Technology | DOI:10.1145/1785414.1785422 Kirk L. Kroeker

Mainstreaming Augmented Reality

Advancements in computer vision, object recognition, and related technologies are leading to new levels of sophistication in augmented-reality applications and presenting new ways for humans to relate to the natural world.

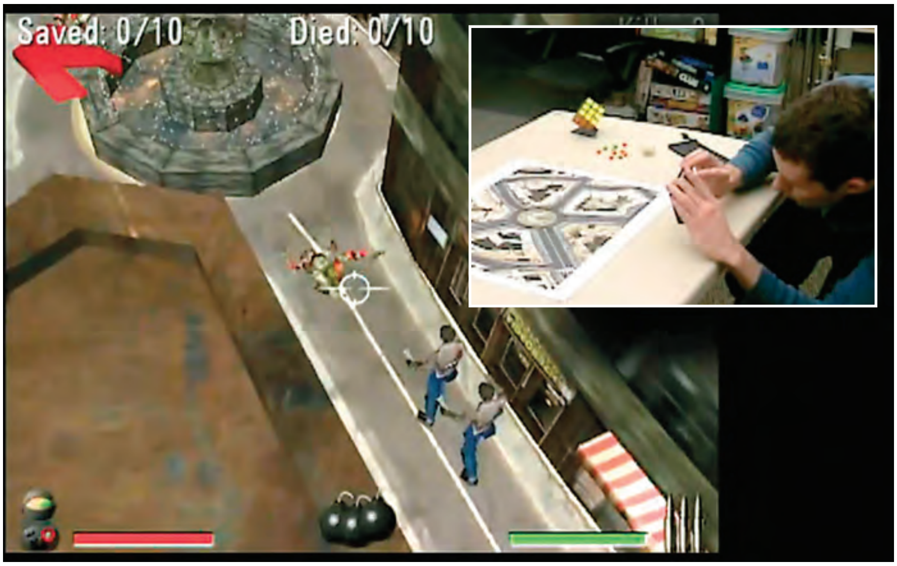

An augmented reality game called *ARhrrrr!* developed at Georgia Tech and the Savannah College of Art and Design. In the game, the graphics are tightly registered to a physical game board using an image-based feature tracker developed at Graz University.

PHOTOGRAPH COURTESY OF BLAIR MACINTYRE AND STEVEN FEINER

Since the emergence of the first augmented-reality applications 20 years ago, the field has drawn a great deal of interest and enthusiasm, not only from researchers working in computer science at the cutting edge of graphics technologies, but also from leaders in aerospace, medicine, the military, and a wide range of other industries and government sectors. In augmented reality (AR), a real-world setting or set of objects is augmented by a computer-generated overlay. Advancements in computer vision, object recognition, and related technologies are increasing the level of sophistication of that overlay, and presenting entirely new ways for humans to relate to the natural world.

While a great deal of research is being conducted in this area, given the promise of the technology to have a major impact in industrial and consumer applications, significant challenges remain, such as the accuracy of Global Positioning System- (GPS-) or compass-based AR applications, the bulkiness of head-mounted displays, and other issues endemic to the sciences and systems upon which AR technologies rely. Still, researchers developing AR systems continue to build increasingly compact and powerful applications, many of which require nothing more than a current-generation smartphone.

Examples of mobile AR applications include Layar, a "reality browser" that retrieves point-of-interest data on the basis of GPS, compass, and camera view, and GraffitiGeo, an application that lets users read and write virtual Twitter-style comments on the walls of restaurants, movie theaters, and cafes. Both applications are available for the iPhone platform. Another example is Goggles, a Google-created application that allows users to search the Web on Android phones simply by capturing photos of landmarks or other objects. The technology also allows users to point the phone's camera at local storefronts to retrieve business information automatically with GPS and compass data.

While the number of such mobile applications is increasing rapidly, AR evangelists say a killer app will be needed to make AR technologies truly catch on in the consumer space. Given the enormous popularity of Web-based social networking, for example, one killer app might come in the form of a mobile facial-recognition application that can automatically link the humans to their social-network profiles. One company, The Astonishing Tribe, has demonstrated an AR interface concept, called Recognizr, to show the possibilities of doing just that.

Another approach to mainstreaming AR is in gaming. One researcher working in this area is Blair MacIntyre, who directs the Augmented Environ-

Augmented-reality applications are increasingly compact and powerful, and many of them require nothing more than a current-generation smartphone.

ments Lab at the Georgia Institute of Technology. MacIntyre says his current work in AR is driven mainly by the desire to understand how to create compelling AR experiences, interfaces, and tools. To that end, he and his team build games and study them, focusing on everything from interactivity and visualization techniques to the feel of game mechanics to the social experiences they foster.

"I'm very driven to create tools and platforms that will give a broad range of people the ability to experiment with the technology," says MacIntyre. "Just as we didn't know what the Web would be used for until people with real problem- and design-driven goals started trying to create applications, the same will be true for AR."

The mainstreaming of augmented reality now largely depends on the ability to manufacture and sell the technology profitably, says Steven K. Feiner.

Tightly Registered AR Games

For now, MacIntyre is focusing on what he calls tightly registered AR games, in which the graphics appear to be locked onto the real world. In the *ARhrrrr!* game, for example, a handheld device's graphics are aligned with the physical game board using an image tracker to determine where the camera on the handheld is located, relative to the board. The system pulls video from the camera, runs it through a vision library, and returns an estimate about the game board's relative position. Using that information, the handheld draws graphics in the camera's view of the board. Those graphics remain locked in place over a wide range of movement by the player.

"We found that if the graphics are unambiguously aligned with features in the world, game players treat the combined physical-virtual view as one merged space," he says. "As a result, they can refer to virtual content smoothly and unambiguously, and can collaborate or compete as they would on a physical board game."

MacIntyre says the biggest challenge he faces is with the limitations of the vision-based tracking technology that signals to the phone what the camera's relation is to the world. "We are constantly struggling with the tension between what we want the games to do and what is technically possible to know about the world and to track and interact with," he says. Because accuracy is directly related to the quality of the inputs, MacIntyre and his team use vision-based tracking technology instead of less-accurate alternatives such as handheld-based GPS, compass, and accelerometer sensors, which might work for large-scale AR applications but lack the precision needed for tightly registered games.

Another researcher working in this area is Steven K. Feiner, director of the Computer Graphics and User Interfaces Laboratory at Columbia University. Feiner began his work in AR by exploring how the technology might be used to assist in maintenance and repair, and has directed projects ranging in focus from restaurant guides and gaming to integrating technical instructions directly into a task domain. "Our overarching goal is to design user interfaces that help people be better at whatever they do," Feiner says, noting that his general approach in these AR projects is to build dynamic virtual worlds that are overlaid on and geo-

Employment

U.S.'s Bright CS Job Forecast

The job outlook for U.S. college students majoring in computer science is very favorable, according to *The Market For Computing Careers*, a report by Joel Adams, a professor of computer science at Calvin College. Adam's report contains an analysis of data from the U.S. Bureau of Labor Statistics, Computing Research Association's Taulbee Survey, and *U.S. News & World Report*.

"The U.S. Bureau of Labor Statistics predicts that computing will be one of the fastest-growing U.S. job markets in science, technology, engineering, and mathematics (STEM) for the foreseeable future," according to the report, with "nearly three out of four new science or engineering jobs in the U.S. going to be in computing." Of these new computing jobs, 27% will be in software engineering, 21% in computing networking, and 10% in systems analysis.

The "demand for software engineers, network administrators, systems analysts, and other computing-related professionals is exploding, but fewer and fewer students are choosing to study what is needed to get these jobs," the report says. The U.S. Bureau of Labor Statistics, for example, predicts, nearly 140,000 new job openings in computing per year through 2018, with less than 50,000 CS graduates vying for those jobs.

Meanwhile, as fewer students enter CS, the salaries for software engineers, network administrators, and systems analysts "are climbing." According to *U.S. News & World Report*, the median salary for a software engineer ranged from $85,000-$92,000 in 2008, with the best-paid 10% of software engineers earning more than $136,000.

"I think the most surprising thing [in the report] is that the U.S. Bureau of Labor Statistics is projecting more than four times as many new jobs in computing than in all the traditional (non-software) engineering areas combined," Adams said in an email interview. "A second surprise was their projection of more than twice as many new computing jobs per year than there are computing graduates at present. The third surprise was that computing is the only STEM discipline where the demand for graduates outstrips the supply.

"Calvin College is in Michigan, which has, I believe, the highest unemployment of any state, but we are already seeing the effects of this imbalance," says Adams. "This past semester, we received an average of three requests per week from local businesses seeking students with significant computing skills. We don't have nearly enough students to meet such demand."

—Jack Rosenberger

An augmented-reality application developed at Columbia University. The "opportunistic controls" shown in this image are virtual buttons on a raised portion of an aircraft engine housing, providing haptic feedback for a maintenance technician.

PHOTOGRAPH COURTESY OF BLAIR MACINTYRE AND STEVEN FEINER

metrically registered with a user's perception of the physical world, providing information that would otherwise be invisible.

In one example of this approach, Feiner and his team explored how users could more effectively control those AR applications in which it's necessary to select and adjust certain physical operating parameters, all without using physical controls and without diverting attention from the task at hand. One of Feiner's graduate students, Steve Henderson, developed a solution to this problem. Called "opportunistic controls," the technology locates virtual controls, such as buttons, sliders, knobs, and so forth, on top of physical elements of the task environment. For example, the AR system might place a virtual button on a protruding bolt or a virtual rotary knob on a rotating cable connector.

To create an opportunistic control, the system overlays a physical object with a 3D widget and associates it with a hand gesture. The depiction of the widget is rendered in a head-worn display, while the hand gestures are recognized through computer-vision algorithms performed in real time on video captured from an overhead camera. A separate camera captures the video so the control can be operated even when the user is not looking directly at it. The shape of the physical objects associated with the controls can help the user distinguish them by touch, as with conventional controls.

Making a Profit

With these and other AR technologies growing increasingly robust and reliable, Feiner says the mainstreaming of AR technology now largely depends on the ability to manufacture and sell the technology profitably. He says he remains convinced that future AR technology will not be a mere novelty; instead, he says, it will be one of the fundamental user interface paradigms through which humans interact with the world. In the future envisioned by Feiner, AR technology will be housed not only in comfortable eyewear, but also in handheld or stationary see-through displays, in projected displays, and even some in surfaces that are themselves displays.

Georgia Institute of Technology's MacIntyre shares a similar view of the future in which humans are constantly immersed in a mixed physical-virtual world. The major challenge on the path to a future of ubiquitous AR technology is, of course, to develop the complete infrastructure, from the necessary technologies to track where users are and what they are looking at to the privacy and security infrastructure to ensure that users can trust the system, and also to ensure that user privacy and safety are not violated. "I think the technology has a long way to go before we can experience such constant immersion," says MacIntyre. "But we will begin getting a taste of it in the very near future."

With the goal of nudging the research community in that direction, one of MacIntyre's projects is a standards-based platform for mobile AR, the aim of which is to do for AR what the early decoupled client-server architecture did for the Web. In contrast to cloud computing, AR applications currently require dedicated programs running on client devices. MacIntyre's idea is to create a general-purpose AR browser and a corresponding collection of cloud-based technologies to allow anyone with a server to create and deploy mobile AR apps without requiring users to install anything.

"We need to start developing open standards for AR applications, so a wide variety of people, companies, and organizations can create and deploy these applications," MacIntyre says. "I believe these application environments and open standards will have the biggest impact on the blossoming of AR as a widely used technology." **C**

Further Reading

Barfield, W., and Caudell, T. (eds.)
Fundamentals of Wearable Computers and Augmented Reality, **Lawrence Erlbaum, Mahwah, NJ, 2001.**

Bimber, O., and Raskar, R.
Spatial Augmented Reality: Merging Real and Virtual Worlds, **A.K. Peters, Natick, MA, 2005.**

Feiner, S.K.
The importance of being mobile: some social consequences of wearable augmented reality systems, *Proceedings of the 2nd IEEE and ACM International Workshop on Augmented Reality*, 1999.

Hainich, R.R.
Augmented Reality and Beyond, **Booksurge, Charleston, SC, 2009.**

Haller, M., Billinghurst, M., and Thomas, B.
Emerging Technologies of Augmented Reality: Interfaces and Design, **Idea Group Publishing, Hershey, PA, 2006.**

Kirk L. Kroeker works in communications and has written extensively about the impact of emerging technologies.