

Localization in Spatiotemporal Fields via Environmental PDEs

Abstract—This paper proposes a localization framework that uses spatiotemporal fields governed by partial differential equations (PDEs) as localization signatures. Two PDE classes are considered: the shallow water equations, which describe free-surface flows in coastal and riverine environments, and the advection-diffusion equation, which models the transport and mixing of scalar quantities such as temperature, salinity, and dissolved oxygen. A numerical PDE solver provides predicted fields over the domain, and multiple field channels are fused as multimodal measurements to improve localization accuracy. We formulate the problem within a Rao-Blackwellized particle filter (RBPF) that partitions the vehicle state into a nonlinear component sampled by particles and a linear sensor bias component tracked analytically via per-particle Kalman filters. This factorization reduces the required number of particles compared to a standard particle filter while accounting for realistic sensor drift. Simulation studies on both PDE scenarios show that the RBPF consistently outperforms a standard particle filter in terms of final position error and Root Mean Square Error (RMSE) across varying particle counts. Field experiments with an autonomous surface vehicle measuring salinity, temperature, and dissolved oxygen validate that PDE-governed environmental fields provide sufficient spatial variability for practical localization. Related experimental videos are available at <https://localization-environmental-pdes.github.io/>, with the open-source code to be released upon publication.

I. INTRODUCTION

Localization in GPS-denied environments is a fundamental challenge for autonomous vehicles operating underwater, indoors, or in contested electromagnetic settings. In such scenarios, the vehicle cannot rely on satellite navigation and must instead exploit alternative information sources. Prior work has considered magnetic field anomalies [1, 2], radio signal strength [3], and visual landmarks [4] as substitutes. A common thread in these approaches is that a spatially varying physical field is matched against onboard sensor readings to constrain the vehicle position. However, many physical fields of practical interest, such as those arising in ocean surveying, pollution tracking, and hazard detection [5], exhibit significant temporal dynamics in addition to spatial structure. Relying on a purely spatial representation ignores temporal correlations that, if modeled, could sharpen localization estimates and extend the range of usable environmental signatures. A spatiotemporal modeling framework is therefore required to capture both the spatial variability and the temporal evolution of such fields.

A small but growing body of work has begun to address localization and mapping in fluid flow environments. In [6, 7] and [8], flow-based localization and mapping frameworks were introduced that exploit flow field structure for simultaneous state and map estimation. However, these methods are

restricted to steady or slowly varying flows, or to specific dynamical models amenable to data-driven decomposition. They cannot accommodate the full range of spatiotemporal fields encountered in environmental applications, particularly those governed by the advection-diffusion equation or the shallow water equations, which exhibit complex wave propagation, diffusive spreading, and nonlinear coupling.

A separate line of research has investigated flow field estimation, adaptive navigation, and source identification in environmental scalar fields. In [9], an active learning framework for unknown flows was proposed using Koopman operators and ensemble methods. In [10], multi-robot dynamic coverage path planning with workload adaptation was addressed. In [11], partial differential equation (PDE) based observers are used to estimate underwater vehicle trajectories. In [12], a learning-based adaptive navigation framework for scalar field mapping was introduced, employing the advection-diffusion PDE and Gaussian Process Regression to estimate fields and their gradients in real time. In [13], a probabilistic framework for multi-robot gas source localization was developed, modeling gas dispersion with the 2D diffusion PDE and using factor graph inference with Sparse Bayesian Learning to estimate source locations and magnitudes. While these works demonstrate the value of PDE-based environmental models for estimation and planning tasks, they all focus on flow modeling, coverage optimization, or source parameter estimation. None of them addresses the inverse problem of localizing the vehicles themselves, given field measurements in those realistic environments.

In the context of vehicle localization, the Rao-Blackwellized particle filter (RBPF) exploits conditional independence structure to factorize the posterior distribution into components estimated by particles and components updated analytically, yielding dramatic reductions in the required number of particles [14]. This technique has been successfully applied across diverse domains, yet it has not been extended to spatiotemporal field-based localization. In [15] developed RBPF-based particle smoothing for SLAM with conditionally linear or linearized models. However, the maps considered are static, and the approach does not address spatiotemporal dynamics governed by environmental PDEs. In [16] applied RBPF to connected vehicle localization but relies on GNSS signal measurements rather than environmental field observations and does not consider spatiotemporal field dynamics. In [17] employed RBPF for 6-DOF localization in robotic peg-in-hole assembly with a static geometric map, bearing no connection to environmental fields.

Our results apply directly to underwater localization, which remains a major challenge. It often relies on costly

sensing modalities that are not universally reliable, such as Doppler Velocity Loggers [18], acoustic positioning infrastructures, or acoustic and optical imaging systems that demand substantial onboard computation and that can fail due to lighting conditions and water turbidity conditions [19, 20], and sonar-based sensors can suffer from low resolution.

Many spatiotemporal environmental phenomena of practical interest are governed by PDEs. The shallow water equation (SWE) describes water flows in coastal, estuarine, and riverine environments. The advection-diffusion equation describes the transport and mixing of scalar quantities such as temperature, salinity, dissolved oxygen, and chemical concentration in water bodies. In each case, the PDE produces fields that naturally vary in space and time, and that can be predicted in future steps with current information. To address the challenges identified above, the main contributions of this paper are as follows:

- A general localization framework that uses environmental fields governed by PDEs as localization signatures for GPS-denied operation. Two PDE classes are formulated: the shallow water equation and the advection-diffusion equation.
- An RBPF formulation that exploits the conditionally linear structure of the PDE measurement model.
- Simulation studies across two PDE scenarios: *i*) coastal wave propagation (SWE) and *ii*) water quality modeling (advection-diffusion), demonstrate that the RBPF consistently outperforms the standard PF.
- Real-world validation using an autonomous surface vehicle measuring salinity, temperature, and dissolved oxygen, demonstrating that PDE-governed fields provide sufficient spatial variability for practical localization.

II. PROBLEM FORMULATION

We consider a vehicle navigating a two-dimensional domain $\mathcal{W} \subseteq \mathbb{R}^2$ that is characterized by M environmental scalar and vector fields governed by PDEs; these fields represent physical quantities such as water velocity, temperature, dissolved oxygen, or water height, and are collectively denoted by $\phi : \mathcal{W} \times \mathbb{R}^+ \rightarrow \mathbb{R}^M$.

The vehicle state at time t is $\mathbf{x}_t \in \mathcal{X}$, where $\mathcal{X}$ is the state space. The vehicle moves through $\mathcal{W}$ according to a kinematic motion model $\mathbf{f} : \mathcal{X} \times \mathcal{U} \rightarrow \mathcal{X}$, driven by a control input $\mathbf{u}_t$ from the set of admissible controls $\mathcal{U}$:

$$\mathbf{x}_{t+1} = \mathbf{f}(\mathbf{x}_t, \mathbf{u}_t). \quad (1)$$

The projection $\Pi : \mathcal{X} \rightarrow \mathcal{W}$ maps the full state to the vehicle position $\mathbf{p} = \Pi(\mathbf{x})$ in the domain. Moreover, the vehicle is equipped with M onboard sensors, each measuring a corresponding channel of the environmental field ϕ . The observation model $\mathbf{h} : \mathcal{X} \rightarrow \mathcal{Y}$ maps the vehicle state to the observation space $\mathcal{Y}$ as

$$\mathbf{h}(\mathbf{x}_t) = \phi(\Pi(\mathbf{x}_t), t) + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(\mathbf{0}, \Sigma), \quad (2)$$

where ε_t is zero-mean Gaussian measurement noise. Each observation therefore, provides a multimodal measurement

vector that simultaneously captures multiple PDE-governed field quantities at the vehicle position.

Problem 1: Given the motion model (1), the multimodal observation model (2), and the PDE-predicted environmental fields ϕ , the objective is to recursively estimate the vehicle state $\mathbf{x}_t$ at each time step t , conditioned on the history of observations $\mathbf{y}_{1:t}$ and the sequence of control inputs $\mathbf{u}_{1:t-1}$. In a nutshell, how to approximate the posterior distribution $p(\mathbf{x}_t \mid \mathbf{y}_{1:t}, \mathbf{u}_{1:t-1})$ by exploiting the spatiotemporal structure of ϕ as a dynamic localization reference?

III. MODELING ENVIRONMENTAL PHENOMENA UTILIZING PDES

The solution of PDEs yields a set of spatially varying field quantities $\phi = (\phi_1(\mathbf{x}, t), \dots, \phi_M(\mathbf{x}, t))^T$ over our domain of interest $\mathcal{W}$. The vehicle is equipped with M onboard sensors that measure the local values of these fields at its current position $\mathbf{x} \in \mathcal{W}$. Because these fields exhibit distinct spatial variation, the resulting sensor readings provide sufficient information to constrain the vehicle's position, facilitating localization without the necessity of GPS.

This formulation remains general; while different PDE classes model various physical phenomena, the underlying localization structure is consistent. We focus on two specific instances relevant to maritime and coastal environments.

1) Shallow Water Equation (SWE): The shallow water equations model free-surface flows in coastal and riverine environments where the horizontal scale is much larger than the depth; it presents a useful framework for coastal or river tidal modeling. This model is an abstraction of the Navier-Stokes equations for fluid dynamics. It assumes that, at a given spatial coordinate $\mathbf{x}$, the column of water is moving as a single entity by averaging the velocities along it. This model represents conserved variables: water mass $\rho\eta : \mathcal{W} \times \mathbb{R}^+ \rightarrow \mathbb{R}^+$, and its momentum $\rho\eta\mathbf{v}$. Here, ρ is the fluid's density, $\mathbf{v} : \mathcal{W} \times \mathbb{R}^+ \rightarrow \mathbb{R}^2$ groups the velocities in the x and y directions, respectively. The SWE, in its conservative form, can be expressed in its block notation:

$$\begin{aligned} \frac{\partial \phi}{\partial t} + \nabla \cdot F(\phi) &= \mathbf{s}(\mathbf{x}, t, \phi), \text{ for } (\mathbf{x}, t) \in \mathcal{W} \times \mathbb{R} \\ F(\phi) &= \eta \begin{bmatrix} \mathbf{v}^T \\ \mathbf{v} \otimes \mathbf{v} + \frac{1}{2} g \eta I \end{bmatrix} \\ \phi &= \phi_0, \text{ for } t = 0, \end{aligned} \quad (3)$$

where $\otimes$ is the exterior product between two vectors, g represents gravitational acceleration, I is the 2×2 identity matrix, ϕ_0 is the initial condition encompassing initial water height and velocity, and $\mathbf{s}$ denotes the forcing term increasing or decreasing any variable. It is worth noting that even though $\mathbf{s}$ still depends on the conserved variables, it is considered a forcing term due to its nature of incorporating external information into the model. Furthermore, we purposely included the variables that $\mathbf{s}$ depend on, this will facilitate introducing typical forcing terms that are presented in Section V. On the other hand, if we assume the fluid is incompressible -as water is close to being- its density can be removed

from the equation, simplifying it. In this context, the vehicle measurement vector is $\phi = (\eta, u, v)^\top$, corresponding to $M = 3$ sensing channels.

2) *Advection-Diffusion Equation*: To model the transport of scalar water properties such as temperature, salinity, chlorophyll, or chemical concentrations, we utilize the advection-diffusion equation; it balances the behavior when the main evolution of such properties is determined by transportation and dissolution factors. For a scalar field $\phi(\mathbf{x}, t)$, the dynamics are governed by:

$$\begin{aligned} \frac{\partial \phi}{\partial t} + \nabla \cdot (\phi \mathbf{v} - \kappa \nabla \phi) &= s, \text{ for } (\mathbf{x}, t) \in \mathcal{W} \times \mathbb{R} \\ \phi &= \phi_0, \text{ for } t = 0, \end{aligned} \quad (4)$$

where $\mathbf{v}$ is the water velocity field, κ is the diffusivity term, and s represents the source term. The diffusion term $-\kappa \nabla \phi$ accounts for the diffusion of the property from regions of high concentration to the regions with low concentration, while the advection term $\phi \mathbf{v}$ describes its transport by the fluid flow.

In oceanographic applications, distinct scalar fields ϕ_i , $i = 1, \dots, M$, that satisfy equation (4) may be coupled in the same region of interest with specific coefficients and source terms by the same velocity field. Similar to the SWE, by vectorizing $\phi = (\phi_1, \dots, \phi_M)^\top$, $\mathbf{s} = (s_1, \dots, s_M)^\top$, $\phi_0 = (\phi_{01}, \dots, \phi_{0M})^\top$, $\boldsymbol{\kappa} = (\kappa_1, \dots, \kappa_M)^\top$ and defining $\text{diag}(\mathbf{z})$ as the diagonal matrix whose entrances are filled with the elements of $\mathbf{z}$; the different features can be written as a coupled PDE as

$$\begin{aligned} \frac{\partial \phi}{\partial t} + \nabla \cdot (\phi \otimes \mathbf{v} - \text{diag}(\boldsymbol{\kappa}) \nabla \phi) &= \mathbf{s}, \text{ for } (\mathbf{x}, t) \in \mathcal{W} \times \mathbb{R} \\ \phi &= \phi_0, \text{ for } t = 0. \end{aligned} \quad (5)$$

In this case, the vehicle measurement vector is ϕ containing M channels.

IV. LOCALIZATION WITH RAO-BLACKWELLIZED PARTICLE FILTER (RBPF)

A. State Partition

Unlike standard particle filter which samples the full state and computes importance weights from the measurement likelihood, the Rao-Blackwellized particle filter (RBPF) exploits posterior factorization when the state admits a conditionally linear substructure, resulting in reducing the variance of the posterior estimate while requiring fewer particles [21, 22] as:

$$p(\mathbf{x}_t^n, \mathbf{x}_t^l \mid \mathbf{y}_{1:t}) = p(\mathbf{x}_t^n \mid \mathbf{y}_{1:t}) p(\mathbf{x}_t^l \mid \mathbf{x}_t^n, \mathbf{y}_{1:t}), \quad (6)$$

where the nonlinear states $\mathbf{x}_t^n$ are represented by particles, and the conditionally linear states $\mathbf{x}_t^l$ are tracked analytically via a Kalman filter conditioned on each particle [23].

In our setup, we partition the full state into a nonlinear component $\mathbf{x}_t^n$ and a linear component $\mathbf{x}_t^l$. The nonlinear state captures the vehicle pose and speed as: $\mathbf{x}_t^n = [p_x, p_y, \theta, v]^\top \in \mathbb{R}^4$, where (p_x, p_y) is the position, θ is the

heading, and v is the forward speed. The linear state models a slowly varying sensor bias on each measurement channel as: $\mathbf{x}_t^l = [\delta\phi_1, \dots, \delta\phi_M]^\top \in \mathbb{R}^M$. With this decomposition, the measurement model becomes conditionally linear:

$$\mathbf{y}_t = \mathbf{h}(\mathbf{x}_t^n) + H \mathbf{x}_t^l, \quad (7)$$

remembering that $\mathbf{h}$ is the nominal PDE evaluated at the particle position (see equation (2)) and $H = I_M$ is the identity observation matrix for the bias states. Conditioned on a given particle position $\mathbf{x}_t^{n,(i)}$, equation (10) is linear in $\mathbf{x}_t^l$, which is exactly the structure required for Rao-Blackwellization.

B. Motion Model

The RBPF maintains N particles. Each particle i is represented by the tuple $\{\mathbf{x}_t^{n,(i)}, \hat{\mathbf{x}}_t^{l,(i)}, P_t^{(i)}, w_t^{(i)}\}$, where $\mathbf{x}_t^{n,(i)}$ denotes the sampled nonlinear state, $\hat{\mathbf{x}}_t^{l,(i)}$ and $P_t^{(i)}$ are the conditional Kalman filter mean and covariance of the linear state, and $w_t^{(i)}$ is the associated importance weight. The nonlinear component evolves according to a unicycle-type motion model driven by the control input $\mathbf{u}_t = (v_{\text{cmd}}, \omega_{\text{cmd}})^\top$:

$$\mathbf{x}_{t+1}^n = \mathbf{f}(\mathbf{x}_t^n, \mathbf{u}_t) + \boldsymbol{\omega}_t, \quad \boldsymbol{\omega}_t \sim \mathcal{N}(\mathbf{0}, \Sigma_f), \quad (8)$$

where Σ_f is the covariance of the additive process noise affecting the nonlinear state. The linear bias state evolves according to an identity state-transition model driven by zero-mean Gaussian white noise:

$$\mathbf{x}_{t+1}^l = \mathbf{x}_t^l + \boldsymbol{\zeta}_t, \quad \boldsymbol{\zeta}_t \sim \mathcal{N}(\mathbf{0}, \Lambda), \quad (9)$$

where Λ is a diagonal covariance matrix modeling slow-varying sensor drift.

C. Measurement Update

At each time step, the RBPF incorporates a new observation $\mathbf{y}_t \in \mathbb{R}^M$ of the M PDE field channels, as defined in (2). The update performs two coupled operations for each particle i , i.e., computing the importance weight and updating the Kalman filter for the linear substate $\mathbf{x}_t^l$.

The linear substate $\mathbf{x}_t^l \in \mathbb{R}^M$ represents a slowly varying additive sensor bias on each field channel. The measurement model at particle i is

$$\mathbf{y}_t = \phi(\mathbf{x}_t^{n,(i)}, t) + H \mathbf{x}_t^{l,(i)} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \Sigma), \quad (10)$$

Innovation: For each particle i , the nominal field vector $\phi(\mathbf{x}_t^{n,(i)}, t)$ is obtained by rounding the particle's continuous position to the nearest grid index and reading off the stored field values. If the cell lies outside the domain boundary or is a dry (masked) cell, the particle receives zero weight. Otherwise, the innovation is

$$\mathbf{e}_t^{(i)} = \mathbf{y}_t - \phi(\mathbf{x}_t^{n,(i)}, t) - H \hat{\mathbf{x}}_{t|t-1}^{l,(i)}, \quad (11)$$

and the innovation covariance is

$$S_t^{(i)} = H P_{t|t-1}^{(i)} H^\top + \Sigma. \quad (12)$$

Weight Update: The importance weight is proportional to the marginal likelihood of $\mathbf{y}_t$ given the particle's nonlinear state, obtained by integrating out $\mathbf{x}_t^l$ analytically. This marginal likelihood is Gaussian because the measurement model is conditionally linear in $\mathbf{x}_t^l$, which is the core benefit of Rao-Blackwellization. The log-weight is computed via the Cholesky factorization $S_t^{(i)} = L_t^{(i)} L_t^{(i)\top}$ for numerical stability:

$$\log \tilde{w}_t^{(i)} = - \sum_{k=1}^M \log [L_t^{(i)}]_{kk} - \frac{1}{2} \left\| L_t^{(i)-1} \mathbf{e}_t^{(i)} \right\|^2 - \frac{M}{2} \log 2\pi, \quad (13)$$

which corresponds to evaluating $\log \mathcal{N}(\mathbf{e}_t^{(i)}; \mathbf{0}, S_t^{(i)})$. Particles in dry or out-of-bounds cells are assigned $\log \tilde{w}_t^{(i)} = -\infty$, enforcing the physical constraint that the vehicle must remain in navigable water. Weights are then normalized using the log-sum-exp identity:

$$w_t^{(i)} = \frac{\exp(\log \tilde{w}_t^{(i)} - c)}{\sum_{j=1}^N \exp(\log \tilde{w}_t^{(j)} - c)}, \quad c = \max_{1 \leq j \leq N} \log \tilde{w}_t^{(j)}. \quad (14)$$

Kalman Filter Update: The Kalman filter for particle i is updated via the standard equations:

$$K_t^{(i)} = P_{t|t-1}^{(i)} H^\top (S_t^{(i)})^{-1}, \quad (15)$$

$$\hat{\mathbf{x}}_{t|t}^{l,(i)} = \hat{\mathbf{x}}_{t|t-1}^{l,(i)} + K_t^{(i)} \mathbf{e}_t^{(i)}, \quad (16)$$

$$P_{t|t}^{(i)} = P_{t|t-1}^{(i)} - K_t^{(i)} S_t^{(i)} K_t^{(i)\top}. \quad (17)$$

Each particle carries its own Kalman filter that tracks the per-channel sensor bias conditioned on that particle's trajectory hypothesis.

D. Prediction and Resampling

The RBPF prediction step differs from a standard particle filter by performing ancestor resampling before propagating the nonlinear state. At the start of each time step, ancestor indices are drawn via systematic resampling from the current weight distribution. Each particle i then inherits the nonlinear state, bias mean, and bias covariance of its chosen ancestor, and all weights are reset to uniform. The nonlinear state is then propagated by the motion model with additive noise:

$$\mathbf{x}_t^{n,(i)} = \mathbf{f}(\mathbf{x}_{t-1}^{n,(a_i)}, \mathbf{u}_t + \boldsymbol{\omega}_t^{(i)}), \quad \boldsymbol{\omega}_t^{(i)} \sim \mathcal{N}(\mathbf{0}, \Sigma_f), \quad (18)$$

where a_i is the ancestor index for particle i and Σ_f is the motion noise covariance acting on the control inputs. The linear state covariance is updated by adding the bias drift noise:

$$\hat{\mathbf{x}}_{t|t-1}^{l,(i)} = \hat{\mathbf{x}}_{t-1}^{l,(a_i)}, \quad P_{t|t-1}^{(i)} = P_{t-1}^{(a_i)} + \Lambda, \quad (19)$$

where Λ is the bias process noise covariance (c.f. equation (9)). The bias mean is unchanged in the prediction step since the bias is modeled as a random walk.

Algorithm 1 RBPF for PDE-based Vehicle Localization

Require: Particle set $\{\mathbf{x}_{t-1}^{n,(i)}, \hat{\mathbf{x}}_{t-1}^{l,(i)}, P_{t-1}^{(i)}, w_{t-1}^{(i)}\}_{i=1}^N$, control $\mathbf{u}_t$, observation $\mathbf{y}_t \in \mathbb{R}^M$

Ensure: Updated particle set $\{\mathbf{x}_t^{n,(i)}, \hat{\mathbf{x}}_t^{l,(i)}, P_t^{(i)}, w_t^{(i)}\}_{i=1}^N$

- 1: **Ancestor resampling:** Draw indices $\{a_i\}$ by systematic resampling from $\{w_{t-1}^{(i)}\}$; reset $w^{(i)} \leftarrow 1/N$
- 2: **for** $i = 1, \dots, N$ **do**
- 3: **Predict nonlinear:** $\mathbf{x}_t^{n,(i)} \leftarrow \mathbf{f}(\mathbf{x}_{t-1}^{n,(a_i)}, \mathbf{u}_t + \boldsymbol{\omega}_t^{(i)})$,
 $\boldsymbol{\omega}_t^{(i)} \sim \mathcal{N}(\mathbf{0}, \Sigma_f)$
- 4: **Predict linear:** $\hat{\mathbf{x}}_{t|t-1}^{l,(i)} \leftarrow \hat{\mathbf{x}}_{t-1}^{l,(a_i)}$, $P_{t|t-1}^{(i)} \leftarrow P_{t-1}^{(a_i)} + \Lambda$
- 5: **Innovation:** $\mathbf{e}_t^{(i)} \leftarrow \mathbf{y}_t - \phi(\mathbf{x}_t^{n,(i)}, t) - H \hat{\mathbf{x}}_{t|t-1}^{l,(i)}$
- 6: **Innovation covariance:** $S_t^{(i)} \leftarrow H P_{t|t-1}^{(i)} H^\top + R$;
factor $S_t^{(i)} = L_t^{(i)} L_t^{(i)\top}$ via Cholesky
- 7: **Log-weight:** $\log \tilde{w}_t^{(i)} \leftarrow \log w_{t-1}^{(i)} + \log \mathcal{N}(\mathbf{e}_t^{(i)}; \mathbf{0}, S_t^{(i)})$ via (13)
- 8: **Kalman update:** Compute $K_t^{(i)}$; update $\hat{\mathbf{x}}_{t|t}^{l,(i)}, P_{t|t}^{(i)}$ via (15)–(17)
- 9: **end for**
- 10: **Normalize:** $w_t^{(i)} \leftarrow \exp(\log \tilde{w}_t^{(i)} - c) / \sum_j \exp(\log \tilde{w}_t^{(j)} - c)$ via (14)
- 11: **if** $N_{\text{eff}} < N/2$ **then**
- 12: SIR resample particles with associated $(\hat{\mathbf{x}}_t^{l,(i)}, P_t^{(i)})$; reset weights to $1/N$
- 13: **end if**

After the measurement update in (14), particle degeneracy is monitored via the effective sample size

$$N_{\text{eff}} = \frac{1}{\sum_{i=1}^N (w_t^{(i)})^2}. \quad (20)$$

A secondary SIR resampling step is triggered when $N_{\text{eff}} < N/2$. During this step, the nonlinear state $\mathbf{x}_t^{n,(i)}$ and the associated Kalman filter $(\hat{\mathbf{x}}_t^{l,(i)}, P_t^{(i)})$ are duplicated together, preserving the per-particle posterior on the linear substate. After resampling, all weights are reset to $1/N$. The complete RBPF cycle is summarized in Algorithm 1.

V. RESULTS

This section aims to showcase two case studies in different areas using the two presented PDE models, given that they require different types of data to be simulated and validated. All experiments and simulations are conducted using a general-purpose laptop computer. For simulations, the measurement and motion noise covariances are set to $\text{diag}(0.2^2, 0.2^2, 0.1^2)$ and $\text{diag}(0.5^2, 0.087^2)$, respectively. The linear process noise is set to $\text{diag}(10^{-8}, 10^{-8}, 10^{-8})$.

A. Case Study 1: Shallow Water Equations

In this scenario, we target the study of equations that arise as the solutions of (3), the robot has access to $M = 3$ measurement channels, that is, the water height η and the water velocities $\mathbf{v} = (u, v)$. As mentioned before, we made

explicit the dependence of the source term $\mathbf{s}$ on ϕ . To have realistic conditions from small- to medium-sized coastal areas, we incorporate two terms to $\mathbf{s}$: bathymetry $b : \mathcal{W} \rightarrow \mathbb{R}$ and wind speed $\omega : \mathcal{W} \times \mathbb{R}^+ \rightarrow \mathbb{R}^2$. Concisely, redefine

$$\mathbf{s} = \begin{bmatrix} 0 \\ \boldsymbol{\tau}(\mathbf{v} - \boldsymbol{\omega}) - g \cdot (\nabla b)^\top \end{bmatrix}, \quad (21)$$

$$\boldsymbol{\tau}(\mathbf{z}) = \frac{\rho_{\text{water}}}{\rho_{\text{air}}} c_c \|\mathbf{z}\| \mathbf{z}.$$

$\boldsymbol{\tau}$ encompasses the wind stress tensor, ρ_{water} , ρ_{air} , and c_c are the water density, the air density, and the drag coefficient, respectively. The data used to estimate the bathymetry were obtained from the General Bathymetric Chart of the Oceans (GEBCO), and a region of interest was selected around a designated coastal area. On the other hand, we obtained the wind velocities from recorded conditions around this area when the wind has a significant impact on the water's velocity $\|\omega\| \approx 12$ m/s. Figure 1 shows multiple snapshots of this model simulating how a water wave propagates towards the coast; in particular, it displays the water height from the surface, defined as the water height corrected by the bathymetry $\eta - b$. Moreover, it shows the robot traversing an elliptical path (ground truth) while using the RBPF to localize itself by measuring the water velocities and water height $\phi = (\eta, u, v)$. The RBPF uses 500 particles, simulation step time is set to 1 s, and a total simulation time of 500 s. As can be observed in Figure 1, the particles start scattered. In the absence of field measurements, the estimated red trajectory deviates from the desired path. However, when the vehicle encounters the field's noisy observations, the particles quickly converge toward the true vehicle state.

B. Case Study 2: Advection-Diffusion

The other motivating scenario includes the incorporation of other key features that are of major importance in multiple environmental surveillance, monitoring, and restoration efforts. We bring again equation (5), which shows that we need to define the diffusivity constants κ , which depend on each water feature and the velocity field $\mathbf{v}$ that couples all the equations together. The diffusivity constants vary between $0.01 \text{ m}^2/\text{s}$ and $1 \text{ m}^2/\text{s}$ and the water velocity field averages $\|\mathbf{v}\| = 1.0 \text{ m/s}$. We utilized data previously acquired to fit the initial condition ϕ_0 of the PDE.

Figure 2 shows multiple snapshots of this model simulating how environmental quantities evolve over a total simulation time of 500 s with data collected at 1 Hz. Here, we chose three water features: chlorophyll ($\mu\text{g/L}$), turbidity in Formazin Nephelometric Units (FNU), and temperature ($^\circ\text{C}$). Similar to the SWE simulation, we use the RBPF with 500 particles. The blue trajectories represent the ground truth collected by the GPS sensor, while the red trajectories are the RBPF estimates. As can be observed in Figure 2, the RBPF accurately estimates the vehicle state when all three measurement channels are fused. In the absence of spatially distinctive measurements, the estimated trajectories deviate from the ground truth.




Fig. 1: RBPF localization at different timestamps showing the estimation of the robot state over time, solving the SWE (3). The SWE incorporates the **water height** (m) corrected by the bathymetry (i.e., water above sea level); the **velocity field** (ms^{-1}) is plotted as a vector field; and the **terrain above sea level** defines the coastal domain. The robot is localizing itself by reconstructing an **estimated trajectory** from an elliptical **ground truth trajectory** using 500 RBPF particles.




Fig. 2: RBPf localization at different timestamps showing the estimation of the robot state over time. The robot estimates the **ground truth path** by estimating its **current trajectory** using the RBPf particles. The robot has access to $M = 3$ fields from left to right: Turbidity (FNU), Temperature ($^{\circ}\text{C}$), and Chlorophyll ($\mu\text{g/L}^{-1}$).

C. Performance Analysis

To evaluate filter performance, we report two complementary metrics: final error and root mean square error (RMSE). The final error captures the position error at the last time step, providing a direct measure of whether the filter has converged by the end of the run, though it reflects only a single snapshot. RMSE, on the other hand, aggregates squared errors across all time steps, providing a measure of average tracking accuracy over the entire trajectory. Reporting both together gives a more complete picture: a filter with low final error but high RMSE performed poorly early and recovered

late, while one with low RMSE but high final error tracked well throughout but diverged near the end.




Fig. 3: Comparison of PF and RBPf localisation accuracy on the shallow-water estimation (SWE) task across five independent runs for varying numbers of particles. Both metrics decrease as the particle count increases, with RBPf consistently outperforms PF at equivalent sample sizes due to its Rao-Blackwellised marginalisation of the linear sub-structure.

For the SWE model in Figure 3, the RBPf consistently outperforms the standard PF in sustained tracking accuracy. At 500 particles, the RBPf achieves an RMSE of 0.80 ± 0.03 m compared to the PF's 1.07 ± 0.04 m—a roughly 25% reduction—while also attaining a lower final error (0.70 ± 0.15 m vs. 0.82 ± 0.02 m). Notably, the RBPf with only 100 particles (RMSE = 0.84 ± 0.04 m) already surpasses the PF at 500 particles, demonstrating substantially greater sample efficiency. Across all particle counts, the PF's RMSE plateaus near 1.07–1.10 m, whereas the RBPf steadily improves from 0.84 m to 0.80 m, indicating that additional particles yield diminishing returns for the PF but continued benefit for the RBPf.




Fig. 4: Comparison of PF and RBPf localisation accuracy on the Advection-Diffusion task across five independent runs for varying numbers of particles. Interestingly, both metrics have steady RMSE values as the particle count increases. Again, RBPf consistently outperforms PF at equal sample size.

For the advection-diffusion model in figure 4, both filters face a considerably harder localization problem due to the


weaker spatial gradients in the scalar concentration field. The PF exhibits remarkably stable but high errors across all particle counts, with RMSE clustered around 2.3–2.4 m and final errors near 2.2–3.0 m, suggesting a fundamental accuracy ceiling. The RBPF shows greater variability—its RMSE ranges from 1.23 ± 0.35 m at 500 particles to 3.1 ± 1.17 m at 400 particles—but achieves its best performance at 500 particles, where it reduces RMSE by nearly 47% relative to the PF (1.23 m vs. 2.32 m). The high variance across runs reflects the sensitivity of the Kalman update to the weaker and more ambiguous measurement gradients in this environment. These results suggest that the RBPF’s analytical correction step offers a clear advantage when sufficient particles are available, but its performance is less predictable than in the richer SWE observation space.

To assess localization sensitivity to multimodal measurement fusion, we conduct a feature ablation study by varying the number of observable state variables. For the SWE model, the full measurement vector comprises three quantities: two velocity components and free-surface elevation, indexed as 0, 1, 2. For the Advection-Diffusion model, the situation is similar, yet we may choose from a wider range of water quality features: temperature, dissolved oxygen, salinity, turbidity, and chlorophyll. We chose at random a number of measurement channels $M = 1, 2, 3$, and the Kalman filter dimension, measurement noise, and bias states are reduced accordingly. Five independent runs per configuration capture variability from both random feature selection and stochastic particle dynamics.

M	SWE (3)		Advection-Diffusion (5)	
	RMSE (m)	Final Error (m)	RMSE (m)	Final Error (m)
1	1.59 ± 0.55	1.35 ± 0.49	$2.63 \pm 1.6.38$	2.11 ± 0.132
2	0.94 ± 0.10	0.90 ± 0.26	$2.27 \pm 1.2.06$	$1.61 \pm 0.3.9$
3	0.82 ± 0.06	0.69 ± 0.08	1.51 ± 0.79	1.62 ± 0.64

TABLE I: Performance Comparison of SWE and Advection-Diffusion Models across varying Feature Counts.

As shown in Table I, the SWE model degrades gracefully: RMSE increases from 0.82 ± 0.06 m with all three features to 1.59 ± 0.55 m with a single feature—a roughly twofold rise. The advection–diffusion model exhibits markedly steeper degradation, with RMSE growing from 1.51 ± 0.79 m to 2.63 ± 1.63 m and substantially higher variance throughout. These results confirm that fusing multiple field channels as multimodal measurements significantly improves localization accuracy.



(a) ASV SeaRobotics Suyveyor, left (b) ASV during the field YSI EXO2, right Sontek M9 ADCP. experiment.

Fig. 5: Hardware utilized during the field experiments; the ASV is deployed in a coastal area where it can use its sensors to measure water quality features, yet it can also measure velocity current speeds, and height from the seafloor.

D. Physical Experiments

As shown in Figure 5, we deploy an ASV in a coastal area where it was configured to sense three water quality parameters satisfying the advection diffusion equation (5): temperature ($^{\circ}\text{C}$), dissolved oxygen (mg/L), and salinity in Parts Per Thousand (PPT). In the same fashion, the diffusivity constants were set in the same range as in the previous simulations. The velocity field was extracted from the ASV’s ADCP.

The vehicle completed a loop around the designated area, measuring these time-variant variables by measuring the three fields as it moved. The ground truth was obtained by recording the robot’s GPS location during its journey. Figure 6 shows the predictions made by the RBPF and how the particles are, in general, grouped around the vehicle’s true location, which shows that a proper water quality field combination, although dynamic and non-stationary, can provide proper information to extract the vehicle’s position with relative precision.

VI. CONCLUSION

This paper presented a localization framework that exploits spatiotemporal environmental fields governed by PDEs as localization signatures for GPS-denied operation. Two PDE classes were formulated, the shallow water equation and the advection-diffusion equation, and fused as multimodal measurements within a Rao-Blackwellized particle filter that maintains closed-form estimates of per-channel sensor bias through conditionally linear Kalman filters. Simulation results confirm that the RBPF consistently outperforms a standard particle filter in both PDE scenarios, achieving lower RMSE and final position error across all particle counts while requiring significantly fewer particles to reach comparable accuracy. The feature ablation study further demonstrated that fusing multiple field channels as multimodal measurements is critical, with localization accuracy degrading substantially as channels are removed. Field experiments with an ASV measuring salinity, temperature, and dissolved oxygen validated that PDE-governed environmental fields provide sufficient spatial variability for practical localization. Future work will extend this framework to the SLAM problem,




Fig. 6: The path traversed by the ASV, it loops around the region of interest. Again, it uses RBPF localization, measuring $M = 3$ water quality features: Salinity (PPT), Temperature ($^{\circ}\text{C}$), and Dissolved oxygen (mgL^{-1}). This shows different timestamps showing the estimation of the robot state over time. The robot estimates the ground truth path by estimating its current trajectory using the RBPF particles.

where the vehicle jointly estimates its own state and the underlying environmental field map in real time.

REFERENCES

- [1] A. Solin, M. Kok, N. Wahlström, T. B. Schön, and S. Särkkä, "Modeling and interpolation of the ambient magnetic field by Gaussian processes," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 1112–1127, 2018.
- [2] A. Solin, S. Särkkä, J. Kannala, and E. Rahtu, "Terrain navigation in the magnetic landscape: Particle filtering for indoor positioning," in *The European Navigation Conference (ENC)*, pp. 1–9, IEEE, 2016.
- [3] B. Ferris, D. Fox, and N. D. Lawrence, "WiFi-SLAM using Gaussian process latent variable models," in *International Joint Conference on Artificial Intelligence*, (Hyderabad, India), pp. 2480–2485, 2007.
- [4] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit, "FastSLAM: A factored solution to the simultaneous localization and mapping problem," in *The AAAI Conference on Artificial Intelligence*, vol. 593598, pp. 593–598, 2002.
- [5] M. Dunbabin and L. Marques, "Robots for environmental monitoring: Significant advancements and applications," *IEEE Robotics & Automation Magazine*, vol. 19, no. 1, pp. 24–39, 2012.
- [6] Z. Song and K. Mohseni, "Towards background flow based auv localization," in *IEEE Conference on Decision and Control*, pp. 6945–6950, IEEE, 2014.
- [7] Z. Song and K. Mohseni, "Concurrent flow-based localization and mapping in time-invariant flow fields," pp. 7205–7210, 2019.
- [8] A. Kumar, T. C. Silva, V. Edwards, and M. A. Hsieh, "Flow-based localization and mapping for multi-robot systems," IEEE, 2025.
- [9] A. K. Li, T. C. Silva, and M. A. Hsieh, "Encode: Active learning of unknown flows with koopman operators," *IEEE Robotics and Automation Letters*, vol. 9, no. 12, pp. 11282–11289, 2024.
- [10] J. Chen, M. Chen, and S. Park, "MDCPP: Multi-robot dynamic coverage path planning for workload adaptation," in *arXiv preprint arXiv:2509.23705*, 2025.
- [11] J. Jouffroy and J. Opderbecke, "Underwater vehicle trajectory estimation using contracting pde-based observers," in *American Control Conference*, vol. 5, pp. 4108–4113, 2004.
- [12] "—," in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1–7, IEEE, 2025.
- [13] T. Wiedemann, C. Manss, D. Shutin, A. J. Lilienthal, V. Karolj, and A. Viseras, "Probabilistic modeling of gas diffusion with partial differential equations for multi-robot exploration and gas source localization," in *European Conference on Mobile Robots*, pp. 1–7, IEEE, 2017.
- [14] K. Murphy and S. Russell, "Rao-Blackwellised particle filtering for dynamic Bayesian networks," in *Sequential Monte Carlo methods in practice*, pp. 499–515, Springer, 2001.
- [15] M. Kok, A. Solin, and T. B. Schön, "Rao-blackwellized particle smoothing for simultaneous localization and mapping," *Data-Centric Engineering*, vol. 5, p. e15, 2024.
- [16] M. Shen, J. Sun, H. Peng, and D. Zhao, "Improving localization accuracy in connected vehicle networks using rao-blackwellized particle filters: Theory, simulations, and experiments," vol. 20, pp. 2255–2266, IEEE, 2018.
- [17] Y. Taguchi, T. K. Marks, and H. Okuda, "Rao-blackwellized particle filtering for probing-based 6-dof localization in robotic assembly," in *IEEE International Conference on Robotics and Automation*, pp. 2610–2617, IEEE, 2010.
- [18] F. C. Teixeira, J. Quintas, and A. Pascoal, "Auv terrain-aided navigation using a doppler velocity logger," *Annual Reviews in Control*, vol. 42, pp. 166–176, 2016.
- [19] J. McConnell, I. Collado-Gonzalez, and B. Englot, "Perception for underwater robots," *Current Robotics Reports*, vol. 3, no. 4, pp. 177–186, 2022.
- [20] M. J. Islam, A. Q. Li, Y. A. Girdhar, and I. Rekleitis, "Computer vision applications in underwater robotics and oceanography," in *Computer Vision*, pp. 173–204, Chapman and Hall/CRC, 2024.
- [21] A. Doucet, S. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statistics and Computing*, vol. 10, no. 3, pp. 197–208, 2000.
- [22] K. Murphy and S. Russell, "Rao-Blackwellised particle filtering for dynamic Bayesian networks," in *Sequential Monte Carlo Methods in Practice*, pp. 499–515, Springer, 2001.
- [23] T. Schön, F. Gustafsson, and P.-J. Nordlund, "Marginalized particle filters for mixed linear/nonlinear state-space models," in *IEEE Transactions on Signal Processing*, vol. 53, pp. 2279–2289, IEEE, 2005.