

Situated Computing: Bridging the Gap between Intention and Action

Anatole V. Gershman, Joseph F. McCarthy, Andrew E. Fano

Center for Strategic Technology Research (CSTaR)

Accenture

3773 Willow Road

Northbrook, IL 60062 USA

{*anatole.v.gershman, joseph.f.mccarthy, andrew.e.fano*}@accenture.com

Abstract

Situated computing represents a new class of computing applications that bridges the gap between people's intentions and the actions they can take to achieve those intentions. These applications are contextually embedded in real-world situations, and are enabled by the proliferation of new kinds of computing devices, expanding communication capabilities and new kinds of digital content. Three types of discontinuities give rise to intention/action gaps and provide opportunities for situated computing applications: physical discontinuities, information discontinuities and awareness discontinuities. Several examples of applications that overcome these discontinuities are presented.

1. Introduction

Most people are aware of the increasing pace, and impact, of technological innovations. We believe that three converging trends – the three C's, if you will – are fueling these innovations: (1) *Computing* and sensory devices are becoming cheaper and smaller. (2) *Connectivity* is becoming more widespread, less expensive and multi-modal: from broadband to wireless. (3) *Digital content* and services are becoming more ubiquitous and abundant. Taken together, these trends open the possibility for very different applications of computing – applications embedded into our physical environment and the everyday things we use. These *situated computing* applications will know who we are, where we are, what we are doing, what we want, and how we can take advantage of the resources available in our physical environment. This knowledge will make the new applications vastly more effective in helping us with our tasks both at home and at work.

Other researchers have talked about situated computing [Hull, *et al.*, 1997], but their primary focus was on an architecture to support situated computing applications. Our aim here is to focus primarily on the applications of situated computing, and to provide a

conceptual framework that describes how and why these applications are potentially so useful.

We argue that one of the primary impacts of these situated computing applications will be to bridge the gap between our *intentions* – the needs and wants we seek to satisfy – and the *actions* we can take to achieve them. We propose three types of discontinuity that give rise to this intention/action gap – physical, informational and awareness – and describe several prototypes we have built to overcome these types of discontinuity at Accenture's Center for Strategic Technology Research (CSTaR).

2. The gap between intention and action

At any given time, each of us has a multitude of intentions, both conscious and unconscious, to which we assign varying levels of both importance and relevance to our current context. There often exists a gap between any given intention and our ability to take action to achieve it. These gaps arise for a number of reasons: we may not be in a place where we can take relevant (and effective) action, we may not have the right information with which to direct our action, and/or we may not be aware that an opportunity to take action lies just beyond our immediate focus of attention. Each of these factors is based on a type of discontinuity: physical, informational and awareness. We illustrate each of these types of discontinuity in a separate scenario below.

We are all familiar with some version of the following situation: we are in a store evaluating how well some product would “fit” with our existing possessions. For example, we may be in a furniture store admiring a couch; we think it will match our living room nicely, but we are not sure. Our living room is in one physical location while the couch is in another. We can't be sure about the match unless we see the couch in the living room. This physical discontinuity creates a gap between our intention to improve our living room and our action to achieve it - to buy a couch in this example. If we could overcome this discontinuity, if we could bring the necessary aspects of these two locations together, we would create what

business people call a *moment of value*, a window of opportunity within which we can act to satisfy our intentions. This example is not a peculiar case. These physical discontinuities occur whenever we need to decide whether to introduce or remove an object in a physical setting, or alter a physical setting in some way. Such decisions will be based on our capacity to effectively anticipate the resulting state and evaluate its appropriateness. Applications designed to overcome physical discontinuities constitute the first type of situated computing applications that we will discuss.

Another familiar station is seeing (or hearing) a product – an article of clothing, a book or CD – and wanting to buy that product ... for the best price, of course. We might be at a friend's house listening to a CD or hearing about their favorite book, and want to get a copy for ourselves, but have to wait until the next time we're at a book or CD store to act on that intention. Or we might even be in a physical book or CD store and wonder if a better price might be found at one of their on-line competitors. In both cases, we know what we want, but we don't have ready access to the information that would help us satisfy our intention. This lack of information where we need it most again creates a gap between our intention and action. In many instances we will not be able to assess the quality of the opportunity without this information. Applications designed to overcome such information discontinuities constitute the second type of situated computing applications.

Finally, the third type of discontinuity that gives rise to an intention/action gap is based on a lack of awareness. For example, we may be walking through a mall, headed for a store to buy a sweater. While doing so we walk right past a store that sells those water filter replacements we've been meaning to get. In this instance we are missing an opportunity to address a longstanding, though not terribly urgent, goal simply because we are not aware the opportunity exists. A key aspect of our planning abilities as humans is to seek and seize opportunities to achieve or further our goals. However due to our various limitations and the constraints imposed by a situation, we will often simply fail to notice and become aware of opportunities that are readily available. Therefore the third class of applications we introduce are those that help address these awareness discontinuities by identifying and highlighting opportunities to address various goals.

3. Overcoming physical discontinuities

In our work we have explored two distinct and complementary approaches to the problem of physical discontinuities. The first approach is embodied in the MAGIC HOME application, which illustrates the furniture shopping example introduced earlier. MAGIC HOME was designed to bring both physical locations – the store and

the home – into a common virtual space. In MAGIC HOME, the layout, furniture and other objects in a particular living room are all represented digitally and stored on a smart card, which a furniture shopper brings with him or her to the furniture store. At the store, a wireless bar code scanner is used to select a couch the shopper is interested in. The store has a digital representation of all its merchandise, each identified by bar code, and can merge the digital representation of the shopper's living room with the digital representation of the selected couch. A flat-panel display can be used to depict how the couch would look in the living room, allowing the shopper to rearrange the virtual furniture and even check to see whether the selected couch would fit through the doorways. Figure 1 shows a screen shot from the MAGIC COUCH application.



Figure 1. Magic Home

The capability to merge the content of two physical locations in one virtual location – a virtual house, in this case – has a great deal of potential for creating moments of value. Unfortunately, there are some obstacles to its widespread deployment in today's world. One obstacle is the availability of digital representations of our houses and goods. However, most new houses and manufactured goods are designed on computers, so their 3D renderings are already created during the natural course of the design process, and older homes and artifacts can rely on a rendering technique based on photographing an object from several angles. Thus, in the future, we anticipate greater availability of these models. Even if we can overcome the first obstacle, we are still faced with the problem of updating the position of every piece of furniture every time it is moved. One solution to this problem would be to insert small positioning and orientation sensors into every piece of furniture.

We are exploring solutions to the problem of creating and maintaining digital representations of one's possessions in another project, MAGIC WARDROBE,

currently under development at CSTaR. In this application, articles of clothing have embedded smart tags, e.g. in the button of a jacket. When you purchase a new jacket and put it into your wardrobe, the wardrobe immediately recognizes that event and asks if you want to register your purchase with the manufacturer. In exchange, the manufacturer will provide you with a 3D rendering of your new jacket. Now the manufacturer has established a direct relationship with you and you have an updated virtual wardrobe. Next time you go to a different store and want to buy a tie, you can call up your virtual wardrobe to the store's screen and see if the tie goes together with your jacket. Currently, such solutions are prohibitively expensive; however, these sensors are becoming increasingly affordable. Such technological solutions to the physical discontinuity problem will provide a key differentiator for those retailers who are willing to invest in new technologies.

We believe the opportunities afforded by a virtual house justify what are, in the long term, relatively inexpensive investments and surmountable obstacles. Such a virtual house has the advantage of being accessible from anywhere it is needed. Moreover, note that the virtual model does not have to be complete or perfectly consistent to be useful. We don't need to represent the location and orientation of every ash tray in the house to see if a couch would match nicely. And we would probably rather access a representation of our house after it has been cleaned and ordered, rather than a faithful representation of its more likely sorry current state.

MAGIC HOME and MAGIC WARDROBE represent one approach to solving the problem of physical continuity, by merging virtual models of objects from two locations in the physical world. Another approach to this problem is illustrated by one of our augmented reality applications [Dempski, 1999]. IN-HOME SHOWROOM superimposes a virtual model of a selected product over the shopper's view of their physical home using a head-mounted display (HMD). IN-HOME SHOWROOM projects a 3D rendering of the couch I saw in the store or in an on-line catalogue. I can fix the imaginary position of that couch in my room and then walk around it. Using position and orientation sensors built into my HMD, my computer re-renders the couch to provide the correct perspective as I move around the room.

An obvious advantage of this augmented reality approach is that it eliminates the need to render the environment into which the new object is placed. On the other hand, it is hard to "subtract" anything from reality. For example, it is difficult to virtually remove the old couch from my living room before "augmenting" it with the new couch.

We believe that these two approaches for overcoming physical discontinuity have complementary strengths and challenges. The "virtual house" approach is predicated on

our ability to automatically synchronize the states of the physical and the virtual worlds. The "augmented reality" approach is predicated on our ability to track our own position and orientation in the physical world.

4. Overcoming information discontinuity

Physical discontinuity involves the problem of merging representations of two or more objects or locations in the physical world. However, another type of discontinuity – information discontinuity – is based on the inaccessibility of information in the digital world in physical contexts where that information would be useful.

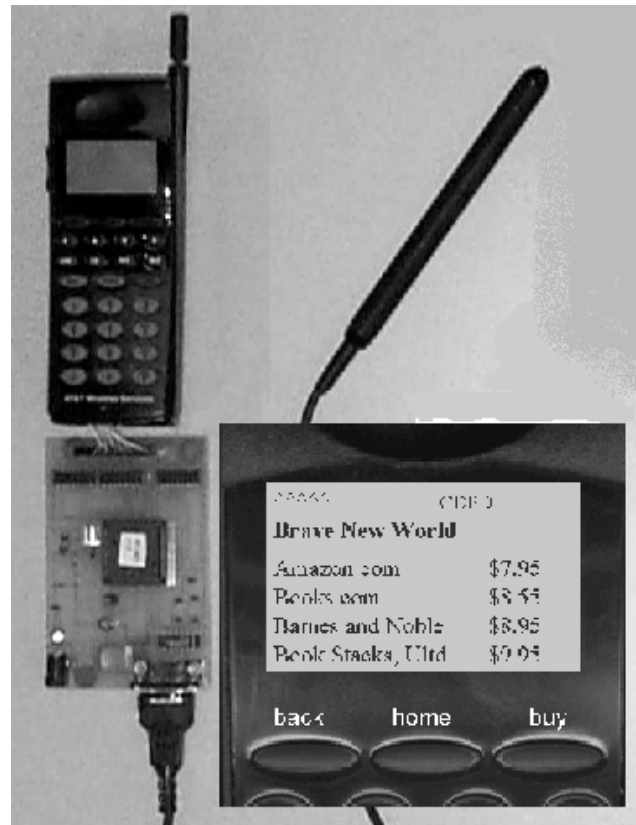


Figure 2. Pocket Bargain Finder

To help the book buyer introduced in Section 2, we developed a device called POCKET BARGAINFINDER [Brody & Gottsman, 1999], which is essentially a portable instantiation of the web-based comparison shopper we pioneered with BARGAINFINDER [Krulwich, 1996], which ran on a desktop machine. POCKET BARGAINFINDER, depicted in Figure 2, is a digital cellular phone with a barcode scanner. You scan the barcode on the back of the book, the device transmits it to a server which checks the prices in several on-line stores and sends them back to you to be displayed on your phone's screen. Technically,

it is a fairly simple system, but it sent shock waves through the merchant community [ABC]. It combines the simplest form of physical object tagging (the barcode on the back of the book), a simple sensor (the barcode scanner), the most common wearable computing device (the cellular phone), and the World Wide Web, to connect the physical object and the online information about it. You can use POCKET BARGAINFINDER to buy a book from an on-line bookstore, effectively making the device into a situated cash register for the on-line store. This capability brought together the intention to buy the book at a good price and the action of buying, creating a moment of value.

We can easily imagine many extensions of the same basic idea. When faced with information discontinuity about physical objects, place some remotely readable tags on these objects. Have a wearable or portable device capable of recognizing the tags. Have information about the object relevant to the current context sent to the device and provide the capability to act on this information through the device. Here are some other examples of this idea that we are exploring.

Imagine that you are wearing a jacket with a button containing a "smart tag" discussed in Section 3. If I like your jacket, I can pull out my "smart phone", read the tag and ask my clothier to find me something similar. In this way, the whole world can become a giant showroom.

Or imagine a worker looking at a compressor in the guts of an oil refinery. The sensor in his hardhat recognizes what he is looking at and puts the maintenance records of the compressor and the correct maintenance procedures on a heads-up display.

The second type of a situation where we experience an information discontinuity is familiar to anyone who travels. You are on your way to an out-of-town meeting. You get off the airplane and immediately call your office. Is the meeting still on? Is it at the same time and same place as originally scheduled? Is there any news relevant to the meeting topic? What is the name of the limousine company that you usually use in this city? In other words, you are engaged in a specific task (preparing for the meeting) that requires a variety of types of information. You are mobile but the information you need is in your office, in the head of a colleague, or in some other location different from where you are. A common solution to this problem today is to make a lot of phone calls and check your e-mail all the time, but this is far from optimal.

Our approach to this problem is based on the way that some well-organized high-level executives maintain awareness of relevant information. The process involves two assistants. One assistant is in the home office collecting and filtering the information the executive might need in his or her current tasks. The potentially relevant information is communicated to the second

assistant who is observing the local context and deciding when it is appropriate to communicate this information to the executive. This is a highly effective system, but prohibitively expensive for many executives.

Our goal was to provide a similar functionality automatically by using a combination of an intelligent agent running on the home server (MUNIN) and a personal digital assistant with a wireless communication capability (the AWARENESS MACHINE). The AWARENESS MACHINE is implemented on a Windows CE handheld computer with a built-in wireless modem. It is responsible for communicating the location of the user and other local changes in the user's schedule back to MUNIN. On the home server, MUNIN continuously collects several kinds of information and messages relevant to all the user's interests and tasks, but sends to the AWARENESS MACHINE only the information and messages that are deemed relevant to the current context of the user. It figures out the context based on the user's schedule, time of the day and the location of the user. You could pull the machine out of your pocket, turn it on and immediately see the most important things you needed to know at that moment. Unfortunately, the current generation of hardware makes continuous use of the machine impractical. The communication link turned out to be slow, unreliable and drained the battery too quickly. Ideally, the information should have been trickling in while the machine was not in use, but the current generation of pocket-size machines is not designed to do this. However, in limited contexts, the AWARENESS MACHINE worked well, providing the right information at the right time and in the right place, bridging the gap between our intentions and the actions we can take.

The AWARENESS MACHINE is an application developed as part of our Active Knowledge Management project, in which we are exploring various mechanisms for actively providing the right information at the right time and place. We have experimented with information delivery vehicles other than a pocket-size wireless PDA. One such alternate delivery vehicle, MAGIC WALL, is a large plasma display installed in a high-traffic public place in our Palo Alto office. Your presence near the display is detected by an RF sensor that can read your ID badge within several feet of the display. MUNIN recognizes your presence near a public display and selects the messages appropriate to that context. For example, in the morning when you pass the display on the way to your office, it reminds you of your next appointment. In the evening, when you are on your way out, it displays the traffic report along your usual route home.

Another application that seeks to overcome the information discontinuity problem is the MAGIC MEDICINE CABINET [Wan, 1999], a situated portal to online health information embedded in a medicine cabinet. One door of the cabinet contains a traditional

mirror, the other door houses a flat screen that displays a personalized selection of health care information. The MAGIC MEDICINE CABINET also has a camera (for face recognition capabilities), a microphone and speakers (for speech-based interaction), a computer with an internet connection, and a variety of devices related to health maintenance. This on-line medicine cabinet also knows something about your health. For example, if you have allergies, it monitors on-line resources that measure the local pollen count and warns you to take your medicine. If you suffer from hypertension, the cabinet reminds you to take a blood pressure reading and automatically communicates the results to the doctor's office. Through the use of sensors and smart labels on pill containers, the cabinet also knows what medicines are inside and which one has been taken out, warning you if you pick up the wrong one. It can establish real-time connections to your doctor, pharmacist and other providers of health care related services when the situation calls for it.

A working prototype of the MAGIC MEDICINE CABINET is depicted in Figure 3. As with many of our applications, the technologies employed in the prototype are still a bit too expensive for widespread deployment, but the price-performance curves in the computing and communications industries will bring the costs within the mass-market range in the next few years.



Figure 3. Magic Medicine Cabinet

5. Overcoming awareness discontinuity

The first two types of discontinuity discussed so far assume that we have already identified the specific focus of our intentions, e.g., we know that we want to buy a couch or a book, we need to get to a specific meeting. A third type of discontinuity arises in connection with opportunistic behavior. At any given time, we have many different goals or intentions. While a small number of them may be in the foreground at a particular time – getting home from work, buying a jacket, writing a business proposal, etc., the other goals or intentions persist in the background. Our immediate goal may be getting home from work, but if we see an interesting shoe sale advertised at a store along the way, we may stop and explore this opportunity. We continuously scan our environment, consciously or unconsciously looking for opportunities to satisfy one or more of our intentions. However, if we are not aware of an opportunity, we cannot take advantage of it. In our laboratory, we are exploring several ways in which technology can greatly expand our awareness of the opportunities available in our environments.

SHOPPER'S EYE is an application that addresses the problem of shoppers' lack of awareness of buying opportunities [Fano, 1998]. The system maintains the user's shopping list and preferences on a wireless handheld computer with a Global Positioning System (GPS) receiver. As the user travels around a shopping mall, the system makes the items on the list available to the local stores, which in turn make bids for the user's business (the latter part was simulated). In essence, SHOPPER'S EYE sends a message to the merchants saying: "There is a customer 75 feet away from your store who wants to buy what you sell. Do you have anything to say to him?" A similar system could be installed in a car, greatly expanding the driver's awareness of the surrounding businesses, such as potentially interesting shoe sales.

While SHOPPER'S EYE focuses on the consumer at the shopping mall, the AQUARIUM application [Bryan & Gershman, 1999] focuses on the consumer at home. The AQUARIUM, depicted in Figure 4, uses a large, flat-panel display to present pictures of your favorite products, animals, places, etc., appearing in the back, floating up to the front and then disappearing. The mix of objects on the screen is constantly changing, governed by the system's perception of your interests. At any time, you can touch an object on the screen and say "more of this," or "less of that." The system will adjust the mix of products based on your preferences. In this way, the AQUARIUM helps consumers "discover" more of their needs and intentions. If you do nothing, the system will gradually sift through its entire database while still showing you mostly the kinds of things you like. If you want to buy something you see, tap on it and say "buy". The AQUARIUM creates

an awareness of buying opportunities for products of interest to the user, bringing their intentions to the surface and enabling them to act immediately.

Obviously, the awareness discontinuity is not limited to shopping. Imagine a technician walking through a plant on the way from lunch. The equipment that needs a quick adjustment may be a few feet away, but the technician would not know it. This would be different if the equipment could sense the technician's presence and opportunistically request service while he or she is nearby.




Figure 4. The Aquarium

So far, we discussed the awareness discontinuity related to specific products or. Another important class of awareness discontinuity arises in a corporate environment with respect to knowledge management. In a large global organization it is very difficult to know if someone else has been working on topics related to yours. For example, if you are writing a proposal for an electronic commerce system and mention computer security considerations, your proposal might be greatly enhanced if you could cite the work done at one of your company's other clients or get advice from your company's security guru. For this to happen, you have to be aware that computer security is an important topic and that your company has the relevant experience. In a large organization, it is very difficult to maintain awareness of all the topics for which the company has one or more experts, or all the topics that may be relevant to a proposal. The INFOSCOUT project [Prasad & Anagnost, 1998] addresses this problem.

INFOSCOUT is a proactive recommender agent designed to watch a user performing a particular task and suggest other people and resources that could potentially be of help. For example, as you are writing your proposal and type "computer security," InfoScout automatically identifies this phrase as an important concept, and collects pointers to relevant documents and people. INFOSCOUT

can also be used to read your incoming correspondence, suggesting documents and people related to the concepts mentioned in the messages.

The last type of awareness discontinuity that we will discuss in this paper is related to the awareness of people within the physical environment. I have a good sense of awareness of when my colleagues that are immediately adjacent to my office are available, but have very little awareness of the availability of colleagues just two or three doors away. Although various methods exist for contacting these people in the electronic realm (electronic mail, telephone, voice mail), there exist certain kinds of issues for which face-to-face interaction is necessary. For example, when I want to ask a colleague about how to solve a problem, and I'm not quite sure how to describe the problem, real-time interaction is necessary.

Several research labs are looking at how to support collaboration across remote sites, but we are looking at how we can better support collaboration among physically proximate colleagues. One application we have developed, ACTIVEMAP [McCarthy & Meidel, 1999], shown in Figure 5, superimposes images of coworkers over the locations in which they were last seen (or rather, the locations in which their infrared badges were last detected). The system uses a variety of mechanisms to represent the freshness of the location information, so that users can make inferences about the potential availability of colleagues. The most common use of the application is in a kiosk installed in a hallway in the CSTaR work area, where people can track someone down when they can't find them in their office.




Figure 5. ActiveMap

In addition to its use as an awareness mechanism for physically proximate colleagues, the tool is also proving useful for distributed workers. I can login into the system from an outside location (e.g., home, another office, etc.) and easily see "what's going on" – who is in the office, who is talking to whom, who is attending a meeting, etc. If I want to reach a person, I can send a message to be

read through the loudspeaker nearest to that person's current location. I can also set up automatic rules that would be triggered by specific events such as someone returning to his or her office.

Although there are privacy issues involved in this work, as well as most other situated computing applications, people are generally willing to give up some of their privacy for commensurate benefits, as evidenced by the widespread use of credit cards. Thus far, we have managed to provide a level of benefit that is perceived to outweigh the loss of privacy in ACTIVEMAP; however, we readily acknowledge the need to be vigilant about this issue in our efforts to develop other situated computing applications.

6. Conclusion

We believe that Situated Computing represents an important new class of contextually embedded applications – providing computing, connectivity and content in physical situations where they are most needed, i.e., when there exists a gap between a person's intentions and ability to act. We have defined three types of discontinuity – physical discontinuity, information discontinuity and awareness discontinuity – that give rise to such gaps, and highlighted a number of applications that attempt to close these gaps. As this collection of applications demonstrates, situated computing has the potential for defining entirely new “moment of value” opportunities, enabling people to achieve their intentions more frequently and more successfully.

While much attention is currently devoted to businesses that deal exclusively in the virtual domain, the most significant business opportunities in the future lay in the convergence of the virtual and physical worlds, in applications that take full advantage of what each world has to offer. Situated computing applications take this one step further, by tailoring the best of this convergence to the specific contexts in which people have opportunities, currently unrecognized or unavailable, to achieve specific intentions.

As we have shown in this paper, we are engaged in a number of projects that explore what kinds of contexts might provide important applications for situated computing. We will continue to push on in this area, identifying new contexts and new technologies that can be used to fit into and take advantage of these contexts. After all, despite all the advances in the virtual world, we still live in the physical world, and as we enhance situations in the physical world with relevant connections to the virtual world, our experiences will be greatly enriched.

7. References

ABC World News Tonight, November 12, 1998.

Adam B. Brody and Edward J. Gottsman. 1999. Pocket BargainFinder: A Handheld Device for Augmented Commerce. To appear in *Proceedings of the International Symposium on Handheld and Ubiquitous Computing (HUC '99)*. Karlsruhe, Germany.

Doug Bryan and Anatole Gershman. 1999. Opportunistic Exploration of Large Consumer Product Spaces. To appear in *Proceedings of the ACM Conference on Electronic Commerce (EC '99)*.

Kelly L. Dempksi. 1999. Augmented Workspace: The World as Your Desktop. To appear in *Proceedings of the International Symposium on Handheld and Ubiquitous Computing (HUC '99)*. Karlsruhe, Germany.

Andrew Fano. 1998. Shopper's Eye: Using Location-based Filtering for a Shopping Agent in the Physical World. *Proceedings of the Second International Conference on Autonomous Agents (Agents '98)*, Minneapolis, MN. pp. 416-421.

Richard Hull, Philip Neaves and James Bedford-Roberts. 1997. Towards Situated Computing. In *Proceedings of the First International Symposium on Wearable Computing (ISWC '97)*. Cambridge, MA.

Bruce Krulwich. 1996. The BargainFinder Agent: Comparison Price Shopping on the Internet. In *Bots and Other Internet Beasties*. J. Williams (ed). SAMS.NET.

Joseph F. McCarthy and Eric S. Meidel. 1999. ActiveMap: A Visualization Tool for Location Awareness to Support Informal Interactions. To appear in *Proceedings of the International Symposium on Handheld and Ubiquitous Computing (HUC '99)*. Karlsruhe, Germany.

M. V. Nagendra Prasad and Theodore Anagnost. 1998. InfoScout: An Active Recommender Agent. In *Proceedings of the AAAI-98 Workshop on Recommender Agents*, Madison, Wisconsin.

Dadong Wan. 1999. Magic Medicine Cabinet: A Situated Portal for Healthcare. To appear in *Proceedings of International Symposium on Handheld and Ubiquitous Computing (HUC'99)*. Karlsruhe, Germany.