

# Multi-sensor ATTenuation Estimation (MATTE): Signal-strength prediction for teams of robots

Johannes Strom and Edwin Olson

**Abstract**—Multi-robot teams are often constrained by communications; better signal-strength models enable more efficient coordination while still maintaining adequate communication. This work discusses several prediction algorithms applicable to this scenario. Whereas previous approaches typically focus on prediction in the presence of deployed base-stations, we consider the more general problem where all nodes in the network can be mobile. Our new algorithm, Multi-sensor ATTenuation Estimation (MATTE), addresses this problem by leveraging other forms of sensor data in combination with signal-strength measurements to infer the locations of attenuating materials in the robots’ environment. We also extend prior tomographic and correlation-based approaches to the multi-robot case, allowing a competitive evaluation. All methods are evaluated on a large corpus of real-world indoor and outdoor environments.

## I. INTRODUCTION

The inherently parallel nature of search-and-rescue missions creates an opportunity for teams of collaborating robots to assist in emergency response. However, robot teams that wish to collaborate must communicate to coordinate effectively. Current approaches to multi-robot coordination typically only incorporate a simple fixed-radius communication model as a planning constraint [1], [2], [3]. In environments with variable attenuation (e.g. urban environments), picking a single radius may not be appropriate since communication will be easier in open spaces, and harder in densely-built neighborhoods.

Attempts to incorporate more complicated models of signal propagation typically focus on the case where robots communicate with an existing fixed base station [4]. Unfortunately, many domains lack a usable communications infrastructure (e.g. disaster zones), forcing robots to deploy their own. Achieving good performance from fully-mobile networks is challenging because more complicated signal-propagation models must be incorporated into the path-planning process to ensure connectivity. Furthermore, existing models for fixed-transmitters do not extend directly to the case where all nodes are mobile. While prediction in the former case is analogous to regression in a two dimensional space, the later requires making predictions for a four dimensional space, but without a corresponding increase in data density. The result is that achieving similar generalization performance from observed signal-strength measurements becomes more challenging. This paper explores how existing methods can be modified to better cope with this reduction

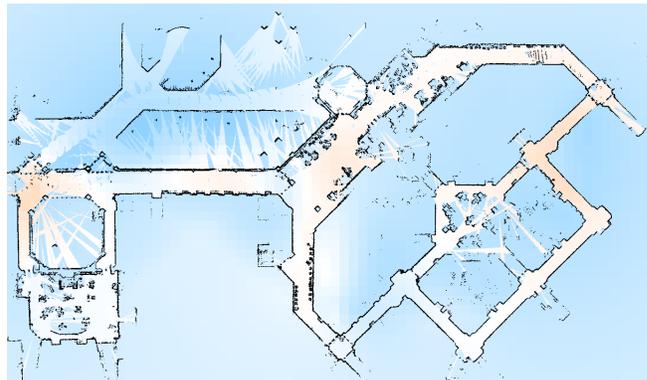


Fig. 1. Attenuation estimation computed from a single traversal of a  $160 \times 100\text{m}$  environment by 3 robots conducting an exploration mission. Blue indicates regions where signals are attenuated, white and orange regions indicate where signals pass more easily. Black denotes building structure. Our algorithm, MATTE, is designed for predicting signal strength in the challenging case where all transmitters and receivers are mobile.

in data density, and explores methods for using additional sensor data to inform a better prior over the locations of significant attenuators in the environment. Specifically, the main contributions of this paper are:

- Extension of tomographic and correlative signal prediction techniques to the case of multiple mobile robots without a base station.
- A new signal-strength prediction method, Multi-sensor ATTenuation Estimation (MATTE), which additionally leverages the robots’ LIDAR data to better predict the location of attenuating objects.
- Extensive evaluation on real world datasets covering over  $40,000 \text{ m}^2$  of both indoor and outdoor environments

## II. RELATED WORK

In the robotics domain, planning with guaranteed communication is a well studied topic [5], [6]. However, communication between two agents cannot be ensured in general, so these methods are limited in their application to the real world. Others have studied collaborative planning under a fixed-radius communication assumption [1], but such methods result in unnecessarily conservative strategies because they fail to exploit long-range links when possible, reducing the effective speed of the robots. More realistic models can be obtained by first predicting the signal strength along a given link, which can then be used to predict packet success rates. There are two main approaches to signal prediction in the literature – the first is *correlative* in nature, that is

This work was supported by U.S. DoD Grant FA2386-11-1-4024.

The authors are with Department of Computer Science and Engineering, University of Michigan, Ann Arbor, MI 48109 USA {jbstrom, ebolson}@umich.edu

broadcasts from nearby positions are assumed to have similar signal strengths, allowing prediction at nearby points using what is essentially a locally-weighted average. The second technique uses principles from *tomography*, where changes in signal-strength are used to infer the presence of attenuating objects, which can in turn be used to predict signal-strength at unknown locations.

Malmirchegini and Mostofi have explored correlative methods of predicting the link quality between a fixed base station and a mobile robot, taking into account the spatial correlations of signal strength measurements [7]. Fink and Kumar use a similar approach to allow a robot to automatically localize a base-station [8]. However, these approaches, also demonstrated in a simulated robotic domain in [4], assume signals are uniformly correlated in all directions without regard to the location of attenuating objects. Nonetheless, this approach has successfully been applied to prediction of signal strength in real-world environments, and forms the basis for one of the methods we benchmark in this paper. Our extension of their approach addresses the more general problem when all nodes are potentially mobile. This makes the problem significantly more challenging because of a substantial reduction in data density.

An alternative approach to predicting signal strength, using principles from the field of tomography, is to explicitly estimate the location and properties of attenuating objects affecting signal propagation. Knowing the location of attenuating objects enables future signal strength to be predicted, even when they are not spatially adjacent to previous measurements. The central challenge with this technique is that directly estimating the positions of attenuating objects from signal-strength measurements results in an ill-posed estimation problem. This is because for any given set of signal measurements, there are many possible world configurations which explain the data. Prior work by Wilson and Patwari has explored the capabilities of RF tomography to recover the motion of moving people inside a region bounded by a large number of regularly placed radios [9]. Their work estimated the derivative of the attenuation over a grid of pixels inside the perimeters of radios to extract the positions of moving people. By only estimating the derivative, their signal model was simplified, allowing a single calibration step to be used to later recover the target’s position. Despite the large number of radios (28), their estimation problem was still ill-posed, requiring application of regularization techniques to make computing the attenuation-derivative possible [10]. However, their approach does not directly apply to the case where multiple robots traverse arbitrary paths, since in such cases no prior calibration is possible, and the number of links constraining the attenuation computed at each pixel can vary considerably. Our signal-strength prediction algorithm MATTE addresses these problems by employing a novel fusion of laser range-finder data with signal-strength measurements to better constrain the attenuations estimated at each pixel. Furthermore, we show that our approach scales to environments hundreds of meters in diameter.

### III. BACKGROUND

Standard macroscopic models of RF propagation describe the expected signal strength, which depends on the location of the receiver  $\vec{r}$  and transmitter  $\vec{t}$ , as having three main components [2]:

$$y_{dBm} = \underbrace{L_0 - a \log_{10}(\|\vec{r} - \vec{t}\|)}_{\text{path-loss}} - \underbrace{g(\vec{r}, \vec{t})}_{\text{shadowing}} - \underbrace{\epsilon}_{\text{multipath}} \quad (1)$$

In simplified, ideal environments, signal-strength can be determined purely by path-loss which has two degrees of freedom:  $L_0$  corresponds to the power of the transmitter and  $a$  is the path-loss exponent that determines how quickly the signal attenuates with distance. In real environments both shadowing and multipath can additionally affect the signal strength: shadowing corresponds to the attenuation a signal experiences as it passes through dense objects, and multipath corresponds to amplification or cancellation which occurs when waves travel multiple paths of different lengths from source to destination. In general, multipath is very difficult to predict because it results from complex reflection and diffraction interactions. Shadowing, on the other hand, is easier to predict, since the scale of its effects are larger and more spatially coherent. For the remainder of this paper, we will focus on models which predict path-loss and shadowing but ignore multipath.

Robots are able to measure the signal-strength from each other robot during the robot’s mission. These noisy values form a vector  $\hat{y}$ , with corresponding vectors of positions  $R$  and  $T$  containing all pairs of receiver and transmitter positions, respectively. We treat predicting  $y$  as a linear regression problem. For the case of the simple path-loss model, estimating the two variables ( $L_0, a$ ) from the data is sufficient to predict signal measurements for arbitrary positions. This can be done using standard least-squares approaches: if  $x = [L_0 \ a]$  and the  $i^{\text{th}}$  row of  $A$  is  $[1 \ \log_{10}(\|\vec{r}_i - \vec{t}_i\|)]$  then the least-squared error estimate for  $x$  is  $\bar{x} = (A^T A)^{-1} A^T \hat{y}$ . Path-loss-only models are useful due to their simplicity and correspondence with theoretical propagation equations: the low degree of freedom reduces the chance of over-fitting. However, as this model ignores shadowing effects, its application in environments with varying attenuation is limited. In our evaluation, we will use this simple log-fit as a baseline.

Correlative and tomographic methods include the same path-loss model, but also explicitly incorporate shadowing, allowing for improved predictive performance. The way in which shadowing is captured varies between the two types of models. In the tomographic case, the shadowing function,  $g(\cdot)$  is computed by integrating the effect of all attenuators between transmitter and receiver. In practice, we model the individual attenuation of a grid of infinitely tall columns, represented by a regular 2D grid of pixels. If the signal passes through a series of  $p$  discretized pixels, the shadowing is computed as [9]:

$$g(\vec{r}, \vec{t}) = \sum_i^p w_i v_i \quad (2)$$

where  $v_i$  is the attenuation in the  $i^{\text{th}}$  pixel and  $w_i$  is the importance weight of that pixel for that signal’s path. (e.g. in our implementation, the weights  $w_i$  correspond to the length of the line between transmitter and receiver which is contained in the pixel.) In other words, we compute attenuation per pixel by assuming that signal strength is reduced linearly according to the sum of the pixels along the path between receiver and transmitter. Given a set of signal strength measurements, we can attempt to find the attenuation values  $x_i$  of each pixel using a similar least-squares approach as described above. However, there are many possible attenuation assignments to the pixels which can adequately explain the signal strength measurements, resulting in an under-determined system of equations. In the next section we will discuss applying regularization techniques to encode a prior that prefers real-world environments, thereby over-constraining the system of equations.

In the correlative case, the shadowing component is considered to be uniformly correlated in all directions, allowing predictions to be made by making inferences from spatially-proximate training points. In particular, prior approaches have had success modeling shadowing using a Gaussian-process (GP) [2], [7].

In the case of a fixed base-station (located at  $\vec{b}$ ), we first use the log-fit as a mean function, and then use a standard squared-exponential kernel to specify the expected covariance between signals at two locations,  $x$  and  $x'$ :  $k(x, x') = \sigma_f^2 \exp\{-\frac{\|x-x'\|}{l^2}\}$ . Using this covariance (kernel) function, we can apply standard GP regression techniques to make predictions for a set of sample points [11]:

- 1) **Compute Log-fit:** Fit  $L_0$  and  $a$  using least-squared approach described above.
- 2) **Fit Hyperparameters:** Choose correlation distance  $l$  and function variance  $\sigma_f$  to maximize likelihood of training data.
- 3) **Compute Covariances:** Compute covariance matrix  $K_y$  of training data, and covariance vector  $k_*$  of sample points with respect to training data.  $K_y = K_f + \sigma_d^2 I$ , where each entry  $k_{i,j} \in K_y = k(x_i, x_j)$ .
- 4) **Evaluate prediction:**  $y_{\text{sample}} = k_*^T K_y^{-1} (y_{\text{observed}} - y_{\text{log}})$

These correlative methods have good predictive performance, especially in the case when many training points are available. The main shortcomings of this method are that prediction assumes signals are uniformly correlated in all directions – an assumption which breaks down in the presence of discrete attenuating objects. Direct extension of this method to the case of mobile nodes is also problematic, since training points are in  $\mathbb{R}^4$ , requiring significantly more data to achieve the same performance. Finally, this method is computationally expensive, requiring the inversion of a matrix whose dimension is determined by the number of training points.

#### IV. METHODS

In this section, we extend both the tomographic and correlative methods to the case of multiple mobile nodes. We will describe our modifications to these existing approaches, and

introduce our new approach, MATTE, which also leverages other sensors to infer the location of attenuating objects.

##### A. Correlative Methods for Mobile Nodes

Extending previous approaches of correlative prediction to the case of moving nodes exacerbates the data-sparsity problem. Instead of producing signal predictions for  $\mathbb{R}^2$  (all points in the plane), we now must produce predictions in  $\mathbb{R}^4$  (all possible *pairs* of points in the plane). Since we can’t increase the number of signal-strength observations that robots make without slowing down the speed of exploration, this means we have significantly reduced data density. However, we were able to mitigate this problem somewhat by recognizing that our signal-strength models are symmetric with regard to where the transmitter and receiver are – that is, we assume the signal strength is the same from robot A to B as it is from robot B to A. This observation allows us to construct a symmetric distance function,  $d_{s4}$ , which effectively doubles the data density:

$$d_{s4}(\vec{r}_a, \vec{t}_a, \vec{r}_b, \vec{t}_b) = \min \left\{ \begin{array}{l} \|\vec{r}_a - \vec{r}_b\| + \|\vec{t}_a - \vec{t}_b\| \\ \|\vec{r}_a - \vec{t}_b\| + \|\vec{t}_a - \vec{r}_b\| \end{array} \right. \quad (3)$$

In other words,  $d_{s4}$  is a distance metric for pairs of lines that is invariant to rotations of 180 degrees.

In practice, many thousands of observations may be available for use in training the Gaussian process. Due to the  $O(n^3)$  computation cost of matrix inversion, it quickly becomes impractical to include all training data. However, prediction performance is improved as more training points are used. In the case where  $K$  is very sparse, the inversion can be done more quickly, enabling use of more training data. However, the squared exponential kernel is not sparse; two measurements will never have a covariance of exactly zero, even if they are very far apart. By modifying the kernel to have compact support – that is, the kernel is exactly zero once some distance  $\Theta$  is reached – the matrix  $K$  becomes sparse [12]:

$$k_s(x, x') = \max(0, 1 - \frac{\|x - x'\|}{\Theta})^\gamma * k(x, x') \quad (4)$$

The resulting sparsity allows computing  $K^{-1}$  much faster, enabling the use of more training points. Although sparse kernels generally have worse performance, we found in practice that the speed improvement enabled the incorporation of enough extra training points to achieve a net improvement in performance. This sparsification also increases the number of hyper-parameters that must be estimated to a total of 5. In principle, we can learn these additional parameters the same way we learn the parameters for the original covariance function. However, increasing  $\Theta$  will always result in a lower training error and an increased computation due to a less-sparse covariance matrix. To limit worst-case computation time, we do an offline parameter sweep to estimate the best  $\Theta$  and  $\gamma$  subject to a fixed CPU budget.

Together, the symmetric distance metric and the forced sparsification of the correlation function enabled us to adapt existing correlative prediction techniques to the case of

multiple mobile robots, and achieve results competitive with our proposed method. These additions to existing methods form the Multi-robot Gaussian Process “MRGP” method which we include in our evaluation.

### B. Multi-robot Tomography (MRT)

Conceptually, extension of the tomographic methods to the case of multiple robots is relatively straight-forward. The attenuation of each pixel through which a signal passed is estimated in least-squares fashion. The number of pixels,  $p$ , is determined by the grid size such that the pixels completely fill the workspace of the robots. For a grid of width  $w$  and height  $h$ ,  $p = w \times h$  and pixels are labeled  $v_{0,0} \cdots v_{w-1,h-1}$ . In addition, the two path-loss parameters,  $L_0$  and  $a$  are also estimated simultaneously, resulting in  $n = p + 2$  unknowns:  $x = [L_0 \ a \ v_{0,0} \cdots v_{w-1,h-1}]$ . Each of the  $m$  signal strength observations  $y_i$  provide one equation partially constraining a subset of the pixels in addition to the path loss parameters (See Eqns 2 and 1). Together these equations are stacked to form the rows of an  $n \times m$  matrix  $A$ . If  $A$  is full rank,  $x = (A^T A)^{-1} A^T y$ . However, even though  $m$  is generally greater than  $n$ ,  $A$  remains rank deficient, resulting in many possible solutions for  $x$ . Related approaches have solved this by using Tikhonov regularization that enforces smooth changes in attenuation between neighboring pixels [9]. This corresponds to constructing a Tikhonov matrix  $\Gamma$  whose rows correspond to equations of the form  $v_{i,j} - v_{i+1,j} = 0$  and  $v_{i,j} - v_{i,j+1} = 0$  for each  $i, j < w, h$ . The over-constrained solution for  $x$  now takes the form

$$x = (A^T A + \lambda^2 \Gamma^T \Gamma)^{-1} A^T y \quad (5)$$

where  $\lambda$  is a parameter to determine the weight of the smoothness constraints in  $\Gamma$  relative to the observation equations in  $A$ . In contrast to the correlative methods, which scale  $O(m^3)$  with respect to the number of observations, the tomographic methods scale  $O(p^3)$  with respect to the number of pixels. However, unlike the correlative methods,  $A^T A$  is naturally sparse since each pixel is only jointly constrained with a small number of other pixels. This means that sparse matrix-inversion methods, such as a Cholesky decomposition, can perform significantly better than  $O(p^3)$ . In practice,  $A^T A$  is still dense enough (95% zeros) that we are limited to solving grids on the order of  $100 \times 100$  (10000 pixels). For the datasets in our evaluation, this translates to a grid size between 2 and 4 meters. Such large grid sizes poorly approximate the sharp spatial changes in attenuation found in real environments, such as at the border between a building and neighboring free space. Furthermore, in order to avoid over-fitting,  $\lambda$  must be set large enough to allow very little spatial gradient in the attenuation of each pixel (See Fig. 1). In areas where attenuation changes quickly, the predictive power is limited. Our implementation of this approach, Mutli-robot Tomography (MRT), is included in our evaluation.

### C. Multi-sensor ATTenuation Estimation (MATTE)

The approaches we’ve discussed so far focus solely on using signal-strength readings for prediction of future mea-

surements. Many robots already carry additional sensors for mapping and navigation. Intuitively, since the physical structure of an environment (the position of walls and other solid surfaces) influences signal propagation, a map of the environment should help predict signal propagation. This is especially true for sensors like laser range-finders, which tend to have a range of at least 10-30 meters. Incorporating a map does not solve the problem completely though, since the attenuation of a wall depends on the material and LIDAR generally can’t distinguish between a reinforced concrete wall and drywall.

Specifically, we propose the use of occupancy grids derived from laser range-finders to provide a more informed regularization constraint to tomographic methods. The occupancy grids collected by our robots label the world with three classes: known free space, known structure and unknown. This information can provide a much better prior about the attenuating properties of the environment – for example, we generally expect areas which are marked as free space to pass signals with little interference. Similarly, knowing the location of structures can provide a prior about where attenuation should increase dramatically. Besides providing a better prior about the magnitude of attenuation, the occupancy grids also provide a much finer view of the environment. Using the MRT method, we are typically limited to coarse grid sizes (e.g. 3 m for our datasets) due to computational constraints. On the other hand, LIDAR-based occupancy grids can be computed cheaply even for fine-grained grid sizes (we used 10 cm grid in our experiments).

Of the many ways of incorporating this data, we explored explicitly estimating the attenuation of each of the three classes separately. A simple approach to this problem is to fit a single attenuation value to all pixel of the same class. For example, known free space might have an attenuation of  $-0.01dBm$  per meter, whereas structure (e.g. walls) could have an attenuation of  $-0.3dBm$  per meter, and unknown space could be approximated as somewhere in between. This approach is notable in its simplicity – it has no parameters to tune and only 5 degrees of freedom to fit the observed data: two for path loss parameters and three more for the attenuation assigned to each class in the occupancy grid. In general, however, it is a poor assumption that all objects detected by the robots as structure will have the same attenuation. For example, a wooden fence and a brick building have very different effects on a signal, but can appear very similar to a robot’s laser range finder. As our initial experiments confirmed, this model performed poorly in explaining real-world data.

Another way to use the occupancy grids is to better inform regularization methods, for example by enforcing smoothness constraints only between neighboring pixels of the same class. Using a more flexible model is difficult, however, since individually estimating the attenuation for each pixel in the fine-grained occupancy grid results in a poorly constrained, computationally prohibitive problem. The largest map we present has nearly 5 million pixels! Downsampling to a coarser resolution can make the problem

more tractable, but reduces the ability to incorporate features such as doors or walls which may be lost by resampling.

---

**Algorithm 1** MATTE\_UPDATE(trainPoints, map)
 

---

```

init  $A, b$  from  $bounds(map)$ 
for all  $p \in \text{trainPoints}$  do
   $r \leftarrow [1, \log_{10}(\|p_r - p_t\|), 0 \dots, 0]$ 
  for all locations  $v_i \in \text{losPath}(p)$  do
     $c \leftarrow \text{map.class}(v_i)$ 
     $r(v_i, c) \leftarrow 1$ 
  end for
   $A \leftarrow \text{appendRow}(r)$ 
   $b \leftarrow \text{appendElement}(p_z)$ 
end for
 $\{A, b\} \leftarrow \text{appendRegularization}(A, b)$ 
 $x \leftarrow \text{cholesky.solve}(A^T A, A b)$ 
store  $x$ 
store  $map$ 
return

```

---



---

**Algorithm 2** MATTE\_PREDICT(testPoints)
 

---

```

retrieve  $x$ 
 $z \leftarrow \text{vector}(\text{testPoints.len})$ 
for all  $p \in \text{testPoints}$  do
   $z_p \leftarrow x(0) + \log_{10}(\|p\|)$ 
   $z_s \leftarrow \text{attenuationIntegral}(x, p, \text{map})$ 
   $p \leftarrow \text{append}(p, z_p + z_s)$ 
end for
return  $p$ 

```

---

Instead, we propose a method that benefits from the detailed resolution of the map, but has computational complexity comparable to the MRT method. Our technique, called MATTE, separately estimates spatially-varying attenuation for each of the three classes in the occupancy grid. These estimates allow querying any point in the environment and determining, for example, *if* there was structure there, what attenuation it would have. When predicting signal-strength, we first examine pixels in the occupancy grid to determine which class they are, and then perform a look-up on the appropriate attenuation estimate. The benefit of this approach is that we can estimate the spatially-varying attenuation at a resolution independent of the resolution required to adequately map the environment structure. Typically grid sizes of 10 to 20 cm are required to preserve the presence of doors or thin walls – but attenuation rates of a particular material, (e.g. building) tend to vary at a much slower rate. The computational implications of this approach are significant – whereas the resolution of the occupancy grid is typically 10 cm, we have achieved good results with a grid size of 4 m for the spatially varying per-class attenuations, resulting in a reduction by a factor of 1600 in the number of unknowns.

Similar to the previous tomographic approach, we simultaneously estimate the path-loss parameters and the spatially

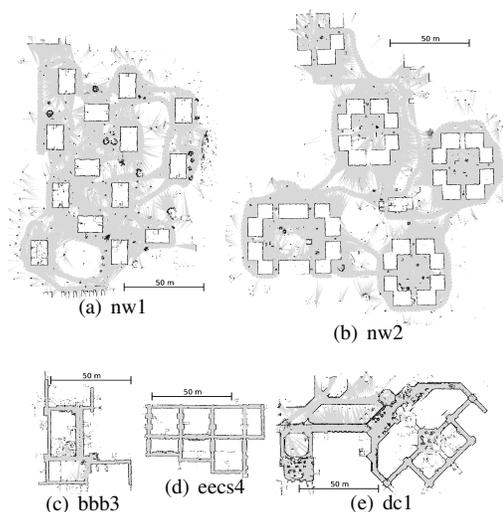


Fig. 2. Maps collected in each of the five datasets we used in our evaluation. From left to right they are: two outdoor sections of the Northwood residential housing complex, the 3rd floor of the Bob and Betty Beyster Building, the 4th floor of the Electrical Engineering and Computer Science Building and the first floor of the Duderstadt Library. White is unexplored, gray is known free-space and black is structure. Each evaluation set consists of back-to-back traversals of three of these environments; meta-parameters are tuned using the remaining two.

varying attenuation. For a workspace  $w \times h$  pixels in dimensions, we estimate a total of  $2 + 3 * w * h$  variables that comprise  $x$ . Although this represents a threefold increase in the number of degrees of freedom over our previous approach,  $A^T A$  is generally more sparse (e.g. 99% sparse in our datasets), ultimately requiring less time to compute a result. An overview of the update and prediction steps are shown in Algorithms 1 and 2. The MATTE algorithm has several important parameters which can either be tuned by hand or estimated automatically from training data. The most important parameter is the relative weight  $\lambda$  given to the smoothing regularizer we described in Eqn. 5. In addition we introduce three parameters to govern the expected attenuation in areas where we have not collected any signal data: a weighting factor  $\phi$  determines the relative weight of this prior, and prior attenuation values in units of  $dBm$  per meter for each class are  $\rho_f, \rho_s, \rho_u$ , for free-space, structure, and unknown space respectively. In practice, we set  $\rho_s = \rho_u$ , allowing free space to have a distinct prior from other areas.

An additional potential advantage of MATTE is that new occupancy grids can be incorporated cheaply, as the underlying spatially varying attenuation does not necessarily need to be updated when the map changes. For example, our approach might encode the knowledge that a structure-class pixels in a particular area tend to have an attenuation of  $X$   $dBm$  per meter. If the map of that area is later expanded, that information will be able to provide an estimate for the attenuation expected there, without needing to recompute  $x$ . This allows our approach to potentially be adapted to run online.

Eval.	Method	Meta-parameters
1	MATTE MRT	Train: {nw2,bbb3} Test: {eecs4,dc1,nw1} $\lambda=2.0$ $\phi=.2$ $\rho_f=-.1$ $\rho_{u,s}=-.3$ $\lambda=0.1375$
2	MATTE MRT	Train: {eecs4,dc1} Test: {bbb3,nw2,nw2} $\lambda=1.225$ $\phi=0.05$ $\rho_f=0.2125$ $\rho_{u,s}=-0.2875$ $\lambda=0.1125$
3	MATTE MRT	Training: {bbb3,eecs4} Test: {nw2,nw1,dc1} $\lambda=2.0$ $\phi=0.275$ $\rho_f=0.2875$ $\rho_{u,s}=-0.4875$ $\lambda=0.1250$
4	MATTE MRT	Train: {nw1,eecs4} Test: {nw2,dc1,bbb3} $\lambda=1.38$ $\phi=0.3$ $\rho_f=-.1$ $\rho_{u,s}=-.3$ $\lambda=0.1125$
5	MATTE MRT	Train: {dc1,nw1} Test: {eecs4,bbb3,nw2} $\lambda=2.0$ $\phi=0.1$ $\rho_f=-0.15$ $\rho_{u,s}=-.3$ $\lambda=0.5875$

TABLE I

AUTOMATICALLY DETERMINED PARAMETER SETTINGS FOR EACH EVALUATION USING COORDINATE DESCENT ON THE TRAINING SET. Some parameters were fixed for performance reasons, including the grid sizes for MATTE and MRT methods, at 4.0 meters and 3.0 meters respectively. For MRGP, the number of training points was set to  $\min(600, .05 * m)$ ,  $\Theta$  was fixed to 20.0, and  $\gamma$  to 3. Also  $\sigma_d$  and  $\sigma_f$  were fixed for numerical stability reasons at 1.5 and 1.0 respectively. For MRGP, coordinate descent fixed  $l$  to 28.75 on all training data.

## V. EVALUATION

We evaluated the three primary methods we have discussed to determine which methods were most successful at predicting real-world signal-strength measurements. The main purposes of our evaluation is to show that previous signal-strength predictions methods using a fixed base station can be extended to the more general case where all nodes are mobile. We also show that in many cases sensor data can be effectively leveraged to further improve signal-strength predictions. Several of our models have many degrees of freedom, so our evaluation also seeks to show that the methods we present can generalize well from previous observations to the prediction of future signal measurements.

### A. Test Apparatus

We used a robot platform our lab custom designed for urban reconnaissance [13]. We outfitted three of our 14 robots with additional 2.4 GHz TP-Link WiFi radios which were programmed to report signal-strength measurements at 20 Hz. Our robots are equipped with 3D laser range finders, in addition to IMUs and odometers, enabling them to produce high-quality globally consistent maps using Simultaneous Localization and Mapping (SLAM) algorithms [14], [15], [16]. This capability allows us to quickly collect large amounts of signal-strength data that is co-registered with a global grid.

We used these three robots to collect a series of 5 datasets in various indoor and outdoor environments, spanning a total area over 40,000  $m^2$ . The datasets consist of three indoor environments, **bbb3**, **eecs4** and **dc1** for short, as well as two

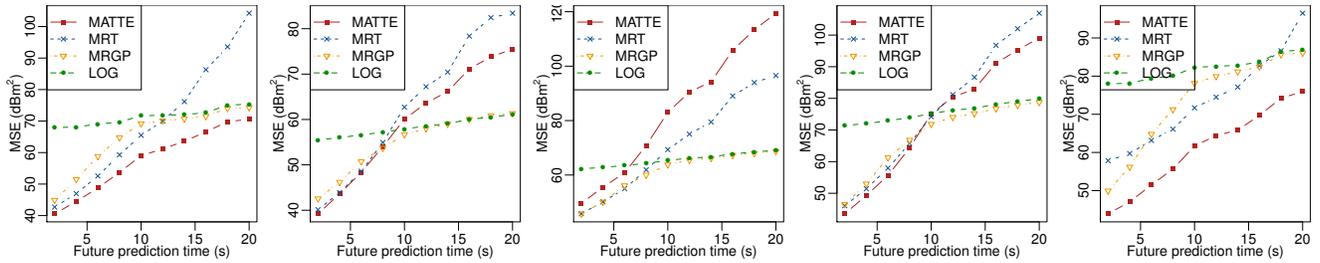
larger outdoor environments, abbreviated **nw1** and **nw2** (See Fig. 2 for details). Real-world rescue robots must operate in mixed indoor-outdoor urban environments. In order to mimic these conditions, we randomly selected a sequence of three of the environments for testing, guaranteeing a mix of indoor and outdoor datasets. Each of the datasets consist of exploration missions so very few sensor measurements can be considered duplicates, since robots do not tend to retrace their steps. This enables us to explicitly test the predictive performance of each method, rather than their recall abilities.

For each randomly-selected set of three areas, we replayed the signal-strength observations and corresponding maps in two second increments. After each step, the methods were tested on their prediction performance of the next 10 intervals of future signal strength data (e.g from 0-2 seconds, 2-4 seconds, 4-6 seconds, etc.). As the traversal played out, methods had access to an increasing amount of training data, but the amount of testing data was relatively similar at each step. Several of the methods also have meta- or hyper- parameters which need to be tuned before the start of a traversal. We used the two remaining areas not selected for the test set to train meta parameters using a compass search. We include results for 5 such randomly generated traversals. The parameters selected by the compass search for each method on each evaluation are shown in Table I. All of the methods we presented were tuned to use comparable amounts of computational resources. While each method exhibits significantly different asymptotic growth, we set parameters which resulted in roughly equal CPU use over the course of an evaluation set. Run-times for MATTE ranged between 2 seconds to process the bbb3 dataset up to 30 seconds for the significantly larger nw2 dataset. For the GP method, times varied between 15 seconds for bbb3 datasets and 20 seconds on nw2.

### B. Results

The testing results for the five evaluation sets are displayed in Fig. 3. As expected, all methods are better at making short-term rather than long-term predictions. This is a result of the fact that measurements nearer in the future are more likely to be similar to existing measurements, or the fact that attenuating objects impacting observations in the near future are likely to be correlated with sensor data the robots have just now collected.

All of the methods we have implemented show competitive performance, especially for predictions between 0 and 10 seconds in the future. However, MATTE significantly outperforms the other methods in evaluations 1 and 5. In evaluations 2 and 4, it performs comparable to the tomographic method. However, in evaluation 3, our method exhibits worse performance than the other methods. An examination of the meta-parameters determined via compass search show evidence of over-fitting the training set, which by nature of our randomly selected test sets, happened to both be exclusively indoors. This is manifest as a positive attenuation prior for free-space ( $\rho_f$ ), consistent with the wave-guide effect sometimes seen in hallways. In the other evaluation



(a) Eval 1 {eecs4,dc1,nw1} (b) Eval 2 {bbb3,nw2,nw1} (c) Eval 3 {nw2,nw1,dc1} (d) Eval 4 {nw2,dc1,bbb3} (e) Eval 5 {eecs4,bbb3,nw2}

Fig. 3. MSE of prediction accuracy for each method evaluated on 5 randomly chosen sequences over the testing datasets. Error bars are bounded by  $0.8 \text{ dBm}^2$  and are omitted for clarity. Our proposed method, MATTE, is compared to extensions of previous methods MRGP and MRT. Our method exploits generally performs better for near-term predictions where the robots’ sensor data provides an informative prior. Poor performance of MATTE in evaluation 3 is attributable to the randomly generated training set which contains only indoor datasets, while the testing set contains both indoor and outdoor sets.

sets, there was at least one outdoor dataset used for training, which helped to mitigate this type of over fitting. In most of the evaluation sets, we see that MATTE can out-perform the correlative method for short-term predictions up to ten seconds in the future. For longer-term predictions, where training data is mostly useless, it is difficult to beat the log baseline.

## VI. CONCLUSION

In this work, we have explored the difficult problem of predicting signal strength when all nodes are actively exploring a variety of real world environments. This problem is challenging because robots typically only sample the possible signal propagation paths very sparsely, making generalization beyond a naive log-fit difficult. We extended several existing signal strength prediction methods to the case of multiple robots conducting exploration-style missions in mixed indoor-outdoor urban environments. We additionally suggested a new method for signal prediction which uses additional sensors already in use by many autonomous robots to better estimate the regions of attenuation in an environment. Our methods performs competitively with the other approaches, and in some cases performs much better. Key to our approach is the ability to estimate the attenuation properties of the environment at a coarse level, but still use a fine-grained spatial model of an environment derived from additional sensors. It is interesting that both correlative and tomographic methods exhibit such similar performance, as they employ very different approaches to the same problem. This suggests that both methods are exploiting similar aspects of signal strength propagation.

Our work has important applications to autonomous robot teams which are collaborating to achieve a joint goal. By introducing more advanced signal-strength modeling techniques, such teams can better predict when robots can expect to communicate, allowing them to plan their future actions

by explicitly including communication as a constraint.

In the future we look to continue exploring methods for effectively predicting signal strength. In particular, it may be possible to leverage additional sensing modalities to improve prediction. Incorporating this predictor in an online planning system would also serve to further validate our approach.

## REFERENCES

- [1] M. Rooker and A. Birk, “Multi robot exploration under the constraints of wireless networking,” *Control Engineering Practise*, vol. 15, no. 4, pp. 435–445, 2007.
- [2] N. Michael, M. Zavlanos, V. Kumar, and G. Pappas, “Maintaining connectivity in mobile robot networks,” *Experimental Robotics*, 2009.
- [3] G. Hollinger and S. Singh, “Multi-robot coordination with periodic connectivity,” in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*. IEEE, 2010, pp. 4457–4462.
- [4] A. Ghaffarkhah and Y. Mostofi, “Channel learning and communication-aware motion planning in mobile networks,” in *American Control Conference (ACC)*, July 2010, pp. 5413–5420.
- [5] W. Burgard, M. Moors, C. Stachniss, and F. Schneider, “Coordinated multi-robot exploration,” *Robotics, IEEE Transactions on*, vol. 21, no. 3, pp. 376–386, June 2005.
- [6] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.
- [7] M. Malmirchegini and Y. Mostofi, “On the spatial predictability of communication channels,” *Wireless Communications, IEEE Transactions on*, vol. 11, no. 3, pp. 964–978, March 2012.
- [8] J. Fink and V. Kumar, “Online methods for radio signal mapping with mobile robots,” in *International Conference on Robotics and Automation*, 2010.
- [9] J. Wilson and N. Patwari, “Radio tomographic imaging with wireless networks,” *Mobile Computing, IEEE Transactions on*, vol. 9, no. 5, pp. 621–632, May 2010.
- [10] J. Wilson, N. Patwari, and F. Vasquez, “Regularization methods for radio tomographic imaging,” in *2009 Virginia Tech Symposium on Wireless Personal Communications*, 2009.
- [11] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2005.
- [12] H. Zhang, M. Genton, and P. Liu, “Compactly supported radial basis function kernels,” Tech. Rep., 2004.
- [13] E. Olson, J. Strom, R. Morton, A. Richardson, P. Ranganathan, R. Goedel, M. Bulic, J. Crossman, and B. Marinier, “Progress towards multi-robot reconnaissance and the MAGIC 2010 competition,” *Journal of Field Robotics*, 2012.
- [14] J. Strom and E. Olson, “Occupancy grid rasterization in large environments for teams of robots,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2011.
- [15] E. Olson, “Real-time correlative scan matching,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Kobe, Japan, June 2009.
- [16] F. Dellaert, “Square root SAM,” in *Proceedings of Robotics: Science and Systems (RSS)*, Cambridge, USA, June 2005.