

# High-level Goal Recognition in a Wireless LAN

Jie Yin and Xiaoyong Chai and Qiang Yang

Department of Computer Science  
Hong Kong University of Science and Technology  
Clearwater Bay, Kowloon, Hong Kong, China  
{yinjie, carnamel, qyang}@cs.ust.hk

## Abstract

Plan recognition has traditionally been developed for logically encoded application domains with a focus on logical reasoning. In this paper, we present an integrated plan-recognition model that combines low-level sensory readings with high-level goal inference. A two-level architecture is proposed to infer a user's goals in a complex indoor environment using an RF-based wireless network. The novelty of our work derives from our ability to infer a user's goals from sequences of signal trajectory, and the ability for us to make a trade-off between model accuracy and inference efficiency. The model relies on a dynamic Bayesian network to infer a user's actions from raw signals, and an N-gram model to infer the users' goals from actions. We present a method for constructing the model from the past data and demonstrate the effectiveness of our proposed solution through empirical studies using some real data that we have collected.

## Introduction

As the pervasive computing technology becomes more and more mature, different context-aware applications have emerged. A variety of techniques are now available for sensing the location of an individual using GPS devices, active badges, motion detectors, mobile phones and PDA's in wireless networks. Using these systems, an individual's location can be tracked in a pervasive computing environment. On observing a user's past and current mobile sensory readings, it is natural to ask: where is the user likely to visit in the next while? What is the user trying to do? What is the ultimate goal of the user?

The above questions are instances of inferring high-level user-behavior patterns from low-level sensory data through location-based plan recognition. Being able to accomplish this task is critical to many applications. For people suffering from various cognitive limitations in hospitals and care facilities, the technique can discover when a person's behavior is out of the norm and provide help in a timely manner (Patterson *et al.* 2003). For shoppers in a busy business environment such as a shopping mall, services and products can

be offered not only according to the people's current location, but also according to their intended actions and goals.

Despite the potentially wide range of applications, location-based plan recognition remains a difficult task. One reason is that the location itself is often inferred from low-level sensory data, but often such data provides ambiguous information, making it impossible to obtain accurate estimation. This is particularly the case in a wireless LAN, in-door environment. Another source of difficulty is that the users' behavior is often inherently ambiguous. For example, a professor in an academic building may follow a route by which he walks through a sequence of hallways and then reaches some location in the office area. However, his intention and subsequent behavior can be quite different depending on not just his current location but his entire trace as well as the time and data of his activities. Depending on his goals he may wish to attend a seminar or to go to a printer.

In this paper, we address the problem of inferring a user's high-level goals from low-level mobile data in an indoor environment, where a wireless LAN is available. The novelty of our work can be seen from several aspects. First, in the pervasive computing literature, an important focus has been to determine and track a user's location from sensory data. Examples include the use of GPS, ultrasonic-based systems, infrared-based systems (Fox *et al.* 2002) and radio frequency (RF)-based systems (Bahl & Padmanabhan 2000; Youssef, Agrawala, & Shankar 2003; Ladd *et al.* 2002). Among these systems, the RF-based systems utilize an underlying wireless network to estimate the location of users, which has gained more attention recently, especially for indoor applications. Unlike infrared-based systems, RF-based techniques can provide more ubiquitous coverage, and do not require additional expensive hardware since many buildings are already equipped with IEEE 802.11b wireless Ethernet. Following this tradition, in our work, we utilize the data gathered from an RF-based wireless LAN. However, so far there has been a lack of study on the problem of inferring a user's high-level goals from low-level sensory data. In fact, as observed in (Patterson *et al.* 2003), having a good understanding about a user's high-level behavior patterns and goals will help in estimating the user's current locations.

Second, in the artificial intelligence area, recognizing complex high-level behavior has traditionally been the focus of plan recognition (Kautz & Allen 1986; Lesh & Et-

zioni 1995). A Bayesian network was used for plan recognition in story understanding in (Charniak & Goldman 1993). In (Blaylock & Allen 2003), a corpus-based N-gram model was proposed to predict the goal from a given sequence of command actions in the UNIX domain. In addition, other advanced stochastic models for recognizing high-level behavior were proposed such as Dynamic Bayesian Networks (DBNs) (Albrecht, Zukerman, & Nicholson 1998) and Probabilistic State Dependent Grammars (Pynadath & Wellman 2000). However, most of the work in plan recognition has been restricted to the high level for inference, and the challenge of dealing with low-level sensor models has not been addressed.

Only in recent years, attempts have been made to integrate high-level behavior models with low-level sensor models. The work of (Patterson *et al.* 2003) presented an approach by applying a Bayesian model to predict a user’s transportation mode based on location readings from GPS devices in an urban environment. One feature of the GPS technology is that it can directly provide the information of locations more accurately. In contrast, in our problem, only RF-based signal-strength values from multiple sensor sources are available. Moreover, the signal itself is highly unstable and is affected by the environment to a large extent. Therefore, it is a challenging task to build a model that is both accurate and efficient using such noisy signals. Similarly, the activity maps developed in (Demirdjian *et al.* 2002) used computer vision to detect the users’ non-transient behavior through spatial-temporal clustering.

In this paper, we present an architecture for inferring a user’s high-level actions and goals from data obtained in the context of a popular RF-based wireless network in complex indoor environments. We build on the work of (Patterson *et al.* 2003), where a DBN is applied for inferring actions from traces of signals. On top of this framework, we provide a novel two-level hierarchy to enable the inference of goals from actions efficiently and accurately. Instead of directly applying a multi-level DBN model, which does not scale well for increasing numbers of goals, we provide a simple N-gram inference method for inferring goals from action sequences. We show that this model is more efficient than a monolithic DBN model, and is almost as accurate as a DBN model. One advantage of our architecture is that high-level, common sense knowledge can be easily incorporated into the reasoning process.

The rest of the paper is organized as follows. We begin by providing an overview of the problem domain. Then we present a DBN solution to the recognition problem. We then propose a two-level model that is more efficient but can achieve almost the same accuracy as the DBN. We describe experimental results in a real office environment. Finally we discuss conclusions and some directions for future work.

### Overview of Problem Domain

In this section, we first describe the problem domain in detail. The layout of the environment is shown in Figure 1. In the figure, several base stations with known Media Access Control (MAC) addresses are marked with double concrete circles.




Figure 1: The Layout of Office Area

After collecting sequences of signals from the base stations and labelling each sequence with its intended goal, we obtain a database of a user’s historic traces to achieve different goals as given in Table 1. Each trace, corresponding to a single goal, records a sequence of observed signals where each element of the trace is a vector. At each time instant  $t_i$ , the vector consists of several pairs of MAC addresses and signals from the corresponding base stations where the signals can be detected by a wireless device. For example, we can see from Table 1 that at the time point  $t_2$ , the signal-strength values of three base stations for the first trace are 56, 30 and 62, respectively.

Trace #	Observed Signal Sequences				Goal
	$t_1$	$t_2$	...	$t_k$	
1	(b1:57)	(b1:56)	(b1:55)	(b1:52)	$G_1$
	(b2:33)	(b2:30)	(b2:36)	(b2:62)	
	(b3:51)	(b3:62)	(b3:56)	(b3:47)	
2	(b1:62)	(b1:39)	(b1:46)	(b1:41)	$G_2$
	(b2:57)	(b2:41)	(b2:45)	(b2:43)	
	(b3:55)	(b3:32)	(b3:43)	(b3:27)	

Table 1: An Example of Trace Database

A user’s behavior model can be built based on these labelled historic sensory data. A user’s behavior is represented as a sequence of actions to achieve a goal in the high level. There are 11 actions and 19 goals of a professor’s behavior in this environment. To illustrate, examples of actions include “Walk-in-HW1”, “Enter-Room1” and “Print”. Out of the 19 goals, we illustrate using four examples of goals.  $G_1$  = “Seminar-in-Room1”: a professor leaves his office, walks through hallways HW1, HW3, HW4, and enters the P&S Room1 to attend a seminar there;  $G_2$  = “Print-in-Room1”: he follows the same route as before, but he goes to get some printed material from a printer in P&S Room1;  $G_3$  = “Seminar-in-Room2”: the professor leaves his office, walks through hallways HW1, HW3 and HW6 and then enters P&S Room2 for attending a seminar;  $G_4$  = “Return-to-Office”: he returns to his office after teaching a course via hallways HW5, HW4, HW3 and HW1.

The high-level goal-recognition problem is, given a new sequence of observed signals received from multiple base stations while a user is moving, infer the most probable high-level goal or intention of the user. For example, when a new sequence of signals  $\langle (b_1 : 48)(b_2 : 60)(b_3 : 32) \rangle \langle (b_1 : 45)(b_2 : 59)(b_3 : 35) \rangle \dots$  is observed, we would like to infer which goal the user is most probably pursuing currently,  $G_1$  or  $G_2$ .

However, subject to reflection, refraction, diffraction and absorption by structures and even human bodies, the signal propagation suffers from severe multi-path fading effects in an indoor environment (Hashemi 1993). As a result, the signal-strength value received from the same base station varies with time even at a fixed location, and the number of base stations covering a location varies with time. Therefore, the goal-recognition problem is complicated due to the *noisy characteristics* of signals.

### DBN for Action and Goal Recognition

We first model the goal-recognition problem based on the framework of Dynamic Bayesian networks (DBNs) (Dean & Kanazawa 1989; Murphy 2002a). A DBN extends Bayesian networks by including a temporal dimension. It consists of a sequence of Bayesian networks to represent the world. At each time slice, exactly the same Bayesian network is used to model the dependencies among variables. In addition to the intra-slice connections in the Bayesian network, inter-slice connections are also required to represent temporal dependencies in consecutive time slices.




Figure 2: Two Time-slice DBN Model

The DBN model shown in Figure 2 is adopted to model our goal-recognition problem. It shows two time slices numbered  $t$  and  $t - 1$  respectively in our behavior model. The shaded nodes  $SS_{i,t}$  represent the strength variables of signals received from multiple base stations, which are directly observable. All the other variables - the physical location  $S_t$  of the user, the action  $A_t$  the user is taking and the goal  $G_t$  the user is pursuing - are hidden, with the values to be inferred from the raw signals. The dependencies among these nodes are shown in two kinds of directed links: solid links for intra-slice connections and dashed links for inter-slice

connections. The network consists of two parts: a *sensor model* and an *action model*. At the bottom, a *sensor model* is used for location estimation based on sensory readings; while on the top, an *action model* is constructed for inferring actions and, subsequently, goals from the estimated locations. In the framework, two models are integrated by the location nodes that bridge the low-level evidence in terms of signal-strength values with the high-level actions and goals.

### Sensor Model

An important component of the DBN framework is the sensor model for estimating locations based on the signals received from multiple base stations. Regarding the DBN structure, in the time slice  $t$ , the links from the node  $S_t$  in the location layer (states) to nodes  $SS_{i,t}$  in the observation layer (signals) mean that the received signals are dependent on the actual physical location of the user in the current environment. The links are associated with probability distributions that reflect the uncertainty involved due to the noisy characteristics of the signals.

To build the sensor model, we model the world as a finite location-state space  $\mathbb{S} = \{s_1, \dots, s_n\}$  with a finite observation space  $\mathbb{O} = \{o_1, \dots, o_m\}$ . The *sensor model* is defined to be a predicted model of the conditional probabilities  $Pr(o_j|s_i)$ , the likelihood of observing some sensory measurement  $o_j \in \mathbb{O}$  at state  $s_i \in \mathbb{S}$ . The state space  $\mathbb{S}$  is defined as a set of physical grid points on the floor map:

$$\mathbb{S} = \{s_1 = (x_1, y_1), \dots, s_n = (x_n, y_n)\}.$$

An observation  $o_j$  in the observation space  $\mathbb{O}$  consists of a set of signal-strength measurements received from  $k$  base stations respectively. We represent each observation  $o_j$  at a particular state  $s_i$  as a vector:

$$o_j = \langle (b_1, ss_{1j}), \dots, (b_k, ss_{kj}) \rangle.$$

where  $b_p$  represents the  $p$ th base station and  $ss_p$  is the average signal-strength value received from the  $p$ th base station.

Our on-site calibration revealed that it would be improper to impose some commonly used assumptions that use Gaussian or some other predefined statistical distributions to fit signal distribution, a point also made by (Roos *et al.* 2002). Therefore, we adopted a simpler but more robust scheme of directly sampling the conditional probabilities. We recorded the signal-strength values at each grid point  $s_i$  and built a histogram of observed signal-strength values for each base station. A conditional probability  $Pr(ss_p|b_p, s_i)$  was calculated for each base station, which is the probability that the base station  $b_p$  has the average signal-strength value  $ss_p$  at the state  $s_i$ . By making an independence assumption among signals from different base stations, we compute the conditional probability of receiving a particular observation  $o_j$  at state  $s_i$  as

$$Pr(o_j|s_i) = \prod_{p=1}^k Pr(ss_p|b_p, s_i).$$

### Action Model

After the sensor model is built, the next issue is to *automatically construct* the *action model* for inferring high-level

actions and goals. In the DBN structure, the links from the node  $G$  to the node  $A$  mean that a goal is achieved by carrying out a sequence of actions, and the links from the node  $A$  to the node  $S$  mean that an action takes place in some locations. The action  $A_t$  that a user takes at the time slice  $t$  depends on his action  $A_{t-1}$  and his location  $S_{t-1}$  at the previous time slice, as well as the goal  $G_t$  he is currently aiming at. By making use of the domain knowledge such as the number of location states, the number of actions and goals, we use an EM algorithm (Dempster, Laird, & B. Rubin 1977) on the training sequences of signals to learn the conditional probabilities among the nodes in the DBN.

### A Two-Level Recognition Model

Despite the power and flexibility of the DBN in providing a coherent modelling framework, a DBN model does suffer from certain disadvantages. Chiefly among them is that the complexity of the model is exponential in the number of hidden states. A large number of parameters need to be estimated and as a result a large data set is required to obtain statistically meaningful results. Therefore, to meet the needs of the real-time goal recognition, we wish to find a tradeoff between the expressive power of DBN and inference efficiency. We achieve this tradeoff in terms of several adaptations.

First, the users' behavior is often inherently ambiguous in our office environment. A typical example is that the activities taking place in the same location cannot be distinguished solely based on locations. Therefore, we introduce the concept of the time duration. In the framework of DBN, a semi-Markov structure can be appended to allow hidden states to have variable durations (Murphy 2002a). The basic idea is that each state emits a sequence of observations and the state duration can be specified. In a DBN, duration nodes can be added to explicitly represent how long a user has been in one state (Murphy 2002b). Although the conditional probability of these nodes is deterministic (counting the time steps spent in one state), inference is still linear in the maximum number of steps spent in a state. Therefore, to achieve the inference efficiency, we have to reduce the number of the sampled time points. However, this in turn has the damaging effect of reducing the accuracy.

Second, in order to achieve the tradeoff between accuracy and computational complexity, we propose a novel two-level architecture that separates the whole DBN model into two parts: a low-level DBN model and a high-level N-gram model. As shown in Figure 3, the low-level DBN model starts from the observation layer to the action layer. The high-level N-gram model is applied to infer the goals from actions. The N-gram-based inference is a simple and yet effective technique that has shown great success in natural language processing (NLP) (see, e.g. (Chen & Goodman 1996)).

Given a sequence of observations  $o_1, o_2, \dots, o_t$  obtained up to time  $t$ , the low-level DBN model is responsible for computing the most possible action sequence  $A_1, A_2, \dots, A_t$ . This task is performed by using the junction tree algorithm (Murphy 2002a). After this, the next task is to infer the most likely goal. We define this task of *goal recognition* as follows: given an estimated action sequence

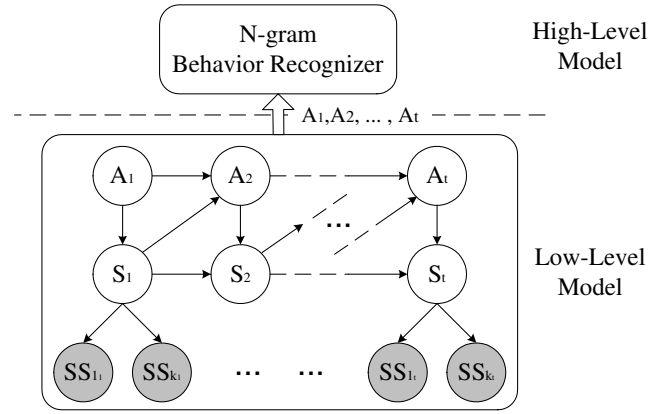


Figure 3: Two-Level Recognition Model

$A_1, A_2, \dots, A_t$  obtained so far, infer the most likely goal  $G^*$ :

$$\begin{aligned} G^* &= \arg \max P(G|A_1, A_2, \dots, A_t) \\ &= \arg \max P(G|A_{1:t}). \end{aligned}$$

By applying the Bayes' Rule, the above formula becomes:

$$\begin{aligned} G^* &= \arg \max \frac{P(A_{1:t}|G)P(G)}{P(A_{1:t})} \\ &= \arg \max P(A_{1:t}|G)P(G) \end{aligned}$$

where the term  $P(A_{1:t})$  is a constant and can be dropped.

Since each action lasts for some duration, we can represent the action sequence in a more compact form:  $A_{D_1}^1, A_{D_2}^2, \dots, A_{D_m}^m$ , where in each time segment  $D_i$ , the same action  $A_i$  is taken. Note that the sum of durations  $D_i$  is equal to the current time  $t$ :  $\sum_{i=1}^m D_i = t$ .

Furthermore, by using the Chain Rule, the calculation of  $P(A_{1:t}|G)$  can be expanded as follows:

$$\begin{aligned} G^* &= \arg \max P(A_{D_m}^m | A_{D_{m-1}}^{m-1}, \dots, A_{D_1}^1, G) \cdot \\ &P(A_{D_{m-1}}^{m-1} | A_{D_{m-2}}^{m-2}, \dots, A_{D_1}^1, G) \cdots P(A_{D_1}^1 | G). \end{aligned}$$

Since the estimation of these conditional probabilities are computationally expensive, we propose to use a simpler N-gram model: the action segment  $A_{D_i}^i$  only depends on the goal  $G$  and the  $n-1$  action segments  $A_{D_{i-n+1}}^{i-n+1}, \dots, A_{D_{i-1}}^{i-1}$  preceding it. By further assuming that the transitions between actions are independent of action durations  $D_i$ , we get a Bigram model when  $n=2$  as follows:

$$\begin{aligned} G^* &= \arg \max P(G)P(A_1|G) \cdot \\ &\prod_{i=2}^m P(A_i|A_{i-1}, G) \prod_{i=1}^m P(D_i|A_i) \end{aligned}$$

where nonparametric nor parametric representations of time distribution can be used for  $P(D_i|A_i)$ . Initially, the goal recognition is solely based on the prior  $P(G_i)$ . After a period of time  $t$ , the expanded action sequence is inferred from the observations, and the goal is recognized by the

above equation. The computational complexity is linear in the number of goals and in the length of an action sequence. An advantage of the proposed N-gram model is that the durations of actions can be explicitly and efficiently modelled.

### Experimental Evaluation

To test the validity of the model, we use the environment shown in Figure 1 as our testbed. The environment is modelled as a space of 99 states, each representing a 2-meter grid cell. Using the device driver and API we developed, the signals from base stations were recorded by an IBM laptop with a standard wireless Ethernet card. In our experiment, 8 out of 25 base stations were selected such that their signals occurred frequently and their signal-strength values were the strongest on average. We first collected 100 samples at each state, one per second, and these samples are used to estimate the sensor model. Then we collected about 570 traces for 19 goals of a professor’s behavior in the office area.




Figure 4: Behavior Recognition Example Using Four Out of the 19 Goals in Our Experiment

Figure 4 illustrates the recognition process of one trace belonging to the goal "Seminar-in-Room1", with respect to three other goals among the 19 goals. As shown in the figure, at the beginning, the probabilities of goals "Seminar-in-Room1", "Print-at-Room1" and "Seminar-in-Room2" are approximately equal since their starting points are the same, e.g. the office. However, one interesting point is that, at time point A, the probability of "Return-to-Office" is not equal to zero although its starting point is quite different from the other three goals. This is because the sensor model dominates the recognition result at the beginning since no much historic movement information can be used for smoothing at this point. As a result, it is unreliable to perform goal recognition solely based on the current location estimated from noisy signals. As time moves on, the probabilities of "Seminar-in-Room1" and "Print-in-Room1" increase when the user begins to take the action "Walk-in-HW4" at time point B.

A user’s behavior patterns are often inherently ambiguous. Consider two goals happening in the same location such as "Seminar-in-Room1" and "Print-in-Room1". Since they can neither be distinguished based on the current location

nor on the historic movement, a time-duration variable is introduced to distinguish them. For example, the time of attending a seminar is much longer than that of fetching some material from a printer. Therefore, the probability of "Seminar-in-Room1" is higher than "Print-in-Room1" *only after a certain period of time*. After time point C, the probability of "Seminar-in-Room1" is always the highest. Following (Blaylock & Allen 2003), we refer to time point C as a *convergence point*, and the recognition process after C is considered to *converge*.

To evaluate our experimental results, we use the following evaluation criteria (Blaylock & Allen 2003).

- *Efficiency*: Efficiency is measured in terms of the average processing time for each observation in our on-line goal recognition.
- *Accuracy*: For a certain goal, accuracy is defined as the number of correct recognition divided by the total number of recognition.
- *Convergence rate*: When applied to a specific goal, this criterion indicates the average number of observations, after which a recognized goal converges to the correct answer, over the average number of observations for those traces which converge. It measures how fast the recognition process of a goal converges to the correct answer.

We ran our experiments by doing a three-fold cross-validation over the collected traces. Figure 5 compares the inference efficiency of the whole DBN solution with that of the DBN+Bigram approach. Here *sampling interval* refers to the time interval during which the received signal-strength values are accumulated as one sample. Therefore, the length of traces decrease as the sampling interval increases. As shown in Figure 5, for both the whole DBN solution and the DBN+Bigram approach, the processing time for each observation decreases as the sampling interval increases. However, the DBN+Bigram approach is more efficient than the whole DBN solution for each sampling interval.

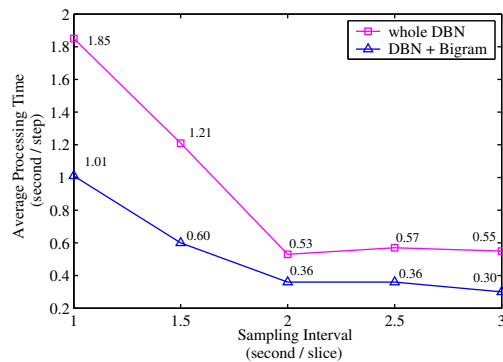


Figure 5: Comparison about the Average Processing Time for Each Observation

Table 2 shows the average recognition accuracy over 19 goals of the whole DBN and DBN+Bigram with respect to the sampling interval. The average accuracy decreases as the sampling interval increases. This is because when the sampling interval increases, the signal-strength values within a

longer period of time will be accumulated into one observation corresponding to one time slice in the trace. This weakens the discriminative power of signals towards different locations, which in turn reduces the recognition accuracy. As can be seen from Table 2, the accuracy of DBN+Bigram is comparable to that of DBN.

Interval (s)	1	1.5	2	2.5	3
Whole DBN (%)	89.5	87.1	84.2	75.4	71.9
DBN+Bigram (%)	90.5	83.2	82.1	74.7	72.6

Table 2: Comparison about the Recognition Accuracy

Table 3 compares the convergence rate of the whole DBN solution with that of the DBN+Bigram approach. Due to the limited space, we only list the convergence rates of three goals and the average convergence rate of 19 goals. As can be seen from Table 3, the average convergence rate of DBN+Bigram is also comparable to that of DBN.

	Whole DBN	DBN+Bigram
Entrance3-to-Office	65.4%	71.5%
Entrance1-to-Entrance3	82.3%	82.0%
Seminar-at-Room2	84.0%	82.2%
Average	75.4%	76.2%

Table 3: Comparison about the Convergence Rate

## Conclusions and Future Work

We addressed the problem of inferring a user's high-level goals from low-level noisy signals in a complex indoor environment using an RF-based wireless network. We first model this problem based on the DBN framework. Then we propose a novel two-level architecture: in the low level, a DBN is used to infer the actions from signals; in the high level, an N-gram is used to infer the user's goals from actions. The experiments demonstrate that the two-level approach is more efficient than the whole DBN solution while the accuracy is comparable.

Our work can be extended in several directions. In this paper we assume that a user carrying out a sequence of actions is only aiming at achieve a single goal. However, a user can accomplish multiple goals within a single sequence of actions. In addition, we also wish to explore how to detect different behavior patterns of multiple users through plan recognition (Devaney & Ram 1998).

## Acknowledgements

The authors are supported by a grant from Hong Kong RGC and the Hong Kong University of Science and Technology. We thank Wong Wing Sing, To Ka Kui and Jin Yuan for helping us to collect the data.

## References

Albrecht, D.; Zukerman, I.; and Nicholson, A. 1998. Bayesian models for keyhole plan recognition in an adventure game. *User Modelling and User-adapted Interaction* 8(1-2):5-47.

Bahl, P., and Padmanabhan, V. N. 2000. RADAR: An in-building RF-based user location and tracking system. In *Proceedings of IEEE INFOCOM2000*, 775-784.

Blaylock, N., and Allen, J. 2003. Corpus-based statistical goal recognition. In *Proceedings of the 8th International Joint Conference on Artificial Intelligence*, 1303-1308.

Charniak, E., and Goldman, R. 1993. A Bayesian model of plan recognition. *Artificial Intelligence Journal* 64:53-79.

Chen, S. F., and Goodman, J. 1996. An empirical study of smoothing techniques for language modeling. In *Proceedings of the Thirty-Fourth Annual Meeting of the Association for Computational Linguistics*, 310-318.

Dean, T., and Kanazawa, K. 1989. A model for reasoning about persistence and causation. *Artificial Intelligence* 93(1-2):1-27.

Demirdjian, D.; Tollmar, K.; Koile, K.; Checka, N.; and Darrell, T. 2002. Activity maps for location-aware computing. In *Proceedings of the IEEE WACV2002*.

Dempster, A. P.; Laird, N. M.; and B. Rubin, D. 1977. Maximum likelihood from incomplete data via EM algorithm. *Journal of the Royal Statistical Society Series B* 39:1-38.

Devaney, M., and Ram, A. 1998. Needles in a haystack: Plan recognition in large spatial domains involving multiple agents. In *Proceedings of AAAI1998*, 942-947.

Fox, D.; Hightower, J.; Liao, L.; and Schulz, D. 2002. Bayesian filtering for location estimation. *IEEE Pervasive Computing* 2(3):24-33.

Hashemi, H. 1993. The indoor radio propagation channel. volume 81, 943-968.

Kautz, H., and Allen, J. F. 1986. Generalized plan recognition. In *Proceedings of AAAI1986*, 32-38.

Ladd, A.; Bekris, K.; Marceau, G.; Rudys, A.; Kavradi, L.; and Wallach, D. 2002. Robotics-based location sensing using wireless ethernet. In *Proceedings of MOBICOM2002*.

Lesh, N., and Etzioni, O. 1995. A sound and fast goal recognizer. In *Proceedings of IJCAI95*, 1704-1710.

Murphy, K. 2002a. *Dynamic Bayesian Networks: Representation, Inference and Learning*. Ph.D. Dissertation, UC Berkeley.

Murphy, K. 2002b. Hidden semi-markov models (HSMMs). Technical report, MIT AI Lab.

Patterson, D. J.; Liao, L.; Fox, L.; and Kautz, H. 2003. Inferring high-level behavior from low-level sensors. In *Proceedings of UBIComp2003*.

Pynadath, D. V., and Wellman, M. P. 2000. Probabilistic state-dependent grammars for plan recognition. In *Proceedings of the Sixteenth Conference on UAI*, 507-514.

Roos, T.; Myllymaki, P.; Tirri, H.; Misikangas, P.; and Sievanen, J. 2002. A probabilistic approach to WLAN user location estimation. *International Journal of Wireless Information Networks* 9(3):155-164.

Youssef, M.; Agrawala, A.; and Shankar, U. 2003. Wlan location determination via clustering and probability distributions. In *Proceedings of IEEE PerCom2003*.