

Bayesian Techniques for Location Estimation

Dieter Fox, Jeffrey Hightower, Henry Kauz, Lin Liao, and Don Patterson

University of Washington, Department of Computer Science & Engineering, Seattle, WA

Keywords

Bayesian filtering, location estimation, uncertainty

1. INTRODUCTION

Location awareness is important to many pervasive computing applications. A fundamental problem in this context is *location estimation*, which is the estimation of a person's location from a stream of sensor data. Since no location sensor takes perfect measurements, it is crucial to represent uncertainty in sensed location information and combine information from different types of sensors. Bayesian filter techniques provide a powerful tool to help manage measurement uncertainty and perform multi-sensor fusion. Their statistical nature makes Bayes filters applicable to arbitrary sensor types and representations of environments. For example, Bayes filters provide a sound approach to location estimation using GPS data along with street maps or signal strength information along with topological representations of indoor environments. Furthermore, they have been applied with great success to a variety of state estimation problems including speech recognition, target tracking, vision, and robotics. In this article, we briefly survey the basics of Bayes filters and their different implementations. Furthermore, we discuss directions for future research in Bayesian techniques for location estimation.

2. BAYESIAN FILTERING

Bayes filters probabilistically estimate the state of a dynamic system from a sequence of noisy sensor observations. In the most basic form of location estimation, the state of interest is the location of a person or object, and observations are provided by sensors either placed in the environment or carried by the person.

2.1 Belief Update

Bayes filters represent the state at time t by random variables x_t . At each point in time, the uncertainty is represented by a probability distribution over x_t called *belief* $Bel(x_t)$. The key idea of Bayes filters is to sequentially estimate such beliefs over the state space conditioned on the information contained in the sensor data. Let us assume that the sensor data consists of a sequence of time indexed sensor observations $z_{1:t}$. The belief $Bel(x_t)$ is then defined by the posterior density over the random variable x_t conditioned on all sensor data available at time t :

$$Bel(x_t) = p(x_t | z_{1:t}) \quad (1)$$

Roughly speaking, the belief provides an answer to the question "What is the probability that the person is at location x

if the history of sensor measurements is $z_{1:t}$?", for all possible locations x . In general, the complexity of computing such posterior densities grows exponentially over time since the number of sensor measurements increases over time. To make the computation tractable, Bayes filters assume the dynamic system is Markov, *i.e.* all relevant information is contained in the current state variable x_t . The update of the Bayes filter is performed in two steps:

Prediction: At each time update, the state is *predicted* according to the following update rule.

$$Bel^-(x_t) \leftarrow \int p(x_t | x_{t-1}) Bel(x_{t-1}) dx_{t-1} \quad (2)$$

Here, the term $p(x_t | x_{t-1})$ describes the *system dynamics*, *i.e.* how the state of the system changes over time. In location estimation, this conditional probability is the motion model – where the person might be at time t , given that she previously was at location x_{t-1} . The motion model strongly depends on the information available to the estimation process. It can range from predicting the next position using estimates of a person's motion velocity to the prediction of when a person will exit the elevator using an estimate of the person's goal.

Correction: Whenever new sensor information z_t is received, the measurement is used to correct the predicted belief using the observation.

$$Bel(x_t) \leftarrow \alpha_t p(z_t | x_t) Bel^-(x_t) \quad (3)$$

$p(z_t | x_t)$, the *perceptual model*, describes the likelihood of making observation z_t given that the person is at location x_t . For location estimation, the perceptual model is usually considered a property of a given sensor technology. It depends on the types and positions of the sensors and captures a sensor's error characteristics. The term α_t in (3) is simply a normalizing constant which ensures that the posterior over the entire state space sums up to one.

$Bel(x_0)$ is initialized with prior knowledge about the location of the person, typically uniformly distributed if no prior knowledge exists. Bayes filters are an abstract concept in that they only provide a probabilistic framework for recursive state estimation. To implement Bayes filters, one has to specify the perceptual model $p(z_t|x_t)$, the dynamics $p(x_t|x_{t-1})$, and the representation of the belief $Bel(x_t)$. The properties of the different implementations of Bayes filters strongly differ in the way they represent probability densities over the state x_t .

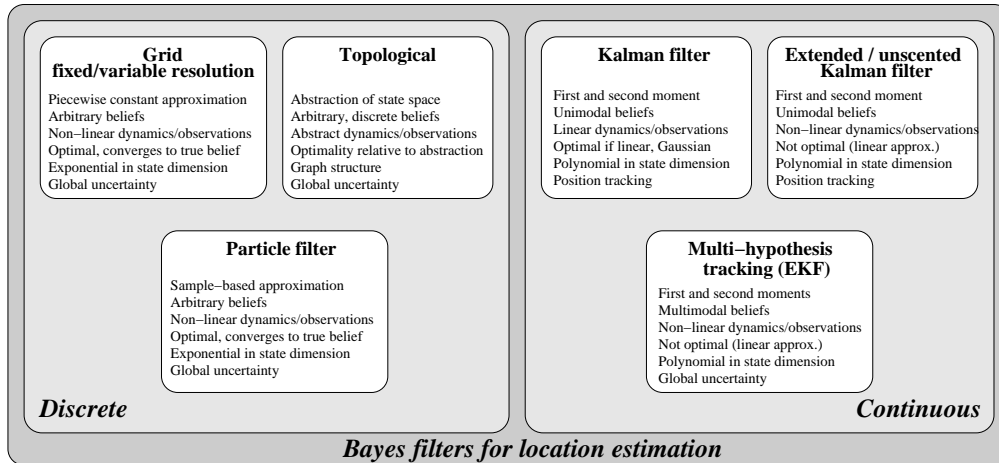


Figure 1: Properties of the most common implementations of Bayes filters for location estimation.

2.2 Belief Representations

This section gives a brief overview of different representations for the beliefs of Bayes filters (see also Figure 1).

Kalman filters are the most widely used variant of Bayes filters [1]. Roughly speaking, these filters approximate beliefs by unimodal Gaussian distributions, represented by their mean and variance. While the mean gives the expected location of the person, the variance represents the uncertainty in the estimate. Even though Kalman filters make strong assumptions about the nature of the sensors and a person’s motion, they have been applied with great success to various estimation problems. The main advantage of Kalman filters is their computational efficiency, which comes at the cost of restricted representational power since Kalman filters can only represent unimodal distributions. Hence, Kalman filters are best if the uncertainty in a person’s location is not too high. Typical sensors used for Kalman filter based estimation are cameras, laser range-finders, and GPS systems.

Multi-hypothesis tracking (MHT) extends Kalman filters to multi-modal beliefs [1]. MHT represent the belief by *mixtures* of Gaussians where each hypothesis is tracked using a Kalman filter. The weights of the hypotheses are determined by how well they predict the sensor measurements. Due to their ability to represent multi-modal beliefs, MHT approaches are more widely applicable than the Kalman filter.

Grid-based approaches overcome the restrictions imposed on Kalman filters by relying on discrete, piecewise constant representations of the belief. For indoor location estimation, grid-based filters tessellate the environment into small patches, typically of size between 10cm and 1m. Each grid cell contains the belief the person is currently in the cell. A key advantage of these approaches is that they can represent arbitrary distributions over the discrete state space. The disadvantage of grid-based approaches is the computational complexity, which makes them applicable to low-dimensional estimation problems only, such as estimating the position and orientation of a person.

The computational complexity of grid-based methods can be

avoided by non-metric representations of an environment. For instance, *graph structures* are well suited to represent the motion of people in buildings [5] or even in cities [8]. Each node in the graph corresponds to a location and the edges describe the connectivity of the environment. The advantage of topological approaches is their efficiency since they represent distributions over small, discrete state spaces. Their disadvantage is the coarseness of the representation which enables only rough information about a person’s location. Topological approaches are typically adequate if the sensors in the environment provide only very imprecise location information.

Particle filters represent beliefs by sets of weighted samples distributed according to the belief [3]. Particle filters realize Bayes filter updates according to a sampling procedure, often referred to as sequential importance sampling with resampling. The key advantage of particle filters is their ability to represent arbitrary probability densities, which makes them applicable to problems for which Kalman filters are not well-suited. Compared to grid-based approaches, particle filters are very efficient since they automatically focus their resources (particles) on regions in state space with high probability. However, since the worst-case complexity of these methods grows exponentially in the dimensions of the state space, one has to be careful when applying particle filters to high-dimensional estimation problems. Recently, Rao-Blackwellised particle filters [2], the combination of particle filters with Kalman filters, have been applied successfully to tracking the locations and identities of multiple people [10].

2.3 Parameter Learning

The parameters of the perceptual and motion models can be learned from data using expectation maximization (EM), a popular approach to parameter estimation from incomplete data [9]. The perceptual model $p(z_t | x_t)$ is typically independent of the person and can be learned beforehand. The motion model, on the other hand, might be different for each person. Learning the parameters of the motion model allows the system to adapt to a specific person, thereby increasing the accuracy and efficiency of the estimation process. For

example, [6] show how to use EM to learn typical motion patterns of a person in indoor environments using a graph-based Bayes filter. [8] use the same technique to learn the navigation patterns of a person through an urban environment.

3. RESEARCH DIRECTIONS

In this section we briefly discuss directions for future research in Bayesian location estimation.

Adaptive Estimation

Most applications of Bayes filters use the same, fixed representation of the state space during the entire estimation process. However, especially in the context of location estimation, this is not appropriate. For example, the location of a person moving through an urban environment can be tracked well using multi-hypothesis tracking along with a GPS sensor and a street-map. However, as soon as the person enters a building, other sensors and representations are needed. Furthermore, even within the same building, different areas might be covered by completely different types of sensors requiring different representations of the belief state. A key question is thus when and how to switch between different representations in a statistically sound way.

High-level Representations

The location of a person provides only very limited information about the person's current activity. Richer representations might include information such as the time of day, the mode of transportation, the destination of the current trip, and the purpose of a specific location. *Dynamic Bayesian networks*, a variant of Bayes filters, provide a sound way of describing and reasoning with such structured, hierarchical information [7]. Some questions remain: What are important locations in a person's life? How can they be described in a general way and learned from sensor data? How can we transfer experience gained from one person to another person? *Relational probabilistic models* [4], which can represent relations between classes of objects, provide a promising framework for addressing these problems.

User Errors

In the context of assisting cognitively impaired people, the detection of when a person seems to be lost is an important aspect of location estimation. *Online model selection* is a technique that can potentially solve this problem. Model selection aims at identifying the model that is best suited to explain the observed data [11]. To apply model selection in the location context, one could generate generic and user-specific Bayes models of activities. Both models are able to track a user's activities, but the specific model 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. 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. If the person fails to exit the

bus at the usual stop, then the general model predicts it with higher probability, thereby triggering the detection of a potential user error. Obviously, such an approach can provide valuable information to user intervention modules.

4. CONCLUSIONS

We presented Bayes filters as a general framework for location estimation, allowing the integration of sensor information over time. The application of Bayes filters goes well beyond location estimation. The generation of hierarchical models allows the seamless integration of location estimation into user activity estimation. We consider Bayesian techniques to be an extremely promising tool for location aware computing.

REFERENCES

- [1] Y. Bar-Shalom and X.-R. Li. *Multitarget-Multisensor Tracking: Principles and Techniques*. Yaakov Bar-Shalom, 1995.
- [2] 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.
- [3] A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte Carlo in Practice*. Springer-Verlag, New York, 2001.
- [4] N. Friedman, L. Getoor, D. Koller, and A. Pfeffer. Learning probabilistic relational models. In *IJCAI*, 1999.
- [5] J. Krumm, L. Williams, and G. Smith. SmartMoveX on a graph – an inexpensive active badge tracker. In *International Conference on Ubiquitous Computing (UbiComp)*, 2002.
- [6] L. Liao, D. Fox, J. Hightower, H. Kautz, and D. Schulz. Voronoi tracking: Location estimation using sparse and noisy sensor data. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2003.
- [7] K. Murphy. *Dynamic Bayesian Networks: Representation, Inference and Learning*. PhD thesis, UC Berkeley, Computer Science Division, 2002.
- [8] D. Patterson, L. Liao, D. Fox, and H. Kautz. Inferring high-level behavior from low-level sensors. In *International Conference on Ubiquitous Computing (UbiComp)*, 2003.
- [9] L. R. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. In *Proceedings of the IEEE*. IEEE, 1989. IEEE Log Number 8825949.
- [10] D. Schulz, D. Fox, and J. Hightower. People tracking with anonymous and id-sensors using Rao-Blackwellised particle filters. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2003.
- [11] M. West and P.J. Harrison. *Bayesian Forecasting and Dynamic Models*. Springer-Verlag, New York, 2nd edition, 1997.