

OASIS: Optimal Arrangements for Sensing in SLAM

Pushyami Kaveti¹, Matthew Giamou², Hanumant Singh¹, and David M. Rosen¹.

Abstract—The number and arrangement of sensors on mobile robot dramatically influence its perception capabilities. Ensuring that sensors are mounted in a manner that enables accurate detection, localization, and mapping is essential for the success of downstream control tasks. However, when designing a new robotic platform, researchers and practitioners alike usually mimic standard configurations or maximize simple heuristics like field-of-view (FOV) coverage to decide where to place exteroceptive sensors. In this work, we conduct an information-theoretic investigation of this overlooked element of robotic perception in the context of simultaneous localization and mapping (SLAM). We show how to formalize the sensor arrangement problem as a form of subset selection under the E-optimality performance criterion. While this formulation is NP-hard in general, we show that a combination of greedy sensor selection and fast convex relaxation-based post-hoc verification enables the efficient recovery of *certifiably optimal* sensor designs in practice. Results from synthetic experiments reveal that sensors placed with OASIS outperform benchmarks in terms of mean squared error of visual SLAM estimates.

I. INTRODUCTION

Advances in sensing apparatus often drive improvements in robotic perception. For example, the availability of lidar sensors that produce accurate and dense point clouds has enabled substantial progress with self-driving vehicles. Similarly, the proliferation of lightweight, low-power, and affordable sensors for mobile phones and other consumer electronics has provided researchers with a growing arsenal of sensing modalities with which to outfit mobile robots. However, in spite (or perhaps because) of this perceptual embarrassment of riches, there is at present no definitive theoretical framework for understanding how sensors should be mounted on a mobile robot designed for some task.

In this paper, we address this significant research gap by developing a formal methodology to optimally arrange the sensors on a robot designed to perform SLAM. While sensor placement also concerns other tasks like tracking and collision detection, we focus on SLAM as it is a fundamental capability that significantly impact a robot’s capacity to perform other downstream tasks.

Our proposed method, *Optimal Arrangements for Sensing in SLAM* (OASIS), can be applied to optimize the design of any mapping or navigation system that fuses independent measurements from multiple sensors, and whose operation in representative environments can be easily simulated. In brief, our method accepts as input a finite (but potentially large) list of candidate sensor mountings, and selects a subset that maximizes an information-theoretic measure of




Fig. 1: Overview of OASIS. At each pose x_i , sensor c_j observes some subset of the landmarks. OASIS maximizes the minimum eigenvalue of the joint Fisher information matrix, which is composed of sub-matrices \mathcal{I}_j from each sensor c_j . In this example, the sensor “budget” limits us to selecting two out of the three candidate sensors. Finally, note that the discrete binary variables indicating sensor selection have been *relaxed* to a convex superset.

localization accuracy evaluated on a sample set of simulated robot trajectories. If a robot is expected to take a wide range of trajectories across many environments, a large number of heterogeneous scenarios can be simulated to find a sensor arrangement that performs well on average. Alternatively, our method can also benefit a robot constrained to specific paths in a fixed environment (e.g. as in warehouse operations) to find an arrangement that is tailored to its specific niche. OASIS has three key components that make it unique:

- 1) a design space consisting of a finite set of candidate sensor mountings, leading to a set function maximization problem that can be expressed as a binary integer program (IP);
- 2) a computationally tractable objective function based on E-optimality (i.e. minimum eigenvalue maximization) of a Schur complement of the Fisher information matrix of landmark-based SLAM; and
- 3) an efficient optimization approach that combines greedy sensor selection with a convex relaxation-based computation of an upper bound to verify optimality.

As we will see in Section II, OASIS provides roboticists with a general and easy-to-use tool for designing sensor rigs tailored to accurate SLAM. Our experiments in Section V demonstrate that OASIS scales gracefully to large numbers of candidate sensor mountings, while recovering sensor designs that are essentially optimal. The code and datasets are open sourced¹.

¹P. Kaveti, H. Singh, and D. Rosen are with Northeastern University, USA {kaveti.p, ha. singh, d.rosen}@northeastern.edu

²M. Gaimou is with McMaster University, CA giamou@mcmaster.ca

¹ https://github.com/PushyamiKaveti/optimal_camera_placement

II. RELATED WORK

In this section we summarize prior work on the design of sensor architectures for mobile robots, and set function optimization methods for perception tasks similar to ours.

A. Design of Perception Systems

Optimal sensor placement is a well-explored topic in the context of placing sensors in fixed positions in an environment [1][2][3]. However, sensor design on mobile agents has received relatively limited attention. Recent studies in this area differ from our approach in either the objective function being optimized or the specific downstream target task. Most methods aim to place sensors to maximize *coverage* of an environment [4], [2], [5], [6], [7], but there is relatively little work on the optimal design of *multi-sensor perception systems* for mobile robot localization and mapping.

Recent research has examined the impact of stereo camera orientation on the accuracy of visual odometry (VO) for aerial [8] and ground vehicles [9], favoring oblique camera angles to the direction of travel for performance improvement. Other work [10] uses a pan-tilt unit to adjust the "gaze" of a stereo camera toward salient features, enabling adaptability to various motion profiles and environments. Lastly, an observability analysis of stereo radar placement is experimentally validated in [11]. In contrast to these works, we address the placement of multiple generic sensors, utilizing an information-theoretic optimality criterion for SLAM.

In SLAM, information-theoretic objectives are commonly applied in active SLAM [12], feature selection [13], and dynamic sensor selection [14]. Previous works that have explored information-theoretic sensor placement deal with bespoke sensors like radar [11], localization with fixed sensors [15], or object detection tasks [16]. The recent literature on *co-design* of complex systems presents a framework for optimization with highly-coupled multidisciplinary constraints and objectives [17]. This approach has been applied to the Pareto-efficient design of autonomous drones [18] and self-driving vehicles [19], including sensor placement on an autonomous vehicle considering candidates formed by discretizing yaw angle [20]. Our method diverges from multidisciplinary design criteria and instead focuses on a much finer discretization of SE(3) to generate a comprehensive set of candidate sensor mountings. Furthermore, our approach replaces the common (but computationally expensive) D-optimality performance criterion with the (much more tractable) E-optimality criterion.

B. Set Function Optimization for Perception

Set function optimization is a natural formulation for many problems, but finding optimal solutions to these challenging combinatorial problems is NP-hard in general. The theory of submodular function maximization provides tractable solutions with *a priori* suboptimality guarantees [21]. Alternatively, convex relaxations can also provide fast approximate solutions and upper bounds on the optimal value.

Set function maximization with submodularity-based guarantees is used in [22] with an observability metric to greedily

select visual features that contribute to accurate solutions of SLAM. In [13], a similar approach is employed with information-theoretic criteria for monocular visual SLAM. This efficient feature selection strategy is extended in [23] to visual-inertial navigation which simulates a short horizon of robot dynamics to *anticipate* future motion and its effect on feature importance.

Set function maximization has also been employed in the pose graph optimization (PGO) formulation of SLAM, especially in multi-agent settings. In [24], graph-theoretic properties of PGO are shown to be intimately related to estimation quality and can be exploited to design greedy measurement selection algorithms with suboptimality guarantees or convex relaxations. A more efficient approach to graph pruning for PGO is used in [25]: instead of the costly D-criterion of [24], the E-criterion is approximately minimized with the Frank-Wolfe method. A related sparsification strategy is applied to the problem of resource-constrained cooperative SLAM (CSLAM) in [26]. Depending on the choice of optimality criterion, suboptimality guarantees exist for greedy solutions to this "resource-aware" formulation of CSLAM [27], [28].

In [29], the sensors mounted on autonomous agents are considered constant parameters while the poses of fiducial markers in an indoor construction environment are instead treated as design variables. Much like the sensor selection literature, a discrete set of candidate marker poses is considered, and an information-theoretic cost function is used. However, a genetic algorithm without any optimality guarantees is used to generate an optimal tag selection.

Our approach is most closely related to the sensor architecture design problem described in [30], which considers multi-objective optimization for a heterogeneous set of sensors and relatively few candidate mountings. OASIS is novel in that it selects a sensor arrangement that is frequently certifiably information-theoretically optimal, even when applied to a large number of candidate sensor mountings.

III. PROBLEM STATEMENT

In this section, we show how to formulate the optimal sensor arrangement problem for SLAM as an integer program (IP) with binary variables.

A. Parameterizing the space of sensor arrangements

In order to "search" for optimal sensor arrangements, we must first parameterize the space of possible designs. Given a model of a robot chassis, a natural approach might be to first decide upon the number and type of sensors to select, and then determine where on the robot to mount them. While intuitively appealing, this approach leads to a very challenging nonconvex problem that requires joint optimization over a product of (discrete) *sensor selection* variables and (continuous) *sensor pose* variables.

Instead, we propose to formulate the *sensor arrangement* problem as an instance of *subset selection*. Specifically, we assume that we are given a finite set S whose elements completely enumerate all possible *sensor mountings* (i.e. a decision to mount a *specific sensor* at a *specific pose* on the

robot chassis). Designing a sensor arrangement then amounts to selecting a *subset* $S \subset \mathcal{S}$ of specific sensor mountings.

Our formulation provides several advantages. First, it straightforwardly captures the fact that, due to mechanical constraints, most robot chassis have only a finite set of locations on which sensors may be mounted. Second, it avoids the need to perform joint optimization over discrete (sensor selection) and continuous (sensor pose) variables; instead, our approach models the set of possible sensor poses by *enumerating* a discrete set of distinct sensor mounting candidates in \mathcal{S} . While this leads to an increase in the number of decision variables versus a hybrid discrete-continuous formulation (whose magnitude depends upon how finely the set of candidate sensor poses is discretized), our approach has the advantage that (as we show in the sequel) we can leverage fast approximate combinatorial optimization algorithms to very efficiently recover high-quality solutions from the set of candidate mountings \mathcal{S} , even when this set is large.

B. Modeling SLAM performance of sensor arrangements

In this section we describe how to model the SLAM performance of a given sensor arrangement.

Let us imagine our mobile robot navigating through an initially unknown environment containing N_l uniquely recognizable landmarks $L \triangleq \{l_i\}_{i=1}^{N_l}$. As the robot explores, it moves through a sequence of poses $X = \{x_i\}_{i=1}^{N_x} \subset \text{SE}(d)$ while collecting measurements from its onboard sensors. Let us denote by $\tilde{Z} = \{\tilde{z}_i\}_{i=1}^M$ the complete set of measurements generated by *all* candidate sensor mountings in \mathcal{S} . We assume that each of these is sampled *independently* from a known sensor model of the form:²

$$\tilde{z}_i \sim p_i(\cdot | X, L) \quad i \in [M]. \quad (1)$$

Now let us consider our robot's SLAM performance under a candidate sensor arrangement $S \subset \mathcal{S}$. Without loss of generality, letting $N \triangleq |\mathcal{S}|$, we may label the candidates in \mathcal{S} by $1, \dots, N$, and then *identify* each subset $S \subset \mathcal{S}$ with a binary vector $s \in \{0, 1\}^N$ defined by:

$$s_i = \begin{cases} 1, & i \in S, \\ 0, & i \notin S. \end{cases} \quad (2)$$

Similarly, let us denote by $\ell: [M] \rightarrow N$ the function that assigns to each $i \in [M]$ the label of the sensor mounting that produced the i th measurement \tilde{z}_i . Using this notation, we can parameterize the joint likelihood of the data that would be available to our robot under sensor arrangement s as:

$$p(\tilde{Z} | X, L; s) \triangleq \prod_{i=1}^M p_i(\tilde{z}_i | X, L)^{s_{\ell(i)}}. \quad (3)$$

In turn, the specific instance of the SLAM maximum-likelihood estimation that our robot would solve under the

²In general each measurement \tilde{z}_i will only *explicitly* depend upon a small *subset* of the latent states (X, L) . While exploiting this fact is essential for solving SLAM problems efficiently in practice (as in e.g. factor-graph-based SLAM [31]), these details are unimportant for understanding the derivation of OASIS, and so we elide them here.

measurements afforded by arrangement s is:

$$(\hat{X}(s), \hat{L}(s)) = \underset{X, L}{\operatorname{argmin}} \left\{ \sum_{i=1}^M -s_{\ell(i)} \log p_i(\tilde{z}_i | X, L) \right\}. \quad (4)$$

C. Fisher Information and the Cramér-Rao Lower Bound

Recall that the *Cramér-Rao Lower Bound* (CRLB) provides a *lower bound* (in the Loewner order sense) on the achievable covariance of any unbiased maximum-likelihood estimator $\hat{\Theta}$ of a fixed but unknown parameter Θ [32]:

$$\operatorname{Cov} \hat{\Theta} \succeq \mathcal{I}(\Theta)^{-1}, \quad (5)$$

where the matrix $\mathcal{I}(\Theta)$ appearing on the right-hand side of (5) is the *Fisher information matrix* (FIM):

$$\mathcal{I}(\Theta) \triangleq \mathbb{E}_{\tilde{Z}} \left[-\frac{\partial^2}{\partial \Theta^2} \log p(\tilde{Z} | \Theta) \right]. \quad (6)$$

For the SLAM likelihood (3), the CRLB takes the form:

$$\mathcal{I}(X, L; s) = \sum_{i=1}^M s_{\ell(i)} \mathbb{E}_{\tilde{z}_i} \left[-\frac{\partial^2}{\partial^2(X, L)} \log p(\tilde{z}_i | X, L) \right]. \quad (7)$$

Note that the conditional independence of the measurements \tilde{z}_i in (1) implies that $\mathcal{I}(X, L; s)$ is the *sum* of the information matrices contributed by *each individual observation* \tilde{z}_i that is available under design s . Equivalently, writing:

$$\mathcal{I}_k(X, L) \triangleq \sum_{\{i \in [M] | \ell(i)=k\}} \mathbb{E}_{\tilde{z}_i} \left[-\frac{\partial^2}{\partial^2(X, L)} \log p(\tilde{z}_i | X, L) \right] \quad (8)$$

for the sum of information matrices from all measurements generated by sensor k , equation (7) is equivalent to:

$$\mathcal{I}(X, L; s) = \sum_{k=1}^N s_k \mathcal{I}_k(X, L). \quad (9)$$

That is: the FIM for the SLAM estimation problem (4) under sensor arrangement s is simply the sum of the information provided by each individual sensor included in s .

D. Performance Criteria for Sensor Arrangements

The CRLB implies that if we want to recover a SLAM estimate $(\hat{X}(s), \hat{L}(s))$ from (4) with a “small” uncertainty, we must choose a sensor arrangement s such that the corresponding FIM $\mathcal{I}(X, L; s)$ is as “large” as possible.

To that end, inspired by [25], we propose to use the *E-optimality* criterion as a performance measure for optimizing the design of sensor arrangements. In brief, this approach requires maximizing $\lambda_1(\mathcal{I})$, the minimum eigenvalue of \mathcal{I} . The advantage of E-optimality versus the more common D-optimality (which maximizes $\log \det(\mathcal{I})$) is that the latter depends upon the *entire spectrum* of \mathcal{I} , while the former requires only the *single smallest* eigenvalue; this can be calculated very efficiently, even for very large matrices. Moreover, since $\lambda_1(\mathcal{I})$ lower-bounds the spectrum of \mathcal{I} , maximizing this quantity can be interpreted as maximizing a lower bound on $\log \det(\mathcal{I})$ itself.

We also note that in many SLAM applications, we are mainly interested in the *robot pose estimates* \hat{X} ; the landmark locations \hat{L} are only interesting insofar as they support accurate robot localization. In this case, our main concern is in minimizing $\text{Cov} \hat{X}$, the marginal covariance of the pose estimates. In light of (5) and the 2x2 block-matrix inversion formula, the relevant form of the CRLB in this case is:

$$\text{Cov} \hat{X}(s) \succeq \text{Schur}(\mathcal{I}(X, L; s)), \quad (10)$$

where here $\text{Schur}(\mathcal{I})$ denotes the generalized Schur complement of $\mathcal{I}(X, L; s)$ with respect to the landmark variables L [33]. Thus, we propose to use the following objective:

$$f_E(X, L; s) \triangleq \lambda_1(\text{Schur}(\mathcal{I}(X, L; s))). \quad (11)$$

E. Optimal sensor arrangement

We are now ready to formalize the optimal sensor arrangement problem. Given a set of candidate sensor mounts \mathcal{S} , a realization of an environment L and robot trajectory X , and a number of sensors K to select, our task is to find the K -cardinality subset $S \subset \mathcal{S}$ that maximizes the objective (11):

Problem 1 (Optimal Sensor Arrangement).

$$f^* = \max_{s \in \{0,1\}^N} f(X, L; s) \quad \text{s.t.} \quad \sum_{i=1}^N s_i = K. \quad (12)$$

IV. FAST APPROXIMATION ALGORITHMS

In this section we describe fast approximate algorithms for solving the sensor arrangement problem (Problem 1). We resort to approximations because finding the global optimum of Problem 1 is NP-hard, and therefore cannot (in general) be solved in polynomial time [34]. Our approach employs both a simple greedy method described in Section IV-A and a convex relaxation developed in Section IV-B. Throughout this section we will drop explicit reference to X and L to ease notation (since these are not variables of optimization).

A. Greedy Maximization

As their name suggests, greedy set maximization algorithms iteratively construct a solution set $S \subseteq \mathcal{S}$ by appending, in each iteration, an element $x \in \mathcal{S} - S$ that leads to the greatest marginal gain in the value of the objective. When the objective f is *normalized monotone submodular* [21], the greedy solution $S_g \subset \mathcal{S}$ is guaranteed to satisfy

$$f(S_g) \geq (1 - \frac{1}{e})f(S^*) \approx 0.63f(S^*), \quad (13)$$

where S^* is a globally maximizer of 12. Unfortunately, while f_E is monotonic and normalized, λ_1 only satisfies *approximate* forms of submodularity [35], [36].

B. Convex Relaxation

Let us relax the (nonconvex) *binary* constraints appearing in Problem 1 to (convex) *Boolean* constraints:

Problem 2 (Boolean Relaxation of Problem 1).

$$\mu^* = \max_{\omega \in [0,1]^N} f(\omega) \quad \text{s.t.} \quad \sum_{i=1}^N \omega_i = K. \quad (14)$$

Algorithm 1 OASIS Algorithm

Input: Concave obj. function f , feasible point ω of (14).

Output: A feasible solution s_g of Problem 1 and an upper bound $\mu^* \geq f^*$ on Problem 1's optimal value.

```

1: function OASIS( $\omega$ )
2:    $s_g \leftarrow \text{GREEDY}(f)$   $\triangleright$  Greedy solution; Sec. (IV-A)
3:    $(\mu^*, \omega^*) \leftarrow \text{FRANKWOLFEAC}(\omega; f)$   $\triangleright$  Solve (14)
4:   return ( $s_g, \mu^*, \omega^*$ )
5: end function

```

Observe that if f is concave, then (14) is a *convex program*. Consequently, the computational tractability of Problem 2 hinges on the concavity of our objective function f . Fortunately, the following proposition (proven in the Appendix³) states that the E-optimality performance criterion f_E defined in (11) is indeed concave.

Proposition 1 (Concavity of f_E). *The function f_E defined in (11) is concave on the domain $[0, 1]^N$.*

It follows that Problem 2 is convex when using the E-optimality criterion (11), and thus can be solved *globally optimally* using standard convex optimization methods. Following the insights of [25], we thus propose to solve Problem 2 using the Frank-Wolfe method.

Let us consider how the optimal values of Problems 1 and 2 compare. Since Problem 2 is a relaxation of (12), its optimal value μ^* provides an upper bound on the optimal value f^* of (12). Conversely, we clearly have $f(s) \leq f^*$ for any feasible s in (12). These inequalities together imply:

$$\mu^* - f(s) \geq f^* - f(s) \quad (15)$$

for any feasible s in (12). The significance of inequality (15) is that it enables us to use the optimal value μ^* of Problem 2 (which we can compute efficiently) in order to bound the suboptimality $f^* - f(s)$ of *any* feasible solution s of Problem 1. In particular, as we will see in Section V, this will provide a practical means of *verifying* the (global) optimality of candidate solutions of Problem 1.

C. The OASIS algorithm

The entire OASIS procedure is summarized in Algorithm 1. In brief, our method applies sequential greedy set maximization to obtain a feasible solution s_g of Problem 1, and then solves the convex relaxation Problem 2 (using the Frank-Wolfe method) to obtain an upper bound μ^* on Problem 2's optimal value that we can use to bound s_g 's suboptimality via (15). As we show in Section V, this simple approach enables us to recover certifiably optimal solutions of the sensor arrangement problem (12) in practice.

V. EVALUATION

Here we describe our experimental evaluation, and present results of greedy and convex optimization approaches in terms of both the information-theoretic criterion f_E in (11) and downstream SLAM performance.

³Due to space constraints, the proof and additional figures are provided in the supplement [37]

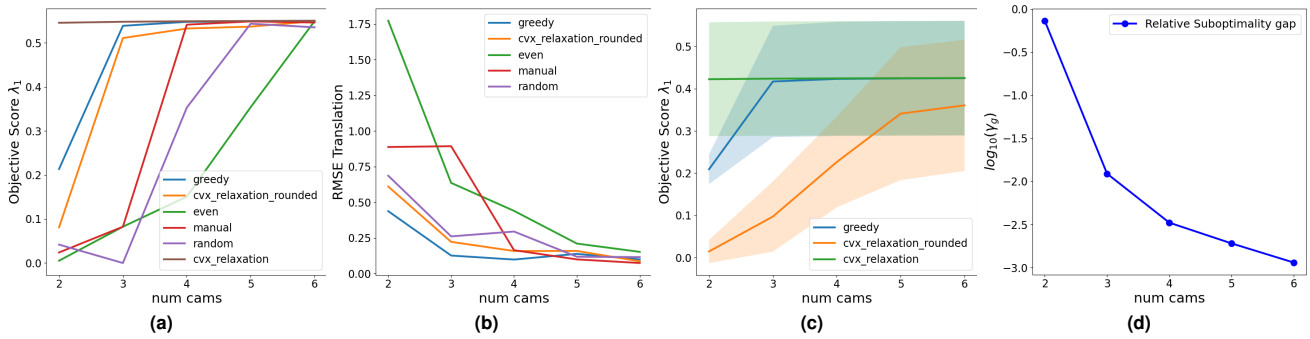


Fig. 2: Experimental Results I: Quantitative results of optimal placement on the synthetic dataset of 50 simulations of a random motion of the sensor rig with varying number of selected cameras, k . (a) and (b) Show the trend of the optimized score of the objective function λ_1 and median RMSE of the translational component of the pose estimates computed from the graph resulting from the optimal camera selection with respect to the ground truth computed across simulations. (c) Gives a closer look at the mean and standard deviation of the objective scores of greedy, convex relaxation solutions before and after k -max rounding across simulations. The greedy optimization results in a near-optimal solution, as demonstrated by its closeness to the score of the unrounded convex relaxation approach, which gives the upper bound on the optimal value, especially for $k > 2$. (d) Shows that we achieve a very tight relative sub-optimality gap, $\gamma^* = f(w^*) - x_g/x_g$ asserting the effectiveness of submodular greedy optimization.

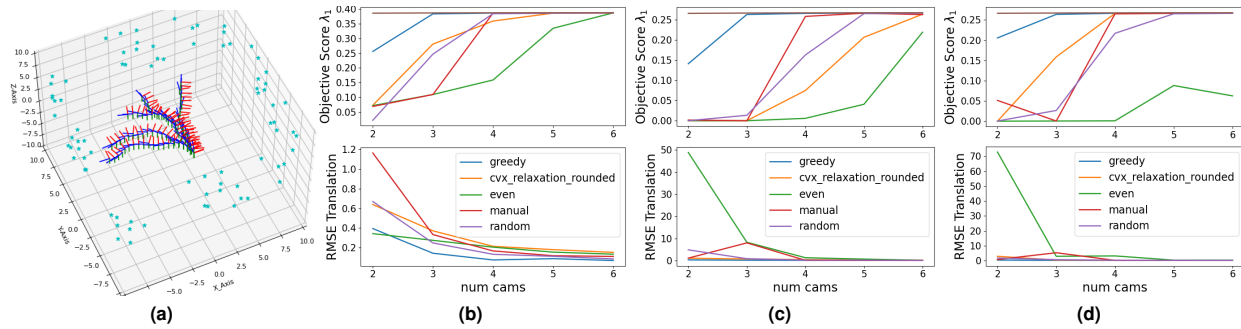


Fig. 3: Experimental Results II: (a) The synthetic data collection setup. A simulated room-like environment with landmarks and random trajectories from the top view. (b-d) Quantitative results showing median score and RMSE of the benchmarking algorithms for Circular, Forward and Lateral motions.

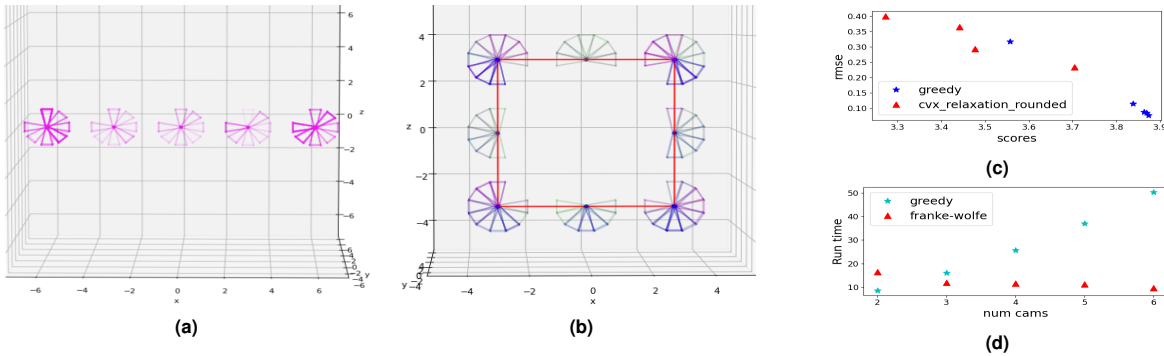


Fig. 4: Experimental Results III. Visualization of the greedy optimal selection across multiple experiments overlaid for a candidate pool (a) lying on a linear array configuration and (b) regularly spaced in both translation and orientation. Deeper/darker colors indicate a higher frequency of selection. (c) Shows that the score λ_1 is inversely related to the RMSE of SLAM pose estimates for both greedy and convex relaxation approaches. Thus, E-optimality improves SLAM performance. (d) Run time comparison of greedy optimization with the convex relaxation approach. Time complexity of the greedy method increases linearly with the number of cameras, while there is not much effect on the convex relaxation approach.

A. Experimental setup

Datasets We employ landmark-based visual SLAM as the evaluation framework for our proposed optimization methods. We developed a Python-based simulation system¹ tailored to generate different synthetic datasets for assessing the performance of various sensor arrangement methods (shown in fig. 3(a)). Within this simulated environment, a room-like setting was constructed, housing a mobile multi-camera sensor rig programmed to exhibit diverse motion patterns, including circular movements, straight-line forward and lateral motions, and random movements. The simulated environment offers several advantages, including the capacity

to conduct a large number of reproducible experiments with fixed parameters related to various aspects such as the number of poses, camera fields of view (FoVs), landmark distribution, environmental characteristics, rig motion patterns, among others. Moreover, it provides support for incremental development and debugging. In our simulations, we generate four distinct trajectory types: circular, lateral, forward, and random motions, using a simulated sensor rig equipped with 68 different candidate camera poses regularly spaced on a square frame as shown in fig. 4(b). We aggregate the results from 50 simulations for each trajectory type, with each simulation involving varying landmark distribution.

Benchmark Algorithms In all experiments, we compare our algorithms with the following benchmarks:

- 1) random sampling from the candidate pool (random);
- 2) a standard configuration from a similar commercially available autonomous vehicle (manual);
- 3) a “uniform coverage” heuristic which evenly spaces the cameras in terms of viewing direction and position throughout the vehicle’s feasible set of SE(3) (even);

These heuristic-based methods act as baselines to demonstrate the effectiveness of set function maximization for the camera placement problem.

Metrics For each dataset, we consider a selection of $k=2,3,\dots,6$ cameras from the candidate pool for the optimal camera placement problem. We evaluate each method by comparing the score of the objective function f_E in (11), and the results of the maximum likelihood estimation (MLE) of the SLAM problem for the camera sets selected by respective methods. To evaluate SLAM performance, we compute the Root Mean Square Error (RMSE) between the MLE estimates of poses obtained by performing optimization (using GTSAM [38]) on the factor graph induced by camera selection and the ground truth poses. We also assess the performance of greedy and convex relaxation approaches in terms of run-time and ability to produce optimal solutions.

B. Results

1) *Comparison with benchmarks:* Fig. 2 and 3 provides a concise summary of quantitative outcomes from the benchmarking algorithms discussed in Section V-A, applied to various camera selections. It is evident that our proposed greedy optimization consistently outperforms the even, manual, and random selection methods. This performance gap becomes more pronounced when dealing with a smaller number of cameras to be positioned (k). This observation underscores the effectiveness of the set function maximization approach for achieving optimal sensor placement.

Fig. 4(c) illustrates the relationship between the value of the objective $f_E(s)$ for our optimized sensor arrangements s , and the empirical RMSE of the SLAM estimates $\hat{X}(s)$ obtained using that design within our experiments. These results demonstrate that improving the value of our proposed information-theoretic criterion f_E in (11) is indeed predictive of improved downstream SLAM performance.

2) *Greedy vs Convex Relaxation:* Fig. 2 compares the objective value achieved by the greedy method with (i) the upper bound $\mu^* \geq f^*$ on Problem 1’s optimal value provided by the Boolean relaxation (14), and (ii) the objective value obtained by *rounding* the solution ω^* of the convex relaxation using K -max rounding.

These results reveal several very interesting trends. First, we observe that the upper bound μ^* on Problem 1’s optimal value f^* provided by the convex relaxation and the lower bound $f^* \geq f(s_g)$ provided by the greedy solution are remarkably close across the majority of our experiments. In light of inequality (15), this shows that (i) the heuristic greedy method is remarkably effective in finding high-quality sensor arrangements in these experiments, and (ii)

the convex relaxation’s optimal value μ^* is a remarkably sharp approximation of Problem 1’s optimal value, enabling us to *certify* the optimality of the greedy method’s solutions. Indeed, examining Fig. 2(d), which plots the upper bound $(\mu^* - f(s_g))/f(s_g) \geq (\mu^* - f(s_g))/f^*$ on s_g ’s *relative* suboptimality, reveals that the greedy method succeeds in finding solutions that are (at most) 1 – 2% suboptimal in the majority of cases. (Note that this suboptimality bound is a substantial improvement on the usual $1 - 1/e$ suboptimality bound obtained from greedy submodular maximization). At the same time, however, the sensor arrangements recovered by rounding the actual *solutions* ω^* returned by Problem 2 are consistently significantly suboptimal; evidently while the *objective value* of Problem 2’s maximizers very tightly approximates f^* , the solutions ω^* are not close to integral.

Altogether these results show that while greedy maximization or convex relaxation would not *individually* suffice to obtain certifiably optimal solutions, their combination is remarkably effective in practice. Finally, we note that one limitation of the greedy method is that its per-iteration complexity is linear in N (see fig. 4(d)). This However, this is not a significant hindrance for sensor rig design, as it can be conducted as an offline task.

3) *Qualitative Analysis:* Figures 4(a) and (b) showcase the qualitative outcomes of the greedy optimization approach. In these visualizations, we have superimposed the candidate pool and the camera selections optimized through the greedy approach across all experimental runs. This graphical representation aids us in identifying hot spots or key areas of interest for camera placement. In Figure 4(a), we depict the results of selecting two cameras from a candidate set arranged linearly. Notably, the extremities of the linear array are consistently favored. This observation aligns with the preference for a larger baseline, a characteristic advantageous when utilizing two cameras. Turning our attention to Fig. 4(b), we consider the scenario of positioning cameras on a mobile robot. Here, we observe a notable preference for the four extreme corners over other potential positions.

VI. CONCLUSION

In this paper we proposed a formal methodology for optimally arranging the sensors on a mobile robot designed to perform SLAM. Our approach formalizes the design task as an optimal subset selection problem under a computationally tractable E-optimality performance measure. While subset selection problems are NP-hard in general, we also developed a fast approximate optimization scheme that combines greedy sensor selection with convex-relaxation-based post-hoc suboptimality bounds. Our experimental evaluations show that our approach is remarkably effective in practice, enabling the efficient recovery of sensor arrangements that are within 1-2% of the optimal value.

While we have currently focused on camera selection for visual SLAM, our future work will explore designing a broader class of sensor rigs with heterogeneous sensors, incorporating a richer set of design constraints (SWAPC) than cardinality, and conducting real-world experiments.

REFERENCES

- [1] Y. Morsly, N. Aouf, and M. S. Djouadi, "On the optimal placement of multiple visual sensor based binary particle swarm optimization," *IFAC Proceedings Volumes*, vol. 42, no. 19, pp. 279–285, 2009.
- [2] S. Nikolaidis, R. Ueda, A. Hayashi, and T. Arai, "Optimal camera placement considering mobile robot trajectory," in *2008 IEEE International Conference on Robotics and Biomimetics*. IEEE, 2009, pp. 1393–1396.
- [3] A. Mavrinac and X. Chen, "Modeling coverage in camera networks: A survey," *International journal of computer vision*, vol. 101, pp. 205–226, 2013.
- [4] J. Zhao, R. Yoshida, S.-c. S. Cheung, and D. Haws, "Approximate Techniques in Solving Optimal Camera Placement Problems," *International Journal of Distributed Sensor Networks*, vol. 9, no. 11, p. 241913, Nov. 2013.
- [5] S. Indu, S. Srivastava, and V. Sharma, "Optimal camera placement and orientation of a multi-camera system for self driving cars," in *Proceedings of the 2020 4th International Conference on Vision, Image and Signal Processing*, 2020, pp. 1–5.
- [6] M. Brandao, R. Figueiredo, K. Takagi, A. Bernardino, K. Hashimoto, and A. Takanishi, "Placing and scheduling many depth sensors for wide coverage and efficient mapping in versatile legged robots," *The International Journal of Robotics Research*, vol. 39, no. 4, pp. 431–460, 2020.
- [7] Z. Liu, M. Arief, and D. Zhao, "Where should we place lidars on the autonomous vehicle?-an optimal design approach," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 2793–2799.
- [8] J. Kelly and G. S. Sukhatme, "An experimental study of aerial stereo visual odometry," *IFAC Proceedings Volumes*, vol. 40, no. 15, pp. 197–202, 2007.
- [9] V. Peretroukhin, J. Kelly, and T. D. Barfoot, "Optimizing camera perspective for stereo visual odometry," in *2014 Canadian Conference on Computer and Robot Vision*. IEEE, 2014, pp. 1–7.
- [10] T. Manderson, A. Holliday, and G. Dudek, "Gaze Selection for Enhanced Visual Odometry During Navigation," in *2018 15th Conference on Computer and Robot Vision (CRV)*. IEEE, 2018, pp. 110–117.
- [11] A. Corominas-Murtra, J. Vallve, J. Sola, I. Flores, and J. Andrade-Cetto, "Observability analysis and optimal sensor placement in stereo radar odometry," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2016, pp. 3161–3166.
- [12] R. Sim and N. Roy, "Global A-Optimal Robot Exploration in SLAM," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*. IEEE, 2005, pp. 661–666.
- [13] Y. Zhao and P. A. Vela, "Good Feature Selection for Least Squares Pose Optimization in VO/VSLAM," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Oct. 2018, pp. 1183–1189.
- [14] H.-P. Chiu, X. S. Zhou, L. Carlone, F. Dellaert, S. Samarasekera, and R. Kumar, "Constrained optimal selection for multi-sensor robot navigation using plug-and-play factor graphs," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 663–670.
- [15] D. Szalóki, S. Kolumbán, K. Csorba, and G. Tevesz, "Camera Placement Optimization in Object Localization Systems," *Acta Cybernetica*, vol. 22, no. 1, pp. 211–228, 2015.
- [16] S. Roos, T. Brühl, M. Pfeffer, L. Ewecker, and W. Stork, "A method for evaluation and optimization of automotive camera systems based on simulated raw sensor data," in *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2022, pp. 1334–1341.
- [17] A. Censi, "Monotone Co-Design Problems; or, Everything is the Same," in *2016 American Control Conference (ACC)*. IEEE, 2016, pp. 1227–1234.
- [18] G. Zardini, A. Censi, and E. Frazzoli, "Co-Design of Autonomous Systems: From Hardware Selection to Control Synthesis," in *2021 European Control Conference (ECC)*. IEEE, 2021, pp. 682–689.
- [19] G. Zardini, D. Milojevic, A. Censi, and E. Frazzoli, "Co-design of Embodied Intelligence: A Structured Approach," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Sep. 2021, pp. 7536–7543.
- [20] A. Collin, A. Siddiqi, Y. Imanishi, Y. Matta, T. Tanimichi, and O. de Weck, "A multiobjective systems architecture model for sensor selection in autonomous vehicle navigation," in *Complex Systems Design & Management: Proceedings of the Tenth International Conference on Complex Systems Design & Management, CSD&M Paris 2019*. Springer, 2020, pp. 141–152.
- [21] A. Krause and D. Golovin, "Submodular Function Maximization," in *Tractability*, 1st ed. Cambridge University Press, Feb. 2014, pp. 71–104.
- [22] G. Zhang and P. A. Vela, "Good Features to Track for Visual SLAM," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Jun. 2015, pp. 1373–1382.
- [23] L. Carlone and S. Karaman, "Attention and Anticipation in Fast Visual-Inertial Navigation," *IEEE Transactions on Robotics*, vol. 35, no. 1, pp. 1–20, 2018.
- [24] 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.
- [25] 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)*. IEEE, 2022, pp. 01–08.
- [26] M. Giamou, K. Khosoussi, and J. P. How, "Talk Resource-Efficiently to Me: Optimal Communication Planning for Distributed Loop Closure Detection," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 3841–3848.
- [27] Y. Tian, K. Khosoussi, M. Giamou, J. P. How, and J. Kelly, "Near-Optimal Budgeted Data Exchange for Distributed Loop Closure Detection," *Robotics: Science and Systems XIV*, 2018.
- [28] Y. Tian, K. Khosoussi, and J. P. How, "A Resource-Aware Approach to Collaborative Loop-Closure Detection with Provable Performance Guarantees," *The International Journal of Robotics Research*, vol. 40, no. 10-11, pp. 1212–1233, Sep. 2021.
- [29] 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.
- [30] 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.
- [31] F. Dellaert and M. Kaess, "Factor graphs for robot perception," *Foundations and Trends in Robotics*, vol. 6, no. 1-2, pp. 1–139, 2017.
- [32] T. Ferguson, *A Course in Large Sample Theory*. Boca Raton, FL: Chapman & Hall/CRC, 1996.
- [33] C.-K. Li and R. Mathias, "Extremal Characterizations of the Schur Complement and Resulting Inequalities," *SIAM Review*, vol. 42, no. 2, pp. 233–246, Jan. 2000.
- [34] M. Shamaiah, S. Banerjee, and H. Vikalo, "Greedy sensor selection: Leveraging submodularity," in *49th IEEE conference on decision and control (CDC)*. IEEE, 2010, pp. 2572–2577.
- [35] L. Chamon and A. Ribeiro, "Approximate supermodularity bounds for experimental design," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [36] A. Hashemi, M. Ghasemi, H. Vikalo, and U. Topcu, "Submodular observation selection and information gathering for quadratic models," in *International Conference on Machine Learning*. PMLR, 2019, pp. 2653–2662.
- [37] P. Kaveti, M. Giamou, H. Singh, and D. M. Rosen, "Oasis: Optimal arrangements for sensing in slam," *arXiv preprint arXiv:2309.10698*, 2023.
- [38] F. Dellaert, "Factor graphs and GTSAM: A hands-on introduction," Institute for Robotics & Intelligent Machines, Georgia Institute of Technology, Tech. Rep. GT-RIM-CP&R-2012-002, Sep. 2012.