Coupled Sensor Configuration and Path-Planning in a Multimodal Threat Field*

Chase L. St. Laurent[1] and Raghvendra V. Cowlagi[2,3]

[1] AMETEK Inc., Peabody MA 01960, USA clstlaurent@wpi.edu
[2] Worcester Polytechnic Institute, Worcester MA 01609, USA rvcowlagi@wpi.edu
https://labs.wpi.edu/ace-lab/

Abstract. A coupled path-planning and sensor configuration method is proposed. The path-planning objective is to minimize exposure to an unknown, spatially-varying, and temporally static scalar field called the *threat* field. The threat field is modeled as a weighted sum of several scalar fields, each representing a *mode* of threat. A heterogeneous sensor network takes noisy measurements of the threat field. Each sensor in the network observes one or more threat modes within a circular field of view (FoV). The sensors are configurable, i.e., parameters such as location and size of field of view can be changed. The measurement noise is assumed to be normally distributed with zero mean and a variance that monotonically increases with the size of the FoV, emulating the FoV v/s resolution trade-off in most sensors. Gaussian Process regression is used to estimate the threat field from these measurements. The 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. Because the TDIG does not have any convenient structural properties, a surrogate function called the self-adaptive mutual information (SAMI) is considered. Sensor configuration based on the TDIG or SAMI introduces coupling with path-planning in accordance with the dynamic data-driven application systems paradigm. The benefit of this approach is that near-optimal plans are found with a relatively small number of measurements. In comparison to decoupled path-planning and sensor configuration based on traditional information-driven metrics, the proposed CSCP method results in near-optimal plans with fewer measurements.

Keywords: sensor networks · trajectory- and path-planning · bayesian methods · sensor configuration

1 Introduction

Consider applications where an autonomous mobile agent learns about its unknown environment using data collected by an exteroceptive sensor network. For example, we envision a situation where the mobile agent - henceforth called an *actor* - needs to find a minimum-threat path in an adverse environment. The nature of the threat is multimodal and correlated, e.g., fire, smoke, and heat. A network of mobile sensors, e.g., unmanned aerial vehicles (UAVs), is available to collect data about the threat, but each sensor may be limited in its ability to detect the different threat modalities, e.g., one UAV may carry a camera that visually detect fire and smoke, whereas another UAV may carry a temperature sensor. In this situation, it may be: (1) possible to *configure* the sensors, e.g., send different types of sensors to different locations, and (2) crucial for the actor to find a plan with high confidence but with *as few measurements as possible.* To this end we ay ask: *how do we optimally configure sensors to find a near-optimal plan with a minimal number of measurements?*

We study the problem of finding a path of minimum threat exposure in an unknown environment. The unknown threat is a sum of multiple scalar fields representing modes. Information about the threat is gained through data collected by a heterogeneous sensor network, where each sensor is able to detect one or more modes of threat.

Related Work: Optimal path-planning and the related field of motion-planning are well studied for metrics such as minimum path length, maximum traversal utility, and obstacle avoidance [1,2]. Classical approaches include artificial potential fields, probabilistic roadmaps, and cell decomposition. Discrete-space methods such as Dijkstra's algorithm [3], A*, and its variants [4] are well-known for path-planning. Probabilistic

* Supported by NSF grant #1646367 and #2126818.

techniques are used for planning under uncertainty [5], including dynamic programming and its variants. [6]. Partially observable Markov Decision process models are typically used in uncertain planning tasks when the agent has active onboard sensing [7].

The sensor configuration literature is focused on sensor *placement*. Optimal sensor placement addresses, for example, minimizing uncertainty, or maximizing spatial coverage and communication reachability [8]. Optimization metrics include set coverage and information theoretical criteria such as Kullback-Leibler divergence, Fisher information, and mutual information [9]. Sensor placement applications include placement for estimation of gaseous plumes [10], cooperative tracking of forest fires [11], and observing dynamics of volcanic ash [12]. Near-optimal sensor placement for linear inverse problems are studied by [13]. Clustering-based algorithms such as k-means [14] and density-based clustering [15] are studied for sensor placement in office spaces.

A comparison of *task-driven* versus *information-driven* sensor placement in tracking applications is discussed [16]. The literature cited above represents information-driven approaches. Task-driven approaches include recent works, for example, on optimal sensor selection for linear quadratic Gaussian feedback control systems [17], and hierarchical path-replanning concurrently with multi-agent data fusion [18]. Target tracking UAVs with limited field of view are studied in applications where the UAVs find optimal paths and optimal sensing position simultaneously [19].

The authors' previous work has addressed coupled sensor configuration and path-planning [20,21,22]. In all of these works, we assumed that the network of extroceptive sensors is homogeneous and that the underlying threat environment is unimodal. In practice, heterogeneous sensor types can be utilized to capture various modalities of an environment that can be measured observable by a particular sensor type. In this work, we address heterogeneous sensor networks that observe multimodal threat.

There are many situations in which a sensing agent (e.g., a UAV) could have a payload with multiple heterogeneous sensor types, allowing for simultaneous sensing of variably correlated threat modalities in an environment. In what follows, we address situations in which a sensor network can be comprised of heterogeneous sensing agents with the following scenarios: (1) every sensing agent payload is equipped with every sensor modality, (2) each the sensing agent's payload is deficient in at least one modality, and (3) the sensing agent payloads each contain a sensor modality.

The details of the threat and sensor models, and the proposed coupled sensing and path-planning technique are provided next.

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 $N \times N$ 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 $\boldsymbol{a}$. $\boldsymbol{I}_{(N)}$ denotes the identity matrix of size N. For $\mu, \sigma \in \mathbb{R}$, $\mathcal{N}(\mu, \sigma^2)$ denotes the normal distribution with mean μ and variance σ^2.

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 of this graph consist of pairs of geometrically adjacent grid points. In a minor abuse of notation, 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. Without loss of generality we consider "4-way" adjacency of points (i.e., adajcent points are top, down, left, and right).

A *threat field* $c : \mathcal{W} \to \mathbb{R}_{>0}$ is a strictly positive temporally static scalar field. A *path* $\boldsymbol{\pi} = (\boldsymbol{\pi}[0], \boldsymbol{\pi}[1], \ldots, \boldsymbol{\pi}[\lambda])$ between prespecified initial and goal vertices $v_0, v_{\text{L}} \in V$ is a finite sequence, without repetition, of successively adjacent vertices such that $\boldsymbol{\pi}[0] = v_0$ and $\boldsymbol{\pi}[\lambda] = v_{\text{L}}$ 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.

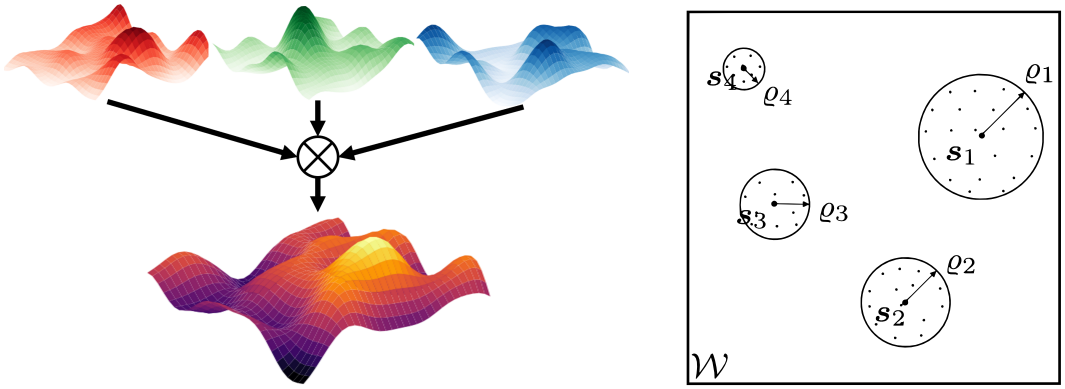

(a) Multimodal threat field.

(b) Sensor FoVs and points of observations (red dots).

Fig. 1. Illustration of the threat and sensor models considered in this work.

We cannot solve this problem as stated because the threat field is unknown. The threat field can be measured by sensors, but each sensor may not be able to measure all of the modalities that constitute the threat. For the sake of a tangible example, sensor modalities may include (1) electro-optical (EO) imaging, (2) infrared (IR) imaging, and (3) a lidar (LI) point cloud. Fusion of these sensors could occur at the *sensor level* as raw data passes through circuitry, at the *data level* as the analog data acquisition becomes digitized, at the *feature level* where the data is combined through a latent embedded space, at the *decision level* which fuses independent output decisions from each sensor type, or at the *mission level* which fuses data with respect to spatial or task relevant correlations. In this work, we attempt to emulate mission level fusion via a statistical field estimation formulation. In terms of the aforementioned sensors, this would mean fusion occurs after the EO, IR, and LI sensor data is digitized, a context-based decision about the data is performed, including spatial association.

The *context* of the multimodal data fusion is of significant importance to how fusion is performed. We define each i^{th} threat field modality as $c^{(i)} : \mathcal{W} \to \mathbb{R}_{>0}$ as a strictly positive temporally static scalar field, as illustrated in Fig. 1(a). We then define a *fused* threat field as $\check{c} := \boldsymbol{m}[1]c^{(1)} + \boldsymbol{m}[2]c^{(2)} + \cdots + \boldsymbol{m}[N_m]c^{(N_m)}$, where $\boldsymbol{m}$ is a user specified weighted fusion vector. The values prescribed to $\boldsymbol{m}$ define the context in which the fusion occurs. The path cost is then calculated as $\mathcal{J}(\boldsymbol{\pi}) := \Delta p \sum_{j=1}^{\lambda} \check{c}(\boldsymbol{p}_{\boldsymbol{\pi}_j})$. Lastly, in what follows we make use of a multimodal vertex set $\tilde{V} := \{N_{\mathrm{g}} N_m\}$ to represent the vertex indices scaled to N_m modalities.

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 $k \in \{N_{\mathrm{s}}\}$. Maximum and minimum FoV radius constraints are specified as $\varrho^{\max}$ and $\varrho^{\min}$, respectively. The set of all sensor parameters is called a *configuration*, which we denote by $\mathcal{C} = \{\boldsymbol{s}_1, \varrho_1, \boldsymbol{s}_2, \ldots, \varrho_{N_{\mathrm{s}}}\}$.

We introduce an observability incidence vector for each k^{th} sensor where $\boldsymbol{o}_k \in \{0,1\}^{N_m}$, which characterizes if the k^{th} sensor can view a threat modality and is specified by the user given each sensing agent's payload. We create an *observed* cover incidence matrix for each k^{th} sensor as $\tilde{\boldsymbol{\nu}}_k := \mathrm{vec}(\boldsymbol{\nu}_k \boldsymbol{o}_k) \in \mathbb{R}^{N_g N_m \times 1}$, where we define $\mathrm{vec}(\cdot) := \mathbb{R}^{N_g \times N_m} \to \mathbb{R}^{N_g N_m \times 1}$. Finally, the combined observed covered incidence vector becomes $\tilde{\boldsymbol{\nu}} := (\tilde{\boldsymbol{\nu}}_1 \vee \tilde{\boldsymbol{\nu}}_2 \vee \ldots \vee \tilde{\boldsymbol{\nu}}_{N_{\mathrm{s}}})$.

We update the sensor observations notation as follows. The collection of sensor data locations for a particular i^{th} sensor type is denoted as $\boldsymbol{X}^{(i)} = \{\boldsymbol{x}_{11}^{(i)}, \ldots, \boldsymbol{x}_{km}^{(i)}, \ldots, \boldsymbol{x}_{N_{\mathrm{s}} M_{N_{\mathrm{s}}}}^{(i)}\} \; \forall \; i \in N_m$. We denote by, $\mathcal{X} = \{\boldsymbol{X}^{(1)}, \boldsymbol{X}^{(2)}, \ldots, \boldsymbol{X}^{(N_m)}\}$, the set of sensor data locations for each sensing modality which is aggregated or updated with each iteration ℓ. The training matrix of the training set augmented by the corresponding modality index is formulated as follows:

$$\tilde{\boldsymbol{X}} := \begin{bmatrix} \boldsymbol{X}^{(1)} & \boldsymbol{X}^{(2)} & \ldots & \boldsymbol{X}^{(N_m)} \\ \boldsymbol{0} & \boldsymbol{1} & \ldots & \boldsymbol{N_m} \end{bmatrix}^{\mathsf{T}} \tag{1}$$

Similarly, we define the collection of vectorized noisy threat field observations and observation noise as $\tilde{\boldsymbol{z}} = \begin{bmatrix}\boldsymbol{z}_1^\intercal \ \boldsymbol{z}_2^\intercal \ \dots \ \boldsymbol{z}_{N_m}^\intercal\end{bmatrix}^\intercal$ and $\tilde{\boldsymbol{\sigma}} = \begin{bmatrix}\boldsymbol{\sigma_1}^\intercal \ \boldsymbol{\sigma_2}^\intercal \ \dots \ \boldsymbol{\sigma_{N_m}}^\intercal\end{bmatrix}^\intercal$, respectively. We say that each of the observations for each i^{th} modality is modeled as $\boldsymbol{z}_{km}^{(i)} = c^{(i)}(\boldsymbol{x}_{km}) + \eta_{km}^{(i)}$. The measurement error η increases monotonically with the sensor FoV.

3 Multimodal Field Estimation

We use Gaussian Process regression (GPR) to construct an estimate of the multimodal threat field. Basic details regarding the GPR approach to unimodal threat estimation are provided in our previous work [22]. The joint posterior distribution is modified as follows to account for the multimodal vectorized output and multimodal threat field estimates:

$$\begin{bmatrix}\tilde{\boldsymbol{z}} \\ \tilde{\boldsymbol{f}}\end{bmatrix} \sim \mathcal{N}\left(\boldsymbol{0}, \begin{bmatrix}\boldsymbol{K}_{\tilde{\boldsymbol{z}}} & \boldsymbol{K}_* \\ \boldsymbol{K}_*^\intercal & \boldsymbol{K}_{**}\end{bmatrix}\right), \tag{2}$$

where $\boldsymbol{K}$ is the kernel function. We modify the kernel structure to enable learning of the cross-correlation between modalities. Namely, we utilize the intrinsic model of coregionalization (ICM) kernel $\boldsymbol{K}^X$ defined as:

$$\boldsymbol{K}^X = (\boldsymbol{\Theta}_w \boldsymbol{\Theta}_w^\intercal + \mathrm{diag}(\boldsymbol{\theta_v})) \tag{3}$$

The ICM kernel is a matrix of size $N_m \times N_m$ and has learnable parameters $\boldsymbol{\Theta_w}$ and $\boldsymbol{\theta_v}$. The matrix of parameters is of size $N_m \times r$, where $r \in \mathbb{N}$ is a small value to emulate low-rank positive definite correlation between modalities. The parameter vector $\boldsymbol{\theta}_v$ is of size $N_m \times 1$ and models the independent scaling factor of each modality. For threat field modeling, we utilize the kernel $\boldsymbol{K}_{ij} = \boldsymbol{K}_{ij}^R \cdot \boldsymbol{K}^X[i,j] \ \forall \ i,j \in N_m$. The resulting kernel has the form:

$$\boldsymbol{K} := K(\tilde{\boldsymbol{X}}, \tilde{\boldsymbol{X}}) = \begin{bmatrix} \boldsymbol{K}_{11} & \boldsymbol{K}_{12} \cdots & \boldsymbol{K}_{1N_m} \\ \boldsymbol{K}_{21} & \ddots \ \vdots & \vdots \\ \vdots & \vdots \ \ddots & \vdots \\ \boldsymbol{K}_{N_m 1} & \cdots \ \cdots & \boldsymbol{K}_{N_m N_m} \end{bmatrix}$$

The collection of hyperparameters to optimize are $\boldsymbol{\theta} = (\boldsymbol{\Theta}_r, \boldsymbol{\Theta}_w, \boldsymbol{\theta}_v)$. The multimodal input kernel with heteroscedastic noise vector is formulated as $\boldsymbol{K}_{\tilde{\boldsymbol{z}}} := \boldsymbol{K} + diag(\tilde{\boldsymbol{\sigma}})$. The diagonal elements of $\boldsymbol{K}$ represent the auto-covariance of points within a modality, whereas the off-diagonal elements represent the cross-covariance which models the latent relationship between modalities.

From the joint distribution, we can obtain the current iteration multimodal threat field estimate and multimodal threat error covariance matrix as:

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

We note that the multimodal threat field estimate and multimodal threat field error covariance matrix are constructed as:

$$\tilde{\boldsymbol{f}}_\ell = \begin{bmatrix}\boldsymbol{f}_\ell^{(1)\intercal} \ \boldsymbol{f}_\ell^{(2)\intercal} \ \dots \ \boldsymbol{f}_\ell^{(N_m)\intercal}\end{bmatrix}^\intercal, \qquad\qquad \tilde{\boldsymbol{P}}_\ell = \begin{bmatrix} \boldsymbol{P}_\ell^{(11)} & \boldsymbol{P}_\ell^{(12)} \cdots & \boldsymbol{P}_\ell^{(N_m)} \\ \boldsymbol{P}_\ell^{(21)} & \ddots \ \vdots & \vdots \\ \vdots & \vdots \ \ddots & \vdots \\ \boldsymbol{P}_\ell^{(N_m 1)} & \cdots \ \cdots & \boldsymbol{P}_\ell^{(N_m N_m)} \end{bmatrix}.$$

The fused threat field estimate and fused threat field error covariance matrix can then be computed using the weighted fusion vector $\boldsymbol{m} \in \mathbb{R}^{N_m}$ as:

$$\check{\boldsymbol{f}}_\ell = \sum_{i=1}^{N_m} \boldsymbol{m}[i] \tilde{\boldsymbol{f}}_\ell^{(i)}, \qquad\qquad \check{\boldsymbol{P}}_\ell = \sum_{i=1}^{N_m} \boldsymbol{m}^2[i] \tilde{\boldsymbol{P}}_\ell^{(ii)} + 2 \sum_{j=2}^{N_m} \sum_{k=1}^{j-1} \boldsymbol{m}[j]\boldsymbol{m}[k] \tilde{\boldsymbol{P}}_\ell^{(jk)}. \tag{4}$$

We may then use the fused threat field and fused threat error covariance matrix to find the estimated optimal path-plan and the estimated path-plan variance.

4 Coupled Sensor Configuration

The coupled sensor configuration and path-planning (CSCP) method for unimodal threats was introduced in our previous works, and is summarized in Fig. 2. The key step in the CSCP approach to configure sensors by maximizing what is called the task-driven information gain (TDIG). Informally, TDIG is a measure of entropy reduction in a region of interest "near" the currently optimal path. Further details are in our previous works. In [23], we introduced a function called *self-adaptive mutual information (SAMI)*, which approximates the TDIG, but has the added benefit of being submodular.

1: Initialize: $\ell := 0$, set initial mean threat estimate to zero, and covariance arbitrarily large
2: Find initial estimated optimal path $\boldsymbol{\pi}_0^* := \arg\min \overline{\mathcal{J}}_0(\boldsymbol{\pi})$
3: **while** $\mathrm{Var}_\ell(\boldsymbol{\pi}_\ell^*) > \varepsilon$ **do**
4: Find optimal sensor configuration $\mathcal{C}_\ell^*$ (as discussed in this section)
5: Record new measurements with the new sensor configuration
6: Increment iteration counter $\ell := \ell + 1$
7: Update mean threat estimate and estimation error covariance $\boldsymbol{P}_\ell$ (see §3)
8: Find $\boldsymbol{\pi}_\ell^* := \arg\min \overline{\mathcal{J}}_\ell(\boldsymbol{\pi})$

Fig. 2. Proposed iterative CSCP method.

The introduction of heterogeneous sensors presents some additional nuances. Since each sensor network can have multiple modalities, we consider three separate situations in which the sensor configuration problem can be formalized. In what follows, we describe situations in which every sensor is equipped with every modality, a mixture of modalities, or only a single modality.

First, we say that a multimodal region of interest is defined as the collection of region of interest vertices with index values biased by the appropriate mode index as $\tilde{\mathcal{R}} := \{\cup_{i=0}^{N_m-1} \mathcal{R} + iN_\mathrm{g}\}$. The multimodal region of interest is used for indexing in the SAMI formulation. The entropy of any vertex $i \in \tilde{V}$ as:

$$h(i) := \frac{1}{2} \ln(2\pi e \tilde{\boldsymbol{P}}_\ell[i,i]) \tag{5}$$

In certain cases the conditional entropy computations can be batched and computed in parallel. For any $i \notin \mathcal{R}$, the batched conditional entropy vector $\boldsymbol{h}(\cdot|\tilde{\mathcal{R}}_{\setminus i})$ is computed as:

$$\boldsymbol{h}(\cdot|\tilde{\mathcal{R}}_{\setminus i}) = \frac{1}{2} \ln((2\pi e) diag(\tilde{\boldsymbol{P}}_\ell - \tilde{\boldsymbol{P}}_\ell[\cdot,\tilde{\mathcal{R}}]\tilde{\boldsymbol{P}}_\ell[\tilde{\mathcal{R}},\tilde{\mathcal{R}}]^{-1}\tilde{\boldsymbol{P}}_\ell[\tilde{\mathcal{R}},\cdot])) \tag{6}$$

We note that the matrix inverse only needs to be computed once. However, for any $i \in \mathcal{R}$, the conditional entropy of any vertex $i \in \{N_g N_m\}$ given $\tilde{\mathcal{R}}_{\setminus i}$ can be computed in parallel as:

$$h(i|\tilde{\mathcal{R}}_{\setminus i}) = \frac{1}{2} \ln((2\pi e)(\tilde{\boldsymbol{P}}_\ell[i,i] - \tilde{\boldsymbol{P}}_\ell[\tilde{\mathcal{R}}_{\setminus i}, i]^\intercal \tilde{\boldsymbol{P}}_\ell[\tilde{\mathcal{R}}_{\setminus i}, \tilde{\mathcal{R}}_{\setminus i}]^{-1} \tilde{\boldsymbol{P}}_\ell[\tilde{\mathcal{R}}_{\setminus i}, i])) \tag{7}$$

Given these equations, we can calculate the mutual information I between the multimodal region of interest and any vertex for a particular mode $i \in \tilde{V}$ as:

$$I(\tilde{\mathcal{R}}_{\setminus i}; i) := h(i) - h(i|\tilde{\mathcal{R}}_{\setminus i}). \tag{8}$$

The SAMI reward term for multimodal threat fields is then:

$$R(i) := (1-\alpha) I(\tilde{\mathcal{R}}_{\setminus i}; i) + \alpha I(\tilde{\mathcal{R}}^c_{\setminus i}; i) \tag{9}$$

In (9), the ROI complement is taken as $\tilde{\mathcal{R}}^c := \tilde{V} \setminus \tilde{\mathcal{R}}$ and $i \in \tilde{V}$. We also update the mutual information reward vector to be $\boldsymbol{\gamma} := \begin{bmatrix}\gamma(1)\ \gamma(2)\ \ldots\ \gamma(N_\mathrm{g})\ \ldots\ \gamma(N_\mathrm{g}N_m)\end{bmatrix}^\intercal$. The reward function given the sensor configuration is $\Gamma(\mathcal{C}_\ell) = \tilde{\boldsymbol{\nu}}^\intercal \boldsymbol{\gamma}$.

The SAMI penalty function is calculated as:

$$\Upsilon(\mathcal{C}_\ell) := -\frac{1}{2}\sum_{i\in\mathcal{F}}\left(\frac{1}{2}\ln(2\pi e) - \ln\sum_{k\in N_\mathrm{s}}(\boldsymbol{\nu}_k \odot \boldsymbol{\sigma}_k^{-2})[\tilde{\mathcal{F}}]\right) \tag{10}$$

In (10), the $\odot$ operator is the Hadamard product, which is used to perform element-wise multiplication between the cover incidence vector and the inverse noise vector $\boldsymbol{\sigma_k}^{-2}$. The entries are indexed by $[\tilde{\mathcal{F}}]$ to ensure only covered vertices are accounted for prior to taking the natural logarithm of each element. The SAMI is the sum of the reward and penalty terms.

Sensors with all modalities $\boldsymbol{o}_k = \mathbf{1}\,\forall\, k \in N_\mathrm{s}$: this problem reduces to maximizing the SAMI surrogate objective function directly over multimodal SAMI surrogate values. Because the SAMI is submodular, we can sequentially optimize sensor configuration for each set of k^{th} sensor parameters $\{\boldsymbol{s}_k, \varrho_k\}$ to find $\mathcal{C}_\ell^* := \arg\max S(\mathcal{C}_\ell)$.

Sensors with overlapping modalities: To illustrate this situation, suppose that we have available three sensors $\mathcal{S} = \{\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3\}$ with the following modalities: $\mathcal{S}_1 = \{\text{EO, IR}\}$, $\mathcal{S}_2 = \{\text{EO, LI}\}$, and $\mathcal{S}_3 = \{\text{IR, LI}\}$. To determine the sequence in which a sensing agent in the sensor network is optimized to adhere to the submodularity property, we need to calculate the total reward for each modality as $\bar{\gamma}_i = \sum_{j\in N_\mathrm{g}}(1 - \tilde{\boldsymbol{\nu}}^{(i)})\boldsymbol{\gamma}^{(i)}$ prior to each sequential sensor configuration optimization. We denote the $\tilde{\boldsymbol{\nu}}^{(i)}$ as the partition of the multimodal sensor cover incidence for the i^{th} modality. Similarly, the i^{th} modality partition of the SAMI reward vector is denoted $\boldsymbol{\gamma}^{(i)}$. The total rewards is then $\bar{\boldsymbol{\gamma}} = \begin{bmatrix}\bar{\gamma}_1 & \bar{\gamma}_2 & \dots & \bar{\gamma}_{N_m}\end{bmatrix}^\intercal$. We may then calculate the total potential reward for each k^{th} sensor as $\bar{\gamma}^{(k)} = \boldsymbol{o}_k\bar{\boldsymbol{\gamma}}$. Therefore, each iteration of sequential sensor configuration, we choose $\mathcal{S}_k = \arg\max\bar{\gamma}^{(k)}$, optimize it with the SAMI objective function, and remove it from the set $\mathcal{S}$ (just for that round of optimization). In our example, if $\bar{R}^{(2)}$ was the maximum value, we would perform sensor configuration optimization with $\mathcal{S}_2$ and remove it from the set of optimizable sensing agents, leaving only $\mathcal{S}_1$ and $\mathcal{S}_3$ to be optimized. We proceed to score the sensing agents again until they have all been optimized.

Sensors with unique modalities: This is a special observability case which has unique sensor configuration optimization implications. We no longer require ranking the sensor modalities as they are entirely separable and additive. Due to this, the sensor configuration for uniquely observable sensor payloads allows for parallelization of sequential sensor configuration for each modality. In fact, each mode reduces to solving independent sensor configuration optimization N_m times in parallel, allowing for computational savings.

5 Results and Discussion

We conducted a study with four various sensor networks in an environment of area $9km^2$, workspace resolution of 21^2, and a desired termination threshold $\varepsilon = 1$. We considered a multimodal threat field with $N_m = 3$ correlated threat modalities. All mobile sensing agents were constrained to $\varrho^{\min} = 0.05$ and $\varrho^{\max} = 0.5$. The region of interest for the experiments was found with $N_a = 3$ alternate path plans.

We considered four various sensor network scenarios for the experiments. The first sensor network $\mathcal{S}_A$ was comprised of 3 sensors with all modalities. $\mathcal{S}_B$ had three sensors each with two modalities such that each threat modality was observable by at least 2 sensors. It also included one sensor with all modalities totaling $N_s = 4$. Sensor network $\mathcal{S}_C$ had the same pairs of modalities as $\mathcal{S}_B$, but instead of the single full-modal sensor, it had three unimodal sensors (one for each modality) totaling $N_s = 6$. Finally, we considered a sensor network $\mathcal{S}_D$ which was comprised of $N_s = 9$ unimodal sensors such that each modality equally had 3 sensors which could make observations. These networks were chosen to study the effect of modifying the degrees of freedom of the sensor network on convergence and sensor configuration optimization search time. The average iterations for convergence and the average sensor configuration optimization time for each sensor network are recorded in Table 1. We make the following observations from the collected results. In summary, the parallel optimization of sensor networks with unique modalities, by way of the separable nature of the multimodal formulation of the SAMI surrogate function, are able to quickly find sensor configurations whilst providing good convergence performance.

Increasing sensor network flexibility improves convergence performance. We may characterize network flexibility in context of the constraint of having multiple modalities with identical FoVs to optimize

Table 1. Comparative study for various sensor payload configurations.

Sensor Network	Iterations	Configuration Time
Network $\mathcal{S}_A$	16.331 ± 4.163	2.313 ± 0.683
Network $\mathcal{S}_B$	16.033 ± 3.436	3.068 ± 0.995
Network $\mathcal{S}_C$	15.667 ± 4.619	4.900 ± 1.472
Network $\mathcal{S}_D$	14.330 ± 4.041	2.878 ± 1.004

on a single sensor. The flexibilities for the experiment sensor networks are in increasing order $\mathcal{S}_A < \mathcal{S}_B < \mathcal{S}_C < \mathcal{S}_D$. The results show that adding the increased flexibility by not confining multiple multimodal sensor payloads to a single sensing agent proportionally decreases the average iterations for path-plan convergence. This answers a typical quandary between, say, uying an expensive UAV equipped with many heterogeneous sensors versus buying multiple inexpensive UAVs each with a single modality. For coupled sensing and path-planning problems the additional flexibility by distributing the sensors across sensing agents within the network is more valuable.

Parallel optimization of unimodal sensor payloads yields low configuration times. Somewhat unsurprisingly, the parallel sequential optimization of sensor network $\mathcal{S}_D$ which had sensing agents with unimodal observability, was able to be optimized faster than sensor networks $\mathcal{S}_B$ and $\mathcal{S}_C$. We note that the reason it still takes longer to configure these sensors over sensor network $\mathcal{S}_A$ is that there are more potential configurations and globally optimizing is made more difficult. However, we note how close the runtime performance is for sensor configuration between $\mathcal{S}_A$ and $\mathcal{S}_D$ and note that the time difference is negligible in comparison to the average iterations until convergence performance gain.

5.1 Demonstrative Example

We demonstrate the MM-CSCP algorithm on an example randomly generated multimodal threat field with $N_m = 3$. The problem definition remains the same as the numerical experiments from the previous section, but we increase the environment area to $25km^2$. The individual threat field modalities and the fused threat field along with the optimal path-plan are shown in Figure 3. We considered a sensor network with 3 sensing agents with unique pairings of observability and 2 sensing agents with unimodal observability for the first two threat modes. Therefore, the first two threat modes have up to three FoV covers per iteration while the third only has two. The initial sensor network configuration is depicted overlaying the SAMI reward function for each vertex in Figure 4.

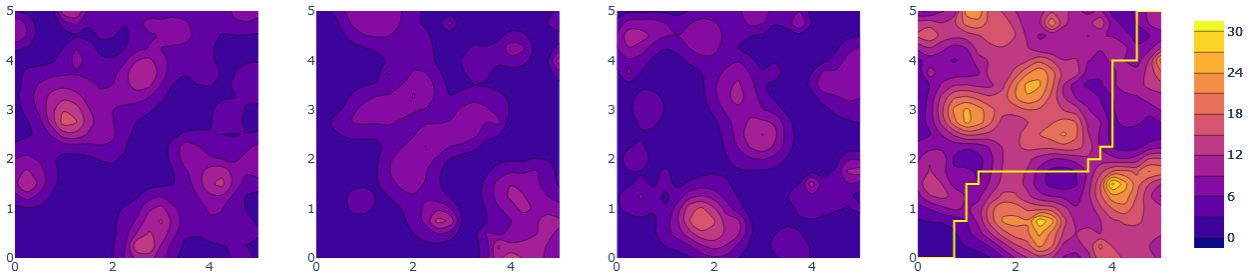

Fig. 3. Ground truth threat field modalities and the fused threat field (right) along with the optimal path-plan.

The data collected from the first iteration sensor configuration is then utilized to determine the multimodal threat field estimate and subsequently the fused threat field estimate as shown in Fig. 5. The resultant estimated path-plan is shown in this figure, and Fig. 6 shows the corresponding multimodal threat error covariance values for each vertex and the fused covariance values. The algorithm proceeds until convergence at iteration $\ell = 22$. Figures 7–18 show the SAMI reward functions and sensor configurations, the multimodal

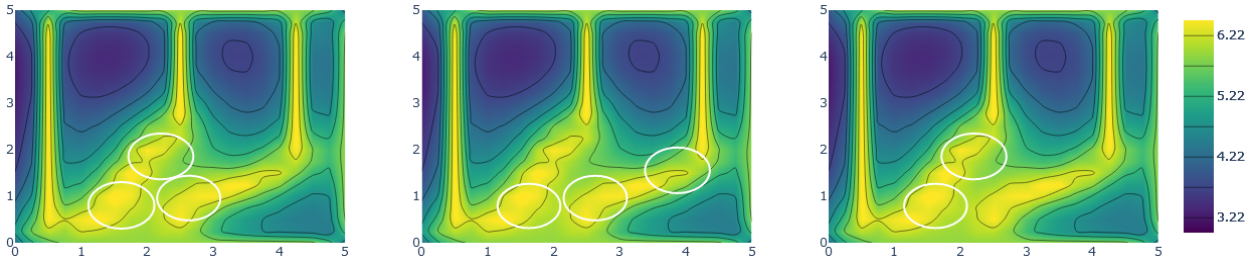

Fig. 4. Initial iteration sensor network configurations overlaying the multimodal SAMI reward values for each mode.

and fused threat field estimates with the estimated path-plan, and the multimodal and fused threat variances at select iterations.

The SAMI reward examples show that in some cases sensing agents with paired observability are configured such that the reward of one modality is considered in conjunction with a potentially low reward modality. We also note that the multimodal threat field estimation and fusion visually does a good job at learning the environment. We also observe from the final fused threat error variance in Fig. 18 that the MM-CSCP still emphasizes learning task-driven regions of interest, rather than the entire environment, in order to plan the path. Ultimately, the estimated path plan is near-optimal.

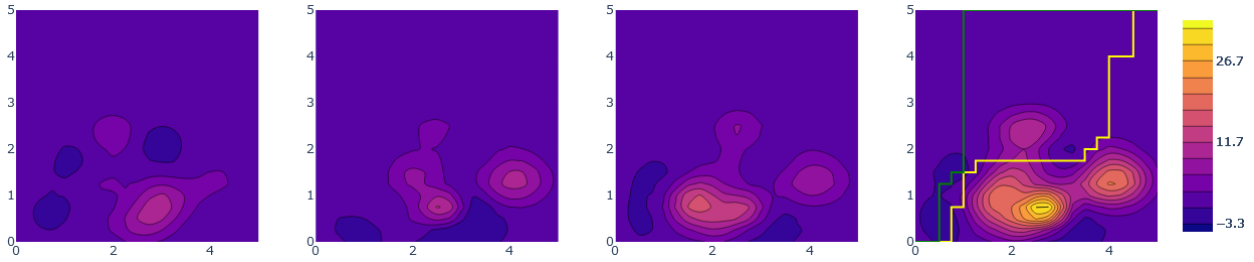

Fig. 5. Initial multimodal threat estimates and fused threat field estimate (right) along with the estimated optimal path plan (green) and the true optimal path plan (yellow).

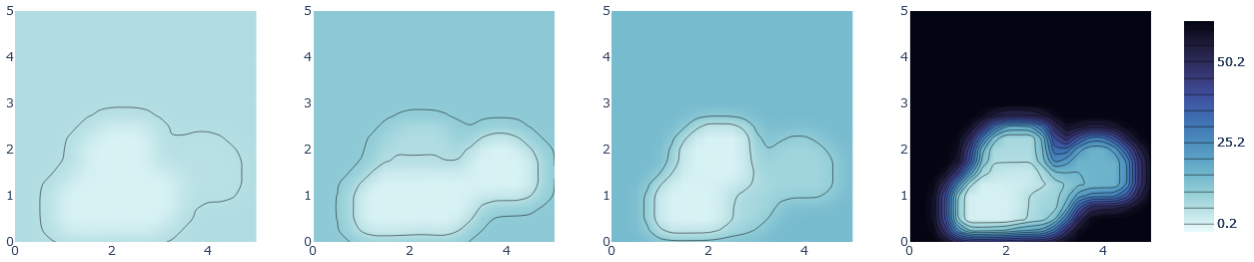

Fig. 6. Initial multimodal threat field error covariance vertex values and fused vertex covariance values.

6 Conclusion

A coupled path-planning and sensor configuration method was proposed. This work was based on the authors' previous work on CSCP in unimodal threats. To extend CSCP to multimodal threats, we introduced different

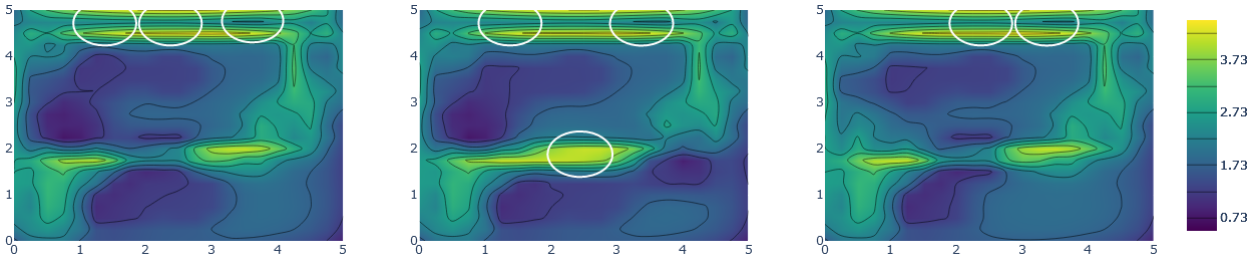

Fig. 7. Sensor network configurations overlaying the multimodal SAMI reward values for each threat mode at $\ell = 3$.

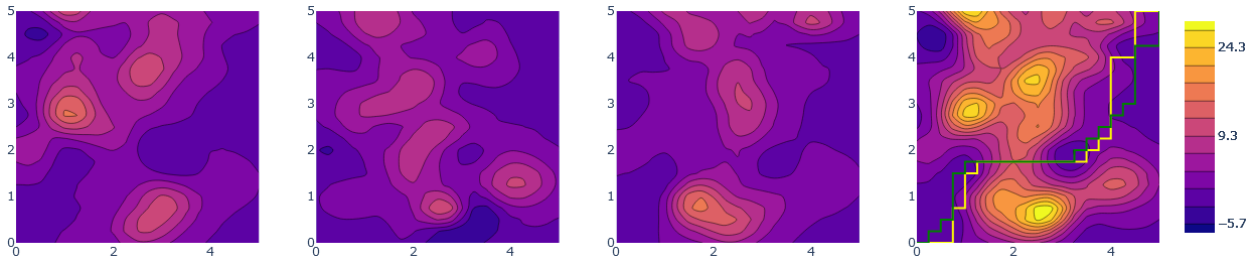

Fig. 8. Multimodal threat field estimates and fused threat field estimate (right) along with the estimated optimal path-plan (green) and the true optimal path-plan (yellow) at iteration $\ell = 3$.

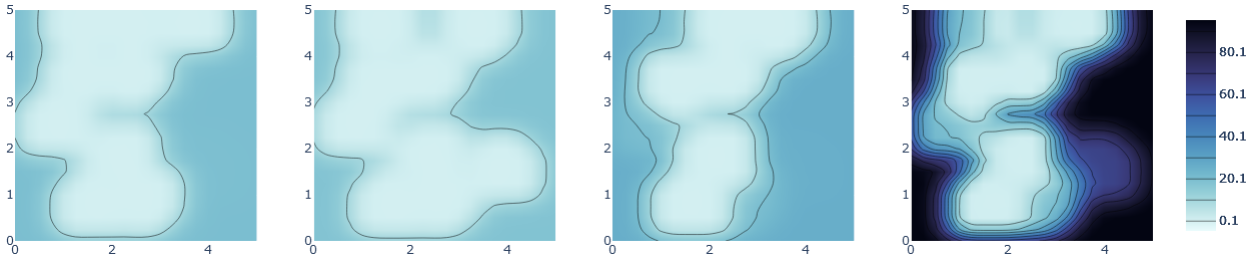

Fig. 9. Multimodal threat field error covariance vertex values and fused vertex covariance values at iteration $\ell = 3$.

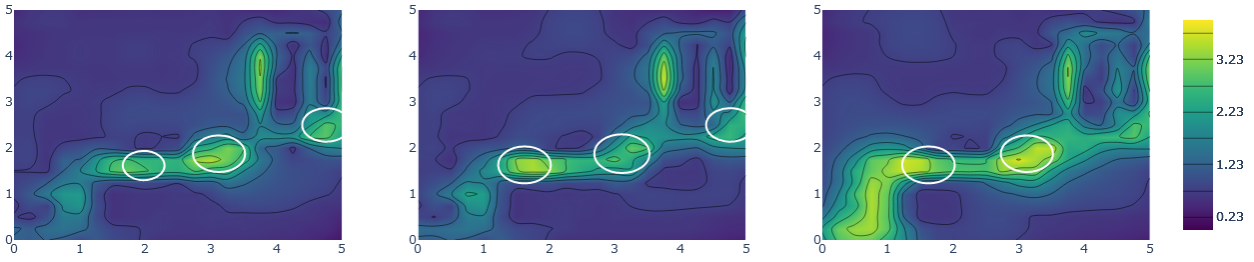

Fig. 10. Sensor network configurations overlaying the multimodal SAMI reward values for each threat mode at $\ell = 7$.

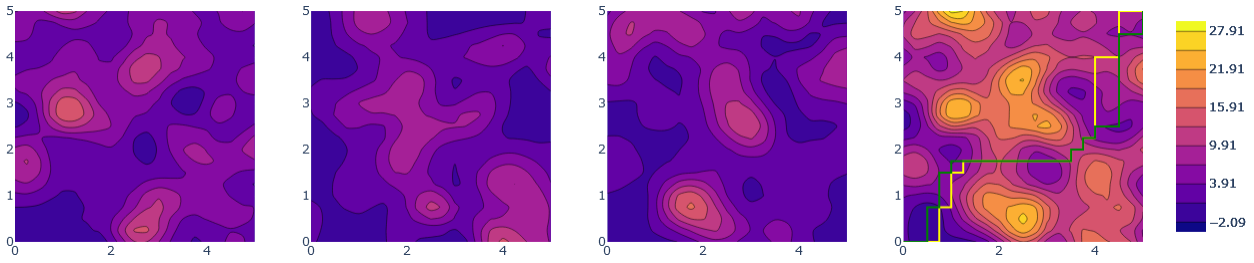

Fig. 11. Multimodal threat field estimates and fused threat field estimate (right) along with the estimated optimal path-plan (green) and the true optimal path-plan (yellow) at iteration $\ell = 7$.

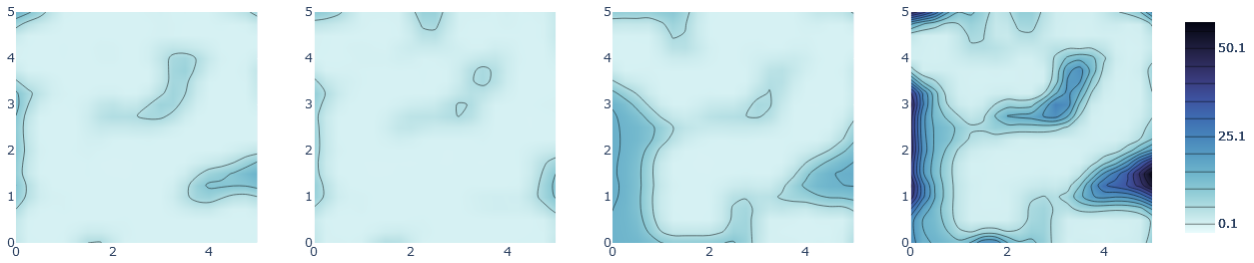

Fig. 12. Multimodal threat field error covariance vertex values and fused vertex covariance values at iteration $\ell = 7$.

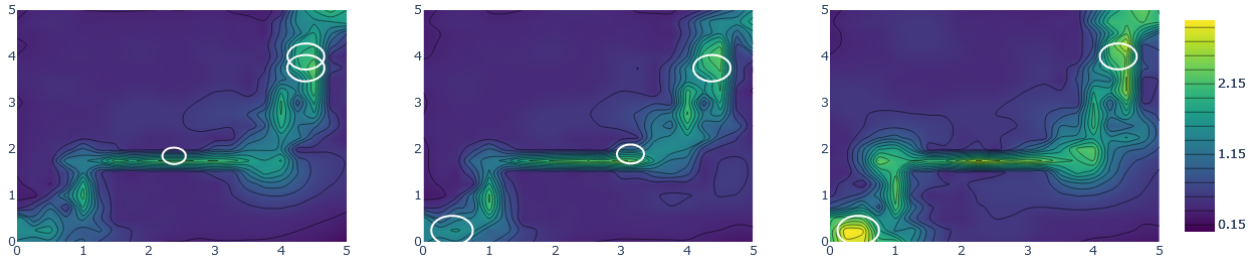

Fig. 13. Sensor network configurations overlaying the multimodal SAMI reward values for each threat mode at $\ell = 13$.

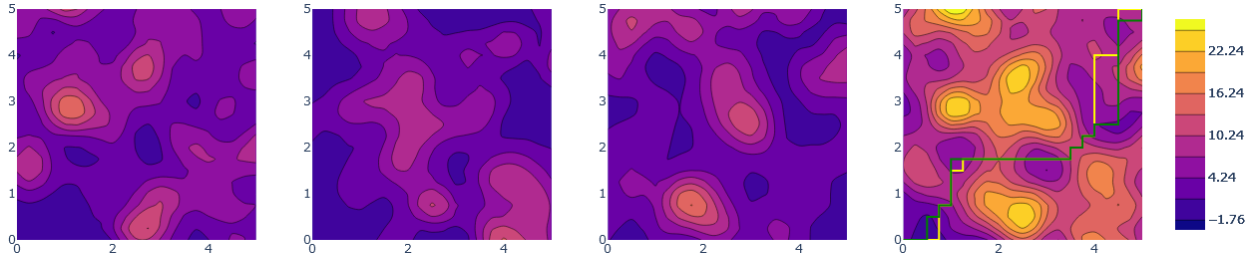

Fig. 14. Multimodal threat field estimates and fused threat field estimate (right) along with the estimated optimal path-plan (green) and the true optimal path-plan (yellow) at iteration $\ell = 13$.

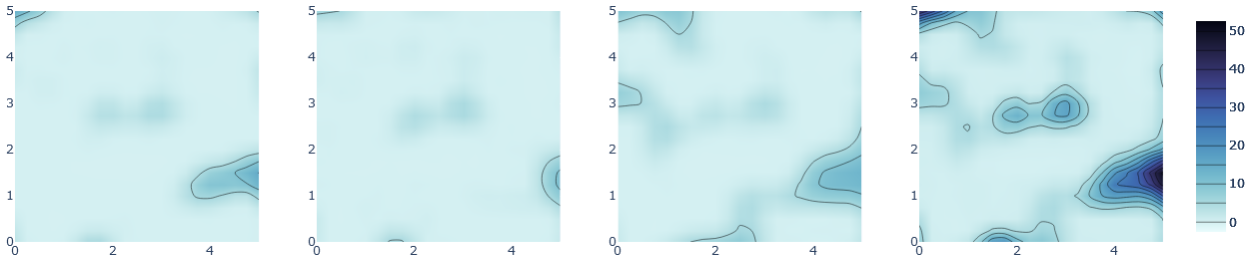

Fig. 15. Multimodal threat field error covariance vertex values and fused vertex covariance values at iteration $\ell = 13$.

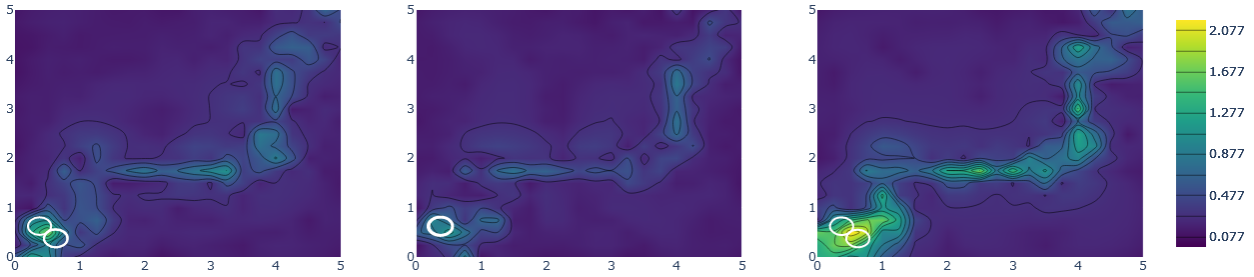

Fig. 16. Sensor network configurations overlaying the multimodal SAMI reward values for each threat mode at $\ell = 22$, the final iteration.

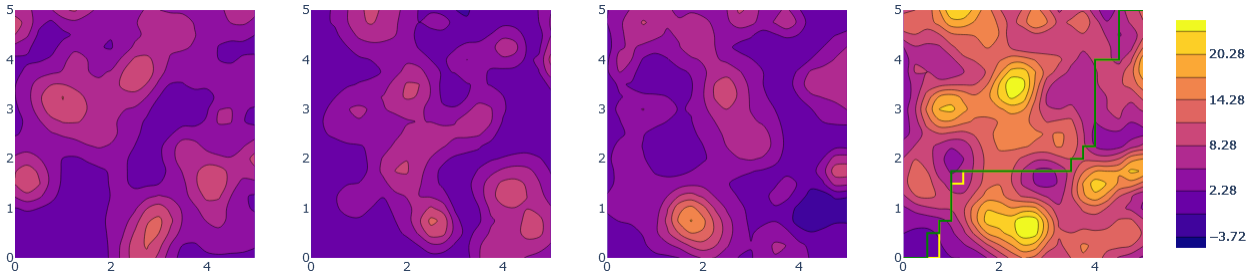

Fig. 17. Multimodal threat field estimates and fused threat field estimate (right) along with the estimated optimal path-plan (green) and the true optimal path-plan (yellow) at iteration $\ell = 22$, the final iteration.

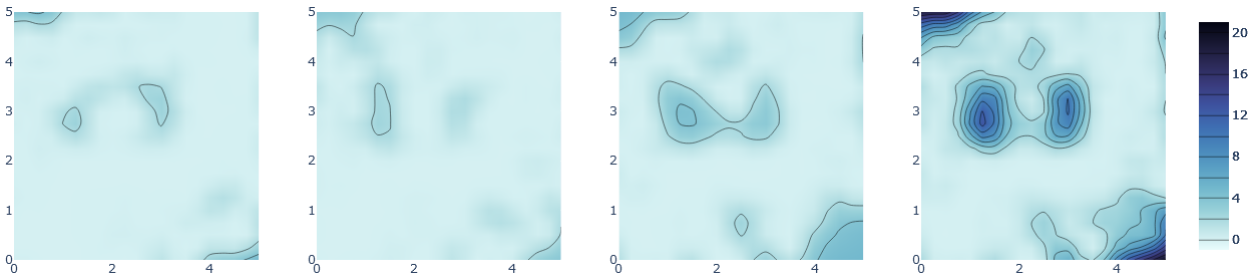

Fig. 18. Multimodal threat field error covariance vertex values and fused vertex covariance values at iteration $\ell = 22$.

ways of calculating and optimizing the SAMI for different cases where each sensor in the network is able to all, some, or one unique modality of the threat. Numerical simulation results were discussed for one demonstrative example and for more thorough studies comparing sensor networks with different levels of flexibilities. Future work includes extension to time-varying multimodal threat fields.

References

1. S. M. LaValle, *Planning Algorithms*. Cambridge University Press, 2006.
2. S. Aggarwal and N. Kumar, "Path planning techniques for unmanned aerial vehicles: A review, solutions, and challenges," *Computer Communications*, vol. 149, pp. 270–299, Jan. 2020.
3. E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische mathematik*, vol. 1, no. 1, pp. 269–271, 1959.
4. 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*, vol. 4, no. 2, pp. 100–107, 1968. [Online]. Available: https://doi.org/10.1109/tssc.1968.300136
5. M. Mohanan and A. Salgoankar, "A survey of robotic motion planning in dynamic environments," *Robotics and Autonomous Systems*, vol. 100, pp. 171–185, Feb. 2018.
6. P. M. Esfahani, D. Chatterjee, and J. Lygeros, "Motion Planning for Continuous-Time Stochastic Processes: A Dynamic Programming Approach," *IEEE Transactions on Automatic Control*, vol. 61, no. 8, pp. 2155–2170, Aug. 2016.
7. H. Kurniawati, T. Bandyopadhyay, and N. M. Patrikalakis, "Global motion planning under uncertain motion, sensing, and environment map," *Autonomous Robots*, vol. 33, no. 3, pp. 255–272, 2012.
8. D. Ramsden, "Optimization Approaches To Sensor Placement Problems," Ph.D. dissertation, 2009.
9. 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*, no. 0, pp. 1049–1052, 2013.
10. 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*, vol. 22, no. 3, pp. 853–867, 2013.
11. 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*, vol. 23, no. 3-4, pp. 165–184, 2006.
12. 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*, vol. 271, pp. 39–59, 2014.
13. J. Ranieri, A. Chebira, and M. Vetterli, "Near-optimal sensor placement for linear inverse problems," *IEEE Transactions on Signal Processing*, vol. 62, no. 5, pp. 1135–1146, 2014.
14. S. Li, H. Zhang, S. Liu, and Z. Zhang, "Optimal sensor placement using FRFs-based clustering method," *Journal of sound and vibration*, vol. 385, pp. 69–80, 2016.
15. D. Yoganathan, S. Kondepudi, B. Kalluri, and S. Manthapuri, "Optimal sensor placement strategy for office buildings using clustering algorithms," *Energy and buildings*, vol. 158, pp. 1206–1225, 2018.
16. 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*, 2005, pp. 4004–4009.
17. V. Tzoumas, L. Carlone, G. J. Pappas, and A. Jadbabaie, "Lqg control and sensing co-design," *IEEE Transactions on Automatic Control*, vol. 66, no. 4, pp. 1468–1483, 2021.
18. 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*, May 2009, pp. 3245–3250.
19. 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*, Oct. 2006, pp. 2436–2442.
20. C. L. St. Laurent and R. V. Cowlagi, "Breadth-first coupled sensor configuration and path-planning in unknown static environments," in *Proceedings of the 60th IEEE Conference on Decision & Control*, 2021.
21. ——, "Depth-first coupled sensor configuration and path-planning in unknown static environments," in *Proceedings of the 2021 European Control Conference*, 2021.
22. ——, "Coupled sensor configuration and path-planning in unknown static environments," in *Proceedings of the 2021American Control Conference*, 2021.
23. C. L. St. Laurent, "Coupled sensor configuration and path-planning in uncertain environments using multimodal sensors," Ph.D. dissertation, Worcester Polytechnic Institute, Worcester, MA, USA, 2022. [Online]. Available: https://digital.wpi.edu/show/n296x2271