

International Workshop on Location- and Context-Awareness (LoCA 2005), May 12-13, 2005, Oberpfaffenhofen, Germany.

Web-Enhanced GPS

Ramaswamy Hariharan¹, John Krumm², and Eric Horvitz²

¹ School of Information and Computer Sciences
University of California, Irvine
444 Computer Science Building
Irvine, CA 92697
USA
rharihar@ics.uci.edu

² Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
USA
{jckrumm, horvitz}@microsoft.com

Abstract. Location-based services like reminders, electronic graffiti, and tourist guides normally require a custom, location-sensitive database that must be custom-tailored for the application at hand. This deployment cost reduces the initial appeal of such services. However, there is much location-tagged data already available on the Web which can be easily used to create compelling location-aware applications with almost no deployment cost. Such tagged data can be used directly in applications as well as to provide evidence in models of activity. We describe three applications that take advantage of existing Web data combined with location measurements from a GPS receiver. The first application, “Pinpoint Search”, finds web pages of nearby places based on GPS coordinates, queries from a Web mapping service, and general Web searches. The second application, “XRay”, uses the mapping service to find businesses in a building by pointing a GPS-equipped electronic compass at the building. The third application is called “Travelogue”, and it builds a map and clickable points of interest to help automatically annotate a trip based on GPS coordinates. Finally, we discuss the use of Web-based data as rich sources of evidence for probabilistic models of a user’s activity, including a means for interpreting the explanation for the loss of Web signals as users enter structures.

1 Introduction

Location-based services use knowledge of a user’s location to index into services and data that are likely useful at that location. For instance, a reminder application like comMotion[1] can give the user relevant information at a given location, like, “You’re near a grocery store, and you need milk at home.” A so-called “electronic graffiti” system, such as Stick-e Notes[2], supports users who want to leave electronic notes for themselves or others that are associated with a particular location, like “There is a

better Thai restaurant one block north of here.” Location-based tour guides such as Cyberguide[3] offer relevant information about the exhibit or site at which the user is standing. These and most other location-based services share a need for a custom database dedicated to storing and serving data for specified locations. Reminder systems must have reminders, electronic graffiti needs digital tags, and tour guides need site information.

While a custom store of location-indexed data can lead to interesting applications, there is already a wealth of location data available on the Web that can be exploited for location services without the data deployment costs of traditional applications of this type. This paper demonstrates three applications that use existing location data on the Web in conjunction with position information from a GPS receiver. In this way, we avoid the deployment costs of creating a database of location information, relying instead on what is already available. These applications show that it is possible to create useful location-based services using existing location data.

The three applications are:

- Pinpoint Search – Convert a measured (latitude, longitude) into search terms for Web searches, giving web pages relevant to the user’s immediate surroundings.
- XRay – Point a pose-sensitive device at a scene of interest and get a list of what businesses are situated along that direction.
- Travelogue – Help a user recall points of interest from a trip logged with GPS data.

Our maps and point-of-interest data come from the Microsoft MapPoint Web service, which requires a subscription fee. However, using this service is much more economical than building a custom store of location data. Other point-of-interest databases could be used as well. We have implemented these three applications on a desktop computer using real GPS data. Their ultimate target is mobile users, and we foresee few problems modifying the applications for use on a PDA or cell phone. The target device would need a Web connection and a GPS receiver.

2 Pinpoint Search

Pinpoint Search is designed to give information about a mobile user’s immediate surroundings. We assume the user is equipped with a GPS receiver and Web-connected mobile computer. Starting with a (latitude, longitude) from the user’s GPS, Pinpoint Search uses MapPoint to compute the nearest street address. This is shown in Figure 1(a) which is a screen shot of a working mockup of Pinpoint Search’s client and server parts. In this offline demonstration program, each stored (latitude, longitude) can be resolved, which triggers the street address lookup, the results of which are shown in a popup window on the map.

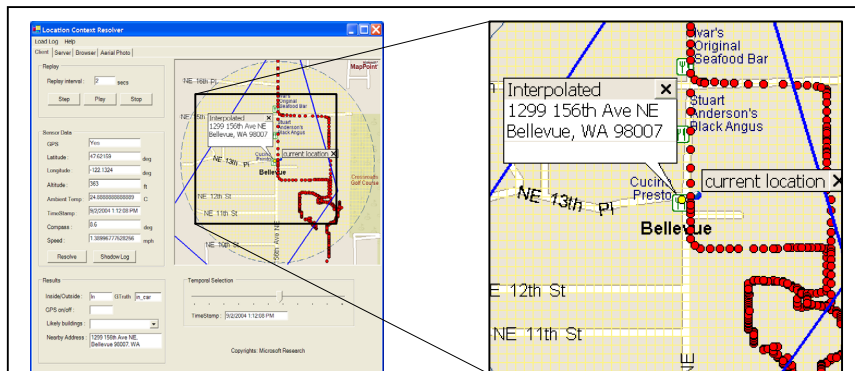
The conversion from (latitude, longitude) to street address is important, because the street address serves as a good search term for Web searches. This is illustrated in Figure 1(b), which shows a screen shot giving the result of an MSN[®] search automati-

cally performed on the resolved street address within the Pinpoint Search program. The search results give links to web pages that contain the nearby street address, resulting in entries for a nearby acupuncturist, chiropractor, and restaurant.

Not all relevant Web pages contain a street address that matches our search term, so we also convert (latitude, longitude) into a list of nearby businesses using the Map-Point Web Service. This service uses a database of business locations from Acxiom[®] that are categorized by type, such as food stores, automobile dealers, and restaurants. As shown in Figure 2(a), this lookup results in three nearby restaurants, one of which is called “Stuart Anderson’s Black Angus”. After the user clicks on this result, Pinpoint Search automatically runs another Web search using the restaurant name as a search term. One of the search results is shown in Figure 2 (b), which gives positive and negative reviews of the restaurant. Starting with GPS coordinates, the user could find this review page with only two mouse clicks using Pinpoint Search.

Our implementation of Pinpoint Search runs on a desktop computer and uses stored GPS coordinates for demonstration. It would be straightforward to port it to PDA or cell phone equipped with GPS.

Pinpoint Search shares similarities with other programs. Google Local[4] and Yahoo! Local[5] allows Web searches around a certain geographic area, specified by parts of a street address. Pinpoint Search adds the important preprocessing step of



(a) Map shows stored locations. Popup titled “Interpolated” shows the nearest street address to the currently selected location.

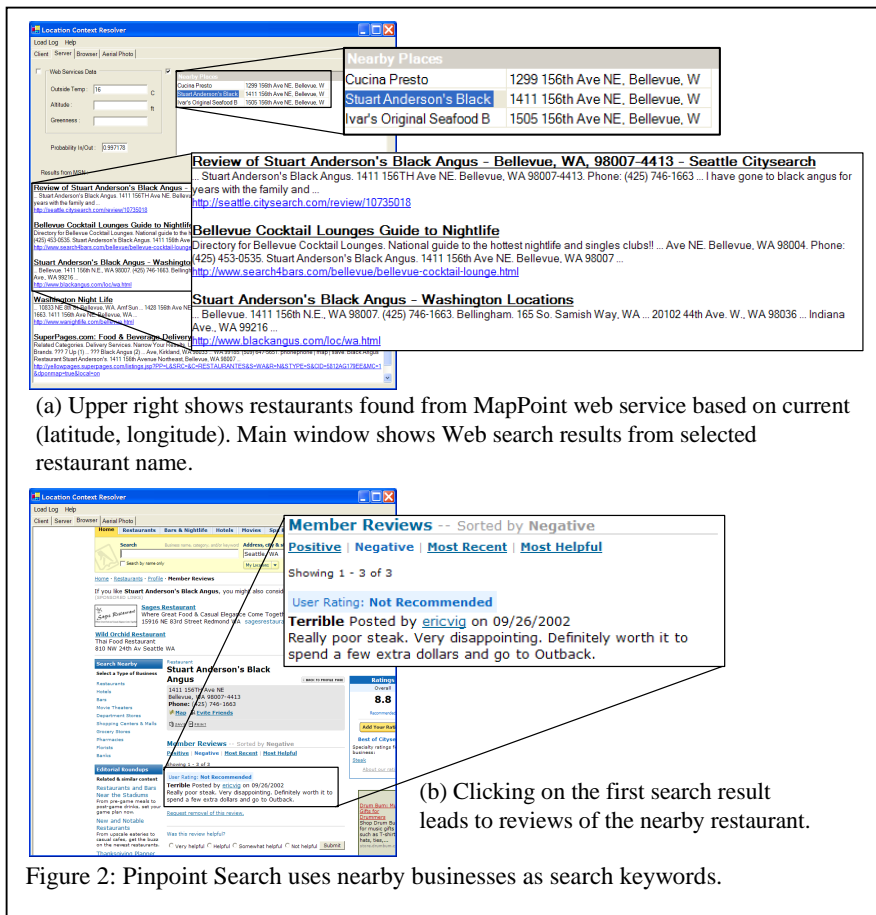


(b) Web search based on found street address from screen above finds nearby acupuncturist, chiropractor, and restaurant.

Figure 1: Pinpoint Search converts (latitude, longitude) to a street address, which is used as a search keyword.

converting a numerical (latitude, longitude) into a street address or business name, which allows the local search to be completely automated based on a GPS receiver. Thus, to use Pinpoint Search, the user does not need to know a street address, postal code, or even the name of the city. Because it starts with a GPS coordinate, the search results are likely to be much more focused on the user's immediate surroundings.

Another similar project is AURA[6], which uses bar codes and other means to create Web search terms, much like we use location data to generate addresses and business names for searches. Like these projects, Pinpoint Search requires no special databases, but instead exploits the extensive amount of data already available on the Web. It is unique in that it uses a measured position to ultimately index into web pages about nearby things.



3 XRay

XRay is a concept for a pose-sensitive query device designed to allow a mobile user to point toward an outside object and discover what is inside or behind it. It is based on a device containing a GPS receiver, electronic compass, and network connection. The user physically points the device at something and issues a query. XRay responds with a list of businesses or other points of interest along the direction of pointing. A working mockup of the program is illustrated in Figure 3. Here XRay has been pointed toward a street that the user cannot yet see, and its "field of view" has been adjusted to query over the length of the unseen block. The list at the right shows different categories of items returned from MapPoint's[®] Acxiom[®] database. The "Apparel" and "Eating and drinking" categories have been expanded to show the names of places inside the query cone. Clicking on one of these places puts an icon on the map at its location. Using this program, users can quickly get a sense of what is around them by simply pointing in a direction of interest. As with Pinpoint Search, XRay relies on a rich, existing database of places, meaning that it works "out of the box" with no deployment cost.

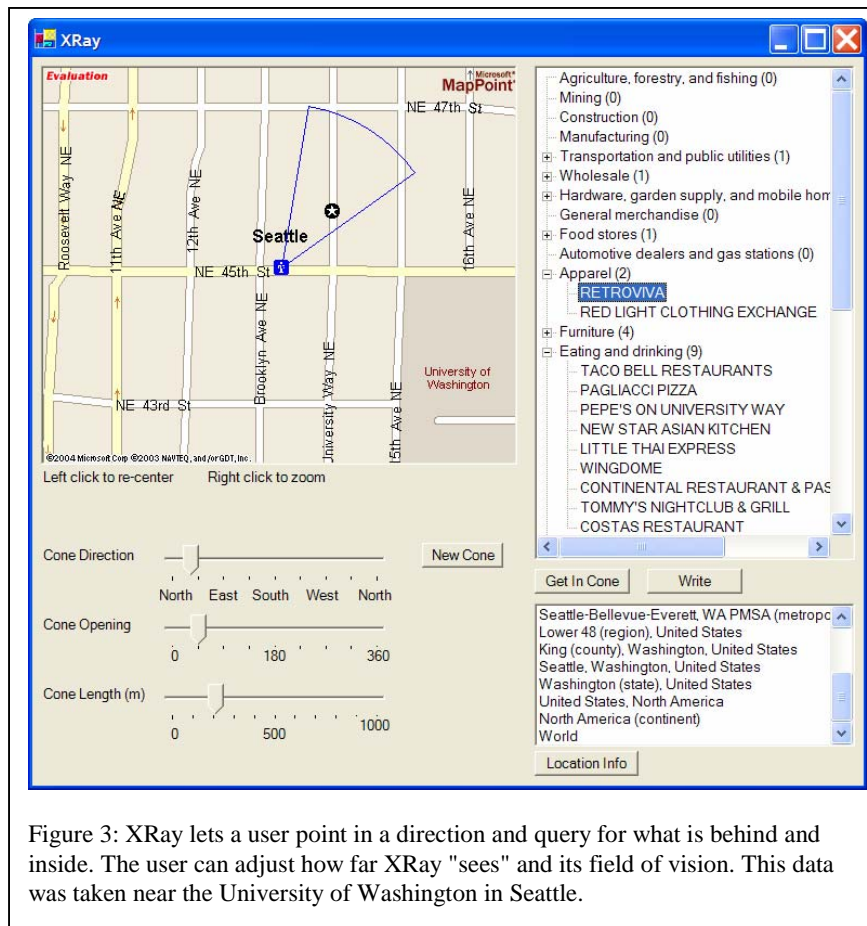


Figure 3: XRay lets a user point in a direction and query for what is behind and inside. The user can adjust how far XRay "sees" and its field of vision. This data was taken near the University of Washington in Seattle.

Although XRay shows real results, we have not yet found a commercially available combination of devices that would let us exercise it in the field. While there are GPS receivers with a built-in compass, we have not yet found one that can interface to a handheld computer with ubiquitous network access. An ideal platform would be a cell phone with a GPS, an electronic compass, and Web access, which we expect will soon be easy to buy or make. XRay is similar to augmented reality systems, e.g. [7, 8] that project extra information into the user's view. In fact, the "Touring Machine"[7] is capable of projecting information from campus buildings that the user can see. Similarly, UCSD's Active Campus project[9] can show information on a handheld computer about nearby buildings based on positioning with Wi-Fi triangulation. XRay shows that this type of application can exploit existing data on the Web instead of custom-authored data as previous systems required.

4 Travelogue

Pinpoint Search and XRay are both designed to provide real time data to the user. In contrast, Travelogue is aimed at analyzing a trip after it happens. Specifically, Travelogue annotates a sequence of (latitude, longitude) points with points of interest. In this way, it helps a user recall what he or she might have visited or seen during a trip without requiring the user to make any notes during the trip. A screenshot of Travelogue showing points of interest is shown in Figure 4.

Travelogue starts with a list of (latitude, longitude) points logged from a GPS receiver. Many GPS receivers have a logging feature built in, so a user can simply set up his or her GPS to periodically save location data during the day for later downloading. Travelogue analyzes this data to find places that might later be of interest to the user. In our implementation, these places are those locations where the number of GPS satellites visible to the receiver dropped below the minimum four that are required for a full GPS fix. This heuristic works well in the urban environment where we tested, because most of the "interesting" places we visited were indoors where the satellites are occluded by the building. Other more sophisticated methods of finding interesting places include algorithms like [10] that finds locations where the user dwells. The points of interest found during one of our short driving trips are shown in Figure 4 as red squares.

Clicking on a point of interest triggers the program to display the (latitude, longitude) of the point, the time of arrival, and the amount of time spent there. Using our heuristic, the time spent is simply the time over which at least four GPS satellites could not be detected. Clicking on a point of interest also populates multiple drop-down lists showing nearby places found in MapPoint's[®] Acxiom[®] database. In Figure 4, the "Eatery" list has been expanded to show nearby restaurants. While the user did not necessarily visit every place on these lists of nearby places, the lists give a simple, fast way of jogging the user's memory for what places he or she actually did visit.

We envision normal users using Travelogue as a way of automatically annotating business and pleasure trips without having to worry about making explicit notes on where they were. Travelogue could also be used as a memory aid for people with

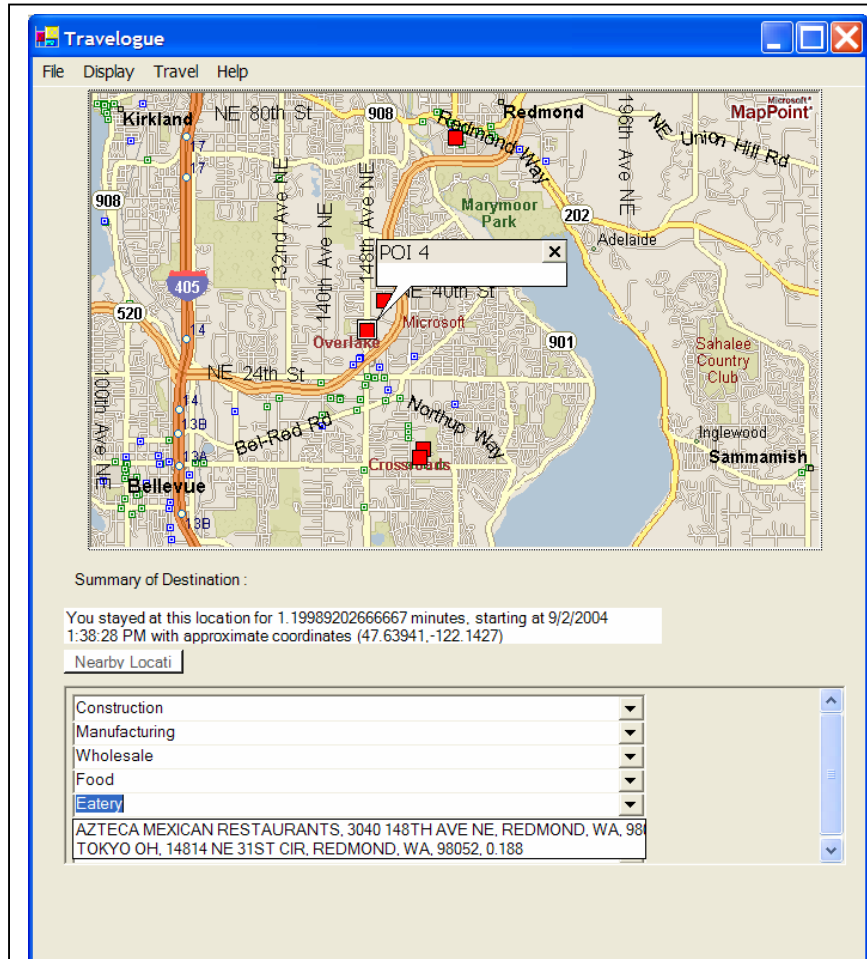


Figure 4: Travelogue discovers points of interest and displays them as red squares. Clicking on these points triggers the program to display the amount of time spent there and a list of nearby places. This can help a user remember what he or she did during a trip.

cognitive decline as a way to help them remember what they have been doing and as a way for their caregivers to assess their activities.

5 Web Data and Models of Activity

Semantic content associated with locations on the web can provide rich sources of evidence about users' activities over time. We have been exploring general probabil-

istic models with the ability to fuse multiple sources of information. Such models can be used to perform inferences about a user's activities from the historical and short-term GPS data, as well as extended sensing with such information as temperature[11], barometric pressure, ambient light and sound, and Web data. Web content can be used to update, in an automated manner, a set of key resources and venues available at different locations, providing Bayesian dependency models with sets of resources that are coupled to an ontology of activities (*e.g.*, shopping, restaurants, recreation, government offices, schools, entertainment etc.). Such information can be used as a rich source of evidence in a probabilistic model that computes the likelihood of different plausible activities. Inferences can further take into consideration the dwelling of a user at a location with zero or small velocities and the complete loss of GPS signals at particular locations for varying periods of time, indicating that a user has entered a structure the blocks receipt of GPS signals. The timing, velocity, and frank loss of signal after a slowing of velocity provide rich evidence about a user's interests or entries into different proximal buildings and structures, as characterized by the content drawn from the Web about resources in the region of the last seen GPS coordinates. Such reasoning can be enhanced by a tagged log of prior activities noted by a user. Reasoning about losses of GPS signal can take into consideration a log of known "GPS shadows," that are not associated with being inside buildings, such as those occurring inside "urban valleys," as GPS access is blocked by tall structures.

We are pursuing rich probabilistic models of activity and location based on multiple sources of information, including information available from the Web and from logs of prior activities and GPS availability. Figure 5 displays a time slice of a more general dynamic Bayesian network model, showing probabilistic dependencies among key measurements and inferences. Server icons signify access of information from the Web about local resources as well as access from a store of known GPS shadows. The model is designed to make inferences about the probability distribution over a user's activities and over the location of a user, even when GPS signals are unreliable or lost temporarily. Sub-inferences include computation about whether a user is indoors or outdoors, employing information about the loss of GPS signals, a log of GPS shadows, information about local resources to the current location, and sensed temperature. We call out two key variables from an adjacent, earlier time slice to highlight the potential value of including dependencies among variables in adjacent time slices.

6 Conclusion

Knowledge of a user's raw (latitude, longitude) is not normally very useful. However, with publicly available databases, location measurements can be converted into useful information. We first reviewed three simple applications. Pinpoint Search takes (latitude, longitude) and finds web pages that are pertinent to the user's immediate surroundings. XRay lets a user point in a direction of interest and retrieve listings of places in that direction. Travelogue helps a user remember interesting places encountered during a trip. These applications share the trait of starting with raw GPS readings and using publicly available Web data to produce useful information. Although each

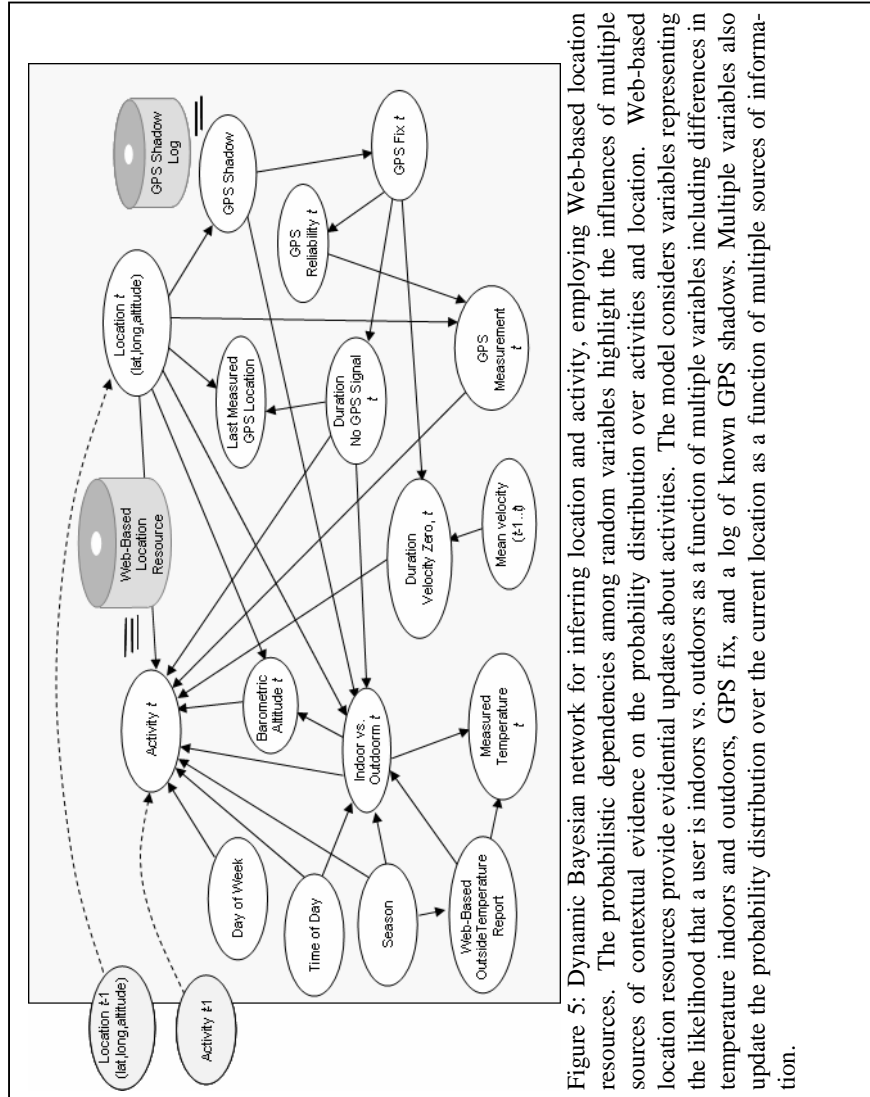


Figure 5: Dynamic Bayesian network for inferring location and activity, employing Web-based location resources. The probabilistic dependencies among random variables highlight the influences of multiple sources of contextual evidence on the probability distribution over activities and location. Web-based location resources provide evidential updates about activities. The model considers variables representing the likelihood that a user is indoors vs. outdoors as a function of multiple variables including differences in temperature indoors and outdoors, GPS fix, and a log of known GPS shadows. Multiple variables also update the probability distribution over the current location as a function of multiple sources of information.

of these applications has the potential for more development, together they show how the Web can be used in a simple way to enhance GPS. Finally, we reviewed the prospect of using the Web, in addition to other sources of information, to support rich probabilistic inferences about a user's activities and location. Such inferences can provide a window into a user's activities as well as access to location information even when GPS fixes become erroneous or are lost completely. Indeed, such models can take losses of GPS signal as valuable evidence for making inferences about activities and location.

References

1. Marmasse, N. *comMotion: A Context-Aware Communication System*. in *CHI 99*. 1999.
2. Pascoe, J. *The Stick-e Note Architecture: Extending the Interface Beyond the User*. in *IUI 97*. 1997. Orlando, Florida, USA.
3. Abowd, G.D., et al., *Cyberguide: A Mobile Context-Aware Tour Guide*. *ACM Wireless Networks*, 1997(3): p. 421-433.
4. <http://local.google.com/>.
5. <http://local.yahoo.com/>.
6. Smith, M., D. Davenport, and H. Hwa. *AURA: A mobile platform for object and location annotation*. in *Fifth International Conference on Ubiquitous Computing - UbiComp 2003*. 2003. Seattle, WA, USA.
7. Feiner, S., et al. *A touring machine: Prototyping 3D mobile augmented reality systems for exploring the urban environment*. in *First IEEE Int. Symp. on Wearable Computers (ISWC '97)*. 1997. Cambridge, MA, USA.
8. Kutulakos, K.N. and J.R. Vallino, *Calibration-Free Augmented Reality*. *IEEE Transactions on Visualization and Computer Graphics*, 1998. **4**(4): p. 1-20.
9. Griswold, W.G., et al., *ActiveCampus - Sustaining Educational Communities through Mobile Technology*. 2002, University of California at San Diego: San Diego, CA, USA.
10. Hariharan, R. and K. Toyama. *Project Lachesis: Parsing and Modeling Location Histories*. in *Third International Conference on GIScience*. 2004. Adelphi, MD, USA.
11. Krumm, J. and R. Hariharan. *TempIO: Inside/Outside Classification with Temperature*. in *Second International Workshop on Man-Machine Symbiotic Systems*. 2004. Kyoto, Japan.