Auton Robot (2012) 32:1–14
DOI 10.1007/s10514-011-9239-y

# Target tracking without line of sight using range from radio

**Geoffrey A. Hollinger · Joseph Djugash · Sanjiv Singh**

Received: 7 December 2010 / Accepted: 21 June 2011 / Published online: 12 July 2011


**Abstract** We propose a framework for utilizing fixed ultra-wideband ranging radio nodes to track a moving target radio node in an environment without guaranteed line of sight or accurate odometry. For the case where the fixed nodes' 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. For the case of unknown fixed node locations, we present a two-step approach that first reconstructs the target node's path using Gaussian Process Latent Variable models (GPLVMs) and then uses that path to determine the locations of the fixed nodes. We present experiments verifying our algorithm in an office environment, and we compare our results to those generated by online and batch SLAM methods, as well as odometry mapping. Our algorithm is successful at tracking a moving target node without odometry and mapping the locations of fixed nodes using radio ranging data that are both noisy and intermittent.

**Electronic supplementary material** The online version of this article (doi:10.1007/s10514-011-9239-y) contains supplementary material, which is available to authorized users.

G.A. Hollinger (✉)
Computer Science Department, Viterbi School of Engineering, University of Southern California, Los Angeles, CA, USA
e-mail: gahollin@usc.edu

J. Djugash · S. Singh
Robotics Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA

J. Djugash
e-mail: josephad@ri.cmu.edu

S. Singh
e-mail: ssingh@ri.cmu.edu

**Keywords** Range sensing · Sensor networks · Target tracking

## 1 Introduction

Our goal is to track a target moving in an indoor environment 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 (Tsai et al. 2010).

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. Unlike many tracking sensors, ranging radios do not require line-of-sight between nodes, and they do not suffer from false detections (i.e., the signal will only be detected if it has been transmitted). In addition, node identification can be transmitted with the signal, solving the data association problem trivially. The locations of the stationary beacons can either be surveyed as part of a pre-installed infrastructure, or they can be determined as part of a tracking algo-

rithm. 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 using a Bayesian map.

We present performance results of our algorithms in an office environment with varying node configurations 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 mapping.

This paper is organized as follows. Section 2 discusses related work in the areas of tracking, localization, 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, Sect. 6 draws conclusions and discusses directions for future work.

This paper is an extended version of a conference paper on the same topic (Hollinger et al. 2008). The journal version contains more explanatory detail of the proposed algorithms, supplementary images and test maps describing the mapping and tracking results, and additional comparisons validating our algorithms against batch SLAM and odometry-based mapping techniques.

## 2 Related work

The framework that we develop in this paper for tracking and mapping using ranging radios is closely related to literature in localization and dimensionality reduction. Kumar et al. (2004) discussed the problem of tracking a human first responder in an urban search and rescue scenario with robots and sensor networks. They outline some open questions and provide the motivation for using range-only devices for human tracking. Liao et al. (2006) extended these ideas by developing a tracking algorithm using range-only measurements and particle filters. 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.

Ranging radio tracking systems have been demonstrated in prior work (Schroeder et al. 2005; Kuhn et al. 2009). Many of these systems operate at a small scale (e.g., the room or even sub-room level) with a high node density of multiple nodes per room. Thus, they do not require probabilistic noise modeling beyond some simple signal processing. A survey of larger-scale ranging radio tracking systems is given by Gezici et al. (2005). They provide fundamental limits for the noise characteristics of various types of ranging sensors.

At the scale of a floor of a building with only a few nodes on the entire floor, it is beneficial to incorporate probabilistic noise models into the tracking system. Gustafsson and Gunnarsson (2005) developed a probabilistic noise model for ranging radios that jointly estimates whether the node is within line-of-sight, and Nicoli et al. (2008) provided a model based on jump Markov particle filters that also considers the line-of-sight characteristics of the measurement. These approaches are particularly useful in environments with few walls where there are distinct differences between the noise characteristics of measurements that do and do not have line-of-sight. In larger environments with many walls, it is more difficult to develop separate measurement models for the line-of-sight and non-line-of-sight cases. This prior work examined the case of known node locations, and they do not provide methods for reconstructing the locations of the radios without prior knowledge or pre-installed infrastructure.

The room-level tracking techniques proposed in the current paper were originally introduced in our previous work as part of a framework for searching a cluttered environment for a mobile target (Hollinger et al. 2009). Our previous work describes an algorithm that incorporates both coordination and tracking. In the current 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 initially unknown, the tracking problem is closely related to simultaneous localization and mapping (SLAM). Kehagias et al. (2006) developed a batch SLAM method for use with range-only sensors. This technique minimizes the total squared error of the

target's path and the locations of the stationary nodes given the range measurements. This technique worked well when a good initial estimate was available, but it was brittle to changes in the initial path. Djugash et al. (2008) proposed a range-only SLAM method using Extended Kalman Filters (EKFs). 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 improvements, 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 Sect. 5, we compare EKF SLAM and batch SLAM with non-linear minimization to our proposed algorithm.

In an application similar to ours, Olson et al. (2006) presented a method for locating ranging sonar beacons with an Autonomous Underwater Vehicle (AUV). They discuss a spectral graph partitioning method for outlier rejection, and they use a discretized voting scheme to attain an approximate estimate of the 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 Sect. 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. (2005) proposed a method for coordinating a mobile node to localize a range-only network. 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 (2004) also discussed localizing sensor networks that contain moving nodes. 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 reconstruct the location of the 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 et al. (2006) looked at tracking humans in office environments using measurements of wireless signal strength. Schwaighofer et al. (2003) also applied Gaussian Processes with the Matern kernel function for localization using cellular phone signal strength. 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 (Lawrence 2005). GPLVMs were later extended by Wang et al. (2007) to incorporate dynamics with applications to human motion modeling. Modeling dynamics allows for the incorporation of simple motion models into the GPLVM framework.

Ferris et al. (2007) applied GPLVMs to solve the problem of localization with wireless signal strength when training data is unavailable. Their algorithm takes advantage of the above tools in a target tracking scenario. Our work builds on this technique 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. In addition, tracking with wireless signal strength typically uses ten or more access points. Our method for reconstructing the node locations allows for tracking with sparse deployments, such as those analyzed in the current paper.

Ranging radio tracking has both advantages and disadvantages in comparison with wireless signal strength tracking (Ferris et al. 2006, 2007). Data from wireless signal strength are often erratic and can be affected by slight changes in the environment, while 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. Ranging radios also have the advantage over many sensors that they do not suffer from false detections. That is, a ranging radio will not believe it has received a range measurement that was not sent (assuming the data packets are properly synchronized and identified between radios). In addition, data association can be handled trivially by transmitting node identification numbers with range packets. On the other hand, wireless access points are already part of the infrastructure of many buildings, and radio nodes would need to be placed in a building before use in tracking. In addition, multipath and line-of-sight blockage can detrimentally affect the accuracy of ranging radio measurements.

Significant research effort has also been applied to tracking and localization using RFID tags (Zhou and Shi 2009). Ranging radios have the advantage of providing a range measurement between nodes. RFID tags typically provide only a passive "on" or "off" signal between detector and tag. In addition, RFID tags can suffer from false detections, since the receiver is not actively transmitting its signal.

## 3 Tracking with known node locations

We now present a room-level tracking approach that allows us to determine the room or hallway location of a target carrying a ranging node using stationary ranging nodes at known locations. In many indoor tracking applications, it may be more important to determine the target's room location than its location in 2D space. For instance, in the case of a lost first responder, a small tracking error could lead to the target being on different sides of a wall. On the contrary, when a room-level approach is used, a search team would know which room to search. In this section, we will assume that the locations of the nodes, as well as the map of the environment, are known. However, the noise modeling techniques are applicable to any discrete representation (e.g., a particle filter or regular grid) when the map is unknown.

### 3.1 Room-level tracking

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 Sect. 5.

Taking into account cell adjacencies in the environment forms an undirected graph of the target's movement options. We can now define a probability distribution $p_t(C)$ over the target's possible location at time $t$ in $m$ cells, where $C = \{c_1, c_2, \ldots, c_m\}$. The cells represent a subset of the map (in this case a subset of $\Re^2$). 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 $t$:

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

The diffusion matrix allows for flexible modeling of target motion, which can account for speed changes, room size, and other factors. If a more sophisticated non-Markovian model is available, it can be incorporated into this framework using sampling-based methods, such as particle filters (Thrun et al. 2005).

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:

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

where $p(z|c)$ is the probability of receiving range measurement $z$ given that the target is in cell $c$, and $\eta$ is a normalizing constant. The measurement $z$ is represented by a range value in $\Re^1$. Note that the measurement space is not required to be discrete, and the locations of the nodes do not need to be reduced to discrete values. We utilize a continuous likelihood function in the following sections.

In the following sections, we assume that the measurement likelihoods can be treated independently. The presence of measurement dependencies relies heavily on the homogeneity of the obstacles in the environment. For instance, a thick wall could potentially break the independence assumption. In the indoor environments of interest, the construction and number of walls is relatively homogeneous. The independence assumption has also been used successfully in prior work (Ferris et al. 2006).

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.

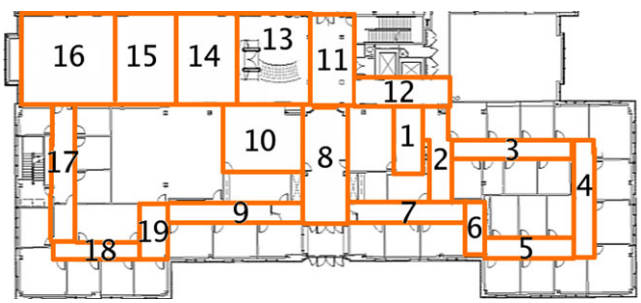


**Fig. 1** Discretization of office environment used for ranging radio tracking. The discretization uses structural aspects of the built environment to provide a room-level tracking with very noisy measurements. This environment was used with three different configurations of ranging radio nodes (see Fig. 4)

### 3.2 Ranging radio noise modeling

Ultra-wideband ranging radios operating in the 6 GHz + frequency band can provide range measurements through walls in indoor environments (Multispectral Solutions 2008). Ranging radios show better accuracy than alternative ranging devices in cluttered environments because the ultra-wideband signal produces sharper transitions and more accurate timing estimates from transmitter to receiver when ranging through obstacles.

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 difficult to model. 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 (3):

$$p(z|c) = \mathcal{N}(z; r_c, \sigma_r^2), \tag{3}$$

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

The range measurement variance $\sigma_r^2$ 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 (Hollinger et al. 2009). 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. (2006, 2007).

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

$$z_i = f(x_i) + \varepsilon. \tag{4}$$

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.[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 $z_i^o$:

$$z_i = z_i^o - r_{x_i}. \tag{5}$$

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 Sect. 4.

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:

$$\begin{aligned}\operatorname{cov}(f(x_i), f(x_j)) &= k(x_i, x_j) \\ &= \sigma_f^2 \exp\left(-\frac{1}{2l^2}|x_i - x_j|^2\right),\end{aligned} \tag{6}$$

where $\sigma_f^2$ and $l$ are hyperparameters.[2]

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

$$\operatorname{cov}(Z^q) = K + \sigma_n^2 I. \tag{7}$$

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

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

where

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

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

$k_*$ 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:

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

If $x_*$ is replaced by a point in a finely discretized subcell $c^b \in c$, (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.

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

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

### 3.2.3 Mixture of Gaussians modeling

A mixture of Gaussians model 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:

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

where $r_q$ is the true range between node $q$ and the target, $\varepsilon$ is zero mean Gaussian noise with variance $\sigma^2$, $\omega_i$ is the weight of the mixture Gaussian $i$, and $\Gamma$ 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 $\hat{z}^\star$ comes from 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 $\omega_i$ can be updated recursively within each iteration as described below:

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

where $\eta$ is a normalization term, which adjusts the weights $\omega_i$ 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:

$$p(z|c, X, Z) = \mathcal{N}\left(z; \left(\sum_i^{\|\Gamma\|} \omega_i s_i\right) * r_c, \sigma_r^2\right). \tag{14}$$

In the general case, the number of Gaussians, their offsets, and their variances can be learned from the data. In the results below, we use a single offset Gaussian, which empirically performs well on the ranging radio data. The Gaussians offset and its variance are learned from the data by maximizing the squared distance between the training data and the model prediction.

### 3.2.4 Outlier rejection

One particular benefit we gain from any of the above models is ease of incorporating a filtering algorithm to reject outliers within the calibration data. We can do this 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:

$$\epsilon_\nu = \nu^T \sigma^{-2} \nu, \tag{15}$$

where $\nu$ is the innovation (i.e., the difference between the measurement $z^o$ and its expected value $\hat{z}^\star$) and $\sigma^2$ is the variance of the accepted measurements. The term $\epsilon_\nu$ has a chi-square distribution, so the gate's bounding values can be read from the chi-square table, and a measurement can be discarded if its $\epsilon_\nu$ value is outside the bounding values. The measurement gating step provides increased robustness to outliers.

## 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 $\Re^2$). 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 (Lawrence 2005). GPLVMs have also been extended to incorporate dynamics (Wang et al. 2007), and they have been used to reconstruct 2D paths using wireless signal strength (Ferris et al. 2007). 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 tracking with ranging radios in indoor environments (Kehagias et al. 2006). For sensors with larger gaps in the data, a more sophisticated approach may be necessary.

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:

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

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 (Wang et al. 2007) or using a more specialized prior using distance and orientation constraints (Ferris et al. 2007). Our results in Sect. 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 the joint distribution in (17). As in prior work (Ferris et al. 2007), we use the Isomap algorithm to generate an initial path for the conjugate gradient optimization (Tenenbaum et al. 2000).

$$p(X^r, Z) = p(Z|X^r)p(X^r). \tag{17}$$

The resulting distribution $p(X^r, Z)$ provides a full generative model of the sensor noise. The path defined by $X^r$ is locally consistent (i.e., the distances between points are correct relative to each other, but not necessarily relative to any global reference frame). 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 the scale of the reconstructed path was often incorrect. While the range measurements encode scale, there is no odometry available to constrain the scale of the path itself, and the sparse nodes make reproduction of scale highly underconstrained. The use of a more informed motion model could potentially provide a better reconstruction of scale. For this paper, we simply 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.

### 4.2 Recovering node locations

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 approach (Thrun et al. 2005). We finely discretize the region in $\Re^2$ in which node $q$ could be located into a grid $X^q_{map}$.[3] Then, we step along the path $X^r$ calculating the following at each cell:

$$p(x^q_{map}) = \prod_{t=1:n} \mathcal{N}(z^q_t; |x^q_{map} - x^r_t|, \sigma_r^2), \tag{18}$$

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 $\sigma_r^2$ is a noise estimate for the radio sensors (estimated from prior data). Having calculated $p(x^q_{map})$ for all cells in $X^q_{map}$, we find the location of node $q$ by setting $l^r_q = \max_x p(x^q_{map})$.

The resulting grid map encodes the posterior probability of the node being located at each grid cell given the target's path and the range measurements received. Note that, while this is a mapping problem, it is simplified from general occupancy grid mapping that takes into account negative information (Mullane et al. 2009). The use of negative information (in this case missed measurements) may improve performance, at the cost of a significantly more complex map update. In addition, since a batch method is used, information from missed measurements would likely be redundant, as they would only give information about large areas in which the node could not be.

Once the node locations have been reconstructed on the 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 (Djugash et al. 2006).

After reconstructing the node locations, we can combine them with the path estimate and range measurements. This

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

yields enough information to calculate any of the noise models presented in Sect. 3, which can then be utilized for future tracking.

## 5 Experimental results

### 5.1 Hardware setup

A test environment was constructed 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 scanner, and a map of the environment was found using laser AMCL-SLAM methods from the Carmen software package (Thrun et al. 2005). Laser localization with the map was used for ground truth comparison and acquiring training data. The laser-based localization and odometry (wheel encoders and a gyro) were used to generate the ground truth, but neither were used for the GPLVM tracking experiments, which models the case of a human target without odometry and a laser scanner. The robot moved at a speed of approximately 0.3 m/s during the experiments.

We utilized five ultra-wideband radio beacons from Multispectral Solutions to provide sensor measurements (Multispectral Solutions 2008). These sensors use time-of-arrival of ultra-wideband signals to provide ranging measurements between nodes through walls. They propagate these measurements back to a base node using wireless communication. They are set 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 ranging through walls. Four radio nodes were placed around the environment in three different configurations (see Fig. 4), and one was placed on the Pioneer robot. Figure 2 shows a photograph of the Pioneer robot with a mounted ranging radio.[4]

### 5.2 Results with known node locations

We tested our methods for tracking with known fixed node locations (Sect. 3) in an office environment shown in Fig. 1. We first ran a calibration test to estimate the variance $\sigma_r^2$ of the ranging radios. The variance with a simple Gaussian was estimated 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. Figure 3 shows the offset Gaussian fit for the smaller loop. Applying this offset

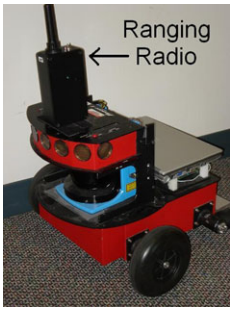


**Fig. 2** Photograph of multispectral ultra-wideband ranging radio mounted on Pioneer robot. The robot was teleoperated around the environment to act as the moving target

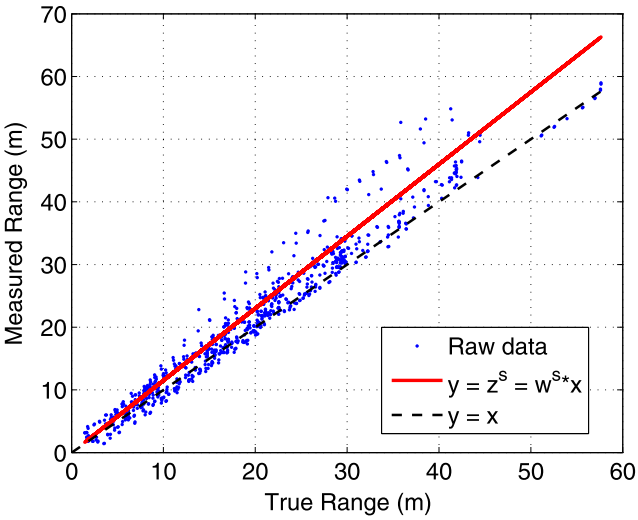


**Fig. 3** Noise modeling using an offset Gaussian. The corrected line (*solid*) using an offset Gaussian provides a better model of the true range than the raw measurements

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. The value of the offset that produced the smallest $\sigma_r$ was consistent throughout all trials, showing it to be robust to changes in node configuration throughout the same environment.

Table 1 shows tracking results with known node locations in both a 45 m × 30 m and a 60 m × 30 m office building loop. Each result is averaged over two separate trials (not including the training trial). The coarse discretization is shown in Fig. 1, and the finer subcells used for tracking were set to 50 cm × 50 cm. All methods estimate that the target is either in the correct cell or in an adjacent cell over 90% of the time. In these trials, a random target diffusion model was used that corresponds to a speed of 0.3 m/s. Varying the speed of the random walk by an order of magnitude did not significantly affect the tracking accuracy for these trials, which shows the method to be fairly robust to changes in target motion modeling.

The results with known node locations show that Gaussian Process modeling improves room-level tracking accu-

[4]This paper has supplementary downloadable material, which includes a movie clip showing playback of the tracking experiments. This material is available at http://www.springer.com/engineering/robotics/journal/10514.

**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

| | Loop size of target path | |
|---|---|---|
| | 45 m × 30 m | 60 m × 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% |

racy 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. This is likely due to the Gaussian Process method's poor performance for areas that it does not have dense training data. In Sect. 6, we discuss several methods for improving the performance of the GP method. The offset Gaussian if particularly useful for correcting the bias towards measurements longer than the true range (see Fig. 3). The tracking methods all had a tendency to estimate the target's cell incorrectly when erratic measurements were received (from multi-path effects or signal interference). These cases tended to produce several erroneously long measurements in a row, which our methods were not able to model or filter out.

We also show results using a standard 2D Kalman filter for comparison. We use a Kalman filter implementation with a constant velocity motion model (no odometry) that linearizes the range measurements in polar space (Djugash et al. 2008). This filter operates in continuous space and 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.

### 5.3 Results with unknown node locations

We now examine the performance of our proposed method when the path is reconstructed using GPLVM dimensionality reduction and the node locations are mapped with a Bayesian grid. For the following experiments, we assume that the floor plan is known a priori. However, we note that the floor plan is only necessary for the room-level tracking component and is not used in the GPLVM dimensionality reduction or node reconstruction phases. If a floor plan were not available, alternative continuous tracking methods could be utilized with any of the proposed noise models and the reconstructed node locations.

Using the same test environment as in the case of known nodes, Fig. 4 shows example paths reconstructed by GPLVM dimensionality reduction and an image of reconstructed nodes on a floor plan. The hyperparameters of the GPLVM were estimated by maximizing the log likelihood of the data from a training run with a different node configuration and known node locations. The approximate hyperparameter values were $L = 5$, $\sigma_n = 1$, and $\sigma_f = 10$, though there was some variation depending on the training run used. The locations of the path at $t = 0$ and $t = 300$ was assumed to be known and used to adjust the scale and rotation to align with the floor plan (see Sect. 4.1). After the path was found using dimensionality reduction, the locations of the nodes were reconstructed on the Bayesian grid. Fig. 5 shows an example Bayesian grid progression for a large loop. The node estimate quickly becomes a circular range annulus and later becomes unimodal as more measurements are incorporated.

Table 2 shows the mapping accuracy of the nodes utilizing the paths reconstructed from the GPLVM. Since the number of nodes was held constant, and their operating range is limited, larger areas covered by the target led to sparser ranging data. The dimensionality reduction problem also becomes more challenging as the 2D area spanned by the target path becomes longer. We additionally show mapping errors from node reconstruction using 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, subject to the accuracy of the measurement noise model. In other words, if the target's path were reconstructed as accurately as ground truth, grid mapping with ranging radio data would yield these errors. Note that a more accurate measurement model could improve this estimate.

Since GPLVM dimensionality reduction uses conjugate gradient descent, it can require significant run times for convergence. Trials were terminated when the iterative log-likelihood increase fell below a threshold, which generally took approximately one hour on a 3.2 GHz Intel i7 processor with 9 GB or RAM. The MATLAB implementation was not heavily optimized, and the running time could likely be reduced through vectorization and optimization for parallel processing. The use of sparse approximations for the GPLVM would also reduce runtime (Snelson 2007).

For comparison, we implemented an Extended Kalman Filter (EKF) SLAM method (Djugash et al. 2008) for mapping unknown node locations. This algorithm updates an online Kalman estimate of the locations of all nodes using a polar representation. The polar representation provides increased accuracy after linearization (see Djugash et al. 2008 for more detail). This method maintains multimodal estimates, which avoids errors from poor initialization. We present results for the EKF method both with and

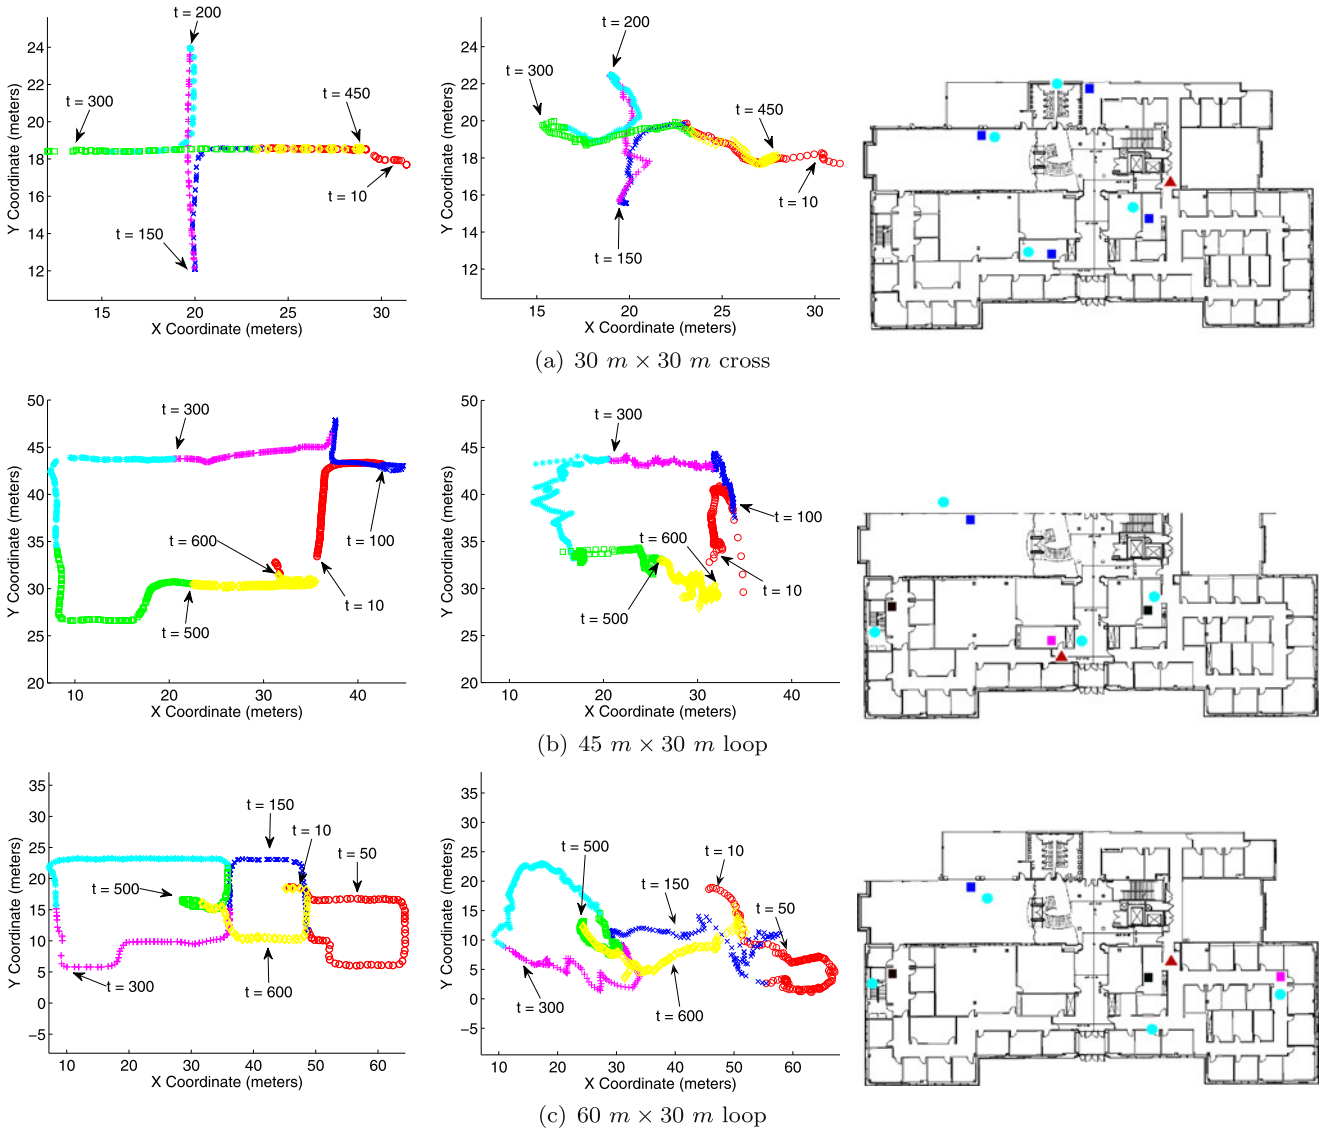


**Fig. 4** (Color online) Example runs of mapping unknown node locations with GPLVM. Ground truth location of target (*left*), reconstructed path from GPLVM (*middle*), and estimated positions of radio nodes after grid mapping (*right*). Sections of the path are colored differently for visual correspondence between the true and reconstructed path. *Squares* show actual stationary node locations, *circles* show reconstructed node locations, and *triangle* shows starting position of mobile node. Though the target path reconstruction is highly imperfect, the resulting reconstructed nodes are sufficient to track the target with room-level accuracy (see Table 3). Note that only four nodes are used for as much as a 60 m × 30 m area

without odometry in Table 2. A constant velocity model for the moving target is assumed for the case without odometry.

The EKF method outperforms the GPLVM method when odometry is used. The robot's odometry is found using both a wheel encoder and gryo, making it quite accurate over short distances. In the case of tracking a human, such accurate odometry would not be available. Without odometry, the GPLVM method reconstructs the node locations 35% more accurately than the EKF method. The accuracy gain is greater for smaller environments. This further demonstrates the appropriateness of our algorithm when odometry is unavailable.

Unlike the proposed GPLVM method, the EKF SLAM method does not perform a batch optimization. To determine the benefit of batch optimization, we also compare to a batch SLAM method proposed by Kehagias et al. (2006). This method encodes the range measurements as constraints and performs a non-linear minimization on both the path and node locations. For the batch SLAM, Isomap was used to produce the initial path, and a motion model was incorporated that penalizes points on the path that would require a speed greater than 0.3 m/s. The squared error of both the

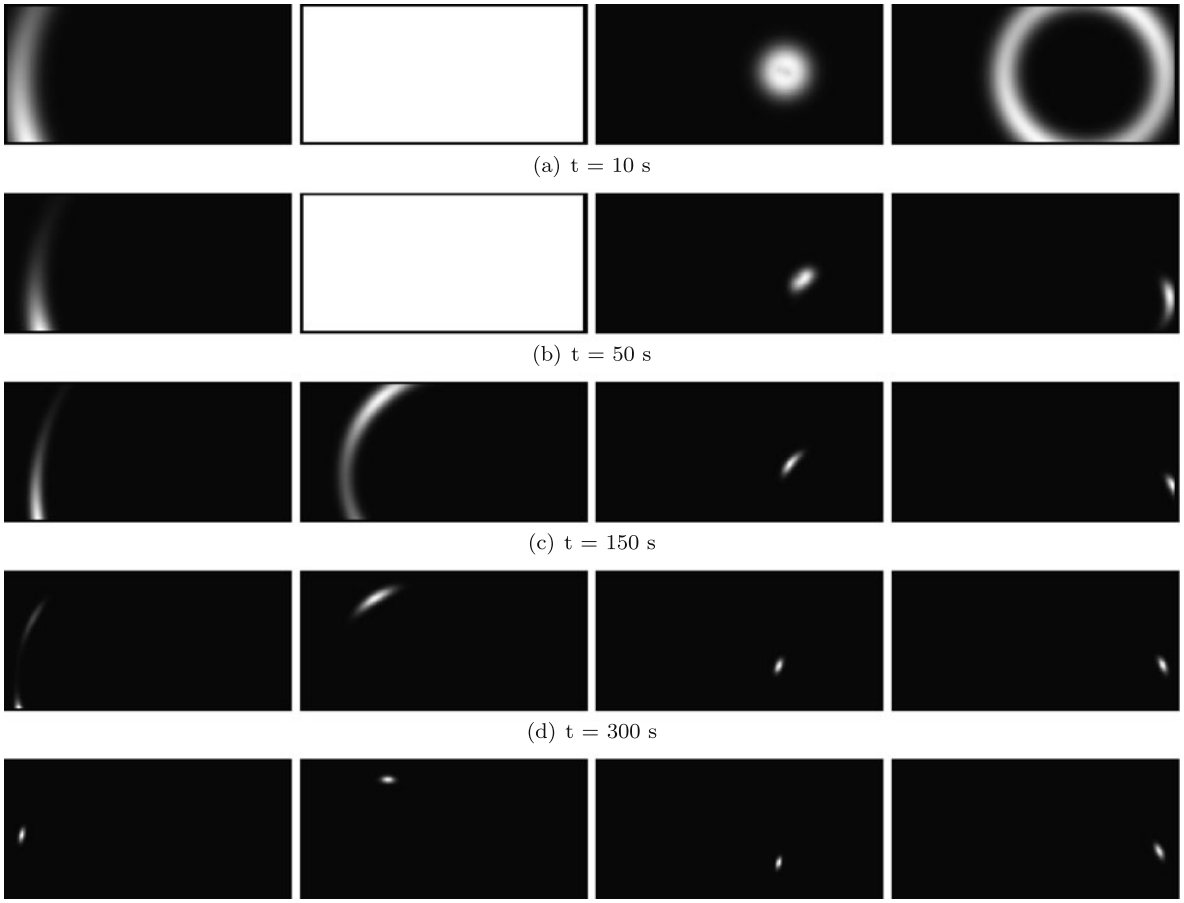


**Fig. 5** Example likelihood map progression using GPLVM path for the nodes in Fig. 4 (60 m × 30 m loop). *The left-to-right columns* in this figure correspond to the left-to-right sequence of nodes (*brown*, *blue*, *green*, *magenta*) on the map. The estimates quickly become range annuli centered 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

**Table 2** Node mapping error comparison in office environment with unknown node locations. Mapping error is average Euclidean error for four nodes. Two types of target paths are considered: the loop moves around in a closed circuit, and the cross returns to a middle point a number of times (see Fig. 4). The mapping errors with a known path use a ground truth path found using laser localization

| Map size | Path type | Mapping error | | | | |
|---|---|---|---|---|---|---|
| | | Known path | EKF SLAM w/ Odom | EKF SLAM w/o Odom | Batch SLAM w/o Odom | GPLVM |
| 30 × 30 | Cross | 2.0 m | 2.5 m | 6.9 m | 2.9 m | 3.0 m |
| 45 × 30 | Loop | 2.7 m | 3.9 m | 6.2 m | 6.7 m | 4.4 m |
| 60 × 30 | Loop | 3.2 m | 3.6 m | 5.6 m | 12.4 m | 4.5 m |

range measurements and the motion penalties was minimized for the three trials, and the results are presented in Table 2.

The batch SLAM method is competitive with our proposed method for the cross trial, but it fails to reconstruct the node locations accurately in the loop trials with longer target paths. We conjecture that the failure arises because the minimization finds a local minimum of the cost function. In contrast, the GPLVM is more robust to local minima and outliers. The batch SLAM method was run until convergence, which yielded comparable running times with the GPLVM (approximately one hour on the large loop).

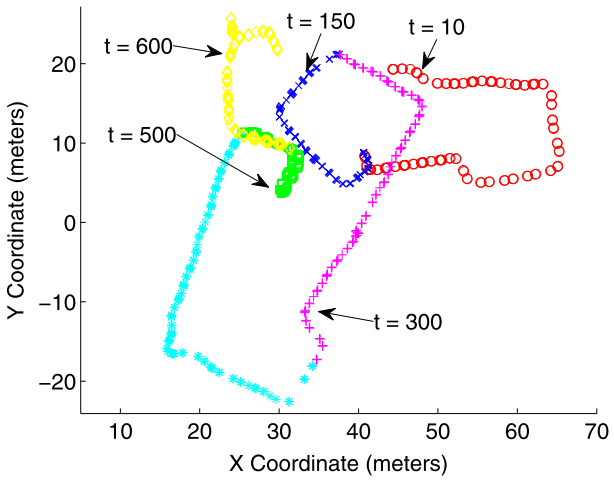


(a) Path from odometry in 60 $m \times 30$ $m$ loop

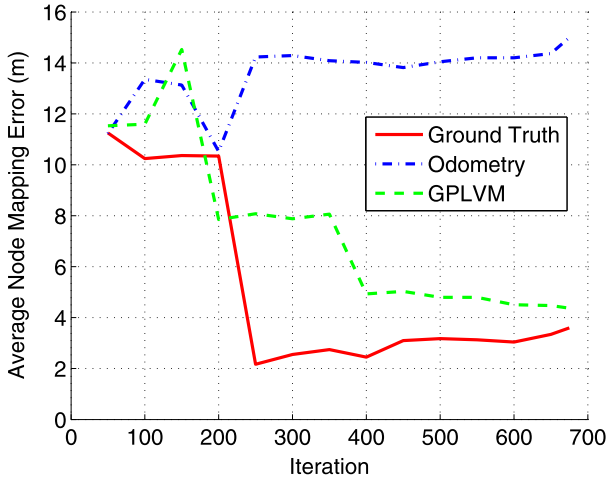


(b) Node mapping errors as measurements are incorporated

**Fig. 6** Comparison of path and node reconstruction for the GPLVM method and a method using only odometry in the 60 m × 30 m environment (ground truth and GPLVM reconstructed paths shown previously in Fig. 4). 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

**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 × 30 m | Loop | 4.4 m | 52.4%/76.4% |
| 60 m × 30 m | Loop | 4.5 m | 52.6%/80.4% |

An alternative to reconstructing node locations using a path from GPLVM dimensionality reduction is to use a path found from odometry (Olson et al. 2006). To explore this possibility, we ran a 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 grid. Figure 6 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.

We also tested our method for tracking using reconstructed node locations in the same office environments. Tracking was performed using an offset Gaussian calculated from the reconstructed path and node locations. The method was able to predict the target in the correct cell or in an adjacent cell approximately 80% of the time (see Table 3). Note that these results are for the case without pre-installed infrastructure and do not utilize any odometry or inertial measurement of the target.

## 6 Conclusions and future work

In this paper, we have shown that it is feasible to track a moving target without line-of-sight 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. Our noise modeling methods utilize a discretized map of the environment, and they are capable of operating without a map through the use of a regular grid or particle filter.

When the node locations are unknown, we have demonstrated that the two-step method of GPLVM dimensionality reduction followed by Bayesian grid mapping can effectively find the positions of radio nodes with on average 35% more accuracy than an EKF SLAM technique. In addition, our method is able to map the node locations when a more

standard batch SLAM method fails. We have also shown that paths from GPLVM reconstruction can locate nodes 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 approximately 80% of the time. The proposed method uses a batch technique to reconstruct the target's path and then utilizes that path to map the locations of the nodes. The path reconstruction and node mapping are decoupled and solved sequentially. This decoupling reduces the computation and complexity of the approach, making it feasible on a standard processor. The performance of the two-step technique demonstrates that this decoupling is justified in this scenario.

Our method for reconstructing unknown node locations outperforms an EKF SLAM method (Djugash et al. 2008) 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 technique maintains an online estimate of the world state. All prior data is folded into this estimate, and a history is not maintained. In target tracking applications, it is clearly beneficial to maintain this history, which leads to improved results. The GPLVM method also succeeds in reconstructing the nodes when a batch SLAM method that uses non-linear minimization fails. The GPLVM is more robust to local minimum, and it is able to reconstruct the node locations in trials with long target paths and sparse measurement data.

An interesting area for future work is to combine the benefits of GPLVMs with those of graph-based SLAM methods (Kaess et al. 2008; Grisetti et al. 2010). Unlike GPLVM dimensionality reduction, graph-based SLAM methods allow for both batch operation and incremental improvement during operation, but they are often brittle to the choice of motion model. To apply such techniques in the ranging radio tracking domain, it would be necessary to develop methods for incorporating range measurements into the SLAM graph without the use of accurate odometry or a target motion model.

Additional future work includes examining alternative Gaussian Process variations for modeling the noise characteristics of the sensors. Specifically, the use of non-stationary kernels or a polar representation could improve performance when little training data is available. For the path reconstruction phase, refining the target's dynamics models to incorporate motion constraints would likely improve performance. More informed dynamics models will help improve the accuracy of the reconstructed node locations and the tracking accuracy after reconstruction. The use of the recently developed GP-Bayes filters (Ko and Fox 2011) would also allow for better incorporation of motion.

One avenue for broadening the applicability of our methods is to extend them to multi-floor tracking, which would require a denser deployment of 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. This paper shows that probabilistic dimensionality reduction provides a powerful tool for solving target tracking problems with ranging radios. We believe that further exploration into these methods will lead to improvements across application domains.

**Acknowledgements** 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 NSF under Grant No. IIS-0426945.

## References

Djugash, J., Singh, S., Kantor, G., & Zhang, W. (2006). Range-only SLAM for robots operating cooperatively with sensor networks. In *Proc. IEEE int. conf. robotics and automation*.

Djugash, J., Singh, S., & Grocholsky, B. (2008). Decentralized mapping of robot-aided sensor networks. In *Proc. IEEE int. conf. robotics and automation*.

Ferris, B., Hahnel, D., & Fox, D. (2006). Gaussian processes for signal strength-based location estimation. In *Proc. robotics: science and systems conf.*

Ferris, B., Fox, D., & Lawrence, N. D. (2007). WiFi-SLAM using Gaussian process latent variable models. In *Proc. int. joint conf. artificial intelligence*.

Gezici, S., Thian, Z., Giannakis, G. B., Kobayashi, H., Molisch, A. F., Poor, H. V., & Sahinoglu, Z. (2005). Localization via ultra-wideband radios. *IEEE Signal Processing Magazine*, *22*(4), 70–84.

Grisetti, G., Kuemmerle, R., Stachniss, C., Frese, U., & Hertzberg, C. (2010). Hierarchical optimization on manifolds for online 2D and 3D mapping. In *Proc. IEEE int. conf. robotics and automation*.

Gustafsson, F., & Gunnarsson, F. (2005). Mobile positioning using wireless networks. *IEEE Signal Processing Magazine*, *22*(4), 41–53.

Hollinger, G., Djugash, J., & Singh, S. (2008). Tracking a moving target in cluttered environments with ranging radios. In *Proc. IEEE int. conf. robotics and automation*.

Hollinger, G., Singh, S., Djugash, J., & Kehagias, A. (2009). Efficient multi-robot search for a moving target. *The International Journal of Robotics Research*, *28*(2), 201–219.

Hu, L., & Evans, D. (2004). Localization for mobile sensor networks. In *Proc. int. conf. mobile computing and networking*.

Kaess, M., Ranganathan, A., & Dellaert, F. (2008). iSAM: Incremental smoothing and mapping. *IEEE Transactions on Robotics*, *24*(6), 1365–1378.

Kehagias, A., Djugash, J., & Singh, S. (2006). *Range-only SLAM with interpolated range data*. Technical Report CMU-RI-TR-06-26, Robotics Institute, Carnegie Mellon Univ.

Ko, J., & Fox, D. (2011). Learning GP-Bayes filters via Gaussian process latent variable models. *Autonomous Robots*, *30*(1), 3–23.

Kuhn, M., Zhang, C., Mahfouz, M., & Fathy, A. E. (2009). Real-time UWB indoor positioning system with millimeter 3-D dynamic accuracy. In *IEEE antennas and propagation society int. symp.*

Kumar, V., Rus, D., & Singh, S. (2004). Robot and sensor networks for first responders. *IEEE Pervasive Computing*, *3*(4), 24–33.
Lawrence, N. D. (2005). Probabilistic non-linear principal component analysis with Gaussian process latent variable models. *Journal of Machine Learning Research*, *6*, 1783–1816.
Liao, E., Hollinger, G., Djugash, J., & Singh, S. (2006). Preliminary results in tracking mobile targets using range sensors from multiple robots. In *Proc. int. symp. distributed autonomous robotic systems*.
Mullane, J., Adams, M. D., & Wijesoma, W. S. (2009). Robotic mapping using measurement likelihood filtering. *The International Journal of Robotics Research*, *28*(2), 172–190.
Multispectral Solutions, Inc. (2008). Company website. http://www.multispectral.com/.
Nicoli, M., Morelli, C., & Rampa, V. (2008). A jump Markov particle filter for localization of moving terminals in multipath indoor scenarios. *IEEE Transactions on Signal Processing*, *56*(8), 3801–3809.
Olson, E., Leonard, J. J., & Teller, S. (2006). Robust range-only beacon localization. *IEEE Journal of Oceanic Engineering*, *31*(4), 949–958.
Priyantha, N. B., Balakrishnan, H., Demaine, E., & Teller, S. (2005). Mobile-assisted localization in wireless sensor networks. In *Proc. IEEE INFOCOM*.
Schroeder, J., Galler, S., & Kyamakya, K. (2005). A low-cost experimental ultra-wideband positioning system. In *IEEE int. conf. ultra-wideband*.
Schwaighofer, A., Grigoras, M., Tresp, V., & Hoffmann, C. (2003). GPPS: a Gaussian process positioning system for cellular networks. In *Proc. 17th conf. on neural information processing systems*.
Snelson, E. (2007). *Flexible and efficient Gaussian process models for machine learning*. Ph.D. Thesis, Gatsby Computational Neuroscience Unit, University College London, University of London.
Tenenbaum, J. B., de Silva, V., & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, *290*(5500), 2319–2323.
Thrun, S., Burgard, W., & Fox, D. (2005). *Probabilistic robotics*. Cambridge: MIT Press.
Tsai, Y.-L., Tu, T.-T., Bae, H., & Chou, P. H. (2010). EcoIMU: a dual triaxial-accelerometer inertial measurement unit for wearable applications. In *Int. conf. body sensor networks*.
Wang, J. M., Fleet, D. J., & Hertzmann, A. (2007). Gaussian process dynamical models for human motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *30*(2), 283–298.
Zhou, J., & Shi, J. (2009). RFID localization algorithms and applications—a review. *Journal of Intelligent Manufacturing*, *20*, 695–707.

**Geoffrey A. Hollinger** is a Postdoctoral Research Associate in the Robotic Embedded Systems Laboratory and Viterbi School of Engineering at the University of Southern California. His current interests are active perception and multi-robot path planning for communication constrained marine robotics. His past projects include autonomous search at Carnegie Mellon University, personal robotics at Intel Research Pittsburgh, active estimation at the University of Pennsylvania's GRASP Laboratory, and miniature inspection robots for the Space Shuttle at NASA's Marshall Space Flight Center. He received his Ph.D. (2010) and M.S. (2007) in Robotics from Carnegie Mellon University and his B.S. in General Engineering along with his B.A. in Philosophy from Swarthmore College (2005).

**Joseph Djugash** is currently a Ph.D. candidate at the Robotics Institute in Carnegie Mellon University. He received his Bachelors in Computer Engineering from the University of Minnesota at the Twin Cities in 2004 and his Masters in Robotics from Carnegie Mellon University in 2007. His research interests are navigation (tracking, localization, and mapping), ranging radios, nonlinear estimation techniques, decentralized and distributed filtering, and active estimation.

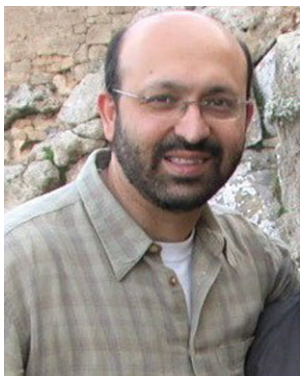

**Sanjiv Singh** is a Research Professor at the Robotics Institute, Carnegie Mellon University. His recent work has two main themes: perception in natural environments and multi-agent coordination. He has led projects in both ground and air vehicles operating in unknown or partially known environments, in applications such as mining, agriculture, emergency response, surveillance and exploration. He is also actively involved in the automation of complex tasks, such as the assembly of large space structures, that cannot be addressed by single agents and must necessarily be performed by teams. Prof. Singh received his B.S. in Computer Science from the University of Denver (1983), M.S. in Electrical Engineering from Lehigh University (1985) and a Ph.D. in Robotics from Carnegie Mellon (1995). He is the founder and Editor-in-Chief of the Journal of Field Robotics.