Near-Optimal Task-Driven Sensor Network Configuration

Chase St. Laurent [a], Raghvendra V. Cowlagi [b],

[a] *AMETEK Inc., Peabody, MA 019060, USA*

[b] *Aerospace Engineering Department, Worcester Polytechnic Institute, Worcester, MA 01609, USA*

Abstract

A coupled path-planning and sensor configuration method is proposed. The path-planning objective is to minimize exposure to an unknown spatially-varying scalar field, called the threat field, measured by a network of sensors. Gaussian Process regression is used to estimate the threat field from these measurements. Crucially, the sensors are configurable, i.e., parameters such as location and size of field of view can be changed. A main innovation of this work is that sensor configuration is performed by maximizing a so-called task-driven information gain (TDIG) metric, which quantifies uncertainty reduction in the cost of the planned path. For computational efficiency, a surrogate metric called the self-adaptive mutual information (SAMI) is introduced and shown to be submodular. The proposed method is shown to vastly outperform traditionally decoupled information-driven sensor configuration in terms of the number of measurements required to find near-optimal plans.

Key words: Sensor Networks; Cooperative Perception; Trajectory and Path Planning; UAVs; Bayesian Methods.

1 Introduction

In a large-scale autonomous system, a mobile agent may leverage information about the environment gained from an exteroceptive sensor network. Consider a post-disaster scenario where many roads are flooded and the problem is to plan a safe route to transport people from a dangerous location to a safe location. A network of unmanned aerial vehicles (UAVs) visually surveys the extent of flooding and identify safe roads. This application example sheds light on two related issues, namely, how to: (1) *configure* sensors, e.g., send the UAVs to different locations, and (2) plan with *as few measurements as possible,* without significantly reducing the quality of the plan. A research question that naturally arises is: *how do we optimally configure sensors to find a near-optimal plan with a minimal number of measurements?* To address this question, we study the problem of finding a path of minimum threat exposure (see Fig. 1(a)) for a mobile agent traversing in an unknown environment. Information about the threat is gained from an exteroceptive sensor network.

Related Work: The two main components of this work, which are traditionally studied independently, are optimal path-planning under uncertainty and sensor configuration. Path-planning is well studied for objectives such as minimum length, maximum traversal utility, and obstacle avoidance (Aggarwal and Kumar, 2020). Graph-based algorithms such as Dijkstra's algorithm, A*, and its variants (Hart et al., 1968) are frequently used. Partially observable Markov Decision process models are typically used in uncertain planning tasks with onboard sensing (Kurniawati et al., 2012).

The literature on sensor configuration is largely restricted to sensor *placement*. Optimal sensor placement addresses minimizing uncertainty, or maximizing spatial coverage and communication reachability (Ramsden, 2009). Optimization metrics include Kullback-Leibler divergence, Fisher information, and mutual information (Cochran and Hero, 2013). Sensor placement applications include estimation of gaseous plumes (Demetriou et al., 2013), cooperative tracking of forest fires (Merino et al., 2006), and observing dynamics of volcanic ash (Madankan et al., 2014). Near-optimal sensor placement for linear inverse problems are studied by Ranieri et al. (2014). Clustering-based algorithms such as k-means (Li et al., 2016) and density-based clustering (Yoganathan et al., 2018) for sensors in office spaces.

The comparison of *task-driven* versus *information-driven* sensor placement in tracking applications were studied by Kreucher et al. (2005). Task-driven methods are those where sensor placement is somehow coupled to a planning or control task, whereas information-driven methods are decoupled. Similarly, Tzoumas et al. (2021)

Email addresses: clstlaurent@wpi.edu (Chase St. Laurent), rvcowlagi@wpi.edu (Raghvendra V. Cowlagi).

address optimal sensor selection for linear quadratic Gaussian feedback control systems. Target tracking UAVs with limited field of view are studied for simultaneous path-planning and optimal sensing position (Skoglar et al., 2006). Allen et al. (2009) study hierarchical path-replanning concurrently with data fusion.

In contrast to the existing literature where planning and sensor configuration are studied independently, we propose a coupled path-planning and sensor configuration (CSCP) method and we show that significant benefits are achieved by introducing this coupling. We consider path-planning to minimize exposure to an unknown, spatially varying, and temporally static scalar field called the *threat field.* Sensors take noisy measurements of the threat field. We use Gaussian Process regression to estimate the threat field from measurements. Sensor configuration is performed by maximizing a so-called *task-driven information gain* metric or its surrogate. Cooper and Cowlagi (2019) first reported such a coupled sensor placement and path-planning approach but considered pointwise sensors, only. We consider sensors with configurable fields of view (FoV). We demonstrate that the proposed CSCP method outperforms decoupled information-driven sensor configuration in terms of the number of measurements required to find a comparable near-optimal plan.

Statement of Contributions: We provide a method to configure sensors in the case where noisy observations are taken over a finite area rather than pointwise, without lossy aggregation of the observations. To the best of our knowledge, no comparable method is available in the literature. Furthermore, we consider two configurations for each sensor, namely, its location and FoV (which may be altitude-dependent). We explicitly model sensor noise dependence on FoV to emulate commonly used EO/IR sensors. In so doing, we quantify the "exploration v/s exploitation" trade-off, where "exploration" is achieved through large FoV (noisy measurements over large areas) and "exploitation" is achieved through small FoV (accurate measurements over small areas). We demonstrate that the proposed method achieves near-optimal paths using fewer observations by *orders of magnitude* compared to traditional sensor placement methods.

Preliminary results of the work presented in this paper were previously reported in conferences (St. Laurent and Cowlagi, 2021a,b,c). The crucial innovation in this paper is a new metric called *self-adaptive mutual information,* which we show is submodular.

2 Problem Formulation

We denote by $\mathbb{R}$ and $\mathbb{N}$ the sets of real and natural numbers, respectively, and by $\{N\}$ the set $\{1, 2, \ldots, N\}$ for any $N \in \mathbb{N}$. For any $\boldsymbol{a} \in \mathbb{R}^N$, $\boldsymbol{a}[i]$ is its i^{th} element, $diag(\boldsymbol{a})$ is the diagonal matrix with the elements of $\boldsymbol{a}$ on the principal diagonal, and $\boldsymbol{a}^{\circ(-1)}$ denotes the vector with reciprocal elements of $\boldsymbol{a}$. For any matrix $A \in \mathbb{R}^{M\times N}$, $A[i, j]$ is the element in the i^{th} row and j^{th} column. For $A \in \mathbb{R}^{N\times N}$ and for the indicator vector $\boldsymbol{a} \in \{0, 1\}^N$, $diag(A)$ is the diagonal vector and $A[\boldsymbol{a}]$ is the submatrix of rows and columns indicated by $\boldsymbol{a}$. Similarly, $A[i, \boldsymbol{a}]$ denotes elements in the i^{th} row and columns indicated by $\boldsymbol{a}$. $\boldsymbol{I}_N$ is the identity matrix of size N. For $\mu, \sigma \in \mathbb{R}, \mathcal{N}(\mu, \sigma^2)$ is the normal distribution with mean μ and variance σ^2. $\mathbb{E}$ is the expectation operator.

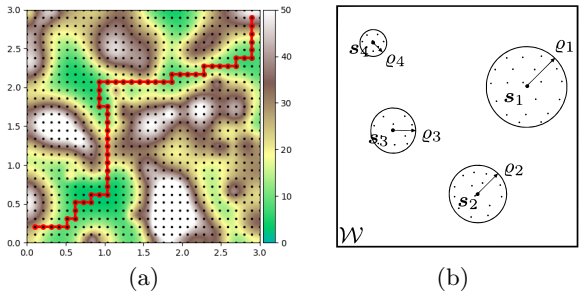

Fig. 1. (a) Example of a threat field and optimal path. (b) Illustrations of sensor FoVs and observations (red dots).

2.1 Path-Planning and Sensor Configuration

The agent operates in a compact square planar region called the *workspace* $\mathcal{W} \subset \mathbb{R}^2$. Consider a uniformly-spaced square grid of points $i = 1, 2, \ldots, N_{\text{g}}$ and a graph $\mathcal{G} = (V, E)$ whose vertices $V = \{N_{\text{g}}\}$ are uniquely associated with these grid points. The set of edges E consists of pairs of geometrically adjacent grid points. We label the vertices the same as grid points. We denote by $\boldsymbol{p}_i = (p_{ix}, p_{iy})$ the coordinates of the i^{th} grid point and by Δp the distance between adjacent grid points.

A *threat field* $c : \mathcal{W} \to \mathbb{R}_{>0}$ is a strictly positive scalar field. A *path* $\boldsymbol{\pi} = (\boldsymbol{\pi}[0], \boldsymbol{\pi}[1], \ldots, \boldsymbol{\pi}[\lambda])$ between prespecified initial and goal vertices $i_{\text{start}}, i_{\text{goal}} \in V$ is a finite sequence, without repetition, of successively adjacent vertices such that $\boldsymbol{\pi}[0] = i_{\text{start}}$ and $\boldsymbol{\pi}[\lambda] = i_{\text{goal}}$ for some $\lambda \in \mathbb{N}$. When the meaning is clear from the context, we also denote by $\boldsymbol{\pi}$ the unordered set of vertices in a path. A *path incidence vector* $\boldsymbol{v}_{\boldsymbol{\pi}} \in \{0, 1\}^{N_{\text{g}}}$ has $\boldsymbol{v}_{\boldsymbol{\pi}}[i] = 1$ if $i = \boldsymbol{\pi}[j]$ for $j \in \{\lambda\}\backslash 0$ and $\boldsymbol{v}_{\boldsymbol{\pi}}[i] = 0$ otherwise. The *cost* of a path $\boldsymbol{\pi}$ is the total threat exposure calculated as $\mathcal{J}(\boldsymbol{\pi}) := \Delta p \sum_{j=1}^{\lambda} c(\boldsymbol{p}_{\boldsymbol{\pi}[j]})$. The main problem of interest is to find a path $\boldsymbol{\pi}^*$ of minimum cost.

The threat field is unknown, but it can be observed by a network of $N_{\text{s}} \in \mathbb{N}$ sensors. Each sensor measures the threat in a circular FoV as shown in Fig. 1(b). The center $\boldsymbol{s}_k \in \mathcal{W}$ and radius $\varrho_k \in \mathbb{R}_{>0}$ of this circular FoV are parameters that we can choose for each $k \in \{N_{\text{s}}\}$. Maximum and minimum FoV radii are prespecified as $\varrho^{\max}$ and $\varrho^{\min}$, respectively. The set of all sensor parameters is called a *configuration*, denoted by $\mathcal{C} = \{\boldsymbol{s}_1, \varrho_1, \boldsymbol{s}_2, \ldots, \varrho_{N_{\text{s}}}\}$.

We denote by $\mathcal{F}_k \subset V$ the set of vertices with grid points

lying within the FoV of the k^{th} sensor. A sensor *cover incidence vector* $\boldsymbol{\nu}_k$ is defined such that $\boldsymbol{\nu}_k[i] = 1$ if $i \in \mathcal{F}_k$ and $\boldsymbol{\nu}_k[i] = 0$ otherwise. The union of all FoVs is $\mathcal{F} := \cup_{k\in\{N_\text{s}\}}\mathcal{F}_k$ and the cover incidence of all sensors in the network is $\boldsymbol{\nu} := (\boldsymbol{\nu}_1 \vee \boldsymbol{\nu}_2 \vee ... \vee \boldsymbol{\nu}_{N_\text{s}})$. Within its FoV, the k^{th} sensor takes $M_k \in \mathbb{N}$ pointwise and noisy measurements $z_{km} = c(\boldsymbol{x}_{km}) + \eta_{km}$, for $m \in \{M_k\}$.

Assumption 1 *The measurement error η_{km} is independent and identically distributed (i.i.d) with $\eta_{km} \sim \mathcal{N}(0, \sigma_k^2)$, where σ_k^2 is a monotonically increasing function of the FoV radius ϱ_k.*

We denote by $\boldsymbol{z} = [z_{11} \dots z_{1M_1} \dots z_{N_\text{s}M_{N_\text{s}}}]^\intercal$ the measurements made by the collection of all sensors. We may use these measurements to construct a stochastic estimate of the threat field, and in turn, to find an optimal path that minimizes the *expected* cost. Furthermore, we would like to reduce the uncertainty in the estimated path cost, i.e., require that the variance of the path cost be less than a certain prespecified threshold. To do so, we must collect a sufficient number of measurements, which may be achieved by repeatedly changing the sensor configuration over multiple iterations. Conceptually, at each iteration $\ell = 0, 1, \dots$, a sensor configuration $\mathcal{C}_\ell^*$ is chosen, the threat field estimate is updated using the new measurements, and an optimal path is computed. The main problem of interest is formulated as follows.

Problem 2 *Let $\varepsilon > 0$ be a prespecified termination threshold. Over finite iterations $\ell = 0, 1, \dots, L$, find sensor configurations $\mathcal{C}_\ell^*$ and a path $\boldsymbol{\pi}^*$ of minimum expected cost $\mathbb{E}[\mathcal{J}(\boldsymbol{\pi}^*)]$ such that $\mathbb{E}[(\mathcal{J}(\boldsymbol{\pi}^*) - \mathbb{E}[\mathcal{J}(\boldsymbol{\pi}^*)])^2] \leqslant \varepsilon$.*

2.2 Threat Field Modeling

We employ a Gaussian Process (GP) model to estimate the threat field from measurements. GP regression (GPR) is a supervised machine learning method that is known to be a universal approximator given sufficient training data and a stationary kernel (Micchelli et al., 2006). For a set of training points $\boldsymbol{X} = \{\boldsymbol{x}_1, \dots, \boldsymbol{x}_M\}$, we define a matrix kernel $\boldsymbol{K} = \kappa(\boldsymbol{X}, \boldsymbol{X}) \in \mathbb{R}^{M\times M}$ using anisotropic scaled radial basis functions as $\boldsymbol{K}[i,j] = \theta_\text{c} \exp(-\frac{1}{2}(\boldsymbol{x}_i - \boldsymbol{x}_j)^\intercal \boldsymbol{\Theta}_r^{-2}(\boldsymbol{x}_i - \boldsymbol{x}_j))$, where θ_c is a scalar coefficient and $\boldsymbol{\Theta}_r$ is a matrix of input dimension hyperparameters, both to be chosen. On the diagonal of $\boldsymbol{\Theta}_r$ are the dimension-specific length scales whereas the off-diagonal are the length scale correlations between dimensions. For a set of training points $\boldsymbol{X}$ and a set of test points $\boldsymbol{X}_*$, we can define kernels $\boldsymbol{K}_* = \kappa(\boldsymbol{X}, \boldsymbol{X}_*)$ and $\boldsymbol{K}_{**} = \kappa(\boldsymbol{X}_*, \boldsymbol{X}_*)$ similarly.

In the present context GPR is applied as follows. At each iteration $\ell = 0, 1, \dots$, the algorithm maintains and recursively updates a pointwise estimated mean threat $\boldsymbol{f}_\ell \in \mathbb{R}^{N_\text{g}}$ and an estimation error covariance matrix $\boldsymbol{P}_\ell \in \mathbb{R}^{N_\text{g}\times N_\text{g}}$. The expected cost of any path $\boldsymbol{\pi}$ is $\mathbb{E}_\ell[\mathcal{J}(\boldsymbol{\pi})] = \Delta p\ \boldsymbol{v}_{\boldsymbol{\pi}}^\intercal \boldsymbol{f}_\ell$. The path cost variance is $\text{Var}_\ell(\boldsymbol{\pi}) := \mathbb{E}_\ell[(\mathcal{J}(\boldsymbol{\pi}) - \mathbb{E}_\ell[\mathcal{J}(\boldsymbol{\pi})])^2] = (\Delta p)^2 \boldsymbol{v}_{\boldsymbol{\pi}}^\intercal \boldsymbol{P}_\ell \boldsymbol{v}_{\boldsymbol{\pi}}$.

At the ℓ^{th} iteration, the sensor configuration $\mathcal{C}_\ell^*$ provides $N_\text{m} := \sum_{k=1}^{N_\text{s}} M_k$ measurements. The training point set $\boldsymbol{X} = \{\boldsymbol{x}_{11}, \dots, \boldsymbol{x}_{km}, \dots, \boldsymbol{x}_{N_\text{s}M_{N_\text{s}}}\}$ represents the locations where these measurements are taken and the test points $\boldsymbol{X}_*$ are the grid points. The associated measurement error vector is $\boldsymbol{\sigma} := (\sigma_1^2, \dots, \sigma_{km}^2, \dots, \sigma_{N_\text{s}M_{N_\text{s}}}^2)$. The hyperparameters $\boldsymbol{\theta} = (\theta_\text{c}, \boldsymbol{\Theta}_r)$ are calculated by minimizing the negative marginal log-likelihood,

$$\log p(\boldsymbol{z}|\boldsymbol{X}, \boldsymbol{\theta}) = -\tfrac{1}{2}\left(\boldsymbol{z}^\intercal \boldsymbol{K}_z^{-1}\boldsymbol{z} + \log|\boldsymbol{K}_z| + N_\text{m}\log 2\pi\right),$$

where $\boldsymbol{K}_z := \boldsymbol{K} + \text{diag}(\boldsymbol{\sigma})$. From the joint distribution of observations $\boldsymbol{z}$ and the mean threat $\boldsymbol{f}$, we can obtain the threat field estimate and error covariance matrix as:

$$\boldsymbol{f}_\ell = \boldsymbol{K}_*^\intercal \boldsymbol{K}_z^{-1}\boldsymbol{z}, \qquad \boldsymbol{P}_\ell = \boldsymbol{K}_{**} - \boldsymbol{K}_*^\intercal \boldsymbol{K}_z^{-1}\boldsymbol{K}_*. \tag{1}$$

The aforementioned kernel is a local approximation kernel because it has the property of regressing to the function mean $\kappa(\boldsymbol{x}, \boldsymbol{x}') \to 0$ as $\|\boldsymbol{x} - \boldsymbol{x}'\| \to \infty$. This property enforces high uncertainty and an "optimistic" mean threat estimate (i.e., approaching zero) in regions far away from training data locations.

3 Coupled Sensor Configuration and Planning

We propose an iterative solution to Problem 2, as described in Fig. 2 and illustrated in Fig. 3. At each iteration, a sensor configuration is found and new measurements are obtained. Next, the threat field estimate is updated using the new measurements. Finally, the path plan is updated according to the updated threat field estimate and the iterations repeat. The iterations stop when the planned path cost variance is reduced below a user-specified threshold ε. We call this method *coupled* sensor configuration and path-planning (CSCP) because the sensor configuration is dependent on the planned path. Note that this coupling is a property of the proposed *solution,* not of Problem 2 itself. The premise of this work is that, whereas traditional solutions to Problem 2 decouple the sensor configuration and planning subproblems, the proposed coupled method can reach near-optimal solutions more efficiently.

The algorithm initializes with an optimistic field representation and high uncertainty by setting $\boldsymbol{f}_0 = \boldsymbol{0}$ and $\boldsymbol{P}_0 = \chi \boldsymbol{I}$, where χ is a large arbitrary number. At any iteration $\ell \in \mathbb{N}$, including initialization with $\ell = 0$, an optimal path $\boldsymbol{\pi}_\ell^*$ of minimum expected cost is found using Dijkstra's algorithm.

3.1 Sensor Configuration

The proposed sensor configuration method relies on the notion of a *region of interest* (ROI) "nearby" the planned path $\boldsymbol{\pi}_\ell^*$. The ROI $\mathcal{R}_\ell \in V$ at any iteration ℓ, is a set of vertices in graph $\mathcal{G}$, which indicates a subregion of the workspace that the FoV of each sensor must overlap.

Initialize: $\ell := 0$, $\boldsymbol{f}_0 := \boldsymbol{0}, \boldsymbol{P}_0 := \chi \boldsymbol{I}$
$\boldsymbol{\pi}_0^* := \arg\min \overline{\mathcal{J}}_0(\boldsymbol{\pi})$
while $\mathrm{Var}_\ell(\boldsymbol{\pi}_\ell^*) > \varepsilon$ **do**
  Find optimal sensor configuration $\mathcal{C}_\ell^*$
  Record new measurements $\boldsymbol{z}$
  Increment iteration counter $\ell := \ell + 1$
  Update $\boldsymbol{f}_\ell$ and $\boldsymbol{P}_\ell$ per §2.2
  Find $\boldsymbol{\pi}_\ell^* := \arg\min \overline{\mathcal{J}}_\ell(\boldsymbol{\pi})$
end while

Fig. 2. Proposed iterative CSCP solution for Problem 2.

Similar to the path incidence vector, we define an ROI incidence vector $\boldsymbol{r} \in \{0,1\}^{N_\mathrm{g}}$.

We consider a *posterior sampling-based* ROI by taking a fixed number of random samples $N_\mathrm{a} \in \mathbb{N}$, from the path incidence vector space $\{0,1\}^{N_\mathrm{g}}$. For each sample $\boldsymbol{g}_i$ for $i \in \{N_\mathrm{a}\}.$, we construct a posterior estimated mean threat as $\hat{\boldsymbol{f}}^{(i)} := \boldsymbol{f}_\ell + \boldsymbol{Q}\boldsymbol{g}_i$, where $\boldsymbol{Q}$ is such that $\boldsymbol{Q}\boldsymbol{Q}^\intercal = \boldsymbol{P}_\ell$. Next, we find a path $\boldsymbol{\pi}_{\boldsymbol{a}}^{(i)} = \Delta p\ \boldsymbol{v}_{\boldsymbol{\pi}}^\intercal\ \hat{\boldsymbol{f}}_\ell^{(i)}$, where $\boldsymbol{\pi}_{\boldsymbol{a}}^{(i)}$ is the i^{th} set of alternate estimated optimal path-plan vertices, and then define the ROI as $\mathcal{R}_\ell = \boldsymbol{\pi}_a := \boldsymbol{\pi}_\ell^* \cup \{\cup_{i \in N_a} \boldsymbol{\pi}_a^{(i)}\}$. Informally, this ROI covers a set of alternative paths that may minimize the expected cost in the next iteration. If the grid resolution is high, i.e., N_g is large, then the posterior sampling-based ROI can have a large cardinality. To mitigate this problem, we may define a union set $\mathcal{I} \subset V$ of vertices covered by the FoVs of sensors placed in previous iterations. The ROI definition is updated as $\mathcal{R}_\ell = \mathcal{R}_\ell \cap \{V \backslash \mathcal{I}\}$, which restricts the ROI to "new" vertices that have not previously been observed by sensors.

We formalize the *coupling* of sensor configuration and path-planning by introducing the *task-driven information gain* (TDIG) metric τ, which depends on the ROI $\mathcal{R}_\ell$ as well as the sensor configuration $\mathcal{C}_\ell$. The TDIG quantifies reduction in threat field estimation uncertainty in the ROI. We may find a sensor configuration $\mathcal{C}_\ell^*$ by maximizing $\tau(\mathcal{C}_\ell, \mathcal{R}_\ell)$, where

$$\tau(\mathcal{C}_\ell, \mathcal{R}_\ell) := \boldsymbol{r}_\ell^\intercal (\boldsymbol{P}_\ell - \boldsymbol{P}_{\ell+1}) \boldsymbol{r}_\ell. \tag{2}$$

The dependence on sensor configuration $\mathcal{C}_\ell$ is through the posterior covariance $\boldsymbol{P}_{\ell+1}$. To estimate $\boldsymbol{P}_{\ell+1}$, we assume that spatial correlations remain fixed or assume spatial independence. The assumption is justified because spatial points within a local neighborhood do not contribute to the maximization of TDIG. We approximate $\hat{\boldsymbol{P}}_{\ell+1} = diag(\boldsymbol{q})$ where $\boldsymbol{q}$ is such that $\boldsymbol{q}^{\circ(-1)} = diag(\boldsymbol{P}_\ell)^{-1} + \sum_{k=1}^{N_s} (\boldsymbol{\nu}_k / \sigma_k^2)$. Recall from §2.1 that $\boldsymbol{\nu}_k$ and σ_k^2 are both sensor configuration parameters. Sensor configuration is achieved by numerical maximization of the TDIG over $\mathcal{C}_\ell = \{\boldsymbol{s}_1, \varrho_1, \boldsymbol{s}_2, \ldots, \varrho_{N_\mathrm{s}}\}$. However, the TDIG does not have convenient properties to establish guarantees on the convergence and/or optimality of numerical optimization. To resolve this problem, we define a surrogate function and show it is submodular.

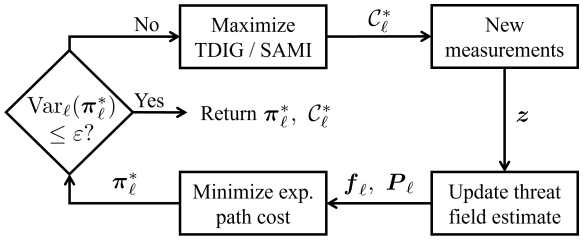

Fig. 3. Illustration of the CSCP iterative method.

3.2 Self-Adaptive Mutual Information

The proposed TDIG surrogate, which we call *self-adaptive mutual information (SAMI)*, is defined as the difference between a reward term and a penalty term as discussed next. The TDIG and SAMI are calculated at each iteration $\ell \in \mathbb{N}$. For clarity in notation we drop the iteration number subscript for various quantities in this subsection, e.g., we write $\boldsymbol{P}$ instead of $\boldsymbol{P}_\ell$. We denote $\boldsymbol{P}[\mathcal{R}, \mathcal{R}] := \boldsymbol{P}[i,j]\ \forall\ i,j \in \mathcal{R}$ and by $\boldsymbol{P}[\mathcal{R}, j] := \boldsymbol{P}[i,j]\ \forall\ i \in \mathcal{R}$ and for any $j \in V$. An ROI $\mathcal{R}$ excluding a point $i \in V$ is denoted $\mathcal{R}_{\backslash i}$. We calculate the entropy of a vertex $i \in V$ as $h(i) := \frac{1}{2}\ln(2\pi e \boldsymbol{P}[i,i])$.

We rely on the conditional entropy of some point i given the set $\mathcal{R}$ which involves matrix inversion. To ease the computational burden, we note that the computation of conditional entropy $h(i|\mathcal{R}_{\backslash i})$ can partially be computed as a one-time batch operation and partially computed in parallel for efficiency. For any $i \notin \mathcal{R}$, we may perform the following vectorized conditional entropy $\boldsymbol{h}(\cdot|\mathcal{R}_{\backslash i})$ calculation which is a one-time batch operation:

$$\boldsymbol{h}(\cdot|\mathcal{R}_{\backslash i}) = \tfrac{1}{2}\ln((2\pi e) diag(\boldsymbol{P} - \boldsymbol{P}[\cdot, \mathcal{R}]\boldsymbol{P}[\mathcal{R}, \mathcal{R}]^{-1}\boldsymbol{P}[\mathcal{R}, \cdot])).$$

When $i \in \mathcal{R}$, we can compute the following conditional entropy equation for any $i \in V$ in parallel batches:

$$\begin{aligned} h(i|\mathcal{R}_{\backslash i}) = \tfrac{1}{2}\ln((2\pi e)(&\boldsymbol{P}[i,i] \\ &- \boldsymbol{P}[\mathcal{R}_{\backslash i}, i]^\intercal \boldsymbol{P}[\mathcal{R}_{\backslash i}, \mathcal{R}_{\backslash i}]^{-1} \boldsymbol{P}[\mathcal{R}_{\backslash i}, i])). \end{aligned} \tag{3}$$

The mutual information I for the ROI $\mathcal{R}$ at vertex $i \in V$ is then calculated as $I(\mathcal{R}_{\backslash i}; i) := h(i) - h(i|\mathcal{R}_{\backslash i})$. Finally, we define the SAMI reward γ for $i \in V$ with $\mathcal{R}$ and its complement $\mathcal{R}^c := V \backslash \mathcal{R}$, as:

$$\gamma(i) := (1-\alpha) I(\mathcal{R}_{\backslash i}; i) + \alpha I(\mathcal{R}_{\backslash i}^c; i). \tag{4}$$

The adaptation parameter α is formulated as the relationship between the average mutual information between the workspace vertices and the ROI $\bar{I}(\mathcal{R}_{\backslash i}; i)$ and the average mutual information between the workspace vertices and the ROI complement $\bar{I}(\mathcal{R}_{\backslash i}^c; i)$:

$$\alpha := \bar{I}(\mathcal{R}_{\backslash i}; i) / \Big(\bar{I}(\mathcal{R}_{\backslash i}; i) + \bar{I}(\mathcal{R}_{\backslash i}^c; i) \Big). \tag{5}$$

Informally, α compares the estimation error within the current ROI compared to the overall estimation error. A small α indicates that the iterative solution is nearing convergence and sensor configuration should exploit the current ROI. We define the vector of rewards for each $i \in V$ as $\boldsymbol{\gamma} := (\gamma(1), \gamma(2), \ldots, \gamma(N_\mathrm{g})) \in \mathbb{R}^{N_\mathrm{g} \times 1}$. The information gain for a configuration $\mathcal{C}_\ell$ is defined as

$$\Gamma(\mathcal{C}_\ell) := \boldsymbol{\nu}^\intercal \boldsymbol{\gamma}. \tag{6}$$

Next, we define a SAMI penalty function as the entropy of the measurement noise for a sensor configuration:

$$\Upsilon(\mathcal{C}_\ell) := -\tfrac{1}{2} \sum_{i \in \mathcal{F}} \left(\tfrac{1}{2} \ln(2\pi e) - \ln \textstyle\sum_{k \in N_\mathrm{s}} \frac{\nu_k}{\sigma_k^2}[\mathcal{F}] \right). \tag{7}$$

The summation over $\mathcal{F}$ ensures that nonzero elements are removed prior to computing the elementwise entropy. Finally, the SAMI is written as the sum of the information gain and penalty: $S(\mathcal{C}_\ell) = \Gamma(\mathcal{C}_\ell) + \Upsilon(\mathcal{C}_\ell)$.

3.3 Optimization of Sensor Configuration

To find an optimal sensor configuration we maximize either the TDIG or SAMI objective functions subject to spatial constraints, e.g., workspace limits on sensor placement and FoV radius bounds. The decision variables for optimization are $\{\boldsymbol{s}_1, \varrho_1, \boldsymbol{s}_2, \ldots, \varrho_{N_\mathrm{s}}\}$, i.e., a total of $3N_\mathrm{s}$ scalar variables in a 2D workspace. If the number of sensors N_s is moderate or high, this optimization problem is high-dimensional.

In previous works we discussed different heuristic methods to mitigate the high dimensionality of the problem. We first applied an evolutionary computation method for global optimization (St. Laurent and Cowlagi, 2021a), and later proposed a depth-first (DF) method (St. Laurent and Cowlagi, 2021c) and a breadth-first (BF) (St. Laurent and Cowlagi, 2021b) method for TDIG maximization. Briefly, the DF method makes use of the set of identified vertices in the acting agent's workspace. The approach first targets unidentified regions and prunes those which have been identified in a prior iteration. Once the region of interest is fully identified, sensor configurations are batched until either the estimated TDIG is reduced below ε. The BF approach weighs the least frequent regions of sampled path plans proportional to the inverse frequency of visits. These methods pursue *joint* optimization, i.e., finding all sensor configurations simultaneously.

Another method to mitigate the high dimensionality is *sequential* optimization, i.e., finding the optimal configuration for one sensor at a time. The main result of this paper is that *the proposed SAMI surrogate function is submodular* (see Prop. 7). Therefore, sequential optimization provides a near-optimal solution (see Prop. 8).

3.4 Theoretical Analysis

Proposition 3 *The CSCP algorithm terminates in a finite number of iterations $L \in \mathbb{N}$ for $N_\mathrm{s} > 0$.*

Proof. At any iteration $\ell \in \mathbb{N}$, by Line 4 of Algorithm 2, each sensor is configured such that its FoV intersects with the ROI $\mathcal{R}_\ell$. By any definition of the ROI in §3.1, $\boldsymbol{\pi}_\ell^* \subset \mathcal{R}$ and therefore, at least one measurement is taken on at least one vertex of path $\boldsymbol{\pi}_\ell^*$. Therefore, $\mathrm{Var}_{\ell+1}(\boldsymbol{\pi}_\ell^*) < \mathrm{Var}_\ell(\boldsymbol{\pi}_\ell^*)$, i.e., the path cost variance monotonically decreases at each iteration. Because the path cost variance is also lower bounded by zero, $\inf_\ell\{\mathrm{Var}_\ell(\boldsymbol{\pi}_\ell^*)\} = 0$, and it follows that there exists a finite $L \in \mathbb{N}$ such that $\mathrm{Var}_L(\boldsymbol{\pi}_L^*) < \varepsilon$, which is when the algorithm terminates. □

Corollary 4 *The CSCP algorithm solves Problem 2.*

Proof. Upon termination the CSCP algorithm returns a path $\boldsymbol{\pi}_L^*$ of minimum expected cost that satisfies $\mathrm{Var}_L(\boldsymbol{\pi}_L^*) = \mathbb{E}_L[(\mathcal{J}(\boldsymbol{\pi}_L^*) - \mathbb{E}[\mathcal{J}(\boldsymbol{\pi}_L^*)])^2] < \varepsilon$ per Line 3, thereby solving Problem 2. □

Corollary 5 *The path $\boldsymbol{\pi}_L^*$ is near-optimal in the that*

$$\mathbb{P}\left[|\mathcal{J}(\boldsymbol{\pi}_L^*) - \mathbb{E}_L[\mathcal{J}(\boldsymbol{\pi}_L^*)]| \leqslant 3\sqrt{\varepsilon}\right] \geqslant 0.9973.$$

Proof. Due to the GPR-based estimation and linearity of $\mathcal{J}$, the path cost is normally distributed, and the result follows immediately from the standard normal table. □

Due to the linearity of $\mathcal{J}$ and the fact that the path cost is normally distributed, a stronger statement about the near-optimality of $\boldsymbol{\pi}_L^*$ follows immediately from the standard normal table.

Corollary 6 *The path $\boldsymbol{\pi}_L^*$ is near-optimal in that*

$$\mathbb{P}\left[|\mathbb{E}_L[\mathcal{J}(\boldsymbol{\pi}_L^*)] - \mathcal{J}(\boldsymbol{\pi}^*)| \leqslant 3\sqrt{\varepsilon}\right] \geqslant 0.9973,$$

where $\boldsymbol{\pi}^*$ is the true optimal path.

Proposition 7 *The SAMI, S, is submodular.*

Proof. By (6), the gain Γ is proportional to the number of vertices lying within the union $\mathcal{F} := \cup_{k \in \{N_\mathrm{s}\}} \mathcal{F}_k$ of the sensor FoVs. Therefore, Γ is a weighted coverage function and consequently submodular.

For the SAMI penalty Υ, first consider two sensors $i, j \in N_\mathrm{s}$ with disjoint FoVs, i.e., $\mathcal{F}_i \cap \mathcal{F}_j = \varnothing$. Then, per (7), $\Upsilon(\mathcal{F}_i) + \Upsilon(\mathcal{F}_j) = \Upsilon(\mathcal{F}_i \cup \mathcal{F}_j)$. Next, consider sensors $i, j \in N_\mathrm{s}$ with overlapping FoVs and note that $\Upsilon(\mathcal{F}_i) + \Upsilon(\mathcal{F}_j) = \Upsilon(\{\mathcal{F}_i \cup \mathcal{F}_j\} \backslash \{\mathcal{F}_i \cap \mathcal{F}_j\}) + 2\Upsilon(\mathcal{F}_i \cap \mathcal{F}_j)$. Let $\Upsilon(\mathcal{F}_m) \leqslant \Upsilon(\mathcal{F}_i \cap \mathcal{F}_j)$ be the reduced cost function due to the weighted update from intersecting sensor covers in (7). Therefore $\Upsilon(\mathcal{F}_i \cup \mathcal{F}_j) + \Upsilon(\mathcal{F}_i \cap \mathcal{F}_j) = \Upsilon(\{\mathcal{F}_i \cup \mathcal{F}_j\} \backslash \{\mathcal{F}_i \cap \mathcal{F}_j\}) + \Upsilon(\mathcal{F}_i \cap \mathcal{F}_j) + \Upsilon(\mathcal{F}_m)$. Therefore, Υ exhibits the following condition for submodularity:

$$\Upsilon(\mathcal{F}_i) + \Upsilon(\mathcal{F}_j) \geqslant \Upsilon(\mathcal{F}_i \cup \mathcal{F}_j) + \Upsilon(\mathcal{F}_i \cap \mathcal{F}_j).$$

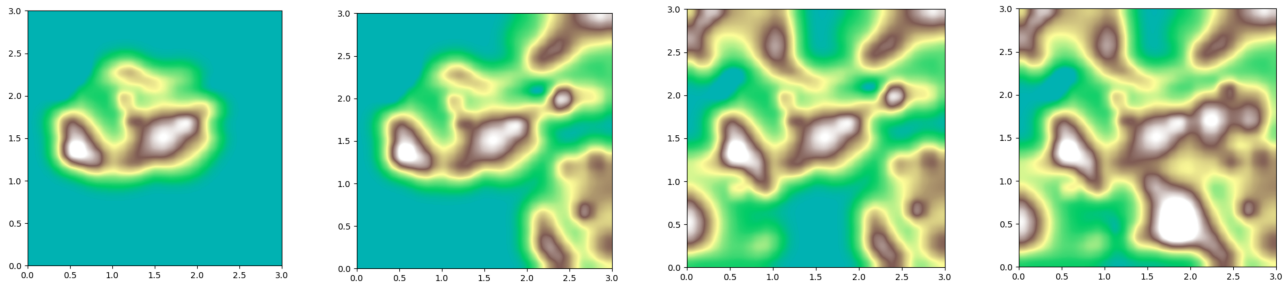

(a) $\ell = 1$ (b) $\ell = 3$ (c) $\ell = 5$ (d) $\ell = 7$

Fig. 4. Threat estimate (mean values) at various intermediate iterations.

Because the SAMI S is the sum of Γ and Υ that are each submodular, it follows that S is submodular. $\square$

3.4.1 Computational Complexity

Jointly optimizing all sensors and their parameters with the TDIG metric has a time complexity of $\mathcal{O}(N_g^2)$ during each step in the optimization process. By comparison, the SAMI metric has a time complexity of $\mathcal{O}(N_g)$ for computing the information gain and penalty functions.

The TDIG requires joint optimization of $N_d N_s$ parameters, where $N_d = 3$ is the dimensionality of each sensor configuration. This can make finding a near-globally-optimal solution difficult. By comparison, due to its submodularity SAMI optimization can be performed sequentially while achieving near-optimal results per the following result, for which the reader interested is referred to Krause et al. (2008b) for a proof.

Proposition 8 *Sequential optimization of submodular functions in general, and S in particular, is guaranteed to be at least $(1 - 1/e)$ the optimal value.*

4 Results & Discussion

In this section, we provide computational results of implementing the CSCP iterative solution, specifically: (1) an illustrative example, (2) results showing the benefits of CSCP via numerical optimization of the TDIG, over traditional sensor configuration methods and (3) results showing that the SAMI is an appropriate and efficient surrogate for the TDIG.

Illustrative Example: Consider the problem of finding a minimum threat path, where the true threat field (unknown to the planner) is as shown in Fig. 1(a). Fig. 4 shows the evolution of the estimate mean $\boldsymbol{f}$ over several CSCP iterations with a termination threshold of $\varepsilon = 1.0$. Fig. 5 (a) shows the threat estimate mean at the last iteration, with the resultant path (indicated in yellow) overlaid the true optimal path (in red). Fig. 5 (b) indicates the estimation error covariance $\boldsymbol{P}$ at the final iteration.

The candidate optimal path (yellow) is spatially close to the true optimal path, and is found to be 0.02% suboptimal. TDIG BF optimization yielded a solution in 9

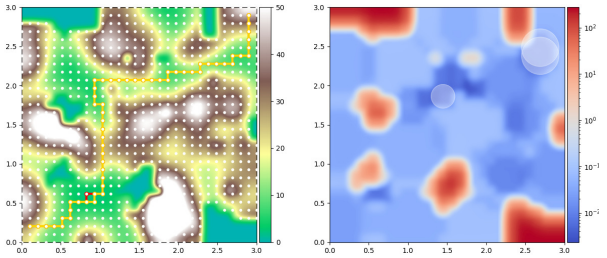

(a) $\boldsymbol{f}_L$ and $\boldsymbol{\pi}_L^*$ (b) $\boldsymbol{P}_L$ and $\mathcal{C}_L^*$

Fig. 5. Final threat estimate and error covariance along with estimated path-plan and sensor FoVs.

iterations of CSCP. If the termination threshold is made more stringent to $\varepsilon = 0.01$, the TDIG BF optimization took 15 iterations. As we show next, *any* of the three task-driven strategies outperforms existing traditional approaches by orders of magnitude.

CSCP vs. Traditional Sensor Configuration: To numerically assess the performance of the CSCP approach, we established a study that randomly generated, for each experiment, 100 threat fields of the form $c(\boldsymbol{x}) = \sum_{n=1}^{N_p} b_n \phi_n(\boldsymbol{x})$. Here N_p represents the number of radial basis functions ϕ_n the support of which together covers the workspace $\mathcal{W}$, which is set to be a square region. The threat intensity $b_n > 0$ is a randomly chosen coefficient. We fixed the number of threat parameters to be $N_p = 50$. The sensor FoV radii were constrained to $\varrho_{\min} = 0.01$ km and $\varrho_{\max} = 0.5$ km, respectively, and the the sensor noise was modeled as $\sigma_k^2 = \frac{1}{2}\log(1+\exp^{\pi\varrho_k^2}) - 0.1505$, which is monotonically increasing for $\varrho_k \geqslant 0$ per Assumption 1. The termination threshold was fixed to $\varepsilon = 0.01$.

We performed numerical experiments (100 trials each) with all combinations of the following problem data: (A) Area of workspace $\mathcal{W} = 9$ km^2 or 25 km^2, (B) Number of sensors $N_s = 3, 5$, or 9, and (C) Number of grid points $N_g = 225$ or 400. In each trial of each experiment, we compared the execution time of the following TDIG maximization methods in CSCP: (1) direct TDIG maximization along the path-plan, (2) the BF method with three different values of the sampling parameter $N_a = 25, 50$, or 100, and (3) the DF method. For com-

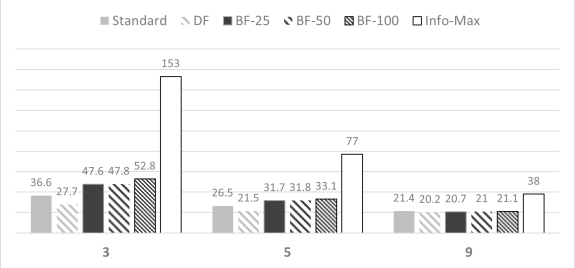

Fig. 6. Comparison of CSCP methods vs. traditional Info–Max method to find paths with cost variance below ε. Sample result for $|\mathcal{W}| = 25$ km^2, and $N_g = 400$.

parison, we also developed a traditional (non-CSCP) method based on maximization of mutual information (Info-Max), as described by Krause et al. (2008a). Because the traditional MI-maximization literature addresses only sensor *placement,* i.e., neglects sensor FoV, for the Info-Max computations we fixed the sensor FoV to 0.1 km. Per Cor. 6, by either approach the resultant path costs are, with high probability, within $3\sqrt{\varepsilon}$ of the minimum.

Figure 6 shows a sample results of these experiments, specifically, the average number of iterations required for the path cost variance to reduce below ε. Notice that the CSCP methods easily outperform the traditional Info-Max approach in all experiments. The advantages of the CSCP methods are particularly evident when the number of sensors is small. For the case shown in Fig. 6, the Info-Max method needs more than thrice as many iterations as the standard CSCP method to converge to a path of similar cost. Because $N_s = 3$, the number of *measurements* required for the Info-Max method is an *order of magnitude* larger than CSCP.

TDIG vs. SAMI Maximization: We demonstrate next that maximizing SAMI is computationally more efficient compared to TDIG. To this end, we performed a series of numerical experiments of 100 trials each with grid resolutions $N_g \in \{11^2, 21^2, 31^2, 41^2, 51^2\}$ and $N_s \in \{1, 3, 5, 7, 9\}$ sensors. The termination threshold was fixed to $\varepsilon = 0.1$. In each experiment, the ratio of the maximum sensor FoV $\varrho_{\max}$ to the workspace area $\mathcal{W}$ was varied. The sampling-based ROI parameter was fixed at $N_a = 10$. The sensor FoVs were constrained to $\varrho^{min} = 0.05$ km and $\varrho^{max} = 0.5$ km.

Figure 7 shows that, in these experiments, the SAMI maximization enables the overall CSCP approach to converge in comparable or fewer iterations compared to TDIG maximization. Figure 8 shows the comparison of computation times per iteration. Note that sequential SAMI maximization is always faster and it is over an order of magnitude faster when N_s is large.

Finally, we conducted a numerical experiment to study the effects of the SAMI parameter α. We studied sequential SAMI optimization for fixed values of the parameter $\alpha \in \{0, .25, .5, .75, 1\}$ against the adaptive value as

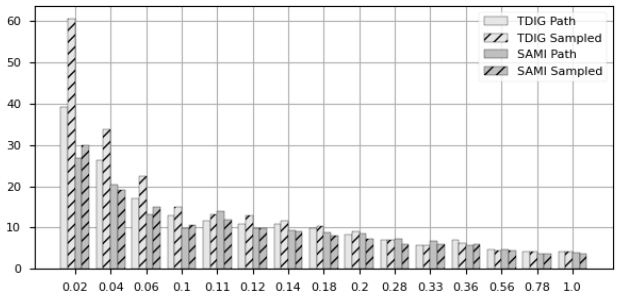

Fig. 7. Comparison of SAMI vs. TDIG maximization for number of CSCP iterations to convergence.

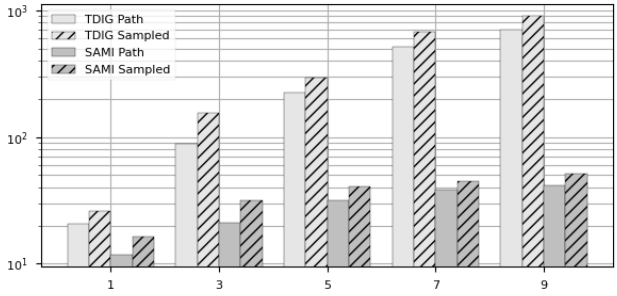

Fig. 8. Average computation time per iteration (in seconds) for TDIG and SAMI maximization with respect to sensor count.

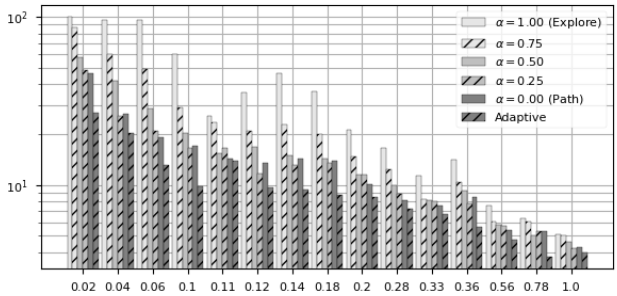

Fig. 9. Average iterations to converge for various fixed α values and self-adaptive α with sequential SAMI.

formulated in (5). An $\alpha = 1$ corresponds to exploration greedy optimization whereas $\alpha = 0$ corresponds to exploitation greedy optimization. Figure 9 shows the performance of the various (manually chosen) values of α in comparison to the self-adaptive method. Performance improves as we move from exploration to exploitation, but the self-adaptive method blends exploration and exploitation and performs the best.

The reader interested may find further detailed results in the first author's dissertation (St. Laurent, 2022). A software implementation of CSCP is available here: https://github.com/Chase1325/CSCP_SAMI

5 Conclusions

We presented an iterative coupled sensor configuration and path-planning method to minimize the number of sensor measurements required to find a near-optimal path. The task-driven information gain (TDIG) metric and its submodular surrogate self-adaptive mutual information (SAMI) were introduced. Due to submodularity, sequential maximization of the SAMI leads to near-optimal sensor configurations. The CSCP method was proven to provide near-optimal path plans in a finite number of iterations. Through numerical simulation experiments, we first demonstrated that sensor configuration via TDIG maximization outperforms traditionally decoupled information maximization. Next, we showed that SAMI maximization is significantly faster compared to TDIG maximization. Finally, we demonstrated that the self-adaptive property of SAMI outperforms any manual choice of the parameter α, which quantifies an exploration versus exploitation trade-off.

Acknowledgements

This research is funded by AFOSR grant #FA9550-17-1-0028 and NSF grants #1646376 and #2126818.

References

S. Aggarwal and N. Kumar. Path planning techniques for unmanned aerial vehicles: A review, solutions, and challenges. *Computer Communications*, 149:270–299, Jan. 2020. ISSN 01403664. doi: 10.1016/j.comcom.2019.10.014.

T. Allen, A. Hill, J. Underwood, and S. Scheding. Dynamic path planning with multi-agent data fusion - The Parallel Hierarchical Replanner. In *2009 IEEE International Conference on Robotics and Automation*, pages 3245–3250, May 2009. doi: 10.1109/ROBOT.2009.5152883.

D. Cochran and A. O. Hero. Information-driven sensor planning: Navigating a statistical manifold. *2013 IEEE Global Conference on Signal and Information Processing, GlobalSIP 2013 - Proceedings*, (0):1049–1052, 2013. doi: 10.1109/GlobalSIP.2013.6737074.

B. S. Cooper and R. V. Cowlagi. Interactive planning and sensing in unknown static environments with task-driven sensor placement. *Automatica*, 105:391–398, Jul 2019. ISSN 00051098. doi: 10.1016/j.automatica.2019.04.014.

M. Demetriou, N. Gatsonis, and J. Court. Coupled controls-computational fluids approach for the estimation of the concentration from a moving gaseous source in a 2-d domain with a Lyapunov-guided sensing aerial vehicle. *IEEE Transactions on Control Systems Technology*, 22(3):853–867, 2013. doi: 10.1109/TCST.2013.2267623.

P. Hart, N. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2): 100–107, 1968. doi: 10.1109/tssc.1968.300136.

A. Krause, A. Singh, and C. Guestrin. Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies. *Journal of Machine Learning Research*, 9:235–284, 2008a. ISSN 15324435. doi: 10.1145/1102351.1102385.

A. Krause, A. Singh, and C. Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9(2), 2008b.

C. Kreucher, A. O. Hero, and K. Kastella. A comparison of task driven and information driven sensor management for target tracking. In *Proc. 44th IEEE Conference on Decision and Control*, pages 4004–4009, 2005.

H. Kurniawati, T. Bandyopadhyay, and N. M. Patrikalakis. Global motion planning under uncertain motion, sensing, and environment map. *Autonomous Robots*, 33(3):255–272, 2012.

S. Li, H. Zhang, S. Liu, and Z. Zhang. Optimal sensor placement using FRFs-based clustering method. *Journal of sound and vibration*, 385:69–80, 2016. ISSN 0022-460X.

R. Madankan, S. Pouget, P. Singla, M. Bursik, J. Dehn, M. Jones, A. Patra, M. Pavolonis, E. B. Pitman, T. Singh, and P. Webley. Computation of probabilistic hazard maps and source parameter estimation for volcanic ash transport and dispersion. *Journal of Computational Physics*, 271:39–59, 2014.

L. Merino, F. Caballero, J. R. Martínez-de Dios, J. Ferruz, and A. Ollero. A cooperative perception system for multiple uavs: Application to automatic detection of forest fires. *Journal of Field Robotics*, 23(3-4):165–184, 2006.

C. A. Micchelli, Y. Xu, and H. Zhang. Universal kernels. *J. Mach. Learn. Res.*, 7:2651–2667, dec 2006. ISSN 1532-4435.

D. Ramsden. *Optimization Approaches To Sensor Placement Problems*. PhD thesis, Rensselaer Polytechnic Institute, 2009.

J. Ranieri, A. Chebira, and M. Vetterli. Near-optimal sensor placement for linear inverse problems. *IEEE Transactions on Signal Processing*, 62(5):1135–1146, 2014. doi: 10.1109/TSP.2014.2299518.

P. Skoglar, J. Nygards, and M. Ulvklo. Concurrent Path and Sensor Planning for a UAV - Towards an Information Based Approach Incorporating Models of Environment and Sensor. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2436–2442, Oct. 2006. doi: 10.1109/IROS.2006.281685.

C. St. Laurent. *Coupled Sensor Configuration and Path-Planning in Uncertain Environments using Multimodal Sensors*. PhD thesis, Worcester Polytechnic Institute, 2022. URL https://digital.wpi.edu/show/n296x2271.

C. St. Laurent and R. V. Cowlagi. Coupled sensor configuration and path-planning in unknown static environments. In *2021 American Control Conference (ACC)*, pages 1535–1540, 2021a. doi: 10.23919/ACC50511.2021.9483071.

C. St. Laurent and R. V. Cowlagi. Breadth-first coupled sensor configuration and path-planning in unknown static environments. In *2021 Conference on Decision and Control (CDC)*, December 2021b.

C. St. Laurent and R. V. Cowlagi. Depth-first coupled sensor configuration and path-planning in unknown static environments. In *2021 European Control Conference (ECC)*, pages 1733–1738, 2021c. doi: 10.23919/ECC54610.2021.9655145.

V. Tzoumas, L. Carlone, G. J. Pappas, and A. Jadbabaie. Lqg control and sensing co-design. *IEEE Transactions on Automatic Control*, 66(4):1468–1483, 2021. doi: 10.1109/TAC.2020.2997661.

D. Yoganathan, S. Kondepudi, B. Kalluri, and S. Manthapuri. Optimal sensor placement strategy for office buildings using clustering algorithms. *Energy and buildings*, 158:1206–1225, 2018. ISSN 0378-7788.