Unified adaptive and cooperative planning using multi-task coregionalized Gaussian processes

Lorenzo Booth Stefano Carpin

***Abstract*— For robots tasked with surveying the temporal dynamics of a changing environment, a choice must be made to observe novel regions of the environment or to re-survey previously visited regions, which may have changed. We present a novel multi-robot informative path planner (IPP) that combines an environmental and task kernel to direct mobile robots to gather samples from regions that would result in the greatest expected improvement in map accuracy. Our planner utilizes a multi-output Gaussian process to unify priors about the spatiotemporal environment along with priors about observational correlations between sensing vehicles. Additionally, we extend our analysis into an adaptive planning scenario and examine the performance under different planning configurations. We find that planning performance is largely driven by the choice of environmental priors, and that unrepresentative priors can be improved through adaptive planning.**

I. Introduction and related work

With the growth in use of antonomous vehicles for agriculture and natural resource management [15], so too grows the burden for distilling actionable information from the observations gathered. In the field of geostatistics, there is a long history of using mathematical methods for modeling environments through sparse measurements [24]. This is particularly relevant when observations require direct sampling, such as for robotic plant phenotyping [6] and robotic sampling of plant tissues [5].

This paper considers the task of modeling a dynamic phenomenon in an environment with a team of robots equipped with point sensors. Given a limited movement budget, it is desirable for the robots to visit locations that will result in the most accurate signal reconstructions when the surveying mission is complete.

This task of *informative path planning* (IPP) aims to find obstacle-free trajectories in an environment that maximize the information gathered during traveling. When applied to environmental monitoring tasks, it is closely related to the task of *optimal sensor placement* [10], and can be considered as an optimization problem subject to constraints such as the physical confines of the rows of an vineyard [25] or the locations of pre-established monitoring stations [3]. When the process to be monitored evolves over time, there is a choice to be made: whether to explore unvisited locations in the environment, or to re-survey previously-visited locations that may have changed. Approaches explored in prior works include a recursive-greedy approach for surveying a time-varying field with a single robot [2] and multi-robot efforts that leverage clustering in order to divide the observation domain among different sampling vehicles. Recent IPP approaches consider robotic path planning in response to multiple objectives, such as variable sensor models [19], and multi-modal sensor configurations [17].

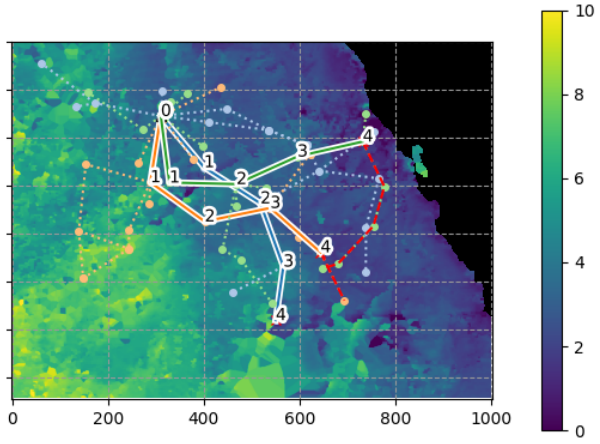

Fig. 1: An example of a planning mission with 3 vehicles in an ocean turbidity monitoring example. Vehicle movement history is displayed with a white outline. In this image, 3 vehicles have acquired 4 observations each. Candidate paths and sampling locations as proposed by our planner are shown in dashed lines. The planner for the green vehicle has just finished, and proposed the best candidate path, shown in the dotted red line.

When surveying unknown environments, it is often desirable to utilize an *adaptive* sampling scheme, where new samples are targeted based on information collected from previous samples in the surveying mission in order to improve the overall surveying ability of the robotic system (subject to the evaluation criteria) [7]. When extended to robotic surveying, this has been called the adaptive informative path planning (*AIPP*) problem. Recent literature has found learning-based approaches to be particularly suited to AIPP across a wide variety of mapping objectives [18] [26] [11] and different approaches have been used to extend the task into multi-robot surveying efforts [14] [22].

In this work we present a multi-robot informative path planner *IIG-Cooperative*, that guides robots to gather observations in regions that result in a higher expected improvement of the environmental map. Our approach is *cooperative* as it is able to incorporate information from other robots involved in the team surveying task. This algorithm builds

L. Booth and S. Carpin are with the Dept. of Electrical Engineering and Computer Science, University of California, Merced, CA, USA. L. Booth is partially supported by the Labor & Automation in California Agriculture UC-MRPI M21PR3417, and by the EDA BBBRC iCREATE grant. S. Carpin is partially supported by the IoT4Ag Engineering Research Center funded by the National Science Foundation under NSF Cooperative Agreement Number EEC-1941529. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the view of the funding agencies.

upon *IIG*, an informative sampling-based planner [13] that utilized mutual information as the basis for establishing when a proposed sampling location would result in an improved model of an environment. The algorithm in turn builds upon *RIG*, which presented an informative planner based on RRTs [16] to find obstacle-free trajectories through an environment that maximize information gathered along the way by leveraging the asymptotic optimality of random-trees. [12]. In a previous effort, we presented *IIG-ST*, which utilizes an information-theoretic stopping criterion with an information function that trades off between exploring new environments and re-surveying previous potentially-stale observations, when monitoring a time-varying spatiotemporal phenomenon [4].

IIG-Cooperative adds an information function that incorporates the location of observations from other sensing vehicles into a robot's evaluation of the information gain at proposed sampling locations. In [4], the performance of the algorithm was dependent on tuning the information function to the environment through expert selection of spatial and temporal priors. In this work, we overcome the limitation by presenting an adaptive planning framework that alternates between planning, execution, and updating the environmental priors in order to generate an improvement of planning performance as a survey proceeds. The main contributions of this work are as follows:

- *IIG-Cooperative*, an informative path planner that integrates multi-vehicle interactions through a separable kernel, reconciled to the surveying objective
- An open-source [1] adaptive planning framework that takes advantage of observations gathered to create a best-effort continuous improvement of model hyperparameters and planning optimality
- An experimental validation of the base planner and a spatiotemporal variant under different configuration and communication scenarios

Figure 1 shows an example of our cooperative planner used in an ocean monitoring example. Our approach is discussed in section III and is enabled by the use of a multi-task Gaussian process, where each robot's observations comprise an input of a multi-input modeling task. We extend this effort in an adaptive planning framework, applied in experiments on simulated and natural environments in section IV. Our framework can accommodate a diverse variety of planning and configuration scenarios, including variable sensor models and variable noise models, all of which will serve as the basis for future investigation (discussed in section V).

II. Problem Formulation

Consider the case where an environmental field is observed by different sensing agents v, comprising a team of N_v agents. The different agents may be equipped with different sensors that may have different noise characteristics, and observations that are unlikely to be coincident in space and time.

[1]https://github.com/ucmercedrobotics/ipp-RRT-family

We can consider three scenarios:

1) Where observations from all sensing agents are shared to form one unified training set, for a unified model that serves as the basis for allocating future sampling locations for all sensing vehicles, by a global planner.
2) Where observations from a given sensing agent v are used to build a representation of the environment that is *only used* to inform future plans for the agent v.
3) Where correlations between sensing agents are captured in a covariance function. Each agent independently builds a representation of the environment, assisted with knowledge of observations collected from other sensing agents.

Case 1 can be addressed through a variety of approaches, including partitioning an environment among a set of sensing agents. Case 2 is applicable if we apply any single-robot IPP algorithm to a team of surveying robots, who independently sample and plan without knowledge of the other agents in the team.

This paper considers Case 3, where each member v of the team of robots is *tasked* with producing an independent, internal representation of the environment to guide its own planning. We denote each task with the letter j. Observations are shared between robots and are used to update each robot's internal representation of map uncertainty.

In our planning framework, robots alternate between: planning, surveying, communicating, updating priors, and re-planning. This occurs in a continuous loop until the end of the allotted survey period, or until the robots consume their movement budget. At the end, observations from all robots are aggregated to form a final, unified model of the environment.

III. Methods

A. Overview

Our planning framework encodes the intuition that it is more desirable to collect observations where a model is deficient, rather than where a model is sufficient. In the following sections we describe: the form of the environmental model and how we derive model adequacy (subsection III-B); how we encode and update prior knowledge about the environment and the sensing vehicles (subsection III-C); and how this knowledge allow us to determine where to sample (subsection III-D). Implementation details of our planner and an algorithmic overview are found in (subsection III-E)

B. Environmental model

Congruent with standard approaches in geostatistics, we extend a 2D spatial regression task (Kriging) to consider a 3D regression task of an unknown scalar-valued environmental process that changes over time (e.g., chemical concentration and distribution, soil moisture content, etc.) represented as the function: $f\colon \mathcal{X} \to \mathbb{R}$ that is modeled on discrete intervals in space ($N_{x,y}$) and time (N_t) where, $\mathcal{X} \subset \mathbb{R}^{N_{x,y}} \times \mathbb{R}^{N_t}$. The phenomenon is observed through noisy measurements $y_{i,j}$ made at location $i \in N_{x,y} \cup N_t$ by sensing agent j where, $y_{i,j} = f(x_i) + \varepsilon_j$. Noisy perturbations ε_j are modeled with

a zero-mean, homoskedastic, additive Gaussian noise model, that is consistent within a given sensor on vehicle v (denoted by task j): $\varepsilon \sim \mathcal{N}\left(0, \sigma_j^2 I\right)$.

1) GP regression: We model f as a realization of a Gaussian process, represented over a space of functions in $\mathbb{R}^3$ and follow the standard notation provided in [20]. We can represent the joint distribution of noisy observations $\mathbf{y}$ $(f(x_1)+\varepsilon_1, \ldots, f(x_n)+\varepsilon_n)$ and predictions of the function f $(f_\star, \ldots, f_{\star^n})$ as:

$$\begin{bmatrix} \mathbf{y} \\ f(x_\star) \end{bmatrix} \sim \mathcal{N}\left(\mu, \begin{bmatrix} k(\mathbf{X},\mathbf{X}) + \sigma^2 I_n & k\left(\mathbf{X}, x_\star\right) \\ k\left(x_\star, \mathbf{X}\right) & k\left(x_\star, x_\star\right) \end{bmatrix}\right) \quad (1)$$

defined over the vector of input observations $\mathbf{X}$ and query points $x_\star$ (where predictions are made). k is a covariance function (or kernel), that establishes a basis of correlation in the environment as a function of distance in the input space. This basis can encode prior knowledge about the environment and sensing agents and is discussed in detail in subsection III-C.

From this joint distribution, we can obtain the expected value of the environmental field ($\mathbb{E}[f_\star]$) and variance ($\mathbb{V}[f_\star]$) evaluated at a query point $x_\star$. For path planning, we are particularly interested in the posterior variance (the joint variance after incorporating observations at ($\mathbf{X}$)), which can be derived as:

$$\sigma = \mathbb{V}\left[f_\star\right] = k\left(x_\star, x_\star\right) - k\left(x_\star, \mathbf{X}\right) \times \left[k(\mathbf{X},\mathbf{X}) + \sigma_n^2 \mathbf{I}_n\right]^{-1} k\left(\mathbf{X}, \mathbf{x}_*\right) \quad (2)$$

From this, we can obtain the conditional probability distribution of f, conditioned on an observation y: $f \mid y = \mathcal{N}(\mu_{f|y}, \Sigma_{f|y})$, which will be used in the utility function described in subsection III-D.

C. Spatiotemproal and task priors

1) Spatiotemporal prior: The kernel (or, covariance function) k is a function that provides the expected correlation between pairs of data points. While arbitrary functions of input pairs are not guaranteed to be valid covariance functions, there exists a considerable amount of choice and discretion in choosing a function that is appropriate to the predictive task. Both the choice of kernel and the function hyperparameters encode assumptions about the property which we wish to predict [20].

Following [4], we establish a base kernel, composed of a spatial and temporal kernel:

$$k((s,t), t(s',t')) = k_s(s,s')k_t(t,t') \quad (3)$$

where s refers to the spatial index (such as, a geographical coordinate) and t refers to a temporal index (a timestamp). For the spatial dimensions, we use a Matérn kernel with $\nu = 3/2$, chosen in part for its use in the geostatistical literature and its ability to capture discontinuities present in natural phenomena. We use a radial basis function kernel to capture smoothly diffusive process in the time dimension. For additional discussion and details about the kernel choice, readers may refer to [24] and [4].

2) Task prior: In section II, we describe how each sensing agent is *tasked* with producing a unique model of the environment, to guide its path planning. Following the terminology from [1], we describe our system of models as *multi-task*, where the process f is observed by different sensing agents j. Our model is *multi-output*, where each output d can produce a unique representation of the environment based on observations collected from a particular sensing agent j.

Following the notation of [1] let us consider a training set constructed of data pairs $S_d = (\mathbf{X_d}, \mathbf{Y_d}), = (\mathbf{x}_{d,1}, y_{d,1}), \ldots (\mathbf{x}_{d,N_j}, y_{d.N_j})$ for outputs d where, $|\mathbf{d}| = N_j$ (where, N_j is the number of sensing agents). From this, we obtain a vector of sampling locations for each sensing vehicle j: $\mathbf{X} = \{\mathbf{X}_j\}_{j=1}^{D} = \mathbf{X}_1, \ldots, \mathbf{X}_D$, where $\mathbf{X}_d = \{\mathbf{x}_{d,n}\}_{n=1}^{N}$. In this general sense, a separate process f_d can be learned by training set S_d, where, $\mathbf{f}(\mathbf{X}) = \left(f_1\left(\mathbf{x}_{1,1}\right), \ldots, f_1\left(\mathbf{x}_{1,N}\right)\right), \ldots, \left(f_D\left(\mathbf{x}_{D,1}\right), \ldots, f_D\left(\mathbf{x}_{D,N}\right)\right)$ and where N represents the number of query points.

We can construct a similar formulation for a vector-valued GP as in Equation 1. The vector-valued kernel $\mathbf{K}$ is an $ND \times ND$ with entries: $\left(\mathbf{K}\left(\mathbf{x}_i, \mathbf{x}_j\right)\right)_{d,d'}$, for $i, j = 1, \ldots, N$ and $d, d' = 1, \ldots, D$.

We can consider a *separable* kernel function, formulated as a sum of products between a kernel function for the *input space alone* (the spatiotemporal kernel in this work) and a *task* kernel function that encodes interactions between the outputs (correlations between the sensor models, in this work). Such a kernel can be defined as the Hadamard product of an input kernel and a task kernel and takes the form:

$$\left(\mathbf{K}\left(\mathbf{x}, \mathbf{x}'\right)\right)_{d,d'} = k(\mathbf{x}, \mathbf{x}')k_T(d, d') \quad (4)$$

where k is the input kernel, k_T is the task kernel, both defined over $\mathcal{X} \times \mathcal{X}$ with $\{1, \ldots, D\} \times \{1, \ldots, D\}$. Equivalently, this can be written as a matrix expression:

$$\mathbf{K}\left(\mathbf{x}, \mathbf{x}'\right) = k\left(\mathbf{x}, \mathbf{x}'\right)\mathbf{B} \quad (5)$$

We establish an index kernel, defined by a lookup table of indices corresponding to the number of tasks. In this paper, the number of tasks is equal to the number of sensing agents (N_v), that is: $N_v = N_j = D$:

$$k(i,j) = \left(BB^\top + (\mathbf{I}v)_{i,j}\right) \quad (6)$$

where B is a low-rank matrix that establishes the variance between tasks and v is a positive constraint on the inter-task variance. Refer to [1] for a detailed treatment of kernel functions for multi-output GPs and [8] for details of the function `gpytorch.kernels.IndexKernel` used in our implementation.

It is important to note that the vehicles do not attempt to construct the posterior variance from the perspective of other vehicles in the team. For each vehicle, only the output corresponding to the ego vehicle is used for planning. The other outputs could be used to infer model variance from the perspective of other sensing agents, and could be used to project their probable next-actions. While an interesting area

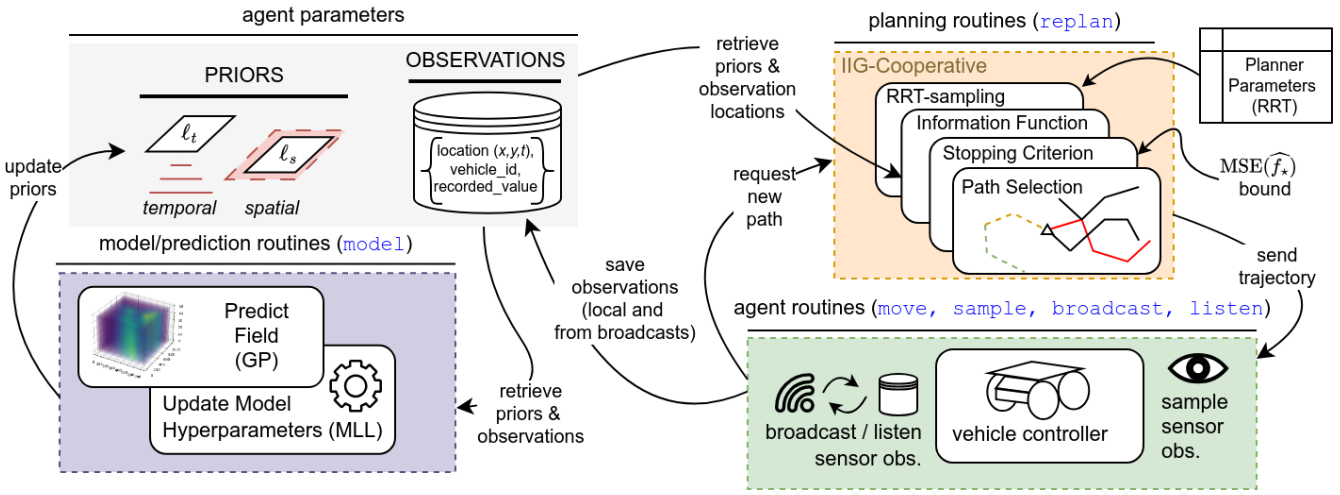

Fig. 2: Schematic overview of the adaptive planning routine. Not pictured: When the vehicle controller has reached its movement budget (time or distance), the modeler is called one last time to produce a final map prediction.

for future work, this consideration is outside of the scope of this study.

D. Utility Formulation

Understanding that the differential entropy of a Gaussian random variable is a monotonic function of its variance, we can construct a utility function based on the reduction of map entropy H, given a new observation Z. We derive this from the posterior variance, which is obtained from Equation 2. Specifically, we evaluate the information gain of a new proposed sampling location, using the mutual information I between the current training set X and a new set X', containing the observation Z.

For a random vector of observations $\mathbf{X} = (X_1, \ldots, X_n)$, for every X_i the mutual information becomes:

$$\hat{I}^{[i]}(X_i; Z) = \frac{1}{2}\left[\log\left(\sigma_{X_i}\right) - \log\left(\sigma_{X_i|Z}\right)\right] \tag{7}$$

and can be calculated as the sum of marginal variances at i: $\hat{I}(X;Z) = \sum_{i=1}^{n} \hat{I}^{[i]}(X_i;Z)$. Refer to [4] for additional context and see [13] for a derivation. Crucially, we are able to obtain from the covariance matrix produced by the kernel $\mathbb{V}[X_i] = K^{[i,i]}$.

E. Path selection and planning

1) General framework: Path planning proceeds according to the procedure described in [13] (*IIG*) and [4] (*IIG-ST*), using the task-aware covariance function described in the subsection III-C. The complete set of parameters used by the planner can be found in the accompanying video and code repository.

2) Convergence criterion and path selection: We utilize the convergence criterion from [4] in the re-planning stage, to establish when the agent should stop adding new proposed locations to the RIG tree and switch to path generation. Path generation is performed using a vote-based heuristic from [13]. We use posterior map variance as a lower bound for mean-square error, given optimal hyperparameters θ for the kernel function.

$$\mathrm{MSE}\left(\widehat{f}_\star\right) \geq \underbrace{\mathbb{V}[f_\star]}_{=\sigma^2_{\star|y}(\theta)} \tag{8}$$

These kernel hyperparameters θ comprise another component of our spatiotemporal and task-priors. Optimal parameters can be chosen in a standard Bayesian approach, using the marginal log likelihood (MLL) of the GP model, when applied to observed data. We employ this approach to update model hyperparameters, as the sensing agents receive observations during the survey mission.

3) Adaptive planning: An overview of the complete adaptive planning routine is visualized in Figure 2. Each vehicle alternates between collecting observations, updating internal parameters and re-planning at a fixed interval until a time budget is elapsed. Each agent broadcasts the value and location of a sample immediately upon collection, and all agents continuously listen for observations from other robots in the team. Each robot re-plans after every 2nd sample.

In the experiments, it assumed that vehicles have access to a wide-area, low-bandwith communication link (such as LoRa [23]) and are able to communicate their observations globally. Prior to the next re-planning procedure, each vehicle evaluates the map expectation $\mathbb{E}$ (posterior mean) from the multi-task model, incorporating all observations $\mathbf{y_j}$ collected by the agent and communicated by other agents. The model output is used to calculate the marginal log likelihood of the GP with respect to the observations collected, which forms the basis of the hyperparameter optimization routine.

IV. Experimental Evaluation and Discussion

In this section, we evaluate our coordinated planner *IIG-Cooperative* against the non-cooperative *IIG-ST* introduced in [4] as applied to the task of surveying and modeling a physical phenomenon that changes in an environment over a fixed period of time. In the first scenario, the surveyed phenomenon advects and diffuses in a simulated fluid environment over time, similar to the dynamics of an environmental contaminant in the soil, water, or air. In the

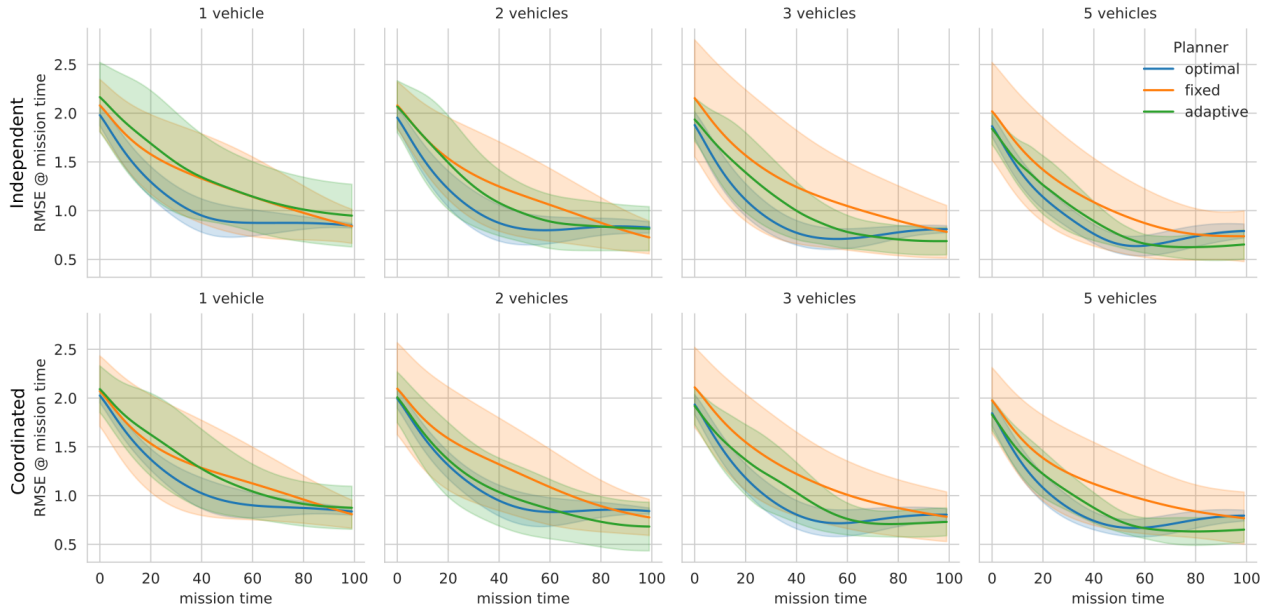

Fig. 3: **Advection/diffusion simulation**: Map error (lower is better) at different moments in the survey mission time for different planning configurations for $IIG - Cooperative$ with varing numbers robots deployed in the sensing task. Planners with fixed priors are run in two configurations, fixed-suboptimal ($\ell_t = 50$ and $\ell_s = 100$) and fixed-optimal ($\ell_t = 20$ and $\ell_s = 30$). The adaptive planner starts with the suboptimal priors and continuously updates model hyperparameters throughout the survey mission. All planners are run with global communication (bottom row) and without communication, in an independent planning scenario (top row).

second scenario an oceanic turbidity dataset is used as the target for the surveying objective.

We evaluate our planner according to the objective that would be most salient to a surveyor; that is, the accuracy of the final map representation at the end of the survey period ($t = 100$) with a root-mean squared error metric (*RMSE*). We also consider the auxiliary objective of making predictions at arbitrary points in time. This is relevant if the operator wishes to reconstruct the dynamics of the system. However, while our spatiotemporal planners incorporate time into the planning objective, the robot and sensor can obviously only travel forward through the temporal dimension.

A. Experimental setting

The simulated world implements advection and diffusion according to the Navier- Stokes equations for an incompressible fluid (forward-differencing discretization). All robots are initialized with the same path planner with fixed planning parameters that are updated independently during the course of a survey. The full table of parameters set for the planner can be found in the accompanying video and the code repository. We executed the experiments in a GNU/Linux environment on a 3.6 GHz Intel i7-4790 computer with 10 GB of RAM available. RRT planning procedures were derived from [21] and GP posterior variance and final map predictions were executed with GPyTorch [8], run without hardware acceleration so as to simulate the resources available on a mobile robot.

B. Comparison of cooperative and adaptive planning

To assess the effectiveness of the cooperative and adaptive planners, we compare the planners in three configurations: A planner with fixed model lengthscales (ℓ) chosen proportional to the size of the world ("fixed" planner, $\ell_s = 100, \ell_t = 50$); an adaptive planner that starts with the same hyperparameters as the fixed planner, but is allowed to update as observations are collected; and a fixed, planner given parameters determined through hyperparameter optimization along a dense, coverage path plan ("optimal" planner, $\ell_s = 30, \ell_t = 20$).

A summary of the main results can be found in Figure 3, which presents the RMSE between the state of the world at time t and the state of the predicted world at time t, constructed after observations have been compiled by *all sensing vehicles* at the end of the survey mission. This aligns with the typical mode of operation in multi-agent surveys. All scenarios are run with the multi-task planner, which degenerates into a single-task problem when $n = 1$ sensing vehicle, and/or when there is no transfer of information between the vehicles (in other words, there is no information for the other input tasks). This later case is labeled as "independent" planning.

In Figure 3, lower error values found on average across nearly the entire mission envelope for the adaptive planner, although it does not produce as accurate a map posterior as if it were given "optimal" hyperparamters. A notable observation is the map accuracy for the final time index of the survey period, when all planners converge to similar performance, with the "optimal" planner producing a slightly less accurate representation at $t = 100$ than the adaptive planner (or even the fixed planner, in some configurations). There are a few possible explanations for this phenomenon: 1. For large numbers of vehicles, by $t = 100$, the environment has been uniformly surveyed, regardless of the direction

			Coordinated		Independent	
Planner	Function	N_v	t_{100}	t_{all}	t_{100}	t_{all}
Adaptive	*IIG*	1	0.888	1.344	0.935	1.367
		2	0.670	1.104	0.827	1.234
		3	0.749	1.109	0.703	1.091
		5	0.664	0.988	0.689	**0.999**
	IIG-ST	1	0.845	1.263	0.983	1.477
		2	0.703	1.174	0.797	1.181
		3	0.683	1.080	0.650	1.143
		5	**0.618**	**0.986**	**0.566**	1.015
fixed	*IIG*	1	0.820	1.321	0.839	1.374
		2	0.787	1.372	0.710	1.225
		3	0.806	1.312	0.782	1.374
		5	0.778	1.197	0.791	1.177
	IIG-ST	1	0.782	1.293	0.841	1.294
		2	0.760	1.247	0.749	1.365
		3	**0.726**	1.203	0.785	1.205
		5	0.745	**1.161**	**0.648**	**1.149**
optimal	*IIG*	1	0.841	1.162	0.841	1.128
		2	0.839	1.093	0.842	1.069
		3	0.825	1.043	0.802	0.986
		5	0.802	**0.946**	**0.787**	**0.969**
	IIG-ST	1	0.833	1.164	0.851	1.113
		2	0.843	1.143	0.812	1.063
		3	**0.783**	0.936	0.818	0.984
		5	0.787	0.963	0.796	0.972

TABLE I: Fluid Simulation: Summary of average map error (*RMSE*) produced by observations collected by all vehicles v at the end of a survey mission. Error is represented across entire survey envelope (t_{all}) and for the last time step of the survey mission ($t = 100$). Lowest values within a given configuration are emphasized in bold.

given by the informative planner. 2. The "optimal" parameters were derived through sampling along a conventional, coverage path, of a distance equal to the average distance traveled by the informed planners. This deterministic route does not guarantee representative samples across the entire spatiotemporal domain (an thus are not strictly optimal). This demonstrates another weakness of traditional surveying procedures. 3. The "optimal" parameters were determined through a maximum-likelihood estimator, optimized for the predictive ability across the *entire* survey envelope, and not *solely* on the final map state. Likely, different results would be obtained for parameters chosen to minimize the error of the final map state. Reconciling these potentially competing objectives within a unified planning framework could be explored in future studies.

A summary of the performance for all planning configurations is presented in Table I, where the two objectives are presented under: t_{100} for the predictive accuracy at $t = 100$ and t_{all}, for the predictive accuracy across the entire survey envelope using the coordinated planner (*IIG-Cooperative*) and the independent planning scenario. We also explore the addition of the spatiotemporal kernel to each configuration ("Function" column). Notably, the adaptive planner often outperforms the non-adaptive configuration, given informed "optimal" priors. This is due to the reasons outlined in the previous paragraph. In multi-robot configurations without coordination, minimal to no gains are found with the spatiotemporally-informed planner across most configurations. This is likely due to the aggregate effect of a more even coverage of the environment that occurs with time-naive robots, at high numbers.

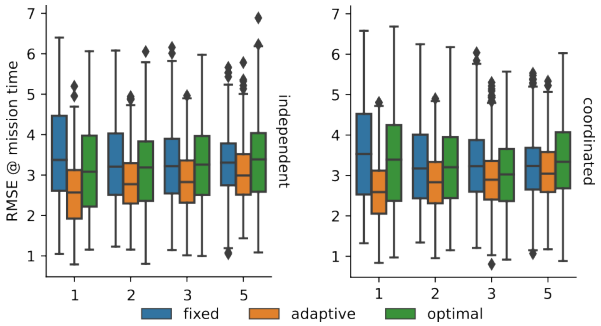

Fig. 4: Ocean turbidity: Summary of average map error produced for different planner configurations and varying number of survey vehicles.

C. Environmental monitoring scenario

As a proof of concept, we demonstrate our planners in a synoptic-scale ocean monitoring experiment, using ocean reflectance off the west coast of California from the Moderate Resolution Imaging Spectroradiometer as a turbidity proxy [9]. Our simulated vehicles are configured with velocities congruent to the Wave Glider autonomous aquatic vehicle, based on the reported long-mission average speed of 1.5 knots (approximately 330 km per week) [**?**]. Figure 4 presents the results for different configuration of the coordinated and independent planners, configured with the spatiotemporal kernel. With a sampling rate of 1/week for a ~50 week interval, inputs were relatively sparse and no distinguishing trends were observed between the performance of the cooperative and independent planners.

As with the synthetic simulation, a basis of comparison for the "optimal" fixed planner was established with hyperparameters collected by a single-vehicle lemniscatic coverage. As with the previous experiments, the adaptive planner results in improved map accuracy for all configurations across the survey envelope.

V. Conclusions

This work presented a novel integration of multi-vehicle informative path planning, informed by both model uncertainty and information from other sensing agents, which may differ in their contribution toward reducing model uncertainty. We also demonstrated how this approach could be utilized in an adaptive planning framework. Finally, we quantified the effectiveness of our planning approach, grounded in map accuracy, the salient objective for the survey operator.

This work illuminates multiple avenues for future work. Given global communication, each robot could directly integrate observations from all other robots, using a distributional or variational kernel. Robots could also attempt to project likely *future* movements by other agents in the team, using a similar utility metric grounded in model uncertainty. Observations collected in these experiments used a point sensor model, where the robot obtained a low-noise observation of the environment directly beneath the vehicle. Multi-task planning scenarios can incorporate variable sensor models, including variable noise, observation windows, and observation frequency.

REFERENCES

[1] Mauricio A. Álvarez, Lorenzo Rosasco, and Neil D. Lawrence. Kernels for Vector-Valued Functions: A Review. *Foundations and Trends in Machine Learning*, 4(3):195–266, 2012.

[2] Jonathan Binney, Andreas Krause, and Gaurav S. Sukhatme. Optimizing waypoints for monitoring spatiotemporal phenomena. *The International Journal of Robotics Research*, 32(8):873–888, July 2013.

[3] Lorenzo Booth and Stefano Carpin. Distributed estimation of scalar fields with implicit coordination. In *Distributed Autonomous Robotic Systems*, Springer Proceedings in Advanced Robotics, Monbellard, FR, 2023. Springer International Publishing.

[4] Lorenzo A. Booth and Stefano Carpin. Informative path planning for scalar dynamic reconstruction using coregionalized Gaussian processes and a spatiotemporal kernel. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, October 2023.

[5] Merrick Campbell, Amel Dechemi, and Konstantinos Karydis. An Integrated Actuation-Perception Framework for Robotic Leaf Retrieval: Detection, Localization, and Cutting. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9210–9216, October 2022.

[6] Felix Esser, Radu Alexandru Rosu, André Cornelißen, Lasse Klingbeil, Heiner Kuhlmann, and Sven Behnke. Field Robot for High-Throughput and High-Resolution 3D Plant Phenotyping: Towards Efficient and Sustainable Crop Production. *IEEE Robotics & Automation Magazine*, 30(4):20–29, December 2023.

[7] Jan N. Fuhg, Amélie Fau, and Udo Nackenhorst. State-of-the-Art and Comparative Review of Adaptive Sampling Methods for Kriging. *Archives of Computational Methods in Engineering*, 28(4):2689–2747, June 2021.

[8] Jacob R. Gardner, Geoff Pleiss, David Bindel, Kilian Q. Weinberger, and Andrew Gordon Wilson. GPyTorch: Blackbox matrix-matrix Gaussian process inference with GPU acceleration. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, pages 7587–7597, Red Hook, NY, USA, December 2018. Curran Associates Inc.

[9] NASA Ocean Biology Processing Group. MODIS-Aqua Level 3 Mapped Particulate Organic Carbon Data Version R2018.0, 2017.

[10] Carlos Guestrin, Andreas Krause, and Ajit Paul Singh. Near-optimal sensor placements in Gaussian processes. In *Proceedings of the 22nd International Conference on Machine Learning - ICML '05*, pages 265–272, Bonn, Germany, 2005. ACM Press.

[11] Gregory Hitz, Enric Galceran, Marie-Ève Garneau, François Pomerleau, and Roland Siegwart. Adaptive continuous-space informative path planning for online environmental monitoring. *Journal of Field Robotics*, 34(8):1427–1449, 2017.

[12] Geoffrey A. Hollinger and Gaurav S. Sukhatme. Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research*, 33(9):1271–1287, August 2014.

[13] Maani Ghaffari Jadidi, Jaime Valls Miro, and Gamini Dissanayake. Sampling-based incremental information gathering with applications to robotic exploration and environmental monitoring:. *The International Journal of Robotics Research*, 38(6):658–685, April 2019.

[14] Yiannis Kantaros, Brent Schlotfeldt, Nikolay Atanasov, and George J. Pappas. Asymptotically Optimal Planning for Non-Myopic Multi-Robot Information Gathering. In *Robotics: Science and Systems XV*. Robotics: Science and Systems Foundation, June 2019.

[15] Jeongeun Kim, Seungwon Kim, Chanyoung Ju, and Hyoung Il Son. Unmanned Aerial Vehicles in Agriculture: A Review of Perspective of Platform, Control, and Applications. *IEEE Access*, 7:105100–105115, 2019.

[16] Steven M. Lavalle. Rapidly-Exploring Random Trees: A New Tool for Path Planning. Technical report, Iowa State University, Computer Science Deptartment, October 1998.

[17] Joshua Ott, Edward Balaban, and Mykel J. Kochenderfer. Sequential Bayesian Optimization for Adaptive Informative Path Planning with Multimodal Sensing. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7894–7901, May 2023.

[18] Marija Popović, Joshua Ott, Julius Rückin, and Mykel J. Kochenderfer. Learning-based methods for adaptive informative path planning. *Robotics and Autonomous Systems*, 179:104727, September 2024.

[19] Marija Popovic, Teresa Vidal-Calleja, Jen Jen Chung, Juan Nieto, and Roland Siegwart. Informative Path Planning for Active Field Mapping under Localization Uncertainty. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10751–10757, Paris, France, May 2020. IEEE.

[20] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, Mass, 2006.

[21] Atsushi Sakai, Daniel Ingram, Joseph Dinius, Karan Chawla, Antonin Raffin, and Alexis Paques. PythonRobotics: A Python code collection of robotics algorithms. August 2018.

[22] Matthew A. Schack, John G. Rogers, Qi Han, and Neil Dantam. Optimizing Non-Markovian Information Gain Under Physics-Based Communication Constraints. *IEEE Robotics and Automation Letters*, 6(3):4813–4819, July 2021.

[23] Jothi Prasanna Shanmuga Sundaram, Wan Du, and Zhiwei Zhao. A Survey on LoRa Networking: Research Problems, Current Solutions, and Open Issues. *IEEE Communications Surveys & Tutorials*, 22(1):371–388, 2020.

[24] Michael L. Stein. *Interpolation of Spatial Data: Some Theory for Kriging*. Springer Series in Statistics. Springer-Verlag, New York, 1999.

[25] Thomas C. Thayer, Stavros Vougioukas, Ken Goldberg, and Stefano Carpin. Multirobot Routing Algorithms for Robots Operating in Vineyards. *IEEE Transactions on Automation Science and Engineering*, 17(3):1184–1194, July 2020.

[26] Samuel Yanes Luis, Manuel Perales Esteve, Daniel Gutiérrez Reina, and Sergio Toral Marín. Deep Reinforcement Learning Applied to Multi-agent Informative Path Planning in Environmental Missions. In Ahmad Taher Azar, Ibraheem Kasim Ibraheem, and Amjad Jaleel Humaidi, editors, *Mobile Robot: Motion Control and Path Planning*, pages 31–61. Springer International Publishing, Cham, 2023.