

A kernel-based learning approach to ad hoc sensor network localization

XUANLONG NGUYEN¹

MICHAEL I. JORDAN^{1,2}

and

BRUNO SINOPOLI¹

¹Department of Electrical Engineering and Computer Science

² Department of Statistics

University of California, Berkeley

We show that the coarse-grained and fine-grained localization problems for ad hoc sensor networks can be posed and solved as a pattern recognition problem using kernel methods from statistical learning theory. This stems from an observation that the kernel function, which is a similarity measure critical to the effectiveness of a kernel-based learning algorithm, can be naturally defined in terms of the matrix of signal strengths received by the sensors. Thus we work in the natural coordinate system provided by the physical devices. This not only allows us to sidestep the difficult ranging procedure required by many existing localization algorithms in the literature, but also enables us to derive a simple and effective localization algorithm. The algorithm is particularly suitable for networks with densely distributed sensors, most of whose locations are unknown. The computations are initially performed at the base sensors and the computation cost depends only on the number of base sensors. The localization step for each sensor of unknown location is then performed locally in linear time. We present an analysis of the localization error bounds, and provide an evaluation of our algorithm on both simulated and real sensor networks.

Categories and Subject Descriptors: C.2.1 [**Network Architecture and Design**]: Network communications, Distributed networks; I.5 [**Pattern Recognition**]: Applications

General Terms: Ad hoc wireless sensor networks, Localization, Kernel methods, Statistical machine learning

1. INTRODUCTION

A sensor network can be viewed as a distributed pattern recognition device. In the pattern recognition approach, rather than transforming sensor locations and sensor readings into Euclidean, world-centric coordinates, we work directly with the (non-Euclidean) coordinate system given by the physical sensor readings themselves. Using the methodology of “kernel functions,” the topology implicit in sets of sensor readings can be exploited in the construction of signal-based function spaces that

Author’s address: XuanLong Nguyen, Computer Science Division, 387 Soda Hall, University of California, Berkeley, CA 94720-1776.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 0000-0000/20YY/0000-0001 \$5.00

are useful for the prediction of various extrinsic quantities of interest, using any of a variety of statistical algorithms for regression and classification. In the current paper we illustrate this approach in the setting of a localization problem [Hightower and Borriello 2000; Bulusu et al. 2000; Savarese et al. 2002].

The localization problem that we study is that of determining the location of a (large) number of sensors of unknown location, based on the known location of a (small) number of base sensors. Let X_1, \dots, X_m denote a set of m sensors, and let x_i denote the position in \mathbb{R}^2 of sensor X_i . Suppose that the locations of the first n sensors are known, i.e., $X_1 = x_1, \dots, X_n = x_n$, where $n \ll m$. We want to recover the positions of X_{n+1}, \dots, X_m solely on the basis of the receive/transmit signals $s(x_i, x_j)$ between pairs of sensors.

An important characteristic of radio or light signal strength is the relationship of the signal attenuation as a function of distance [Seidel and Rappaport 1992]. For instance, for radio signals in an idealized environment, given that the sending and receiving antennas are focused on the same radio frequency, we have:

$$s \propto Pd^{-\eta}, \quad (1)$$

where $\eta > 2$ is a constant, and P is the sending signal voltage. Such relationships provide the basis for a variety of localization algorithms in the literature, which consist of two main steps: (1) a ranging procedure which involves estimating the distance from a sensor to another sensor based on the signal strength of the signals transmitted/received between the two, and (2) a procedure that recovers the locations of the sensors based on their pairwise distance estimates either by triangulation or by least-squares methods [Priyantha et al. 2000; Girod and Estrin 2001; Savvides et al. 2001; Whitehouse 2002]. Unfortunately, however, the idealized model in Eq. (1) can be highly inaccurate due to variability caused by multipath effects and ambient noise interference as well as device-specific factors such as the frequencies of node radios, physical antenna orientation, and fluctuations in the power source [Bulusu et al. 2000; Priyantha et al. 2000]. Methods based on ranging inherit these inaccuracies and improvements are possible only if difficult problems in signal modeling are addressed.

In this paper we propose a method that bypasses the ranging step altogether. We show that it is possible to pose a coarse-grained localization problem as a discriminative classification problem that can be solved using tools from the statistical machine learning literature. Fine-grained localization is then achieved by a second application of the coarse-grained localization technique. Our localization algorithm thus involves two phases. First, there is a training phase that chooses discriminant functions for classifying positions using arbitrarily constructed target regions. This phase is performed either on-line at the base stations, or taken off-line, and takes $O(n^3)$ computational time, where n is the number of base sensors. Hence, our assumption is that the base sensors have sufficient power and processing capability (indeed, these are also the nodes that might have GPS-capability to determine their own exact locations). Second, once the training phase is completed, other location-unknown low-power sensors can determine their own position locally, and the computation takes only $O(n)$ time for each of these sensors.

Our approach makes use of kernel methods for statistical classification and regression [Schölkopf and Smola 2002], an example of which is the “support vector

machine (SVM).” Central to this approach is the notion of a *kernel function*, which provides a generalized measure of similarity for any pair of entities (e.g., sensor locations). The functions that are output by the SVM and other kernel methods are sums of kernel functions, with the number of terms in the sum equal to the number of data points. Kernel methods are examples of *nonparametric* statistical procedures—procedures that aim to capture large, open-ended classes of functions.

Kernel functions typically used in practice include Gaussian kernels and polynomial kernels. A technical requirement of these functions is that they are positive semidefinite, which is equivalent to the requirement that the $n \times n$ *Gram matrix* formed by evaluating the kernel on all pairs of n data points is a positive semidefinite matrix. Intuitively, this requirement allows a kernel function to be interpreted as a generalized measure of similarity. The kernel function imposes a topology on the data points which is assumed to be useful for the prediction of extrinsic quantities such as classification labels.

Given that the raw signal readings in a sensor network implicitly capture topological relations among the sensors, kernel methods would seem to be particularly natural in the sensor network setting. In the simplest case, the signal strength would itself be a kernel function and the *signal matrix* $(s(x_i, x_j))_{ij}$ would be a positive semidefinite matrix. Alternatively, the matrix may be well approximated by a positive semidefinite matrix (e.g., a simple transformation that symmetrizes the signal matrix and adds a scaled identity matrix may be sufficient). More generally, and more realistically, derived kernels can be defined based on the signal matrix. In particular, inner products between vectors of received signal strengths necessarily define a positive semidefinite matrix and can be used in kernel methods. Alternatively, generalized inner products of these vectors can be computed—this simply involves the use of higher-level kernels whose arguments are transformations induced by lower-level kernels. In general, hierarchies of kernels can be defined to convert the initial topology provided by the raw sensor readings into a topology more appropriate for the classification or regression task at hand. This can be done with little or no knowledge of the physical sensor model.

Our focus is on the discriminative classification problem of locating sensors in an ad hoc sensor network. It is worth noting that similar methods have been explored recently in the context of tracking one or more objects (e.g., mobile robots) that move through a wireless sensor field.¹ Systems of this type include Active Badge [Want et al. 1992], [Ward et al. 1997], RADAR [Bahl and Padmanabhan 2000], Cricket [Priyantha et al. 2000], and UW-CSP (cf. [Li et al. 2002]). In Bahl and Padmanabhan [2000], the authors describe a simple nearest neighbor classification algorithm to obtain coarse localization of objects. Most closely related to our approach is the work of Li et al. [2002] in which a number of classification algorithms are used for tracking moving vehicles, including k -nearest neighbor and support vector machines. We elaborate on the connections between this work and ours in the description of our algorithm.

¹The alternative to discriminative classification is classification using *generative* probabilistic models. This is a well-explored area that dates back to contributors such as Wiener and Kalman. Recent work in this vein focuses on the distributed and power-constrained setting of wireless sensor networks (e.g. [Sheng and Hu 2003; D’Costa and Sayeed 2003]).

The paper is organized as follows. We begin with a brief background of classification using kernel methods, and motivate our application of kernel methods to the localization problem based on sensor signal strength. Next, the localization algorithm and its error analysis are described. We then present details of the implementation of the algorithm and its computational cost, followed by an evaluation of our algorithm with simulated and real sensor networks. Finally, we present our conclusions in the final section.

2. CLASSIFICATION USING KERNEL METHODS

In a classification algorithm, we are given as training data n samples $(x_i, y_i)_{i=1}^n$ in $\mathcal{X} \times \{\pm 1\}$, where \mathcal{X} denotes the input space. Each y_i specifies whether the data point $x_n \in \mathcal{X}$ lies in a class $C \subseteq \mathcal{X}$ (if $y_i = 1$) or not. A classification algorithm involves finding a discriminant function $y = \text{sign}(f(x))$ that minimizes the classification error $P(Y \neq \text{sign}(f(X)))$.

Central to a kernel-based classification algorithm (e.g., the SVM) is the notion of a kernel function $K(x, x')$ that provides a measure of similarity between two data points x and x' in \mathcal{X} . Technically, K is required to be a symmetric positive semidefinite function.² For such a function, Mercer's theorem implies that there must exist a feature space \mathcal{H} in which K acts as an inner product, i.e., $K(x, x') = \langle \Phi(x), \Phi(x') \rangle$ for some mapping $\Phi(x)$. The SVM and related kernel-based algorithms choose a linear function $f(x) = \langle w, \Phi(x) \rangle$ in this feature space. That is, they find a vector w which minimizes the loss

$$\sum_{i=1}^n \phi(y_i f(x_i))$$

subject to $\|w\| \leq B$ for some constant B . Here ϕ denotes a convex function that is an upper bound on the 0-1 loss $\mathbb{I}(y \neq \text{sign}(f(x)))$.³ In particular, the SVM algorithm is based on the hinge loss $\phi(yf(x)) = (1 - yf(x))_+$.⁴ By the Representer Theorem (cf. [Schölkopf and Smola 2002]), it turns out that the minimizing f can be expressed directly in terms of the kernel function K :

$$f(x) = \sum_{i=1}^n \alpha_i K(x_i, x) \quad (2)$$

for an optimizing choice of coefficients α_i .

There are a large number of kernel functions that satisfy the positive semidefinite property required by the SVM algorithm. Examples include the Gaussian kernel:

$$K(x, x') = \exp(-(\|x - x'\|^2/\sigma))$$

as well as the polynomial kernel:

$$K(x, x') = (\gamma + \|x - x'\|)^{-\sigma},$$

²For a translation-invariant kernel, i.e., $K(x, x') = h(x - x')$ for some function h , K is a positive semidefinite kernel if the Fourier transform of h is non-negative.

³The indicator function is defined as $\mathbb{I}(A) = 1$ if A is true, and 0 otherwise.

⁴The subscript $+$ notation means that $x_+ = \max(x, 0)$.

for parameters σ and γ . Both of these kernel functions decay with respect to the distance $\|x-x'\|$, a property that is shared by most idealized signal strength models. In particular, the radio signal model (1) has a form similar to that of a polynomial kernel. In Sheng and Hu [2003], the authors justify the use of an acoustic energy model for localization that has the form of the Gaussian kernel above. These relationships suggest a basic connection between kernel methods and sensor networks. In particular, a naive usage of kernel methods could be envisaged in which signal strength is used directly to define a kernel function. In general, however, signal strength in real sensor networks need not define a positive semidefinite function. Nonetheless, it is the premise of this paper that signal strength matrices provide a useful starting point for defining kernel-based discriminant functions. We show how to define derived kernels which are stacked on top of signal strength measurements in the following section.

Finally, it is worth noting that multi-modal signals are naturally accommodated within the kernel framework. Indeed, suppose that we have D types of sensory signals, each of which can be used to define a kernel function $K_d(x, x')$ for $d = 1, \dots, D$. Then any conic combination of K_d yields a new positive semidefinite function:

$$K(x, x') = \sum_{d=1}^D \beta_d K_d(x, x').$$

There are methods for choosing the parameters $\beta_d > 0$ based on empirical data [Lanckriet et al. 2004].

3. LOCALIZATION IN AD HOC SENSOR NETWORK

3.1 Problem statement

We assume that a large number of sensors are deployed in a geographical area. The input to our algorithm is a set of m sensors, denoted by X_1, \dots, X_m . For each i we denote by x_i the position in \mathbb{R}^2 of sensor X_i . Suppose that the first n sensor locations are known, i.e., $X_1 = x_1, \dots, X_n = x_n$, where $n \ll m$. For every pair of sensors X_i and X_j , we are given the signal $s(x_i, x_j)$ that sensor X_j receives from X_i . We want to recover the positions of X_{n+1}, \dots, X_m .

3.2 Algorithm description

We first aim to obtain a coarse location estimate for X_{n+1}, \dots, X_m . Given an arbitrarily constructed region $C \subseteq \mathbb{R}^2$, we ask whether $X_i \in C$ or not, for $i = n+1, \dots, m$. This can be readily formulated as a classification problem. Indeed, since the location of the base sensors X_1, \dots, X_n are known, we know whether or not each of these base sensors are in C . Hence we have as our training data n pairs $(x_i, y_i = \text{sign}(x_i \in C))_{i=1}^n$. For any sensor X_j , $j = n+1, \dots, m$, we can predict whether $X_j \in C$ or not based on the sign of the discriminant function $f(x_j)$:

$$f(x_j) = \sum_{i=1}^n \alpha_i K(x_i, x_j). \quad (3)$$

We emphasize that the value of $f(x_j)$ is known because the values of the kernels, $K(x_i, x_j)$, are known, despite the fact that we do not know the position x_j per se.

Next, we turn to the definition of the kernel matrix $K = (K(x_i, x_j))_{1 \leq i, j \leq m}$. In general we envision a hierarchy of kernels based on the signal matrix. An example of such a hierarchy is as follows:

- (1) We might simply define $K(x_i, x_j) = s(x_i, x_j)$. We call this naive choice a *first-tier* kernel. If the signal matrix $S = (s(x_i, x_j))_{1 \leq i, j \leq m}$ is a symmetric positive semidefinite Gram matrix then this approach is mathematically correct although it may not yield optimal performance. If S is not symmetric positive semidefinite, then a possible approximation is $(S + S^T)/2 + \delta I$. This matrix is symmetric, and is positive semidefinite for sufficiently large $\delta > 0$ (in particular, for δ larger in absolute value than the most negative eigenvalue of $(S + S^T)/2$).
- (2) Alternatively, define $K = S^T S$, to be referred to as a *second-tier* linear kernel. K is always symmetric positive semidefinite. This kernel can be interpreted as an inner product for a feature space \mathcal{H} which is spanned by vectors of the form:

$$\Phi(x) = (s(x, x_1), s(x, x_2), \dots, s(x, x_m)).$$

Specifically, we define:

$$K(x_i, x_j) = \sum_{t=1}^m s(x_i, x_t) s(x_j, x_t).$$

Intuitively, the idea is that sensors that are associated with similar vectors of sensor readings are likely to be nearby in space.

- (3) Finally, it is also possible to evaluate any kernel function (e.g., Gaussian) on the feature space \mathcal{H} induced by the second-tier kernel. This yields a symmetric positive semidefinite matrix, to be referred to as a *third-tier* kernel. Specifically, a third-tier Gaussian kernel has the following form, for a parameter σ :

$$\begin{aligned} K(x_i, x_j) &= \exp\left\{-\frac{\|\Phi(x_i) - \Phi(x_j)\|^2}{\sigma}\right\} \\ &= \exp\left\{-\frac{\sum_{t=1}^m (s(x_i, x_t) - s(x_j, x_t))^2}{\sigma}\right\}. \end{aligned}$$

Given training data $(x_i, y_i)_{i=1}^n$ and a kernel function K , we apply the SVM algorithm to learn a discriminant function $f(x)$ as in Eq. (2). The algorithmic details and computational costs are discussed in Section 4.

It is important to note that our classification formulation has several unique characteristics. The training points correspond to the base sensors, and thus may be limited in number, making the learning problem nominally a difficult one. However, because we are free to choose the target region C , the problem can in fact be made easy. This ability to design the geometry of the boundary to fit the geometry of the classifier distinguishes this problem from a traditional pattern recognition problem.

We now turn to the fine-grained estimate of sensor positions. We use the coarse-grained solution presented above as a subroutine for a localization algorithm for sensor $X_j (j = n + 1, \dots, m)$. The idea is as follows: We fix a number of overlapping regions $C_\beta (\beta = 1, \dots, U)$ in the geographical region containing the sensor network. For each β , we formulate a corresponding classification problem with respect to class C_β and predict whether or not $X_j \in C_\beta$. Hence, X_j has to be in the intersection

of regions that contain it. We might, for example, assign its location x_j to be the centroid of such an intersection. Given an appropriate choice of granularity and shapes for the regions C_β , if most of the classification labels are correct we expect to be able to obtain a good estimate of x_j .

As we have seen in our experiments on both simulated data (using a Gaussian or polynomial kernel) and real sensor data (using kernels that are constructed directly from the signal matrix), given a sufficient number of base sensors (i.e., training data points), the SVM algorithm can fit regions of arbitrary shape and size with reasonable accuracy. When the number of base sensors is limited, it is found that the SVM algorithm can still fit elliptic shapes very well. This can be turned to our advantage for fine-grained localization: By picking appropriate regions C_β such as ellipses that are easy to classify, we do not need many base sensors to achieve reasonable localization performance for the entire network. In the sequel, we will show that this intuition can be quantified to give an upper bound on the expected (fine-grained) localization error with respect to the number of base sensors.

3.3 Localization error analysis

Suppose that the sensor network of size $L \times L$ is covered uniformly by k^2 discs with radius R . Then any given point in the sensor network is covered by approximately $\pi(Rk/L)^2$ discs. Each of these discs are used to define the region for our region classification problem. To obtain a fine-grained location estimate for all remaining sensors X_j for $j = n + 1, \dots, m$, we need to solve k^2 region classification problems. Let e_β be the training error for each of these problems, for $\beta = 1, \dots, k^2$. That is,

$$e_\beta = \sum_{i=1}^n \phi(\text{sign}(x_i \in C_\beta) f(x_i)).$$

Since the size and shape of the regions are ours to decide, it is reasonable to assume that the training error for these classification problems are small. For instance, the circle/elliptic shape is particularly suited for Gaussian or polynomial kernels. Define ϵ to be the upper bound for all training errors:

$$\epsilon = \max_{1 \leq \beta \leq k^2} e_\beta.$$

From statistical learning theory [Vapnik 1998], for each $\beta = 1, \dots, k^2$, the probability of misclassification for each new sensor X_j and region C_β is $e_\beta + O(1/\sqrt{n})$, where n is the number of training points (i.e., number of base sensors). Since each location is covered by $\pi R^2 k^2 / L^2$ discs, the probability of misclassification for at least one these covering discs is, by the union bound, less than $\frac{\pi R^2 k^2}{L^2} (\epsilon + O(1/\sqrt{n}))$. If a given sensor is correctly classified with respect to all of its covering discs, then we assign the sensor's location to be the center of the intersection of all these discs, in which case the localization error is bounded by $O(L/k)$.

Hence, the expectation of the localization error is bounded by

$$O\left(\frac{L}{k}\right) + \frac{\pi R^2 k^2}{L} (\epsilon + O(1/\sqrt{n})).$$

This asymptotic bound is minimized by letting $k \propto L^{2/3} R^{-2/3} (\epsilon + O(1/\sqrt{n}))^{-1/3}$. The bound then becomes $O(L^{1/3} R^{2/3} (\epsilon + O(1/\sqrt{n}))^{1/3})$. Thus, we have proved:

PROPOSITION 3.1. *Assume that all sensor locations are independently and identically distributed according to an (unknown) distribution. For any sensor location x , let \hat{x} be the location estimate given by our algorithm, then:*

$$\mathbb{E}\|x - \hat{x}\| \leq O(L^{1/3}R^{2/3}(\epsilon + O(1/\sqrt{n}))^{1/3}).$$

If the base sensor locations are linearly separable in the feature space for all k^2 discs, i.e. $\epsilon = 0$, the following corollary is then immediate.

COROLLARY 3.2. *Assume that the base sensor locations are linearly separable in the feature space for all k^2 discs, then the localization error is $O(L^{1/3}R^{2/3}n^{-1/6})$.*

The foregoing analysis gives a guarantee that the localization performance improves as the sensor network becomes more densely distributed. In addition, the localization error increases with the size of the network, but this increase is at most linear. Our analysis also sheds light on the effects of R and k on the localization error. In particular, the localization error decreases as R decreases; on the other hand, as R decreases, the optimum value of k increases, resulting in additional computational effort, because there are k^2 discs to classify. Hence, variation in R yields a tradeoff between localization accuracy and computational complexity. This tradeoff is confirmed in our simulations.

Our analysis only considers a uniform distribution for the locations of regions C_β and circular shapes. While the analysis can be readily generalized to other specific choices, it would be of substantial interest to develop a general optimization-theoretic approach to choosing the regions.

4. ALGORITHM DETAILS AND COMPUTATIONAL COST

During the training phase associated with each coarse localization subroutine, i.e., classification with respect to a fixed region C_β , we construct the training data set based on the locations of the base sensors as described in the previous section. This is achieved by having all base stations send the signal matrix entries $s(x_i, x_j)$ and their known locations to a central station, a procedure which involves receiving and storing $n^2 + n$ numbers at the central station. The central station then solves the following optimization problem:

$$\min_{\mathbf{w}} \|\mathbf{w}\|^2 + \frac{c}{n} \sum_{i=1}^n \phi(y_i f(x_i)),$$

where $f(x) = \langle \mathbf{w}, \Phi(x) \rangle$, $\phi(yf(x)) = (1 - yf(x))_+$, and c is a fixed parameter.⁵ This is a convex optimization problem, which has the following dual form [Schölkopf and Smola 2002]:

$$\max_{0 \leq \alpha_i \leq c} 2 \sum_{i=1}^n \alpha_i - \sum_{1 \leq i, j \leq n} \alpha_i \alpha_j y_i y_j K(x_i, x_j). \quad (4)$$

The algorithm finds optimizing values of $\{\alpha_i\}$, which are then used to form the discriminant function in Eq. (2).

⁵The parameter c is a regularization parameter associated with the SVM algorithm. In all our experiments, we fix $c = 10$.




Fig. 1. Summary of the coarse localization algorithm.

It is known that the solution to this optimization problem can be found in the worst case in $O(n^3)$ computational time. Thus if there are k^2 regions to classify, this result suggests a total training time of $O(n^3 k^2)$. However, this worst-case estimate is overly conservative in our setting. Indeed, an estimate based on the number of support vectors n_s returned by each classification algorithm (those X_i such that $\alpha_i \neq 0$) reveals that the computational complexity is $O(n_s^3 + n_s^2 n)$ instead of $O(n^3)$. Usually $n_s \ll n$. Our simulation experience (to be presented in the next section) shows that when discs with radius R are used, the support vectors reside mostly along the boundaries of the discs, hence $n_s \approx O(\min(n\pi R^2/L^2, 2\pi R))$, in which case the overall training phase takes only $O(R^2 n k^2)$ time. Note also that this training phase is the most expensive part of our algorithm and is performed at a central station.

Once the training phase is complete, each base sensor is required to store the n parameters $(\alpha_1, \dots, \alpha_n)$ for the purpose of classification of the remaining (location-unknown) sensors. If the first-tier kernel is used, a new sensor X_j for $j = n + 1, \dots, m$ records the signal $s(x_i, x_j)$ from the n_s base sensors $i \in \{1, \dots, n\}$, and combines these with the non-zero values α_i , resulting in a cost of $O(n_s)$ in time and storage. If a second-tier linear kernel or a third-tier Gaussian kernel is used, a new sensor X_j records n -element signal vectors $(s(x_i, x_1), \dots, s(x_i, x_n))$ from the n_s base stations, resulting in a $O(n_s n)$ cost in time and storage. The kernel values $K(x_i, x_j)$ are then readily computable from the received signals $s(x_i, x_j)$ in $O(1), O(n), O(n^2)$ time for the first-tier, second-tier and third-tier kernel, respectively. Then a simple computation (Equation 3) determines for sensor X_j whether it resides in the region C or not. The attractive feature of the localizing step is that it is done locally (in a distributed fashion), taking only linear (for the first-tier and second-tier kernels) or quadratic (for the third-tier Gaussian kernel) time and storage space (in terms of n). Since the localization is done on an as-needed basis, its time and storage cost do not depend on the total number of sensors m in the network. A summary of our algorithm is provided in Figure 1.

Now we turn to fine-grained localization. At both algorithmic and system levels, this involves invoking the coarse localization subroutine k^2 times with respect to regions C_1, \dots, C_{k^2} . Therefore, for each region $\beta = 1, \dots, k^2$ we have a set of parameters $(\alpha_i)_{i=1}^n$. Each sensor X_j can then determine its location by setting x_j to be the centroid of the intersection of all regions C_β that it finds itself residing

in. In the case in which C_β are discs with centers c_β , this yields:

$$x_j := \frac{\sum_{\beta=1}^{k^2} c_\beta \mathbb{I}(X_j \in C_\beta)}{\sum_{\beta=1}^{k^2} \mathbb{I}(X_j \in C_\beta)}.$$

Other methods for determining the shapes of C_β and for estimating x_j are certainly worthy of investigation. Clearly, the computational cost of a fine-grained localization algorithm is k^2 times as much as the computational cost of each coarse localization step. In summary, our fine-grained localization algorithm is shown in Figure 2.




Fig. 2. Summary of the fine-grained localization algorithm.

5. EXPERIMENTAL RESULTS

We evaluate our algorithm on simulated sensor networks in the first two subsections, and then on a real network using Berkeley sensor motes.

5.1 Coarse localization

Simulation set-up: We consider a network of size 10×10 square units. The base sensors are distributed uniformly in a grid-like structure. There are a total of n^2 such sensors.⁶ We are concerned with recognizing whether a sensor position x , characterized by the signal reading $s(x_i, x)$ for $i = 1, \dots, n$, lies in a region C or not.

We first define a signal model: Each sensor location x is assumed to receive from a sensor located at x' a signal value following a fading channel model: $s(x, x') = \exp - \frac{\|x-x'\|^2 + N(0, \epsilon)}{\sigma}$, where $N(0, \epsilon)$ denotes an independently generated normal random variable with standard deviation ϵ . This signal model is a randomized version of a Gaussian kernel. We have also experimented with a signal strength model that is a randomized version of the polynomial kernel: $s(x, x') = (\|x - x'\| + N(0, \epsilon))^{-\sigma}$.

⁶In this section, we use n to denote the number of base sensors along each dimension, while in the previous sections, n is used to denote the total number of base sensors.




Fig. 3. Simulation results with (randomized) Gaussian models. The x -axis shows the number of sensors employed along each dimension of the network. The y -axis shows the ratio between the number of misrecognized points and the number of points inside the area to be recognized.




Fig. 4. (a) Illustration of a simulated sensor network with 15×15 base sensors, and the recognized boundary in black (with $R = 2$) using a Gaussian kernel with $\sigma = 1$. The black squares are the support vector base sensors. The test error in this figure is 0.27. (b) Plots show the effect of the sensor fading signal parameter σ and signal noise parameter ϵ on coarse localization performance.

The results for the polynomial kernels are similar to the Gaussian kernels, and are not presented here. It is emphasized that although the use of these models have been motivated elsewhere [Seidel and Rappaport 1992; Sheng and Hu 2003], they

are used here merely to generate the signal matrix S , which is used as the input to our algorithm. We do *not* assume any knowledge of the signal model that generates S .

Next, we define a region C to be recognized. In particular, C consists of all locations x that satisfy the following equations: $(x - v)^T H_1 (x - v) \leq R$ and $(x - v)^T H_2 (x - v) \leq R$, where $v = [5 \ 5]^T$, $H_1 = [2 \ -1; -1 \ 1]$ and $H_2 = [2 \ 1; 1 \ 1]$. The radius R is used to describe the size of C . For each simulation set-up (n, R, σ, ϵ) , we learn a discriminant function f for the region C using the training data given by the base sensor positions. Once f is learned, we test the classification at 100×100 sensor locations distributed uniformly in the region containing the network.

Figure 4(a) illustrates C as a green shaded region, for $R = 2$, while the black boundary represents the region learned by our localization algorithm. Qualitatively, the algorithm has captured the shape of the target region C . We now we present a quantitative analysis on the effects of n, R, σ , and ϵ on the localization (i.e., classification) performance:

Effects of n : The plots in Figure 3 show the localization (test) error with respect to the number of base sensors deployed in the network. The test error is defined to be the ratio between the number of misclassified points and the number of points located within the area $C(R)$ (out of 100×100 locations distributed uniformly in the grid). In this set of simulations, we fix the noise parameter $\epsilon = 0$, and let $\sigma = 1$ and $\sigma = 7$, while varying n . The plots confirm that the localization error decreases as the sensor network become more densely distributed. Note that if we need to recognize a particular area, we only need to plant base sensors in the area around the boundary, because these are the likely locations of support vectors. Of course, in our context coarse-grained localization is only a subroutine for fine-grained localization, and it is our interest to have base sensors spread throughout the whole geographical area.

Effects of σ , and signal noise ϵ : Parameter σ is used to describe the sensitivity of the signal strength with respect to the sensor distance. For instance, for Gaussian models, a small value of σ implies that the signal strength fades very quickly for distant sensors. Parameter ϵ describes the noise of the signal strength apart from a deterministic functional relationship based on distance. The plots in Figure 4(b) show the effects of both ϵ and σ on the localization performance. In this set of simulations, we fix the number of base sensors along each dimension to be 10, the radius of C to be $R = 2$, while varying σ and ϵ . It is seen that the localization performance degrades as we increase the noise parameter ϵ , and the degradation is more severe for the more sensitive signal, i.e., when σ is small. In addition, when fixing ϵ , the localization performance improves as σ gets smaller.

5.2 Fine-grained localization

Simulation set-up: The network set-up is the same as the previous section, except that the n^2 base sensors are now distributed approximately uniformly at random in the whole area. By this we mean that each base sensor is initially planted at a grid point in the $L \times L$ square, where $L = 10$, and then perturbed by Gaussian noise $N(0, L/(2n))$. There are 400 other sensors whose locations are to be determined using our algorithm. These 400 sensors are deployed uniformly in the grid. Again, we assume the signal strength follows a Gaussian signal model, with noise parameter




Fig. 5. Left figure shows the effect of the number of base sensors n^2 on fine-grained localization error mean and standard deviation (for all nodes). Right figure shows the effects of the size of discs (by radius R) and the number of discs (k^2) distributed uniformly on the field. The means and variances are collected after performing the simulation on 20 randomly generated sensor networks.




Fig. 6. Localization results for a simulated sensor network of size 10×10 square units with 25 base sensors (left figure) and 64 base sensors (right figure). The base sensors are the black squares. Each blue line connects a true sensor position (in circle) and its estimate. The signal model is Gaussian. The mean localization error is 0.4672 in the left figure and 0.3877 in the right figure.

Method	Mean	Median	Std
Two-step ranging-based	6.99	5.28	5.79
First-tier signal kernel	6.67	4.60	7.38
Second-tier linear kernel	3.65	2.51	4.29
Third-tier Gaussian kernel	3.53	2.63	3.50

Table I. Comparison between a two-step ranging-based algorithm and our kernel-based localization algorithm in a sensor network with 25 base sensors covering a 40×40 square-inch area. The localization error mean, median and standard deviation are taken over all position estimates, and measured in inches.

$\epsilon = 0.2$.

We applied the algorithm described in Section 3 for fine-grained localization. The algorithm involves repeated application for coarse localization with respect to a number of regions that cover the whole network area. We choose these regions to be discs of radius R , and distributed uniformly over the whole network. Let k be the number of discs along each dimension, hence there are a total k^2 discs to recognize. In this simulation we study the effects of R , k and the number of base sensors n on the localization performance. Specifically, we examine the tradeoff between computational cost and localization accuracy of our algorithm by varying these parameters, as anticipated in the theoretical analysis in Section 3.3.

Effects of n : Figure 5(a) shows that the localization error mean (averaged over all sensor networks and over all sensors) decreases monotonically as more base sensors added to the network. This agrees with the theoretical analysis presented in Section 3.3. Figure 6 illustrates the localization results for each node in the networks with 25 and 64 base sensors. The mean localization error (averaging over all sensors) for these two networks are 0.46 and 0.38, respectively.

Effects of R and k : Figure 5(b) shows the effects of R and k on the localization performance. In this set of simulations, we fix $\epsilon = 0.2$, $\sigma = 2$, and $n = 10$, while varying R and k . The analysis in Section 3.3 suggests that for each value of R , there exists an optimum value for k that increases as R decreases. Since there are k^2 classification problems to solve, the computational cost generally increases as R decreases. On the other hand, the optimum mean localization error decreases as R decreases. This is proven in Corollary 3.2, and confirmed in the plots of Figure 5(b). We clearly observe a tradeoff between computational cost and localization accuracy as manifested by the behavior of R and k .

5.3 Localization with Berkeley sensor motes

Experiment set-up: We evaluated our algorithm on a real sensor network using Berkeley tiny sensor motes (Mica motes) as the base stations. The goal of the experiment is to estimate the positions of light sources given the light signal strength received by a number of base sensors deployed in the network. Our hardware platform consists of 25 base sensors placed 10 inches apart on a 5×5 grid in an indoor environment. Each sensor mote is composed of one Atmel ATmega 103 8-bit processor running at 4MHz, with 128Kb of flash and 4Kb of RAM, RFM TR1000 radio, EEprom and a sensor board which includes light, temperature, microphone sensors and a sounder. Our experiment makes use of light sensor data received by the motes’s light receiver antennas (using random orientation settings). The




Fig. 7. Panel (a) shows the noisy relationship between signal strength received by sensors and the distances. Few sensors exhibit a clear signal strength-distance functional pattern as in panel (a), while most are like those in panels (c) and (d). Note that only data points marked with x in red are available for regression training.

measured signals are a scalar field produced by a light source shining over the sensor network. Only the position of light sources placed at the base sensors are given as training data. To be estimated are 81 light source positions distributed uniformly in a 9×9 grid spread over the whole network.

A range-based algorithm: We compared our algorithm to a state-of-the-art algorithm that epitomizes a majority of localization algorithms in the literature. This algorithm was described in Whitehouse [2002], and consists of two main steps: (1) a ranging procedure aimed at establishing a mapping between the signal strength received by a base sensor and the distance to the light source, and (2) a localization procedure giving the distance estimates using least-squares methods.

Figure 7 illustrates the difficulty of the ranging problem—the functional relationship between distances and signal strengths is very noisy. Much of this noise is device-specific; as shown in Figure 7, a few sensors exhibit a clear distance-to-signal-strength pattern, while most others exhibit a very noisy pattern. As shown in Whitehouse [2002], improvement in the ranging step can be achieved by ac-




Fig. 8. Localization result for a real sensor networks covering a 40×40 square-inch area. There are 25 base sensors (Berkeley notes) spaced in a 5×5 grid. Each line connects a true position (in circle) and its estimate. Panel (a) shows the results given by a traditional 2-step localization algorithm, while panels (b,c,d) show the localization results obtained by our algorithm using three different kernels (the first-tier, second-tier and third-tier Gaussian kernel, respectively).

counting for properties specific base sensors. This is done by introducing regression coefficients for each of these base sensors. Once the ranging step is completed, we have estimates of the distance between the base sensors and the positions of the light source. The initial position estimates are obtained using the Bounding-Box algorithm, and are then iteratively updated using a least-squares method (see [Whitehouse 2002; Savvides et al. 2001]). Figure 8(a) shows the localization results for this algorithm.

Results for the kernel-based algorithm: Three different kernels are used in our algorithm. The first is a first-tier symmetric positive semidefinite approximation
ACM Journal Name, Vol. V, No. N, Month 20YY.

of the signal matrix. In particular, as discussed in Section 3, given a signal matrix S , we define $S' := (S + S^T)/2 + \delta I$. The remaining kernels are a second-tier linear and third-tier Gaussian kernel, with the parameter σ fixed to 0.5 in the latter case. For fine-grained localization, coarse localization is repeatedly applied for discs of radius $R = L/2 = 20$ inches that cover part of the network area. The centers of these discs are five inches apart in both dimensions, and there are 10 discs along each dimension (i.e., $k = 10$).

Table I shows that the localization error achieved by the kernel-based approach is smaller than that of the two-step algorithm. Among the three choices of signal kernels, the second-tier kernels are much better than the simple first-tier kernel. The localization results are depicted spatially in Figure 8. Note that the minimum distance between two neighboring base sensors is about 10 inches, and the localization error of our algorithm (using second-tier kernels) is slightly over one third of that distance.

6. DISCUSSION

We have presented an algorithm for coarse-grained and fine-grained localization for ad hoc wireless sensor networks. Our approach treats the signal strength as measured by sensor nodes as a natural coordinate system in which to deploy statistical classification and regression methods. For the localization problem, this approach avoids the ranging computation, a computation which requires accurate signal models that are difficult to calibrate. Instead, we use signal strength either directly to define basis functions for kernel-based classification algorithms, or indirectly via derived kernels that operate on top of the signal strength measurements. We show how a kernel-based classification algorithm can be invoked multiple times to achieve accurate localization results, and we present an error analysis for the accuracy that can be achieved as a function of base sensor density. Our algorithm is particularly suitable for densely distributed sensor networks, and is appealing for its computational scaling in such networks: The preprocessing computations are performed at the base sensors, which are assumed to have sufficient processing and power capability, while the localizing step at location-unknown sensors can be achieved in linear time.

We have argued for a simple approach to localization that dispenses with ranging computations and sensor modeling. We do not necessarily believe, however, that our statistical approach is always to be preferred. In particular, the level of accuracy that we appear to be able to obtain with our approach is on the order of one third the distance between the nodes. While this accuracy is sufficient for many potential applications of sensor networks, in some applications higher accuracy may be required. In this case, ranging-based approaches offer an alternative, but only in the setting in which highly accurate models of the relationship between sensor signals and distances are available.

REFERENCES

- BAHL, P. AND PADMANABHAN, V. N. 2000. RADAR: An in-building RF-based user location and tracking system. In *INFOCOM (2)*. 775–784.
- BULUSU, N., HEIDEMANN, J., AND ESTRIN, D. 2000. GPS-less low cost outdoor localization for
- ACM Journal Name, Vol. V, No. N, Month 20YY.

- very small devices. Tech. Rep. 00-729, Computer Science Department, University of Southern California.
- D’COSTA, A. AND SAYEED, A. 2003. Collaborative signal processing for distributed classification in sensor networks. In *2nd International Workshop on Information Processing in Sensor Networks (IPSN)*. 193–208.
- GIROD, L. AND ESTRIN, D. 2001. Robust range estimation using acoustic and multimodal sensing. In *IEEE/RSI International Conference on Intelligent Robots and Systems (IROS)*.
- HIGHTOWER, J. AND BORRIELLO, G. 2000. Real-time error in location modeling for ubiquitous computing. In *Location Modeling for Ubiquitous Computing—UbiComp 2001 Workshop Proceedings*. 21–27.
- LANCKRIET, G., CRISTIANINI, N., GHAOUI, L. E., BARTLETT, P., AND JORDAN, M. 2004. Learning the kernel matrix with semi-definite programming. *Journal of Machine Learning Research* 5, 27–72.
- LI, D., WONG, K., HU, Y., AND SAYEED, A. 2002. Detection, classification, and tracking of targets. *IEEE Signal Processing Magazine* 19, 17–29.
- PRIYANTHA, N., CHAKRABORTY, A., AND BALAKRISHNAN, H. 2000. The Cricket location-support system. In *ACM International Conference on Mobile Computing and Networking*. ACM Press, New York.
- SAVARESE, C., RABAAY, J., AND LANGENDOEN, K. 2002. Robust positioning algorithms for distributed ad-hoc wireless sensor networks. In *USENIX Annual Technical Conference, Monterey, CA*. 317–327.
- SAVVIDES, A., HAN, C., AND SRIVASTAVA, M. 2001. Dynamic fine grained localization in ad-hoc sensor networks. In *Proceedings of the Fifth International Conference on Mobile Computing and Networking*. 166–179.
- SCHÖLKOPF, B. AND SMOLA, A. 2002. *Learning with Kernels*. MIT Press, Cambridge, MA.
- SEIDEL, A. AND RAPPAPORT, T. 1992. 914MHz path loss prediction models for indoor wireless communications in multi-floored buildings. *IEEE Transactions on Antennas and Propagation* 40, 207–217.
- SHENG, X. AND HU, Y. 2003. Energy based acoustic source localization. In *2nd International Workshop on Information Processing in Sensor Networks (IPSN)*. 285–300.
- VAPNIK, V. 1998. *Statistical Learning Theory*. John Wiley, New York.
- WANT, R., HOPPER, A., FALCAO, V., AND GIBBONS, J. 1992. The active badge location system. *ACM Transactions on Information Systems* 10, 91–102.
- WARD, A., JONES, A., AND HOPPER, A. 1997. A new location technique for the active office. *IEEE Personnel Communications* 4, 42–47.
- WHITEHOUSE, C. 2002. The design of Calamari: An ad-hoc localization system for sensor networks. M.S. thesis, Department of Electrical Engineering and Computer Science, University of California, Berkeley.