Building Personal Maps from GPS Data

LIN LIAO, DONALD J. PATTERSON, DIETER FOX, AND HENRY KAUTZ

Department of Computer Science & Engineering, University of Washington, Seattle, Washington 98195, USA

ABSTRACT: In this article we discuss an assisted cognition information technology system that can learn personal maps customized for each user and infer his daily activities and movements from raw GPS data. The system uses discriminative and generative models for different parts of this task. A discriminative relational Markov network is used to extract significant places and label them; a generative dynamic Bayesian network is used to learn transportation routines, and infer goals and potential user errors at real time. We focus on the basic structures of the models and briefly discuss the inference and learning techniques. Experiments show that our system is able to accurately extract and label places, predict the goals of a person, and recognize situations in which the user makes mistakes, such as taking a wrong bus.

KEYWORDS: personal map; GPS; Relational Markov Network (RMN); dynamic Bayesian network (DBN)

INTRODUCTION

A typical map consists of significant places and road networks within a geographic region. In this article, we present the concept of a *personal map*, which is customized based on an individual's behavior. A personal map includes personally significant places, such as home, a workplace, shopping centers, and meeting places and personally significant routes (i.e., the paths and transportation modes, such as foot, car, or bus, that the person usually uses to travel from place to place). In contrast with general maps, a personal map is customized and primarily useful for a given person. Because of the customization, it is well suited for recognizing an individual's behavior and offering detailed personalized help. For example, in this article we use a personal map to:

Ann. N.Y. Acad. Sci. 1093: 249–265 (2006). ${\ensuremath{\mathbb C}}$ 2006 New York Academy of Sciences. doi: 10.1196/annals.1382.017

An earlier version of this chapter appeared in the working notes of the *Workshop on Modeling Others from Observations* at the *International Joint Conference on Artificial Intelligence (IJCAI-2005)*, Edinburgh, Scotland, 2005.

Address for correspondence: Lin Liao, AC101 Paul G. Allen Center, Box 352350, 185 Stevens Way, Seattle WA 98195-2350. Voice: 206-543-1795; fax: 206-543-2969.

e-mail: liaolin@cs.washington.edu

- Discriminate a user's activities (Is she dining at a restaurant or visiting a friend?);
- Predict a user's future movements and transportation modes, both in the short term (Will she turn left at the next street corner? Will she get off the bus at the next bus stop?) and in terms of distant goals (Is she going to her workplace?);
- Infer when a user has *broken his ordinary routine* in a way that may indicate that he has made an error, such as failing to get off his bus at his usual stop on the way home.

We describe a system that builds personal maps automatically from raw location data collected by wearable Global Positioning System (GPS) units. Many potential applications can be built upon the system. A motivating application for this work is the development of personal guidance systems that helps cognitively impaired individuals move safely and independently throughout their community (Patterson *et al.* 2004). Other potential applications include *customized* "just in time" information services (for example, providing the user with current bus schedule information when she is likely to need it or real time traffic conditions on her future trajectories), intelligent user interface (instructing a cell phone not to ring when in a restaurant or at a meeting), and so on (Golledge 2003).

Our work spans two of the four areas that have been highlighted in the Converging Technologies conferences, information technology and cognitive science (Roco and Bainbridge, 2002). Because cognitive disabilities are usually biological in origin, the envisioned application of the work touches upon a third area, biotechnology, as well.

This chapter is focused on the fundamental techniques of learning and inference. We develop probabilistic models that bridge low-level sensor measurements (i.e., GPS data) with high-level information in the personal maps. Given raw GPS data from a user, our system first finds a user's set of significant places, then a Relational Markov Network (RMN) is constructed to recognize the activities in those places (e.g., working, visiting, and dining out); as discriminative models, RMNs often outperform their corresponding generative models (e.g., hidden Markov models) for classification tasks (Taskar *et al.* 2002). The system then uses a dynamic Bayesian network (DBN) model (Murphy 2002) for learning and inferring transportation routines between the significant places; such a generative model is well suited for online tracking and real-time user error detection.

RELATED WORK

Over the last years, estimating a person's activities has gained increased interest in the artificial intelligence (AI), robotics, and ubiquitous computing communities. One approach learns significant locations from logs of GPS measurements by determining the time a person spends at a certain location (Ashbrook and Starner 2003; Hariharan and Toyama 2004). For these locations, researchers use frequency counting to estimate the transition parameters of Markov models. Their approach then predicts the next goal based on the current and the previous goals. Our system goes beyond their work in many aspects. First, our system not only extracts places, but also recognizes activities associated with those places. Second, their models are not able to refine the goal estimates using GPS information observed when moving from one significant location to another. Furthermore, such a coarse representation does not allow the detection of potential user errors. In contrast, our hierarchical generative model is able to learn more specific motion patterns of transportation routines, which also enables us to detect user errors.

In the machine learning community, a variety of *relational probabilistic mod*els were introduced to overcome limitations of propositional probabilistic models. Relational models combine first-order logical languages with probabilistic graphical models. Intuitively, a relational probabilistic model is a *template* for propositional models such as Bayesian networks or Markov random fields (similar to how first-order logic formulas can be instantiated to propositional logic). Templates are defined over object classes through logical languages such as Horn clauses, frame systems, Structural Query Language (SQL), and full first-order logic. Given data, these templates are then *instantiated* to generate propositional models (typically Bayesian networks or Markov random fields), on which inference and learning is performed. Relational probabilistic models use high-level languages for describing systems involving complex relations and uncertainties. Because the structures and parameters are defined at the level of classes, they are shared by the instantiated networks. Parameter sharing is particularly essential for learning from sparse training data and for knowledge transfer between different people. As a popular relational probabilistic model, RMNs define the templates using SOL, a widely used query language for database systems, and the templates are instantiated into (conditional) Markov networks, which are undirected models that do not suffer the cyclicity problem and are thereby more flexible and convenient. Since their introduction, RMNs have been used successfully in a number of domains, including web page classification (Taskar et al. 2002), link prediction (Taskar et al. 2003), and information extraction (Bunescu and Mooney 2004).

In the context of probabilistic plan recognition, Bui and colleagues (2002) introduced the abstract hidden Markov model, which uses hierarchical representations to efficiently infer a person's goal in an indoor environment from camera information. Bui (2003) extended this model to include memory nodes, which enables the transfer of context information over multiple time steps. Bui and colleagues introduced efficient inference algorithms for their models using Rao-Blackwellised particle filters (RBPF). Because our model has a similar structure to theirs, we apply the inference mechanisms developed by Bui. Our work goes beyond the work of Bui's group in that we show how to learn the

parameters of the hierarchical activity model, and their domains, from data. Furthermore, our low-level estimation problem is more challenging than their indoor tracking problem.

The task of detecting abnormal events in time series data (called *novelty detection*) has been studied extensively in the data-mining community (Guralnik and Srivastava 1999), but remains an open and challenging research problem. We present the results on abnormality and error detection in location and transportation prediction using a simple and effective approach based on comparing the likelihood of a learned hierarchical model against that of a prior model.

EXTRACTING AND LABELING PLACES

In this section, we briefly discuss place extraction and activity labeling. Full technical details of the activity labeling are provided elsewhere (Liao *et al.* 2005).

Place Extraction

Similar to the work of Ashbrook and Starner (2003) and of Hariharan and Toyama (2004), our current system considers significant places to be those locations where a person typically spends extended periods of time. From the GPS data, it first looks for locations where the person stays for a given amount of time (e.g., 10 min), and then these locations are clustered to merge spatially similar points. An extension of the approach that takes into account more complex features is discussed later.

Activity Labeling

We build our activity model based on the RMN framework (Taskar *et al.* 2002). RMNs describe specific relations between objects using clique templates specified by SQL queries: each query C selects the relevant objects and their attributes, and specifies a *potential* function, or clique potential, Φ_C , on the possible values of these attributes. Intuitively, the clique potentials measure the "compatibility" between values of the attributes. Clique potentials are usually defined as a log-linear combinations of *feature* functions, that is, $\Phi_C(\mathbf{v}_C) = \exp\{\mathbf{w}^T_C \cdot \mathbf{f}_C(\mathbf{v}_C)\}$, where \mathbf{v}_C are the attributes selected in the query, $\mathbf{f}_C(I)$ is a feature vector for *C*, and \mathbf{w}^T_C is the transpose of the corresponding weight vector. For instance, a feature could be the number of different homes defined using aggregations.

To perform inference, an RMN is *unrolled* into a Markov network, in which the nodes correspond to the attributes of objects. The connections among the

nodes are built by applying the SQL templates to the data; each template C can result in several cliques, which share the same feature weights. Standard inference algorithms, such as belief propagation and Markov Chain Monte Carlo (MCMC), can be used to estimate the conditional distribution of hidden variables given all the observations.

Relational Activity Model

Because behavior patterns can be highly variable, a reliable discrimination between activities must take several sources of evidence into account. More specifically, our model defines the following templates:

- i. *Temporal* patterns: Different activities often have different temporal patterns, such as their duration or the time of day. Such local patterns are modeled by clique templates that connect each attribute with the activity label.
- ii. *Geographic* evidence: Information about the types of businesses close to a location can be extremely useful to determine a user's activity. Such information can be extracted from geographic databases like http://www.microsoft.com/mappoint. Because location information in such databases is not accurate enough, we consider such information by checking whether, for example, a restaurant is within a certain range from the location.
- iii. *Transition* relations: First-order transitions between activities can also be informative. For example, going from home to work is very common while dining out twice in a row is rare.
- iv. *Spatial* constraints: Activities at the same place are often similar. In other words, the number of different types of activities in a place is often limited.
- v. *Global* features: These are soft constraints on the activities of a person. The number of different home locations is an example of the global constraints. Such a constraint is modeled by a clique template that selects all places labeled as home and returns how many of them are different.

Inference

In our application, the task of inference is to estimate the labels of activities in the *unrolled* Markov networks. FIGURE 1 offers an example with six activities. Solid straight lines indicate the cliques generated by the templates of temporal, geographic, and transition features; bold solid curves represent spatial constraints (activity 1 and 4 are associated with the same place and so are 2 and 5); dashed curves stand for global features, which are label-specific cliques (activity 1 and 4 are both labeled as "AtHome" or "AtWork" at this moment).



FIGURE 1. An example of unrolled Markov network with six activities.

Inference in our relational activity models is complicated by the fact that the structure of the unrolled Markov network can change during inference. This is due to the fact that, in the templates of global features, the label of an object determines to which cliques it belongs. We call such cliques *label-specific cliques*. Because the label values are hidden during inference, such cliques potentially involve all the labels, which makes exact inference intractable.

We perform approximate inference using MCMC (Gilks *et al.* 1996). We first implemented MCMC using basic Gibbs sampling. Unfortunately, this technique performs poorly in our model because of the strong dependencies among labels. To make MCMC mix faster, we developed a mixture of two transition kernels: the first is a block Gibbs sampler and the second is a Metropolis sampler (see Liao *et al.* 2005 for details). The numbers of different homes and workplaces are stored in the chains as global variables. This allows us to compute the *global* features *locally* in both kernels. In order to determine which kernel to use at each step, we sample a random number *u* between 0 and 1 uniformly, and compare *u* with the given threshold γ ($\gamma = 0.5$ in our experiments).

Learning

The parameters to be learned are the weights **w** of the features that define the clique potentials. To avoid overfitting, we perform maximum a posterior (MAP) parameter estimation and impose an independent Gaussian prior with constant variance for each component of **w**. Because the objective function for MAP estimation is convex, the global optimum can be found using standard numerical optimization algorithms (Taskar *et al.* 2002). We apply the quasi-Newton methods to find the optimal weights because they have been found to be very efficient for conditional random fields (Sha and Pereira 2003). Each iteration of this technique requires the value and gradient of the objective function computed at the weights returned in the previous iteration. Elsewhere (Liao *et al.* 2005) we presented an algorithm that *simultaneously* estimates at each iteration the value and its gradient using MCMC.



FIGURE 2. Hierarchical activity model representing a person's outdoor movements.

Although parameter learning in RMNs requires manually labeled training data, parameter sharing makes it easy to transfer knowledge. For example, in our system, we can learn a *generic* model from people who have manually labeled data, and then apply the model to people who have no labeled data. Generic models in our system can perform reasonably well, as we will show in the experiments.

LEARNING AND INFERRING TRANSPORTATION ROUTINES

We estimate a person's activities using the three-level DBN model shown in FIGURE 2 (Liao *et al.* 2004). This is an hierarchical activity model representing a person's outdoor movements during everyday activities. The upper level estimates the user mode, the middle layer represents goals and trip segments, and the lowest layer is the flat model, estimating the person's location, velocity, and transportation mode. The individual nodes in such a temporal graphical model represent different parts of the state space, and the arcs indicate dependencies between the nodes (Murphy 2002). Temporal dependencies are represented by arcs connecting the two time slices k-1 and k. The highest level of the model indicates the user mode, which could be typical behavior, user error, or

deliberate novel behavior. The middle level represents the person's goal (i.e., next significant place) and trip segment (defined below). The lowest level is the *flat* model, which estimates the person's transportation mode, location and motion velocity from the GPS sensor measurements. In this section, we explain the model from bottom up.

Locations and Transportation Modes

We denote by $\chi_k = (l_k, v_k, c_k)$ the location and motion velocity of the person, and the location of the person's car (subscripts *k* indicate discrete time). We include the car location because it strongly affects whether the person can switch to the car mode. In our DBN model, locations are estimated on a graph structure representing a street map. GPS sensor measurements, z_k , are generated by the person carrying a GPS sensor. Because measurements are given in continuous *xy*-coordinates, they have to be "snapped" to an edge in the graph structure. The edge to which a specific measurement is "snapped" is estimated by the association variable θ_k . The location of the person at time k depends on his previous location, l_{k-1} , the motion velocity, v_k , and the vertex transition, τ_k . Vertex transitions τ model the decision a person makes when moving over a vertex in the graph, for example, to turn right when crossing a street intersection.

The mode of transportation can take on four different values $m_k \in \{BUS, FOOT, CAR, BUILDING\}$. Similar to (Patterson *et al.* 2003), these modes influence the motion velocity, which is picked from a Gaussian mixture model. For example, the walking mode draws velocities only from the Gaussian representing slow motion. *BUILDING* is a special mode that occurs only when the GPS signal is lost for significantly long time. Finally, the location of the car only changes when the person is in the *CAR* mode, in which the car location is set to the person's location.

An efficient algorithm based on RBPFs has been developed to perform online inference for the flat model (Doucet *et al.* 2000). In a nutshell, the RBPF samples transportation mode $m_k^{(i)}$, transportation mode switch $f_k^{(i)}$, data association $\theta_k^{(i)}$, edge transition $\tau_k^{(i)}$, and velocity $v_k^{(i)}$, then it updates the Gaussian distribution of location $I_k^{(i)}$ using a one-dimensional Kalman filter. After all components of each particle are generated, the importance weights of the particles are updated. This is done by computing the likelihood of the GPS measurement z_k , which is provided by the update innovations of the Kalman filters.

We apply expectation maximization (EM) to learn the model parameters. Before learning, the model has no preference for when a person switches mode of transportation, or which edge a person transits to when crossing a vertex on the graph. However, information about bus routes, and the fact that the car is either parked or moves with the person, already provide important constraints on mode transitions. At each iteration of EM, the location, velocity, and mode of transportation are estimated using the RBPF of the flat model. In the E-step, transition counts of a forward and a backward filtering pass through the data log are combined, based on which we update the model parameters in the M-step. In our Rao-Blackwellised model, edge transitions are counted whenever the *mean* of a Kalman filter transits the edge. The learned flat model encodes information about typical motion patterns and significant locations by edge and mode transition probabilities.

After we estimate the mode transition probabilities for each edge, we find *mode transfer locations*, that is, usual bus stops and parking lots, by looking for those locations at which the mode switching exceeds a certain threshold.

Goals and Trip Segments

A trip segment is defined by its start location, t^s_k , end location, t^e_k , and the mode of transportation, t^m_k , the person uses during the segment. For example, a trip segment models information such as: "She gets on the bus at location t^s_k and takes the bus up to location t^e_k , where she gets off the bus." In addition to transportation mode, a trip segment predicts the route on which the person gets from t^s_k to t^e_k . This route is not specified through a deterministic sequence of edges on the graph but rather through transition probabilities on the graph. These probabilities determine the prediction of the person's motion direction when crossing a vertex in the graph, as indicated by the arc from t_k to τ_k .

A goal represents the current target location of the person. Goals include the significant locations extracted using our discriminative model. The transfer between trip segments and goals is handled by the boolean switching nodes f_k^t and f_k^g , respectively.

To estimate a person's goal and trip segment, we apply the inference algorithm used for the abstract hidden Markov memory models (Bui 2003). More specifically, we use an RBPF both at the low level and at the higher levels. Each sample of the resulting particle filter contains the discrete and continuous states described in the previous section, and a joint distribution over the goals and trip segments. These additional distributions are updated using exact inference.

Because we have learned the set of goals using the discriminative model and the set of trip segments using the flat model, we only need to estimate the transition matrices at all levels: between the goals, between the trip segments given the goal, and between the adjacent streets given the trip segment.

Again, we use EM in the hierarchical model, which is similar to that in the flat model. During the E-steps, smoothing is performed by tracking the states both forward and backward in time. The M-steps update the model parameters using the frequency counts generated in the E-step. All transition parameters are smoothed using Dirichlet priors.

User Modes

To detect user errors or novel behavior, we add the variable u_k to the highest level, which indicates the user's behavior mode ε {*Normal,Novel,Erroneous*}. Different values of u_k instantiate different parameters for the lower part of the model. When user mode is typical behavior, the model is instantiated using the parameters learned from training data. When a user's behavior is *Erroneous*, the goal remains the same, but the trip segment is set to a distinguished value "unknown" and as a consequence the parameters of the flat model (i.e., transportation mode transitions and edge transitions) are switched to their *a priori* values: An "unknown" trip segment cannot provide any information for the low-level parameters. When a user's behavior is *Novel*, the goal is set to "unknown," the trip segment is set to "unknown," and the parameters of the flat model are again set to their a priori values.

To infer the distribution of u_k , we run two trackers simultaneously, and at each time their relative likelihood is used to update the distribution. The first tracker uses the hierarchical model with learned parameters and second tracker uses the flat model with a priori parameters. When a user is following her ordinary routine, the first tracker has higher likelihoods, but when the user makes error or does something novel, the second tracker becomes more likely. Unless the true goal is observed, the system cannot distinguish errors from novel behavior, so the precise ratio between the two is determined by hand selected prior probabilities. In some situations, however, the system knows where the user is going, for example if the user asks for directions to a destination, or if a caregiver indicates the "correct" destination, and thus the goal is fixed, treated as an observed, and therefore *clamped*. After we have clamped the goal, the probability of novel behavior becomes zero and the second tracker just determines the probabilities of an error.

EXPERIMENTS

To evaluate our system, we collected two sets of location data using wearable GPS units. The first data set contains location traces from a single person over a time period of 4 months. FIGURE 3 shows part of the locations contained in the data set of a single person, collected over a period of 4 months, the x-axis is 8 miles long. It includes about 400 activities at 50 different places. The second data set consists of 1 week of data from five different people. Each person has 25–30 activities and 10–15 different significant places. We extracted from the logs each instance of a subject spending more than 10 min at one place. Each instance corresponds to an activity. We then clustered the nearby activity locations into places.



FIGURE 3. Some of the locations contained in the data set of a single person.

Evaluating the RMN Model

For training and evaluation, the subjects manually labeled the data with their activities from the following set: {**AtHome, AtWork, Shopping, DiningOut, Visiting, Other**}. Then, we constructed the unrolled Markov networks using the templates described above, trained the models, and tested their accuracy. Accuracy was determined by the activities for which the most likely labeling was correct.

In practice, it is of great value to learn a *generic* activity model that can be immediately applied to new users without additional training. In the first experiment, we used the data set of multiple users and performed leave-onesubject-out cross-validation: We trained using data from four subjects, and tested on the remaining one. FIGURE 4 A shows error rates of cross-validation of the generic models and customized models using different feature sets. The average error rates are indicated by the white bars. By using all the features, the generic model achieved an error rate of 20%. Note that the global features and the spatial constraints are very useful. To gauge the impact of different habits on the results, we also performed the same evaluation using the data set of single subject. In this case, we used 1-month data for training and the other 3-month data for test, and we repeated the validation process for each month. The results are shown by the gray bars in FIGURE 4 A. In this case, the model achieved an error rate of only 7%. This experiment shows that it is possible to learn good activity models from groups of people. It also shows that



FIGURE 4. Results showing improvement in the error rates.

if the model is learned from more "similar" people, then higher accuracy can be achieved. This indicates that models can be improved by grouping people based on their activity patterns.

When estimating the weights of RMNs, a prior is imposed in order to avoid overfitting. Without additional information, a zero mean Gaussian is typically used as the prior (Taskar *et al.* 2002). Here we show that performance can also be improved by estimating the *hyper-parameters* for the means of the weights using data collected from other people. Similar to the first experiment, we want to learn a customized model for a person **A**, but this time we also have labeled data from others. We could simply ignore the others' data and use the labeled data from **A** with a zero-mean prior. Or we can first learn the weights from the other people and use that as the mean of the Gaussian prior for **A**. We evaluate the performance of the two approaches for different amounts of training data from person **A**. FIGURE 4 B shows results, zero-mean prior versus learned model as prior mean (showing the error rates over the new places only). We can see that using data from others to generate a prior boosts the accuracy significantly, especially when only small amounts of training data are available.

Using the Bayesian prior smoothly shifts from generic to customized models: On one end, when no data from the given subject is available, the approach returns the generic (prior) model; on the other end, as more labeled data become available, the model adjusts more and more to the specific patterns of the user and we get a customized model.

Evaluating the DBN Model

The learning of the generative model was done completely unsupervised without any manual labeling. FIGURE 5 shows the learned trip segments and



FIGURE 5. Learned model in the area around the workplace.

street transitions zoomed into the workplace, and the very likely transitions (probability above 0.75). Dashed lines indicate car mode, solid lines indicate taking a bus, and dashed-dotted lines indicate foot travel. In FIGURE 5 A, the goal is the workplace, whereas in FIGURE 5 B the goal is home. The model successfully discovered the most frequent trajectories for traveling from home to the workplace and vice-versa, as well as other common trips, such as to the homes of friends.

As we described, an important feature of our model is the capability to capture user errors and novel behavior using a parallel tracking approach. To demonstrate the performance of this technique, we did the following two experiments.

In the first experiment, shown in parts (A) and (B) of FIGURE 6, a user took the wrong bus home. For the first 700 s, the wrong bus route coincided with the correct one and the system believed that the user was in $\{u_k = Normal\}$ mode. But when the bus took a turn that the user had never taken to get home, the probability of errors in the clamped model dramatically jumped (FIG. 6 A). In contrast, the unclamped model cannot determine a user error because the user, while on the wrong bus route, was on a bus route consistent with other previous goals (FIG. 6 B). When the goal is unclamped, the prior ratio of typical behavior, user error and novel behavior is 3:1:2. When goal is clamped, the probabilities of novel behavior are always zero.

The second experiment was a walking experiment in which the user left his office and proceeded to walk away from his normal parking spot. When the destination was not specified, the tracker had a steady level of confidence in the user's path (FIG. 6 D), because there are lots of previously observed paths from his office. But when the goal was specified, the system initially saw behavior consistent with walking toward the parking spot, and then as the user turned away at time 125, the tracker's confidence in the user's success dropped (FIG. 6 C).



FIGURE 6. The probabilities of user mode in two experiments.

CONCLUSIONS AND FUTURE WORK

In this chapter we have described a system that can build personal maps automatically from GPS sensors. More specifically, the system is able to: recognize significant locations of a user and activities associated with those places, infer transportation modes and goals, and detect user errors or novel behavior. The system uses an RMN for place classification and a hierarchical DBN for online tracking and error detection. This technique has been used as the basis for both experimentation and for real context-aware applications including an automated transportation routing system that ensures the efficiency, safety, and independence of individuals with mild cognitive disabilities (see Patterson *et al.* 2004).

In our future work we plan to improve the place extraction. The current approach relies only on measuring the time periods a person stays at each place and uses a fixed threshold to distinguish significant places from insignificant ones. However, it is hard to find a fixed threshold that works for all significant places. If we set the threshold too big (say 10 min, as in our experiments), some places could be missed (e.g., places a user stops by to get coffee or pick up his kids); if we set the value too small (e.g., 1 min), some trivial places (such as traffic lights) may be considered significant. Therefore, to extract more places accurately, we will take into account more features besides stay duration. For example, transportation mode is a very useful indicator: If a user switches to *foot* at some place during a *car* trip, that place is likely to be significant. Because transportation mode itself has to be inferred, we must design a model that considers all these uncertainties comprehensively. In order to do that, we plan to extend the existing relational probabilistic languages so that we can model complex relations and still perform efficient inference and learning.

The application of this work to personal guidance systems for persons with cognitive disabilities is only one example of what can be called an "assisted cognition system" (Kautz *et al.* 2003). Assisted cognition systems, in general, combine artificial intelligence and ubiquitous sensing technology to create systems that observe and understand human behavior, infer and predict a person's needs and difficulties, and provide proactive guidance, prompts, and other forms of help. We believe that systems that provide such cognitively-aware assistance will play a vital role in the future of healthcare, and enable many individuals to lead active and productive lives (Turkle 2003; Wolbring and Golledge 2003).

ACKNOWLEDGMENTS

This research is supported in part by NSF under grant number IIS-0433637 and SRI International subcontract 03–000225 under DARPA's CALO project. We also thank anonymous reviewers for their helpful comments.

REFERENCES

- ASHBROOK, D., and T. STARNER. 2003. Using GPS to learn significant locations and predict movement across multiple users. In *Personal and Ubiquitous Computing*. Available online at http://www.springerlink.com/content/ t41dukuu8p2ulek9/fulltext.html.
- BUI, H.H. 2003. A general model for online probabilistic plan recognition. In Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI); August 9–15, Acapulco, Mexico.
- BUI, H.H., S. VENKATESH, and G. WEST. 2002. Policy recognition in the abstract hidden Markov models. *Journal of Artificial Intelligence Research*, 17. Available online at http://www.jair.org/.
- BUNESCU, R., and R.J. MOONEY. 2004. Collective information extraction with relational Markov networks. In *Proceedings of the 42nd Annual Meeting of the Association* for Computational Linguistics (ACL). East Stroudsburg, Pennsylvania: Association for Computational Linguistics.

- DOUCET, A., N. DE FREITAS, K. MURPHY, and S. RUSSELL. 2000. Rao-blackwellised particle filtering for dynamic Bayesian networks. In *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence (UAI2000)*, June 30- July 3, Stanford California: Stanford University. Available online at http://www.auai.org/.
- GILKS, W.R., S. RICHARDSON, and D.J. SPIEGELHALTER. 1996. *Markov Chain Monte Carlo in Practice*. Boca Raton, FL: Chapman and Hall/CRC.
- GOLLEDGE, R. 2003. Spatial cognition and converging technologies. Pp. 122–40 in M.C. Roco and W.S. Bainbridge (eds.), *Converging Technologies for Improving Human Performance*. Dordrecht, Netherlands: Kluwer.
- GURALNIK, V., and J. SRIVASTAVA. 1999. Event detection from time series data. Pp. 33–42 in 5th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Diego, California, August 15–18.
- HARIHARAN, R., and K. TOYAMA. 2004. Project Lachesis: parsing and modeling location histories. In M.J. Egenhofer, C. Freksa, and H.J. Miller (eds.), *Proceedings of the Third International Conference on Geographic Information Science*, Adelphi, Maryland, October 20–23. Available online at http://research.microsoft.com/~toyama/hariharan_toyama_GIScience2004.doc.
- KAUTZ, H., O. ETZIONI, D. FOX, and D. WELD. 2003. Foundations of Assisted Cognition Systems. University of Washington Department of Computer Science & Engineering Technical Report. Available online at http://www.cs.washington.edu/ homes/kautz/papers/ac03tr.pdf.
- LIAO, L., D. FOX, and H. KAUTZ. 2004. Learning and inferring transportation routines. In D.L. McGuinness and G. Ferguson (eds.), *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, San Jose, California, July 25–29. Available online at http://www.cs.washington.edu/homes/kautz/papers/gps-tracking.pdf.
- LIAO, L., D. FOX, and H. KAUTZ. 2005. Location-based activity recognition using relational Markov networks. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, Edinburgh, Scotland, July 30 – August 5. Available online at http://www.cs.washington.edu/homes/kautz/papers/ijcai05rmn-liao-fox-kautz.pdf.
- MURPHY, K. 2002. *Dynamic Bayesian Networks: Representation, Inference and Learning*. Ph.D. thesis, Berkeley: University of California.
- PATTERSON, D.J., L. LIAO, D. FOX, and H. KAUTZ. 2003. Inferring high-level behavior from low-level sensors. In *International Conference on Ubiquitous Computing (UbiComp)*. Available online at http://www.cs.washington.edu/homes/ kautz/papers/High-Level-140.pdf.
- PATTERSON, D.J., L. LIAO, K. GAJOS, M. COLLIER, N. LIVIC, K. OLSON, S. WANG, D. FOX, and H. KAUTZ. 2004. Opportunity knocks: a system to provide cognitive assistance with transportation services. Pp. 433–350 in I. Siio, N. Davies, and E. Mynatt (eds.), *Proceedings of UBICOMP 2004: The Sixth International Conference on Ubiquitous Computing*. Berlin: Springer. Available online at http://www.cs.washington.edu/homes/kautz/papers/UBICOMP2004Patterson.pdf.
- ROCO, M.C., and W.S. BAINBRIDGE, eds. 2002. Converging Technologies for Improving Human Performance: Nanotechnology, Biotechnology, Information Technology and Cognitive Science. Arlington, Virginia: National Science Foundation. Available online at http://www.wtec.org/ConvergingTechnologies/Report/NBIC_report.pdf.
- SHA, F., and F. PEREIRA. 2003. Shallow parsing with conditional random fields. In Proceedings of the Conference of the North American Chapter of the

Association for Computational Linguistics on Human Language Technology, Edmonton, Canada, May 31.

- TASKAR, B., P. ABBEEL, and D. KOLLER. 2002. Discriminative probabilistic models for relational data. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*. Available online at http://ai.stanford.edu/~koller/ Papers/Taskar+al:UAI02.pdf.
- TASKAR, B., M. WONG, P. ABBEEL, and D. KOLLER. 2003. Link prediction in relational data. In *Advances in Neural Information Processing Systems (NIPS)*. Available online at http://ai.stanford.edu/~koller/Papers/Taskar+al:NIPS03b.pdf.
- TURKLE, S. 2003. Sociable technologies: enhancing human performance when the computer is not a tool but a companion. Pp. 150–158 in M.C. Roco, and W.S. Bainbridge (eds.), *Converging Technologies for Improving Human Performance*. Dordrecht, Netherlands: Kluwer.
- WOLBRING, G., and R. GOLLEDGE. 2003. Improving quality of life of disabled people using converging technologies. Pp. 270–273 in M.C. Roco, and W.S. Bainbridge (eds.), *Converging Technologies for Improving Human Performance*. Dordrecht, Netherlands: Kluwer.