

Enabling Ad-Hoc Collaboration Through Schedule Learning and Prediction

Daniel Ashbrook and Thad Starner
College Of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280 USA
{anjiro, thad}@cc.gatech.edu

Abstract

The transferal of the desktop interface to the world at large is not the goal of ubiquitous computing. Rather, ubiquitous computing strives to increase the responsiveness of the world at large to the individual. A large part of this responsiveness is improved communication with other individuals. In this paper we describe a system that can enable ad-hoc collaboration between several people by creating a model of the daily schedules of individuals and by performing predictions based on this model. Using GPS data we learn to distinguish locations and track the times that these locations are visited. In addition, we use Markov models to predict which locations might be visited next based on the user's previous behavior.

1 Introduction

For any user-assisting technology to be useful and not merely irritating, it must have some knowledge of the user to be assisted: it must understand—or at least predict—what the user will do, when she will do it, and, ideally, the reason for her actions. User modeling is a necessary step toward answering some of these questions. Csinger defines user modeling as “. . . the acquisition or exploitation of explicit, consultable models of either the human users of systems or the computational agents which constitute the systems” [3].

As we use stochastic processes for user modeling, our system requires the collection of large amounts of data. Ideally, we will constantly collect relevant data on our user, and use this data to continuously adapt and refine the model. Practically, however, privacy and technological considerations must be taken into account. A fully ubiquitous computing infrastructure is still many years away, so we must rely on wearable sensors for a full range of data collection abilities. Additionally, wearable technology tends to be more

privacy-sensitive: data is usually stored with the sensor, on the body, and only uploaded later. This gives the user the opportunity to remove any data she does not wish to have integrated into the model.

For the purposes of our research, we have chosen to use a wearable global positioning system (GPS) receiver with an attached data logger. This device satisfies the constraints of privacy and ubiquity; in addition, its data is easily interpretable and the hardware itself is quite small and lightweight. GPS has the further advantage that it is globally available, and even the absence of its signal may provide useful information, as we will discuss.

GPS provides several useful pieces of information along with latitude and longitude. The satellite transmits date and time, and the receiver calculates altitude and heading. As we will demonstrate, position and date/time information is sufficient to perform simple user modeling with a fairly high degree of accuracy. These two pieces of data allow us to get an idea of where the user will go next, given her current location.

2 Potential Applications

2.1 Individual Benefits

As noted in the *comMotion* system by Marmasse and Schmandt [6], an agent that learns locations of importance to the user may become very valuable. “To-do” lists and reminders can be associated with specific physical locations. If the agent can predict which location the user will visit next, reminders can be presented to the user in a timely fashion. One example Marmasse provides is that of reminding the user of her shopping list as she nears a grocery store. Reminding the user as she passes the store would be frustrating and distracting. However, reminding the user several miles in advance, or even as she enters her car, would be more productive. In addition, if

the agent can determine the user’s goal in advance, it may suggest alternative routes based on current traffic conditions.

With mobile computers, such a predictive agent of the user’s travel can also help with systems issues. For example, due to lack of infrastructure, radio shadows due to buildings, and power requirements, wireless networks are often inaccessible. However, with a wearable computer or PDA, this lack may be hidden from the user by caching. For example, if the user composes an e-mail while in the subway, the wearable may add the message to its outgoing queue and wait until the network is available to send it. On the other hand, if the message is urgent, this behavior may not be appropriate. The user’s agent may predict that the user will not be within range of the wireless network and alert the user to travel out of her way to guarantee connection. For less urgent e-mail and Internet services, the agent may delay transmission even when network connection is available. Energy is one of the most precious resources for mobile devices, and the amount of energy needed to transmit a message may go up with the fourth power of distance in some situations [2]. In addition, the cost of transmitting a message may vary with the time of day and type of service used. Thus, a predictive agent could optimize its transmissions based on its prior knowledge of how the user moves throughout the day.

2.2 Collaborative benefits

When user models are shared, other benefits can be realized. User models may be shared by replicating the relevant models on each individual’s computer, querying a given user’s model each time information is needed, or by delegating the coordination of user models to a central service. For convenience, we will assume the last case and will call this service SAM: the Service for Accessing Models.

The most straightforward service that SAM might offer is determining potential meetings times for a given set of individuals. By having access to a model of users’ locations during a typical week, SAM can provide a first pass at this common problem.

More interestingly, SAM might encourage serendipitous meetings. For example, if SAM predicts that Don’s friend Cindy will be traveling near his office around lunch time and that Don will be available, it may suggest to Cindy that she call Don to make arrangements to meet.

SAM may also be able to suggest more complicated interactions that are not immediately apparent to even a cooperative group. The following scenario

illustrates this point.

Alice and Bob have a weekly Tuesday meeting with the rest of the research group. One week, the subject turns to encryption technology, and Alice promises to bring Bob one of her books on cryptography. Later, Alice realizes she won’t see Bob until the next meeting but would like to get him the books sooner. She asks SAM for help, and it discovers that Eve, with whom Alice eats dinner on Wednesdays, will see Bob on Friday. SAM asks Eve if she will deliver the book for Alice. When Eve agrees, Alice is notified, and brings the book to dinner. Eve delivers the book to Bob on Friday, and he is able to spend the weekend reading in preparation for a discussion with Alice on Tuesday.

Such a series of interactions alludes to a more complicated favor-transfer mechanism. In the above scenario, Eve does not directly benefit from the service she provides Alice and Bob except perhaps in a building of goodwill between friends. Suppose instead that Eve was unknown to Alice and Bob but offered her services in general to SAM in exchange for some form of reward. This reward may be an increase in “karma” points that can be exchanged for similar favors from SAM’s pool of participants at a latter time. While such an idea may seem at first far-fetched (indeed, the concept has been the basis of the science fiction book *Distraction* by Bruce Stirling), it may be appropriate in the corporate world. Web sites Advogato [4] and Slashdot [5] have already developed similar trust systems for the evaluation of new postings and software development.

3 Previous Work

Location prediction systems have become of interest in the cellular network community in recent years. The United States government wants to be able to locate people who place 911 calls from cell phones, and various location-based contextual services are being discussed. Bhattacharya and Das [1] describe a cellular-user tracking system for call delivery that uses transitions between wireless cells as input to a Markov model. As users move between cells, or stay in a cell for a long period of time, the model is updated and the network has to try fewer cells to successfully deliver a call.

As the number of cellular subscribers grows, it becomes ever more imperative to manage the admission

of telephones to cell sites. Yu and Leung [9] have developed a system using Bhattacharya and Das’ mobility prediction to guess which cell a user is most likely to visit next and reserve the appropriate resources.

The problem of identifying significant locations has been explored, both by Wolf et. al. [8], who used resting time to mark starting and ending points of trips, and by Marmasse and Schmandt [6], who modeled locations based on loss of GPS signal to deliver contextually-useful information.

4 GPSModel

In order to begin investigating the kind of serendipitous agents we describe above, we constructed a program to model an individual’s travel. The program, GPSModel, reads in a file of GPS coordinates and outputs a list of destinations, a list of trips taken, and probabilities for each destination.

This allows queries by an agent such as “My user is currently at home. What is the most likely place she’ll go next?” and “How likely is my user to stop at the grocery store on the way home from work?” The answers to these questions for several users can lead to serendipitous meetings: if the response to “Where is my user most likely to go next?” is the same for two colleagues, they can be alerted that they’re likely to meet each other.

4.1 Data collection and manipulation

Data collection was performed with a standard Garmin portable GPS receiver and the GeoStats GeoLogger [7]. The GeoLogger records GPS data at one-second intervals, but only if the receiver is moving at over one mile per hour and has line-of-sight to the GPS satellites. Data was collected from one user over a period of four months, with travel mainly in the Atlanta area. One can clearly discern the user’s home and work areas in Figure 1 by the higher density of dots in those locations.

While it might seem strange to throw away data points, logging data only while the user is traveling at greater than one mph actually helps pre-process our data. Since we are largely interested in the locations where the user spends her time, rather than how she gets there, we can consider any point with a certain amount of time between it and the previous logged point as significant. If the point is an aberration, such as a long wait on the highway due to an accident or construction, then it will have only one occurrence against all the other recorded points, and

hence the probability of it being counted as a destination will be quite low. If, on the other hand, the user stops at the same place regularly, such as home, work, or the coffee shop, these points will have a high relative occurrence rate and the probability of being considered a destination will be correspondingly high.




Figure 1: *Significant points in the Atlanta area, as determined by the 10-minute rule*

After some experimentation, we found that 10 minutes was a good amount of stopping time to consider a GPS coordinate as a place of interest (Figure 1). Because a GPS measurement taken in the same physical location can vary by as much as 10 meters, the recorder will not log exactly the same GPS point for a location even if the user stops for 10 minutes at precisely the same point every day. For this reason, we used a variant on the k-means clustering algorithm. The basic idea is to take one of the significant points and a radius, such as a half mile. All the points within the radius are marked, and the mean of these points is found. The mean is then taken as the new center point, and the process is repeated. This continues until the mean stops changing. When the mean no longer moves, all points within its radius are removed from consideration, and the process repeats until no points remain and we are left with a collection of clusters of points.

Because our original radius was fairly arbitrary, informed only by what we know about the limitations of GPS and the general scale of buildings and parking lots, we process the data a second time to find any sub-groups we may have missed the first time. This is

Transition	Relative Frequency	Probability
$A \rightarrow B$	14/20	0.700000
$A \rightarrow B \rightarrow A$	3/14	0.214286
$A \rightarrow B \rightarrow D$	3/14	0.214286
$A \rightarrow B \rightarrow E$	1/14	0.071429
$A \rightarrow B \rightarrow F$	1/14	0.071429
$A \rightarrow B \rightarrow G$	1/14	0.071429
$A \rightarrow B \rightarrow C$	2/14	0.142857
$A \rightarrow B \rightarrow H$	1/14	0.071429
$A \rightarrow B \rightarrow I$	1/14	0.071429
$B \rightarrow A$	16/77	0.207792
$B \rightarrow A \rightarrow B$	13/16	0.812500
$B \rightarrow A \rightarrow J$	3/16	0.187500
$B \rightarrow C$	10/77	0.129870
$B \rightarrow C \rightarrow A$	6/10	0.600000
$B \rightarrow C \rightarrow K$	4/10	0.400000
$D \rightarrow B$	5/7	0.714286
$D \rightarrow B \rightarrow A$	2/5	0.400000
$D \rightarrow B \rightarrow L$	2/5	0.400000
$D \rightarrow B \rightarrow M$	1/5	0.200000

Table 1: First and second order probabilities using data recorded from Sep–Dec 2001

accomplished by taking every cluster and running it through the same clustering algorithm as above, but several times. Each time the radius is changed by a small amount, and the result is plotted on a graph. We then look for a knee in the graph—this signifies a radius just before the number of clusters begins to converge to the number of points. If a knee exists, its corresponding radius is used to form sub-clusters. This has the effect of allowing multiple scales for our data—a human, or walking scale, such as a college campus; and an automobile scale, as for a city. There could also be larger, state-scale and country-scale clusters.

Once significant clusters of locations are found, the original stream of GPS data is analyzed to what clusters the user moved between. A Markov model is generated with each node representing one cluster, and the probabilities of each transition is calculated. The probabilities can be calculated for any order of model—that is, looking at the probability of visiting a certain location given a previous location (order one), or given two previous locations (order two) and so on. For some order one and order two predictions based on our data, see Table 4.1.

5 Future Work

While we have gained some useful results thus far, we would like to improve the performance of GPSModel.

One problem is lack of data - we only have about two months worth of data, or 97,000 data points; out of these, we get just 147 significant locations. We have not yet implemented time-of-day or day-of-week data, e.g, a trip to work on Monday is compared with a trip to church on Sunday.

While we do find clusters at multiple scales, we do not currently perform prediction on each of these scales. We would like to implement schedule learning on multiple levels, so movement between buildings on a college campus could be predicted as easily as movement between areas of a city.

Many of our examples involved combination of several users’ data; currently, we have only recorded data from one user and have thus been unable to experiment with multiple data sets. We plan to track several users and perform full-scale experiments on collaboration.

References

- [1] Amiya Bhattacharya and Sajal K. Das. Lezi-update: An information-theoretic approach to track mobile users in PCS networks. In *Mobile Computing and Networking*, pages 1–12, 1999.
- [2] Jae-Hwan Chang and Leandros Tassiulas. Energy conserving routing in wireless ad-hoc networks. In *INFOCOM (1)*, pages 22–31, 2000.
- [3] A. Csinger. *User Models for Intent-based Authoring*. PhD thesis, 1995.
- [4] Raph Levien. Advogato. <http://advogato.org>, 2002.
- [5] Rob Malda. slashdot.org. <http://slashdot.org>, 2002.
- [6] Natalia Marmasse and Chris Schmandt. Location-aware information delivery with commotion. In *HUC*, pages 157–171, 2000.
- [7] Jean Wolf. Geostats. <http://www.geostats.com>, 2002.
- [8] Jean Wolf, Randall Guensler, and William Bachman. Elimination of the travel diary: An experiment to derive trip purpose from gps travel data.
- [9] Fei Yu and Victor C.M. Leung. Mobility-based predictive call admission control and bandwidth reservation in wireless cellular networks. In *IEEE International Conference on Communications*, 2001.