

Multi-robot Multi-source Localization in Complex Flows with Physics-Preserving Environment Models

Benjamin Shaffer^{1†}, Victoria Edwards², Brooks Kinch¹, Nathaniel Trask¹, and M. Ani Hsieh¹

Abstract—Source localization in a complex flow poses a significant challenge for multi-robot teams tasked with localizing the source of chemical leaks or tracking the dispersion of an oil spill. The flow dynamics can be time-varying and chaotic, resulting in sporadic and intermittent sensor readings, and complex environmental geometries further complicate a team’s ability to model and predict the dispersion. To accurately account for the physical processes that drive the dispersion dynamics, robots must have access to computationally intensive numerical models, which can be difficult when onboard computation is limited. We present a distributed mobile sensing framework for source localization in which each robot carries a machine-learned, finite element model of its environment to guide information-based sampling. The models are used to evaluate an approximate mutual information criterion to drive an infotaxis control strategy, which selects sensing regions that are expected to maximize informativeness for the source localization objective. Our approach achieves faster error reduction compared to baseline sensing strategies and results in more accurate source localization compared to baseline machine learning approaches.

I. INTRODUCTION

A core challenge in environmental monitoring is *source localization*: recovering the location of an unobservable source from sparse, indirect measurements [1], [2]. A robot team equipped with sensors is a promising solution because the team can adaptively select informative sampling locations and are robust for operations in hazardous conditions [3]. However, the robot team is often observing a field (e.g. concentration or temperature) resulting from the transport of a source under some physical process [4]. For example, in hydrothermal vent identification, autonomous surface vehicles (ASVs) can directly detect the presence of “tracer” components on the ocean surface, but not the presence of an undersea vent [5], e.g., Fig. 1. In this work, we address the problem of source localization through indirect measurements of a source field in the presence of a complex flow.

The source localization problem is common in nature [6]; the behavior of source-seeking moths inspired the infotaxis algorithm [7], where sensing locations are selected by greedily maximizing the expected informativeness for the source location, under a known prior model. The method

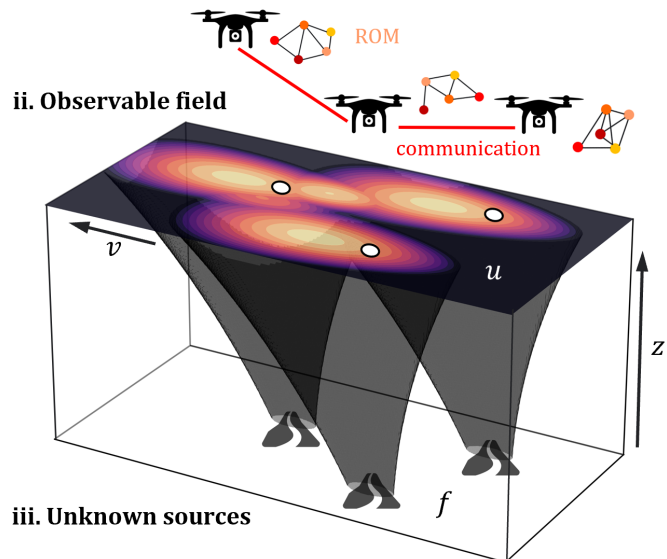
*M. A. Hsieh’s work is supported by NSF Award 2121887. B. Shaffer’s work is supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-2236662. N. Trask’s work is supported by SEA-CROGS (Scalable, Efficient and Accelerated Causal Reasoning Operators, Graphs and Spikes for Earth and Embedded Systems), a Mathematical Multifaceted Integrated Capability Center (MMICCs) and the Department of Energy early career program.

¹ Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, PA


² Interdisciplinary Computing, College of the Atlantic, Bar Harbor, ME

† Corresponding author, B. Shaffer: ben31@seas.upenn.edu

i. Mobile sensor network with Reduced Order Model



ii. Observable field



iv. Source estimate

Fig. 1: We consider a networked team of robots (i.) equipped with sensors measuring an observable field (ii.) resulting from the transport of an unobservable source (iii.). We equip each robot with a physics based reduced order model (i.) to perform source inference which are suitable for computing an estimate of the informativeness of potential sensing locations. The robots iteratively update their estimate and move to areas of higher sensing value to give a well resolved source estimate (iv.).

has been extended to robotics [8], in a distributed setting [9], for multiple sources [10], and in continuous time by treating the information as a potential field [11]. The key limitation of these approaches is the assumption of a *prior* environmental models [12], and explicit consideration of a single source. Solving complex fluid dynamics requires substantial computational resources and is often not feasible on resource-constrained robotic platforms.

To address this challenge, we utilize a structure-preserving machine learning framework [13], [14], which serves as an adaptable, real-time environmental model, while maintaining

key properties of the source transport dynamics (conservation, source to field coupling) [13]. Structure preserving methods in machine learning are techniques which exactly preserve physical properties, for instance symmetries, invariances, or conservations by construction. The Conditional Neural Whitney Form (CNWF) method we use combines the flexibility of transformer-based operator learning and the numerical properties of traditional numerical methods and guarantees global conservation. The conservation framing captures the relation between the observable concentration field and the source location through a learnable flux. We use this model to derive an approximate mutual information metric, which can be computed locally onboard robots. This information metric captures the expected reduction in uncertainty of the source locations due to a measurement at a given location. To control the robot team, we employ an infotaxis-style information-gradient ascent controller [11]. Our approach builds on [9] to enable a distributed implementation with inter-robot communication. We model the source as a density which naturally allows for extension to an arbitrary number of local sources, or more complex distributed source.

We demonstrate that our approach produces more interpretable and effective sensing strategies than baseline sampling approaches in simulated settings. More broadly, we view the integration of principled reduced order modeling and robot control as an example of onboard physical reasoning and highlight the future potential of well-designed surrogates for robot applications. To that end, our contributions are:

- A structure preserving reduced order environment model for source localization;
- An infotaxis-style information-based acquisition framework built around the learned environment model; and
- A distributed implementation using onboard computation.

II. RELATED WORK

A common simplifying assumption in the source localization problem is that the source location coincides with the maximal value of the observable concentration field. This results in gradient-based [15], [16], gradient approximation [17], and non-gradient based-methods for robot allocation [18], [19]. Factor graphs have also been used as the underlying environment model [20]. These approaches do not typically: address the effect of complex flows, provide direct estimates of the source term, and may not converge to the source, for instance if the concentration field is advected downwind, as in the hydrothermal vent example.

There are several recent approaches that incorporate explicit effects of advection-diffusion on the source signal. For example, robots teams perform plume tracking using infotaxis-style solutions with different formulations [9], [21], [22]. Various forms of Kalman filtering are combined with formation control to model spatiotemporal fields [23], [24] or perform level curve tracking [25]. Proper Orthogonal Decomposition (POD) has been used to model the field and information theoretic planning for robot team sampling [26]. These works all make strict model assumptions for the underlying spatiotemporal process. Our method is flexible

to general nonlinear transport and learned from data while preserving the necessary source structure.

Machine learning presents general methods to improve environmental monitoring [27]. Methods such as Physics Informed Neural Networks (PINNs) [28] or operator learning [29], [30] cover a spectrum from incorporating explicit physics to purely data driven approaches and have been implemented for various inverse problems including source estimation [31]–[36]. However, these methods capture physical structure weakly (by penalty) if at all, and the coupling between the sampled field and the source term must be implicitly deduced from training [37], [38].

Structure preserving machine learning has emerged as methods which incorporate physical or geometric structure directly, by construction. Broadly, these techniques can be understood as strictly enforcing a more fundamental level of physical properties (as opposed to particular equations as in PINNs), and have been shown to mitigate weaknesses of purely data driven approaches, particularly with data scarcity. We build on recent methods designed to offer conservation by design, interpretable structure, and principled integration with mesh-based physical domains [39]–[41].

III. PROBLEM STATEMENT

We consider a team of N mobile robots operating in a bounded workspace $\Omega \subset \mathbb{R}^d$ with kinematics given by

$$\dot{x}_i = c_i, \quad x_i \in \Omega,$$

where $i = 1, \dots, N$ denotes the i^{th} robot and c_i is the control input. Each robot carries a sensor measuring a scalar field $u : \Omega \rightarrow \mathbb{R}$ with additive noise:

$$d_i = u(x_i) + \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, \sigma^2).$$

The field u results from the distortion and displacement of an unknown source distribution $f : \Omega \rightarrow \mathbb{R}$ via a background flow. This process is defined as the forward operator $u = G(f)$. The set of measurements at a given time is $\mathcal{D} = \{(x_i, d_i)\}_{i=1}^N$. The source estimation task is to identify f from \mathcal{D} while actively choosing robot motions to collect informative new measurements. This is also known as the *inverse source identification* problem, and the operator $f = S(\mathcal{D})$ is known as the inverse operator.

We aim to design a control policy

$$\dot{x}_i = \mathcal{P}(x_i, \mathcal{D})$$

that drives the robots to collect data that enables accurate recovery of the unknown source field f quickly and efficiently. Doing so requires knowledge of both the forward and inverse process, where S is used to estimate the source from current or potential measurements, and G is needed to estimate the measurement along a potential robot trajectory for path planning.




Fig. 2: Method overview. We aim to infer the location of unobservable sources via targeted observations of the resulting scalar field, with a physics-based surrogate model driving the information based robot control strategy (shown in a.). The surrogate model is built on Conditional Neural Whitney Forms to provide a physically interpretable source model, discrete conservation by construction, and a geometric basis for handling complex geometries, in b.

IV. METHODS

We adopt an infotaxis-based approach to define the robot control policy. At each iteration, we score a candidate measurement location x by its mutual information with the source field:

$$\mathcal{U}(x) := I[f; u(x) | \mathcal{D}].$$

Finding paths for all robots to maximize the total information of future samples is a combinatorially hard collaboration problem. As a computable proxy, we consider information gradient ascent infotaxis [11], [42] as a method for directing robots to locally informative regions. We, therefore, define the control policy as a gradient ascent over the information density

$$\mathcal{P}(x, \mathcal{D}) = \frac{1}{\gamma} \nabla \mathcal{U}(x, \mathcal{D})$$

where γ is a friction coefficient, defined as $\gamma = 1/\|\nabla \mathcal{U}\|$.

Computing $\mathcal{U}(x)$ requires two key model components: (i) a forward operator G to propagate candidate sources f to predicted measurements u , and (ii) an inverse operator S to infer a posterior distribution over f from data \mathcal{D} . As exact PDE-based implementations of G and S are computationally prohibitive for online planning, we learn reduced order models $\hat{u} = G_\theta(f)$ and $\hat{f} = S_\theta(\mathcal{D})$ from data to enable real-time evaluation of $\mathcal{U}(x)$ and closed-loop control. This process is shown in Figure 2 (a).

A. Bayesian Active Learning by Disagreement

In order to evaluate the mutual information objective, we express it as the expected posterior entropy, and rearrange as [43]

$$\mathcal{U}(x) = \mathbb{H}[u | x, \mathcal{D}] - \mathbb{E}_{f \sim p(f|\mathcal{D})} [\mathbb{H}[u | x, f]], \quad (1)$$

which can be interpreted as seeking the x for which the source prediction under the posterior leads to the most disagreement through the forward model and is known as Bayesian active learning by disagreement (BALD).

To compute the mutual information objective, we maintain a posterior over the sources $p(f | \mathcal{D}) = S_\theta(\mathcal{D})$ and conditional forward $p(u | f) = G_\theta(f)$. The composition of the inverse model and the forward operator defines the induced predictive posterior $p(u | \mathcal{D}) = G_\theta \circ S_\theta(\mathcal{D})$. We utilize dropout on the neural network models to obtain an approximately Bayesian model. While this is an imperfect method [44], it has significant practical value in that it has theoretical backing [45] and is very cheap to implement for inference in a batched context. In this sense, we prioritize compatibility with robot deployment and representation flexibility over more limited models that maintain a fully Bayesian posterior (Gaussian process), or ensemble methods which are expensive to evaluate in real time.

We approximate the BALD acquisition utility for a candidate location $x \in \Omega$ using Monte Carlo sampling from the learned probabilistic inverse and forward models. An example evaluation of \mathcal{U} is shown in Figure 3.

B. Structure Preserving Reduced Order Model

Using the recently developed Conditional Neural Whitney Forms (CNWF) method [13] we define the forward operator G_θ and inverse operator S_θ needed for robot motion planning and source estimation. This approach is built on finite element methods to provide an adaptable PDE model with exact conservation, conditioned on real-time sensor measurements [14]. An overview of this architecture is shown in Figure 2 (b).

There are four learnable components. First, an encoder,

$$z = \mathcal{E}_\theta(\mathcal{D}),$$

which maps sensor readings to a latent representation. Then three networks conditioned on z , a *source prediction* model, $\hat{f} = S_\theta(z)$, a *flux* model, $F = \mathcal{F}_\theta(z, u)$, which describes the differential operator, and an *adaptable reduced basis*, $W = W_\theta(z)$, which all quantities are represented on. Together,




Fig. 3: We demonstrate the evaluation of \mathcal{U} for a representative configuration. Arrows indicate the direction of the information gradient at each robot location (red) which is used in the local and distributed implementation. The information based acquisition function identifies upstream and unobserved locations as the most informative, matching expectations. The observable scalar field is shown in a contour.

these describe a discrete conservation of the form

$$\mathcal{H}(u, z) := \nabla \cdot \mathcal{F}_\theta(u, z) - S_\theta(z) = 0. \quad (2)$$

where the encoded observations, z , act to calibrate; the learned differential operator, adaptive basis, and source term in real-time. This results in a PDE model which best matches the observed data. Through a careful discretization using finite element exterior calculus (FEEC) (a framework developed for electromagnetism) we can ensure this conservation is satisfied exactly [46]. The benefits of this framework for our application are

- The reduced basis allows each robot to evaluate their **lightweight finite element model**.
- The learned flux **ouples scalar field observations to the source distribution** during training.
- The PDE operator defines a forward map, G_θ , for computing information objective.
- The learned source term is **physically interpretable**.
- Offline training allows models to be preloaded onto robots.
- Flexible online adaptation via conditioning allows for **distributed implementation**.

For a given encoded observation input z , the models are evaluated and (2) is solved using a standard nonlinear FEM Newton solver to yield the solution \hat{u} [13]. We identify this forward solve as the parameterized forward operator,

$$G(f) := \hat{u} \quad \text{such that} \quad \mathcal{F}(\hat{u}, z) = f$$

(a solve of the learned PDE) and the source model f_θ as S_θ .

1) *Training procedure:* We establish an offline-training, online-calibration procedure. Prior to use we assume each robot is equipped with a fully trained model (consisting

of each trainable component), which is trained offline on representative data on the same domain. Online we perform inference by calibrating the learned PDE to sensor measurements and evaluate this model to generate the mutual information metric.

We define the offline, supervised training process. Given a training set

$$\mathcal{D}_{\text{train}} = \{(\hat{\mathcal{D}}^{(j)}, f^{(j)}, u^{(j)})\}_{j=1}^N,$$

where $\hat{\mathcal{D}}^{(j)}$ are synthetic observations derived from $u^{(j)}$, we jointly learn $G_\theta : f \mapsto \hat{u}$ and $S_\theta : \mathcal{D} \mapsto \hat{f}$ by minimizing a reconstruction and consistency loss, subject to the learned constraint:

$$\min_{\theta} \|\hat{u} - u_{\text{true}}\|^2 + W_{2,\epsilon}^2(\hat{f}, f_{\text{true}}) \quad (3)$$

$$\text{subject to } \mathcal{H}(\hat{u}, \mathcal{D}) = 0, \quad (4)$$

where $W_{2,\epsilon}^2$ is the Wasserstein transport metric. This is a PDE constrained optimization where \hat{u} is defined implicitly through solving the constraint, ensuring conservation is satisfied exactly, independent of data scarcity.

2) *Implementation details:* CNWFs are a mesh-based method and therefore we assume access to a discretization of the domain Ω^h , under basis functions $\{\varphi_i\}_{i=1}^{n_f}$, where n_f is the discrete problem dimension. Following [14], we use an encoding self-attention transformer model to provide a permutation invariant and rank adaptable representation of the sensing inputs, $z = \mathcal{E}_\theta(\mathcal{D})$. We define the parameterized flux in (2) as a combination of a discrete Laplacian and a learned flux, $\mathcal{N}_\theta(u, z)$, and a parameterized source as $S_\theta(z)$ and obtain the learned PDE form:

$$-\eta \Delta u + \nabla \cdot \mathcal{N}_\theta(u, z) = S_\theta(z) \quad \text{in } \Omega, \quad u|_{\partial\Omega} = g, \quad (5)$$

where η is a user chosen parameter to ensure stability. We define coarse basis functions $\{\psi_j\}_{j=1}^{n_c}$ as convex combinations of the fine-scale basis $\varphi_i(x)$:

$$\psi_j(x) = \sum_{i=1}^{n_f} W_{ji}(z) \varphi_i(x), \quad (6)$$

where $W(z) \in \mathbb{R}^{n_c \times n_f}$ is a learned, nonnegative matrix. This exposes a fully connected graph structure, with the learned fluxes operating on the edges. All quantities are projected into this reduced basis. The system is shown to be well-posed under reasonable continuity and regularity assumptions [13].

C. CNWF-BALD

Combining the CNWF model defined in Section IV-B and the mutual information computation in Section IV-A, we define our robot control policy using information gradient ascent. Algorithm 1 describes the combination of these methods. The spatial gradient can be readily computed on the finite element basis used for the CNWF model. In Section V-B we demonstrate that this method results in good robot placement under global waypoint control and local gradient ascent policies.

Algorithm 1 CNWF–BALD Active Sensing

Require: Pretrained CNWF model $(\mathcal{E}_\theta, S_\theta, G_\theta)$

Require: Initial robot locations $X^{(0)}$, initial measurement history $\mathcal{D}^{(0)}$, iteration count K , time step Δt

Require: Control policy $\mathcal{P}(\cdot)$

Ensure: Final source estimate \hat{f}_K and robot trajectories $\{X^{(k)}\}_{k=0}^K$

- 1: **for** $k = 0$ **to** $K - 1$ **do**
 - 2: **Sensing:** $d_i^{(k)}$ at $x_i^{(k)}$, update $\mathcal{D}^{(k+1)}$
 - 3: **Encode:** $z^{(k+1)} \leftarrow \mathcal{E}_\theta(\mathcal{D}^{(k+1)})$
 - 4: **Source estimate:** $\hat{f}_k \leftarrow S_\theta(z^{(k+1)})$
 - 5: **Compute MI:** $\mathcal{U}(x)$, using (1)
 - 6: $X^{(k+1)} \leftarrow \Delta t \mathcal{P}(\mathcal{U}, X^{(k)})$
 - 7: **end for**
 - 8: $\hat{f}_K \leftarrow S_\theta(\mathcal{E}_\theta(\mathcal{D}^{(K)}))$, final source estimate
 - 9: **return** $\hat{f}_K, \{X^{(k)}\}_{k=0}^K$
-

1) *Distributed implementation:* Centralized coordination is often infeasible in multi-robot settings due to limited bandwidth, latency, and single points of failure. We therefore use a fully distributed architecture in which each robot maintains its own local copy of the CNWF surrogate, performs local inference, and communicates information with neighbors. We perform one communication step per active sensing iteration, k .

Let $\mathcal{H}^{(k)}$ be the set of neighboring robots at iteration k . We define neighbors as all robots within a fixed communication radius r_c :

$$\mathcal{H}_i^{(k)} := \{j \mid \|x_i - x_j\|_2 < r_c, j = 1, \dots, N\}. \quad (7)$$

Each robot maintains a local history of known observations consisting of both collected and communicated data, $\mathcal{D}_i^{(k)}$. During a communication step, each robot shares its current measurement $\{d_i^{(k)}, x_i^{(k)}\}$, and its history of measurements $\mathcal{D}_i^{(k)}$, with all neighbors. Incoming communications are assimilated into each robot's history:

$$\mathcal{D}_i^{(k+1)} \leftarrow \mathcal{D}_i^{(k)} \bigcup_{j \in \mathcal{H}_i^{(k)}} \mathcal{D}_j^{(k)} \cup \{d_j^{(k)}, x_j^{(k)}\}. \quad (8)$$

We assume single-hop, synchronous communication which means information can only propagate one edge per time step, preserving causality in the network. Each robot then solves its own forward problem using the CNWF surrogate conditioned on $\mathcal{D}_i^{(k+1)}$ to obtain estimates $\hat{u}_i(x)$, $\hat{f}_i(x)$ and an information field $\mathcal{U}_i(x)$. This process is described in Algorithm 2

This protocol yields a monotone growth of each robot's information set over time, approaching the global dataset after sufficient communication steps while preserving strictly local, causal updates at each round. The approach is robust to intermittent communication failures and scales naturally to large robot teams [47], [48]. In Section V-D, we empirically verify that i.) measurements effectively propagate through the robot network over the adaptive sampling run, and ii.) the average local prediction with partial observation inputs

Algorithm 2 Distributed CNWF–BALD with Local History Fusion

Require: Pretrained CNWF model $(\mathcal{E}_\theta, S_\theta, G_\theta)$

Require: Initial robot locations $X^{(0)}$, initial measurement histories $\mathcal{D}_i^{(0)}$, iteration count K , communication radius r_c , time step Δt

Require: Local control policy $\mathcal{P}(\cdot)$

Ensure: Final source estimates \hat{f}_K and robot trajectories $\{X^{(k)}\}_{k=0}^K$

- 1: **for** $k = 0$ **to** $K - 1$ **do**
 - 2: **for** $i = 1$ **to** N **do**
 - 3: **Sensing:** $d_i^{(k)}$ at $x_i^{(k)}$
 - 4: **Establish Communication:** Update $\mathcal{H}_i^{(k)}$ by (7)
 - 5: **Communication:** Update $\mathcal{D}_i^{(k+1)}$ by (8)
 - 6: **Encode:** $z^{(k+1)} \leftarrow \mathcal{E}_\theta(\mathcal{D}_i^{(k+1)})$
 - 7: **Source estimate:** $\hat{f}_{k,i} \leftarrow S_\theta(z^{(k+1)})$
 - 8: **Compute MI:** $\mathcal{U}_i(x)$, using (1)
 - 9: **Control Update:** $x_i^{(k+1)} \leftarrow \Delta t \mathcal{P}(\mathcal{U}_i, x_i^{(k)})$
 - 10: **end for**
 - 11: **end for**
 - 12: $\hat{f}_{i,K} \leftarrow S_\theta(\mathcal{E}_\theta(\mathcal{D}_i^{(K)}))$, final local source estimates
 - 13: **return** $\hat{f}_{i,K}, \{X_i^{(k)}\}_{k=0}^K, i = 1, \dots, N$
-

converges to an idealized centralized prediction made with all measurements.

D. Experimental setup

For an example system we consider two dimensional steady-state advection-diffusion. While this problem is a linear PDE, for an unknown velocity, the corresponding inverse problem is nonlinear. The PDE is:

$$-\frac{1}{\text{Pe}} \nabla^2 u + \mathbf{v} \cdot \nabla u = f \quad \text{in } \Omega, \quad u|_{\partial\Omega} = g, \quad (9)$$

where $\text{Pe} > 0$ is the Peclet number, $\mathbf{v} : \Omega \rightarrow \mathbb{R}^d$ is a potentially unknown, divergence-free velocity field with $|\mathbf{v}| = 1$, g is a fixed boundary condition. To generate data we use a prescribed velocity and diffusivity. The system is solved on the discretized domain used in the CNWF model, Ω^h .

We consider two scenarios in this work, first a circular geometry with a constant velocity with a randomly selected orientation (randomly chosen for each trial), and second a complex geometry with a variable stokes-flow velocity profile based on two inlets and six possible outlets, any combination of which may be open for a given sample (chosen randomly). Both the discretization Ω^h and boundary conditions g are provided to initialize the CNWF model. We generate data by solving the advection diffusion equations on the problem domain using finite elements. For each solution, we consider $n \sim U[1, 5]$ sources randomly distributed throughout the domain. We use a constant Peclet number of $1e1$.

V. RESULTS

A. Source Reconstruction

We evaluate the ability of the CNWF reduced order model to reconstruct the source field f from sparse observations.

We compare our method against a transformer-based baseline trained to regress from sensor data to the source field without imposing any conservation structure. The baseline has an identical architecture to the CNWF encoder \mathcal{E}_θ and a prediction head, for fair comparison.

Table I reports the mean Wasserstein-2 distance between the predicted source field and ground truth over 100 random test instances. CNWF achieves significantly lower reconstruction error, indicating that embedding physical structure into the surrogate improves accuracy and generalization in data-sparse regimes.

Geometry	Model	$W_2(\rho_\theta^0, \rho_{true}) \downarrow$
Circle	CNWF (<i>ours</i>)	7.12e-3
	Transformer (<i>baseline</i>)	2.92e-2
Hallway with Rooms	CNWF (<i>ours</i>)	2.48e-1
	Transformer (<i>baseline</i>)	3.26e-1

TABLE I: Source prediction accuracy comparison between CNWF model and baseline transformer

B. Active sensing via mutual information

We compare our information-based robot control strategy against a naive variance maximizing baseline:

$$\mathcal{U}_{var}(x) = \sigma_u(x).$$

This is a popular method due to its simplicity to implement, which treats uncertainty as a proxy for informativeness, and does not address the inverse source identification problem, i.e. this approach only requires a model for the scalar field u . Most previous source seeking approaches (e.g. gradient or coverage-based) assume direct observation of the field of interest and do not provide direct comparison to our method which addresses the inverse source identification problem. For a global acquisition strategy, we simply select the optimal sensing location as $x^* = \operatorname{argmax}_x \mathcal{U}(x)$. We compare the information and variance based acquisition functions under the global waypoint robot control policy. Both methods use an identical CNWF model. We record the average source reconstruction error according to the Wasserstein metric over 8 iterations for 128 initial configurations in the circular domain in Figure 4. Our method consistently reduces the reconstruction error more rapidly than the variance based baseline, demonstrating the value of information-driven placement. The variance based sampling methods requires 5 iterations to improve over a single iteration of the BALD method.

1) *Deployment in complex geometries:* We demonstrate the ability of the CNWF model to recover the source in a closed loop prediction context for the complex geometry with the CNWF-BALD global acquisition function in Figure 5. Interestingly, while most sensors are allocated near the source, some are also placed in the other intersection in the hallway, which qualitatively would seem to be an informative location due to the flow variability.




Fig. 4: We compare the average source reconstruction error for 128 trials over K iterations using global sensor placement strategies for the \mathcal{U} and a naive variance based baseline $\operatorname{Var}(u)$ on the circular domain. Our approach results in more informative sensor placement which in turn substantially reduces source error compared to the baseline.




Fig. 5: The global placement strategy on the hall geometry results in accurate localization of the source from an ill-posed initial configuration. The color indicates the source predictions and the contour indicates the observable scalar field.

C. Local Exploration via Information Gradient Ascent

We also evaluated the local, information-gradient ascent implementation of the infotaxis algorithm. Figure 6 shows an example of this strategy for simultaneously localizing four sources. In only three adaptive sensing iterations the source reconstruction error was reduced by 86%. Qualitatively, the information-based control policy results in a combination of both exploration and exploitation like behaviors, where some robots move directly to a nearby source and others move towards unsampled regions of the domain. The final configuration shows a spatially distributed robot configuration, even without the inclusion of an explicit repulsion in the control. Empirically, this result holds over all cases where the sensors are initially distributed. The model successfully recovers all sources distributed over the domain with only 15 total sampling locations. While this local method does not guarantee globally optimal placement, it achieves competitive

source reconstruction with reduced sampling overhead and realistic constraints.

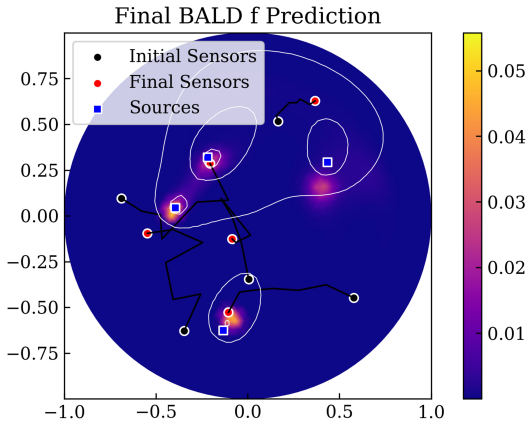


Fig. 6: A sample gradient ascent demonstration for four unknown sources. The background color indicates the predicted source term, the contour indicates the scalar field, and the black lines shows the robot trajectories. In the final configuration (after 6 steps) the CNWF model correctly identifies the presence of four sources and accurately places three of them. The gradient ascent results in well distributed sensors and qualitatively good sensing locations near sources.

D. Distributed Implementation

Our experiments in Figure 7 demonstrate that the distributed CNWF-BALD algorithm achieves performance comparable to the centralized planner. With some exceptions, the collection and exchange of measurements results in increasingly consistent predictions. However, isolated robots may fail to reconnect, a known pitfall of these types of collaboration strategies. We also observe that early measurements effectively propagate through the sensor network in Figure 8.




Fig. 7: We compare the prediction error for each agent in the distributed implementation with an idealized centralized model. While initially the sparse communication results in degraded performance for the distributed implementation, error greatly improves as information propagates the sensor network.




Fig. 8: At the final iteration K , we evaluate the density of information from each previous active sampling step $-k$, that has reached each agent. While recent measurements $k = 0$ are only about 50% dense, while nearly all observations made at time $k = -9$, have reached all robots at the current time. This shows the effectiveness of our simple communication protocol for collating observations.

VI. CONCLUSIONS

We presented a framework for distributed multi-robot, multi-source localization in which each robot is equipped with a machine-learned reduced-order model of its environment. Leveraging the PDE-structured nature of these models, we derived an information-theoretic objective that drives informative path planning via gradient-based control. We further developed a communication and control strategy suitable for fully distributed operation. This approach unifies ideas from Bayesian optimization, PDE-constrained modeling, operator learning, and robotics, achieving improved source estimation and more informative sampling compared to baseline machine learning and planning methods. More broadly, equipping robots with onboard, structure-preserving finite element surrogates demonstrates the potential for robots to act as physically informed, model-driven agents in scientific exploration and sensing tasks.

REFERENCES

- [1] A. Daw, A. Karpatne, K. Yeo, and L. Klein, "Source identification and field reconstruction of advection-diffusion process from sparse sensor measurements." *Conference on Neural Information Processing Systems*, 2022.
- [2] Y. Hon, M. Li, and Y. A. Melnikov, "Inverse source identification by green's function," *Engineering Analysis with Boundary Elements*, vol. 34, no. 4, pp. 352–358, 2010.
- [3] Y. Sung, Z. Chen, J. Das, P. Tokekar *et al.*, "A survey of decision-theoretic approaches for robotic environmental monitoring," *Foundations and Trends® in Robotics*, vol. 11, no. 4, pp. 225–315, 2023.
- [4] M. B. Moghaddam, M. Mazaheri, and J. M. V. Samani, "Inverse modeling of contaminant transport for pollution source identification in surface and groundwaters: a review," *Groundwater for Sustainable Development*, vol. 15, p. 100651, 2021.
- [5] D. R. Yoerger, A. M. Bradley, M. Jakuba, C. R. German, T. Shank, and M. Tivey, "Autonomous and remotely operated vehicle technology for hydrothermal vent discovery, exploration, and sampling," *Oceanography*, vol. 20, no. 1, pp. 152–161, 2007.
- [6] G. Reddy, V. N. Murthy, and M. Vergassola, "Olfactory sensing and navigation in turbulent environments," *Annual Review of Condensed Matter Physics*, vol. 13, no. 1, pp. 191–213, 2022.
- [7] M. Vergassola, E. Villermaux, and B. I. Shraiman, "'infotaxis' as a strategy for searching without gradients," *Nature*, vol. 445, no. 7126, pp. 406–409, 2007.

- [8] E. Martin Moraud and D. Martinez, "Effectiveness and robustness of robot infotaxis for searching in dilute conditions," *Frontiers in neurorobotics*, vol. 4, p. 1213, 2010.
- [9] H. Hajieghrary, D. Mox, and M. A. Hsieh, "Information theoretic source seeking strategies for multiagent plume tracking in turbulent fields," *Journal of Marine Science and Engineering*, vol. 5, no. 1, p. 3, 2017.
- [10] B. Chen, X. Zhang, Y. Ji, Y. Zhao, and Z. Zhu, "Dlw-ci: A dynamic likelihood-weighted cooperative infotaxis approach for multi-drone cooperative multi-source search," *Journal of Safety Science and Resilience*, 2025.
- [11] C. Barbieri, S. Cocco, and R. Monasson, "On the trajectories and performance of infotaxis, an information-based greedy search algorithm," *EuroPhysics letters*, vol. 94, no. 2, p. 20005, 2011.
- [12] J. D. Rodríguez, D. Gómez-Ullate, and C. Mejía-Monasterio, "Limits on the performance of infotaxis under inaccurate modelling of the environment," *arXiv preprint arXiv:1408.1873*, 2014.
- [13] B. Kinch, B. Shaffer, E. Armstrong, M. Meehan, J. Hewson, and N. Trask, "Structure-preserving digital twins via conditional neural whitney forms," 2025. [Online]. Available: <https://arxiv.org/abs/2508.06981>
- [14] B. D. Shaffer, B. Kinch, J. Klobusicky, M. A. Hsieh, and N. Trask, "Physics-informed sensor coverage through structure preserving machine learning," 2025. [Online]. Available: <https://arxiv.org/abs/2509.10363>
- [15] P. Ogren, E. Fiorelli, and N. E. Leonard, "Cooperative control of mobile sensor networks: Adaptive gradient climbing in a distributed environment," *IEEE Transactions on Automatic control*, vol. 49, no. 8, pp. 1292–1302, 2004.
- [16] D. Talwar, T. Lu, D. Sobti, and W. Wu, "Multi-robot source seeking and field reconstruction of spatial-temporal varying fields," in *2024 9th International Conference on Automation, Control and Robotics Engineering (CACRE)*, 2024, pp. 101–105.
- [17] N. Atanasov, J. Le Ny, N. Michael, and G. J. Pappas, "Stochastic source seeking in complex environments," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 3013–3018.
- [18] S. Al-Abri and F. Zhang, "A distributed active perception strategy for source seeking and level curve tracking," *IEEE Transactions on Automatic Control*, vol. 67, no. 5, pp. 2459–2465, 2021.
- [19] T. Zhang, V. Qin, Y. Tang, and N. Li, "Distributed information-based source seeking," *IEEE Transactions on Robotics*, vol. 39, no. 6, pp. 4749–4767, 2023.
- [20] C. E. Denniston, O. Peltzer, J. Ott, S. Moon, S.-K. Kim, G. S. Sukhatme, M. J. Kochenderfer, M. Schwager, and A.-a. Agha-mohammadi, "Fast and scalable signal inference for active robotic source seeking," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 7909–7915.
- [21] J.-W. Wang, Y. Guo, M. Fahad, and B. Bingham, "Dynamic plume tracking by cooperative robots," *IEEE/ASME Transactions on Mechatronics*, vol. 24, no. 2, pp. 609–620, 2019.
- [22] S. H. Singh, F. van Bruegel, R. P. Rao, and B. W. Brunton, "Emergent behaviour and neural dynamics in artificial agents tracking odour plumes," *Nature machine intelligence*, vol. 5, no. 1, pp. 58–70, 2023.
- [23] Z. Zhang, S. T. Mayberry, W. Wu, and F. Zhang, "Distributed cooperative kalman filter constrained by discretized poisson equation for mobile sensor networks," in *2023 American Control Conference (ACC)*, 2023, pp. 1365–1370.
- [24] S. Mayberry, Z. Zhang, W. Wu, and F. Zhang, "Soft-constrained distributed cascaded cooperative kalman filter for mobile robots in unknown advection-diffusion field," *IEEE Robotics and Automation Letters*, 2025.
- [25] S. Guruswamy and W. Wu, "Cooperative level curve tracking in advection-diffusion fields," in *2020 5th International Conference on Automation, Control and Robotics Engineering (CACRE)*, 2020, pp. 434–438.
- [26] R. Khodayi-mehr, W. Aquino, and M. M. Zavlanos, "Model-based active source identification in complex environments," *IEEE Transactions on Robotics*, vol. 35, no. 3, pp. 633–652, 2019.
- [27] R. Khodayi-mehr and M. M. Zavlanos, "Physics-based learning for robotic environmental sensing," *arXiv preprint arXiv:1812.03894*, 2018.
- [28] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational physics*, vol. 378, pp. 686–707, 2019.
- [29] L. Lu, P. Jin, G. Pang, Z. Zhang, and G. E. Karniadakis, "Learning nonlinear operators via deeponet based on the universal approximation theorem of operators," *Nature machine intelligence*, vol. 3, no. 3, pp. 218–229, 2021.
- [30] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar, "Fourier neural operator for parametric partial differential equations," *arXiv preprint arXiv:2010.08895*, 2020.
- [31] M. V. de Hoop, M. Lassas, and C. A. Wong, "Deep learning architectures for nonlinear operator functions and nonlinear inverse problems," *Mathematical Statistics and Learning*, vol. 4, no. 1, pp. 1–86, 2022.
- [32] S. Mishra and R. Molinaro, "Estimates on the generalization error of physics-informed neural networks for approximating a class of inverse problems for pdes," *IMA Journal of Numerical Analysis*, vol. 42, no. 2, pp. 981–1022, 2022.
- [33] S. Kamyab, Z. Azimifar, R. Sabzi, and P. Fieguth, "Deep learning methods for inverse problems," *PeerJ Computer Science*, vol. 8, p. e951, 2022.
- [34] V. V. Vesselinov, B. S. Alexandrov, and D. O'Malley, "Contaminant source identification using semi-supervised machine learning," *Journal of contaminant hydrology*, vol. 212, pp. 134–142, 2018.
- [35] Y. N. Kontos, T. Kassandros, K. Perifanos, M. Karampasis, K. L. Katsifarakis, and K. Karatzas, "Machine learning for groundwater pollution source identification and monitoring network optimization," *Neural Computing and Applications*, vol. 34, no. 22, pp. 19 515–19 545, 2022.
- [36] A. Reiter and M. A. L. Bell, "A machine learning approach to identifying point source locations in photoacoustic data," in *Photons Plus Ultrasound: Imaging and Sensing 2017*, vol. 10064. SPIE, 2017, pp. 504–509.
- [37] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nature Reviews Physics*, vol. 3, no. 6, pp. 422–440, 2021.
- [38] S. Wang, H. Zhang, and X. Jiang, "Physics-informed neural network algorithm for solving forward and inverse problems of variable-order space-fractional advection–diffusion equations," *Neurocomputing*, vol. 535, pp. 64–82, 2023.
- [39] N. Trask, A. Huang, and X. Hu, "Enforcing exact physics in scientific machine learning: a data-driven exterior calculus on graphs," *Journal of Computational Physics*, vol. 456, p. 110969, 2022.
- [40] J. A. Actor, X. Hu, A. Huang, S. A. Roberts, and N. Trask, "Data-driven whitney forms for structure-preserving control volume analysis," *Journal of Computational Physics*, vol. 496, p. 112520, 2024.
- [41] S. Jiang, J. Actor, S. Roberts, and N. Trask, "A structure-preserving domain decomposition method for data-driven modeling," *arXiv preprint arXiv:2406.05571*, 2024.
- [42] B. J. Julian, S. Karaman, and D. Rus, "On mutual information-based control of range sensing robots for mapping applications," *The International Journal of Robotics Research*, vol. 33, no. 10, pp. 1375–1392, 2014.
- [43] N. Houlsby, F. Huszár, Z. Ghahramani, and M. Lengyel, "Bayesian active learning for classification and preference learning," *arXiv preprint arXiv:1112.5745*, 2011.
- [44] Y. Liu, M. Pagliardini, T. Chavdarova, and S. U. Stich, "The peril of popular deep learning uncertainty estimation methods," *arXiv preprint arXiv:2112.05000*, 2021.
- [45] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *international conference on machine learning*. PMLR, 2016, pp. 1050–1059.
- [46] D. N. Arnold, *Finite element exterior calculus*. SIAM, 2018.
- [47] A. G. Dimakis, S. Kar, J. M. Moura, M. G. Rabbat, and A. Scaglione, "Gossip algorithms for distributed signal processing," *Proceedings of the IEEE*, vol. 98, no. 11, pp. 1847–1864, 2010.
- [48] F. Bullo, J. Cortés, and S. Martinez, *Distributed control of robotic networks: a mathematical approach to motion coordination algorithms*. Princeton University Press, 2009.