A Learning Algorithm for Localizing People Based on Wireless Signal Strength that Uses Labeled and Unlabeled Data

Mary Berna, Brennan Sellner, Brad Lisien, Sebastian Thrun, Geoffrey Gordon, Frank Pfenning
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA

Abstract

This paper summarizes a probabilistic approach for localizing people through the signal strengths of a wireless IEEE 802.11b network. Our approach uses data labeled by ground truth position to learn a probabilistic mapping from locations to wireless signals, represented by piecewise linear Gaussians. It then uses sequences of wireless signal data (without position labels) to acquire motion models of individual people, which further improves the localization accuracy. The approach has been implemented and evaluated in an office environment.

1 Introduction

This paper addresses the problem of localizing people using a wireless network. Various researchers have recently developed techniques using the signal strength of IEEE 802.11b access points for people localization [1; 3; 6; 9; 11]. In doing so, these techniques require measurement data labeled by the position at which the data was acquired.

Our approach builds on this idea, but augments it by a learned probabilistic motion model of people. The data used for training the motion model is acquired as people walk through the environment. We employ the EM algorithm to train the resulting hidden Markov model.

2 Signal Strength Maps From Labeled Data

At the most basic level, our approach acquires a map of the wireless signal strength at different locations in the environment. For that, it requires the availability of signal strength measurements labeled by the true position. Our implementation represents this map by a piecewise linear Gaussian, that follows the (mostly) one-dimensional manifold of the center of each corridor in the building. Figure 1 illustrates the signal strength measurements for one of the access points in our test building. The measurement noise is modeled as a Gaussian, whose covariance is estimated from the data using the maximum likelihood estimator.

3 Probabilistic Localization

The signal strength map is sufficient to coarsely localize people. Following [6], our approach implements a continuous-state hidden Markov model (HMM), in which the person's position is the internal state, and the signal strengths are the measurements. The signal strength map provides the measurement probabilities of the HMM. The next state transition, or motion model, may be as simple as a probabilistic model of Brownian motion; a more sensible choice will be discussed below.

Our approach uses this HMM to track the location of the person, by calculating an approximate posterior. This is achieved using the well-known Monte Carlo localization algorithm [10], a variant of particle filters [5].

4 Motion Models From Unlabeled Data

The principal limitation of the approach thus far is its reliance on labeled data. Such data is difficult to collect, as it requires an independent means for generating correct position labels when carrying a wireless receiver through a building. Our approach specifically addresses the use of unlabeled data. Such data is acquired as people walk through the building.

This unlabeled data consists of signal measurements *without* a pose estimate. We use this data to learn a motion model of a person. This motion model is realized by mixtures of Gaussians that characterize the relative change of position within a fixed time interval. Each linear sub-piece of the signal strength map possesses its own Gaussian mixture; in this way, the approach can learn location-specific predictions of people's motion. The motivation to use a Gaussian mixture model arises from the fact that people often engage in one out of a small number of different motions (e.g., turning left or right at an intersection; see Figure 2). In our approach, the number of Gaussians is fixed, as is their variance. Our approach learns the means of the Gaussians and the mixture weights as a function of people's location.

Learning from unlabeled data is achieved using the EM algorithm (see [2; 7; 8] for related work). Beginning with an initial motion model, position probabilities are computed based on the signal strength, using the particle filter described above. These position probabilities are then used to calculate a distribution over the mixture components of the Gaussian, and subsequently to calculate the mean and new mixture weights. This methodology is a straightforward application of the EM algorithm, assuming that both a person's pose and the index of the Gaussian mixture are latent variables [4]. In our implementation, we bias the initial motion model (before learning) towards typical motions people might take, such as going forward, turning left or right, and standing in place.

5 Experimental Results

Systematic experimental results were conducted in an office environment equipped with approximately 50 802.11b wireless access points, distributed over an area of size 120 by 22.5

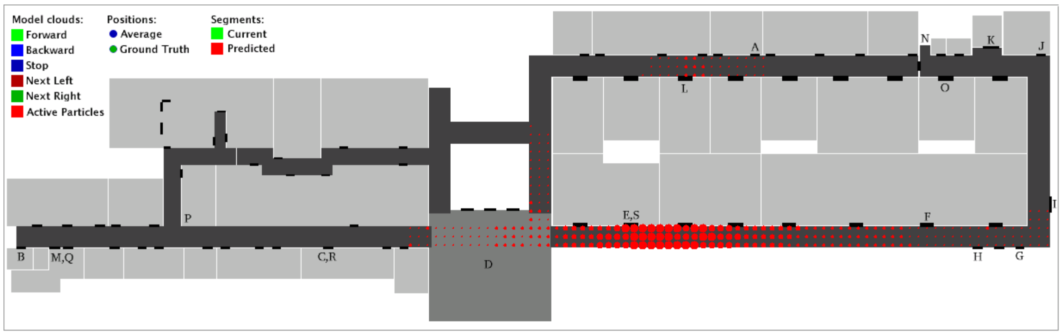

Figure 1: Access point strength as a function of measured location.

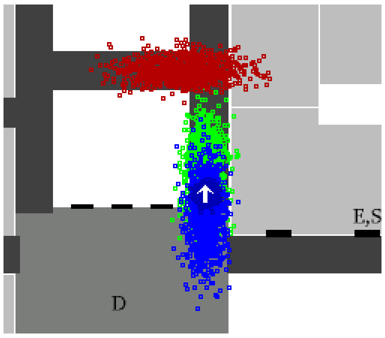

Figure 2: An example motion model. Clearly visible are some of the modes induces by the different actions a person may take.

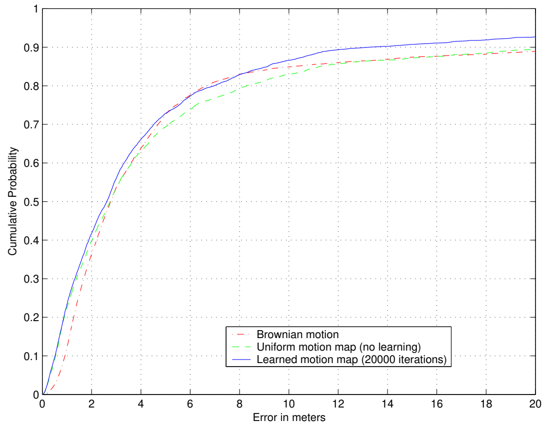

Figure 3: Cumulative probability versus error in meters.

meters. On average, five access points are within communication range of each location in the environment.

All labeled data was acquired by using a laser range finder for localization, similar to work in mobile robotics [10]. A set of labeled data was collected by simply walking about the environment, traversing each corridor at least three times in each direction (along the center, left, and right of the hallway). We also collected a large set of unlabeled training data by traversing paths characteristic of the movements of an individual (all paths originated or terminated at a specific office). Finally, we collected an independent test set with labels.

Figure 3 shows a key result, obtained for 200 particles. This figure shows the cumulative probability function over the average localization error from our labeled test set for different techniques. The top blue curve corresponds to the combined labeled-unlabeled data approach described above, utilizing the learned mixture models. The green dashed curve reflects the performance of the same system with no learning; all mixture weights are uniform. The red dot-dashed curve represents the performance of a pure Brownian motion model, rather than a mixture of Gaussians.

A localization error of 2.25 meters is achieved 70% of the time when only using the labeled data. This number appears to be well in tune with results reported in [6] while requiring a factor of 50 less labeled training data. This error is reduced by 20% as the motion model is trained with the unlabeled data. This illustrates the relative improvement that can be achieved by adding unlabeled measurements to the pool of training data.

References

[1] P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *INFOCOM (2)*, pages 775–784, 2000.

[2] M. Bennewitz, W. Burgard, and S. Thrun. Using EM to learn motion behaviors of persons with mobile robots. *IROS-02*.

[3] P.K. Chrysanthis. On indoor position location with wireless LANs. *ISPIRC-02*.

[4] A.P. Dempster, A.N. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1), 1977.

[5] A. Doucet, J.F.G. de Freitas, and N.J. Gordon, editors. *Sequential Monte Carlo Methods In Practice*. Springer, 2001.

[6] A. M. Ladd, K.E. Bekris, G. Marceau, A. Rudys, D. Wallach, and L. Kavraki. Using wireless ethernet for localization. *ISPIRC-02*.

[7] K. Nigam, A. McCallum, S. Thrun, and T. Mitchell. Learning to classify text from labeled and unlabeled documents. *Machine Learning*, 39(2/3), 1998.

[8] B. Shahshahani and D. Landgrebe. The effect of unlabeled samples in reducing the small sample size problem and mitigating the Hughes phenomenon. *IEEE Trans. on Geoscience and Remote Sensing*, 32(5), 1994.

[9] J. Small, A. Smailagic, and D. Siewiorek. Determining user location for context aware computing through the use of a wireless LAN infrastructure. http://www.cs.cmu.edu/∼aura/docdir/small00.pdf, 2003.

[10] S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust monte carlo localization for mobile robots. *Artificial Intelligence*, 128(1-2), 2000.

[11] M. Wallbaum. Wheremops: An indoor geolocation system. *ISPIRC-02*.