

Tracking a Moving Target in Cluttered Environments with Ranging Radios: Extended Results

Geoffrey Hollinger

Joseph Djughash

Sanjiv Singh

CMU-RI-TR-08-07

February 2008

Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

© Carnegie Mellon University

Abstract

In this technical report, we propose a framework for utilizing fixed, ultra-wideband ranging radio nodes to track a moving target node through walls in a cluttered environment. We examine both the case where the locations of the fixed nodes are known as well as the case where they are unknown. For the case when the fixed node locations are known, we derive a Bayesian room-level tracking method that takes advantage of the structural characteristics of the environment to ensure robustness to noise. We also develop a method using mixtures of Gaussians to model the noise characteristics of the radios. For the case of unknown fixed node locations, we present a two-step approach that first reconstructs the target node’s path and then uses that path to determine the locations of the fixed nodes. We reconstruct the path by projecting down from a higher-dimensional measurement space to the 2D environment space using non-linear dimensionality reduction with Gaussian Process Latent Variable Models (GPLVMs). We then utilize the reconstructed path to map the locations of the fixed nodes using a Bayesian occupancy grid. We present experimental results verifying our algorithm in an office environment, and we compare our results to those generated by methods using Extended Kalman Filter (EKF) SLAM and odometry mapping. Our algorithm is successful at tracking a moving target node and mapping the locations of fixed nodes using radio ranging data that are both noisy and intermittent.

Contents

1	INTRODUCTION	1
2	RELATED WORK	2
3	TRACKING WITH KNOWN NODE LOCATIONS	3
3.1	Room-Level Tracking	3
3.2	Ranging Radio Noise Modeling	4
3.2.1	Simple Gaussian Modeling	5
3.2.2	Gaussian Process Modeling	5
3.2.3	Mixture of Gaussians Modeling	6
4	TRACKING WITH UNKNOWN NODE LOCATIONS	7
4.1	Path Reconstruction with GPLVMs	7
4.2	Recovering Node Locations with Occupancy Grids	9
5	EXPERIMENTAL RESULTS	9
5.1	Hardware Setup	9
5.2	Results with Known Node Locations	10
5.3	Results with Unknown Node Locations	11
6	CONCLUSIONS AND FUTURE WORK	14
7	ACKNOWLEDGMENTS	15

1 INTRODUCTION

Our goal is to track a target moving in an environment like an office building without the requirement of pre-installed infrastructure or accurate odometry. As a sub-goal, we seek to map the unknown locations of sensors in the environment so they can be utilized for future tracking. The problem of tracking a moving target in a cluttered environment is one that is prevalent in many robotics applications. In the dynamic world of mobile robotics, rarely do targets remain stationary, but often we can rely on some motion model or odometry information from the target to assist in tracking.

The specific application of tracking a human in an indoor environment is particularly challenging because human targets often do not have reliable odometry. Furthermore, human motion is often erratic and difficult to predict using a simple motion model. Wearable inertial measurement units (IMUs) are either inaccurate, expensive, or bulky. Even industrial grade IMUs inevitably drift after extended operation.

Alternatively, if the human carries a ranging radio, fixed radio nodes can provide sensor measurements for tracking as well as anchors into the environment that prevent drift. The locations of these beacons can either be surveyed as part of a pre-installed infrastructure, or they can be determined as part of a tracking algorithm. The techniques that we discuss in this paper examine both of these possibilities.

In this paper, we present a framework for tracking a moving target in a cluttered environment using range measurements from ultra-wideband radios. We examine two variations of the tracking problem:

1. The locations of the radio nodes are known a priori
2. The locations of the radio nodes are unknown

For the first scenario, we discuss a room-level tracking approach that uses a discretized version of the floor plan to take advantage of structural characteristics of an indoor environment. We use a Bayesian filter on the discretized floor plan to track the target between rooms. We present two methods for modeling noise in the range measurements using either Gaussian Processes or mixtures of Gaussians. For the scenario where the node locations are unknown, we derive a two-step method that first reconstructs the path of the target using non-linear dimensionality reduction with Gaussian Process Latent Variable Models (GPLVMs) and then uses the reconstructed path to determine the locations of the nodes on a Bayesian occupancy grid. For both scenarios, we assume that a floor plan of the environment is known a priori.

We present performance results of our algorithms in office environments of varying sizes for both the known and unknown node location scenarios. The novelties of this paper include the application of Bayesian room-level tracking techniques to ranging radio data, the use of GPLVMs with ranging radios, and the development of a two-step tracking and mapping method using dimensionality reduction and Bayesian occupancy grid mapping.

This paper is organized as follows. Section 2 discusses related work in the areas of tracking, simultaneous localization and mapping (SLAM), and dimensionality reduction. Sections 3 and 4 present our framework for utilizing range-only measurements with known and unknown radio

node locations. Section 5 discusses results from experiments with ranging radios in an office environment. Finally, Section 6 draws conclusions and discusses directions for future work.

This technical report is an extended version of a conference paper on the same topic [1]. This version contains more explanatory detail and further results verifying the performance of our algorithm.

2 RELATED WORK

The framework that we develop in this paper for localization and mapping using ranging radios is closely related to literature in tracking, simultaneous localization and mapping, and dimensionality reduction. Kumar, Singh, and Rus discussed the problem of tracking a human first responder in an urban search and rescue scenario with robots and sensor networks [2]. They outline some open questions and provide the motivation for using range-only devices for human tracking. Liao et al. extended these ideas by developing a tracking algorithm using range-only measurements and particle filters [3]. This work does not examine the noise characteristics of ranging radios, and it does not discuss techniques for when the locations of the nodes are unknown.

Our room-level tracking techniques were originally introduced in our previous work as part of a framework for searching a cluttered environment for a mobile target [4]. Our previous work describes an algorithm that incorporates both coordination and tracking. In this paper, we extend our tracking sub-algorithm by modeling the noise characteristics of ranging radios and deriving a method for reconstructing unknown node locations. As a result, our techniques using ranging radios can easily be incorporated into the coordinated search framework developed in previous work.

When the locations of the nodes are unknown, the tracking problem is closely related to simultaneous localization and mapping (SLAM). Djugash et al. proposed a range-only SLAM method using Extended Kalman Filters (EKF) [5]. This EKF approach projects range measurements into polar space and uses a multi-modal representation to avoid errors from poor initialization, ambiguities, and noise. Despite these innovations, EKFs do not handle outliers as well as techniques using dimensionality reduction because they require linearization, and they are prone to error when odometry is poor or nonexistent. In Section 5, we run this EKF-SLAM algorithm on our data and compare the results to those obtained by our proposed algorithm.

In an application similar to ours, Olson et al. presented a method for locating ranging sonar beacons with an Autonomous Underwater Vehicle (AUV) [6]. They discuss a spectral graph partitioning method for outlier rejection, and they use a discretized voting scheme to approximately locate beacons before running an EKF. Their beacon initialization method relies heavily on the dead-reckoning (odometry) data from the AUV. In our application, reliable odometry is not available to assist in beacon initialization. In Section 5, we provide a comparison of our method to a method similar to theirs that reconstructs the node locations using robot odometry.

Researchers in the sensor network community have also worked on localizing networks containing both moving and stationary nodes. Priyantha et al. proposed a method for coordinating a mobile node to best localize a range-only network [7]. Their method does not use a probabilistic formulation, which makes it poorly suited for sensors with nonstandard noise models. Their application is also different from ours in that they have control over the mobile node. Hu and Evans also

discussed localizing sensor networks that contain moving nodes [8]. Their work is innovative in that it shows that moving nodes can help localize a sensor network, but they do not directly apply their method to range-only data. We present a fully probabilistic method that focuses on tracking a mobile node in a range-only sensor network. Our method is able to simultaneously locate stationary nodes (if necessary) for later use in tracking.

Recent research has explored the use of Gaussian Processes for modeling the noise characteristics of non-linear sensors. Ferris, Hahnel, and Fox looked at tracking humans in office environments using measurements of wireless signal strength [9]. Schwaighofer et al. also applied Gaussian Processes with the Matern kernel function for localization using cellular phone signal strength [10]. We present results for modeling ranging radios with Gaussian Processes. We also derive a mixture of Gaussians modeling technique that approximates the Gaussian Process solution and allows for outlier removal.

Without reliable odometry, recreating a path from range measurements becomes a non-linear dimensionality reduction problem. Gaussian Process Latent Variable Models (GPLVMs) were introduced by Neil Lawrence as a probabilistic framework for non-linear dimensionality reduction [11, 12]. GPLVMs were later extended by Wang, Fleet, and Hertzmann to incorporate dynamics with applications to human motion modeling [13]. Modeling dynamics allows for the incorporation of simple motion models into the GPLVM framework.

Ferris, Fox, and Lawrence applied GPLVMs to solve the problem of localization with wireless signal strength when training data is unavailable [14]. Their algorithm takes advantage of the above tools in a target tracking scenario. Our work goes one step further by using the reconstructed path to map the locations of ranging radio beacons. This step may not be possible with wireless signal strength due to difficulty in modeling their noise characteristics.

Ranging radio tracking has both advantages and disadvantages in comparison to wireless signal strength tracking [9, 14]. Data from wireless signal strength are often erratic and can be affected by slight changes in the environment. Ranging radio measurements more closely follow the true range between nodes. This allows us to reconstruct the positions of the radio nodes more easily than one would reconstruct the positions of wireless access points. However, wireless access points are already part of the infrastructure of many buildings. Radio nodes would need to be placed in a building before use in tracking.

3 TRACKING WITH KNOWN NODE LOCATIONS

3.1 Room-Level Tracking

To solve the problem of tracking a moving target in a cluttered environment, we present a room-level approach on a discretized version of the floor plan. In an indoor environment, maps often have implicit discretizations into rooms and hallways. In the case of tracking a moving target in these scenarios, it is often advantageous to know in which room the target is located. Drawing from these observations, we discretize the floor plan into convex rooms and hallways. This can be done either by hand or by arbitrarily collapsing regions found using a convex region finding algorithm. Figure 1 shows an example floor plan used in our experiments in Section 5.




Figure 1: Discretization of office environment used for ranging radio tracking. The discretization uses structural aspects of the built environment to provide a room-level tracking with very noisy measurements.

Taking into account cell adjacencies in the environment forms an undirected graph of the target’s movement options. Now, define a probability distribution $p_t(C)$ over the target’s possible location at time t in m cells where $\{c_1, c_2, \dots, c_m\} \in C$. The resulting probability distribution in addition to the cell adjacency matrix forms a Markov Chain. The target’s motion can now be modeled on the Markov Chain using diffusion matrices. For any time $t + 1$, the target’s predicted location is a function of the diffusion matrix D and the distribution at time t :

$$p_{t+1}^-(C) = p_t(C)D. \quad (1)$$

The diffusion matrix allows for flexible modeling of target motion, which can easily account for speed changes, room size, and other factors.

Having properly modeled the target’s motion, we now incorporate information from range-only measurements. In a Bayesian framework, the posterior distribution is given by:

$$p_t^+(c) = \eta p(z|c)p_t^-(c), \quad (2)$$

where $p(z|c)$ is the probability of receiving range measurement z given that the target is in cell c , and η is a normalizing constant.

Since the discretization of the environment into cells is very coarse, it is often advantageous to calculate $p(z|c)$ at a finer resolution. For this purpose, we more finely divide each cell c into subcells $c^b \in c$. We then calculate the probability $p(z|c^b)$ at each subcell and recalculate the larger probability by summing over these subcells.

3.2 Ranging Radio Noise Modeling

Ultra-wideband ranging radios operating in the 6GHz+ frequency band can provide range measurements through walls in indoor environments [15]. Ultra-wideband radios show better accuracy than alternative ranging devices in cluttered environments because their spread spectrum is more likely to find a direct path from transmitter to receiver.

Our experiences show that the environment has a large effect on the noise characteristics of ranging radios. When ranging radio signals move through occlusions, the peak in the signal becomes less pronounced. In a time-of-arrival system, this often leads to the peak being detected

after it actually occurred. This creates a tendency towards measurements that are longer than the actual range between targets, and this elongation tends to increase as the signal travels through more occluded space. Thus, the noise characteristics are non-linear and are very difficult to model using simple techniques. This section describes some methods to model this noise and to estimate the target’s location at the cell level.

3.2.1 Simple Gaussian Modeling

To determine the probability that a target is in cell c given a measurement z , it is necessary to calculate the likelihood $p(z|c)$ as described above. If one assumes that the noise is Gaussian, noting that this is often not the case with ranging radios, the likelihood can be calculated as in Equation 3:

$$p(z|c) = N(z; r_c, \sigma_r^2), \quad (3)$$

where z is the observed range value, r_c is the true range from cell c to the ranging node, and σ_r^2 is the variance of the noisy range measurement.

The range measurement variance σ_r^2 must be estimated from the data, but this often can be done with very little calibration data. To calculate the true range r_c for a cell c , that cell must be reduced to a point. Methods for doing this include using the cell’s centroid or subdividing the cell into a fine grid of cells c^b and then summing the likelihoods over these subcells [4]. We take the latter approach due to the coarseness of our room-level cell discretization.

3.2.2 Gaussian Process Modeling

If calibration data is available for the environment, we can use a learning method to estimate $p(z|c)$. As described above, the noise characteristics of ranging radios are often non-linear. Gaussian Processes offer a non-parametric Bayesian solution to modeling non-linear noise given training data. Our formulation of Gaussian Processes for ranging radio noise estimation closely follows that given by Ferris et al. [9, 14].

A Gaussian Process models a noisy process with an underlying noise model as below:

$$z_i = f(x_i) + \varepsilon. \quad (4)$$

We are given some training data of the form $D = [(x_1, z_1), (x_2, z_2), \dots, (x_n, z_n)]$ where $x_i \in \mathfrak{R}^d$ and $z_i \in \mathfrak{R}$. In the case of ranging radios, x_i is a point in the 2D plane ($d = 2$), and z_i represents a range measurement from a single node to this point.¹ Since z_i is a measurement of range between nodes, we have a strong model that z_i should follow. To utilize this, we subtract off the true range r_{x_i} from all observed measurements z_i^o :

$$z_i = z_i^o - r_{x_i}. \quad (5)$$

Subtracting off the range offset allows the Gaussian Process to learn the deviation from the true range rather than learning the underlying range function. Note that it is necessary to know the positions of the nodes to determine r_{x_i} . We relax this constraint in Section 4.

¹Note that this necessitates learning separate Gaussian Processes for each fixed node.

For n training points, refer to the $n \times d$ matrix of x_i values as X and the $n \times 1$ vector of z_i values as Z^q . Note that there is a Z^q vector for each of the Q nodes (used to learn separate Gaussian Processes). The next step in defining a Gaussian Process is to choose a covariance function to relate points in X . We choose the commonly used squared exponential kernel:

$$\text{cov}(f(x_i), f(x_j)) = k(x_i, x_j) = \sigma_s^2 \exp\left(-\frac{1}{2l^2}|x_i - x_j|^2\right), \quad (6)$$

where σ_s^2 and l are hyperparameters.²

Combining the covariance values for all points into a matrix K and adding a Gaussian observation noise hyperparameter σ_o^2 yields:

$$\text{cov}(Z^q) = K + \sigma_o^2 I. \quad (7)$$

We now wish to predict the function value $f(x_*)$ at a selected point x_* given the training data:

$$p(f(x_*)|x_*, X, Z^q) = N(f(x_*); \mu_{x_*}, \sigma_{x_*}^2), \quad (8)$$

where

$$\mu_{x_*} = k_*^T (K + \sigma_o^2)^{-1} Z^q \quad (9)$$

$$\sigma_{x_*}^2 = k(x_*, x_*) - k_*^T (K + \sigma_o^2)^{-1} k_*, \quad (10)$$

k_* is the covariance vector between the selected point x_* and the training inputs X .

When estimating the likelihood of a measurement z_* at a selected point x_* , we must also add the observation noise:

$$p(z_*|x_*, X, Z^q) = N(z_*; \mu_{x_*}, \sigma_{x_*}^2 + \sigma_o^2). \quad (11)$$

If x_* is replaced by a point in a finely discretized subcell $c^b \in c$, Equation 11 represents the likelihood $p(z|c^b, X, Z^q)$. This can now be used to fold in information from measurements in our room-level tracking framework.

3.2.3 Mixture of Gaussians Modeling

A mixture of Gaussians model when coupled with a filtering algorithm offers another alternative to modeling the sensor noise in the presence of calibration data. We recursively refine a mixture of Gaussians noise model of the sensor as calibration data is fed into the system sequentially. In this approach, the expected value of an observed range measurement z^o is as described below:

$$\hat{z}^* = \left(\sum_i^{\|\Gamma\|} \omega_i s_i \right) * r_q + \varepsilon, \quad (12)$$

where r_q is the true range between node q and the target, ε is zero mean Gaussian noise with variance σ , ω_i is the weight of the mixture Gaussian i , and Γ is the set of s_i , each representing a different Gaussian used within the mixture model. The s_i terms scale the true range measurements

²Hyperparameters for each GP can be determined by maximizing the log likelihood of the measurements given the locations and the hyperparameters.

to model the expected mean offset within the observed range measurements. The term \hat{z}^* is a mixture of Gaussians, each of which has a scaled mean around the true expected range r_q . This model works under the assumption that while the observed range measurements can be scaled when compared to the true measurements, they still contain a noise term that is independent of the scaling offset.

The weights ω_i can be updated recursively within each iteration as described below:

$$\omega_i^+ = \eta \omega_i^- \exp\left(-\frac{(\omega_i^- r_q - z^o)^2}{2\sigma^2}\right), \quad (13)$$

where η is a normalization term, which adjusts the weights ω_i such that $\sum_i^{\|\Gamma\|} \omega_i = 1$.

One particular benefit we gain from this model is ease of incorporating a filtering algorithm to reject outliers within the calibration data. We can do this easily using a chi-squared test as a measurement validation gate. Intuitively, it can be expected that outliers will be occasional and more importantly uncorrelated. If recent measurements to a specific node are similar, we will be confident about the expected measurement, and the gate should uphold a higher criteria for incoming measurements. Using this method, we compute the normalized innovation squared as:

$$\epsilon_\nu = \nu^T \sigma^{-1} \nu, \quad (14)$$

where ν is the innovation (i.e., the difference between the measurement, z^o , and its expected value, \hat{z}^*) and σ is the variance of the accepted measurements. The term ϵ_ν has a chi-square distribution, so the gates bounding values can be read from the chi-square table, and a measurement can be discarded if its ϵ_ν value is outside the bounding values. The measurement gating step provides increased robustness to outliers when calculating the mixture of Gaussian model.

After calibration, when we receive a new measurement, the likelihood of the measurement can be computed as follows:

$$p(z|c, X, Z) = N\left(z; \left(\sum_i^{\|\Gamma\|} \omega_i s_i\right) * r_c, \sigma_r^2\right). \quad (15)$$

4 TRACKING WITH UNKNOWN NODE LOCATIONS

We now present a tracking method that relaxes the constraint that the locations of the nodes must be known a priori and does not require calibration data. This tracking scenario would occur if one entered an environment without pre-installed infrastructure. In this section, we assume that we are given an ordered $n \times Q$ matrix Z of n measurements from a moving target to a number of ranging nodes Q . We are not given any information regarding the locations of those nodes in the environment, nor are we given a vector of ground truth locations that correspond to these measurements.

4.1 Path Reconstruction with GPLVMs

Given a matrix Z of range measurements, it is straightforward to frame the reconstruction the target's path X^r as a dimensionality reduction problem. The problem becomes one of projecting

from the Q dimensional data space to the A dimensional latent space. In our case, Q is the number of radio nodes, and A is two (the target’s path is in \mathfrak{R}^2). Since the measurements from the ranging radios are highly non-linear, we need a method that handles these non-linearities. Also, we wish to utilize the information that Z is ordered, and the points corresponding to Z_t and Z_{t+1} are near each other in latent space. This is the problem of incorporating dynamics.

Gaussian Process Latent Variable Models (GPLVMs) provide a probabilistic method for non-linear dimensionality reduction [11, 12]. GPLVMs have also been extended to incorporate dynamics [13], and they have been used to reconstruct 2D paths using wireless signal strength [14]. Below, we show how GPLVMs can be used to reconstruct the path of a target using information from ranging radios.

Data from ranging radios are often sparse, and we do not receive measurements from each radio at every time step. In fact, it is often the case that a radio reading might be absent when traveling through an entire hallway. To prevent missing data in Z , we linearly interpolate across time steps. This was shown in previous work to be an effective strategy for range-only SLAM [16].

Bayesian dimensionality reduction consists of maximizing both the marginal likelihood of the observations $p(Z|X^r)$ and a prior on the underlying positions $p(X^r)$. For a matrix Z with Q columns:

$$p(Z|X^r) = \prod_{q=1:Q} N(Z^q; 0, K + \sigma_o^2 I), \quad (16)$$

where K is the covariance matrix as described above.

The prior value $p(X^r)$ can be used to model the dynamics in an ordered data stream. This can be done using the standard auto-regressive prior [13] or using a more specialized prior using distance and orientation constraints [14]. Our results in Section 5 use the standard auto-regressive prior, and we leave the derivation of a more informed prior to future work.

Having defined both $p(Z|X^r)$ and $p(X^r)$, the values for X^r can be recovered by running conjugate gradient ascent on:

$$p(X^r, Z) = p(Z|X^r)p(X^r) \quad (17)$$

The resulting path defined by X^r is locally consistent. To recover a globally consistent path, the values may need to be rotated or flipped along an axis. We assume that the locations of two points on the path are known, and we use these points to rotate into a global frame consistent with our environment map. Knowing these points would be as simple as knowing when the target entered a building and when it passed a landmark midway through the run.

We also found that, despite modeling dynamics, the scale of the reconstructed path was often off, and we use the two known points to readjust scale. It is important to note that knowing two points on the target’s path is *not* equivalent to knowing the locations of two stationary nodes. Even if the entire target path were known, locations of all nodes in the environment would still need to be reconstructed from noisy ranging data.

Using a GPLVM to reconstruct the path gives a full Gaussian Process model of the sensor noise, $p(X^r, Z)$. This model includes hyperparameters, which provide an estimate of the Gaussian noise variance σ_r^2 . We utilize this estimate when recovering node locations in the next step.

4.2 Recovering Node Locations with Occupancy Grids

Having reconstructed a globally consistent path X^r , our next step is to reconstruct the radio node locations L^r . Given a reconstructed path X^r and a corresponding vector of ranging measurements Z , we estimate the locations of each node using a Bayesian occupancy grid approach [17]. We finely discretize the region in \mathbb{R}^2 in which node q could be located into a grid X_{occ}^q .³ Now, step along the path X^r calculating the following at each cell:

$$p(x_{occ}^q) = \prod_{t=1:n} N(z_t^q; |x_{occ}^q - x_t^r|, \sigma_r^2), \quad (18)$$

where n is the number of range measurements from node q to the target, z_t^q is the range measurement from node q at time t , $|\cdot|$ is Euclidean distance, and σ_r^2 is a noise estimate for the radio sensors (estimated from the GPLVM hyperparameters). Having calculated $p(x_{occ}^q)$ for all cells in X_{occ}^q , we find the location of node q by setting $l_q^r = \max_x p(x_{occ}^q)$.

Once the node locations have been reconstructed on the Bayesian occupancy grid, we can construct a Kalman update to incorporate internode measurements. The Kalman update adjusts the reconstructed locations so that internode measurements are closer to those measured. In practice, we found that incorporating internode measurements did not lead to much improvement, and we direct the reader our previous work for the formulation of the range Kalman update [18].

After reconstructing the node locations, we can combine them with the path estimate and range measurements. This yields enough information to calculate any of the noise models presented in Section 3, which can then be utilized for future tracking.

5 EXPERIMENTAL RESULTS

5.1 Hardware Setup

We setup a test environment using a Pioneer robot and five radio nodes to examine the performance of our tracking algorithms. The Pioneer carried a radio node and acted as the target in these trials. The Pioneer also carried a SICK laser rangefinder, and a map of the environment was found using laser AMCL-SLAM methods from the Carmen software package [17]. Laser localization with the map was used for ground truth comparison as well as acquiring training data. The odometry of the robot was used to generate the ground truth, but it was thrown out for GPLVM tracking experiments. This models the case of a human target without odometry. The robot moved at a speed of approximately 0.2 m/s during the experiments.

We utilized five ultra-wideband radio beacons from Multispectral Solutions to provide sensor measurements [15]. These sensors use time-of-arrival of ultra-wideband signals to provide internode ranging measurements through walls. They propagate these measurements back to a base node using wireless communication. They are setup to operate continuously, and a full set of measurements between five nodes is received approximately every five seconds. In our experiments, we found that the Multispectral radio nodes have an effective range of approximately 30 m when

³If more information is known about where the nodes are located, then an informed initial distribution can also be used.

ranging through walls. Four radio nodes were placed around the environment, and one was placed on the Pioneer robot. Figure 2 shows a photograph of the Pioneer robot with a mounted ranging radio.



Figure 2: Photograph of Multispectral ultra-wideband ranging radio mounted on Pioneer robot. The robot was teleoperated around the environment to act as the moving target.

5.2 Results with Known Node Locations

We tested our methods for tracking with known fixed node locations (Section 3) in an office environment as shown in Fig. 1. We first ran a calibration test to estimate the variance σ_r^2 of the ranging radios. We estimated the variance with a simple Gaussian to be $\sigma_r^2 = 3.85 m^2$. Running a mixture of Gaussians, we found that a single Gaussian with an offset modeled the data well. Fig. 3 shows the offset Gaussian fit for the smaller loop. Applying this offset to that data yielded $\sigma_r^2 = 1.03 m^2$ showing that the offset significantly reduces the measurement variance. These high variances demonstrate the noisiness of non-line-of-site ranging sensors, which makes tracking with them a challenging problem.

Table 1 shows tracking results with known node locations in both a $45 m \times 30 m$ and a $60 m \times 30 m$ office building loop. Each result is averaged over two separate trials (not including the training trial). All methods estimate that the target is either in the correct cell or in an adjacent cell over 90% of the time.

The results with known node locations show that Gaussian Process modeling improves room-level tracking accuracy slightly over the simple Gaussian method. This improvement would likely be more significant if the original measurement variance were larger than the size of most cells in the environment. The offset Gaussian derived using a mixture of Gaussians outperforms both the simple Gaussian and the Gaussian Process modeling methods. The offset Gaussian is particularly useful with ultra-wideband ranging radios because of their bias towards measurements longer than the true range (see Fig. 3). The tracking methods all had a tendency to incorrectly estimate the target's cell when erratic measurements were received from multi-path effects or signal interference when passing through metal objects in walls. These cases tended to produce several erroneously long measurements in a row, which our methods were not able to properly model or filter out.

Table 1: Tracking accuracy with known node locations in an office environment. Format: (xx/yy), where xx is percentage of estimates in correct cell (room or hallway), yy is percentage of estimates in correct or adjacent cell. All room-level tracking methods estimate that the target is either in the correct cell or in an adjacent cell over 90% of the time.

	Environment Size	
	45 m x 30 m	60 m x 30 m
Kalman Filter	44.8% / 55.5%	60.7% / 66.1%
Simple Gaussian	71.9% / 97.6%	60.9% / 93.9%
Gaussian Process	75.0% / 95.5%	64.4% / 90.8%
Offset Gaussian	76.3% / 97.6%	68.0% / 94.1%

We also show results using a standard 2D Kalman filter for comparison. We use a Kalman filter implementation with a constant motion model (no odometry) that linearizes the range measurements in polar space [5]. This filter does not use a room-level discretization, so the estimate often falls outside of rooms on the map. Room-level tracking prevents this deviation from the map and improves tracking accuracy.




Figure 3: Noise model from using the mixture of Gaussians model where $z^s = \hat{z}^*$ and $w^s = \sum_i^{|\Gamma|} \omega_i s_i$.

5.3 Results with Unknown Node Locations

Table 2 shows mapping accuracy for our GPLVM node reconstruction method. Since the number of nodes was held constant, and their operating range is limited, larger environments led to sparser ranging data. The dimensionality reduction problem also becomes more challenging as environments get larger. We also added a test environment in which the target moved along a $30\text{ m} \times 30\text{ m}$ cross. This trial shows that our method works without closing the loop. Since

Table 2: Node mapping error comparison in office environment with unknown node locations. Mapping error is average Euclidean error for four nodes.

Map Size	Path Type	Mapping Error			
		Ground Truth	EKF-SLAM w/ Odom	EKF-SLAM w/o Odom	GPLVM
30 x 30	Cross	2.0 m	2.5 m	6.9 m	3.0 m
45 x 30	Loop	2.7 m	3.9 m	6.2 m	4.4 m
60 x 30	Loop	3.2 m	3.6 m	5.6 m	4.5 m

Table 3: Tracking results in office environment with unknown node locations. Mapping error is average Euclidean error for four nodes. Tracking accuracy is the room-level localization accuracy using reconstructed nodes from GPLVM path (xx/yy metric same as in Table 1).

Map Size	Path Type	Mapping Error	Tracking Accuracy
45 m x 30 m	Loop	4.4 m	52.4% / 76.4%
60 m x 30 m	Loop	4.5 m	52.6% / 80.4%

GPLVM dimensionality reduction requires conjugate gradient descent, it can often require large run times for convergence. Trials were terminated when the iterative log-likelihood increase fell below a threshold, which generally took approximately two hours on a standard desktop PC.

We also tested our method for tracking using reconstructed node locations in the same office environment. Table 3 presents results from tracking with unknown node locations. Tracking was performed using an offset Gaussian calculated from the reconstructed path and node locations.

For comparison, we implemented an EKF-SLAM method [5] for mapping unknown node locations. This algorithm updates an online Kalman estimate of the locations of the moving and stationary nodes in polar coordinates. It maintains multi-modal estimates, which avoids errors from poor initialization. We present results for the EKF-SLAM method both with and without odometry in Table 2. A constant motion model for the moving target is assumed for the case without odometry. The EKF-SLAM method outperforms our GPLVM method when odometry is used. However, without odometry our GPLVM method reconstructs the node locations 35% more accurately than the EKF-SLAM method. The accuracy gain is greater for smaller environments. This further demonstrates the appropriateness of our algorithm when odometry is unavailable.

Table 2 shows mapping errors from running an occupancy grid on the target’s ground truth path from laser localization. These errors, as high as 3 m on the large map, can be considered a gold standard. In other words, if the target’s path were reconstructed as accurately as ground truth, occupancy grid mapping with ranging radio data would yield these errors. This further demonstrates the non-Gaussian noise characteristics of the ranging sensors.

An alternative to reconstructing node locations using a path from GPLVM dimensionality reduction is to simply use a path found from odometry [6]. To explore this possibility, we ran a

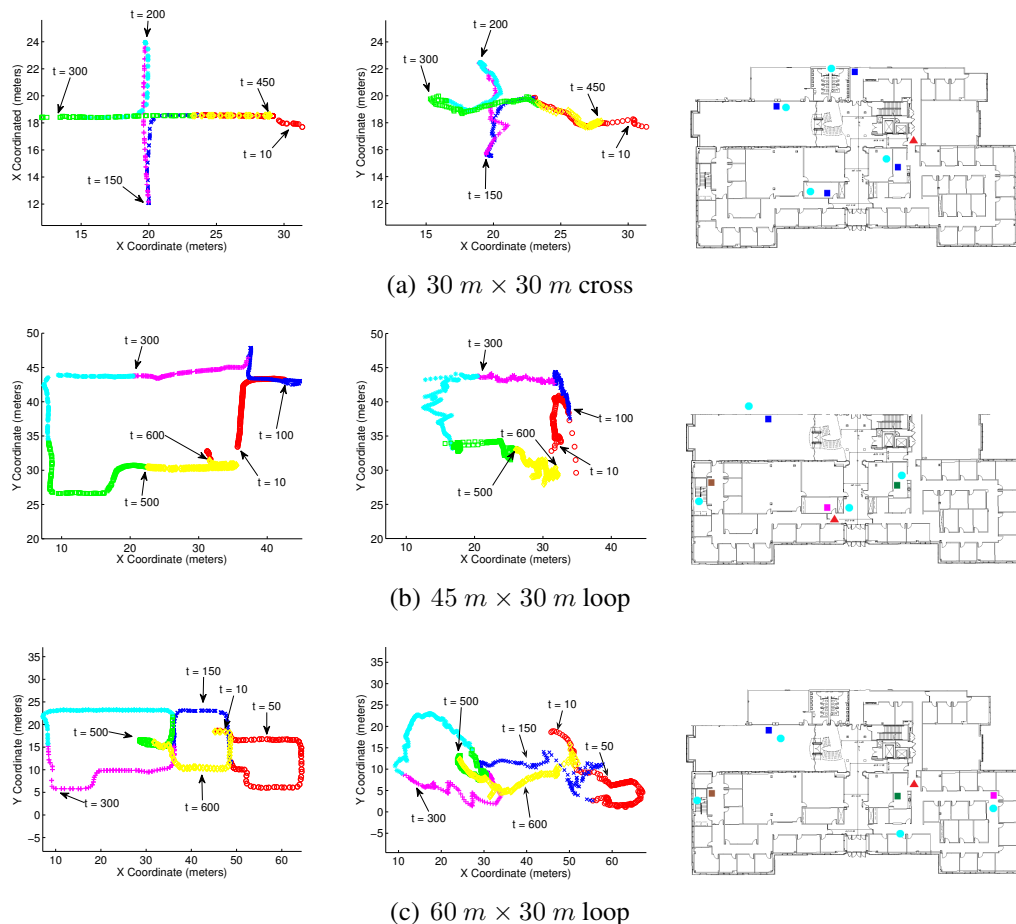


Figure 4: Example runs in different environments of mapping unknown node locations with GPLVM. Ground truth location of target (left), reconstructed path from GPLVM (middle), and estimated positions of radio nodes after occupancy grid mapping (right). Squares show actual stationary node locations, circles show reconstructed node locations, and triangle shows starting position of mobile node.

forward integration of the Pioneer robot’s odometry to generate a path through the environment, and we used this path to reconstruct the node locations with a Bayesian occupancy grid. Fig. 5 shows a comparison of node reconstruction accuracies for our GPLVM method and the method using odometry integration. The results show that, even using relatively accurate robot odometry, the path from odometry drifts far from ground truth. By the end of the run, the node estimates are over 14 m away from their correct locations. The GPLVM method, on the other hand, reconstructs a path that does not drift over time. This yields average final node estimates within 4.5 m of their correct locations. The reconstruction error using the ground truth path (found from laser localization) is also presented. These results show that the path reconstructed from GPLVM can be used to reconstruct the node locations nearly as accurately as the ground truth path.

Fig. 4 shows example paths reconstructed by GPLVM dimensionality reduction and an image of reconstructed nodes on a floor plan. The floor plan demonstrates the number of walls that the

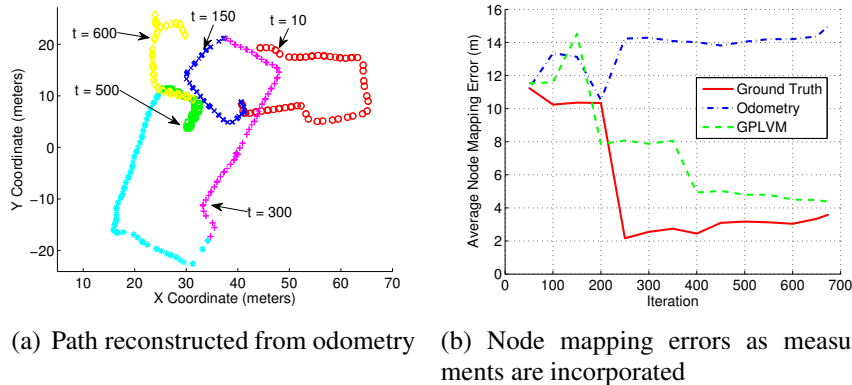


Figure 5: Comparison of path and node reconstruction for the GPLVM method and a method using only odometry in the $60\text{ m} \times 30\text{ m}$ environment. Drift in the path from odometry prevents the nodes from being accurately mapped. The GPLVM method does not use any odometry to reconstruct the path or the node locations.

nodes must travel through to track in this environment. Fig. 6 shows an example occupancy grid progression for a large loop. The node estimate quickly becomes a circular range annulus and later becomes unimodal as more measurements are incorporated.

6 CONCLUSIONS AND FUTURE WORK

In this paper, we have shown that it is feasible to track a moving target in a cluttered environment using very few ranging radio nodes and no odometry information. We have presented tracking methods for both known and unknown node locations, and we have demonstrated these methods in a complex environment. We have incorporated our methods into a room-level tracking framework that outperforms standard tracking methods, and we have presented a method using mixtures of Gaussians that removes outliers and yields better tracking results than both simple Gaussian modeling and Gaussian processes. These methods correctly locate the target in the correct cell or an adjacent cell up to 98% of the time.

When the node locations are unknown, we have demonstrated that GPLVM dimensionality reduction without odometry followed by Bayesian occupancy grid mapping can effectively find the positions of radio nodes with on average 35% more accuracy than an EKF-SLAM technique. We have also shown that paths from GPLVM reconstruction can locate nodes long after paths from odometry have drifted. Our GPLVM method provides a course estimate of the target’s position that puts it in the correct or adjacent cell 80% of the time.

Our method for reconstructing unknown node locations outperforms EKF-SLAM [5] when odometry is unavailable. These gains are likely due to the resilience of batch techniques in the face of outliers and measurement bias. Since the GPLVM framework reconstructs the path using a batch process, outliers in the data and the consistent bias of the range-only measurements are mitigated by the influence of the other data. In contrast, the EKF-SLAM technique maintains an online estimate of the world state. All prior data is folded into this estimate, and a history is not

maintained. In this application, it is clearly beneficial to maintain this history, and this leads to improved results with GPLVM.

For future work, we plan to refine the target’s dynamics models to better incorporate motion constraints and to better regain scale. More informed dynamics models will help improve the accuracy of the reconstructed node locations and the tracking accuracy after reconstruction. To improve sensor noise modeling, it may also be possible to utilize information from the floor plan. Since outliers tend to occur around highly cluttered areas, these characteristics of the environment may be incorporated into the algorithm.

Room-level tracking is a particularly useful tool in built environments. Our formulation utilizes a known map of the environment and serves to filter out noisy and erratic measurements. This technique shows gains over standard unconstrained 2D tracking methods. Previous approaches to map-based tracking with particle filters [9] also use the map to improve tracking results. We leave further comparison to future work.

We also plan to extend our work to multi-floor tracking. This will require more ranging radios and algorithmic adjustments to a 3D space. Another extension is to generalize our two-step method for recovering unknown node locations to sensors other than ranging radios. We believe that probabilistic dimensionality reduction provides a powerful tool for tracking and SLAM problems, which we plan to further explore.

7 ACKNOWLEDGMENTS

The authors gratefully acknowledge Bradley Hamner, Benjamin Grocholsky, Brian Ferris, and Neil Lawrence for their insightful comments. This work made use of Neil Lawrence’s publicly available GPLVM software. We thank Lester Foster at Multispectral Solutions, Inc. for ranging radios and hardware support. This work is funded by the National Science Foundation under Grant No. IIS-0426945.

References

- [1] G. Hollinger, J. Djughash, and S. Singh, “Tracking a Moving Target in Cluttered Environments with Ranging Radios,” in *Proc. Int’l Conf. on Robotics and Automation*, May 2008.
- [2] V. Kumar, D. Rus, and S. Singh, “Robot and Sensor Networks for First Responders,” *Pervasive Computing*, pp. 24–33, 2004.
- [3] E. Liao, G. Hollinger, J. Djughash, and S. Singh, “Preliminary Results in Tracking Mobile Targets Using Range Sensors from Multiple Robots,” in *Proc. 8th Int’l Symp. on Distributed Autonomous Robotic Systems*, 2006, pp. 125–134.
- [4] G. Hollinger, J. Djughash, and S. Singh, “Coordinated Search in Cluttered Environments Using Range from Multiple Robots,” in *Proc. 6th Int’l Conf. on Field and Service Robotics*, 2007.
- [5] J. Djughash, S. Singh, and B. Grocholsky, “Decentralized Mapping of Robot-Aided Sensor Networks,” in *Proc. Int’l Conf. on Robotics and Automation*, 2008.

- [6] E. Olson, J. Leonard, and S. Teller, “Robust Range-Only Beacon Localization,” *IEEE Journal of Oceanic Engin.*, vol. 31, no. 4, pp. 949–958, Oct. 2006.
- [7] N. Priyantha, H. Balakrishnan, E. Demaine, and S. Teller, “Mobile-Assisted Localization in Wireless Sensor Networks,” in *Proc. IEEE INFOCOM*, 2005, pp. 172–183.
- [8] L. Hu and D. Evans, “Localization for Mobile Sensor Networks,” in *Proc. 10th Int’l Conf. on Mobile Computing and Networking*, 2004, pp. 45–57.
- [9] B. Ferris, D. Hahnel, and D. Fox, “Gaussian Processes for Signal Strength-Based Location Estimation,” in *Proc. 2nd Robotics Science and Systems Conf.*, 2006, pp. 303–310.
- [10] A. Schwaighofer, M. Grigoras, V. Tresp, and C. Hoffmann, “GPPS: A Gaussian Process Positioning System for Cellular Networks,” in *Proc. 17th Conf. on Neural Information Processing Systems*, 2003.
- [11] N. Lawrence, “Probabilistic Non-Linear Principal Component Analysis with Gaussian Process Latent Variable Models,” *Journal of Machine Learning Research*, vol. 6, pp. 1783–1816, 2005.
- [12] N. Lawrence and J. Quiñonero-Candela, “Local Distance Preservation in the GP-LVM through Back Constraints,” in *Proc. 23rd Int’l Conf. on Machine Learning*, 2006, pp. 513–520.
- [13] J. Wang, D. Fleet, and A. Hertzmann, “Gaussian Process Dynamical Models for Human Motion,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, June 2007.
- [14] B. Ferris, D. Fox, and N. Lawrence, “WiFi-SLAM using Gaussian Process Latent Variable Models,” in *Proc. 20th Int’l Joint Conf. on Artificial Intelligence*, 2007, pp. 2480–2485.
- [15] “Multispectral Solutions, Inc: <http://www.multispectral.com/>,” 2008.
- [16] A. Kehagias, J. Djughash, and S. Singh, “Range-Only SLAM with Interpolated Range Data,” Robotics Institute, Carnegie Mellon University, CMU-RI-TR-06-26, Tech. Rep., 2006.
- [17] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, MA: MIT Press, 2005.
- [18] J. Djughash, S. Singh, G. Kantor, and W. Zhang, “Range-Only SLAM for Robots Operating Cooperatively with Sensor Networks,” in *Proc. Int’l Conf. on Robotics and Automation*, 2006.




Figure 6: Example likelihood map progression using GPLVM path for the nodes in Fig. 4 ($60 m \times 30 m$ loop). The left-to-right columns in this figure correspond to the left-to-right sequence of nodes (brown, blue, green, magenta) on the map. The estimates quickly become range annuli centered around the middle of the map (the moving node's starting location) and later collapse into a single mode as more range measurements are incorporated. Measurements are not received from the middle-left node until $t = 150$ s. Thus, the estimate is a uniform distribution until that time.