

Mobile Opportunistic Planning: Methods and Models

Eric Horvitz¹, Paul Koch¹, Muru Subramani¹

¹ Microsoft Research, One Microsoft Way
Redmond, Washington USA 98052
{horvitz,paulkoch,murus}@microsoft.com

Abstract. We present a study exploring the promise of developing computational systems to support the discovery and execution of opportunistic activities in mobile settings. We introduce the challenge of mobile opportunistic planning, describe a prototype named *Mobile Commodities*, and focus on the construction and use of probabilistic user models to infer the cost of time required to execute opportunistic plans.

1 Introduction

We believe that computing systems may one day provide great value to people by continuing to identify feasible plans for achieving standing goals in an opportunistic manner—in stream with ongoing activities. We shall explore here the promise of developing methods that can make people aware of opportunities and means for achieving goals in mobile settings. The fundamental idea is straightforward: During the progression of a planned trip, we consider a set of standing goals and preconditions specified by a mobile traveler, perform a search over a space of feasible waypoints for satisfying the goals, and seek to identify and alert the traveler about options for achieving one or more standing goals at minimal cost.

We present a prototype system, named *Mobile Commodities* (MC), which performs a search over the locations of shops, points of interest, and services, and then deliberates about the time and distance added to trips that include waypoints through these locations. MC attempts to minimize the cost of acquiring a product, service, or experience, including a consideration of the cost of time required to include the goal-satisfying waypoint. The MC prototype consists of three programs, one running as a client application on Windows Mobile Pocket PC that accesses GPS information via a Bluetooth puck, the second program running as a desktop companion for assessing preferences, configuring and inspecting policies, and the third, a server-based system that engages in two-way communication with mobile devices via GPRS.

We review the challenge of mobile opportunistic planning, and discuss how distinct subproblems are addressed by different components of MC. We shall focus on the key problem of finding the time to carry out unplanned activities opportunistically, when such activities are overlaid on the execution of existing plans. We present details on the construction and evaluation of probabilistic user models to infer the context-sensitive cost of allocating time to satisfying additional goals, and describe how the models are used in MC to guide the search for opportunistic plans.

2 Opportunistic Planning Challenge

Performing background analyses to identify feasible opportunistic plans requires (1) a means for encoding background goals, (2) a method for generating feasible plans for achieving such goals, and (3) a method for evaluating the economic value of alternate plans. A critical aspect of the economic value of opportunistic plans is the context-sensitive cost of the additional time required to satisfy secondary goals. We shall focus on predictive user modeling of the cost of time in Section 3.

Figure 1 displays the main components of mobile opportunistic planning that reflect the core competencies implemented in the MC prototype. The destination analysis component ascertains the intended destination of a user in motion. Methods for identifying a driver's destination includes (1) acquiring the destination from a user, (2) using user-specified rules that identifies a destination from a set of previously encoded set destinations classified by time of day and day of week, (3) use of a location linked to a forthcoming meeting, drawn from an online calendar, and (4) the inference of a probability distribution over forthcoming destination based on a driver's partial trajectory. We have explored the use of all four methods in MC. Space limitations limit our review of probabilistic models of destinations here; we refer readers to detailed discussion in [5]. The current implementation of MC allows users to specify destinations directly, to specify destinations as a function of the time of day and day of week, or to use the locations of forthcoming meetings drawn from an electronic calendar.

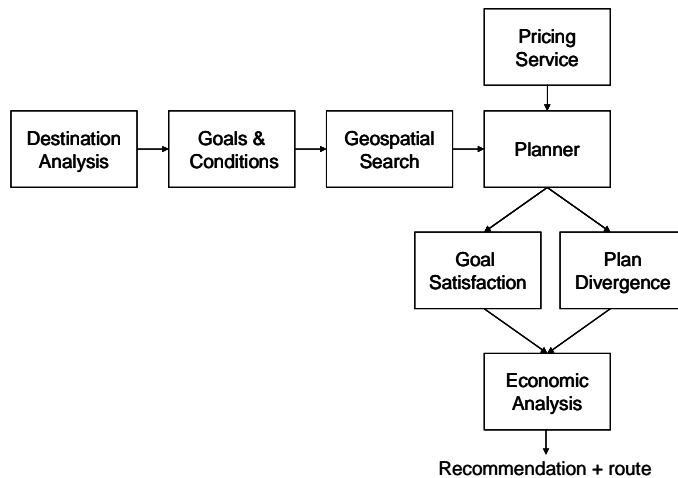


Fig. 1. Flow of analysis for mobile opportunistic planning implemented within the Mobile Commodities (MC) prototype.

A second component of MC contains a representation of standing, background goals asserted by users and the preconditions that specify when goals should be activated. We formulated a sample ontology of products and services and we seeded the system with several classes of products, services, and social goals. Products and ser-

VICES encoded in the system include such goals as obtaining groceries, gasoline, meals, haircuts, and oil changes. Users can specify specific retailers or service providers by name. Social goals allow for the specification of locations of friends and family. For each goal, we allow users to specify preconditions on a form that defines when the goal should be activated.

Users can express policies in terms of recurrent needs to acquire items that require cyclic replenishment or satiation. For example, for the goal of refueling their vehicles, users can specify a threshold amount of gasoline remaining in their car's fuel tank. When the amount of gas remaining drops below the threshold level, a background search for opportunities to seek gasoline is triggered. The frame-based specification of a policy for purchasing gasoline allows users to include the capacity of their fuel tank, the average miles per gallon, and the fuel remaining. Figure 2a shows a view of the goals and conditions specification tool for the goal of refueling. For recurrent goals such as replenishing groceries or getting haircuts, users provide a target duration between each purchase or receipt of service. Goals and preferences are specified via the MC desktop program, which synchronizes with a server that communicates with the MC mobile application.

A geospatial search component identifies locations that can satisfy active goals. MC uses the Microsoft MapPoint database to identify locations of shops and services. This subsystem takes the user's current location and target destination, computes an efficient route to the destination, then identifies candidate locations that can satisfy the active goals should they be added as waypoints on the way to the destination. For enhancing the tractability of MC's search, we limit the number of locations of opportunistic waypoints to those within a maximal tolerated distance of locations from points on the expected path that a user will take.

We will highlight the operation of MC with the example of the system computing recommendations for opportunistic gasoline purchases. The MC server has access to all of the gas stations in the Seattle area via the MapPoint database. The system also has access to a gasoline pricing service being developed at our organization. The service provides prices updated daily for all stations in major cities. Figure 2b shows the locations of gas stations in the Greater Seattle region. Figure 2c shows the overlay of prices for different qualities of fuel.

The planning component attempts to satisfy active goals and to minimize the cost of diverging from the efficient path to the primary destination. The planner also performs an economic analysis, seeking to minimize the expected cost of satisfying active goals. The planner first examines the efficient path of the user to their primary destination and considers active goals and their associated candidate locations. It then performs an exhaustive search over alternate routes that include locations that satisfy goals as waypoints on the path to the destination. For each path, it caches the path, the goals satisfied, the available prices of the desired items or services available at the waypoints, a set of directions that routes the driver from the current location through the identified locations, and the total number of miles and time required for each modified route. The economic analysis subsystem provides a context-sensitive *cost of time* for the user, and seeks to minimize the total cost to the user of diverting off of the most efficient path to the primary destination, based on the additional costs of time and of transportation.

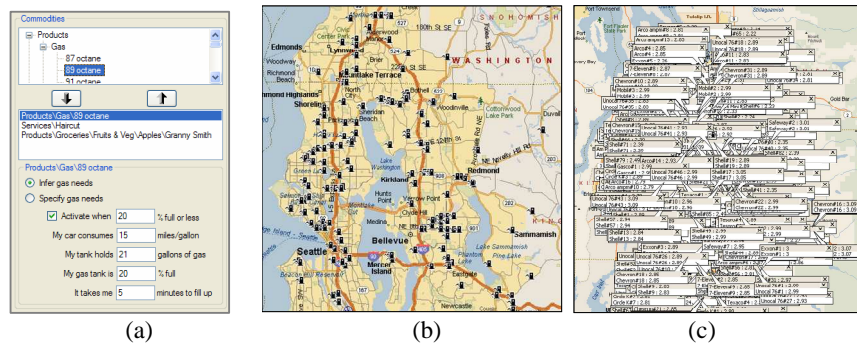


Fig. 2. a: Form within the desktop MC client that allows users to specify goals and preconditions, focused here on the gasoline purchase example; b: view of filling stations for the Seattle area; c: overlay of current prices for fuel by type of fuel at each location

3 Considering Cost of Divergence

It is not always possible to take time out, even if only a short time, to add a new destination to a trip—especially in opportunistic situations, where time may not have been allocated ahead of time for making stops. Informal interviews with potential users of MC highlighted the need for opportunistic planners to represent and reason in a sophisticated manner about the cost of time and other additional travel costs, considering the preferences of the users they support. We pursued the challenge of endowing the MC with a sense for the cost of time in different contexts.

A user’s time is indeed a precious and limited resource—in many cases the most precious resource handled by the opportunistic planner. Reasoning about the cost of time is especially important in a system designed to trade off increasing amounts of distance and time on a trip for accessing increasingly better “deals.” We focused on methods that could allow MC to consider the cost of arriving at a primary destination later than an initially intended or *target* arrival time.

In the general case, the planner needs to consider multiple properties of a destination and the overall context to assign a cost of delay associated with an unplanned stop. We explored the use of the Microsoft Outlook calendar as a means for representing and accessing properties of destinations. We view the use of an online calendar as a transitional representation for context and forthcoming events; we foresee future versions of MC relying on richer representations of patterns of daily activity. Beyond using the calendar in a standard way to represent business appointments and special social events, MC users set up recurrent appointments that capture daily patterns of activity, such as target times for arriving at work and for returning home. We gave the MC desktop application the ability to access such daily life events and more traditional appointments via an interface to Microsoft Exchange.

3.1 Assessment of Costs of Time

The computation of the cost of time in MC makes use of several assessments that are used in conjunction with probabilistic inference to generate the expected costs of time under uncertainty about context. We found that the assessments and training

required approximately a half-hour session of working with forms generated by the MC desktop client. Users first indicate on a seven-day by twenty-four hour spreadsheet-style palette, swaths of time associated with a *low*, *medium*, or *high* cost of arriving at a destination after a target arrival time. Users then directly assess a cost of time in dollars per hour for each of the three states. As we shall see, these background costs are considered by MC when no information is noted on a user's calendar. Users also assess costs of delay for contexts where a calendar is showing a forthcoming appointment. Users are asked to also consider appointments as being associated with low, medium, and high costs contexts, and assess a similar cost of delayed arrival for each of the contexts. Users can optionally enter a *tardy penalty*, a dollar value representing what users would be willing to pay to avoid being late at all. After assessment of background and meeting-centric time costs, the MC desktop application uploads a database of costs by time to the MC server. The server application uses these costs in doing cost analysis during opportunistic planning.

3.2 Learning Predictive Models for the Cost of Time

MC includes a subsystem for constructing probabilistic user models that are used at run time to infer context-sensitive costs of delays. The user models in MC infer (1) the probability that a meeting is associated with a low, medium, or high cost context, and (2) the probability that a target time drawn from a forthcoming appointment on the user's calendar is relevant. The first step in building the predictive models is that extraction of a time-sorted list of appointments from a user's online calendar.

A form displaying the list is composed for user tagging. The form contains two sets of radio selection buttons, adjacent to each appointment item. Users indicate for each meeting whether it is appropriate to consider the start time listed in the appointment as a relevant deadline, and, if so, whether the meeting should be associated with a high, medium, or low cost of being late. Given a database of tagged appointments, the system prepares a training set composed of appointments annotated with tags from the users, and also a set of properties associated with each Outlook appointment. The properties include the day and time of the appointment, meeting duration, strings from the subject and location fields, information about the organizer, the number and nature of the invitees, the response status of the user to an online invitation, whether the meeting is a recurrent meeting or not, and whether the time was marked as busy versus free on the user's calendar. We also include the role of the user, whether the user was the organizer of the meeting versus listed as a required or optional attendee by another organizer. We employ the Microsoft Active Directory service to recognize and annotate organizational relationships among the user, the organizer, and the other attendees. As an example, the system recognizes whether the organizer and attendees are peers, managers, or direct reports. Finally, we note whether the attendees, organizer, or location is "atypical" given the other meetings in the users data base; that is, we identify whether they are present in less than a predefined small fraction of all meetings in the training set.

Given the library of cases, the desktop MC application employs Bayesian structure learning to build Bayesian networks that predict relevancies and cost-of-delay functions. The system constructs models by performing heuristic search over feasible probabilistic dependency models, guided by a *Bayesian score* to rank candidate mod-

els. The Bayesian structure search method we use employs both global and local search [2,3]. For each variable, the method creates a tree containing a multinomial distribution at each leaf, exploiting the local structure search methods.

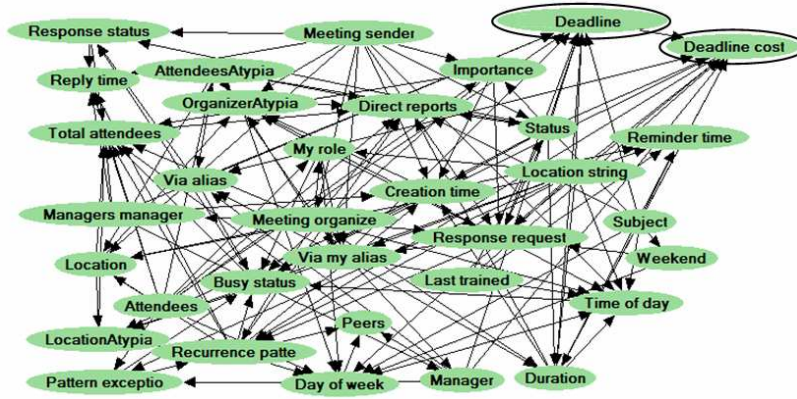


Fig. 3. Bayesian model learned from library of tagged appointments. The model predicts relevancy of target times and the cost function associated with arrival after the specified start time.

A sample predictive model for the cost of time constructed from training data from a subject testing the MC system is displayed in Figure 3. The subject tagged appointments from February 21, 2005 to March 4, 2006, a task which took the subject approximately 45 minutes. We performed a holdout cross validation, using 85 percent of the case library for training the model, and the remaining 15 percent of data to test the predictive accuracy of the models on the remaining 15 percent of holdout data. Target variables for the association of a deadline with the start of an appointment and the cost of being delayed are highlighted as circled nodes. Table 1 displays the accuracies of the inferences about the likelihood of deadlines being associated with calendar items and the probability distribution over the meeting being in the class of low, medium, or high cost of arriving after the target arrival time. The table displays significant predictive lifts over the marginal models for both of these inferences, showing the value of using the model over background statistics.

Table 1. Classification accuracy of predictive model when tested on a holdout set. The accuracy of the respective marginal models are listed beneath the accuracies of the learned models

	Relevant deadline	Cost of delayed arrival
Learned model	0.90	0.88
Marginal model	0.52	0.44

3.3 Integrating Cost of Time into Opportunistic Planning

Adding waypoints to a pre-existing trip in an opportunistic manner adds time and fuel costs to an overall trip. Reasoning about the best options for addressing background

goals requires computing the additional costs for each plan option. In the simple, non-probabilistic case, we consider the additional time and miles incurred with the divergence off the most efficient path for each opportunistic plan. We refer to this cost as the cost of divergence (CD) associated with each candidate plan.

MC makes use of the inferences from the predictive model to generate the cost of diverging from the ideal route to the destination. The analysis considers (1) the assessed background default cost of time in different situations, (2) inferences from the predictive model, based on the properties of a forthcoming appointment, about the probability distribution over the cost of arriving at progressively later times after a target start time, and (3) the likelihood that the target start time of an active appointment is relevant. We use these quantities to compute the *expected cost of divergence* (ECD) associated with any amount of time and/or distance added to the trip by considering the costs associated with the deadline relevant and not relevant situations, and combining the two situations together weighted by the likelihood of relevance and its complement.

We focus now on details of how the predictive model is used in MC to compute the cost of divergence. We shall use S^b and S^a to refer to background, non-appointment situations and special appointment contexts, respectively. We use Δd to refer to the additional miles associated with the plan and Δt to refer to the additional time of the trip due to the inclusion of one or more opportunistic waypoints to achieve standing goals, in addition to the estimate of the time required to execute the goal once at the appropriate location. We decompose Δt into the time until a target time, t^b , and the time, t^a , that falls after the target start time. For the case where the time for executing the modified travel plan leads to arrival after the target time, the cost is the sum of the background time cost incurred before the deadline, $C(S^b, t^b)$, the penalty for arriving after the target time, $C^p(S^a, t^a > 0)$, and cost of arriving at the primary destination at increasingly later times after the target time has passed, $C(S^a, t^a)$. MC also considers the additional transportation costs (fuel and wear and tear on the vehicle), C^f , associated with the divergence from the efficient path to the primary destination. This transportation cost, $C^f(\Delta d)$ is a function of the difference in distance in miles between the opportunistic plan and the primary trip, Δd , for each plan. S_i^a refers to the cost context (low, medium, and high) for an active target time.

We use $p(A|E)$ to refer to the probability that a deadline for the appointment A is relevant conditioned on evidence E , a set of properties of a forthcoming appointment S^a . The probability that a deadline is not relevant is simply the complement, $1 - p(A|E)$. For the case where the deadline is not relevant, the cost of time is just the background default cost based on the default context or situation, $C(S^a, \Delta t)$. For the case where the deadline is active, we have the background time cost incurred before the deadline, $C(S^b, t^b)$, the penalty for being tardy, $C^p(S^a, t^a > 0)$, and the growing cost of lateness, capturing the increasing cost with arriving late for the appointment, $C(S^a, t^a)$. MC is uncertain as to the cost functions associated with arriving after the target times associated with specific appointments, so the system computes an expectation by summing over the probability distribution of time cost functions inferred by the predictive model. Putting these terms together, weighting the influences of the appointment and non-appointment scenarios by the appropriate likelihoods, and adding the transportation cost, we compute the ECD for each alternate route as,

$$\begin{aligned}
ECD = C^f(\Delta d) & \\
& + (1 - p(A|E))C(S^b, \Delta t) \\
& + p(A|E)\left(\sum_i p(S_i^a|E)(C(S_i^a, t^a) + C^p(S_i^a, t^a)) + C(S^b, t^b)\right)
\end{aligned} \tag{1}$$

where $p(S_i^a/E)$ is the probability that each appointment cost context (low, medium, and high) is active. The costs of divergence described in this section are used to identify the best opportunistic plans in the MC prototype.

4 Operation of MC

When MC users get into their automobiles, the MC mobile client recognizes a Bluetooth puck in the car. A signal is sent from the mobile device to the MC server, identifying the server that the user is beginning a trip. Let us consider the example of opportunistically purchasing gasoline. When MC begins to work to satisfy the goal of identifying a best location to purchase gasoline, the system executes a cycle of analysis on the server every 10 minutes. In each cycle, the system identifies the driver's location. When planning is active, the server component accesses the user's assessments of the cost of time for the default period of time and for appointments. The system also accesses a database of the user's forthcoming appointments and examines the appointment properties. It then computes the cost of time with Equation 1.

For each cycle of opportunistic planning, the server application first computes an ideal path from the user's current location to the assumed destination, using the MapPoint route planner. As an example, when the gasoline goal is active, the application identifies all filling stations within the greater Seattle region and loads current gas prices. The system exhaustively searches through alternative routes from the current location to the destination, going through each candidate waypoint.

For each candidate route and waypoint, a divergence in miles and time for the new route, by taking the difference in miles and in time associated with the new trip and the original trip, as well as the cost assumed for the time required to stop and fill up. An overall dollar value cost of divergence is computed for each candidate trip. This cost is added to the cost of the intended purchase, computed as the price of the gas and the number of gallons required to fill the driver's tank. The system then prioritizes the alternate routes from low net cost to higher costs and sends the top five candidates to the MC mobile client, along with summary information about each candidate, including turn-by-turn directions for each. The directions divert the user off of the current path through the way point and then back to the final destination. Drivers can configure an alerting policy to limit the number of notifications during each trip.

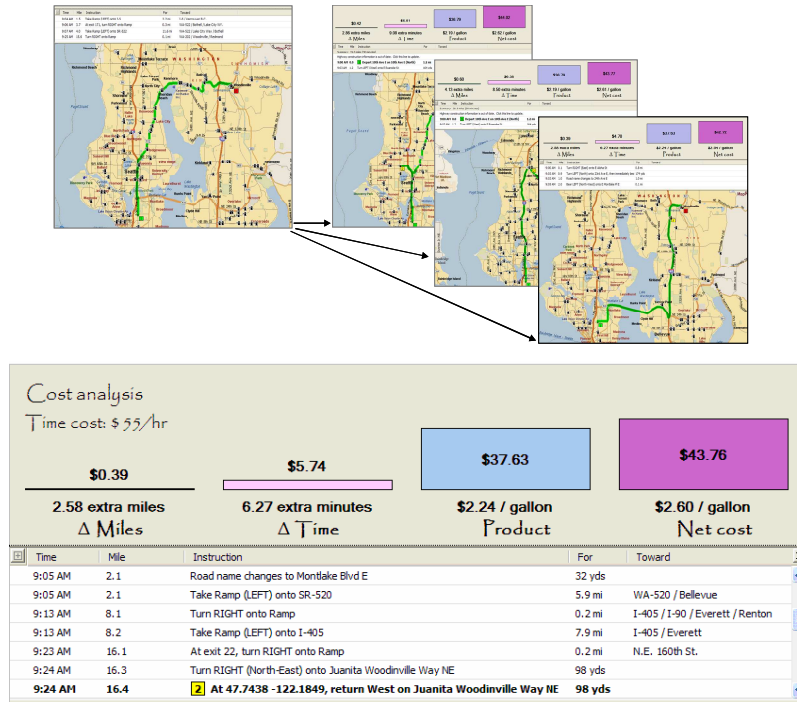


Fig. 4. Top: Portion of MC's deliberation about the best waypoint to stop for fuel. Three alternate plans of the larger search space that satisfy the goal are displayed. Bottom: Economic summary of the cost of diverging from the route to the primary destination for best plan.

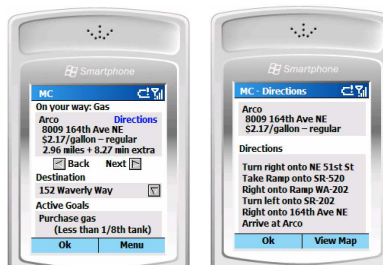


Fig. 5. Mobile views of a notification about a best candidate for opportunistic fueling.

To illustrate how MC operates, we present screens generated by a visualization utility that we created to step through the results of MC's searches. The system displays the original route, as well as candidate routes and locations for purchasing gas ordered from lowest cost to highest cost candidates. The top portion of Figure 4 shows a sequence of views displayed by the system. Each view shows a candidate opportunistic plan. A summary of the divergence analysis, including the net cost, as well as break-outs for the cost of time and for the purchase, is displayed for each candidate plan. A

summary analysis is displayed in the lower portion of Figure 4. Views rendered on a prototype MC mobile client of a notification about an opportunity, and of the directions for the path including the recommended waypoint, are displayed in Figure 5.

5 Related Work

Several prior studies are relevant to the work on MC. Patterson, *et al.* [6] examined an application that identifies when cognitively compromised people have likely strayed off of expected paths, and that works to route them back to a primary destination. Bohnenberger, *et al.* explored the recommendation of paths through a shopping mall based on representations of shoppers' interests [1]. In other related work, Horvitz, *et al.* [4] described the use of machine learning from tagged case libraries of Outlook appointments to construct probabilistic models to predict whether users will attend meetings or not and models of the cost of interruption for those meetings.

6 Summary and Directions

The Mobile Commodities project has focused on identifying challenges and opportunities for building opportunistic planning systems that work continuously to address goals encoded by people. We presented methods and models used in MC, a prototype that highlights key components and challenges with mobile opportunistic planning. On future directions, we are pursuing four extensions: (1) bundling of opportunities and simultaneous search over multiple goals, coupled with an exploration of more sophisticated planning techniques to address the combinatorial challenges; (2) integration of destination prediction services and the generalization of the methods by considering probability distributions over drivers' destinations, (3) moving to a more comprehensive cost-benefit analysis of opportunities, including the development of the ability to learn over time to recognize special offers and anomalously low prices, and (4) development of pricing systems and mechanisms that allows retailers to post standing and time-limited offers to people in a manner sensitive to preferences and context, potentially negotiating directly with peoples' opportunistic planners. We hope that our initial efforts will stimulate the user modeling community to focus more attention on challenges with mobile opportunistic planning. We see great opportunities ahead.

References

1. Bohnenberger, T., Jameson, A., Kruger, A., Butz, A. Location-aware shopping assistance: Evaluation of a decision-theoretic approach, *Mobile HCI 2002*, pp. 155-169.
2. Chickering, D.M., Heckerman, D. and Meek, C. A. Bayesian approach to learning Bayesian networks with local structure. *UAI 1997*, pp. 80-89.
3. Friedman, N. and Goldszmidt, M. Learning Bayesian networks with local structure. *UAI 1996*, pp. 252-262.
4. Horvitz, E., Koch, P., Kadie, K. & Jacobs, A. Coordinate: Probabilistic forecasting of presence and availability. *UAI 2002*, pp. 224-233.
5. Krumm, J. and Horvitz, E., Predestination: Inferring destinations from partial trajectories, *UbiComp 2006*, pp. 243-260.
6. Patterson, D.J., Liao, L., Gajos, K., Collier, M., Livic, N., Olson, K., Wang, S., Fox, D. and Kautz, H. Opportunity knocks: A system to provide cognitive assistance with transportation services. *UbiComp 2004*, pp. 433-450.