

Opportunity Knocks: a System to Provide Cognitive Assistance with Transportation Services

Donald J. Patterson, Lin Liao, Krzysztof Gajos, Michael Collier, Nik Livic,
Katherine Olson, Shiaokai Wang, Dieter Fox, and Henry Kautz

University Of Washington, Seattle WA 98195, USA,
{djp3,liaolin,kgajos,kolson,fox,kautz}@cs.washington.edu
{nlivic,mjc6,shiaokai}@u.washington.edu

Abstract. We present an automated transportation routing system, called “Opportunity Knocks,” whose goal is to improve the efficiency, safety and independence of individuals with mild cognitive disabilities. Our system is implemented on a combination of a Bluetooth sensor beacon that broadcasts GPS data, a GPRS-enabled cell-phone, and remote activity inference software. The system uses a novel inference engine that does not require users to explicitly provide information about the start or ending points of their journeys; instead this information is learned from users’ past behavior. Furthermore, we demonstrate how route errors can be detected and how the system helps to correct the errors with real-time transit information. In addition we present a novel solution to the problem of labeling positions with place names.

1 Introduction

For many individuals, mobility in the community means using public transportation. It is key to their social life, their employment, and their ability to receive goods and services. Unless they can successfully move through their community they cannot lead an independent life. Public transportation, however, can be daunting for anyone who is born with below average cognitive abilities or whose cognitive abilities have begun to decline, however slightly. There is often no choice but for them to give up their potential future independence and be under direct supervision of their care givers or family members; a healthy individual is needed to detect situations where a mistake made by a cognitively disabled person may cause distress or harm. Thus, the inability to safely use public transportation harms their quality of life as well as that of their formal and informal support network [1–3]. However, if impaired individuals had effective compensatory cognitive aids to help them use public transportation, their independence and safety would improve, they would have new opportunities for socialization and employment, and stress on their families and care givers would be reduced.

We developed a ubiquitous computing system, called “Opportunity Knocks,” (*OK*) to explore the feasibility of just such a cognitive aid. This system targets

mentally retarded individuals and individuals with traumatic brain injury, who are generally high functioning but unable to use public transportation due to short-term confusion or memory lapses. These individuals generally show stable levels of cognitive ability over time, are employed, and are either using specialized transportation services or using public transportation with marginal efficacy.

While our system is immediately targeting mentally retarded people and people with traumatic brain injury, it also has promise for other classes of people who exhibit occasional cognitive lapses such as populations with age-related memory loss and even high functioning people who inevitably make mistakes.

The name of our system is derived from the desire to provide our users with a source of computer generated opportunities from which they can learn more efficient transportation routes and correct simple errors before they become dangerous errors. When the system has determined that an especially important opportunity has made itself available, it plays a sound like a door knocking to get the user's attention. Less critical opportunities are simply displayed if the user expresses interest. We desire to support existing cognitive capacities, not replace them, by helping users to remain engaged in their transportation decisions.

The system is implemented on a cell-phone platform and differentiates itself from more familiar web-based route planning or car navigation systems in several ways. First, the system requires no explicit input from the individual: it generates its path planning advice and destination predictions in an unsupervised manner entirely from past observations of the user. Second *OK* is centered on the individual. As such, it travels with the user and provides value to the user across multiple transportation modalities. Finally, *OK* can detect novel and explicitly erroneous user behavior.

In this paper we present three main contributions. First we developed a system architecture which could support our goals. This architecture is described in Section 3. Secondly, we present an elegant method of circumventing the onerous task of labeling positions with place names in Section 4. Thirdly, and most importantly, we designed and implemented an inference engine that supports explicitly reasoning about destinations and detecting user errors, as described in Section 5.

Finally, although in this paper we focus on a system which assists cognitively impaired people, the techniques we present can be applied to any user-centric location-based service that would benefit from probabilistically predicted location information (e.g., just-in-time traffic information for specific routes, home climate and appliance control, or reminders for errands-of-convenience).

2 Scenario

In order to ground our system, we present a fictitious running example that will help illustrate the most important features of our system.

Eileen has a physical therapist at a nearby university campus, whom she visits on a bi-weekly basis. After one such visit, Eileen finds herself exiting the building uncertain of which way to proceed. After a few minutes of hesitation,

she reaches for her phone and invokes Opportunity Knocks. *OK* offers images of three destinations that she typically travels to after the therapist visit: her home, a grocery store, and the house of her friend Ted. Eileen selects her home and the system suggests her typical route: it provides instructions to find the nearest bus stop and tells her to wait for bus number 372.

Bus number 68 arrives first. Since this is the bus that Eileen normally takes to the *grocery store*, she accidentally boards it instead. Its route initially coincides with that of number 372; while *OK* can identify that she is on *a* bus, it is unable to detect the *identity* of the bus. It remains silent as it observes that Eileen is moving toward home in the expected manner. When the bus suddenly turns west after some time, following the bus route to the grocery store, her phone makes a knocking sound and alerts her that she should get off at the next stop. At that point, it directs her back a few hundred feet to a bus stop where she can board the next 372 bus. This time she gets on a correct bus and arrives home safely.

3 System Architecture

In order to support Eileen in the way we describe in the previous scenario, several technical pieces have to be composed. First, we describe the overall architecture of the system before discussing the individual components in detail.




Fig. 1. Architectural Diagram of Opportunity Knocks.

Figure 1 diagrams our overall system architecture. The data flow of our system starts at a sensor beacon which is carried by a user. The sensor samples the environmental context of the user and forwards this information over a secure Bluetooth connection to the cell phone. The cell phone initially acts as a network access point and again forwards the context information to a remote server over the high-speed GPRS data network. The remote server, which is running the *OK* software, uses the sensor information in conjunction with Geographic Information Systems' (GIS) databases to localize the user. When the software has sufficient confidence in the position of the user, it is then able to suggest opportunities about which the user may want to know. These opportunities are sent back to the cell phone for display through the user interface. If an urgent opportunity, such as a plan for recovering from boarding the wrong bus, is recognized, the phone proactively alerts by making a door-knock sound; otherwise

the phone remains passive with information available for reference by the user. If the user selects an opportunity, such as a route to a frequent destination, the cell phone requests supporting information from the server, which may require referencing real-time information about bus schedules.

3.1 Cell Phone

We chose a cell phone as our client hardware because of its role as a defacto standard for a portable computing device. It has inherent value that is related to its primary function as a phone and for many people it is as common to carry as a wallet or a purse. As a result, it is likely to be a familiar, non-stigmatizing method of delivering assistive services. In the cell phone market there is a spectrum of products available, which spans from a traditional phone on one end, to a Personal Digital Assistant (PDA) on the other end. We opted for devices which were more like traditional phones rather than “smart-phones” because of their ubiquity, simple interface, and limited maintenance requirements. Cell phones also offer the promise of a cross-platform development environment which would enable an application written for the J2ME (Java 2 Micro-Edition) platform to work on any compliant phone.

Our system currently uses a Nokia 6600 cell phone. The Nokia 6600 phone is a GSM phone that has a wide-range of features required by our system. First it supports the J2ME Mobile Information Device Profile (MIDP) 2.0 that provides support for secure networking, serial port connection support, and the Application Management System — a push registry that enables authorized applications to be launched remotely. Some model specific features of the phone that we utilize include a high-resolution (176 x 208 pixels), high color (16-bit) screen, a digital camera, Bluetooth support, and high-speed data network capabilities. Under continuous operation our system lasts approximately 4 hours.

3.2 Sensor Beacon

The sensor beacon, which our users are required to carry, is a physically separate unit from the phone. We intend for users to place the sensor beacon in a purse, on their belt, in a backpack or on a wheelchair while transiting. In the future, it appears imminent that at least a simple GPS sensor will be incorporated into the phone itself [4] eliminating the need for a separate sensor beacon.

Currently, however, *OK* utilizes two different beacon implementations, both of which broadcast exclusively GPS information. One is a commercial package available from Socket Communications Inc. [5], shown at the top of Figure 2. This device measures 50x90x16mm, and contains a rechargeable six-hour battery.

The second is a custom-made device shown at the bottom of Figure 2 that utilizes a Bluetooth serial profile broadcaster and connects to an ATmega 128 processor. The ATmega processor functions as a communication gateway, controlling the multiplexing of several sensors, packaging of the data and sending it to the cell phone via Bluetooth. Our custom system will enable prototyping

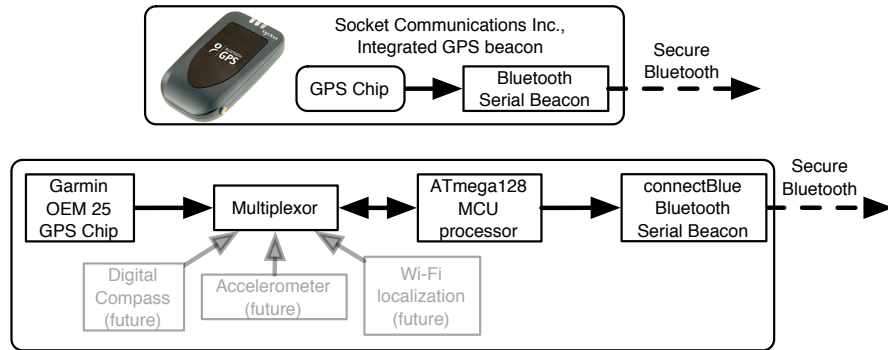


Fig. 2. Top: picture and schematic of off-the-shelf sensor beacon. Bottom: schematic of custom-built sensor beacon.

new sensors (e.g., digital compass, accelerometer, Wi-Fi localizer) in response to new research and our user studies.

4 User Interface

4.1 Concept

Based on exploratory interviews with members of our target community, we have focused on a simple user experience. When the user desires transportation assistance, she refers to her phone and observes up to four images of predicted destinations (in Section 5 we describe how this selection is made). If she would like to go to one, she selects it. If the system has observed the user going to this destination in different ways, for example by foot and by bus, it will prompt her for the method she would currently like to take. The previously observed route is then provided in text form.

The system will not present destinations to which the user has not previously traveled, but it will allow the user to select a familiar destination even if it has never observed the user getting there from the current location. In this case *OK* presents a route that is based on a real-time bus route planning service provided by the local transit authority (e.g., [6, 7]).

Notably in the course of this interaction the user did not have to provide any information about where she was and only a very small amount of information about where she wanted to go, yet the system was still able to route her effectively.

4.2 Position to Place

Our system interfaces with the user by suggesting destinations that it has high confidence she is heading toward, and then routing her to that destination. It

would be insufficient to present destinations as GPS latitude/longitude positions, and infeasible to require the user to enter a description for every interesting position on a cell-phone keypad.

Ideally, we would like to produce place descriptions automatically. This, however, is recognized as a difficult open problem [8, 9]. When attempting to create a meaningful label for a place, it is clear that the purpose of the labeling and the perspective of the labeler quickly dominate the proposed ontology. Should a description of place focus on demographics, land use, administrative use, functional use, or personal memories of the place? What happens when multiple ontologies define a region in different ways, or don't even separate the region in the same way? And which way is the best way for the current user?

To solve this problem we investigated a novel use of the camera phone. Since our system is monitoring the user's location, it is able to recognize when she has spent a sufficient amount of time in a location to call it significant. When this condition is met, the camera phone alerts the user to take a picture that captures her location. In the future, whenever the system wants to refer to that location, rather than trying to call it something in particular, it simply uses the photo to identify the spot. The advantage of such a system is that the user can decide what is meaningful about their location and can take a picture which reflects that.

5 Inference Engines

In previous sections we have described the desired behavior of our system — a behavior that depends on the system being able to learn and reason about its user's transportation routines. In particular, we require the following:

- the system should learn about its user's transportation routines in an unsupervised and unobtrusive manner;
- the system should be able to predict likely destinations the user may want to go to at any given moment in time;
- the system should be able to recognize anomalous behavior; in particular, if told where the user is going (by the user who requests directions or by a care taker or job coach who specifies the destination), the system should be able to detect, as early as possible, when the user strays from one of the usual paths that lead to that destination.

Because of the inherent uncertainties about human behavior as well as the possible errors from the maps and GPS measurements, we adopt a probabilistic approach that can handle potential errors and uncertainties in a statistically sound way. Two probabilistic models have been proposed in the recent literature for describing outdoor movement routines of a user [10, 11]. We will briefly review them and point out their fundamental inadequacies with respect to the set of requirements laid out above. Then we will discuss a more comprehensive model that subsumes the other two and provides new functionality. In Section 6 we will evaluate the new model with respect to our needs.

5.1 Previous Models

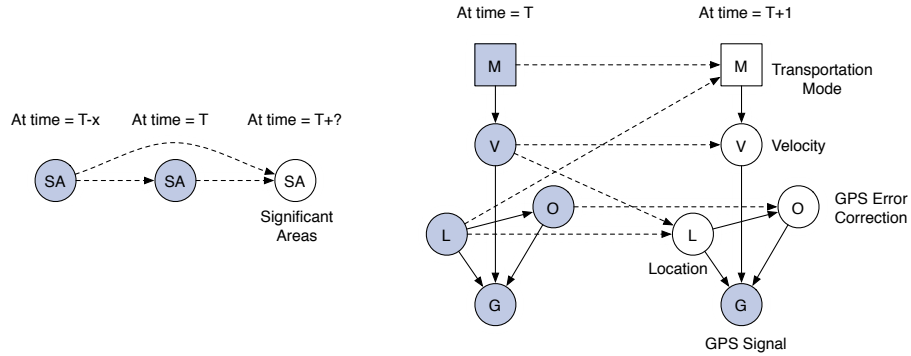


Fig. 3. On the left, a second order Markov Model (2MM) from [10] showing dependencies between observed and hidden variables. Shaded nodes are observed. All links are inter-temporal. On the right, a two-slice Dynamic Bayesian Network (2TBN) from [11] showing dependencies between observed and hidden variables. Shaded nodes are observed. Intra-temporal causal links are solid, inter-temporal links are dashed..

Ashbrook, et al. have proposed using a second order Markov model (2MM – see Figure 3-left) as a predictive tool for reasoning about likely destinations toward which a user may next be traveling [10]. The system logs continuous GPS signals, extracts places where the user seems to have stopped for a significant amount of time and then clusters them into significant locations. The optimal radius for a significant location is chosen after manual inspection of results for different radii. These results become the basis for training a second order Markov model. The authors have demonstrated that given the last two significant locations visited by the user, the system was able to generate a small and accurate set of the next most likely destinations.

In contrast to our desired behavior, this model is not able to refine estimates of the current goal using GPS information observed when moving from one significant location to another. Since significant locations might be long distances away this causes an unacceptable lag in noticing unusual behavior and significant amounts of GPS information are disregarded. This model also has no timing mechanism, so there is no way to judge when destinations will be reached or to react when too much time has passed. Finally, since the model only considers two previous locations, complex plans involving multiple significant locations cannot easily be reasoned about.

Patterson, et al. have proposed a two slice Dynamic Bayesian Network (2TBN – see Figure 3-right) for inferring a user’s transportation mode from continuous GPS signals [11]. A Dynamic Bayesian Network is an extension of Bayesian Networks which allows for time-changing variables (details in [12]). Given a

representation of the street maps, the system was able to accurately infer a user's most likely position, compensating for GPS sensor errors. The system was also able to infer locations of parking lots accessed by the user as well as bus routes and bus stop locations, all of which improved its accuracy. Finally, it estimated a user's street to street transition probabilities in an unsupervised manner and was able to use the information to further improve its accuracy.

The 2TBN could easily be adapted to detect when the user strays from a frequently traveled path. But the biggest shortcoming of this model stems from the fact that the system does not explicitly reason about the ultimate goals of a trip. Therefore, the system cannot predict the likely destinations toward which the user may be heading. Moreover, neither model, even if told a user's destination, can reason about the likely paths the user might take and, subsequently, cannot detect when the user strays from a correct path.

5.2 A New Hierarchical Model

To account for the inability of these models to meet our desiderata we have introduced a *hierarchical* Dynamic Bayesian Network model representing transportation routines [13]. The new model subsumes the capabilities of the previous models and bridges the gap between the raw sensor measurements and the abstract goal intentions of a user. A brief discussion of this model follows; refer to [13] for full technical details of the model structure, inference and training.

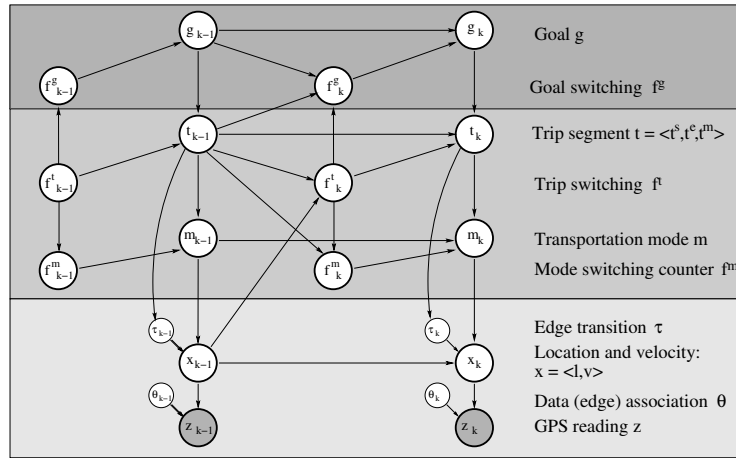


Fig. 4. Hierarchical activity model representing a person's outdoor activities. The top level estimates the current goal, the middle layer represents segments of a trip and mode of transportations and the lowest layer estimates the person's location on the street map.

Figure 4 shows the graphical structure of the new model. At the very highest level of this model, goals g_k (subscript k indicates the discrete time step) are explicitly represented as significant locations. Transitions between goals have specific probability distributions independently of the routes by which they are reached. Each goal destination influences the choice of which *trip segments* the user takes. Trip segments are sequences of motion in which the transportation mode is constant. Each trip segment t includes its start location t_k^s , end location t_k^e , and the mode of transportation, t_k^m , the person uses during the segment. Each trip segment biases the expectation over the mode of transportation and the changes in location. The mode of transportation m , in turn, determines the location and velocity distribution of the user. At the bottom level, we denote by $x_k = \langle l_k, v_k \rangle$ the location and motion velocity of the person. Edge transition τ_k indicates the next street when passing an intersection and data association θ_k “snaps” a GPS measurement onto some streets around it. The switching nodes f_k^g , f_k^t and f_k^m indicate when changes in a variable’s value can happen.

An efficient algorithm based on Rao-Blackwellised particle filters [14–16] has been developed to perform online inference for this model. At the lowest level, location tracking on the street map is done using graph-based Kalman filtering that is more efficient than the grid-based Bayesian filter and traditional particle filtering [17], used for the 2TBN model. At the highest level, the joint distribution of goals and trip segments is updated analytically using exact inference techniques. As a result, this model makes it possible to reason about high level goals (or significant locations) explicitly. The contribution of this model is that it considers not only previous significant locations visited but also the current location and the path taken so far to reason about likely destinations.

The parameters in the model are estimated in an unsupervised manner. This is a three step process. In a first pass through the data, the possible goals for a user are discovered by observing when the user stays at a location for a long time. Then in a second pass, the usual parking spots and bus stops are inferred using an Expectation-Maximization algorithm [18] similar to the learning of the 2TBN in [11]. Finally, the transition matrices at all levels are re-estimated simultaneously using a second Expectation-Maximization procedure with the full model. The learning process does not require any labeled data and therefore, requires no intervention from the user.

To detect abnormal events, the approach uses two models with different transition parameters. The first tracker assumes the user is behaving according to his personal historical trends and uses the learned hierarchical model for tracking. The second tracker assumes a background model of activities and uses an untrained prior model that accounts for general physical constraints but is not adjusted to the user’s past routines. The trackers are run in parallel, and the probability of each model given the observations is calculated. When the user is following his ordinary routine the learned hierarchical model should have a higher probability, but when the user does something unexpected the second model should become more likely.

To compute the probability of each model, we use the concept of *Bayes factors*, which are a standard tool for comparing the quality of dynamic models based on measurements [19]. The Bayes factor H_k is computed recursively as

$$\begin{aligned}
H_k &\equiv \frac{P(z_{1:k} | M_{prior})}{P(z_{1:k} | M_{learned})} \\
&= \frac{P(z_k | z_{1:k-1}, M_{prior})}{P(z_k | z_{1:k-1}, M_{learned})} \cdot \frac{P(z_{1:k-1} | M_{prior})}{P(z_{1:k-1} | M_{learned})} \\
&= \frac{P(z_k | z_{1:k-1}, M_{prior})}{P(z_k | z_{1:k-1}, M_{learned})} \cdot H_{k-1},
\end{aligned} \tag{1}$$

where $P(z_k | z_{1:k-1}, M_{prior})$ and $P(z_k | z_{1:k-1}, M_{learned})$ are the likelihoods of the observation z_k given the untrained and the hierarchical model (the likelihoods are extracted as a side-product of tracking). From a Bayes factor, we can compute the probability of abnormal behavior:

$$\begin{aligned}
P_k(\text{Abnormal}) &\equiv P_k(\text{correct model} = M_{prior}) \\
&= \frac{P(z_{1:k} | M_{prior})}{P(z_{1:k} | M_{prior}) + P(z_{1:k} | M_{learned})} = \frac{H_k}{H_k + 1}
\end{aligned} \tag{2}$$

The last step follows directly from (1).

5.3 Errors versus Novel Behavior

The above approach can detect unexpected events, but cannot distinguish errors from deliberate novel behavior. An important contribution of *OK*, however, is the ability to differentiate these cases using knowledge of the user’s destination. This is possible because there are times when the system knows where the user is going, *e.g.*, if the user asks for directions to a destination, if a care-giver or job coach indicates the “correct” destination, or if the system has access to a location enabled date-book. In those situations we can *clamp* the value of the goal node in our model and reinterpret the low level observations. When the observations diverge significantly from the clamped high level predictions, the system is able to signal a possible error. Unlike in the 2TBN model, this model is capable of spotting anomalous behavior even if the user is following a well-trodden path, provided that path does not lead to the specified destination. This is what enabled us, in Section 2, to alert Eileen that she should get off bus number 68 and switch to 372, even though she takes both routes frequently.

Similarly to Equation (2), the probability of erroneous behavior given the user’s input g (*i.e.*, the true goal and/or true trip segment) is

$$P_k(\text{Error} | g) = \frac{\hat{H}_k}{\hat{H}_k + 1} \tag{3}$$

where the Bayes factor \hat{H}_k is now defined as

$$\hat{H}_k \equiv \frac{P(z_{1:k} | M_{prior})}{P(z_{1:k} | M_{learned}, g)}. \tag{4}$$

Here, $P(z_{1:k} | M_{learned}, g)$ is the likelihood given the clamped model.

In practice, when we track users for a long time, the probability of an error can grow very small and it can take too long for an observed error to cause this probability to cross the recognition threshold. To combat this lag, one could specify a floor that limits the error probability (e.g., 0.01 in our experiments) or compute the Bayes factors using the n most recent measurements.

6 Experiment

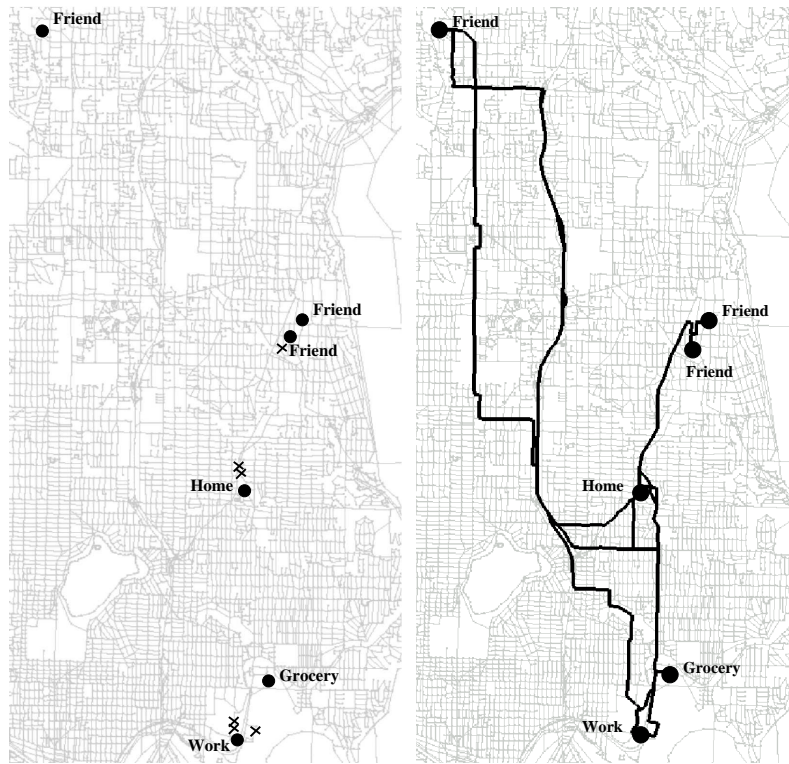


Fig. 5. The left picture shows the street map along with the goals (dots), and usual bus stops and parking lots (cross marks). The right picture shows the most likely trajectories between the goals.

6.1 Experimental Methodology

In order to test our system, we had a user carry a WAAS-enabled GPS logger with him continuously for 24 hours a day for 30 days. We then performed the

three stage training on the data without any manual labeling. The learned model correctly identifies six common goals, frequently used bus stops and parking lots, as shown in Figure 5 (left). Furthermore, our system is able to estimate the transition probabilities between goals, trip segments and streets. Using those transition matrices, we calculate the most likely trajectories on the street map between the goals, as shown in Figure 5 (right).

We tested our system using the learned model on a scenario similar to that in Section 2. The results are shown in Figures 6 and 7. These figures present a sequential panel of experimental results. The top of each panel displays a representation of the reasoning process that the inference engine is undertaking. The center portion of each panel displays what the users saw at each stage of the experiment, and the bottom portion holds a text description of the frame.

6.2 Model Clamping for Error Detection

In Figures 6 and 7, we have shown that *OK* is able to detect errors even when the user was on a frequently taken route. The system achieves this by letting the user explicitly select a destination, which we call *model clamping*. Figure 8 shows the impact of model clamping on inference results.

On the left we use the same data as in Section 6.1. In this example for the first 700 seconds both models have approximately equal belief that the user is not making an error, but when the bus took a turn that the user had never taken to get home, the probability of errors in the clamped model instantly and dramatically jumped. In contrast, the unclamped model cannot determine an error occurred because the user had taken that route to get to other destinations.

On the right is the foot experiment in which the user left his office and proceeded to walk in a direction away from the parking lot. When the destination is not specified, the tracker has a fairly steady level of confidence in the user's path (there are lots of previously observed paths from his office), but when the destination is specified, the system initially sees behavior consistent with walking toward the parking lot, and then as the user begins to turn away at time 125, the tracker begins to doubt the success of the user's intentions.

7 Related Work

There is a large body of work centered on localization and location based services, much of which originates with the pioneering work at Xerox PARC and the PARCtab [20–22] platform. It would be impossible to credit it all, but what follows is a collection that inspired and informed our research.

An important source of localization technology is research on using the known positions of radio frequency beacons to ascertain location. The RADAR system [23] presents results on indoor tracking which was improved by user motion modeling in the SmartMoveX system [24]. There are a number of outdoor wireless localization systems that track and predict movement for the purposes of providing tour guide services. A vision and discussion of this class of applications was

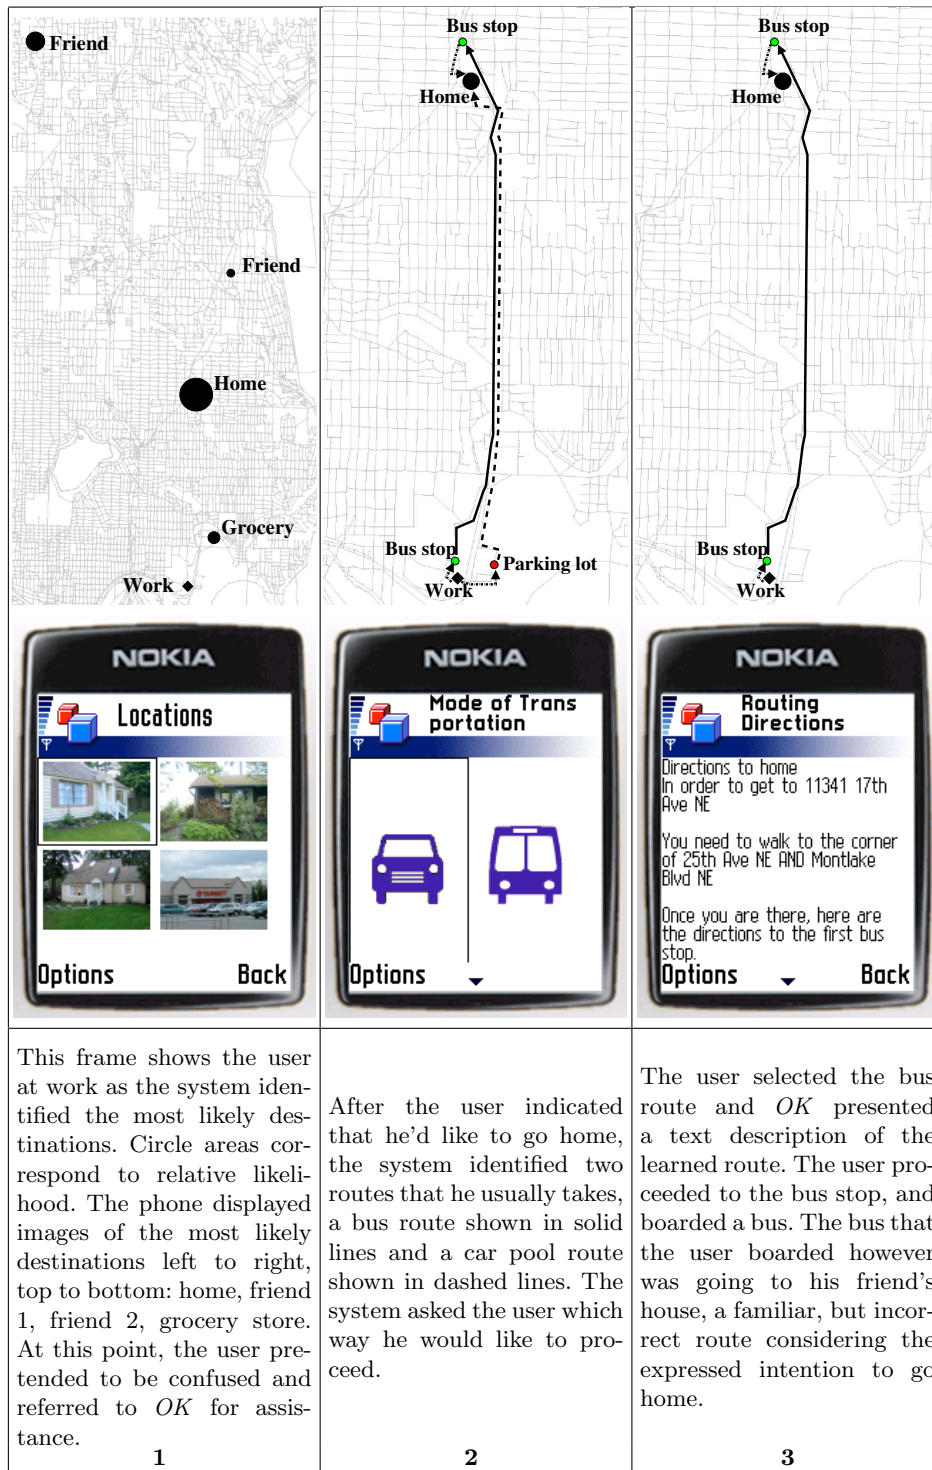


Fig. 6. Experimental Results Part 1




Fig. 7. Experimental Results Part 2




Fig. 8. Impact of Model Clamping on Error Detection. Probability of errors is shown in comparison to time in the presence and absence of a known destination. The dotted line shows the probability of errors when the user’s destination is not known. The solid line shows the probability of errors when the destination is known (and clamped). The bus experiment is on the left (error made at time 700) and the foot experiment is on the right (with a gradual error beginning at time 125).

titled Cyberguides [25], and several systems of this class have been attempted including Campus-Aware [26] and the GUIDE project [27].

The Place Lab initiative [28] is a recent proposal for making outdoor Wi-Fi localization ubiquitous through mass collaboration so that location services such as those explored by the ActiveCampus project [29, 30] are broadly available.

More generally, outdoor localization on highly resource constrained devices based on radio signals has been proposed and explored in the RightSPOT project [31] and work on abstracting and merging many different sources of localization information is being done in the Location Stack [32, 33].

Another class of related work is probabilistic plan (goal) recognition in the AI community. Bui, et al. [15] introduced the abstract hidden Markov model which uses hierarchical representations to efficiently infer a person’s goal in an indoor environment from camera information. Later, Bui [16] extended this model to include memory nodes, which enables the transfer of context information over multiple time steps. Our work goes beyond their work in that we show how to handle a challenging low-level position estimation problem, how to learn the significant transit points, and how to detect errors.

OK itself represents an evolutionary change to an existing system concept [34], which used a graphical “compass” to point a user in the direction he or she should walk on a moment-by-moment basis. Based on preliminary studies and expert feedback, we determined that the compass interface was distracting and required an unavailable resolution of localization. This earlier system did not reason about different modes of transportation, a key feature of *OK*.

Finally, as mentioned in Section 5, our work subsumes related work in user modeling and movement by Ashbrook [10] and Patterson [11].

8 Conclusion

We have presented a system called “Opportunity Knocks” (*OK*) which utilizes a rich model of user motion and behavior based on GPS sensor information to provide transportation assistance to people with mild cognitive disabilities. The primary function of the system is to route an individual from their current location to a chosen destination, but unlike existing route planning systems, it is user-centric, not vehicle-centric and requires very little user input. Instead it relies on observed user history as a basis for predicting likely destinations and identifying novel and erroneous behavior.

Our system utilizes a Bluetooth GPS beacon that talks to a cell-phone, which in turn exchanges information with a remote inference engine. The software on the remote engine runs a new hierarchical dynamic Bayesian network which is able to explicitly reason about how high-level destinations will affect many levels of transportation decisions by the user, down to the street level.

We are able to use the camera function of the phone as a method of labeling places to eliminate the need for a user to manually translate positions to places before the system can communicate about them with the user.

Finally, we have experimentally shown that this system, in conjunction with real-time transit information, has promise for effectively providing transportation assistance in the face of mild confusion, memory lapses, and inattention.

9 Future Work

We are expanding our current system in several ways to address its shortcomings. First we would like to improve power management by lowering duty-cycles, and shutting down power in response to an accelerometer incorporated into the system. Secondly, many of our directions require cardinal compass point orientation which suggest inclusion of a digital compass in the sensor beacon. Thirdly, a Wi-Fi based localizer would help us to handle indoor environments.

We are currently obtaining permission to run a formal user study with mentally retarded individuals. This will be a three stage study. First, we will conduct a user study with normal functioning individuals. Second, we will conduct a user study with mentally retarded people accompanied by a normal functioning safety monitor and, finally, we intend to conduct unassisted user studies. In particular, we will investigate a user interface that employs synthetic speech in addition to, or in place of, graphics.

To support this work we have formed an organization called *Project ACCESS* (Assisted Cognition in Community, Employment and Support Settings) [35] to help address the practical issues of conducting such a study. On the advisory board of this committee are lawyers, care-givers, parents and their children with mental retardation; all of whom are assisting in navigating the social and privacy issues associated with a device like *OK*.

Acknowledgments

This work was funded in part by National Science Foundation grants IIS-0120307, IIS-0307906, and IIS-0093406, National Institute on Disability and Rehabilitation Research grant H133A031739, Office of Naval Research grant N00014-02-1-0932, DARPA's SDR Program (grant number NBCHC020073), and Intel Corporation.

References

1. Lawton, M.P.: Environment and Other Determinants of Well-Being in Older People. *The Gerontologist* **23** (1983) 349–357 NLM:0375327;PMID:6352420.
2. Lawton, M.P.: Aging and Performance of Home Tasks. *Human Factors* **32** (1990) 527–536 NLM:0374660;PMID:2074107.
3. Consolvo, S., Roessler, P., Shelton, B., LaMarcha, A., Schilit, B., Bly, S.: Technology for Care Networks of Elders. *IEEE Pervasive Computing Mobile and Ubiquitous Systems: Successful Aging* **3** (2004) 22–29 http://www.intel-research.net/Publications/Seattle/022320041335_230.pdf.
4. F.C.C.: U.S. Federal Communication Commission E-911 Website (2004) <http://www.fcc.gov/911/enhanced>.
5. Socket Communications Inc.: Socket website (2004) <http://www.socketcom.com>.
6. King County Department of Transportation: Trip planner. <http://triplanner.metrokc.gov/> (2004)
7. Washington Metro Area Transit Authority: The ride guide. <http://rideguide.wmata.com> (2004)
8. Crabtree, A., Rodden, T., Hemmings, T., Benford, S.: Finding a Place for UbiComp in the Home. In Dey, A., Schmidt, A., McCarthy, J.F., eds.: *Proc. of UBICOMP 2003*. Volume LNCS 2864., Springer-Verlag (2003) 208–226
9. Hightower, J.: From Position to Place. In: *Proc. of the 2003 Workshop on Location-Aware Computing*. (2003) 10–12 part of the 2003 Ubiquitous Computing Conf.
10. Ashbrook, D., Starner, T.: Learning Significant Locations and Predicting User Movement with GPS. In: *International Symposium on Wearable Computing*, Seattle, WA (2002)
11. Patterson, D.J., Liao, L., Fox, D., Kautz, H.: Inferring High-Level Behavior from Low-Level Sensors. In Dey, A., Schmidt, A., McCarthy, J.F., eds.: *Proc. of UBICOMP 2003*. Volume LNCS 2864., Springer-Verlag (2003) 73–89
12. Murphy, K.: *Dynamic Bayesian Networks: Representation, Inference and Learning*. PhD thesis, UC Berkeley, Computer Science Division (2002)
13. Liao, L., Fox, D., Kautz, H.: Learning and inferring Transportation Routines. In: *Proc. of the 19th National Conf. on Artificial Intelligence (AAAI)*. (2004)
14. Doucet, A., de Freitas, J., Murphy, K., Russell, S.: Rao-Blackwellised Particle Filtering for Dynamic Bayesian Networks. In: *Proc. of the Conf. on Uncertainty in Artificial Intelligence*. (2000)
15. Bui, H., Venkatesh, S., West, G.: Policy Recognition in the Abstract Hidden Markov Model. In: *Journal of Artificial Intelligence Research*. (2002)
16. Bui, H.: A General Model for Online Probabilistic Plan Recognition. In: *Proc. of the International Joint Conf. on Artificial Intelligence*. (2003)
17. Doucet, A., de Freitas, N., Gordon, N., eds.: *Sequential Monte Carlo in Practice*. Springer-Verlag, New York (2001)
18. Rabiner, L.R.: A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. In Waibel, A., Lee, K.F., eds.: *Readings in Speech Recognition*. Kaufmann, San Mateo, CA (1990) 267–296

19. West, M., Harrison, J.: Bayesian Forecasting and Dynamic Models. second edn. Springer-Verlag (1999)
20. Want, R., Schilit, B., Adams, N., Gold, R., Petersen, K., Ellis, J., Goldberg, D., Weiser, M.: The PARCTAB Ubiquitous Computing Experiment. Technical Report CSL-95-1, Xerox Palo Alto Research Center (1995)
21. Schilit, B.N., Adams, N., Gold, R., Tso, M., Want, R.: The PARCTAB Mobile Computing System. In: Proc. Fourth Workshop on Workstation Operating Systems (WWOS-IV), IEEE (1993) 34–39
22. Adams, N., Gold, R., Schilit, B.N., Tso, M., Want, R.: An Infrared Network for Mobile Computers. In: Proc. USENIX Symposium on Mobile & Location-independent Computing, USENIX Association (1993) 41–52
23. Bahl, P., Padmanabhan, V.N.: RADAR: An In-Building RF-Based User Location and Tracking System. In: INFOCOM (2). (2000) 775–784
24. Krumm, J., Williams, L., Smith, G.: SmartMoveX on a Graph - An Inexpensive Active Badge Tracker. In: Proc. of the 4th international conference on Ubiquitous Computing, Springer-Verlag (2002) 299–307
25. Abowd, G.D., Atkeson, C.G., Hong, J., Long, S., Kooper, R., Pinkerton, M.: Cyberguide: a Mobile Context-Aware Tour Guide. *Wirel. Netw.* **3** (1997) 421–433
26. Burrell, J., Gay, G.K., Kubo, K., Farina, N.: Context-Aware Computing: A Test Case. In: Proc. of the 4th international conference on Ubiquitous Computing, Springer-Verlag (2002) 1–15
27. Cheverst, K., Davies, N., Mitchell, K., Friday, A., Efstratiou, C.: Developing a Context-Aware Electronic Tourist Guide: Some Issues and Experiences. In: CHI. (2000) 17–24
28. Schilit, B., LaMarca, A., Borriello, G., Griswold, W., McDonald, D., Lazowska, E., Balachandran, A., Hong, J., Iverson, V.: Challenge: Ubiquitous Location-Aware Computing and the Place Lab Initiative. In: Proc. of WMASH 2003: The First ACM International Workshop on Wireless Mobile Applications and Services on WLAN. (2003)
29. Griswold, W.G., Boyer, R., Brown, S.W., Truong, T.M.: A Component Architecture for an Extensible, Highly Integrated Context-Aware Computing Infrastructure. In: Proc. of the 25th International Conf. on Software Engineering, IEEE Computer Society (2003) 363–372
30. Griswold, W.G., Shanahan, P., Brown, S.W., Boyer, R., Ratto, M., Shapiro, R.B., Truong, T.M.: ActiveCampus - Experiments in Community-Oriented Ubiquitous Computing. Technical Report CS2003-0750, UC San Diego (2003)
31. Krumm, J., Cermak, G., Horvitz, E.: RightSPOT: A Novel Sense of Location for Smart Personal Object. In Dey, A., Schmidt, A., McCarthy, J.F., eds.: Proc. of UBICOMP 2003. Volume LNCS 2864., Springer-Verlag (2003)
32. Hightower, J., Brumitt, B., Borriello, G.: The Location Stack: A Layered Model for Location in Ubiquitous Computing. In: Proc. of the Fourth IEEE Workshop on Mobile Computing Systems and Applications, IEEE Computer Society (2002)
33. Hightower, J., Borriello, G.: Accurate, Flexible and Practical Location Estimation for Ubiquitous Computing: A Case Study of Particle Filters. In: Proc. of UBICOMP 2004, Springer-Verlag (2004)
34. Patterson, D., Etzioni, O., Fox, D., Kautz, H.: The Activity Compass. In: Proc. of UbiCog '02: First International Workshop on Ubiquitous Computing for Cognitive Aids, Gothenberg, Sweden (2002)
35. Harniss, M.: Project ACCESS website (2004) <http://cognitivetech.washington.edu>.