of mobility for augmented reality? This brief paper explores some of the issues that I believe will be raised by the development and future commonplace adoption of mobile, wearable, augmented reality systems. These include: social influences on tracking accuracy, the importance of appearance and comfort, an increase in collaborative applications, integration with other devices, and implications for personal privacy.

1. Introduction

Over the past decade, the reality of mobile computing has begun to embrace the potential of wearable computing. In the process, several researchers have attempted to clarify what distinguishes wearable computing from mobile computing. Rhodes [10] suggests five criteria for wearable systems: portable while operational, needing minimal manual input, sensitive to the user’s surrounding environment, always on, and able to attract the user’s attention even when not actively in use. Mann [8] cites three desirable properties for wearable computers: situated physically such that the user and others consider it part of the user, controlled by the user (as opposed to, for example, an employer’s externally administered monitoring device), and having negligible operational delays. (His expansion of this last property is essentially equivalent to Rhodes’s last two criteria: always active, and always having some output channel that the wearer can perceive.)

Both of these sets of criteria formulate wearability in terms that attempt to differentiate wearable computers from non-wearable ones, in physical form and human interaction. And both are intentionally designed to rule out some systems that would otherwise fulfill a strict dictionary definition of the term “wearable”; for example, systems that rely

heavily on traditional manual input devices or are controlled by someone other than the wearer.

Just as these criteria have been used to frame discussions of systems that embody them, here I suggest and discuss some of the social consequences of mobility for wearable augmented reality systems. However, rather than adopting one of these specialized definitions, I am comfortable with the generic definition of “wearable” as meaning capable of being worn, with the implication that it refers to something that actually is worn. With similar generality, I use “augmented reality” to mean overlaying material (often, but necessarily, visual) on the user’s experience of the real world. Simple examples of augmented reality are routinely accomplished with displays that are neither head-worn nor head-tracked, such as vehicle instrumentation superimposed on the user’s view of the road through means of a windshield combiner. However, much augmented reality research concentrates on head-worn, head-tracked systems, because of their potential for providing personalized, spatially registered information to a mobile user [1]. Therefore, I will assume such systems to be typical of wearable augmented reality.

2. Social Consequences

2.1. Social influences on tracking accuracy

Augmented reality systems rely on tracking technologies to determine the position and orientation of the user’s head and parts of their body, and of other people and objects. Current approaches to tracking span a wide range of accuracy. For example, current indoor position-tracking technologies include the relatively high accuracy of tethered systems and fiducial-based vision [13, 15], and the coarse approximation of active badges [14]. Outdoor position-tracking systems are currently less accurate, ranging from centimeter-level real-time-kinematic GPS (Global

Positioning System) and submeter differential GPS [3, 5] to regular GPS and dead reckoning (accomplished, for example, with a pedometer, magnetometer, and altimeter [9]).

Hybrid technologies that incorporate vision may ultimately result in highly accurate tracking of users and what they can see indoors and outdoors [16]. However, no matter how accurate these technologies may become, social conventions may influence the accuracy with which we can track others, and, at times, even ourselves. For example, if several users are conversing remotely, then sharing tracking information among them could enable a better approximation of a face-to-face conversation, and therefore may become the “polite” thing to do. Even if the users are collocated, or can otherwise see each other directly, sharing tracking information could provide a literally more stable interaction by improving the overall accuracy of gestural references and adding information about temporarily occluded features. For example, a user may release information about her hand, so that others can see where she is pointing, whether or not her hand is obscured. Furthermore, if another user cannot see his own hand because of an obstruction and her computer can, then she may provide her own vision-based hand-tracking information to him and to others who also are blocked.

Outside of the conventions that may be imposed by direct conversation, the tracking information a user is willing to release and the parties to whom it is released may also depend on social protocols. In general, a person may be comfortable releasing their precise tracking information to a small, select group of friends and colleagues; may be willing to allow others to know their location with less accuracy (e.g., that they are in a specific room or building), and still others with much less accuracy (e.g., that they are in a particular city); and may perhaps be intentionally willing to deceive yet others. Similarly, the recipient of this information may also be more or less interested in the accuracy with which it is provided, if only to suppress some or all of it to decrease the clutter in their own virtual environment that would be produced by displaying it all.

2.2. Appearance and comfort

High-quality tracking, and the precise registration that it could make possible, might seem to be the most important technical issue confronting the development of practical augmented reality systems. There are some applications, however, that do not require registration. For example, suppose that a clock could be discreetly overlaid on one’s visual field whenever desired. This could make it possible to determine the time during a conversation, without the need for an embarrassing glance at one’s watch. Or consider being able to scroll through a set of notes while giving a talk, without breaking eye contact, essentially using the equivalent of a head-worn teleprompter. Informal conversations with many

people indicate their interest in having these capabilities, and the underlying computational framework needed is trivial. What is missing is a wearable display technology that is sufficiently attractive, comfortable, optically transparent, and inexpensive for most people to be willing to wear on a daily basis.

Despite the effort being directed at achieving precise registration, along with wide field of view, high resolution, and high brightness and contrast, it is unclear whether head-worn displays will achieve success in anything other than niche applications until appearance and comfort are adequately addressed. Even the common term “head-mounted display” is a telling one: someone may be far more interested in “wearing” a display, as they do eyeglasses or a hat, than in “mounting” it on their head. One promising direction is represented by the beta version of a head-worn display made by MicroOptical Corporation [11]. The optics for one version of this see-through display are embedded in a conventional eyeglass frame made by Safilo, a vendor whose advertisements are normally found in the fashion pages, rather than the technology pages.

2.3. Mobility breeds collaboration

When users are mobile, they no longer spend their computational time in front of a solitary desk. However, mainstream mobile computing currently supports only those collaborative activities within which it is appropriate to devote one’s attention temporarily to a laptop or hand-held device. In contrast, mobile augmented reality can make it possible for computation to be integrated with essentially any activity. Thus, new collaborative applications could support face-to-face conversation without the need to divert one’s attention to a PDA, let alone pause to flip it open.

Current mobile collaboration relies on error-prone conventional hardwired networks, explicit connections to wireless services, or tedious preparations to “beam” information between PDAs. Mobile wearable augmented reality systems will make it possible for active users to move into and out of the presence of others, as regularly and smoothly as we currently do when we are not computing. Bringing another user into the fold computationally will therefore need to be as easy and natural as including them in a normal conversation—just as we can sense another person’s physical presence, our systems should be able to sense their computational presence.

The technological underpinnings of such ad hoc mobile networks are already being addressed by research on capacitively coupled “personal area networks” [17] and by commercial initiatives such as the Bluetooth radio network specification [2]. However, user interfaces will need to change to take advantage of these capabilities. For example, inviting a new participant to join us might be accomplished with the same voice, face, hand, and body language that we already

use to initiate an unaugmented conversation.

2.4. Integration with other devices

No matter how good head-worn displays may become, there may well be other mobile and stationary displays with which users will want to interact. Combining together heterogeneous displays and devices can create “hybrid user interfaces” that benefit from the complementary capabilities of each [4]. For example, shared wall-sized displays may have the highest resolution available, true 3D volumetric displays may properly couple binocular vergence and accommodation, and large haptic displays may present certain kinds of information more effectively. If see-through (and, similarly, hear-through or feel-through) displays are used in conjunction with other displays, such as these, it could be possible for common information presented to many users on one or more shared displays to be privately overlaid with each user’s individual supplementary information.

Large numbers of users and of displays produce a challenging user interface problem—explicit “direct manipulation” control of such systems may be quite unwieldy. Even taking responsibility for determining the exact size and position of each window on a single unshared display, as supported in current window managers, can be onerous if there are many windows. Imagine, instead, having to maintain explicit control over hundreds or thousands of virtual objects, on dozens of displays—some personal and some shared—as a changing group of users and their displays move about. We have referred to this problem as “environment management” [7], in analogy to window management, and believe that well-designed knowledge-based approaches may help ease the burden.

2.5. Implications for personal privacy

Researchers have pointed out the privacy issues raised by individual wearable tracking devices [14] and individual wearable computers with audio or video recording capabilities [8, 12]. Refusing to wear a monitoring device, or insisting that a recording device be off when one interacts with its wearer, would appear to circumvent these problems, although this may not always be feasible. While the use of a single such system may be invasive, what are the consequences of a world in which networked personal wearable computers with recording devices are commonplace? Future pressures in this direction may be strong, given the potential advantages of wearable vision-based tracking systems that mandate cameras and wearable context-sensitive systems that rely on audio and visual information [12], decreasing size and cost of mass storage, increasing wireless network throughput, and the utility of maintaining a detailed personal diary [6].

Mann [8] paints the reassuring picture of a “safety net” of wearable computer users who watch out for each other,

monitoring information that their fellow group members transmit, and coming to their aid in case of danger. But, consider instead the following dystopian scenario: Suppose that some organization were willing to pay individuals a small, but adequate, sum to acquire real-time access to their recorded experiences. To make such an arrangement more attractive, access might be controlled as the individual saw fit; for example, automatically expurgating material captured at work and home, or interactions with the user’s friends and colleagues, effectively addressing that user’s concerns about maintaining some amount of privacy. An individual user’s material may be of little or no value by itself, but consider what might be done with the aggregate time-stamped and position-stamped recordings of a very large number of users.

Massively parallel image and audio processing could make it possible to reconstruct a selected person’s activities from material recorded by others who have merely seen or heard that person in passing. Imagine a private two-person conversation, recorded by neither participant. That conversation might be reassembled in its entirety from information obtained from passersby, who each overhead small snippets and who willingly provided inexpensive access to their recordings. The price paid for such material, and the particular users to whom that price is offered, might even change dynamically, based on a user’s proximity to events or people of interest to the buyer at that time. This could make it possible to enlist temporarily a well-situated user who may normally refuse access at the regular price, or to engage a user’s wearable computer in a “bidding war” among competing organizations, performed without any need for the user’s attention.

While unauthorized commercial use of the actual material reconstructed this way (e.g., appropriation of a user’s audio or video likeness) could face legal restrictions, merely acting on information derived from this material may not be illegal. Nevertheless, the possibility of having one’s actions traced through the sum of many small and unavoidable interactions is an unsettling prospect.

Imagine, for example, that someone were interested in the activities of a competitor at a large conference. “Tailing” the competitor oneself or hiring someone else to do it may be expensive and ineffective, especially if the competitor is cautious. In contrast, contracting a third party to harvest electronic material from nearby attendees and analyze it may make it possible to piece together significant portions of the competitor’s interactions. The only parts excluded might be interactions that were fully private or overheard only by observers unwilling to part with their information (perhaps those whose silence was already bought by the competitor). However, many people might overhear isolated bits and pieces of just a single conversation, and some may have already assigned exclusive rights to what

their systems record, and will not be able to be temporarily silenced. Thus, ensuring the confidentiality of even the most mundane interactions might become extraordinarily difficult.

Acknowledgments

Research on mobile augmented reality systems at Columbia is supported in part by the Office of Naval Research under Contracts N00014-97-1-0838, N00014-99-1-0249, and N00014-99-1-0394, and by gifts from IBM, Intel, Microsoft, and Mitsubishi. The ideas expressed here have benefited from conversations with many colleagues, including Andreas Butz, Tobias Höllerer, Blair MacIntyre, Steve Mann, and Thad Starner.

References

- [1] R. Azuma. A survey of augmented reality. *Presence*, 6(4):355–385, August 1997.
- [2] Ericsson, IBM, Intel, Nokia, and Toshiba. Bluetooth technology for wireless communication. <http://www.bluetooth.com>, 1999.
- [3] S. Feiner, B. MacIntyre, T. Höllerer, and A. Webster. A touring machine: Prototyping 3D mobile augmented reality systems for exploring the urban environment. *Personal Technologies*, 1(4):208–217, 1997.
- [4] S. Feiner and A. Shamash. Hybrid user interfaces: Breeding virtually bigger interfaces for physically smaller computers. In *Proc. UIST '91 (ACM Symp. on User Interface Software and Technology)*, pages 9–17, Hilton Head, SC, November 11–13 1991.
- [5] T. Höllerer, S. Feiner, and J. Pavlik. Situated documentaries: Embedding multimedia presentations in the real world. In *Proc. ISWC '99 (Third Int. Symp. on Wearable Computers)*, San Francisco, CA, October 18–19 1999.
- [6] M. Lamming and M. Flynn. Forget-me-not: Intimate computing in support of human memory. In *FRIEND21 Symposium on Next Generation Human Interfaces*, Tokyo, Japan, 1994.
- [7] B. MacIntyre and S. Feiner. Future multimedia user interfaces. *Multimedia Systems*, 4(5):250–268, 1996.
- [8] S. Mann. Smart clothing: The wearable computer and wearcam. *Personal Technologies*, 1(1), March 1997.
- [9] Point Research Corporation. PointMan dead reckoning module. <http://www.pointresearch.com>, 1999.
- [10] B. Rhodes. The wearable remembrance agent: A system for augmented memory. *Personal Technologies*, 1(4):218–224, December 1997.
- [11] M. Spitzer, N. Rensing, R. McClelland, and P. Aquilino. Eyeglass-based systems for wearable computing. In *Proc. ISWC '97 (First Int. Symp. on Wearable Computers)*, pages 48–51, Cambridge, MA, October 13–14 1997.
- [12] T. Starner. *Wearable Computing and Context Awareness*. PhD thesis, Program in Media Arts and Sciences, MIT, Cambridge, MA, June 1999.
- [13] A. State, G. Hirota, D. Chen, W. Garrett, and M. Livingston. Superior augmented reality registration by integrating landmark tracking and magnetic tracking. In *Proc. SIGGRAPH '96*, pages 429–438, New Orleans, LA, August 4–9 1996.
- [14] R. Want, A. Hopper, V. Falcao, and J. Gibbons. The active badge location system. *ACM Trans. on Information Systems*, 10(1):91–102, January 1992.
- [15] G. Welch and G. Bishop. SCAAT: Incremental tracking with incomplete information. In *Proc. ACM SIGGRAPH '97*, pages 333–344, Los Angeles, CA, August 3–8 1997.
- [16] S. You, U. Neumann, and R. Azuma. Hybrid inertial and vision tracking for augmented reality registration. In *Proc. IEEE Virtual Reality '99*, pages 260–267, Houston, TX, March 13–17 1999.
- [17] T. Zimmerman. Personal area networks: Near-field intrabody communication. *IBM Systems Journal*, 35(3&4):609–617, 1996.