

Towards Optimal Beacon Placement for Range-Aided Localization

Ethan Sequeira, Hussein Saad, Stephen Kelly, and Matthew Giamou

Department of Computing and Software

McMaster University

Hamilton, Canada

{sequeie, saadh, spkelly, giamoum}@mcmaster.ca

Abstract—Range-based localization is ubiquitous: global navigation satellite systems (GNSS) power mobile phone-based navigation, and autonomous mobile robots can use range measurements from a variety of modalities including sonar, radar, and even WiFi signals. Many of these localization systems rely on fixed anchors or beacons with known positions acting as transmitters or receivers. In this work, we answer a fundamental question: given a set of positions we would like to localize, how should beacons be placed so as to minimize localization error? Specifically, we present an information-theoretic method for optimally selecting an arrangement consisting of a few beacons from a large set of candidate positions. By formulating localization as maximum a posteriori (MAP) estimation, we can cast beacon arrangement as a submodular set function maximization problem. This approach is probabilistically rigorous, simple to implement, and extremely flexible. Furthermore, we prove that the submodular structure of our problem formulation ensures that a greedy algorithm for beacon arrangement has suboptimality guarantees. We compare our method with a number of benchmarks on simulated data and release an open source Python implementation of our algorithm and experiments.

Keywords—range-based localization; nonlinear optimization

I. INTRODUCTION

Range-based localization has been successfully deployed at many scales and with a variety of different sensing modalities. Prominent successes include the Global Positioning System (GPS) [1], and indoor positioning systems based on WiFi signal strength [2]. Accurate localization is essential for safe operation of autonomous mobile robots, and recent work has explored the development of robust and accurate systems that rely on range sensors [3], [4], [5]. In many industrial and commercial settings, engineers have the opportunity to design an environment which aids state estimation algorithms with fiducial markers or static sensors. In the case of range-aided localization, budgetary and infrastructural constraints may limit the number or placement of sensing beacons installed in a robot’s operational environment.

Inspired by the work of Kayhani et al. [6] on how to optimally place fiducial markers for visual localization in construction environments, we introduce a solution to the *optimal beacon placement* (OBP) problem depicted in Figure 1. Our problem formulation assumes that the designers of the environment are provided with or able to produce




Figure 1: We present an information-theoretic approach to arranging the position \mathbf{a}_j of K sensing beacons in an environment in which an agent relies on range measurements for localization at points \mathbf{x}_i for $i \in [M]$. Using a finite set of candidate beacon positions \mathcal{S} and an objective function derived with D-optimal Bayesian experiment design, we cast the optimal beacon placement (OBP) problem as submodular set function maximization with suboptimality guarantees.

a set of positions at which localization is important. These may be samples from a trajectory a service robot has to follow, or positions where localization from other sensing modalities is unavailable (e.g., an indoor area without GPS reception). Furthermore, our formulation incorporates a prior distribution for each position estimate, allowing designers to specify the fact that some positions will have inaccurate onboard estimation (e.g., corridors where visual odometry is highly uncertain in one spatial dimension, or low-friction surfaces that degrade wheel odometry). Finally, we assume that beacons may only be placed in one of finitely many candidate positions specified by the designer. In some scenarios, this limitation reflects infrastructural realities, such as the presence or absence of discrete architectural features onto which beacons may be mounted. In cases without these constraints, our method avoids issues which may arise in a continuous parameterization of OBP’s feasible set, while still supporting a fine discretization involving hundreds or thousands of candidate beacon positions.

When combined with a Gaussian model for range mea-

surements, our novel problem formulation suggests a simple greedy algorithm with a provable suboptimality guarantee. Our novel contributions are as follows:

- 1) an information-theoretic formulation of OBP for maximum a posteriori (MAP) localization;
- 2) a fast greedy solution with a suboptimality guarantee derived from the theory of submodular set function maximization;
- 3) simulated experiments comparing the effectiveness of our approach with benchmark algorithms; and
- 4) a publicly available Python repository containing an implementation of our method and experimental results.¹

The remainder of this paper is organized as follows: Section II summarizes related work. In Section III, we describe our MAP formulation of range-aided localization. Section IV presents OBP, our set function maximization approach to placing a limited number of beacons in an environment. In Section V, we demonstrate that a greedy algorithm can find a solution to OBP with a solution within a constant factor of the global optimum, and we test this method against benchmarks in the simulated experiments of Section VI. Section VII summarizes our findings and identifies promising directions for future work.

II. RELATED WORK

State estimation for mobile robots is a rich topic with a long history [7]. In this section, we briefly review recent results in range-based localization and survey prior work related to optimal sensing in robotics and vision.

A. Range-based Localization

Throughout this paper, we make a distinction between *range-only* localization methods [8], which don't include information from other sources, and *range-aided* localization, which incorporates other sensing models. Range-only localization plays an important role in simultaneous localization and mapping (SLAM) for autonomous underwater vehicles [9], which differ from aerial and terrestrial vehicles in that they cannot rely on optical sensing and are therefore often equipped with acoustic range sensors. Recently, convex relaxations [10] have been applied to the range-based localization problems. In [11] and [4], convex initialization strategies are applied to range-only localization and range-aided SLAM, respectively. Similarly, convex relaxations methods are used in [5] and [12] to directly solve these problems and provide certificates of global optimality. In order to define the optimal beacon placement problem, the present paper uses a formulation of range-aided localization which is compatible with many of the aforementioned solution methods. However, the benefits of our approach are agnostic to the actual method used to perform localization.

Range-only localization has a great deal of overlap with the related problem of sensor network localization (SNL) [13]. In this work, we use a probabilistic and graph-theoretic problem formulation very similar to the one in [14], which proposes a distributed maximum likelihood solution to SNL. Additionally, like the SNL literature, we restrict our attention to position estimation, and note that placing multiple range sensors on an agent enables full six degree of freedom pose estimation through range sensing alone [3].

B. Optimal Sensor Placement

The algorithm developed in this work relies on the machinery of submodular set function maximization described in [15]. This machinery has been successfully applied to a variety of state estimation problems related to robotics including visual feature selection [16], Kalman filtering [17], camera placement [18], and measurement selection for SLAM [19], [20]. Surprisingly, there is very little work on the optimal arrangement of sensing infrastructure for localization or SLAM accuracy. Kayhani et al. [6] uses a genetic algorithm to optimally place fiducial tags for visual sensing, and Kaveti et al. [21] consider the inverse problem of optimally arranging cameras on a robot moving through an environment with known static visual features. The greedy algorithm in [22] additionally considers the visual information already present in a scene to determine the best locations for fiducial markers. While our approach uses an information-theoretic approach and discrete design space similar to these works, we focus exclusively on range-aided localization, and make use of submodularity inherent in our MAP problem formulation to enable a greedy algorithm with suboptimality guarantees.

Most prior work on optimal placement of beacons for range sensing focuses on coverage of the environment. For two-dimensional localization, Andrews et al. [23] propose a regular lattice which ensures unique localization coverage with a minimal number of beacons. Rajagopal et al. [24] minimize a metric combining localization coverage with geometric dilution of precision (GDOP) [25], a quantity proportional to the Cramér-Rao lower bound used in this work. However, their method does not use an efficient greedy algorithm with suboptimality guarantees. In [26], Wang et al. also consider a discretized version of the beacon placement problem that is similar to ours. They also apply greedy methods, but use a sophisticated submodular set function related to the number of points uniquely localized by a set of beacons. In contrast, we focus on a simple information-theoretic objective function. Additionally, none of the work discussed in this section considers the specific range-aided localization problem we present in Section III-C.

Finally, in the realm of SNL, Shames and Summers [27] study a problem closely related to ours involving the selection of a subset of vertices in a sensor network to specify as anchors with known position. By using concepts from

¹<https://github.com/ARCO-Lab/localization-simulator>

rigidity theory, they identify a set function maximization formulation of their problem that is *modular*, meaning a greedy algorithm produces a globally optimal solution. Our problem differs in that we are selecting which anchors to use with a fixed set of vertices that require localization, and we use a purely information-theoretic objective function. However, analyzing the influence of rigidity on our problem formulation is an intriguing direction for future work.

III. BACKGROUND: RANGE-AIDED LOCALIZATION

In this section we begin by describing our probabilistic formulation of range-only localization in Section III-A. In Section III-B, we make mild assumptions about range measurement models to describe a maximum likelihood approach to range-only localization. In Section III-C, we incorporate prior estimates of positions to extend the MLE in Section III-B to a range-aided localization problem. By restricting error models to additive Gaussian noise, we reduce range-aided localization to a simple nonlinear least squares problem.

A. Range-Only Localization

Consider the problem of estimating an unknown position $\mathbf{x} \in \mathbb{R}^d$ using noisy range measurements with N anchors or *beacons* with known positions $\mathbf{a}_j \in \mathbb{R}^d \forall j \in [N]$, where $[n] \triangleq \{1, 2, \dots, n\}$ for any integer $n > 0$. We restrict our attention to noisy range measurement models of the form

$$g_j(\mathbf{x}) = \|\mathbf{x} - \mathbf{a}_j\|_2 + \eta_j, \quad (1)$$

where $\eta_j \sim \mathcal{N}(0, \sigma_j^2)$ is additive zero-mean Gaussian noise with variance σ_j^2 that may be a function of the particular beacon j and its position \mathbf{a}_j with respect to \mathbf{x} .² Additionally, we may have a prior estimate encoded as a multivariate Gaussian distribution with mean $\check{\mathbf{x}} \in \mathbb{R}^d$ and covariance $\check{\Sigma} \in \mathbb{S}_{++}^d$,³ where \mathbb{S}_{++}^n (\mathbb{S}_+^n) is the space of $n \times n$ symmetric positive (semi)definite matrices.

B. Maximum Likelihood Estimation

In this section, we use the probabilistic measurement model in Equation (1) to develop a maximum likelihood estimate (MLE) $\hat{\mathbf{x}}$ of \mathbf{x} [7]. We begin by noting that the conditional probability of observing a particular measurement is

$$\tilde{d}_j \sim p_j(\cdot|\mathbf{x}) = \mathcal{N}(\|\mathbf{x} - \mathbf{a}_j\|_2, \sigma_j^2). \quad (2)$$

Since the measurements are independent, their joint conditional distribution is the product

$$p(\tilde{D}|\mathbf{x}) \triangleq \prod_{j=1}^N p_j(\tilde{d}_j|\mathbf{x}), \quad (3)$$

²As an example, the noise may depend on the range $\|\mathbf{x} - \mathbf{a}_j\|_2$.

³We denote prior estimates with $\check{\cdot}$ and posterior estimates gleaned from measurements with $\hat{\cdot}$.




Figure 2: Range-aided localization of a single unknown position \mathbf{x}_i with prior mean $\check{\mathbf{x}}_i$ and covariance $\check{\Sigma}_i$. For each beacon \mathbf{a}_j in range, a noisy measurement \tilde{d}_{ij} is received.

where $\tilde{D} \triangleq \{\tilde{d}_j\}$ is the set of all measurements. The MLE estimator of \mathbf{x} is therefore

$$\hat{\mathbf{x}} \triangleq \underset{\mathbf{x} \in \mathbb{R}^d}{\operatorname{argmax}} p(\tilde{D}|\mathbf{x}). \quad (4)$$

Taking the negative log of the likelihood, we can obtain $\hat{\mathbf{x}}$ by solving the following equivalent nonlinear least squares minimization problem:

Problem 1 (MLE for Range-Only Localization). *For noisy range measurements $\tilde{d}_j \in \tilde{D}$ with associated variance σ_j^2 , find $\hat{\mathbf{x}}$ that solves*

$$\min_{\mathbf{x} \in \mathbb{R}^d} \sum_{j=1}^N \frac{1}{\sigma_j^2} (\|\mathbf{x} - \mathbf{a}_j\|_2 - \tilde{d}_j)^2, \quad (5)$$

where each \mathbf{a}_j is a sensing beacon with a known fixed position.

The Cramér-Rao lower bound (CRLB) provides a lower bound on the covariance of the MLE estimator $\hat{\mathbf{x}}$ [7]:

$$\operatorname{Cov}(\hat{\mathbf{x}}) \succeq \mathcal{I}(\mathbf{x})^{-1}, \quad (6)$$

where we have used the *Fisher information matrix* (FIM):

$$\mathcal{I}(\mathbf{x}) \triangleq \mathbb{E}_{\tilde{D}} \left[\nabla_{\mathbf{x}} \log p(\tilde{D}|\mathbf{x}) (\nabla_{\mathbf{x}} \log p(\tilde{D}|\mathbf{x}))^\top \right]. \quad (7)$$

In Section IV-B, we use the FIM to define an information-theoretic objective function of beacon positions \mathbf{a}_j .

C. Maximum a Posteriori Estimation

Suppose that we also have access to a prior distribution $p(\mathbf{x}) \sim \mathcal{N}(\check{\mathbf{x}}, \check{\Sigma})$ for our unknown position \mathbf{x} . This may come from some other independent noisy localization system (e.g., wheel odometry), or can function as a way of specifying that an agent is expected to be in a certain area at a certain time as part of its operation. Using Bayes' rule, we obtain the following MAP estimation problem [7]:

Problem 2 (Range-Aided MAP Localization). *Range-aided*

MAP localization with a Gaussian prior takes the form

$$\min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{x} - \tilde{\mathbf{x}}\|_{\Sigma}^2 + \sum_{j=1}^N \frac{1}{\sigma_j^2} (\|\mathbf{x} - \mathbf{a}_j\|_2 - \tilde{d}_j)^2, \quad (8)$$

where $\|\cdot\|_{\mathbf{A}}$ denotes the Mahalanobis distance with respect to $\mathbf{A} \in \mathbb{S}_+^n$ [7].

Problem 2 is essentially Problem 1 with a regularizing term defined by the prior. This regularizing term is important to the method developed in Section IV. Both the MLE and MAP problems are smooth, unconstrained, and nonconvex optimization problems which can be efficiently and reliably searched for local minima by Newton's method and its many variants.

IV. OPTIMAL BEACON PLACEMENT

In this section, we use the probabilistic range-aided localization problem of Section III to introduce the optimal beacon placement problem. Section IV-A describes two ways of parameterizing this problem's design space, and in Section IV-B we define an information-theoretic notion of optimality. Finally, Section IV-C introduces OBP as a set function maximization problem which we will solve in Section V.

A. Feasible Beacon Positions

Suppose that we are designing an environment for an agent equipped with a sensor capable of measuring its range to fixed beacons with known position. Furthermore, we are constrained to a budget of $K \geq 1$ beacons. Physical constraints may limit the placement of beacons to some subset $\Omega \subset \mathbb{R}^d$. We can therefore completely specify all feasible arrangements of beacons via $\mathbf{a}_j \in \Omega \forall j \in [K]$.

Alternatively, we may be constrained to mount beacons on some finite $\mathcal{S} \subset \Omega$, where $|\mathcal{S}| = N \gg K$. In this *discrete* formulation, there are $\binom{N}{K} \triangleq \frac{N!}{K!(N-K)!}$ candidate beacon arrangements. This scenario may reflect real infrastructural limitations on the placement of beacons, but we will see in Section V that this combinatoric view of the problem also serves as a practical algorithmic contrivance.

In our beacon arrangement problem, we aim to estimate M unknown positions $\mathbf{x}_i \forall i \in [M]$, each with a prior Gaussian distribution $p_i(\mathbf{x}_i) \sim \mathcal{N}(\tilde{\mathbf{x}}_i, \tilde{\Sigma}_i)$. Due to range limits or occlusions, only a subset of the beacons \mathbf{a}_j can be used at each position i (see Figure 1). We encode this information in a bipartite graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} \triangleq [M] \cup [N]$ and each edge $i, j \in \mathcal{E}$ is identified with a measurement d_{ij} . The neighbourhood function which maps position i to the beacons it can measure is denoted

$$\nu : [M] \rightarrow 2^{[N]}. \quad (9)$$

Just as in Equation (1), each position i has an associated set of measurement models

$$g_{ij}(\mathbf{x}_i) = \|\mathbf{x}_i - \mathbf{a}_j\|_2 + \eta_{ij} \forall j \in \nu(i), \quad (10)$$

and we denote the distribution of measurement \tilde{d}_{ij} with $p_{ij}(\cdot|\mathbf{x}_i)$. The estimation of each unknown position \mathbf{x}_i is independent of the other positions, but the placement of each beacon j affects the localization quality for multiple positions.

For each $S \subseteq \mathcal{S}$, consider the binary vector $s \in \{0, 1\}^N$ indicating membership in S :

$$s_j = \begin{cases} 1, & \mathbf{a}_j \in S, \\ 0, & \text{otherwise.} \end{cases} \quad (11)$$

This vector allows us to conveniently write our MAP likelihood as a function of included beacons:

$$p(\tilde{D}|X; s) \triangleq p(\mathbf{x}) \prod_{i=1}^M \prod_{j \in \nu(i)} p_{ij}(\tilde{d}_{ij}|\mathbf{x}_i)^{s_j}, \quad (12)$$

where we abuse our notation and use $\mathbf{x} \in \mathbb{R}^{dM}$ to refer to the concatenation of all position \mathbf{x}_i in Equation (12) onwards. We will also make the simplifying assumption that the prior belief $p(\mathbf{x}_i)$ for position \mathbf{x}_i is independent of the others, i.e.,

$$p(\mathbf{x}) = \prod_{i=1}^M p(\mathbf{x}_i). \quad (13)$$

This assumption may not be true for prior estimates originating from visual SLAM or a similar method, but can still serve as a useful approximation. Furthermore, this assumption of independence for each position's prior enables a very efficient solution method to the optimal beacon placement problem in Section V.

B. An Information-Theoretic Objective

In order to define an *optimal* placement of beacons over Ω or \mathcal{S} , we need to introduce a sensible objective function. Recall the FIM of Equation (7), whose inverse is a lower bound to the MLE $\hat{\mathbf{x}}$. This motivates us to choose an objective function that *maximizes* the eigenvalues of $\mathcal{I}(\mathbf{x})$. Since $\mathcal{I}(\mathbf{x}) \in \mathbb{S}_+^{dK}$, we can select the *D-criterion* as our objective function [16]:

$$f_D(S) \triangleq \log \det (\mathcal{I}(S) + \mathcal{I}_0), \quad (14)$$

where

$$\mathcal{I}_0 \triangleq \text{Diag}(\{\tilde{\Sigma}_i^{-1}\}_{i=1}^M), \quad (15)$$

and Diag constructs a block-diagonal matrix with its input. In Equation (14), \mathcal{I} has been overloaded to take $S \subseteq \mathcal{S}$ as its argument:

$$\mathcal{I}(S) \triangleq \sum_{s \in S} \mathbb{E}_{\tilde{D}} \left[\nabla_{\mathbf{x}} \log p(\tilde{D}|\mathbf{x}; s) (\nabla_{\mathbf{x}} \log p(\tilde{D}|\mathbf{x}; s))^{\top} \right], \quad (16)$$

and

$$\begin{aligned} \log p(\tilde{D}|\mathbf{x}; s) &= \sum_{i=1}^M \sum_{j \in \nu(i)} s_j \log p(\tilde{d}_{ij}|\mathbf{x}_i) \\ &= \sum_{i=1}^M \sum_{j \in \nu(i)} \frac{s_j}{\sigma_{ij}^2} (\|\mathbf{x}_i - \mathbf{a}_j\|_2 - \tilde{d}_{ij})^2. \end{aligned} \quad (17)$$

Without the M terms corresponding to the information in the prior distribution, maximizing f_D is akin to minimizing the lower bound (via the CRLB) of the *volume* of a fixed level set of the MLE solution to Problem 1's distribution. By incorporating the prior, we instead minimize uncertainty of the posterior normal approximation in the MAP formulation [28, Sec. 4].

When actually computing Equation (16), we will approximate the expectation with the one-sample mean consisting of $\nabla_{\mathbf{x}} \log p(\tilde{D}|\mathbf{x}; s)$ evaluated at the observed measurements \tilde{D} [21]. Since we are proposing an offline design procedure in simulation, we will also have access to the ground truth value of the parameter $\bar{\mathbf{x}}$. Finally, since the measurements and prior distribution for each position with index $i \in [M]$ are independent of all other positions, the expression for the FIM in Equation (16) has a block diagonal structure, and the evaluation of Equation (14) can be efficiently implemented as the sum of the log-determinant of each diagonal submatrix.

C. D-Optimal Beacon Placement

We are now ready to introduce our problem of interest:

Problem 3 (Optimal Beacon Placement (OBP)). *For a finite set of candidate beacon positions $\mathcal{S} \subset \mathbb{R}^d$, solve*

$$\begin{aligned} \max_{S \subseteq \mathcal{S}} f_D(S) \\ \text{s.t. } |S| \leq K, \end{aligned} \quad (18)$$

where f_D is the information-theoretic cost function defined in Equation (14).

Problem 3 is a *set function maximization* problem and is NP-hard in general [15].

V. ALGORITHMS

In this section, we use the properties of f_D to derive a greedy algorithm with suboptimality guarantees.

A. Normalized, Monotone, and Submodular Set Functions

We begin with a few simple definitions for nonnegative set functions $f : 2^{\mathcal{S}} \rightarrow \mathbb{R}_+$ over a finite set \mathcal{S} . A set function is *normalized* if $f(\emptyset) = 0$, where $\emptyset = \{\}$ is the empty set. A set function is *monotone* if $f(A) \leq f(B)$ for any $A \subseteq B$. Finally, a set function is *submodular* if

$$f(A \cup \{e\}) - f(A) \geq f(B \cup \{e\}) - f(B) \quad (19)$$

Algorithm 1 Greedy Set Function Maximization

Input: A set function f , the set \mathcal{S} , a budget $K < |\mathcal{S}|$

Output: An approximate solution to Problem 3

```

1: function GREEDY
2:    $S \leftarrow \emptyset$  ▷ Initialize the solution set
3:   while  $|S| \leq K$  do ▷ Greedily select  $K$  elements
4:      $e^* = \operatorname{argmax}_{e \in \mathcal{S} \setminus S} f(S \cup \{e\})$ 
5:      $S \leftarrow S \cup \{e^*\}$ 
6:   end while
7:   return  $S$ 
8: end function

```

for any $A \subseteq B \subset \mathcal{S}$ and $e \in \mathcal{S} \setminus B$ [15]. When f is normalized, monotone, and submodular (NMS), the following helpful result holds.

Theorem 1 (Greedy Maximization of NMS Functions [15]). *Let $S_g \subset \mathcal{S}$ be the greedy solution given by Algorithm 1 to a set maximization problem like Problem 3 for some $K = |S_g|$ and a set function f . If f is NMS, then*

$$f(S_g) \geq (1 - 1/e)f(S^*) \approx 0.63f(S^*), \quad (20)$$

where S^* is the optimal solution.

Theorem 1 gives a suboptimality bound for the greedy solution of NMS set function maximization problems. Recall that a brute-force solution to Problem 3 requires the evaluation of $\binom{N}{K}$ subsets, whereas the greedy Algorithm 1 has a linear time complexity of $O(KN)$.

B. Analyzing f_D

In this section, we demonstrate that a constant modification of f_D is in fact NMS and therefore benefits from Theorem 1. First note that

$$\begin{aligned} f_D(\emptyset) &= \log \det \mathbf{I}_0 \\ &= \log \det \left(\operatorname{Diag}(\{\check{\Sigma}_i^{-1}\}_{i=1}^M) \right) \\ &= \log \det \left(\sum_{i=1}^M \check{\Sigma}_i^{-1} \right), \end{aligned} \quad (21)$$

which is well-defined so long as $\sum_{i=1}^M \check{\Sigma}_i^{-1} \succ 0$. Therefore, the biased objective function

$$f_D^+(S) \triangleq f_D(S) - f_D(\emptyset) \quad (22)$$

is normalized and has the same maximizer as f_D with respect to Problem 3.

Next, we observe that our approximate computation of the FIM involves adding a set of rank-one updates for each selected $e \in S \subset \mathcal{S}$. Shamaiah et al. [29] demonstrated that the log det of sums of this form are in fact monotone submodular functions. Therefore, we conclude that f_D^+ is NMS and can be greedily maximized for a solution that satisfies the suboptimality bound in Theorem 1.

C. CMA-ES Adaptation

Inspired by the use of a genetic algorithm in [6] for fiducial tag placement, in this section we consider an alternative to greedy set function maximization from the evolutionary optimization literature. Specifically, we apply the covariance matrix adaptation evolution strategy (CMA-ES) to OBP. CMA-ES is a robust evolutionary algorithm for nonlinear, nonconvex optimization problems [30]. Given an objective function, CMA-ES searches for an optimal solution by iteratively updating a multivariate normal distribution $\mathcal{N}(\mathbf{a}, \Sigma_a)$ of candidate solutions $\mathbf{a} \in \mathbb{R}^{dK}$, based on the covariance matrix Σ_a of the best solutions found in previous iterations. For the OBP, the objective function is represented by the information gain function f_D in Equation (14). The algorithm halts primarily based on convergence criteria that determine whether further iterations are likely to lead to significant improvements in the objective function. Since CMA-ES operates on continuous domains, we adapt its solutions to our discrete problem space by selecting the K beacons in \mathcal{S} closest (in a Euclidean sense) to the positions returned in its output \mathbf{a} .

VI. EXPERIMENTS

In this section, we systematically evaluate the performance of various algorithms for solving the OBP. All experiments were performed on a machine with an Intel(R) Xeon(R) W-11955M CPU @ 2.60 GHz, and 64.0 GB of RAM. We begin by describing benchmark algorithms in Section VI-A. Section VI-B contains a small empirical demonstration of our greedy algorithm's guaranteed suboptimality bounds. We conclude our experiments with larger-scale comparisons of our greedy method, CMA-ES, and the benchmark algorithms in Section VI-C, presenting statistics for runtime and the localization accuracy resulting from beacons selected by each method.

The value of $\tilde{\sigma}_i^2$ varies between experiments, but within each experiment it is the same for all $i \in [M]$. Finally, for each experimental trial, the ground truth positions $\bar{\mathbf{x}}_i$ are randomly sampled from their prior distribution $p_i(\mathbf{x}_i)$.

A. Benchmarks

In addition to the greedy information-theoretic method (henceforth called *Greedy*) and CMA-ES algorithm introduced in Section V, we include the following benchmark algorithms in our experiments:

- 1) *Random*: Randomly selects a subset of K beacons.
- 2) *Brute-Force*: Exhaustively searches all possible subsets of K beacons and returns the subset with highest information gain.
- 3) *Measurement Greedy*: Greedily selects beacons so as to maximize the total number of distance measurements received across all positions \mathbf{x}_i .
- 4) *Coverage Greedy*: Begins by greedily selecting beacon positions to ensure coverage of all positions \mathbf{x}_i .




Figure 3: Blueprints of a factory setting retrieved from [31] overlaid with a simulated robot trajectory in cyan and candidate beacon positions in red. Environments using this trajectory and random samplings of candidate beacon positions are used in the experiments of Section VI-C.

Once all positions are covered, it follows an approach identical to *Measurement Greedy*, greedily selecting beacons which maximize the number of distance measurements. This algorithm ensures (if possible) that all positions receive at least one measurement before optimizing for total measurements received by all positions.

Once a subset of beacons is selected, Problem 2 is solved with Newton's method [32] to yield $\hat{\mathbf{x}}$. Since we are not concerned with evaluating solution methods for Problem 2 and want to avoid local minima in its nonconvex objective function, we use the ground truth value of $\bar{\mathbf{x}}$ to initialize Newton's method. The root mean squared error (RMSE) is used to compare and evaluate the accuracy of the MLE solution:

$$\text{RMSE}(\hat{\mathbf{x}}) \triangleq \sqrt{\frac{1}{M} \sum_{i=1}^M \|\hat{\mathbf{x}}_i - \bar{\mathbf{x}}_i\|^2}. \quad (23)$$

B. Comparing Brute-Force and Greedy

Our first experiment compares the performance of the *Brute-Force* and *Greedy* algorithms for $K = 1, \dots, 7$ in a scenario with $N = 20$ candidate beacons and $M = 10$ positions in $d = 3$ dimensions. Because of the combinatoric runtime complexity of the *Brute-Force* method, we were limited to small values of K and N . Figure 4 empirically verifies that the suboptimality guarantees of Theorem 1 do in fact apply to our information theoretic cost function f_D^+ . The results in Figure 4 are for a cherrypicked set of candidate beacon positions which demonstrated the potential for a gap in performance between the *Greedy* and *Brute-Force* methods: the vast majority of randomly sampled beacon positions result in the *Greedy* algorithm recovering the globally optimal solution (in a fraction of the time required by *Brute-Force*).




Figure 4: Information-theoretic objective function f_D^+ with respect to the number of selected beacons K . The dashed blue line is the suboptimality bound of Theorem 1. For all values of K tested, our *Greedy* method satisfies the suboptimality bound and is very close to the globally optimal solution computed by *Brute-Force*.

C. Large Scale Experiment

In this section, we compare our algorithms for solving OBP on a larger 2D problem instance. Figure 3 illustrates a plausible trajectory for a robot within a factory. The red cross markers are $N = 50$ candidate beacon positions $\mathbf{a}_j \in \mathcal{S}$, the blue dots are $M = 30$ robot positions \mathbf{x}_i forming a trajectory, and the blue circles around positions are equiprobable contours corresponding to one standard deviation for their isotropic prior distributions with variance $\tilde{\Sigma}_i = \tilde{\sigma}_i^2 \mathbf{I}$. We use a constant range measurement variance of $\sigma_{ij}^2 = 25 \text{ m}^2$ across all trials, and we introduce a “cutoff” radius parameter $C > 0$ which adds sparsity to the bipartite measurement graph \mathcal{G} by preventing range measurements between positions \mathbf{x}_i and beacons \mathbf{a}_j separated by a Euclidean distance greater than C .

For each set of the parameter values K , C , and $\tilde{\sigma}_i$, 50 random trials (i.e., realizations of random ground truth samples $\bar{\mathbf{x}}$, random beacon positions \mathbf{a}_j , and noisy measurements \tilde{D}) in the environment displayed in Figure 3 were generated, and each of the four algorithms described in Section VI-A was used to solve the resulting instance of Problem 3.

Figure 5 summarizes the distribution of the RMSE for each algorithm for a variety of parameter values. Where unspecified, baseline values of $K = 5$, $C = 250$, and $\tilde{\sigma}_i = 8$ were used. The RMSE and runtime of the same set of experiments are displayed in Table I, with parameters changed from their baseline values indicated in the leftmost column.

Examining RMSE values in Figure 5, we note that *Greedy* improves relative to the other algorithms as prior uncertainty increases. We also see it outperform the heuristic greedy algorithms as K increases. In addition, we observe that

Table I: Experimental results for Section VI-C.

Algorithm		RMSE		Runtime	
		mean [m]	std [m]	mean [s]	std [s]
$K = 5$	<i>Random</i>	3.30	0.374	4.35×10^{-3}	9.57×10^{-4}
	<i>Greedy</i>	3.19	0.275	0.527	0.0670
	<i>Measurement Greedy</i>	3.14	0.384	4.87×10^{-3}	5.25×10^{-4}
	<i>Coverage Greedy</i>	3.21	0.336	6.08×10^{-3}	1.34×10^{-3}
$K = 10$	<i>Random</i>	2.82	0.359	0.0168	4.32×10^{-3}
	<i>Greedy</i>	2.67	0.257	3.27	0.432
	<i>Measurement Greedy</i>	2.82	0.324	0.0191	4.98×10^{-3}
	<i>Coverage Greedy</i>	3.23	0.350	9.91×10^{-3}	3.48×10^{-3}
$K = 15$	<i>Random</i>	2.37	0.358	0.0330	2.75×10^{-3}
	<i>Greedy</i>	2.25	0.250	10.0	0.899
	<i>Measurement Greedy</i>	2.44	0.366	0.0368	3.88×10^{-3}
	<i>Coverage Greedy</i>	3.07	0.518	0.0141	0.0107
$C = 150$	<i>Random</i>	3.73	0.370	4.03×10^{-3}	1.09×10^{-3}
	<i>Greedy</i>	3.63	0.319	0.501	0.123
	<i>Measurement Greedy</i>	3.52	0.348	4.48×10^{-3}	1.08×10^{-3}
	<i>Coverage Greedy</i>	3.44	0.290	6.36×10^{-3}	1.57×10^{-3}
$C = 300$	<i>Random</i>	3.11	0.418	5.05×10^{-3}	1.25×10^{-3}
	<i>Greedy</i>	3.00	0.318	0.612	0.132
	<i>Measurement Greedy</i>	3.04	0.373	5.62×10^{-3}	1.24×10^{-3}
	<i>Coverage Greedy</i>	3.21	0.362	5.68×10^{-3}	1.61×10^{-3}
$C = 450$	<i>Random</i>	2.59	0.305	5.58×10^{-3}	9.56×10^{-4}
	<i>Greedy</i>	2.55	0.231	0.689	0.134
	<i>Measurement Greedy</i>	2.70	0.319	6.28×10^{-3}	1.12×10^{-3}
	<i>Coverage Greedy</i>	2.88	0.361	5.10×10^{-3}	2.11×10^{-3}
$\tilde{\sigma}_i = 5$	<i>Random</i>	2.64	0.300	4.74×10^{-3}	1.29×10^{-3}
	<i>Greedy</i>	2.63	0.242	0.557	0.129
	<i>Measurement Greedy</i>	2.54	0.269	5.26×10^{-3}	1.27×10^{-3}
	<i>Coverage Greedy</i>	2.64	0.251	6.59×10^{-3}	1.97×10^{-3}
$\tilde{\sigma}_i = 10$	<i>Random</i>	3.68	0.396	4.32×10^{-3}	6.73×10^{-4}
	<i>Greedy</i>	3.45	0.360	0.525	0.0449
	<i>Measurement Greedy</i>	3.45	0.336	5.14×10^{-3}	1.09×10^{-3}
	<i>Coverage Greedy</i>	3.54	0.375	6.23×10^{-3}	1.59×10^{-3}
$\tilde{\sigma}_i = 15$	<i>Random</i>	4.36	0.523	4.23×10^{-3}	5.73×10^{-4}
	<i>Greedy</i>	3.98	0.411	0.527	0.0840
	<i>Measurement Greedy</i>	4.13	0.396	4.93×10^{-3}	8.46×10^{-4}
	<i>Coverage Greedy</i>	4.28	0.479	5.93×10^{-3}	1.12×10^{-3}

Measurement Greedy and *Coverage Greedy* perform well for small cutoff values, which is explained by their explicit consideration of coverage. As the cutoff parameter C increases, *Greedy* and *Random* perform similarly. Overall, *Greedy* demonstrates strong performance across various different experimental setups, suggesting that its D-optimal cost function is a reliable proxy for localization accuracy. Finally, while Table I shows that *Greedy* takes considerably longer to run than its simpler competitors, it is still tractable for the offline design of sensing algorithms and has the potential to run much faster if implemented with optimized code or hardware acceleration.

VII. CONCLUSION

We have presented a novel formulation and solution to the problem of optimally placing range beacons in an autonomous mobile robot’s workspace. Our approach leverages D-optimal Bayesian experimental design to reveal a submodular set function maximization problem with suboptimality guarantees for a fast greedy solver. Preliminary results show that this method is able to improve the accuracy of localization over simple baseline methods in many simulated scenarios.

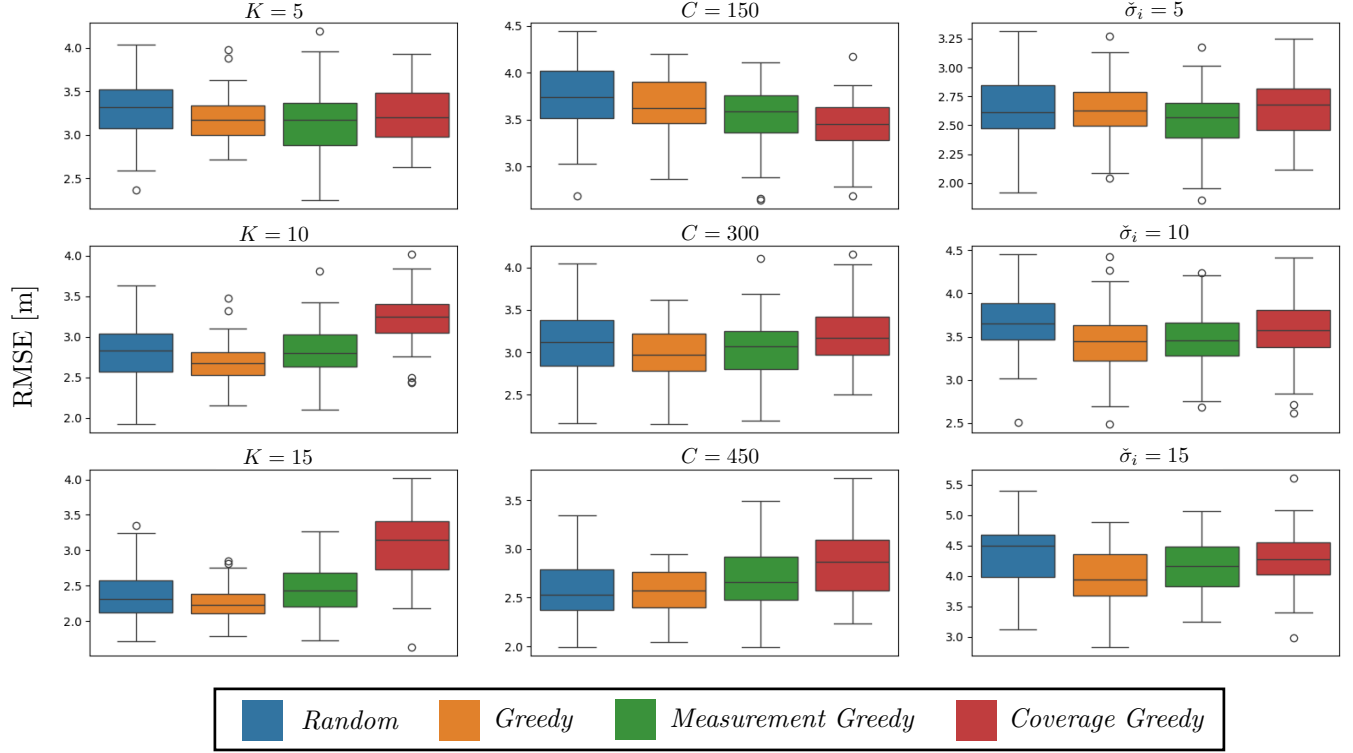


Figure 5: Boxplots of RMSE in meters for experiment in Section VI-C. Each subplot summarizes the distribution of the RMSE over 50 random trials. Where unspecified, parameters are set to their default values of $K = 5$, $C = 250$, and $\tilde{\sigma}_i = 8$.

In addition to conducting experiments with real data, future work will examine multiple measurement noise models, try more complex position priors than the uniform isotropic distributions considered in this work, and consider full pose estimates as in [3]. Additionally, if we remove the requirement that each position’s prior is independently distributed or introduce a relative motion model representing noisy odometry measurements between consecutive positions, we can realistically model a larger variety of localization problems. However, this change will remove the block diagonal structure of the Fisher information matrix exploited in the efficient evaluation of Equation (14). To make our approach tractable, we can investigate the convex relaxations used with E-optimal Bayesian experimental design in [21] and [33].

REFERENCES

- [1] B. Hofmann-Wellenhof, H. Lichtenegger, and J. Collins, *Global positioning system: theory and practice*. Springer Science & Business Media, 2012.
- [2] A. Raza, L. Lolic, S. Akhter, and M. Liut, “Comparing and Evaluating Indoor Positioning Techniques,” in *2021 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. Lloret de Mar, Spain: IEEE, Nov. 2021, pp. 1–8.
- [3] A. Goudar, T. D. Barfoot, and A. P. Schoellig, “Continuous-Time Range-Only Pose Estimation,” in *2023 20th Conference on Robots and Vision (CRV)*. Montreal, QC, Canada: IEEE, Jun. 2023, pp. 29–36.
- [4] A. Papalia, J. Morales, K. J. Doherty, D. M. Rosen, and J. J. Leonard, “SCORE: A Second-Order Conic Initialization for Range-Aided SLAM,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. London, United Kingdom: IEEE, May 2023, pp. 10 637–10 644.
- [5] F. Dumbgen, C. Holmes, and T. D. Barfoot, “Safe and Smooth: Certified Continuous-Time Range-Only Localization,” *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 1117–1124, Feb. 2023.
- [6] N. Kayhani, A. Schoellig, and B. McCabe, “Perception-Aware Tag Placement Planning for Robust Localization of UAVs in Indoor Construction Environments,” *Journal of Computing in Civil Engineering*, vol. 37, no. 2, pp. 1–32, Mar. 2023.
- [7] T. D. Barfoot, *State Estimation for Robotics*, 1st ed. Cambridge University Press, Jul. 2017.
- [8] G. Kantor and S. Singh, “Preliminary results in range-only localization and mapping,” in *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No.02CH37292)*, vol. 2. Washington, DC, USA: IEEE, 2002, pp. 1818–1823.

- [9] P. Newman and J. Leonard, "Pure range-only sub-sea SLAM," in *2003 IEEE International Conference on Robotics and Automation (Cat. No.03CH37422)*. Taipei, Taiwan: IEEE, 2003, pp. 1921–1926.
- [10] S. P. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, UK ; New York: Cambridge University Press, 2004.
- [11] A. Goudar, F. Dümbgen, T. D. Barfoot, and A. P. Schoellig, "Optimal Initialization Strategies for Range-Only Trajectory Estimation," *IEEE Robotics and Automation Letters*, pp. 1–8, 2024.
- [12] A. Papalia, A. Fishberg, B. W. O'Neill, J. P. How, D. M. Rosen, and J. J. Leonard, "Certifiably Correct Range-Aided SLAM," Sep. 2023.
- [13] J. Wang, R. K. Ghosh, and S. K. Das, "A Survey on Sensor Localization," *Journal of Control Theory and Applications*, vol. 8, no. 1, pp. 2–11, Feb. 2010.
- [14] A. Simonetto and G. Leus, "Distributed Maximum Likelihood Sensor Network Localization," *IEEE Transactions on Signal Processing*, vol. 62, no. 6, pp. 1424–1437, Mar. 2014.
- [15] A. Krause and D. Golovin, "Submodular Function Maximization," in *Tractability*, 1st ed., L. Bordeaux, Y. Hamadi, and P. Kohli, Eds. Cambridge University Press, Feb. 2014, pp. 71–104.
- [16] L. Carlone and S. Karaman, "Attention and Anticipation in Fast Visual-Inertial Navigation," *IEEE Transactions on Robotics*, vol. 35, no. 1, pp. 1–20, Feb. 2019.
- [17] V. Tzoumas, A. Jadbabaie, and G. J. Pappas, "Sensor placement for optimal Kalman filtering: Fundamental limits, submodularity, and algorithms," in *2016 American Control Conference (ACC)*. Boston, MA, USA: IEEE, Jul. 2016, pp. 191–196.
- [18] A. Collin and A. T. Espinoza, "Resilient Sensor Architecture Design and Tradespace Analysis for Autonomous Vehicle Localization and Mapping," *arXiv preprint arXiv:1907.08541*, no. arXiv:1907.08541, Jul. 2019.
- [19] K. Khosoussi, M. Giamou, G. S. Sukhatme, S. Huang, G. Dissanayake, and J. P. How, "Reliable Graphs for SLAM," *The International Journal of Robotics Research*, vol. 38, no. 2-3, pp. 260–298, Mar. 2019.
- [20] Y. Tian, K. Khosoussi, M. Giamou, J. How, and J. Kelly, "Near-Optimal Budgeted Data Exchange for Distributed Loop Closure Detection," in *Robotics: Science and Systems XIV*. Robotics: Science and Systems Foundation, Jun. 2018.
- [21] P. Kaveti, M. Giamou, H. Singh, and D. M. Rosen, "OASIS: Optimal Arrangements for Sensing in SLAM," Sep. 2023.
- [22] Q. Huang, J. DeGol, V. Fragoso, S. N. Sinha, and J. J. Leonard, "Optimizing Fiducial Marker Placement for Improved Visual Localization," *IEEE Robotics and Automation Letters*, vol. 8, no. 5, pp. 2756–2763, May 2023.
- [23] Z. C. Andrews and D. T. Ngo, "New Wireless Beacon Deployment for Maximum Coverage Efficiency in Indoor Localization," *IEEE Sensors Letters*, vol. 8, no. 2, pp. 1–4, Feb. 2024.
- [24] N. Rajagopal, S. Chayapathy, B. Sinopoli, and A. Rowe, "Beacon placement for range-based indoor localization," in *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. Alcalá de Henares, Spain: IEEE, Oct. 2016, pp. 1–8.
- [25] A. Moragrega, P. Closas, and C. Ibars, "Supermodular game for energy efficient TOA-based positioning," in *2012 IEEE 13th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*. Cesme, Turkey: IEEE, Jun. 2012, pp. 35–39.
- [26] H. Wang, N. Rajagopal, A. Rowe, B. Sinopoli, and J. Gao, "Efficient Beacon Placement Algorithms for Time-of-Flight Indoor Localization," in *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. Chicago IL USA: ACM, Nov. 2019, pp. 119–128.
- [27] I. Shames and T. H. Summers, "Rigid Network Design Via Submodular Set Function Optimization," *IEEE Transactions on Network Science and Engineering*, vol. 2, no. 3, pp. 84–96, Jul. 2015.
- [28] K. Chaloner and I. Verdinelli, "Bayesian Experimental Design: A Review," *Statistical Science*, vol. 10, no. 3, pp. 273–304, 1995.
- [29] M. Shamaiah, S. Banerjee, and H. Vikalo, "Greedy sensor selection: Leveraging submodularity," in *49th IEEE Conference on Decision and Control (CDC)*. Atlanta, GA, USA: IEEE, Dec. 2010, pp. 2572–2577.
- [30] N. Hansen, "The CMA evolution strategy: A tutorial," *CoRR*, vol. abs/1604.00772, 2016. [Online]. Available: <http://arxiv.org/abs/1604.00772>
- [31] CADENGINEER. Factory building layout plan and elevation cad template dwg. [Online]. Available: <https://cadtemplates.org/factory-building-layout-plan-and-elevation-cad-template-dwg>
- [32] J. Nocedal and S. J. Wright, *Numerical Optimization*, second edition ed., ser. Springer Series in Operation Research and Financial Engineering. New York, NY: Springer, 2006.
- [33] K. J. Doherty, D. M. Rosen, and J. J. Leonard, "Spectral Measurement Sparsification for Pose-Graph SLAM," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Kyoto, Japan: IEEE, Oct. 2022, pp. 01–08.