2008 IEEE International Conference on
Robotics and Automation
Pasadena, CA, USA, May 19-23, 2008

Tracking a Moving Target in Cluttered Environments with Ranging Radios

Geoffrey Hollinger, Joseph Djugash, and Sanjiv Singh

Abstract— **In this paper, 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 methods in an office environment. Our methods are successful at tracking a moving target node and mapping the locations of fixed nodes using radio ranging data that are both noisy and intermittent.**

I. 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.

G. Hollinger, J. Djugash, and S. Singh are with The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA gholling@andrew.cmu.edu, ssingh@ri.cmu.edu, robojoe@cmu.edu

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.

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 II discusses related work in the areas of tracking, simultaneous localization and mapping (SLAM), and dimensionality reduction. Sections III and IV present our framework for utilizing range-only measurements with known and unknown radio node locations. Section V discusses results from experiments with ranging radios in an office environment. Finally, Section VI draws conclusions and discusses directions for future work.

An extended version of this paper is available as a Technical Report [1]. The extended version contains more explanatory detail and further results verifying the performance of our algorithm.

II. 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.

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 (EKFs) [3]. 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 V, we run this EKF-SLAM algorithm on our data and compare the results to those obtained by our proposed algorithm.

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 [4]. Their method does not use a probabilistic formulation, which makes it poorly suited for sensors with nonstandard noise models. Hu and Evans also discussed localizing sensor networks that contain moving nodes [5]. 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.

Recent research has explored the use of Gaussian Processes for modeling the noise characteristics of non-linear sensors. Ferris et al. looked at tracking humans in office environments using measurements from wireless signal strength [6]. Schwaighofer et al. also applied Gaussian Processes with the Matern kernel function for localization using cellular phone signal strength [7]. 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 Lawrence as a probabilistic framework for non-linear dimensionality reduction [8]. GPLVMs were later extended by Wang et al. to incorporate dynamics [9]. Modeling dynamics allows for the incorporation of simple motion models into the GPLVM framework. Ferris et al. applied GPLVMs for use in solving the problem of localization with wireless signal strength when training data is unavailable [10]. 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.

III. TRACKING WITH KNOWN NODE LOCATIONS

A. Room-Level Tracking

When tracking a moving target in a cluttered environment, it is often advantageous to know in which room the target is located. Drawing from this observation, we discretize the floor plan into convex rooms and hallways and use this discretization to perform room-level tracking. 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 V.

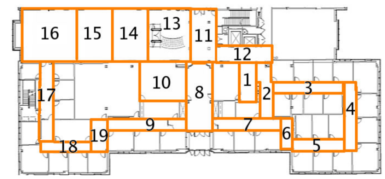

Fig. 1. Discretization of office environment used for ranging radio tracking. The discretization uses structural aspects of the built environment to provide a room-level tracking with very noisy measurements.

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 $\{c_1, c_2, \ldots, c_m\} \in C$. The resulting probability distribution in addition to the cell adjacency matrix forms a Markov Chain, which can be used to model the target's motion [11]. We can now incorporate information from range-only measurements. In a Bayesian framework, the posterior distribution is given by:

$$p_t^+(c) = \eta \; p(z|c) p_t^-(c), \tag{1}$$

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.

B. Ranging Radio Noise Modeling

Ultra-wideband ranging radios operating in the 6GHz+ frequency band can provide range measurements through walls in indoor environments [12]. These sensors 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.

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. 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.

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 $p(z|c)$. 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 2. The true range r_c can be calculated at a point by diving each cell into smaller subcells. The value $p(z|c)$ can then be recovered by summing over all subcells $c^b \in C$.

$$p(z|c) = N(z; r_c, \sigma_r^2), \tag{2}$$

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.

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. This section has been reduced due to space constraints. The longer version of this paper gives the full Gaussian Process derivation [1].

We are given some training data of the form $D = [(x_1, z_1), (x_2, z_2), \ldots, (x_n, z_n)]$ where $x_i \in \Re^d$ and $z_i \in \Re$. 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. 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:

$$z_i = z_i^o - r_{x_i}. \tag{3}$$

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 IV.

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. This is enough to define a Gaussian Process, which provides $p(z_*|x_*, X, Z^q)$ of receiving a range z_* at an arbitrary point x_*. If x_* is replaced by a point in a finely discretized subcell $c^b \in c$, this represents the likelihood $p(z|c^b)$. This can now be used to fold in information from measurements in our room-level tracking framework.

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 define this model recursively as calibration data is fed into the system. In this approach, the expected value of an observed range measurement z^o is as described below:

$$\hat{z}^\star = (\sum_i^{\|\Gamma\|} \omega_i s_i) * r_q + \varepsilon, \tag{4}$$

where r_q is the true range measurement between the 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 to model the expected mean offset within the observed range measurements. The term $\hat{z}^\star$ is a mixture of Gaussians, each of which has a scaled mean around the true expected range r_q.

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

$$\omega_i^+ = \eta\ \omega_i^- \exp(-\frac{(\omega_i^- r_q - z^o)^2}{2\sigma^2}), \tag{5}$$

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

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

$$p(z|c, X, Z) = N(z; (\sum_i^{\|\Gamma\|} \omega_i s_i) * r_c, \sigma_r^2). \tag{6}$$

We can also combine this mixture model with a filtering algorithm to reject outliers within the calibration data. We can do this easily by applying a measurement validation gate, such as a chi-squared test, to the update. The measurement gating step provides increased robustness to outliers when calculating the mixture of Gaussian model.

IV. TRACKING WITH UNKNOWN NODE LOCATIONS

We now present a tracking method that does not require calibration data and allows for initially unknown node locations. 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.

A. 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 $\Re^2$). Since the measurements from the ranging radios are 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. Data from ranging radios are often sparse, and we do not receive measurements from each radio at every time step.

To prevent missing data in Z, we linearly interpolate across time steps.

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:

$$p(Z|X^r) = \prod_{q=1:Q} N(Z^q; 0, K + \sigma_o^2 I), \tag{7}$$

where K is the squared exponential covariance matrix and σ_o^2 is a Gaussian observation noise hyperparameter.

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 [9] or using a more specialized prior using distance and orientation constraints [10]. Our results in Section V 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 Equation 8. The resulting model includes hyperparameters, which provide an estimate of the Gaussian noise variance. We utilize this estimate when recovering node locations in the next step.

$$p(X^r, Z) = p(Z|X^r)p(X^r) \tag{8}$$

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.

B. 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 [13]. We finely discretize the region in $\Re^2$ in which node q could be located into a grid X^q_{occ}. Now, step along the path X^r calculating the following at each cell:

$$p(x^q_{occ}) = \prod_{t=1:n} N(z^q_t; |x^q_{occ} - x^r_t|, \sigma_r^2), \tag{9}$$

where n is the number of range measurements from node q to the target, z^q_t 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 $p(x^q_{occ})$ for all cells in X^q_{occ}, we find the location of node q by setting $l^r_q = \max_x p(x^q_{occ})$.

Once the node locations have been reconstructed, 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 III, which can then be utilized for future tracking.

V. EXPERIMENTAL RESULTS

A. 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 [13]. 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 tracking experiments. This better models the case of a human target without odometry. The robot moved at a speed of approximately $0.2m/s$ during the experiments.

We utilized five ultra-wideband radio beacons from Multispectral Solutions to provide sensor measurements [12]. These sensors use time-of-arrival of ultra-wideband signals to provide inter-node ranging measurements through walls. 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 $30m$ when ranging through walls. Four radio nodes were placed around the environment, and one was placed on the Pioneer robot.

B. Results with Known Node Locations

We tested our methods in the office environment shown in Fig. 1. We first ran a calibration test to estimate the ranging radio noise variance with a simple Gaussian as $\sigma_r^2 = 3.85m^2$. Running a mixture of Gaussians, we found that a single Gaussian with an offset modeled the data well. Fig. 2 shows the offset Gaussian fit for the smaller loop. Applying this offset to that data yielded $\sigma_r^2 = 1.03m^2$ 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 I shows tracking results with known node locations in two office building loops. 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

TABLE I

ROOM-LEVEL TRACKING ACCURACY IN 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.

Known Node Locations	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%
Reconstructed Node Locations	45 m x 30 m	60 m x 30 m
Offset Gaussian	52.4% / 76.4%	52.6% / 80.4%

TABLE II

NODE MAPPING ERROR COMPARISON IN OFFICE ENVIRONMENT WITH UNKNOWN NODE LOCATIONS. MAPPING ERROR IS AVERAGE EUCLIDEAN ERROR FOR FOUR NODES.

		Mapping Error			
Map Size (m x m)	Path Type	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

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 if particularly useful with ultra-wideband ranging radios because of their bias towards measurements longer than the true range (see Fig. 2).

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 [3]. 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.

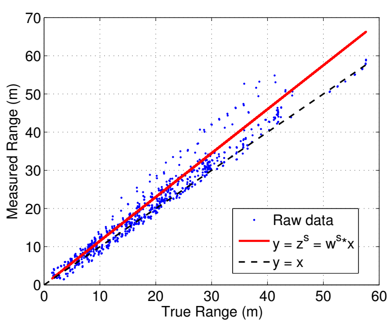

Fig. 2. Noise model from using the mixture of Gaussians model where $z^s = \hat{z}^\star$ and $w^s = \sum_i^{\|\Gamma\|} \omega_i s_i$.

C. Results with Unknown Node Locations

We also tested our method for tracking using reconstructed node locations in the same office environment. We added a test environment in which the target moved along a $30m \times 30m$ cross. This trial shows that our method works without closing the loop. Tables I and II present tracking and mapping results with unknown node locations. Path reconstruction with the GPLVM took approximately two hours per trial on a standard desktop PC. Optimization was terminated when the iterative log-likelihood increase fell below a threshold.

Table II also shows mapping errors from running an occupancy grid on the target's ground truth path from laser localization. These errors, as high as $3m$ 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. These errors further demonstrate the non-Gaussian noise characteristics of the ranging sensors.

For further comparison, we implemented an EKF-SLAM method [3] 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 II. 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. This demonstrates the appropriateness of our algorithm in situations where odometry is unavailable.

The extended version of this paper [1] compares our node reconstruction algorithm to a method that uses solely odometry to generate an initial path. Our algorithm outperforms this method because inevitable drift in odometry leads to inaccurate path reconstructions.

Fig. 3 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 nodes must travel through to track in this environment. The video included with the conference proceedings also shows an animation of our tracking algorithms. Fig. 4 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.

(a) t = 10 (b) t = 50 (c) t = 300

Fig. 4. Likelihood map progression using GPLVM path for the rightmost node (magenta) in Fig. 3. The estimate quickly becomes a range annulus 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.

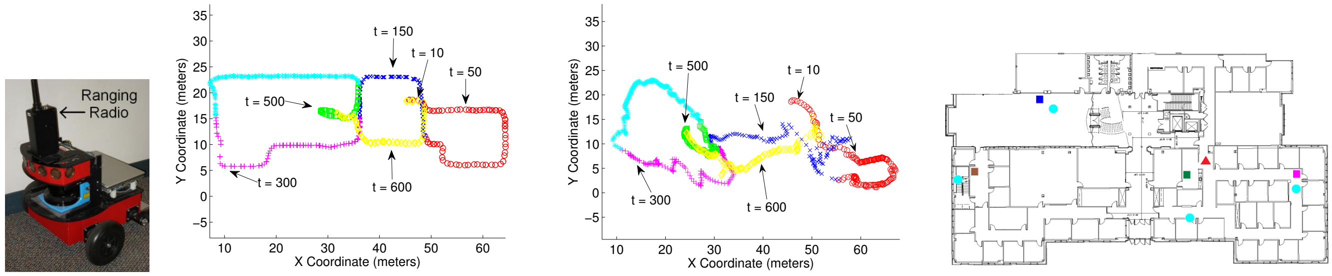

Fig. 3. Photograph of Multispectral ultra-wideband ranging radio mounted on Pioneer robot (far left). The robot was teleoperated around the environment to act as the moving target. Example run in $60m \times 30m$ loop of mapping unknown node locations with GPLVM: ground truth location of target (left), 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.

VI. CONCLUSIONS AND FUTURE WORK

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 locate radio nodes with on average 35% more accuracy than an EKF-SLAM technique. 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 [3] 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 better model the sensor noise, 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. Another extension is to generalize our decoupled 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 explore further.

VII. 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 radio hardware support. This work is funded by National Science Foundation under Grant No. IIS-0426945.

REFERENCES

[1] G. Hollinger, J. Djugash, and S. Singh, "Tracking a Moving Target in Cluttered Environments with Ranging Radios: Extended Results," Robotics Institute, Carnegie Mellon University, CMU-RI-TR-08-07, Tech. Rep., 2008.
[2] V. Kumar, D. Rus, and S. Singh, "Robot and Sensor Networks for First Responders," *Pervasive Computing*, pp. 24–33, 2004.
[3] J. Djugash, S. Singh, and B. Grocholsky, "Decentralized Mapping of Robot-Aided Sensor Networks," in *Proc. Int'l Conf. on Robotics and Automation*, 2008.
[4] N. Priyantha, H. Balakrishnan, E. Demaine, and S. Teller, "Mobile-Assisted Localization in Wireless Sensor Networks," in *Proc. IEEE INFOCOM*, 2005, pp. 172–183.
[5] 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.
[6] 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.
[7] 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.
[8] 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.
[9] J. Wang, D. Fleet, and A. Hertzmann, "Gaussian Process Dynamical Models for Human Motion," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, June 2007.
[10] 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.
[11] G. Hollinger, J. Djugash, 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.
[12] "Multispectral Solutions, Inc: http://www.multispectral.com/," 2008.
[13] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, MA: MIT Press, 2005.