Andrew Glassner's Notebook

http://www.glassner.com

Everyday Computer Graphics

Andrew Glassner

Today we use augmented reality (AR) to analyze complex molecular geometry, perform surgery, and figure out the weather. Some researchers are also experimenting with AR in everyday settings. Current AR technology is frequently pretty awkward, heavy, and obtrusive. Soon the day will come where we can all pop on a pair of glasses embedded with a computer display, and they'll be in wireless contact with a fancy portable computer in our pocket or backpack, or even sewn into our belt. The glasses will also have cameras on them, so the computer can see what we're seeing.

Our pocket computers will know everything we're seeing and hearing, they'll know where we are thanks to built-in GPS receivers, and they'll know all about our personal history and preferences (because, after all, these are truly personal computers). Hopefully they'll protect this private information from others.

If AR becomes universal, then we can use it for casual, small things that will make our daily lives a bit easier. Of course, I'm hardly the first person to think about using AR for everyday tasks. You can catch up with the current state of the art by checking out the sidebar "References for Augmented Reality," which represents only the tip of the iceberg. Entire conferences are devoted to wearable computers, which frequently talk about some of these issues.

One of the more logical uses for AR is communication. For example, as many authors have observed, it would be handy to leave notes for each other in cyberspace, yet attach them to physical objects or locations.

Here I'll suggest some of the applications that I'd like to have available for personal or individual use on an everyday basis. They communicate to users without text. Instead, graphics do all the work.

On the road

Here's my first potential scenario for everyday AR use. There's nothing quite like the joy of driving around a parking lot looking for a space. You crawl along, looking for spots, avoiding the occasional pedestrian, and then always seeing a space open up where you were a moment ago just in time for someone else to grab it. When you're driving around inside a garage, you get the added pleasure of inhaling the exhaust fumes of the people in front of and behind you, who are also trolling for spaces.

Figure 1a shows a parking lot I needed to park in a few years ago. In this instance, a picture of the lot itself (see Figure 1b) would be helpful so that we could at least see where we're going. But even better, we could indicate which spaces are open, as in Figure 1c. Just go to one of the red spots and park.

How will the signboard know which spaces are free? If the day is clear, it can use satellite imagery. Alternatively, little magnetic sensors under each space could report back to the signboard whether there's a car over them. Such sensors already exist, and are used today for everything from counting how many cars get onto a ferry to enforcing the speed limit on some highways.

If the weather is awful, it might be hard to make out the signboard. Then we could put the signboard into the car. Into the dashboard of Figure 2a let's install a little

References for Augmented Reality

The following are a few key references to bring you up to date on the latest in AR technology:

- H. Beadle et al., "Location-Aware Mobile Computing," *Proc. IEEE/IEE Int'l Conf. Telecomm* (ICT 97), IEEE Press, 1997, pp. 1319-1324.
- W. Broll et al., "Interface with Angels: The Future of VR and AR Interfaces," *IEEE Computer Graphics and Applications*, vol. 21, no. 6, Nov./Dec. 2001, pp. 14-17.
- T. Höllerer et al., "Exploring MARS: Developing Indoor and Outdoor User Interfaces to a Mobile Augmented Reality System," *Computers and Graphics*, vol. 23, no. 6, Dec. 1999, pp. 779-785.
- T. Höllerer, S. Feiner, and J. Pavlik, "Situated Documentaries: Embedding Multimedia Presentations in the Real World," *Proc. 3rd Int'l Symp. Wearable Computers* (ISWC 99), IEEE CS Press, 1999, pp. 79-86.
- S. Mann, "Wearable Computing: A First Step Toward Personal Imaging," *Computer*, vol. 30, no. 2, Feb. 1997, pp. 25-32.
- J. Spohrer, "WorldBoard—What Comes After the WWW?" 1997, http://www.worldboard.org/pub/spohrer/wbconcept/default.html.
- B. Thomas et al., "A Wearable Computer System with Augmented Reality to Support Terrestrial Navigation," *Proc. 2nd Int'l Symp. Wearable Computers* (ISWC 98), IEEE CS Press, 1998, pp. 168-171.

 Published by the IEEE Computer Society

1 Finding a place to park. (a) Where can I park in this slushy, filled lot? (b) This map is helpful. (c) Even better, an open space is marked in red.

2 (a) If bad weather obscures the sign of Figure 1, parking is still a hassle. (b) Perhaps you can open a dynamic map inside.

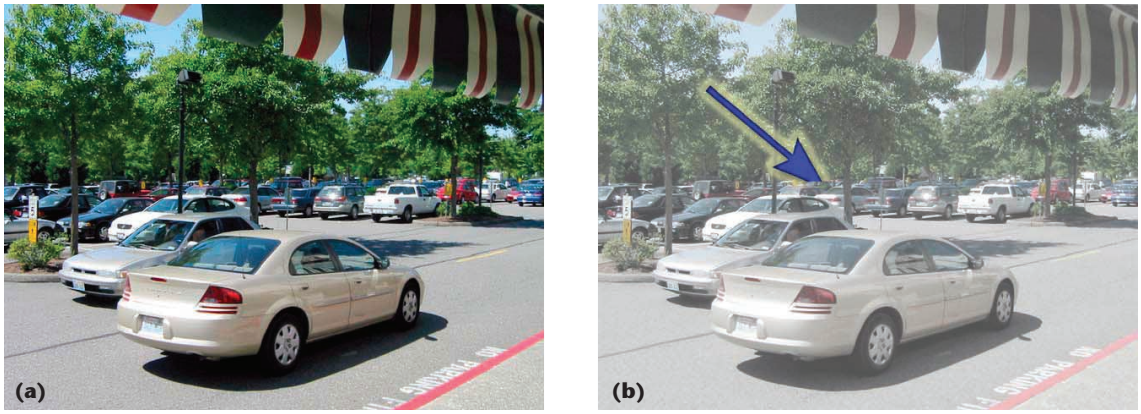

3 (a) Where did I park? (b) Oh, right over there.

pullout drawer, as in Figure 2b. It's got a small display that can receive a signal from the signboard. Of course, if the car has any other internal displays, like an onboard map readout, we could just use that display temporarily for parking information.

Getting home

The remote-control door opener is the modern dowsing rod. Once you've parked your car somewhere and spent the day, you have to remember where you parked to get home again. I spend too much time hunting for my car. Other people have the same problem, and I'm increasingly hearing cars honk and beep as their owners walk around a garage, pushing their remotes' lock and unlock buttons in hopes of hearing a familiar sound they can follow to its source. When you get a few people walking around a concrete parking garage with cars honking and blinking their lights calling to their owners, everyone's individual search gets more difficult, and the whole scene becomes surreal and somewhat deafening.

However, if the computer in your pocket can determine when you parked, then it knows where you were when you got out. Then anytime you return to that endless parking lot (as in Figure 3a) you could just ask the system where your car was when you left it, and get an immediate and precise answer (as in Figure 3b).

Lost and found

As Lady Bracknell said in Oscar Wilde's *The Importance of Being Earnest*, "To lose one parent, Mr.

4 (a) Where was I sitting? (b) Oh, right over there.

5 (a) I'm in time for the big 4th of July fireworks, but my friends got there before me. Where are they? (b) Oh, right over there.

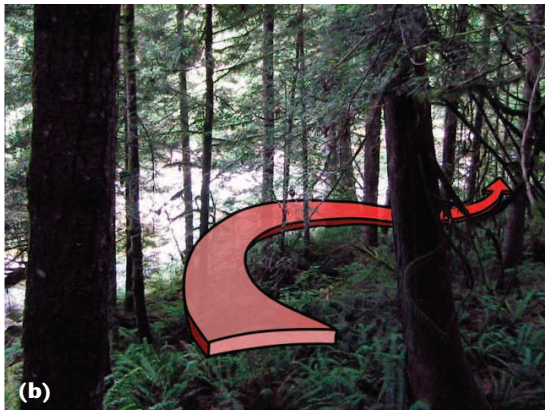

6 (a) Where is my campsite? (b) Ah, it's that way.

Worthing, may be regarded as a misfortune; to lose both looks like carelessness." How much worse than losing just your car, then, is losing your friends.

At any big outdoor concert, like the one at Seattle's Woodland Park Zoo shown in Figure 4a, there are lots of people. If you need to leave your friends for some reason, getting back to them can be difficult. If people are sitting close together, walking around and looking for your friends becomes a tough proposition. However, because your pocket computer has a built-in GPS receiver, you can ask it where you were a few minutes ago, and it can point the way immediately, as in Figure 4b.

Sometimes you show up a little late to an event, and your friends are already there. Even if you climb up to the top of a hill, as in Figure 5a, it's tough to find your friends. Even binoculars might not help. But if your pocket computer could talk to theirs, the two machines can compare locations, and you can go right to your friends (see Figure 5b).

There are times when it's not enough to just know where to go—you also need help figuring out the best route to get there. Road atlas programs find the fastest or most convenient routes using highway and road systems. However, what if you're in a more natural environment, like Figure 6a, and you want to find your campsite? There might be significant topographical information that's relevant to you (for example, you might not want to climb up too steep a hill, or have to jump across a deep ravine). Your pocket computer could combine your location, your goal, and its knowledge of

7 (a) Where's the room for my session? (b) Ah, it's upstairs and to the right.


(a)


(b)


(c)


(d)

8 (a) The endless snack aisle. (b) Which ones have a lot of salt? The less obscured, the saltier. (c) Which ones have a lot of fat? (d) Show me the calories. (Note that these examples don't represent a real nutritional analysis.)

the intervening topography to suggest a course (see Figure 6b).

This is also useful in man-made environments. Figure 7a shows a photo of the Los Angeles Convention Center during Siggraph a few years ago. I arrived early in the morning to prepare for a course, and needed to find the right room. Although Siggraph has unusually good signage for a conference, you can still get lost or confused in the caverns of the convention center. Following the Pac-Man dots in Figure 7b, though, it's easy to navigate the maze from here to there.

Shopping

Snack-food manufacturers just love it when we go to the grocery store hungry. We're tempted to grab some of their richly flavored and textured wares, temporarily unconcerned with the nutritional choices they represent.

AR can help us out when we wander down a snack-food aisle (see Figure 8a). If we're in the mood for something really salty, we could ask the system to rank the foods by how much salt they contain. In Figure 8b the snacks with a lot of salt are more visible than those with less salt. If we're feeling a little health conscious, we could also ask the system to show us how much fat is in one bag relative to another, as in Figure 8c. Finally, we might ask how many total calories we'll consume by eating one bag relative to another. I don't know why stores stock shelves the way they do, but Figure 8d shows us the caloric value of some of the elements at the end of the aisle.

Readers of this column know that I like to make greeting cards using techniques from origami to pop-up devices. But sometimes a store-bought card is a good plan. One problem with buying a card in the store is that there are often hundreds of cards, and you can only see a little strip along the top of each one. Faced with a display like Figure 9a (next page), how can you find the cards you like?

Many card manufacturers make their catalogs available online. Our pocket computer could scan the display, match the visible bit of each card with the catalogs, and guess at what the cards are. Then we could ask for, say, a picture of a moose, or maybe a card without a prewritten message inside. The system could match those criteria against the catalogs and the visible cards, and point us to the ones we like (see Figure 9b).

Buying fruit is a common activity: if we're going to eat it right away, we want something ripe, but if it's going to sit around for a while, we might want something that has a few days to go. It's easy to tell how ripe a banana is just

9 (a) Do they have any cards with a photo on the front and no message inside? (b) Ah, these two fit my criteria.

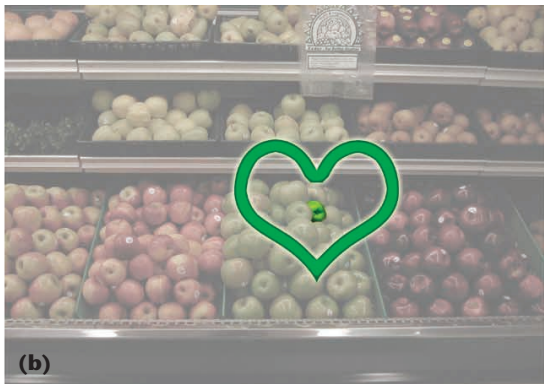

10 (a) Are there any Granny Smith apples of the size and ripeness I like? (b) Ah, right over here.

11 (a) Where does that Dr. John CD go back on the shelf? (b) Ah, right there.

by looking at it, but other fruits are best estimated by feeling them. Some fruits go through subtle visual changes as they ripen. These changes might be too small for us to spot, but a good AR sensor could pick up on them. If we're looking for a Granny Smith apple that's just right for eating for our tastes from the display of Figure 10a, a pocket computer could spot the best one, as in Figure 10b.

Home improvement

Keeping track of lots of little things can be hard, even if you're careful. Suppose you're visiting a friend, and you take down a CD from his collection in Figure 11a. You're in a mood for New Orleans jazz, so you choose a fine CD called *Dr. John Plays Mac Rebennack*. When you go to put it back, you realize the discs are shelved in alphabetical order. Where does this go—under "D" for Doctor, or "J" for John? Or maybe it should go under "R" for Rebennack, because after all Dr. John and Mac Rebennack are two names used by the same guy. Your pocket computer could easily remind you just where you took the CD from, as in Figure 11b, and help you put it right back in its place.

The same problem goes for books. Let's suppose this time that you're searching for a particular book, but because your shelves are in disarray and look something like Figure 12a, you're not sure where to find it. You could simply tell the system what you want (perhaps as easily as by saying the book's name), and the book you want gets highlighted (see Figure 12b).

How many times have you opened the refrigerator looking for something to eat? Each time you open the door, you let the cold air out and that burns some energy. To the normal eye, most home refrigerators look like a bulletin board, like the one in of Figure 13a. If you had x-ray vision (courtesy of your pocket computer and AR glasses) you could see just what was inside, as in Figure 13b. If you were planning a trip to the store, you could even ask the system to show you just how much is left inside various containers, as in Figure 13c.

Moving a household is no fun. It's inevitable that you'll lose track of something in one of the dozens of boxes that seem to accompany even the smallest or most local of

12 (a) Where is that biography of Thomas Jefferson? (b) Ah, right there.

13 (a) I'd like to see what's in the fridge without opening it up and wasting energy. (b) Easily done. But how much milk is left? (c) Here's how much milk, water, and lime juice we have left.

14 (a) Where did I pack those items? (b) Ah, in those boxes.

moves. Suppose you've moved and you're now unpacking, but you can't find your plates and silverware, or some other specific items you want. Figure 14a shows a typical pile of moving boxes. Even those with printed information on the box rarely contain what they originally did. However, because your pocket computer and glasses were active and tracking objects during the whole packing process, you could simply ask the system where some of your items are located, and as in Figure 14b it could immediately show you where to find them.

Telephones are a big part of everyday life. Many people now have caller ID to help them separate wanted calls from unwanted ones. A caller ID box usually has a little text display on the front. When the phone rings the box shows the originating call's phone number; sometimes it also names the person or company associated with the phone number.

It would be even better to literally see the person who's calling. If the phone of Figure 15a (next page) rings, we could look over and see who's calling (see Figure 15b and 15c). This strikes me as a lot more natural than trying to remember a number or read a name.

Let's take a slight detour from using just AR glasses. Suppose that we can put displays in the environment.

15 (a) The phone's ringing—who is it? (b) Oh, the doctor. (c) It's the vet.

16 (a) It would be nice to have another view. (b) Over a lush forest. (c) Underwater.

17 (a) I wonder what the emotional weather is like inside? (b) I probably should avoid the kitchen right now.

Many companies are pursuing technologies that will ultimately give us flexible displays. Think of these as wallpaper displays: you have a flexible sheet of stuff, you put it up on the wall, and voila, it's a computer display!

Home offices are a great place to get work done, and you might even have a nice view, like the one in Figure 16a. If we put up some display wallpaper on the windows, then we could change the view to anything we fancy—from a still image or a movie to a synthetic animation or live video coming from a remote camera or Web cam. Figures 16b and 16c show a couple of possibilities.

Avoiding pressure

Many of us check the weather report before we leave the house in the morning, but perhaps we could get some value from a report before we come back again at the end of the day.

Returning to the house in Figure 17a, you might think that everything's just fine and you're in for a warm welcome and relaxing evening. But of course there's no way of knowing what sort of day the other people inside have had, and they might be feeling a lot of pressure right now. Your dog might have eaten your spouse's favorite book, and he or she might be wearing a pretty frosty expression right now.

Enter the pocket computer emotional weather map. Figure 17b shows us where the high- and low-pressure zones are in the house as we approach. Perhaps on the basis of this information we'll decide that it's a good idea to avoid the kitchen for a little while, and just go hide out in the living room until the pressure front moves through.

Wrapping up

With a pocket computer and a nice head-mounted display, lots of everyday tasks can become easier. None of these are nearly as important as helping surgeons perform operations, or analyzing the safety of a nuclear power plant, but if the technology is around we may find that its everyday, casual uses are just as desirable.

Some of the suggestions in this article were originally presented at talks I gave at IEEE Visualization 2002 and ACM Siggraph 2002. ■

Readers may contact Andrew Glassner at andrew @glassner.com.