# Foundations of Assisted Cognition Systems

Henry Kautz, Oren Etzioni,
Dieter Fox, Dan Weld
Department of Computer Science & Engineering
University of Washington
Seattle, WA 98195 and Lokendra Shastri
International Computer Science Institute
Berkeley, CA 94707

Technical Report CSE-02-AC-01

## 1 Introduction

The Assisted Cognition Project at the University of Washington develops novel representation and reasoning techniques in order to dramatically advance the capacity of ubiquitous computing environments to augment and enhance human capabilities, with a particular emphasis on increasing the independence of people suffering from cognitive limitations.

Assisted Cognition systems (*i*) sense aspects of an individual's location and environment, both outdoors and at home, relying on a wide range of sensors such as global positioning systems (GPS), active badges, motion detectors, and other ubiquitous computing infrastructure; (*ii*) learn to interpret patterns of everyday behavior, and to recognize user errors, confusion, and distress, using techniques from state estimation, plan recognition, and machine learning; and (*iii*) offer proactive help at appropriate times to users through prompts, warnings, and other kinds of interventions.

This overview focuses on the key problems in computer science and engineering that lay the technical foundations for Assisted Cognition. After briefly describing the broad, long-term benefits to society from work in this area, we will define the scientific challenges we address, review relevant previous results, and discuss our specific current and future research.

### 1.1 Motivation and Broader Impact

Errors in memory and problem solving have many causes, ranging from situation-specific information overload to brain injury. The mechanisms of the mind are such that even unimpaired individuals exhibit systematic cognitive errors [133]. The basic technology we are developing will be broadly applicable to many settings. However, the single most pressing social motivation for Assisted Cognition comes from the need to promote the well-being and independence of people suffering from cognitive limitations due to aging and Alzheimer's disease.

The US is facing an epidemic of Alzheimer's disease [150]. Today, approximately four million Americans suffer from Alzheimer's disease; by 2050, the number is expected to rise to 15 million people, out of 80 million people worldwide [111]. For many sufferers today high quality, personalized care is unaffordable [127], despite the fact that a conservative estimate of the cost of the disease to business and the government alone is over $100 billion a year [3]. In the coming decades, the demographic shift toward an elderly population will cause severe shortages of caretakers, regardless of whatever resources are made available.

Assisted Cognition systems will have significant positive impact on this crisis by extending the time that individuals with the disease can live at home, and relieving some of the stress suffered by family caregivers — stress that takes a severe medical and financial toll [97].

In addition to this general benefit to society and the elderly and disabled communities, our project promotes teaching, training, and understanding among researchers and practitioners in computer science, nursing, and gerontology — fields which, to this point, have rarely interacted. Additional educational benefits include the support of 6 PhD graduate students. We currently offer a graduate seminar on Assisted Cognition that is attended by computer science, medical, and nursing students and faculty, and plan to offer a mixed undergraduate / graduate course on the topic in the 2003-2004 academic year.

### 1.2 Example Applications

We are designing and prototyping several Assisted Cognition systems in order to concretely explore the technical challenges in AI and ubicomp this search area presents.

- The ACTIVITY COMPASS is a proactive personal guidance system that learns the typical patterns of its user's movements throughout the day, predicts when the user is likely to be lost or off schedule (for example, about to miss a bus), and helps the user become reoriented. The system has a novel user interface based on the metaphor of a traditional navigation compass.

- The ADL MONITOR is embedded in a home and tracks the user's performance of activities of daily living (ADL's) such as meal preparation, socializing, and personal grooming, and detects user errors and abnormal patterns of behavior.

- The ADL PROMPTER works in conjunction with the monitor to help guide the user to successfully complete multi-step tasks (such as cooking) that would be difficult or impossible for the user to perform without assistance.

Our choice and design of these and other prototype applications is performed in conjunction with our collaborators from the UW schools of Medicine and Nursing, the Alzheimer's Disease Research Center (ADRC), and the private organizations Elite Care and CareWheels.

## 2 Research Challenges

Our current research addresses the fundamental scientific challenges of developing algorithms and representations that connect low-level sensor data to high-level models of movement, behaviors, intentions, disabilities, and interventions. At every level, contextual information from higher levels must be able to feed back to lower-level reasoning processes. To reach this goal we are developing a series of increasingly sophisticated statistical models, extending the dynamic Bayesian models representations, which we will use to implement a well-founded theory of cognitive performance and errors. We will develop a decision-theoretic approach to interventions that takes into account the changing costs and benefits of actions over time.

### 2.1 Location Estimation

Location estimation is the task of extracting information about a person's location from a sequence of noisy sensor readings. In contrast to state-of-the-art location systems, location estimation for Assisted Cognition *cannot* be treated in isolation from higher level activity and goal estimation. For example, knowing that a person is on her way to the grocery store can help to predict her motion path, thereby increasing the accuracy of location estimation. The hierarchical probabilistic models described in the next section will allow us to model and utilize the dependencies between low-level location estimation and high-level abstract reasoning. However, such an integration requires us to bridge the gap between continuous sensor data and discrete, high-level representations. We will develop techniques that can bridge this gap by clustering a user's data into discrete *activity segments*, which are path segments during which a user shows highly predictable behavior. Such segments can be learned for both indoor and outdoor environments using expectation maximization (EM) along with location data and an annotated map of the environment.

Another challenge in location estimation is the fact that users frequently switch between environments with vastly different structure and sensor coverage. For example, a location system has to be able to track a person in her home, on her way to the coffee shop, and during her walk through the park. We will develop *adaptive* filters that are able to switch among or use simultaneously different map representations and different representations of uncertainty using discrete grid filters, particle filters, and / or Kalman filters whenever appropriate.

### 2.2 Behavior, Goal, and Plan Recognition

Behavior recognition explains physical motion in terms of sets of meaningful activities that extend over time. Goal recognition relates behaviors to the performer's intention to bring about some state of affairs. Behaviors and goals can be hierarchically organized into plans.

Representing and reasoning about plans and goals requires us to advance the state of the art in probabilistic dynamic models in two directions:

- Representing hierarchically structured complex activities with relative and absolute metric temporal constraints between substeps leads us to develop new algorithms for hierarchical statistical representations. We are focusing in particular on extensions to the expressive hidden semi-Markov models (HHSMM) formalism. As an example, to track a person preparing breakfast, the system needs to reason about the relative order and duration of subactivities such as setting the table, preparing tea, toasting bread, *etc.*

- Many of the entities in our domain are inherently *relational*. For example, in order to determine how much time a person spends socializing (an important measure of health in the elderly) we need to reason about relationships over sets of individuals. Actions that may be performed in different ways at different times are most naturally captured by slot/filler relationships between an instance of the action and its parameters. Even for handling locations it can be useful to explicitly represent and reason about the functional relationship between an object and a position in space (*e.g.*, "the place I left my car, wherever that is"). Our work on relational Markov models (RMMs) and dynamic probabilistic relational models (DPRMs) will allow us to represent and reason about such entities in a scalable manner; see Fig. 2(c) and Sections 3.3 and 4.3

The combination of hierarchical and relational probabilistic models provides a language and basic inference algorithms for probabilistic plan recognition. Fundamental research challenges remain, however, in defining the *content* of those models. In particular, Assisted Cognition systems need to be able to recognize when the user has made an *error* in performing (or failing to perform) a task and may require assistance.

### 2.3 Modeling Cognitive Impairments

One general strategy for error detection we are developing is based on *online model selection*, which can be used to detect surprising or unusual behaviors. Part of this approach may involve creating error models that target the particular kinds of cognitive errors the user population may make, based on an underlying theory of behavioral impairments resulting from specific neurological deficits. Such models support precise error detection and intervention strategies.

To create such error models by hand would require developing separate models for each combination of activity (*e.g.*, "taking a shower") and neurological deficit (*e.g.*, mild Alzheimer's disease). Moreover, user models would have to be fine-tuned to take into account the idiosyncrasies of an individual user's behavior. A better approach we are pursuing is to develop general neurologically-based models of cognitive deficits, and then use these models to predict likely deviations from the *normative* domain-specific models of daily activities.

### 2.4 Providing Interventions

A well-founded theory of interventions is critical for Assisted Cognition systems, because unwanted or unneeded assistance from the system can have serious negative consequences on the user, including increased loss of independence and motivation, anger, or injury. Beyond domain-level costs, perception of poor performance of a device could lead to early rejection of the system. Furthermore, the tradeoffs for interventions change over time. For example, if the user normally eats lunch at noon but has not eaten by 12:05, a prompt from the system may be intrusive; by 12:45, however, the value of prompting the user to eat may be greater than the potential cost of irritating him.

We will meet this challenge by exploiting and extending work on decision-theoretic control of action in uncertain contexts, so that the intervention strategies can take into account the probability distributions over observed behaviors, inferred mental state, and predicted reactions to system actions. Particular intervention strategies may also need to take into account the inferred neurological cause of the user error, using the models mentioned in the previous section. The same overt behavior may have distinct (and even complementary) neurocognitive causes, and the intervention required to compensate for the deficit may be quite different depending on the cause.

## 3 Prior Results

### 3.1 Location Estimation

Estimating the location of a person using sensor information along with a map of an indoor or outdoor environment is extremely similar to mobile robot localization. In robot localization, the task is to estimate the position of a mobile robot based on a map and sensor data collected by the robot [51]. Over the last years, we investigated several variants of dynamic Bayes filters in the context of robot localization and people tracking [21, 20, 49, 47, 51, 137]. In various experiments we demonstrated the advantages of rich, non-parametric representations over more restricted representations such as Gaussians used in Kalman filters [58, 59, 98]. As a consequence of this research, we introduced particle filters [41] as a powerful tool for state estimation in robotics [47, 48, 51]. The idea of particle filters is to represent probability densities by sets of samples. This representation allows particle filters to efficiently represent arbitrary, multi-modal distributions over a state space. By sampling in proportion to observation likelihood, particle filters focus computational resources on regions with high probability, where things really matter. Due to these advantages, particle filters have been applied with great success to complex, non-linear estimation problems including speech recognition [158], mobile robot navigation [51], computer vision [71], system monitoring [91, 103], and fault diagnosis [156, 33]. In addition to robot localization, we demonstrated the superior performance of particle filters for tracking the locations of people using a mobile robot equipped with a laser range-finder [136, 137].

Despite these encouraging results, our research shows that there is not a single, best representation for uncertainty in localization. The choice of representation strongly depends on the environment (map), the noise in the sensor data, and the uncertainty of the belief. For example, we recently showed that the efficiency of particle filters can be greatly improved by *adapting* the filtering process to the estimation problem [46, 86, 87]. In this project, we will build on our knowledge of sensing technologies for people tracking [60] and our experience in robot localization to develop highly robust and efficient methods for estimating the locations of people in different environments.

### 3.2 Bayesian User Modeling

Over the last decade, there has been significant research and development on *user modeling*, centering on the creation and use of predictive models of user goals and needs conditioned on observed behavior. Several investigators have highlighted the importance of statistical models in user modeling [77, 64]. Compelling applications have included the modulation of the display of information in high-stakes settings [63], the identification of the ideal nature and timing of assistance to provide users working with software [64, 62], reasoning about the next moves of a user in online games [1], the reasoning about ideal pedagogical strategies [29], the ideal timing and modality for alerting users about important messages [65], and the availability and presence of users over time. In all of these applications, researchers have noted the importance of representing uncertain relationships among the changing goals of users and sequences of events, given the typically small keyhole onto a user's plans provided by sensors.

Static and dynamic Bayesian network models have been employed as the predictive heart of several systems. As an example, in our Lumiere project [64], temporal Bayesian user models are employed to infer beliefs about a user's goals and needs as they work with Office productivity software. A background competency model, that persists across sessions, takes into consideration the changing abilities of users with the use of software features with ongoing experiences. The prior beliefs about assistance are updated with ongoing streams of observations provided by an event monitoring system that reports predicates about activity over time, including the patterns of a user's dwells between actions. If a user generates an active textual query, a Bayesian analysis of the words input by the user is gracefully folded in with the ongoing analysis of actions. Beyond beliefs about assistance, a Bayesian model is also used to predict the likelihood that a user desires assistance at different times.

Beyond reasoning about likelihoods, user modeling research has also pursued the linking of Bayesian networks to online decision analyses of the expected costs and benefits of action. For example, our Lookout system [62] learns a statistical predictive model over time to compute the likelihood that a user will desire a calendar service. The inference about the goals of the user is coupled with a decision-theoretic analysis of the context-dependent costs of taking alternate actions. The system considers a range of actions at different levels of precision. The Priorities and Notification Platform systems [65, 66] consider the costs and benefits of alerting users in different ways versus deferring messages to later, by considering a user's context and the time-dependent utility of different messaging actions.

### 3.3 User Modeling with Relational Markov Models

As part of our work on Adaptive Websites [119, 120, 5] we have developed several learning algorithms for acquiring predictive models of user behavior from observations. Our initial work [4] evaluated the suitability of previous techniques, including Naive Bayes mixture models, first-order, second-order, and propositional Markov models. While a mixture of Markov models performed relatively well, we noted that HMMs are quite limited as a representation language, because their notion of state lacks the structure that exists in most real-world domains. Furthermore, because the number of parameters of a first-order Markov model is quadratic in the number of states (and higher for higher-order models), learning Markov models is feasible only in relatively small state spaces.

Dynamic Bayesian networks (DBNs) generalize Markov models by allowing states to have internal structure [152]. If the DBN dependency structure is sufficiently sparse, it is possible to successfully learn and reason about much larger state spaces than with Markov models. However, DBNs are still limited, because they assume that all states are described by the same variables with the same dependencies. In many applications, states naturally fall into different classes, each described by a different set of variables.

For example, we might model a building as a state space where each location is a state. Classes of locations for a home might include: hallways, regions of the kitchen, areas of the living room, *etc.*.

To exploit this insight, we devised *relational Markov models (RMMs)*, a generalization of Markov models that allows states to be of different types, with a different set of variables associated with each type [7].[1] In an RMM, a set of similar states is represented by a predicate or relation, with the state's variables corresponding to the arguments of the predicate. The domain of each argument can in turn have a hierarchical structure, over which shrinkage is carried out [99]. RMMs compute the probability of a transition as a function of the source and destination predicates and their arguments; they excel in large state spaces where training data is, by necessity, sparse. Preliminary experiments show that RMMs require *four orders of magnitude* less training data than traditional Markov models [7].

### 3.4 Plan Recognition and Temporal Reasoning

Plan recognition is the task of identifying an actor's plans and goals given a partial view of that actor's behavior [135]. Plan recognition was initially studied in the context of natural language understanding [2], and was largely descriptive in nature. More principled approaches have been investigated by the authors of this report. Kautz [79] modeled plan recognition logically in a manner that allowed goals and plans to be described at various levels of abstraction. Etzioni *et al.* [94, 95, 92, 93] developed a version space algorithm for plan recognition that is provably sound and polynomial time [94, 93]. Weld *et al.* developed goal recognition algorithms using inductive logic programming [90] and version-space algebra [89, 168, 88] in the context of programming by demonstration.

An important aspect of plan recognition in realistic domains is reasoning about quantitative as well as qualitative temporal constraints, as we discuss in Sec. 4.2. Weld *et al.* [116, 151] developed efficient algorithms and data structures for handling metric constraints in the context of temporal planning. Kautz *et al.* established the first complexity results on reasoning with qualitative temporal constraints [159] and developed efficient algorithms for integrating qualitative and quantitative constraints [81].

### 3.5 Cognitive Models of Memory and Reasoning

Our research on modeling user errors and providing appropriate cognitive intervention to patients will also draw upon the past work of Shastri and his collaborators on neurally plausible models of reasoning [147, 139, 149], decision-making [165], planning [54], and episodic memory [138, 146, 143, 145]. The work on reasoning and decision-making has developed a neurally plausible computational model that demonstrates how a suitably structured neural network can encode a large body of probabilistic, relational knowledge and yet perform a broad class of predictive and explanatory inferences with extreme efficiency.

Neuropsychological [153, 27, 126] and imaging data [39, 134, 169] strongly suggests that the hippocampal system (HS) consisting of the hippocampal formation and neighboring cortical areas in the medial temporal lobe plays a critical role in episodic memory function. Moreover, HS subregions are some of the areas first affected by Alzheimer's disease [61, 57]. Shastri has developed a representationally adequate and anatomically and physiologically plausible computational model that predicts (i) the functional roles of each HS component and the cortical areas interacting with the HS, (ii) the properties of cortically expressed event schemas/frames underlying episodic memories, and (iii) memory deficits that would result from cell loss in the hippocampus and high-level cortical circuits encoding semantic knowledge. It also offers biologically grounded explanations for behavioral findings about human memory such as the fan-effect.

[1] RMMs are an example of a relational probabilistic representation, combining elements of probability and predicate calculus. Other representations of this type include probabilistic relational models [52], probabilistic logic programs [112] and stochastic logic programs [104].

### 3.6 Ubiquitous Computing

The University of Washington's CSE Department has played a leading role in the development of this interdisciplinary research community [162]. Since 1997, the University's Portolano project has created a cutting edge testbed for designing and deploying a wide range of ubiquitous computing technologies and devices (*e.g.*, touch-based communication, RF-based localization, zero-configuration sensor networks, etc.) [42]. Another relevant ubiquitous computing project at UW is Labscape [8], which has produced a sensor-based laboratory assistant for cell biologists This laboratory assistant tackles the problem of tracking progress through a biologist's experimental plan based on sensory data (for example, movement of test tubes and pipettes), in order to automate the process of keeping a laboratory notebook. This form of tracking is a simple form of plan recognition, where the user specifies a single plan in advance and is willing to cooperate with the system by explicitly disambiguating his actions in response to on-screen prompts.

### 3.7 Related Projects

We now briefly relate our work to other projects involving AI, ubiquitous computing, and caregiving.

The Aware Home Research Initiative at Georgia Tech is developing technologies to create a home that can perceive and assist its occupants [131]. Most of their work has focused on the low-level sensing infrastructure and HCI issues [38, 114, 110], although [102] uses stochastic grammars to "parse" videos.

The MIT Media Lab hosts a number of interrelated ubiquitous computing efforts, including Smart Rooms and Learning Humans [117, 118, 14, 78]. Their work, like ours, is based on probabilistic behavior models, but to date they have not considered hierarchies of plans and goals or user errors. Media lab projects specific to problems of the elderly include wearable medical monitors [129] and Memory Glasses [37]. Some of these systems will be tested at The Center for Future Health [128].

Affective Computing [121, 82] studies how a computer can recognize a person's *emotional* state based on various kinds of sensor readings. Such work could be applied to Assisted Cognition systems as a way to help determine if a patient is becoming agitated.

The Nursebot Project at CMU [10] aims at developing personal robots to help elderly people during their everyday lives. Work led by Pollack [100, 28, 122] proposes using planning, temporal reasoning, and voice prompts to help a user remember when to take medication, eat, and take care of personal hygiene.

## 4 Current Work

We now discuss the specific research that falls within the Assisted Cognition project, focusing on the fundamental contributions to computer science.

### 4.1 Location Estimation

Estimating the location of people plays a fundamental role in Assisted Cognition systems. We address the location estimation problem as an instance of Bayesian filtering, which represents the state of a dynamic system by random variables and estimates posterior distributions recursively using probabilistic models of sensors and dynamics [51, 84, 40, 91, 103]. The key problem in connecting location estimation to high-level reasoning is to bridge the gap between continuous, noisy sensor data and discrete, high-level representations. Consider, for example, the task of tracking a person's outdoor location using GPS data. Figure 1 shows raw GPS readings collected by a student over a period of three months (daytime only, manually determined modes of transportation are indicated by different colors). Each GPS reading provides an estimate of the person's $xy$-location and motion velocity. Most state of the art GPS tracking systems simply project GPS readings onto maps such as the one shown in the figure. Several researchers suggested using Kalman filters to integrate GPS sensor readings over time [155, 15, 132]. While such approaches work reasonably well for

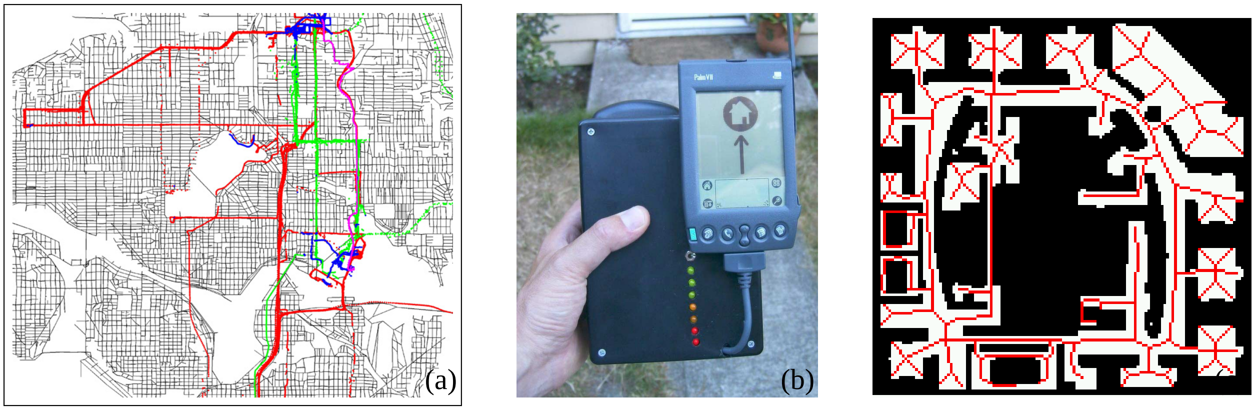


Figure 1: (a) Data log of GPS positions collected by a student over a period of three months. The map shows an area of 40 square miles around the UW campus. Colors of GPS positions indicate the different modes of transportation. The data indicates that it is possible to learn highly predictive models of daily activities. (b) An early prototype of the Activity Compass utilizing a Palm VII hand-held computer with a custom-built battery-pack and GPS receiver. Convergence of hand-held computers, cellular telephones and GPS technology promises to make adding external devices unnecessary. (c) Map of an indoor environment along with the corresponding Voronoi graph.

most applications, they are clearly not sufficient in our context. First, they can not adequately handle sensor loss due to insufficient satellite coverage in urban environments [23, 31]. Second, they do not provide any information beyond a person's location and, finally, they are not able to utilize user-specific information such as patterns of typical behavior.

Along with a person's location we will estimate her mode of transportation using GPS velocity readings, the estimated location of her bike and car, and map information about streets, bus routes, walkways, *etc.*. Initial experiments indicate that it is possible to reliably estimate transportation modes using velocity profiles learned from the data shown in Figure 1. Such data, annotated with semantic location information, will also be used to learn patterns of daily activities like "every morning around 9am, the person takes the bus to the university". To extract such information from sequences of GPS readings, we will discretize paths into *activity segments* during which the user shows highly predictable behavior. For example, the GPS data in Figure 1 shows that the person always takes the same bike trail to the university, which suggests representing the complete path as one abstract action. Clustering continuous trajectories into discrete segments can be formulated as an incomplete data problem, as shown *e.g.* by [96, 12, 106] for the case of linear Gaussian models. Hence, using expectation maximization (EM) [36], we can extract discrete actions from user data. By organizing locations and objects in an abstraction hierarchy (*e.g.*, a specific bus instance, the buses on one route, the class of all buses, means of transportation) and using our relational Markov model (RMM) and dynamic probabilistic relational model (DPRM) learning methods (Sections 3.3 and 4.3) we anticipate vastly improved learning accuracy — even with sparse data. The resulting user-specific model will be able to accurately track and predict a user's daily activities. As explained below, we will detect unusual behavior by comparing the predictive ability of the user model to a general (person-independent) model.

The approaches discussed for outdoor estimation can be transferred to activity estimation in indoor environments. Here, the map is given by an outline of the building and a person's location can be estimated using sensors such as laser range-finders and infrared or ultrasound badge systems [161, 123, 136, 60]. To combine information collected by a variety of location sensors, we currently investigate the application of Rao-Blackwellised particle filters to people tracking. Rao-Blackwellised particle filters combine Kalman filters with the representational richness of particle filters by sampling the non-linear parts of the state space and solving the linear part analytically conditioned on the samples [40, 108]. In our context, we use particle filters to address the data association problem, *i.e.* the problem of determining which sensor measurement is caused by which person. For each sample, we can assume that data association is solved, thereby allowing us to efficiently track people using Kalman filters. First experiments using data collected with six people moving through an indoor environment show very promising results [50]. In this project we will extend this work to tracking *groups* of people using the relational models described in Section 3.3.

Many indoor environments targeted in this project are equipped with only a minimum set of location sensors. Hence, we will also investigate how to estimate the location of people based on very sparse information from inaccurate location sensors such as infrared and ultrasound badge systems [161, 123, 60]. While such sensors are not accurate enough to estimate the exact location of people, they provide enough information to estimate locations on a topological level. For example, infrared sensors can detect the presence of a specific person in a room, and ultrasound sensors give rough estimates of the position of a person. In such situations of sparse sensor coverage, we propose to estimate the location of people using Voronoi graphs, a structure well known in the robotics community [22, 24]. As shown in Fig. 1(c), Voronoi graphs represent indoor environments by a skeleton of the free-space, similar to street maps for outdoor environments. Even with sparse sensor information, we are confident that it is possible to estimate a person's location on the Voronoi graph using particle filters projected onto the graph structure. An additional advantage of Voronoi graphs is the fact that they provide a natural way to discretize the continuous space of locations, thereby simplifying the generation of discrete action models (just like street maps for outdoor locations).

The different approaches discussed above will be combined into adaptive Bayes filters that are able to switch between different models (*e.g.*, outdoor map, indoor Voronoi graph, indoor blueprint) and representations (grid-based, Kalman filter, particle filter) depending on the uncertainty and the availability of sensors.

## 4.2 Representing Hierarchy and Metric Time

As noted in Sec. 2.2, linking location and movement information to an agents behaviors, plans, and goals requires reasoning about the hierarchical relationship between abstract actions and subactions, and both qualitative and quantitative metric constraints. For example, if the system determines that the user places a kettle on the stove and turns it on, it might infer that she is executing a plan for making a hot beverage, and that she will (or should) remove the kettle from the stove at a point between 5 and 10 minutes from now.

Dynamic probabilistic models that include events that occur over metric time are called semi-Markov models [69, 26, 124], and have been mainly studied in the Operations Research community. Hierarchical dynamic models have been a focus of work in AI [45, 107, 19]. In a hierarchical HMM the world is modeled by a sequence of hidden Markov models of higher and higher levels of abstraction. A node in a particular model may include a transition to a node in a less abstract model. Taking such a transition is treated as a function call that suspends the higher level model, until the called model reaches a termination state. Metric, non-exponential time is added to the model by associating a probability distribution over the duration that the system remains in a node before taking a transition, resulting in a hierarchical hidden semi-Markov model (HHSMM). Murphy [106, 107] proposed a scheme for translating HHSMM's into dynamic Bayes nets (DBN) and then applying ordinary exact and approximate DBN inference algorithms.

We are adapting Murphy's approach to implement a HHSMM to track the daily activities of residents in an assisted living community run by the Elite Care organization.[2] The Elite Care homes in Oakfield estates are equipped throughout with motion and smart badge sensors, that provide (noisy) information on the whereabouts of the residents. Fig. 2 shows the probability for the top level node of a generic, three-layer HHSMM for a resident's day. Fig. 2(a) plots left the probabilities during filtering, and Fig. 2(b) the output of fixed lag smoothing (60 minutes lag). Inference was done using Murphy's Bayes Net Toolbox [105] running the junction tree algorithm with sparse matrix representations for memory conservation [13, 70]. It can be seen that especially the smoothed output of the network provides reasonable estimates of a resident's activities. The semi-Markov structure allowed the network to distinguish different activities solely based on

[2] More information on our collaboration with Elite Care appears in Sec. 5.2 below.

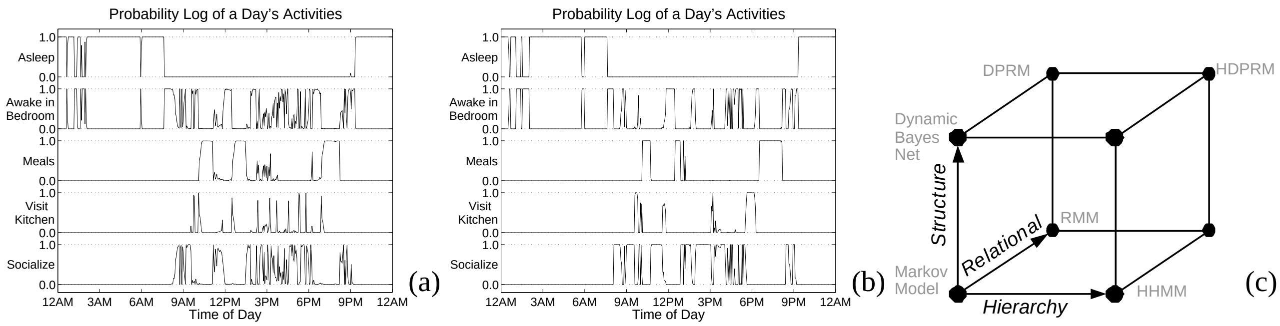


Figure 2: Example output from (a) filtering and (b) fixed lag smoothing of a three-layer HHSMM modeling the daily activities of an Elite Care resident. The three-layered structure of the network was created by hand, and transition probabilities generated for each layer individually from a log of three weeks data. Shown are the variables for sleeping, being awake in one's room, having meals, briefly going to the kitchen, and socializing (spending time in the common area). The lower two layers of the network model traveling between locations, the relationship between locations and sensor measurements, and the duration of different activities. (c) Depicts the space of possible representations of stochastic sequential processes. Our RMM representation (Section 3.3) adds relation structure to Markov Models. Our current work explores dynamic probabilistic relational models (DPRMs) and hierarchical DPRMs as described in Section 4.3.

their duration. Note that the Elite Care data contains only very inaccurate location information indicating the presence of a specific person in a rather large area.

While confirming the basic soundness of the approach, our experiments also clearly show that better algorithms and representations are needed in order to scale up to larger, more detailed models and finer-grained low-level data (the example here required about an hour of CPU time). Handling of action durations leaves many opportunities for improvement. The current model discretizes metric time, yielding an unattractive tradeoff: small units are computationally intractable, while coarse units yield poor-quality results. One approach we will consider is to use parametric representations of time distributions [18]. However, a more promising approach would be to develop variable-resolution methods in which a coarse scale time model is piecewise refined until further splitting of time intervals does not cause significant changes in the predictive distributions. Such an approach is similar to tree-based density estimation [113, 83, 101]. Variable time-resolution will also lend itself to an anytime implementation; instead of fixing a significance threshold, the system will continually refine the discretization so as to maximize accuracy given available computational resources. Finally, we also expect drastic improvements from the application of DPRMs discussed in the next section.

### 4.3 Dynamic Probabilistic Relational Models

The most powerful representation available for nondeterministic sequential phenomena is dynamic Bayesian networks [34], but DBNs are still unable to compactly represent many real-world domains. In particular, domains can contain multiple objects and classes of objects, as well as multiple kinds of relations among them; and objects and relations can appear and disappear over time. For example, an elderly cook might assemble a dinner from myriad ingredients using a variety of appliances, utensils, and operations. Capturing such a domain in a DBN would require exhaustively representing all possible objects and relations among them. This raises two problems. The first one is that the computational cost of using such a DBN would likely be prohibitive. The second is that reducing the rich structure of the domain to a very large, "flat" DBN would render it essentially incomprehensible to human beings. Our work will address these two problems by introducing an extension of DBNs that exposes the domain's relational structure, and by developing methods for efficient inference in this representation.

Recently, significant progress has been made by combining formalisms that can represent objects and relations with a principled treatment of uncertainty. Fig. 2(c) depicts the space of possible representations of stochastic sequential processes. Probabilistic relational models or PRMs [52] are an extension of Bayesian networks that allows reasoning with classes, objects and relations. A *relational schema* is a set of classes where each class $C$ is associated with sets of propositional and relational attributes, the latter called *reference slots* $\mathcal{R}(C)$. A *probabilistic relational model (PRM)* encodes a probability distribution over the set of all possible instantiations *I* of a schema [52]; the semantics of a PRM is defined by unrolling it into a large Bayesian network with one variable for each attribute of each object in the skeleton. We plan to extend PRMs to handle dynamic systems in the same way that DBNs extend Bayesian networks. We define two-time-slice PRMs and dynamic PRMs as follows.

**Definition.** *A* two-time-slice PRM (2TPRM) *for a schema $\mathcal{S}$ is defined as follows. For each class $C$ and each propositional attribute $A \in \mathcal{A}(\mathcal{C})$, we have: (1) A* set of parents $Pa(C.A) = \{Pa_1, Pa_2, ..., Pa_l\}$*, where each $Pa_i$ has the form $C.B$ or $f(C.\tau.B)$, where $\tau$ is a slot chain containing the special attribute* previous *at most once, and $f()$ is an aggregation function. (2) A* conditional probability model *for $P(C.A|Pa(C.A))$.*□

**Definition.** *A* dynamic probabilistic relational model (DPRM) *for a relational schema $\mathcal{S}$ is a pair $(M_0, M_\rightarrow)$, where $M_0$ is a PRM over $I_0$, representing the distribution $P_0$ over the initial instantiation of $\mathcal{S}$, and $M_\rightarrow$ is a 2TPRM representing the transition distribution $P(I_t|I_{t-1})$ connecting successive instantiations of $\mathcal{S}$.* □

For any $T$, the distribution over $I_0, \ldots, I_T$ is then given by $P(I_0, \ldots, I_T) = P_0(I_0) \prod_{t=1}^{T} P(I_t|I_{t-1})$. DPRMs are extended to the case where only the object skeleton for each time slice is known in the same way that PRMs are, by adding to the definition above a set of parents and conditional probability model for each relational attribute, where the parents can be in the same or the previous time slice. When the object skeleton is not known (*e.g.*, if objects can appear and disappear over time), the 2TPRM includes in addition a Boolean existence variable for each possible object, again with parents from the same or the previous time slice. As with DBNs, we distinguish between observed and unobserved attributes. In addition, we can consider an *Action* class whose domain is the set of actions that can be performed by some agent. The distribution over instantiations in a time slice can then depend on the action performed in that time slice.

Just as a PRM can be unrolled into a Bayesian network, so can a DPRM be unrolled into a DBN. However, this DBN may in general contain different variables in different time slices. In principle, we could perform inference on this DBN using particle filtering. However, the particle filter will likely perform poorly, because its state space will be huge. We plan to overcome this by adapting Rao-Blackwellisation to the relational setting [109]. Let $U_t$ be the propositional attributes of all objects and $V_t$ be their relational attributes. A Rao-Blackwellised particle is composed of sampled values for all propositional attributes of all objects, plus a probability vector for each relational attribute of each object. The vector element corresponding to $obj.R[i]$ is the probability that relation $R$ holds between $obj$ and the $i$th object of the target class, conditioned on the values of the propositional attributes in the particle.

Rao-Blackwellis-ing the relational attributes can vastly reduce the size of the state space which particle filtering needs to sample. However, if the relational skeleton contains a large number of objects and relations, storing and updating all the requisite probabilities can still become quite expensive. We believe that this can be ameliorated if context-specific independencies exist, *i.e.*, if a relational attribute is independent of some propositional attributes given assignments of values to others [17]. We can then replace the vector of probabilities with a tree structure (which may grow during the course of prediction) whose leaves represent probabilities for entire sets of objects. Space precludes a precise definition of this *abstraction tree* data structure, but preliminary experiments show that it is very effective.

After refining our inference algorithms for DPRMs, we will investigate learning methods, so we can use DPRMs to generate predictive models of human behavior. Our approach will be to extend the work of Friedman et al. [53, 56]. Given the success of our shrinkage-based approach to learning RMMs [7], we will attempt to adapt these techniques as well. Finally, we will extend DPRMs to allow hierarchy and metric time, building upon the techniques described in the previous section.

## 4.4 User Errors and Impairments

User errors must be detected by different means, depending on the knowledge available to the user model. If explicit models of user errors are available, we can incorporate such errors into the state of the user, as is often done in fault detection for dynamic systems [157, 33]. An error can then be detected whenever the probability of an error state exceeds a specific threshold. For example, if a specific person typically takes a wrong turn on the way to the bus stop, then this turn action can be explicitly incorporated as an error into the user model, thereby allowing an appropriate intervention as soon as the turn action is detected. Unfortunately, in many cases such explicit error models do not exist since it is impossible to predict all errors a user might make. Another common approach in dynamic systems is to monitor the residuals of observations, thereby testing the appropriateness of the underlying model assumptions [11, 166].

We propose an alternative, more capable approach to overcome the limitations of these methods. Our technique is based on *online model selection*, which aims at identifying the model that is best suited to explain the observed data [166, 125]. The quality of a model is given by its predictive performance, *i.e.* the likelihood of the observed data given the model. Multiple models can be compared using Bayes factors [55, 166], which then yield the ratio of posterior model probabilities. To apply model selection in our context, we will generate generic *and* user-specific models of activities. Both models are able to track a user's activities, but the specific one is tuned towards the typical actions of one particular user. The specific model additionally contains all errors that are typical for the user. The idea is that as long as the user performs her usual activities, the tuned model will be much better in predicting these activities. Typical user errors can also be detected by this specific model. Surprising actions, *i.e.* potential errors, however, are not well predicted by the specific model, in which case the generic model receives higher probability. For example, if a person exits the bus every morning at the same bus stop, then the specific model predicts this action with very high probability. The general model, however, predicts exiting the bus with a similar probability for all possible exits. If the person fails to exit the bus at the usual stop, then both models are able to track this action, but the general model predicts it with higher probability, thereby triggering the detection of a potential user error. The key advantage of this model-based approach is that it will be able to track even unusual behavior, *i.e.* it can generate information such as "the person seems to be lost but I still know where she is". Obviously, such an approach can provide valuable information to the user intervention module.

Let us now focus on the problem of learning and modeling user specific errors. Clearly hand-crafting a separate model for each combination of activity (*e.g.*, a shower) and neurological deficit (*e.g.*, mild Alzheimer's disease) will not scale. Instead we propose to use findings from neuropsychology and imaging to develop a general neurologically-based theory of cognitive and behavioral impairments, and then use this theory to predict likely deviations from the *normative* domain-specific models of daily activities.

Consider a user schema (model) for walking to the neighborhood drugstore. Over the duration of the walk, some component of this schema must remember that the destination of the walk is the drugstore. A second component of the schema must control and coordinate limb movements, *etc.*. If we annotate the first component as an "episodic memory operation" and the second as a "sensorimotor operation," it becomes possible to map the neurological deficit associated with, say, mild Alzheimer's disease (AD) to a loss/degradation in the first component (here a loss/degradation could mean a change in transition probabilities). The modified model can easily predict that a mild AD patient is likely to take a wrong turn or keep walking past the store.

To ascertain the underlying causes of a specific cognitive impairment one must know the cell-loss *profile* associated with the neurological deficit (a specification of the loci and rate of cell-loss the cell types primarily affected by the loss). We will carry out large-scale simulations of Shastri's neurally plausible models to examine the behavioral impact of progressive cell-loss associated with neurodegenerative disorders such as AD. These simulation studies will enable us to predict the intensity and nature of behavioral deficits resulting from different types of cell-loss and help us understand the causes of these behavioral deficits in functional and computational terms. Specifically, these simulation studies will (*i*) help in developing behavioral models of patients suffering from AD, (*ii*) help in recognizing the underlying causes of errors made by such patients, and detecting atypical behavior on the part of such patients, and (*iii*) enable Assisted Cognition systems to provide targeted and situation specific intervention that is appropriate given the underlying causes of a patient's cognitive impairment.

### 4.5 Decision Theoretic Intervention Strategies

Assistive applications pose special challenges for research and prototyping in that we assume that users are carrying out their normal activities, and that computing is primarily performing background monitoring and analysis. However, to be effective, we must develop methods that can support effective decision making about if, when, and how to come forward to intervene in an elegant manner.

Poor advice–and the poor timing of even good advice—can lead to the rapid rejection of automated services for assisting both healthy and impaired users. Prior research has noted that even subtle corrections in the timing of a service can be critical. As an example, our research on the Lookout project [62] demonstrated that adjusting the timing of an automated service could qualitatively change the value perceived in the automation. Coming forward to provide assistance at the wrong time was found to be disruptive, even when the advice was useful. Recent studies on the psychological studies of the timing of interruptions in healthy users have pursued the characterization of good and poor times for interrupting users doing different kinds of tasks [32]. Such research can be leveraged in assistive technologies.

We shall take a decision-theoretic perspective on interventions, focusing on representations and inference methods for deliberating about the costs and benefits of taking different actions at different times. We need to develop the means for computing the *expected* costs and benefits under uncertainty about the world and user state at hand.

The success of an assistive system that either provide critical information or be disruptive–or irrelevant–to impaired users hinges critically on machinery for reasoning about the expected value of taking immediate action, versus waiting for more data to better understand a situation, versus engaging in the active collection of data [67]. Also, beyond designing appropriate representations and machinery, we must capture a user's preferences about alternative services at different times and propagate such preferences into the reasoning fabric. Finally, it can be important to allow users to directly invoke services. Effective device-user interaction thus typically requires mechanisms that support an elegant mix of initiatives by the user and the system [62].

We shall explore key research challenges with both developing machinery for supporting effective decision-theoretic intervention strategies, a mix of user and computer initiatives, and methods for learning about the preferences of users. Key challenges include: (1) the assessment and representation of time-dependent utility models that capture expected costs and benefits over time for different settings, (2) the development of methods for learning a user's preferences by watching, (3) development of principles and architectures for deliberating about taking immediate action, versus waiting for more data to better understand a situation, versus engaging in the active collection of data before taking action in assistive settings, (4) design of user interface conventions linked to underlying machinery for supporting mixed-initiative interaction. We expect innovations in these four open challenge areas to lead to valuable results for assistive systems for impaired users as well as for technologies that can provide automated services in mainstream settings.

## 5 Datasets and Evaluation

We will test our core learning and reasoning algorithms on a broad suite of datasets; we already have access to the first four data sets, and the fifth should be available Q4 2003.

- Historical Global Position System (GPS) readings (containing both position and velocity information) acquired from handheld units over a period of months.
- Elite Care data indicating the movement of people, operation of doors and lights, and other low-level data from a set of high-tech assisted living residences in Portland Oregon.
- CareWheels data from dwellings of disabled people whose apartments are wired with motion sensing technology.
- Data from Intel Research Seattle, obtained from active badges, crickets, and laser range-finders.
- In August 2003, the CSE Department will move into a new building, which has been designed in many ways to support work on ubiquitous computing. In addition to the building-wide wireless network, sensors, and support for wearable computing, we are creating a specially designed lab space for capturing human activity in fine granularity using a wide variety of sensors: *e.g.*, RF, infrared and ultrasound sensors, laser range-finders, pressure-sensitive floor panels, motion-capture cameras, and microphone arrays.

In all work on ubiquitous computing, privacy is a central concern. At a minimum all of our datasets are anonymized and we are completing a full Human Subjects Review.

### 5.1 Experimental Evaluation and User Studies

We have described an array of research on problems from location estimation to plan recognition. In each case, our data sets provide the opportunity to empirically evaluate the performance of isolated system modules in a series of controlled tests. Due to space limitations, we provide an example of one such an experiment instead of a comprehensive list. Consider, for example, the GPS readings that we have collected in the context of the activity compass. Because we have annotated the data with meaningful locations and transportation modes, we plan to measure the extent to which our approach outperforms traditional Kalman filters. We will also compute the negative log-likelihood of our predictions to determine the absolute quality of our predictions.

In addition to detailed evaluation of our core technology, We also plan to test advanced versions of three prototypes (see below) by giving them to nurses and other care givers. Based on their experience, these professionals will give us feedback on how impaired individuals are likely to react to the devices. Of course, such professionals are also participating in the design of the Compass and Prompter so we are optimistic as to our chances of success in this difficult evaluation.

All systems we develop will undergo a full Human Subjects Review before deployment to ensure that they safeguard the privacy of the clients as well as their physical safety.

### 5.2 The Activity Compass, ADL Monitor and Prompter

The Activity Compass is designed to address the spatio-temporal confusion that accompanies ordinary age-related memory loss as well as early-stage Alzheimer's disease. The system has access to real-time GPS data and a historical database of several weeks of readings. The Activity Compass chooses when to alert, and outputs an arrow and an icon that instructs the user on how to reach a destination — an interface that minimizes cognitive load on the user. An important human computer interface (HCI) question we are studying is the most appropriate method for negative feedback. The least intrusive approach would simply take the user's lack of compliance with the system's suggestion as negative feedback. However, if the user's disease is such that he is sometimes unaware that the device is trying to give him advice, it may be more appropriate to require the user to explicitly indicate that he is ignoring the suggestion, for example, by tapping on the screen.

In conjunction with Elite Care and CareWheels, we are developing the ADL MONITOR and ADL PROMPTER. Deployed in a home, the ADL MONITOR tracks the user's performance of activities of daily living (ADL's), and detects user errors and abnormal patterns of behavior. Such a system needs to learn distinct patterns of what is "normal" for each client — for example, for some clients sleeping during the day is normal, while for others it is not; our approach is describe in Section 4.4.

The ADL PROMPTER works in conjunction with the monitor to help guide the user to successfully complete multi-step tasks (such as cooking) that would be difficult or impossible for the user to perform without assistance. Elite Care provides a perfect testbed for the ADL PROMPTER since the facility allows residents full access to kitchen facilities, and they often find great personal satisfaction in helping with meal preparation. Residents suffering from cognitive loss often wish to continue working in the kitchen, but are unable to safely and reliably complete multi-step tasks. The ADL PROMPTER will be designed to monitor and prompt (only as necessary) a user through the creation of a simple meal.

## References


[1] D.W. Albrecht, I. Zukerman, A.E. Nicholson, and A. Bud. Towards a bayesian model for keyhole plan recognition in large domains. In *Proceedings of the Sixth International Conference on User Modeling, Sardinia, Italy*, pages 365–376. User Modeling, Springer-Verlag, June 1997.

[2] James F. Allen and C. Raymond Perrault. Analyzing intention in utterances. *Artificial Intelligence*, 15:143–178, 1980.

[3] *Alzheimers Disease: the Costs to US Businesses in 2002*. Alzheimers Association, 2002.

[4] C. R. Anderson, P. Domingos, and D. S. Weld. Adaptive web navigation for wireless devices. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, 2001.

[5] C. R. Anderson, P. Domingos, and D. S. Weld. Personalizing web sites for mobile users. In *Proceedings of the Tenth International World Wide Web Conference*, 2001.

[6] C. R. Anderson, A. Y. Levy, and D. W. Weld. Declarative web-site management with tiramisu. In *Proceedings of the International Workshop on The Web and Databases (WebDB)*, 1999.

[7] Corin R. Anderson, Pedro Domingos, and Daniel S. Weld. Relational markov models and their application to adaptive web navigation. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-2002)*, 2002.

[8] L. F. Arnstein, S. Sigurdsson, and R. Franza. Ubiquitous computing in the biology laboratory. *Journal of Lab Automation (JALA)*, 6(1), 2001.

[9] D. Azari, J.K. Burns, K. Deshmukh, D. Fox, D. Grimes, C.T. Kwok, R. Pitkanen, P. Shon, and P. Tressel. Team description: UW Huskies-01. In A. Birk, S. Coradeschi, and S. Tadokoro, editors, *RoboCup-2001: Robot Soccer World Cup V*. Springer Verlag, 2001.

[10] G. Baltus, D. Fox, F. Gemperle, J. Goetz, T. Hirsh, D. Magaritis, M. Montemerlo, J. Pineau, N. Roy, J. Schulte, and S. Thrun. Towards personal service robots for the elderly. In *Proc. of the Workshop on Interactive Robotics and Entertainment (WIRE-2000)*, 2000.

[11] Y. Bar-Shalom, X.-R. Li, and T. Kirubarajan. *Estimation with Applications to Tracking and Navigation*. John Wiley, 2001.

[12] M. Bennewitz, W. Burgard, and S. Thrun. Using EM to learn motion behaviors of persons with mobile robots. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2002.

[13] J. Binder, K. Murphy, and S. Russell. Space-efficient inference in dynamic probabilistic networks. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, 1997.

[14] Aaron F. Bobick, Stephen S. Intille, James W. Davis, Freedom Baird, Claudio S. Pinhanez, Lee W. Campbell, Yuri A. Ivanov, Arjan Schütte, and Andrew Wilson. Perceptual user interfaces: the kidsroom. *Communications of the ACM*, 43(3):60 – 61, 2000.

[15] P. Bonnifait, P. Bouron, P. Crubillé, and D. Meizel. Data fusion of four ABS sensors and GPS for an enhanced localization of car-like vehicles. In *Proc. of the IEEE International Conference on Robotics & Automation*, 2001.

[16] G. Borriello and R. Want. Embedded computation meets the world-wide-web. *Communications of the ACM*, 43(5), 2000.

[17] C. Boutilier, N. Friedman, M. Goldszmidt, and D. Koller. Context-specific independence in Bayesian networks. In *Proc. 12th Conf. on Uncertainty in Artificial Intelligence*, pages 115–123, Portland, OR, 1996.

[18] J.A. Boyan and M.L. Littman. Exact solutions to time-dependent mdps. In *Advances in Neural Information Processing Systems 13 (NIPS)*. MIT Press, 2001.

[19] H.H. Bui, S. Venkatesh, and G. West. Policy recognition in the abstract hidden markov model. *Journal of Artificial Intelligence Research (JAIR)*, 17, 2002.

[20] W. Burgard, A. Derr, D. Fox, and A. B. Cremers. Integrating global position estimation and position tracking for mobile robots: the Dynamic Markov Localization approach. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1998.

[21] W. Burgard, D. Fox, D. Hennig, and T. Schmidt. Estimating the absolute position of a mobile robot using position probability grids. In *Proc. of the National Conference on Artificial Intelligence*, 1996.

[22] J.F. Canny and B. Donald. Simplified Voronoi diagrams. *Discrete and Computational Geometry*, 3, 1988.

[23] K. Chadha. The global positioning system: Challenges in bringing GPS to mainstream consumers. In *Proc. of the IEEE International Solid-State Circuits Conference*, 1998.

[24] H. Choset. *Sensor Based Motion Planning: The Hierarchical Generalized Voronoi Graph*. PhD thesis, Caltech, 1996.

[25] P. Chou, K. Hines, R. Ortega, K. Partridge, and G. Borriello. ipchinook: An integrated ip-based design framework for distributed embedded systems. In *Proc. of the 36th ACM/IEEE Design Automation Conference*, 1999.

[26] E. Cinlar. *Introduction to stochastic processes*. Prentice-Hall, 1975.

[27] N.J. Cohen and H. Eichenbaum. *Memory, Amnesia, and the Hippocampal System*. MIT Press, Cambridge, MA, 1995.

[28] B. Peintner Colbry and M. Pollack. Execution monitoring with quantitative temporal bayesian networks. In *Proceedings of the 6th International Conference on AI Planning and Scheduling (AIPS 2002)*, 2002.

[29] C. Conati, A.S. Gertner, K. VanLehn, and M.J. Druzdzel. Online student modeling for coached problem solving using bayesian networks. In *Proceedings of the Sixth International Conference on UserModeling, Sardinia, Italy*, pages 231–242. User Modeling, Springer-Verlag, June 1997.

[30] Z. Crisman, E. Curre, C.T. Kwok, L. Meyers, N. Ratliff, L. Tsybert, and D. Fox. Team description: UW Huskies-02. In *RoboCup-2002: Robot Soccer World Cup VI*. Springer Verlag, 2002.

[31] Y. Cui and S.S. Ge. Autonomous vehicle positioning with GPS in urban canyon environments. In *Proc. of the IEEE International Conference on Robotics & Automation*, 2001.

[32] M. Czerwinski, E. Cutrell, and E. Horvitz. Instant messaging: Effects of relevance and time. In *People and Computers XIV: Proceedings of HCI 2000*, pages 71–76. British Computer Society, British Computer Society, September 2000.

[33] N. de Freitas. Rao-Blackwellised particle filtering for fault diagnosis. *IEEE Aerospace,*, 2002.

[34] T. Dean and K. Kanazawa. A model for reasoning about persistence and causation. *Computational Intelligence*, 5(3):142–150, 1989.

[35] Kate Deibel, Dieter Fox, and Henry Kautz. Hierarchical hidden semi-markov models for activity tracking. Technical report, Dept CSE, U. Washington, 2003.

[36] Dempster, Laird, and Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1):1–38, 1977.

[37] Rich DeVaul, Steve Schwatz, and Sandy Pentland. The memory glasses project. online at http://www.media.mit.edu/wearables/mithril/memory-glasses.html, 2000.

[38] Anind K. Dey, Daniel Salber, and Gregory D. Abowd. A context-based infrastructure for smart environments. In *Proceedings of the 1st International Workshop on Managing Interactions in Smart Environments (MANSE '99)*, Dublin, Ireland, December 1999.

[39] R.J. Dolan and P.F. Fletcher. Encoding and retrieval in human medial temporal lobes: an empirical investigation using functional magnetic resonance imaging (fmri). *Hippocampus*, 9:25–34, 1999.

[40] A. Doucet, J.F.G. de Freitas, K. Murphy, and S. Russell. Rao-Blackwellised particle filtering for dynamic bayesian networks. In *Proc. of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2000.

[41] A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte Carlo in Practice*. Springer-Verlag, New York, 2001.

[42] M. Esler, J. Hightower, T. Anderson, and G. Borriello. Next century challenges: Data-centric networking for invisible computing: the portolano project at the university of washington. In *Proceedings Mobicom-99*, 1999.

[43] O. Etzioni, S. Hanks, T. Jiang, R. Karp, O Madani, and O. Waarts. Efficient information gathering on the internet. In *Foundations of Computer Science (FOCS)*, pages 234–243, 1996.

[44] O. Etzioni and D. Weld. A softbot-based interface to the Internet. *C. ACM*, 37(7):72–6, 1994.

[45] S. Fine, Y. Singer, and N. Tishby. The hierarchical hidden Markov model: analysis and applications. *Machine Learning*, 32, 1998.

[46] D. Fox. KLD-sampling: Adaptive particle filters. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14 (NIPS)*, Cambridge, MA, 2002. MIT Press.

[47] D. Fox, W. Burgard, F. Dellaert, and S. Thrun. Monte Carlo Localization: Efficient position estimation for mobile robots. In *Proc. of the National Conference on Artificial Intelligence*, 1999.

[48] D. Fox, W. Burgard, H. Kruppa, and S. Thrun. A probabilistic approach to collaborative multi-robot localization. *Autonomous Robots*, 8(3):325–344, 2000.

[49] D. Fox, W. Burgard, and S. Thrun. Markov localization for mobile robots in dynamic environments. *Journal of Artificial Intelligence Research (JAIR)*, 11:391–427, 1999.

[50] D. Fox, J. Hightower, and D. Schulz. People tracking with anonymous and id sensors. Technical Report UW-CSE-03-02-02, Dept. of Computer Science & Engineering, University of Washington, 2003. `http://www.cs.washington.edu/homes/fox/tracking`.

[51] D. Fox, S. Thrun, F. Dellaert, and W. Burgard. Particle filters for mobile robot localization. In Doucet et al. [41].

[52] N. Friedman, L. Getoor, D. Koller, and A. Pfeffer. Learning probabilistic relational models. In *Proc. 16th Intl. Joint Conf. on Artificial Intelligence*, pages 1300–1309, Stockholm, Sweden, 1999.

[53] Nir Friedman, Kevin Murphy, and Stuart Russell. Learning the structure of dynamic probabilistic networks. In *Proceedings of the Fourteenth Conference on Uncertainty in Artifical Intelligence*, pages 139–147, 1998.

[54] M. Garagnani, L. Shastri, and C. Wendelken. A connnectionist model of planning via back-chaining search. In *Proceedings of the Twenty-Fourth Conference of the Cognitive Science Society*, 2002.

[55] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian Data Analysis*. Chapman and Hall/CRC, 1997.

[56] L. Getoor, N. Friedman, D. Koller, and B. Taskar. Learning probabilistic models of relational structure. In *Proc. 18th Intl. Conf. on Machine Learning*, pages 170–177, Williamstown, MA, 2001.

[57] T. Gomez-Isla, J.L. Price, D.W. McKeel, J.C. Morris, J.H. Growdon, and B.T. Hyman. Profound loss of layer ii entorhinal cortex neurons occurs in very mild alzheimer's disease. *Journal of Neuroscience*, 16:4491–4500, 1996.

[58] J.S. Gutmann, W. Burgard, D. Fox, and K. Konolige. An experimental comparison of localization methods. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1998.

[59] J.S. Gutmann and D. Fox. An experimental comparison of localization methods continued. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2002.

[60] J. Hightower and G. Borriello. Location systems for ubiquitous computing. *Computer*, 34(8), 2001. IEEE Computer Society Press.

[61] G.W. Van Hoesen and B.T. Hyman. Hippocampal formation: anatomy and patterns of pathology in alzheimer's disease. *Progress in Brain Research*, 83:445–457, 1990.

[62] E. Horvitz. Principles of mixed-initiative user interfaces. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*, pages 159–166. ACM Press, May 1999.

[63] E. Horvitz and M. Barry. Display of information for time-critical decision making. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, pages 296–305, Montreal, Canada, August 1995. Morgan Kaufmann, San Francisco.

[64] E. Horvitz, J. Breese, D. Heckerman, D. Hovel, and K. Rommelse. The Lumiere project: Bayesian user modeling for inferring the goals and needs of software users. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, pages 256–265. Morgan Kaufmann, San Francisco, July 1998.

[65] E. Horvitz, A. Jacobs, and D. Hovel. Attention-sensitive alerting. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, pages 305–313. AUAI, Morgan Kaufmann, August 1999.

[66] E. Horvitz, C.M. Kadie, T. Paek, and D. Hovel. Models of attention in computing and communication: From principles to applications. *Communications of the ACM*, 46(3), March 2003.

[67] E. Horvitz and G. Rutledge. Time-dependent utility and action under uncertainty. In *Proceedings of Seventh Conference on Uncertainty in Artificial Intelligence, Los Angeles, CA*, pages 151–158. Morgan Kaufmann, San Francisco, July 1991.

[68] Eric Horvitz, Yongshao Ruan, Henry Kautz, Carla Gomes, Bart Selman, and Max Chickering. A bayesian approach to tackling hard computational problems. In *Proceedings the 17th Conference on Uncertainty in Artificial Intelligence (UAI-2001)*, 2001.

[69] R.A. Howard. *Dynamic Probabilistic Systems*. John Wiley & Sons, New York, 1971.

[70] C. Huang and A. Darwiche. Inference in belief networks: A procedural guide. *Intl. J. Approx. Reasoning*, 15(3), 1996.

[71] M. Isard and A. Blake. Condensation – conditional density propagation for visual tracking. *International Journal of Computer Vision*, 29(1):5–28, 1998.

[72] Zachary Ives, Alon Halevy, and Dan Weld. Integrating network-bound xml data. *Data Engineering Bulletin 24(2)*, pages 20–26, 2001.

[73] Zachary Ives, Alon Halevy, and Dan Weld. An xml query engine for network-bound data. *VLDB Journal, Special Issue on XML Query Processing*, 2003.

[74] Zachary Ives, Alon Levy, Dan Weld, Daniela Florescu, and Marc Friedman. Adaptive query processing for internet applications. *Data Engineering Bulletin 23(3)*, 2000.

[75] Zachary G. Ives, Daniela Florescu, Marc Friedman, Alon Levy, and Daniel S. Weld. An adaptive query execution system for data integration. In *Proceedings of SIGMOD Conference on Management of Data*, Philadelphia, PA, USA, June 1999.

[76] Zachary G. Ives, Alon Y. Halevy, and Daniel S. Weld. Integrating network-bound XML data. *IEEE Data Engineering Bulletin Special Issue on XML*, 24(2), June 2001.

[77] A. Jameson. Numerical uncertainty management in user and student modeling: An overview of systems and issues. *User Modeling and User-Adapted Interaction*, 5:193–251, 1996.

[78] Tony Jebara and Alex Pentland. Action reaction learning: Automatic visual analysis and synthesis of interactive behaviour. In *Proceedings of the International Conference on Vision Systems (ICVS'99)*, Las Palmas de Gran Canaria, Spain, 1999.

[79] Henry Kautz. A formal theory of plan recognition and its implementation. In J.F. Allen, H.A. Kautz, R.N. Pelavin, and J.D. Tennenberg, editors, *Reasoning About Plans*, pages 69–126. Morgan Kaufmann Publishers, 1991.

[80] Henry Kautz, Eric Horvitz, Yongshao Ruan, Carla Gomes, and Bart Selman. Dynamic restart policies. In *Proceedings AAAI-2002*. AAAI Press, 2002.

[81] Henry Kautz and Peter Ladin. Integrating metric and qualitative temporal reasoning. In *Proceedings of AAAI-91*, pages 241–246. AAAI Press, 1991.

[82] Jonathan Klein, Youngme Moon, and Rosalind W. Picard. This computer responds to user frustration: Theory, design, results, and implications. *Interacting with Computers – Special Issue on Cognition and Emotion*, (to appear), 2002.

[83] D. Koller and R. Fratkina. Using learning for approximation in stochastic processes. In *Proc. of the International Conference on Machine Learning*, 1998.

[84] D. Koller and U. Lerner. Sampling in factored dynamic systems. In Doucet et al. [41].

[85] Cody C. T. Kwok, Oren Etzioni, and Daniel S. Weld. Scaling question answering to the web. In *World Wide Web*, pages 150–161, 2001.

[86] C.T. Kwok, D. Fox, and M. Meilă. Real-time particle filters. In *Advances in Neural Information Processing Systems 15*, 2002.

[87] C.T. Kwok, D. Fox, and M. Meilă. Adaptive real-time particle filters for robot localization. In *Proc. of the IEEE International Conference on Robotics & Automation*, 2003.

[88] T. Lau, S. Wolfman, P. Domingos, and D. Weld. Programming by demonstration using version space algebra. *Machine Learning*, 2003.

[89] Tessa Lau, Pedro Domingos, and Daniel S. Weld. Version space algebra and its application to programming by demonstration. In *Proceedings of the Seventeenth International Conference on Machine Learning*, pages 527–534, June 2000.

[90] Tessa Lau and Daniel S. Weld. Programming by Demonstration: an Inductive Learning Formulation. In *Proceedings of the 1999 International Conference on Intelligent User Interfaces (IUI '99)*, pages 145–152, Redondo Beach, CA, USA, January 1999.

[91] U. Lerner, M. Brooks, M. Scott, S. McIlraith, and D. Koller. Monitoring a complex physical system using a hybrid dynamic bayes net. In *Proc. of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2002.

[92] Neal Lesh. Adaptive goal recognition. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*. San Francisco, CA: Morgan Kaufmann, 1997.

[93] Neal Lesh. *Scalable and Adaptive Goal Recognition*. PhD thesis, University of Washington, 1998.

[94] Neal Lesh and Oren Etzioni. A sound and fast goal recognizer. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 1704–1710. San Francisco, CA: Morgan Kaufmann, 1995.

[95] Neal Lesh and Oren Etzioni. Scaling up goal recognition. In *Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning*, pages 178–189, 1996.

[96] Y. Li, T. Wang, and H.Y. Shum. Motion textures: A two-level statistical model for character motion synthesis. In *Proc. of the International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, 2002.

[97] N. L. Mace and P. V. Rabins. *Managing Symptoms in the 36-Hour Day*. The Johns Hopkins University Press, Baltimore, 1981.

[98] P. S. Maybeck. The Kalman filter: An introduction to concepts. In I. J. Cox and G. T. Wilfong, editors, *Autonomous Robot Vehicles*. Springer Verlag, 1990.

[99] Andrew K. McCallum, Ronald Rosenfeld, Tom M. Mitchell, and Andrew Y. Ng. Improving text classification by shrinkage in a hierarchy of classes. In Jude W. Shavlik, editor, *Proceedings of ICML-98, 15th International Conference on Machine Learning*, pages 359–367, Madison, US, 1998. Morgan Kaufmann Publishers, San Francisco, US.

[100] C.E. McCarthy and M. Pollack. A plan-based personalized cognitive orthotic. In *Proceedings of the 6th International Conference on AI Planning and Scheduling (AIPS 2002)*, 2002.

[101] A. W. Moore, J. Schneider, and K. Deng. Efficient locally weighted polynomial regression predictions. In *Proc. of the International Conference on Machine Learning*, 1997.

[102] Darnell Moore and Irfan Essa. Recognizing multitasked activities from video using stochastic context-free grammar. In *Proceedings of AAAI-02*, 2002.

[103] R. Morales-Menéndez, N. de Freitas, and D. Poole. Real-time monitoring of complex industrial processes with particle filters. In *Advances in Neural Information Processing Systems 15*, 2002.

[104] S. Muggleton. Stochastic logic programs. In L. De Raedt, editor, *Proceedings of the 5th International Workshop on Inductive Logic Programming*, page 29. Department of Computer Science, Katholieke Universiteit Leuven, 1995.

[105] K. Murphy. The bayes net toolbox for matlab. *Computing Science and Statistics*, 33, 2001.

[106] K. Murphy. *Dynamic Bayesian Networks: Representation, Inference and Learning*. PhD thesis, UC Berkeley, Computer Science Division, 2002.

[107] K. Murphy and M. Paskin. Linear time inference in hierarchical HMMs. In *Advances in Neural Information Processing Systems (NIPS)*, 2001.

[108] K. Murphy and S. Russell. Rao-Blackkwellised particle filtering for dynamic Bayesian networks. In Doucet et al. [41].

[109] K. Murphy and S. Russell. Rao-Blackwellised particle filtering for dynamic Bayesian networks. In A. Doucet, N. de Freitas, and N. Gordon, editors, *Sequential Monte Carlo Methods in Practice*, pages 499–516. Springer, New York, 2001.

[110] E. Mynatt, I. Essa, and W. Rogers. Increasing the oppurtunities for aging-in-place. In *Proceedings of ACM Conference on Universal Usability*, 2000.

[111] Progress report on alzheimer's disease 2000. online at http://www.alzheimers.org, 2000.

[112] Liem Ngo and Peter Haddawy. Answering queries from context-sensitive probabilistic knowledge bases. *Theoretical Computer Science*, 171(1–2):147–177, 1997.

[113] S. M. Omohundro. Bumptrees for efficient function, constraint, and classification learning. In R. P. Lippmann, J. E. Moody, and D. S. Touretzky, editors, *Advances in Neural Information Processing Systems 3*. Morgan Kaufmann, 1991.

[114] Robert J. Orr and Gregory D. Abowd. The smart floor: a mechanism for natural user identification and tracking. In *Proceedings of the 2000 Conference on Human Factors in Computing Systems (CHI 2000)*, The Hague, Netherlands, 2000.

[115] Donald J Patterson, Oren Etzioni, and Henry Kautz. The activity compass. Technical report, Dept CSE, U. Washington, 2003.

[116] J.S. Penberthy and D. Weld. Temporal planning with continuous change. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, volume 2, pages 1010–1015, Seattle, Washington, USA, July 1994. Menlo Park, CA: AAAI Press.

[117] A. Pentland. Machine understanding of human action. In *Proceedings 7th International Forum on Frontier of Telecom Technology*, Tokyo, Japan, 1995.

[118] A. Pentland. Smart rooms. *Scientific American*, 274(4):68–76, 1996.

[119] M. Perkowitz and O. Etzioni. Adaptive web sites: an AI challenge. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, 1997.

[120] M. Perkowitz and O. Etzioni. Towards adaptive web sites: Conceptual framework and case study. *Artificial Intelligence*, To appear, 2000.

[121] R. W. Picard. *Affective Computing*. MIT Press, Cambridge, 1997.

[122] M. E. Pollack, C. E. .McCarthy, S. Ramakrishnan, I. Tsamardinos, L. Brown, S. Carrion, D. Colbry, C. Orosz, and B. Peintner. Autominder: A planning, monitoring, and reminding assistive agent. In *Proceedings of the 7th International Conference on Intelligent Autonomous Systems*, 2002.

[123] Nissanka B. Priyantha, Anit Chakraborty, and Hari Balakrishnan. The cricket location-support system. In *Proceedings of MOBICOM 2000*, pages 32–43, Boston, MA, August 2000. ACM, ACM Press.

[124] Martin L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley and Sons, 1994.

[125] A. Raftery. Hypothesis testing and model selection via posterior simulation. In W.R. Gilks, S. Richardson, and D.J. Spiegelhalter, editors, *Markov Chain Monte Carlo in Practice*. Chapman & Hall/CRC, London, 1996.

[126] N.L. Rempel-Clower, S.M. Zola, L.R. Squire, and D.G. Amaral. Three cases of enduring memory impairment after bilateral damage limited to the hippocampal formation. *Journal of Neuroscience*, 16:5233–5255, 1996.

[127] D. Rice, P. J. Fox, and W. Max. The economic burden of alzheimer's disease care. *Health Affairs*, 12(2):164–176, 1993.

[128] Facing a healthier future, October 2001.

[129] Tracie Rozhon. The medicine chest pumps up. *The New York Times*, February 2001.

[130] Yongshao Ruan, Eric Horvitz, and Henry Kautz. Restart policies with dependence among runs: A dynamic programming approach. In *Proceedings of the Eighth International Conference on Principles and Practice of Constraint Programming (CP-2002)*. Springer-Verlag, 2002.

[131] Jane M. Sanders. Sensing the subtleties of everyday life. *Research Horizons*, Winter 2000.

[132] J.Z. Sasiadek and Q. Wang. Sensor fusion based on fuzzy kalman filtering for autonomous robot vehicle. In *Proc. of the IEEE International Conference on Robotics & Automation*, 1999.

[133] Daniel L. Schacter. *Searching for memory: The brain, the mind, and the past*. Basic Books, 1996.

[134] D.L. Schacter and A.D. Wagner. Medial temporal lobe activations in fmri and pet studies of episodic encoding and retrieval. *Hippocampus*, 9:7–24, 1999.

[135] C. F. Schmidt, N. S. Sridharan, and J. L. Goodson. The plan recognition problem: an intersection of psychology and artificial intelligence. *Artificial Intelligence*, 11:45–83, 1978.

[136] D. Schulz, W. Burgard, and D. Fox. Tracking multiple moving objects with a mobile robot. In *Proc. of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2001.

[137] D. Schulz, W. Burgard, and D. Fox. People tracking with a mobile robot using joint probabilistic data association filters. *International Journal of Robotics Research (IJRR)*, 2003. Accepted for publication.

[138] L. Shastri. A computational model of rapid memory formation in the hippocampal system. In *Proceedings of the Nineteenth Conference of the Cognitive Science Society*, Stanford University, CA, August 1997.

[139] L. Shastri. Advances in SHRUTI - a neurally motivated model of relational knowledge representation and rapid inference using temporal synchrony. *Applied Intelligence*, 11, 1999.

[140] L. Shastri. An extended local connectionist manifesto: Embracing relational and procedural knowledge. *Behavioral and Brain Sciences*, 23:492–493, 2000.

[141] L. Shastri. Types and quantifiers in shruti — a connectionist model of rapid reasoning and relational processing. In S. Werter and R. Sun, editors, *Hybrid Neural Symbolic Integration*, pages 28–45. Springer-Verlag, Heidelberg, 2000.

[142] L. Shastri. A biological grounding of recruitment learning and vicinal algorithms. In J. Austin, S. Wermter, and D. Wilshaw, editors, *Emergent neural computational architectures based on neuroscience*, pages 348–367. Springer-Verlag, 2001.

[143] L. Shastri. A computational model of episodic memory formation in the hippocampal system. *Neurocomputing*, 38-40:889–897, 2001.

[144] L. Shastri. A computationally efficient abstraction of long-term potentiation. *Neurocomputing*, 44-46:33–41, 2002.

[145] L. Shastri. Episodic memory and cortico-hippocampal interactions. *Trends in Cognitive Sciences*, 6:162–168, 2002.

[146] L. Shastri. From transient patterns to persistent structure: A model of episodic memory formation via cortico-hippocampal interactions. *Behavioral and Brain Sciences*, In revision. http://www.icsi.berkeley.edu/ shastri/psfiles/shastri_ em.pdf.

[147] L. Shastri and V. Ajjanagadde. From simple associations to systematic reasoning: a connectionist encoding of rules, variables and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences*, 16(3):417–494, 1993.

[148] L. Shastri and C. Wendelken. Knowledge fusion in the large - taking a cue from the brain. In *Proceedings of the Second International Conference on Information Fusion*, Sunnyvale, CA, 1999.

[149] L. Shastri and C. Wendelken. Seeking coherent explanations - a fusion of structured connectionism, temporal synchrony, and evidential reasoning. In *Proceedings of the Twenty-Second Conference of the Cognitive Science Society*, Philadelphia, August 2000.

[150] David Shenk. *The Forgetting: Alzherimer's: Portrait of an Epidemic*. Doubleday, New York, NY, 2001.

[151] D. Smith and D. Weld. Temporal graphplan with mutual exclusion reasoning. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, Stockholm, Sweden, Aug 1999. San Francisco, CA: Morgan Kaufmann.

[152] Padhraic Smyth. Clustering sequences with hidden markov models. In Michael C. Mozer, Michael I. Jordan, and Thomas Petsche, editors, *Advances in Neural Information Processing Systems*, volume 9, page 648. The MIT Press, 1997.

[153] L.R. Squire. Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review*, 99:195–231, 1992.

[154] Igor Tatarinov, Zachary Ives, Alon Halevy, and Dan Weld. Updating XML. In *Proc. of SIGMOD*, pages 413–424, 2001.

[155] R. Thrapp, C. Westbrook, and D. Subramanian. Robust localization algorithms for an autonomous campus tour guide. In *Proc. of the IEEE International Conference on Robotics & Automation*, 2001.

[156] S. Thrun, J. Langford, and V. Verma. Risk sensitive particle filters. In *Advances in Neural Information Processing Systems 14*. MIT Press, 2001.

[157] V. Verma, J. Langford, and R. Simmons. Non-parametric fault identification for space rovers. In *International Symposium on Artificial Intelligence and Robotics in Space (iSAIRAS)*, 2001.

[158] J. Vermaak, C. Andrieu, A. Doucet, and S.J. Godsill. Particle methods for Bayesian modelling and enhancement of speech signals. *IEEE Transactions on Speech and Audio Processing*, 2002. Accepted for publication.

[159] M. Vilain, H. Kautz, and P. van Beek. Constraint propagation algorithms for temporal reasoning. In *Readings in Qualitative Reasoning about Physical Systems*. Morgan Kaufmann, 1990.

[160] R. Want and G. Borriello. Tutorial on information appliances. *IEEE Computer Graphics & Applications*, 20(3), 2000.

[161] R. Want, A. Hopper, V. Falcao, and J. Gibbons. The active badge location system. *ACM Transactions on Information Systems*, 10(1), 1992.

[162] University of washington/microsoft research summer institute technologies of invisible computing, July 1999.

[163] C. Wendelken. The role of mid-dorsolateral prefrontal cortex in working memory: a connectionist model. *Neurocomputing*, 44-46:1009–1016, 2002.

[164] C. Wendelken and L. Shastri. Probabilistic inference and learning in a connectionist causal network. In *Proceedings of the Second International Symposium on Neural Computation*, May 2000.

[165] C. Wendelken and L. Shastri. Combining belief and utility in a structured connectionist agent architecture. In *Proceedings of the Twenty-Fourth Annual Conference of the Cognitive Science Society*, 2002.

[166] M. West and P.J. Harrison. *Bayesian Forecasting and Dynamic Models*. Springer-Verlag, New York, 2nd edition, 1997.

[167] S. Wolfman and D. Weld. The LPSAT System and its Application to Resource Planning. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pages 310–317, Stockholm, Sweden, Aug 1999. San Francisco, CA: Morgan Kaufmann.

[168] Steven A. Wolfman, Tessa Lau, Pedro Domingos, and Daniel S. Weld. Mixed initiative interfaces for learning tasks: Smartedit talks back. In *Proceedings of the 2001 International Conference on Intelligent User Interfaces*, pages 167–174, Santa Fe, USA, January 2001.

[169] M.M. Zeineh, S.A. Engel, P.M. Thompson, and S.Y. Bookheimer. Dynamics of the hippocampus during encoding and retrieval of face-name pairs. *Science*, 299:577–580, 2003.