INVITED PAPER

Here We Are! Where Are We? Locating Mixed Reality in The Age of the Smartphone

This paper is about mixed reality and how the real-and-synthetic-scene hybrid is moving from the laboratory to the real world, technology to capability, space to place, and vision to perception.

By Evan Barba, Blair MacIntyre, and Elizabeth D. Mynatt

ABSTRACT | Mixed reality (MR), from location- and context-aware systems through immersive augmented reality applications, has been studied in research labs for more than two decades. With the advent of the modern smartphone, and the rapid adoption of wireless broadband technologies, this research is moving out of the lab and into the real world. Beyond just providing a platform for MR to reach a broad range of users, the capabilities and limitations of the smartphone provide a set of constraints on what kinds of applications can and cannot be delivered to consumers that have fundamentally changed how researchers think about MR. We examine previous definitions of MR, and reimagine the term as a class of experiences occurring in an ecosystem consisting of the smartphone, the cloud, and the user. Using a selection of current MR applications as a lens, we identify three key areas of continuing evolution and suggest how the development of a next-generation MR environment can help us channel the future growth of MR in research, industrial, and consumer communities.

KEYWORDS | Augmented reality (AR); mixed reality (MR); smartphone

I. MOTIVATION

After more than two decades of research, augmented reality and mixed reality (ARMR) technologies are now available to mainstream smartphone users. Applications such as Google Local Search or Yelp's Monocle represent the first wave of consumer technologies that aggregate data from sensors and network services, and present contextually relevant information that integrates with users' current activities and surroundings. While the power of the smartphone and the connectivity of high-speed and broadband wireless have made ARMR technically feasible in consumer devices for the first time, there are still limits. The ubiquity of the smartphone is owed, in part, to its emergence as the "Swiss army knife" of handheld computing. It is capable of many things, but ideal for none of them, and so embodies a series of tradeoffs that make it a useful, but not ideal, platform for ARMR research. It is clear, however, that the smartphone and the kind of mobile computing it enables is here to stay, and so it behooves us to look at its abilities and limitations in regard to ARMR, to both understand how best to exploit these and how the platform might continue to evolve as ARMR applications become more popular.

II. BACKGROUND

An oft-cited definition of mixed reality (MR) was published by Milgram and Kishino in 1994, and defines MR in terms of a "virtuality continuum" [1]. The continuum consists of combinations of real and virtual elements (mixed in different proportions), excluding only purely physical and purely virtual realities at either end of the spectrum. This definition was formulated at a time when mobile computing was in its infancy and virtual reality (VR) was still popular, and well before the paradigm-shifting growth in mobile computing seen in the last decade. It therefore inherits some of the assumptions about the future and nature of computing common to that era. In particular, it focuses heavily on the use of head-worn, 3-D, display technologies (e.g., video see-through, optical see-through,

Manuscript received September 7, 2011; revised September 28, 2011; accepted October 20, 2011. Date of publication February 20, 2012; date of current version March 21, 2012.
The authors are with College of Computing, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: ebarba3@gatech.edu; blair@cc.gatech.edu).

Digital Object Identifier: 10.1109/JPROC.2011.2182070

Fig. 1. ***Early examples of MR displays and scenarios. (a) Custom built head-mounted display (courtesy Department of Computer Science, University of North Carolina Chapel Hill). (b) Application for creating airplane wire bundles (courtesy Mizell/Janin). (c) Touring machine mobile AR application (courtesy Columbia University).***

etc.) to mix virtual content with the user's view of the world around them. Furthermore, it classifies MR based on different ways of visually and spatially mixing physical and virtual content using those displays.

Soon after Milgram and Kishino contributed their definition of MR, Azuma provided a simple and specific definition for augmented reality (AR) [2], which he took to denote interactive spaces created through the 3-D registration of computer-generated imagery with a user's view of the physical world around them. Within this tradition, AR was taken as a subset of MR, rather than a wholly separate field, and so contains many of the same assumptions, especially the reliance on visual sensing and display, that was the focus of early MR.

In 1998, Mackay offered an alternative notion of AR [3], one that focused on the user experience rather than the visual display. While her view attempts to encompass the same kinds of AR that can be found in the MR continuum, the kinds defined by Azuma, it also includes other categories of AR as well. Augmenting the environment, or objects in the environment, to make them networked and interactive is seen as a workable alternative to relying on user-equipped displays. These techniques include things such as networked paper and other "smart" objects, as well as algorithmic analysis of video or other sensor data common to many other forms of MR. In this tradition, AR was understood to connote technologies that took everyday objects and activities and enhanced them in some way. Making activities more collaborative or more personalized, and giving objects memory or awareness, were all ways of "augmenting" user experiences.

Not surprisingly, this definition also builds upon its own historical tradition, and is a product of its time. As the VR vision of computing was fading into the background, a new vision of ubiquitous computing (UC), offered by Weiser [4], was taking its place. This view was, in many ways, the antithesis of the VR view. Instead of inserting ourselves into the virtual world of the computer, UC had us inserting computers into everything around us. UC saw the world as a rich environment of hidden information and capabilities, waiting to be made available to us, and responding to our needs, both hidden and obvious.

Fig. 2. ***Video mosaic was an example of one of Mackay's AR applications (courtesy Wendy Mackay).***

It is clear now that neither of these visions for the future of computing have come to pass, at least not in their extreme versions. Mass-market technology has evolved quite differently from what was imagined by researchers a decade ago. Instead of donning head-worn displays, gloves, and other sensors to enter virtual worlds, or walking through interactive forests of smarter versions of everyday objects, the sensing apparatus and computing power has been compressed into the smartphone, and the intelligence of the environment is piped in "on-demand" from the vast data and processing facilities of the cloud.

Although the question of why computing evolved the way it did is beyond the scope of our work, the implications of this evolution for the future of MR, which we believe is the future of computing, is something we care deeply about; and so, we want to take a moment to speculate about how we arrived here. One likely possibility is that consumer expectations (fed by imagination and pop culture) about what immersive 3-D and smart object technologies could deliver far outpaced what the technology could deliver, and the half-way-there technologies (not just of the 1990s, but even of the present) do not quite measure up to those expectations. A second possibility might be that users simply do not have much use for immersive 3-D everything, or smart lamps and desks, when they do not fit our lifestyles, meet our needs, or provide any perceived benefit. Some combination of both reasons is likely.

The point of this speculation though is not to determine why VR or UC did not come to pass, but rather to suggest that what did come to pass, the age of the smartphone, is as much a function of user needs and expectations as it is of technology. Therefore, we believe that the success of MR in this new era of mobile computing will also rely on understanding three components of a larger socio-technical ecosystem: the smartphone, the cloud, and the user.

In order to recalibrate our conception of ARMR for this new paradigm of computing, there are some elements of past definitions that we want to take with us as we move forward. Naturally, the basic notion that all ARMR aims to integrate physical and virtual elements into a new hybrid reality goes without saying; but there are more subtle ideas. One distinction that we have found useful for both discussion and research is the notion of AR as a subset of MR. This relationship follows Milgram and Kishino's definition, and the definition clearly defined by Azuma to describe the specific technique of visually aligning virtual content with the view of the physical world, in 3-D. This approach has been implemented in many current smartphone apps, and so we prefer to reserve its use for these, visual and display focused, interfaces.

MR, on the other hand, is a term that we feel needs to be broadened beyond its original use by Milgram and Kishino to include many of the forms of augmentation discussed by Mackay. There are many ways to mix realities beyond the purely visual, and many of these can be found in the smartphone ecosystem. We therefore suggest that the most rhetorically and descriptively useful way to look at ARMR in the age of the smartphone is as follows: for a definition of MR, we look to Mackay; for a definition of AR, we look to Azuma; and Milgram and Kishino describe the relationship between them.

III. LEVELING UP

Updating the notion of MR from its beginnings as a collection of graphics and display technologies to its current role as a mode of technological experience is actually fairly straightforward. This evolution is already happening in the world around us; we need only observe and describe it. We have come to see change as an adaptation and expansion of three central concerns in traditional MR research. The first is MR's focus on "vision," by which we mean the predominance of displays, graphics, cameras, and the effects of these on the human visual system. A second fundamental component of traditional MR is "space," meaning the proper alignment and registration of objects, physically and conceptually, within a given environment, which is essential to creating a convincing MR. The final aspect is the technology itself, the apparatus and processes that comprise the MR systems, and which have, until the smartphone, been custom and specifically built for MR.

These three core themes of MR research are the seeds of our knowledge of MR, and they can help us grow our understanding with a little effort. We therefore suggest that each of these core concepts of MR can be adapted for investigation of the smartphone ecosystem in the following way: *vision becomes perception, space becomes place*, and *technologies become capabilities*. We will discuss each of these in turn.

A. Vision Becomes Perception

MR has always been concerned with notions of perception and representation in hybrid physical–virtual spaces. Although the traditional focus of this is almost completely on visual representations, a broader notion of MR needs to look at perception as a function of all the various processes of human cognition, rather than vision alone. There is, after all, much more to comprehending a scene, and the elements and relationships in it, than can be accounted for by what is seen.

Information is perceived and mixed in a number of ways, for a number of purposes, and at various scales, ranging from tabletop surfaces to geographic regions [5]. One form of small scale MR, tabletop AR games [6], works very well technically [Fig. 3(a)], and this functionality has been added to the latest generation of handheld gaming devices such as Nintendo's 3DS. Likewise, at the very highest scale, MR-enabled experiences such as geocaching or simple GPS navigation use 2-D representations of virtual information that are adequate enough for it to provide some considerable benefit in the physical world [Fig. 3(c)].

Fig. 3. *(a) NerdHerder handheld AR game (courtesy Georgia Institute of Technology). (b) Monocle mode in Yelp (courtesy Georgia Institute of Technology. (c) A geocache (courtesy fhwrdh).*

However, the "icons-on-video panning-and-scanning" applications in-between these two extremes, such as the Monocle function embedded in Yelp's iPhone application, tell a different story [Fig. 3(b)]. Monocle is essentially an AR mode of the location-based search tool in which listings for local restaurants are displayed in direction-oriented bubbles rather than as a simple list. For sure this mode of the application garnered much attention, not least because of the way it was cleverly "hidden" by the developers to sneak it into the iTunes App Store (you had to shake the phone to activate it), but one must question the utility of such a feature.

The task at hand for any Yelp user is to find a location to eat, presumably in the next few minutes. This is a task that requires more cognitive operations on any given representation than a floating icon can accommodate. Environmental perception must be integrated with information about personal and group tastes, physical states such as fatigue and hunger levels, and other concerns such as cost and time. Searching a directory of local establishments within a walk-able radius, comparing reviews, getting directions, wait times, and/or menus, are all mission critical features that enhance the user experience. This information can be represented and manipulated most effectively in a list, which supports tasks such as searching and sorting based on many of the needs noted above. Having the data superimposed over a view of the world, in the general direction you need to walk, is of dubious added value. This value is further called into question because the accurate alignment of physical and virtual information, which is the core feature of this mode, is so error prone. Often labels float and move around and appear to be aligned with nothing in particular. This lack of visual alignment then hinders any deeper perceptual or conceptual alignment, creating more confusion than clarity. Of course, the Yelp developers very likely recognized this, and added the Monocle mode as a tantalizing look at what the future might hold, rather than making it the primary mode of representation.

This example points out a critical issue with MR, and AR in particular. The kinds of data we have to represent, and the means with which we represent it, need to match both the requirements of the task and the capabilities of the technology. All hybrid spaces have distinct characteristics that make them more suitable for some kinds of human activities, yet sometimes incompatible with others. What is more, although many activities that MR techniques have been applied to (e.g., navigation or furniture assembly) have visual components, the tasks themselves are not primarily visual. For MR research then, the challenge is to move beyond toy demonstrations of what we might use ARMR for, and begin to focus more explicitly on understanding the elements of perception and representation that are needed to accomplish each of these tasks. Such elements are necessarily internal, to both the human mind and the technological system, and external, or out in the world (physically and virtually). Knowing how to blend and adapt information, on an activity-by-activity basis, and how to represent it with the technologies on hand to help users accomplish their goals, are going to be critical skills for all MR designers and researchers alike.

B. Space Becomes Place

The pace of technological change and adoption has never been faster, but people also grow and adapt along with technology. It is common now for any individual to store and access information on the cloud using a smartphone. Whether e-mail, social networking, simple web queries, or fully connected utilities and data storage mechanisms such as Dropbox, users are accustomed to having personal (and personally meaningful) information accessible wherever and whenever they want. While some of this technology has allowed for the notion of "placeness" to disappear from computing (e.g., you can no longer guarantee that a person is at home or work when accessing e-mail), it has also simultaneously moved placeness and context to the forefront of users' wants and needs. Where computing once demanded that users be in a certain place to access information (a terminal), it is now the users who demand that the information they access be relevant to their place and current needs. This is a tectonic shift whose reverberations have only begun to be felt.

If asked to choose one defining characteristic of MR, it is likely that most people would key into its ability to create hybrid spaces; but there is more to creating a *place* than the space itself. The geographer, Yi Fu Tuan, captures this idea nicely in his seminal and aptly titled work, "Space and Place" [7]. Although it does not do Tuan's work justice, one might summarize the critical distinction between space and place by saying, "place is space with meaning." HCI researchers have used this same distinction to argue that seemingly inconsequential elements, such as furniture and artwork, are actually vital to effective collaboration and telepresence technologies [8] because they help create a sense of place that frames appropriate behaviors. It is also a distinction that can help us adapt MR research for the future.

The notion that location and place can be used as computational inputs was a motivation for Weiser's conception of UC, and is beginning to catch on. Existing cloud-based services, (Flickr and Google Maps) have leveraged their geolocated content, in the form of geotagged photos and search data, to provide simple MR capabilities. Twitter, Facebook, and other predominantly social applications have also begun to include location-based capabilities in the form of geotagged tweets, and "Facebook Places." Newer services, such as Foursquare, have been designed specifically to leverage the location capabilities of these devices.

The evolution from existing web-based content accessed on mobile devices, to web-based services tailored to mobile users, and then to web-based services that are inherently built for location awareness and context awareness, has highlighted a crucial problem for MR researchers. Fundamentally, the kinds of data we capture and the way we structure and access it need to evolve. It is one thing to make a webpage that users will access from their desktop; it is another to format that data for mobile devices; and it is still another to customize that data, and its representation, moment by moment, user by user, and place by place. For the time being it is enough to view tweets and images in the places they happened, but as these data grow we will need smarter tagging and search technologies that are capable of sifting through mountains of in-place data to find what is relevant to any given context.

Currently, *ad hoc* technologies such as geocoded data are enough to provide one level of filtering, but these technologies were not built for the massive pools of data and increasing demands for relevance that individuals are placing on them. New data structures and storehouses must emerge to handle this place-centric data. However, storing and accessing data is actually the easier part, as it is largely a known problem and can be tackled with brute computing force. The more difficult challenge is going to be determining what data are relevant, how to collect data, when to retrieve data, and how to represent data when we are done. These are the elements needed to convert a space into a place; they are fundamentally questions of understanding humans rather than technologies, and we need architectures to support them. To a degree, we can continue to rely on and adapt existing technologies and strategies, but hybrid realities require hybrid information, and at some point these issues will need to be tackled explicitly.

C. Technologies Become Capabilities

In the past, MR researchers were forced to use expensive and cumbersome equipment as testbeds for their experiments [9]–[11]. These systems were often purpose-built for specific tasks or scenarios and demonstrated key MR functionality. However, the mainstream MR that we are seeing today is being done with devices that, while capable of delivering an MR experience, where not explicitly designed for them. This makes these devices less than ideal for many of the things MR researchers have studied in the past, but opens the door to new lines of inquiry.

Despite its power and unprecedented integration, the smartphone is still many generations from delivering "canonical" MR experiences in the seamless, automated, and consistent way that early researchers envisioned them. For sure, there are tremendous success stories. Most automobile drivers cannot remember or imagine what they did before GPS mapping and navigation, and the convenience and immediacy of location-based search has enhanced the day-to-day value of computing for many people. There have also been moderate successes, with apps such as Foursquare having small and loyal followings, but not catching on with a critical mass of users.

There are also some apparent failures in this domain. AR applications have been pushed very hard into the consumer experience by those hoping to capitalize on their novelty, but have largely failed to be recognized for more than that novelty. Countless smartphone applications are able to virtually add information around the user, overlaid on their view of the physical world as seen through the "magic window" of their smartphone camera, using no more than the sensors built in to the device (compass, GPS, accelerometers, gyroscopes). We have already discussed one critical problem with this approach in the example of Monocle; the sensors are of such poor quality (relative to the technical requirements for AR) that the virtual information does not appear to be integrated with the physical world. Contrast the typical sensor-based smartphone app with the level of integration seen during sporting events on television: the first down line in American Football and the flags in the lanes of the pool during the summer Olympics are virtually impossible to distinguish from reality. This is a limitation that will take multiple generations of technology to eliminate, yet the expectation of it being close at hand remains.

While total visual integration is a long way off, computer-vision-based tracking offers the possibility of tight spatial registration. Real-time tracking of 2-D images is possible on today's smartphones, and soon, 3-D objects and outdoor architectural scenes will be trackable, making stable AR possible. Unfortunately, the distinction between applications that require graphics to be tightly registered with the user's view of the physical world, and those that can provide benefit without tight registration, is often unclear to those without experience creating such applications. The pop-culture dream of immersing users in a hybrid physical/virtual world is tantalizing, but the combination of limited knowledge of the physical world with the relatively low accuracy of GPS means that some MR applications attempt to do much more than they can currently deliver. So, it is little surprise that some forms of MR have not caught on the way many imagined they would, not yet at least. Designers and researchers alike need to accept the fact that although smartphones are *capable* of MR, they are not yet ideal for many forms of MR.

Understanding the limitations of the multipurpose technologies built into the smartphone ecosystem is clearly one element of thinking about MR capabilities rather than MR-specific technologies. Another has to do with the additional and simultaneous demands being placed on those technologies by their users. In contrast to earlier purpose-built MR systems that did MR extremely well but were not called upon to do much of anything else, today's MR-capable devices are asked to accomplish quite a bit more. Not only do these devices receive calls, texts, and e-mails in the middle of navigation and search activities, as examples, but they also need to be able to integrate these activities with MR when the context is appropriate. For now this functionality is left to the user, who switches in and out of apps and decides when and where to integrate data into various activities. However, we are fast approaching the limits of this model both in terms of the cognitive capacity of users to context shift, but also in terms of what kinds of MR experiences this model can support.

It is becoming increasingly clear that future generations of the smartphone ecosystem will need new interaction and application models to facilitate new forms of communication and meet increasing user expectations. To continue using our example above, usefully presenting AR information in the panoramic space around a user, requires that it be aligned rigidly with their view of the physical space. Tight alignment at this scale requires fusing the sensors in the devices with computer vision processing of the camera's video [12], [13], which is beyond the capabilities of current technologies. It is clear the capabilities of the smartphone will evolve, but for now both researchers and developers need to focus more on the capabilities of ARMR technology, and treat these as tools that could be added to their applications, rather than thinking of their applications as "AR applications" or "MR applications" in the whole.

IV. ARGON: A MODEL FOR MR APPLICATIONS

Over the past few sections, we have talked about the issues and potential of MR on smartphones. Beyond individual

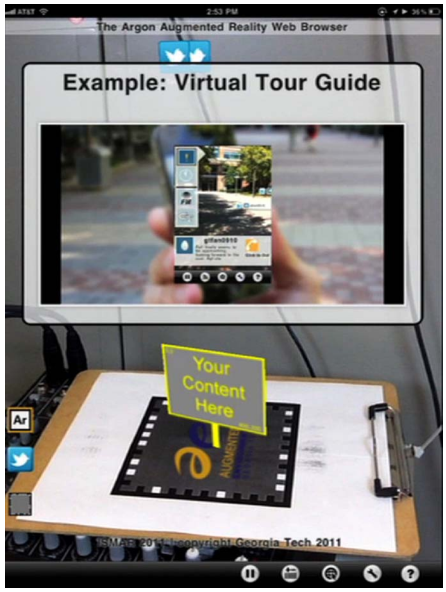

Fig. 4. ***Examples of using Argon for local search and marker-based tracking (courtesy Georgia Institute of Technology).***

application and cloud service requirements and capabilities, MR does present some more significant challenges to the smartphone ecosystem. A foundational idea of MR is that of immersing the user in a mixed physical/virtual world, where relevant virtual information is mixed into the world around them.

First, the application-centric model of modern smartphone ecosystems (e.g., iOS, Android, and Windows-Phone) is problematic. While much non-MR data are reasonably partitioned by the applications that manipulate and access these data, MR data are more logically partitioned by location and context. When users do a mobile search, or look at the status of a nearby location-based game they are playing, they should reasonably expect to see all relevant virtual content at the same time. While each of the bits of nearby virtual information is interesting, none forms a "killer application" for MR; together, however, they may form a kind of "killer experience" that makes the MR ecosystem thrive, with each new service or kind of information added seamlessly to the whole.

A second problem is the physical form factor of the phone itself. Implicit in many MR ideas is the potential for serendipitous information discovery. As users move through their daily life, they discover things of value that they did not know to search for (e.g., a nearby sale on some item they wanted, notes from friends, geocoded snippets from social media, hidden historic tourist information). Such serendipity is impractical with the current smartphone: nobody will periodically stop, take out his/her phone, unlock it, and run a sequence of MR applications hoping to discover something interesting. Even if the smartphone is out and unlocked, holding the phone up to look through" an AR application is awkward, physically and socially. To enable such serendipity and more comfortable interactions, in addition to having all data in one application, new forms of display (such as head-worn displays paired with a user's smartphone) must be developed.

To experiment with ARMR applications on smartphones, we have developed Argon, a mobile ARMR client based on modern web technologies [14]. Argon is design to address many of the issues raised in this paper, and allow ourselves and others to experiment with ARMR applications. While we cannot describe Argon in detail here, we will highlight how some of the design decisions we made relate to the issues discussed in this paper.

Argon takes the idea of a web site, and turns it into the notion of a *channel* of content. The content is hosted on standard web servers, and communicated to Argon using standard web protocols. However, because an MR application is not just a single 2-D page, but a combination of 2-D information on the screen with information in the world around the user, an Argon channel uses a modified version of KML (the markup language for Google Earth and Google Maps) as a container for multiple HTML5 content elements that live on the screen and in the world around the user. The long-term goal of the channel architecture is to bootstrap the creation of standards for MR data markup on the web, supporting the creation of rich MR web services and applications.

A second key design feature of Argon is that multiple channels (authored independently, and living on different web servers) can be displayed in parallel. Architecturally, this is accomplished by rendering each channel in its own HTML web view, and compositing these views over the video from the camera phone. The long-term goal of this architecture is to give users the experience of interacting with all of their MR content in one place, rather than in separate applications. Furthermore, we hope to encourage developers to design content with the idea that it is viewed and interacted with as part of an MR ecosystem, not by itself, leading to a focus on the user experience of adding the developers content to their "MR view."

A final goal of Argon is to create an easy-to-use development platform that supports the full range of capabilities needed by ARMR-enabled mobile applications. One of the problems with developing ARMR applications for smartphones is that it is difficult to deal with the technologies and create compelling applications on top of them. By merging ARMR technology into the web ecosystem, and enabling current web developers to experiment with ARMR ideas, we hope to push beyond the current limited uses of ARMR and discover the true power of the technology.

V. HERE WE GO

We are in the middle of a very exciting stage in the growth of MR, and it is natural that developers experience some growing pains. MR experiences delivered on smartphones have had some obvious successes, some mixed ones, and some failures. However, there is one success story that is easy to overlook, particularly if one is focused closely on the technology. Only when one considers the larger human and social dimensions of MR research, particularly in the context of other historical "game-changing" technologies, does this contribution become clear.

What seems to have really caught on about MR, the impact that two decades of MR research has clearly had on the mainstream consumer, is an enthusiasm for MR and the expectation of its success. Clearly, it is impractical to use many current ARMR apps, and so many people do not. But the fact that you can download an app for your smartphone, from virtually anywhere, that lets you see contextualized information in place is no less than amazing. For the first time users can not only see this advanced technology, but also they can hold it, use it, experiment and develop with it, and experience it for themselves in a way that they could not with the computing paradigms that led to the current MR ecosystem, such as VR or UC; there is definitely something real about the reality of MR.

This also means that MR research will continue to evolve. As researchers, we must stop imagining what is possible and trying to build it in isolation, but instead look at what is being built and try to understand it and influence its direction. That is what we are doing with the Argon platform, both by presenting a new model for the delivery of ARMR applications and also by making the technology accessible to millions of web developers.

Infants need their parents to provide for them, but adolescents need their parents to listen to them. When MR was in its infancy researchers were just trying to keep it alive and help it grow. Now that it has reached adolescence we need to look at how it is interacting with the world. We need to see where it is going and either do what we can to help it get there, or do what we can to redirect it toward more constructive goals. ■

REFERENCES

[1] P. Milgram and F. Kishino, "A Taxonomy of mixed reality visual displays," *IEICE Trans. Inf. Syst.*, vol. E77-D, no. 12, pp. 1–15, 1994.

[2] R. T. Azuma, "A survey of augmented reality," *Media*, vol. 4, pp. 355–385, Aug. 1997.

[3] W. Mackay, "Augmented reality: Linking real and virtual worlds: A new paradigm for interacting with computers," in *Proc. Working Conf. Advanced Visual Interfaces*, 1998, pp. 13–21, DOI: 10.1145/948496.948498.

[4] M. Weiser, "The computer for the 21st century," *Sci. Amer.*, vol. 265, no. 3, pp. 94–104, 1991.

[5] E. Barba and B. Macintyre, "A scale model of mixed reality," in *Proc. Conf. Creativity Cogn.*, 2011, DOI: 10.1145/2069618.2069640.

[6] Y. Xu, E. Barba, I. Radu, R. Shemaka, and B. MacIntyre, "Pre-patterns for designing embodied interactions with handheld augmented reality games," in *Proc. IEEE Int. Symp. Mixed Augmented Reality*, 2011, pp. 19–28.

[7] Y.-F. Tuan, *Space and Place: The Perspective of Experience*. Twin Cities, MN: Univ. Minnesota Press, 1977, p. 496.

[8] S. Harrison and P. Dourish, "Re-place-ing space?: The roles of place and space in collaborative systems," *Proc. Conf. Comput. Supported Cooperative Work*, 1996, DOI: 10.1145/240080.240193.

[9] G. D. Abowd, C. G. Atkeson, J. Hong, S. Long, R. Kooper, and M. Pinkerton, "Cyberguide?: A mobile context-aware tour guide," *Science*, vol. 3, pp. 421–433, 1997.

[10] S. Feiner, B. MacIntyre, T. Höllerer, and A. Webster, "A touring machine: Prototyping 3D mobile augmented reality systems for exploring the urban environment," *Personal Technol.*, vol. 1, no. 4, pp. 208–217, Dec. 1997.

[11] B. N. Schilit, N. Adams, R. Gold, M. M. Tso, and R. Want, "The PARCTAB mobile computing system," in *Proc. IEEE 4th Workshop Workstation Oper. Syst.*, 1993, pp. 34–39.

[12] U. Neumann, "Tracking with omni-directional vision for outdoor AR systems," in *Proc. Int. Symp. Mixed Augmented Reality*, 2002, pp. 47–56.

[13] G. Reitmayr and T. Drummond, "Going out: Robust model-based tracking for outdoor augmented reality," in *Proc. IEEE/ACM Int. Symp. Mixed Augmented Reality*, 2006, pp. 109–118.

[14] B. Macintyre, A. Hill, H. Rouzati, M. Gandy, and B. Davidson, "The Argon AR web browser and standards-based AR application environment," *Proc. Int. Symp. Mixed Augmented Reality*, 2011, pp. 65–74.

ABOUT THE AUTHORS

Evan Barba received the Sc.B. and A.B. degrees in neuroscience and modern culture and media studies, respectively, from Brown University, Providence, RI, in 1999 and the M.S. degree in computer science from New York University, New York City, in 2006. Currently, he is working towards the Ph.D. degree in the Human-Centered Computing program at the College of Computing, Georgia Institute of Technology, Atlanta.

His research examines the relationship between technology and human cognition, with an emphasis on the role of spatial scale in the design and analysis of mixed-reality experiences.

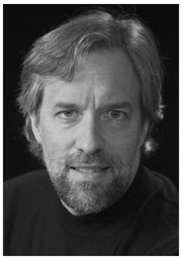

Blair MacIntyre received the B.Math and M.Math degrees in computer science from the University of Waterloo, Waterloo, ON, Canada, in 1989 and 1991, respectively, and the Ph.D. degree from Columbia University, New York City, NY, in 1998.

He is an Associate Professor in the College of Computing, Georgia Institute of Technology, Atlanta. He directs the Augmented Environments Lab, whose research focuses on the design and implementation of interactive mixed-reality (MR) and augmented-reality (AR) environments. His research is focused on enabling nontechnical designers to work with AR, and understanding the potential of AR as a new medium for games, entertainment, education, and work. He has been doing research in AR since 1991.

Dr. MacIntyre is actively involved with numerous conferences and workshops, and has been a Program Committee Chair for the 2000 International Symposium on Wearable Computing (ISWC), the 2003 International Symposium on Mixed and Augmented Reality (ISMAR), and The 2003 ACM Symposium on User-Interface Software and Technology (UIST). He has served on the editorial boards of *The International Journal of Human-Computer Studies* and the journal *Virtual Reality*. He also served as the guest editor of a "Mixed Reality" special issue of the IEEE COMPUTER GRAPHICS AND APPLICATIONS. He is the recipient of a National Science Foundation (NSF) CAREER award.

Elizabeth D. Mynatt received the B.S. degree in computer science from North Carolina State University, Raleigh, in 1988 and the M.S. and Ph.D. degrees from the Georgia Institute of Technology, Atlanta, in 1989 and 1995, respectively.

She is the Executive Director of the GT Institute for People and Technology, and Professor in the College of Computing, Georgia Institute of Technology, Atlanta. Her research program Everyday Computing examines the human–computer interface implications of having computation continuously present in many aspects of everyday life. Her research contributes to ongoing work in personal health informatics, computer-supported collaborative work, and human–computer interface design.

Dr. Mynatt is a member of the ACM SIGCHI Academy, a Sloan and Kavli Research Fellow, and serves on Microsoft Research's Technical Advisory Board. She is also a member of the Computing Community Consortium, a National Science Foundation (NSF)-sponsored effort to engage the computing research community in envisioning more audacious research challenges. Named Top Woman Innovator in Technology by *Atlanta Woman Magazine* in 2005, she has created new technologies that support the independence and quality of life of older adults "aging in place," that help people manage diabetes, and that increase creative collaboration in workplaces.