Informative Planning for Worst-Case Error Minimisation in Sparse Gaussian Process Regression

Jennifer Wakulicz, Ki Myung Brian Lee, Chanyeol Yoo, Teresa Vidal-Calleja and Robert Fitch

***Abstract*— We present a planning framework for minimising the deterministic worst-case error in sparse Gaussian process (GP) regression. We first derive a universal worst-case error bound for sparse GP regression with bounded noise using interpolation theory on reproducing kernel Hilbert spaces (RKHSs). By exploiting the conditional independence (CI) assumption central to sparse GP regression, we show that the worst-case error minimisation can be achieved by solving a posterior entropy minimisation problem. In turn, the posterior entropy minimisation problem is solved using a Gaussian belief space planning algorithm. We corroborate the proposed worst-case error bound in a simple 1D example, and test the planning framework in simulation for a 2D vehicle in a complex flow field. Our results demonstrate that the proposed posterior entropy minimisation approach is effective in minimising deterministic error, and outperforms the conventional measurement entropy maximisation formulation when the inducing points are fixed.**

I. INTRODUCTION

Reconstructing a spatial field from sparse, noisy measurements is an important fundamental problem in robotics. The problem naturally arises in many practical applications, such as oceanography [1] and agriculture [2], and in general tasks such as robot navigation [3, 4]. We are interested in informative path planning that enables robots to collect measurements for spatial field reconstruction with quality guarantees, such as minimising worst-case error. We present an approach using sparse Gaussian process (GP) regression that is inspired by results in interpolation theory on reproducing kernel Hilbert spaces (RKHS).

GP regression [5] is a powerful machine learning technique for modelling spatially correlated phenomena. It has been widely used in the robotics community to estimate a variety of spatial fields including obstacles [3, 4], infrastructure [6], and agricultural [2] or oceanographic data [1, 7]. A well-known challenge in robotics application is that the computational complexity of GP regression scales cubically with the size of the input data.

Sparse GP approaches mitigate this computational challenge by adopting simplifying approximations [8–10]. One such approximation is the *inducing points* formulation [8], where the target function is assumed to be conditionally independent given the function values at a fixed set of inducing points. A recent advance in this direction is that incoming sensor measurements can be 'fused' via recursive Bayesian estimation of a latent Gaussian state of *fixed* dimensionality, reminiscent of Kalman filtering [2, 6, 11–14].

In this paper, we show that the belief maintained by a recursive Bayesian estimator is sufficient for planning paths that minimise the worst-case error in sparse GP regression with bounded noise. This result arises from interpolation theory on RKHSs, and thus assumes that the target function resides in an RKHS. This approach can be viewed as reducing the information gathering problem to a Gaussian belief space planning problem.

We present our algorithm for active spatial field reconstruction and initially demonstrate the error bound in an abstract 1D example of sparse GP regression. Then, we present results from a simulated example of path planning for an underwater robot operating in a flow field and collecting scalar-valued measurements. For comparison, we demonstrate the behaviour of a typical measurement entropy maximisation approach and note that the error bound is non-decreasing over time. The significance of these results is to illuminate the limitations of existing informative path planning approaches in terms of solution quality, and to contribute a new method that achieves a worst-case solution quality guarantee for actively reconstructing spatial fields.

II. RELATED WORK

Path planning for optimal reconstruction of a GP is typically posed as an information gain or marginal entropy maximisation problem [15]. These classify as a submodular maximisation problem, which is NP-hard [16]. When the measurement locations are subject to a dynamics constraint, a non-myopic search is necessary [17], which can be achieved by, e.g., growing a search tree with an appropriate pruning condition [4, 18, 19].

Instead of the abstract information-theoretic quantities, we present an orthogonal approach that minimises the worst-case error in a *deterministic* sense, inspired by the interpolation theory on RKHSs [20, 21]. Nonetheless, we show that the worst-case error minimisation problem admits an information-theoretic analogue that is an instance of the Gaussian belief space planning [22–24]. Gaussian belief space planning is a more restricted class of problems than general submodular maximisation, and more solution algorithms are available with stronger performance guarantees. In particular, we adopt the approach of [23] without loss of their guarantees.

This work is supported by an Australian Government Research Training Program (RTP) Scholarship and the University of Technology Sydney.

Authors are with the University of Technology Sydney, Ultimo, NSW 2006, Australia {jennifer.wakulicz, brian.lee}@student.uts.edu.au, {chanyeol.yoo,teresa.vidalcalleja, rfitch}@uts.edu.au

III. PROBLEM FORMULATION

Suppose we have a mobile robot r that operates in D-dimensional environment $\mathcal{X} \subseteq \mathbb{R}^D$. The motion of robot r is described by a discrete-time non-linear dynamic model

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

where $\mathbf{x}_t \in \mathcal{X}$ is the D-dimensional state of the robot in the environment and $\mathbf{u}_t \in \mathcal{U} \subseteq \mathbb{R}^{D_u}$ is the D_u-dimensional control input at time t. A sequence of robot states is denoted as $\mathbf{X} = \mathbf{x}_0 \ldots$. Likewise, a sequence of N control actions is denoted as $\mathbf{U} = \mathbf{u}_0 \ldots$.

Using an onboard sensor, robot r at $\mathbf{x}_t$ can make a noisy measurement $y_t \in \mathbb{R}$ of a scalar spatial phenomenon of interest. With measurement noise $\epsilon_t \in \mathbb{R}$,

$$y_t = s(\mathbf{x}_t) + \epsilon_t. \tag{2}$$

Given a sequence of measurements from time 0 to t denoted as $\mathbf{y}_t$, the estimate of spatial phenomenon s over $\mathcal{X}$ is denoted as $\hat{s}(\mathbf{x} \mid \mathbf{y}_t)$. The deterministic error between true spatial phenomenon s and estimate $\hat{s}$ is the point-wise absolute difference, defined as

$$E(\mathbf{x} \mid \mathbf{y}_N) = |s(\mathbf{x}) - \mathbb{E}[\hat{s}(\mathbf{x} \mid \mathbf{y}_N)]|. \tag{3}$$

The objective of this paper is to find a sequence of control actions $\mathbf{U}$ over time horizon N that minimises overall deterministic error E over entire domain $\mathcal{X}$. The formal problem statement is found below.

Problem 1. *Given the dynamic model in* (1) *and the measurement model in* (2)*, find a sequence of control actions* $\mathbf{U}^* = \{\mathbf{u}_0 \ldots \mathbf{u}_{N-1}\}$ *that minimises the total deterministic error* (3) *after time-step* N *over the domain* $\mathcal{X}$*:*

$$\mathbf{U}^* = \underset{\mathbf{U}=\mathbf{u}_0 \ldots \mathbf{u}_{N-1}}{\arg\min} \int_{\mathcal{X}} E(\mathbf{x} \mid \mathbf{y}_N) d\mathbf{x} \tag{4}$$

Difficulty arises in solving Problem 1 in realistic scenarios as often full knowledge of the ground truth $s(\mathbf{x})$ is not available. Deterministic error and thus the integral in Problem 1 cannot be evaluated directly in such cases.

IV. WORST-CASE ERROR MINIMISATION AND INFORMATION GATHERING

The aforementioned difficulty in solving Problem 1 can be side-stepped by bounding deterministic error (3) with an expression independent of ground truth knowledge. Here we leverage sparse GP estimators and kernel-based interpolation theory to find such a bound, in turn reducing Problem 1 to a new entropy-based minimisation problem that is tractable.

A. Sparse GP Regression

Given a sequence of noisy measurements $\mathbf{y}_N$, we generate an estimate $\hat{s}(\mathbf{x} \mid \mathbf{y}_N)$ of the spatial phenomenon $s(\mathbf{x})$ using GP regression with sparse approximation. GP is a generalisation of multivariate Gaussian random variables to *random functions*. A GP $s(\mathbf{x}) \sim GP(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$ is characterised by the mean and covariance, or *kernel* functions [5]:

$$\begin{aligned} \mathbb{E}[s(\mathbf{x})] &= m(\mathbf{x}), \\ \mathrm{Cov}[s(\mathbf{x}), s(\mathbf{x}')] &= k(\mathbf{x}, \mathbf{x}'). \end{aligned} \tag{5}$$

We impose a zero-mean GP prior on the scalar field of interest $s(\mathbf{x}) \sim GP(0, k(\mathbf{x}, \mathbf{x}'))$. Let $\mathbf{y}_N$ be a vector containing noisy measurements up to time-step N as per measurement model (2), i.e. $[\mathbf{y}_N]_i = y_i$. With the zero-mean prior, the estimate $\hat{s}(\mathbf{x} \mid \mathbf{y}_N)$ given the measurements $\mathbf{y}_N$ is given by another GP [25]:

$$\begin{aligned} \hat{s}(\mathbf{x} \mid \mathbf{y}_N) &\sim GP(\mu(\mathbf{x} \mid \mathbf{y}_N), \sigma^2(\mathbf{x}, \mathbf{x}' \mid \mathbf{y}_N)), \\ \mu(\mathbf{x} \mid \mathbf{y}_N) &= \mathbf{k}_{\mathbf{X}}^{\mathrm{T}}(\mathbf{x}) K_{\mathbf{X}}^{-1} \mathbf{y}_N, \\ \sigma^2(\mathbf{x}, \mathbf{x}' \mid \mathbf{y}_N) &= k(\mathbf{x}, \mathbf{x}') - \mathbf{k}_{\mathbf{X}}^{\mathrm{T}}(\mathbf{x}) K_{\mathbf{X}}^{-1} \mathbf{k}_{\mathbf{X}}(\mathbf{x}'), \end{aligned} \tag{6}$$

where $[\mathbf{k}_{\mathbf{X}}(\mathbf{x})]_i = k(\mathbf{x}, \mathbf{x}_i)$, and $[K_{\mathbf{X}}]_{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)$. We use the same notation for other sets throughout the paper.

We use the *inducing point*-based approximation of the regression (6) introduced in [8]. Intuitively, this formulation assumes that there is a set of inducing points $\mathbf{Z} = \{\mathbf{z}_i\}_{i=1}^M$ whose function values $[\mathbf{y}_{\mathbf{Z}}]_i = s(\mathbf{z}_i)$ 'summarise' the entire set of measurements $\mathbf{y}_N$. We refer to these function values $\mathbf{y}_{\mathbf{Z}}$ as *inducing measurements*. One way to assert this is to impose that $\mathbf{y}_N$ are *conditionally independent* given $\mathbf{y}_{\mathbf{Z}}$. This means correlation between any two measurements is indirect and is limited by their correlation to the inducing measurements.

A GP satisfies the CI property if the kernel function satisfies:

$$k(\mathbf{x}, \mathbf{x}') = \mathbf{k}_{\mathbf{Z}}^T(\mathbf{x}) K_{\mathbf{Z}}^{-1} \mathbf{k}_{\mathbf{Z}}(\mathbf{x}') \quad \forall \mathbf{x} \neq \mathbf{x}', \tag{7}$$

which follows from asserting that the conditional cross-covariance vanishes given inducing measurements $\mathbf{y}_{\mathbf{Z}}$, i.e. $\sigma^2(\mathbf{x}, \mathbf{x}' \mid \mathbf{y}_{\mathbf{Z}}) = 0$.

We consider two inducing point-based approximations that satisfy CI: the subset of regressors (SoR) and fully independent conditional (FIC) approximations. As noted in [8], the SoR and FIC approximations are equivalent to replacing the kernel $k(\mathbf{x}, \mathbf{x}')$ with approximate ones as follows:

$$\begin{aligned} &\hat{k}_{SoR}(\mathbf{x}, \mathbf{x}') = \mathbf{k}_{\mathbf{Z}}^{\mathrm{T}}(\mathbf{x}) K_{\mathbf{Z}}^{-1} \mathbf{k}_{\mathbf{Z}}(\mathbf{x}'), \\ &\hat{k}_{FIC}(\mathbf{x}, \mathbf{x}') = \hat{k}_{SoR}(\mathbf{x}, \mathbf{x}') + \delta(\mathbf{x}, \mathbf{x}')(k(\mathbf{x}, \mathbf{x}') - \hat{k}_{SoR}(\mathbf{x}, \mathbf{x}')). \end{aligned} \tag{8}$$

Here, $\delta(\cdot)$ is the Kronecker delta function.

B. Kernel-based Interpolation and GP Regression

To derive bounds on the deterministic error (3) using a sparse GP estimator, we leverage kernel-based interpolation theory [21] and its connections to GP regression [25]. Kernel-based interpolation and GP regression are built on the theory of reproducing kernel Hilbert spaces (RKHSs) alike. We therefore provide a brief introduction to RKHSs.

Any positive-definite kernel $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ uniquely induces an RKHS $\mathcal{H}_k$. An RKHS is a linear space of real-valued functions on the set $\mathcal{X}$ equipped with an inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}_k}$ and a unique kernel function $k(\cdot, \mathbf{x})$ such that: 1)

$k(\cdot, \mathbf{x}) \in \mathcal{H}_k \; \forall \mathbf{x} \in \mathcal{X}$, and 2) $\langle f, k(\cdot, \mathbf{x}) \rangle_{\mathcal{H}_k} = f(\mathbf{x})$. That is, the kernel function $k(\cdot, \mathbf{x})$ is a linear function that evaluates all other functions in $\mathcal{H}_k$ at $\mathbf{x}$. The inner product induces a norm $||g||_{\mathcal{H}_k} = \sqrt{\langle g, g \rangle_{\mathcal{H}_k}}$.

With the RKHS framework at hand, one may draw connections between kernel-based interpolation and GP regression. In particular, the equivalence of the minimum-norm kernel-based interpolant and the predictive mean of a GP regressor [26] is of interest here. It can be shown that the predictive mean of GP regression (6) is equivalent to RKHS interpolation with noise-free observations in the form $\min_{g \in \mathcal{H}_k} ||g||_{\mathcal{H}_k}$ such that $\forall \mathbf{x}_i \in \mathbf{X}$, $g(\mathbf{x}_i) = s(\mathbf{x}_i)$ [25]. Note that Gaussian noise in GP regression can be included in the kernel.

C. Worst-Case Error Bounds

The equivalence described above gives a kernel-based framework to bound the deterministic error of a full GP regressor. The following is an extension of the results of [21, 25] to the case of bounded measurement error.

Theorem 1. *Suppose $s \in \mathcal{H}_k$ with arbitrary positive definite kernel k. With bounded measurement noise $\epsilon^2 < \sigma_\epsilon^2$,*

$$E(\mathbf{x} \mid \mathbf{y}_N) \leq ||s||_{\mathcal{H}_k} P_{\mathbf{X}}(\mathbf{x}) + \sqrt{\sigma_\epsilon^2 N \Lambda_k^2(\mathbf{x})}, \tag{9}$$

where $P_{\mathbf{X}}(\mathbf{x}) = \sqrt{\sigma^2(\mathbf{x}, \mathbf{x} \mid \mathbf{y}_N)}$ is called the power function of $\mathbf{X}$ and $\Lambda_k(\mathbf{x}) = ||K_{\mathbf{X}}^{-1} \mathbf{k}_{\mathbf{X}}(\mathbf{x})||$.

From Theorem 1, it is clear that Problem 1 can be solved by choosing a set of measurement points $\{\mathbf{x}_t\}$ that minimise $P_{\mathbf{X}}(\mathbf{x})$ for all possible $\mathbf{x}$. A common approach in interpolation theory is to use the relationship [21]:

$$P_{\mathbf{X}}(\mathbf{x}) = \sqrt{\frac{\det K_{\mathbf{X} \cup \{\mathbf{x}\}}}{\det K_{\mathbf{X}}}}, \tag{10}$$

and reduce $P_{\mathbf{X}}(\mathbf{x})$ by maximising the denominator $\det K_{\mathbf{X}}$, which is independent of query point $\mathbf{x}$. Because Theorem 1 holds for an arbitrary kernel k, the same approach should hold true for SoR and FIC approximations by using the approximate kernels (8). In fact, maximising $\det K_{\mathbf{X}}$ is equivalent to measurement entropy maximisation from the informative path planning literature, e.g., [15]. However, this is still an unsatisfying answer, because 1) the numerator $\det K_{\mathbf{X} \cup \{\mathbf{x}\}}$ still varies with the choice of measurements, and 2) the choice of inducing points also affects $K_{\mathbf{X}}$.

To mitigate this issue, we exploit the CI property of sparse approximations (7). Because CI kernels can be viewed as interpolants to the true kernel [9], the interpolation of $s(\mathbf{x})$ given $\mathbf{y}_N$ can be decomposed into two stages: 1) the interpolation of inducing measurements $\mathbf{y}_{\mathbf{Z}}$ given $\mathbf{y}_N$, and 2) the interpolation of $s(\mathbf{x})$ given $\mathbf{y}_{\mathbf{Z}}$. Then, it is natural to ask if the deterministic error (9) or the power function (10) admits a similar decomposition. The following theorem confirms that there is such a decomposition.

Theorem 2. *Suppose a kernel k satisfies the CI assumption (7). Then, the power function $P_{\mathbf{X}}(\mathbf{x})$ satisfies:*

$$P_{\mathbf{Z}}(\mathbf{x}) \leq P_{\mathbf{X}}(\mathbf{x}) \leq P_{\mathbf{Z}}(\mathbf{x}) \exp H(\mathbf{y}_{\mathbf{Z}} \mid \mathbf{y}_N), \tag{11}$$

where $P_{\mathbf{Z}}(\mathbf{x})$ is the power function of $\mathbf{Z}$ as per (10) and $H(\mathbf{y}_{\mathbf{Z}} \mid \mathbf{y}_N)$ is the posterior entropy of $\mathbf{y}_{\mathbf{Z}}$ given $\mathbf{y}_N$:

$$H(\mathbf{y}_{\mathbf{Z}} \mid \mathbf{y}_N) = \frac{1}{2} \log \left((2\pi e)^M \frac{\det K_{\mathbf{Z} \cup \mathbf{X}}}{\det K_{\mathbf{X}}} \right). \tag{12}$$

Moreover, assuming $s \in \mathcal{H}_k$, the deterministic error (3) can be further bounded as:

$$E(\mathbf{x} | \mathbf{y}_N) \leq ||s||_{\mathcal{H}_k} P_{\mathbf{Z}}(\mathbf{x}) \exp H(\mathbf{y}_{\mathbf{Z}} \mid \mathbf{y}_N) + \sqrt{\sigma_\epsilon^2 N \Lambda_k^2(\mathbf{x})} \tag{13}$$

Note that, in the upper bound of (11), $P_{\mathbf{Z}}(\mathbf{x})$ is independent of the robot's actions, and $H(\mathbf{y}_{\mathbf{Z}} \mid \mathbf{y}_N)$ is independent of the query point $\mathbf{x}$. Therefore, we can minimise $P_{\mathbf{X}}(\mathbf{x})$ for all $\mathbf{x}$ by solving the following surrogate problem:

Problem 2. *Given the dynamic model (1) and the measurement model (2), find a sequence of control actions $\mathbf{u}_1 ... \mathbf{u}_N$ that minimises the posterior entropy of inducing measurements:*

$$\min_{\mathbf{u}_1, \ldots, \mathbf{u}_N} H(\mathbf{y}_{\mathbf{Z}} \mid \mathbf{y}_N). \tag{14}$$

It is worth noting that Problem 2 becomes equivalent to Problem 1 as $H(\mathbf{y}_{\mathbf{Z}} \mid \mathbf{y}_N) \to 0$ assuming the noise-free case. This is because the upper bound on power function (11) is 'tight' in the sense that $P_{\mathbf{X}}(\mathbf{x}) \to P_{\mathbf{Z}}(\mathbf{x})$ as $H(\mathbf{y}_{\mathbf{Z}} \mid \mathbf{y}_N) \to 0$, and $P_{\mathbf{Z}}(\mathbf{x})$ is independent of the control actions. Another interesting observation is that the power function $P_{\mathbf{X}}(\mathbf{x})$ can only be reduced down to $P_{\mathbf{Z}}(\mathbf{x})$. This enforces the importance of selecting good inducing points $\mathbf{Z}$ with low $P_{\mathbf{Z}}(\mathbf{x})$.

V. PLANNING FRAMEWORK

A. Recursive Sparse GP Regression

To minimise the posterior entropy as per Problem 2, we perform recursive estimation to maintain a belief over $\mathbf{y}_{\mathbf{Z}}$, instead of direct computation using (12). The belief can be used not only for computing the posterior entropy, but also to fully recover the posterior GP (6) as long as the conditional independence assumption (7) holds. Note that a similar algorithm was already presented in [11], and was proven to be equivalent to sparse GP regression. Thus, we simply state the process without derivation.

The recursive estimation algorithm tracks the mean and covariance of $\mathbf{y}_{\mathbf{Z}}$ given measurements up to time t:

$$\boldsymbol{\mu}_t = \mathbb{E}[\mathbf{y}_{\mathbf{Z}} \mid \mathbf{y}_t], \quad \Sigma_t = \mathrm{Cov}[\mathbf{y}_{\mathbf{Z}} \mid \mathbf{y}_t]. \tag{15}$$

Given the belief $\boldsymbol{\mu}_t$ and Σ_t, we can recover the posterior GP (6) under the SoR and FIC approximations by a simple linear transform. With $\mathbf{q}(\mathbf{x}) = K_{\mathbf{Z}}^{-1} \mathbf{k}_{\mathbf{Z}}(\mathbf{x})$,

$$\begin{aligned} \mu(\mathbf{x} \mid \mathbf{y}_t) &= \mathbf{q}^T(\mathbf{x}) \boldsymbol{\mu}, \\ \sigma_{SoR}^2(\mathbf{x}, \mathbf{x}' \mid \mathbf{y}_t) &= \mathbf{q}^T(\mathbf{x}) \Sigma_t \mathbf{q}(\mathbf{x}'), \\ \sigma_{FIC}^2(\mathbf{x}, \mathbf{x}' \mid \mathbf{y}_t) &= \sigma_{SoR}^2(\mathbf{x}, \mathbf{x}' \mid \mathbf{y}_t), \\ &\quad + \delta(\mathbf{x}, \mathbf{x}')(k(\mathbf{x}, \mathbf{x}') - \hat{k}_{SoR}(\mathbf{x}, \mathbf{x}')), \end{aligned} \tag{16}$$

Because the of the linear relationship between the query points $\mu(\mathbf{x} \mid \mathbf{y}_t)$ and the inducing measurements in (16), the

belief is updated analogously to a Kalman filter for linear time-varying systems. Initially, the belief is set to $\boldsymbol{\mu}_0 = 0$ and $\Sigma_0 = K_{\mathbf{Z}}$. Given measurement y_{t+1} at $\mathbf{x}_{t+1}$, we generate the predictive mean, variance and cross-covariance of the measurement to perform the Kalman update:

$$\begin{aligned} \hat{y}_{t+1} &= \mu(\mathbf{x}_t \mid \mathbf{y}_t), \\ \Sigma^{yy}_{t+1} &= \sigma^2_*(\mathbf{x}_t, \mathbf{x}_t) + \sigma^2_\epsilon, \\ \Sigma^{y\mathbf{Z}}_{t+1} &= \mathbf{q}(\mathbf{x}_t)\Sigma_t. \end{aligned} \tag{17}$$

Using the predictions (17), we perform a Kalman update:

$$\begin{aligned} \boldsymbol{\mu}_{t+1} &= \boldsymbol{\mu}_t + \Sigma^{y\mathbf{Z}}_{t+1}(\Sigma^{yy}_{t+1})^{-1}(y_{t+1} - \hat{y}_{t+1}), \\ \Sigma_{t+1} &= \Sigma_t - \Sigma^{y\mathbf{Z}}_{t+1}(\Sigma^{yy}_{t+1})^{-1}(\Sigma^{y\mathbf{Z}}_{t+1})^T. \end{aligned} \tag{18}$$

B. Receding Horizon Planning

Given the belief maintained by recursive GP, we plan a path that minimises the cost function $c(\Sigma_t) = \log\det\Sigma_t$ using an adapted version of the RVI algorithm [23]. The RVI algorithm maintains a search tree $\mathcal{T}$ over possible trajectories and belief. Each node $v \in \mathcal{T}$ is associated with a candidate robot position and predicted posterior covariance $(\mathbf{x}^v_t, \Sigma^v_t)$.

An overview of a robot under operation is shown in Alg. 1. The tree is initialised with the robot's initial position and belief $(\mathbf{x}_0, \Sigma_0)$ (line 1). We perform an initial, offline search over the search horizon N by iterating the RVI algorithm (Alg. 2) N times (line 2). Each RVI iteration (Alg. 2) expands the search tree by one timestep and adds the corresponding layer of leaves.

During the online stage, the robot first extracts the control by searching over the set of leaves $\mathcal{L}(\mathcal{T})$ for a node $l^* \in \mathcal{L}(\mathcal{T})$ with the lowest cost $c(\Sigma^l)$. The corresponding control action $\mathbf{u}^*$ is executed, and the robot reaches a new state $\mathbf{x}_{t+1}$ and obtains a new measurement y_{t+1} (lines 7- 8). The belief $(\boldsymbol{\mu}_{t+1}, \Sigma_{t+1})$ is updated using the recursive GP equation (18) (line 9). To generate a new plan, we re-use the subtree $\mathcal{T}_{l^*}$ rooted at the chosen node l^* and perform the RVI iteration (line 10). Because the RVI iteration adds a new layer of leaves, the depth of the tree is always equal to time horizon N.

The RVI iteration proceeds as follows. First, a new layer of leaves are expanded from the current leaves $\mathcal{L}(\mathcal{T})$ in lines 1-5 by sampling the control space $\mathcal{U}$ and propagating the state $\mathbf{x}_{N-1}$ and posterior covariance $\boldsymbol{\Sigma}_{N-1}$ forward. Note that the propagation of posterior covariance $\boldsymbol{\Sigma}_{N-1}$ only depends on $\mathbf{x}_{N-1}$, and does not require the measurement y_{N-1}.

Next, we iterate over the newly added leaves and extract the set of nodes Q whose state are within δ-distance of each other (line 8). If such nodes exist, we check for ϵ-algebraic redundancy (ϵ-alg. red.) (line 9). A node l is ϵ-alg. red. iff there exists a set of coefficients $\{\alpha_q\}_{i=1}^{|Q|}$ such that $\sum \alpha_q = 1$ and $\Sigma_T + \epsilon I \succeq \sum_{i=1}^{|Q|} \alpha_i \Sigma_i$. A candidate node is pruned if it is within δ-distance of other leaf nodes, and is also ϵ-alg. red. with respect to those nodes (line 10).

The benefit of RVI is the strong suboptimality bound that accompanies it. The cost $c(\Sigma^{RVI}_N)$ returned by RVI and the optimal cost $c(\Sigma^*_N)$ satisfy $0 < c(\Sigma^*_N) - c(\Sigma^{RVI}_N) < C(\epsilon, \delta)$, where C is a problem-specific function. In particular, with $\epsilon, \delta = 0$ the result is optimal [23]. Note that the same bound holds for Problem 2, because it is of the same form as in [23].

As noted in [27], the ϵ-alg. red. check is an instance of LMI feasibility problem, and poses computational challenge as the number of inducing points grows. The challenge can be circumvented by setting $\epsilon = \infty$. In this case, the RVI iteration (Alg. 2) only adds the lowest cost nodes that are not within δ distance of each other, owing to the ascending order of iteration (line 7). While there are no bounds in this case, it produces practically viable solutions.

Algorithm 1 Receding horizon planning for worst-case error minimisation

1: $\mathcal{T} \leftarrow \{(\mathbf{x}_0, \Sigma_0)\}$
2: **for** $t = 1, ..., N$
3: $\quad\mathcal{T} \leftarrow \texttt{RVI}(\mathcal{T})$
4: **while** robot is operational
5: $\quad l^* \leftarrow \arg\min_{l \in \mathcal{L}(\mathcal{T})} c(\Sigma^l)$
6: $\quad\mathbf{u}^* \leftarrow \texttt{backtrace}(l^*)$
7: $\quad\mathbf{x}_{t+1} \leftarrow$ execute control $\mathbf{u}^*$
8: $\quad y_{t+1} \leftarrow$ sample measurement at $\mathbf{x}_{t+1}$
9: $\quad$Update $\boldsymbol{\mu}_{t+1}, \Sigma_{t+1}$ with y_{t+1} using (18)
10: $\quad\mathcal{T} \leftarrow \texttt{RVI}(\mathcal{T}_{l^*})$

Algorithm 2 Reduced value iteration

for $\forall l \in \mathcal{L}(\mathcal{T})$
$\quad$**for** $\forall \mathbf{u}_{N-1} \in \mathcal{U}$
$\quad\quad\mathbf{x}^u_N \leftarrow \mathbf{f}(\mathbf{x}^l_{N-1}, \mathbf{u}_{N-1})$
$\quad\quad\Sigma^u_N \leftarrow$ update Σ^l_{N-1} with (18)
$\quad\quad\mathcal{N} \leftarrow \mathcal{T} \cup \{(\mathbf{x}^u_N, \Sigma^u_N)\}$
$S_{\min} \leftarrow \{l \in \mathcal{L}(\mathcal{T}) \mid \Sigma^l_N = \arg\min c(\Sigma_N)\}$
for $l \in \mathcal{L}(\mathcal{T}) \setminus S_{\min}$ in ascending order of $c(\Sigma_N)$
$\quad Q \leftarrow \{\Sigma^u_N \mid u \in \mathcal{L}(\mathcal{T}), d(\mathbf{x}^l_N, \mathbf{x}^u_N) \le \delta\}$
$\quad$**if** Q is not empty and Σ_N is ϵ-alg. red. w.r.t. Q
$\quad\quad\mathcal{T} \leftarrow \mathcal{T} \setminus l$
Return $\mathcal{T}$

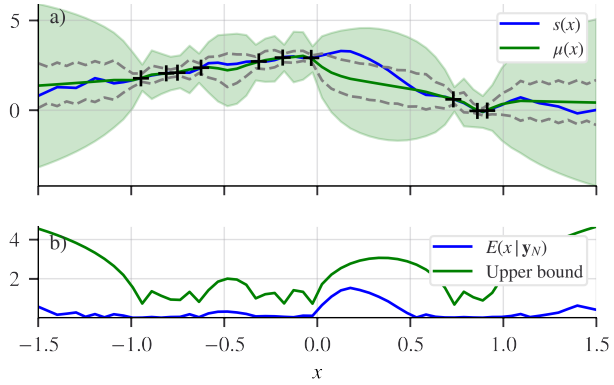

Fig. 1. a) One-dimensional sparse Gaussian process regression of target function in the RKHS. Measurement locations are indicated with black crosses. The proposed bound on deterministic error is shown in the green shaded area. The area enclosed by the grey dotted lines is the standard 1σ-confidence interval. b) Demonstration of the upper bound on deterministic error point-wise over the domain.

VI. EXPERIMENTAL RESULTS

A. Characterisation of the Error Bound

To corroborate the error bound proposed in Theorem 1 we study a sparse GP regression problem with a CI kernel. As the RKHS norm required for computing the error bound may be typically hard to compute, particularly in higher dimensional settings, we consider a one-dimensional regression problem where the target function is designed to take simplified form $s(x) = \sum_{i=1}^{m} \alpha_i k(x, x_i)$. This way we ensure $s \in \mathcal{H}_K$, and the RKHS norm is easily reduced to the Euclidean norm. Such simplifications are enough to illustrate the tightness and behaviour of the proposed bound.

Figure 1a) depicts the outcome of regression with such a target function. With sparse and noisy measurements, the GP regressor is able to reconstruct the target function well. Importantly, the proposed bounds on deterministic error are reasonably tight on the predicted mean and follow the intuitive behaviour of decreasing near measurement locations. When compared to the 1σ-confidence interval obtained from the posterior covariance, the proposed bounds have overall higher value. However, in certain regions of sparse or no measurement the target function is greater than the 1σ-confidence interval and yet remains within our error bound. This demonstrates that while confidence intervals may be broken, the error bound may not.

In Fig. 1b), we show the deterministic error (3) against the proposed upper bound, confirming the deterministic error lies below the bound for all x in the domain. This result corroborates that our bound does indeed give a reasonably tight upper bound on deterministic error.

B. Flow Field Case Study

To demonstrate capability of the proposed algorithm and problem formulation in information-theoretic path planning, we consider a simplified underwater glider operating in a double-gyre flow field. The dynamics are given by:

$$\begin{bmatrix} x_{t+1} \\ y_{t+1} \end{bmatrix} = \begin{bmatrix} x_t \\ y_t \end{bmatrix} + \Delta t (V_g \begin{bmatrix} -\sin(\pi x)\cos(\pi y) \\ \cos(\pi x)\sin(\pi y) \end{bmatrix} + V \begin{bmatrix} \cos u_t \\ sinu_t \end{bmatrix}). \tag{19}$$

The robot aims to solve Problem 2 equipped with the planning algorithm given in Alg. 1 in order to reconstruct a scalar field of interest (such as level of salinity) over the flow field. The scalar field is shown in Fig. 2 as a colour map superimposed on the flow field.

To evaluate the performance with respect to the deterministic error minimisation Problem 1, we evaluate the average absolute error $\frac{1}{N_S} \sum E(\mathbf{x}_i \mid \mathbf{y}_t)$ sampled over a 30×30 grid. We vary the horizon N between $\{1, 5, 10\}$ and examine the behaviour of average absolute error over time. Figure 4 shows the result over 20 randomised initial starting locations in the same environment (Fig. 2). It can be seen that for all choices of search horizon, the average absolute error decreases over time. The rate of reduction is greater with larger search horizon.

To better understand this behaviour, we show example trajectories after 100 time steps in Fig. 3. For the myopic greedy horizon case in Fig. 3a, 3b we observe poor coverage of the spatial domain. Additionally, the myopic nature of this planner results in trajectories that often get "stuck" in

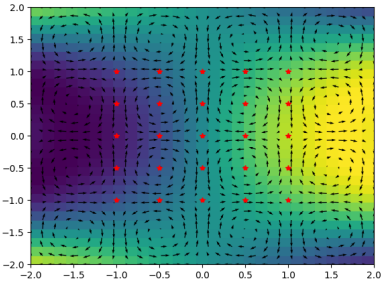

Fig. 2. The ground truth environment: brighter colour indicates higher value, inducing points used for sparse GP regression are shown in red.

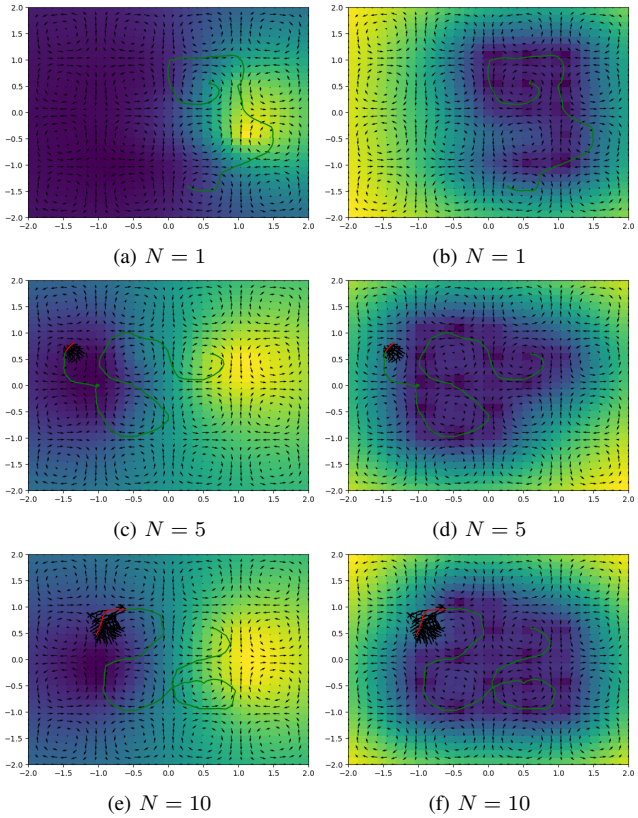

(a) $N = 1$ (b) $N = 1$ (c) $N = 5$ (d) $N = 5$ (e) $N = 10$ (f) $N = 10$

Fig. 3. Example trajectories after 100 steps with varying search horizon. Green: executed trajectory. Left column: mean ($\mu(\mathbf{x} \mid \mathbf{y}_N)$). Right column: variance, or square of power function ($\sigma^2(\mathbf{x} \mid \mathbf{y}_N) = P^2_{\mathbf{X}}(\mathbf{x})$). Red: current plan. Black: search tree. Brighter colour means higher value. Non-myopy leads to better coverage and reconstruction of the function of interest.

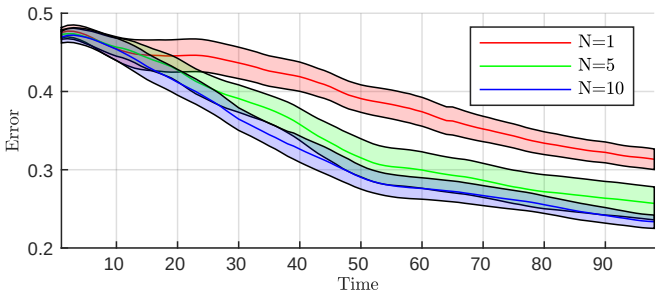

Fig. 4. Average absolute reconstruction error with varying search horizon. Shaded areas represent 95% confidence interval over 20 trials.

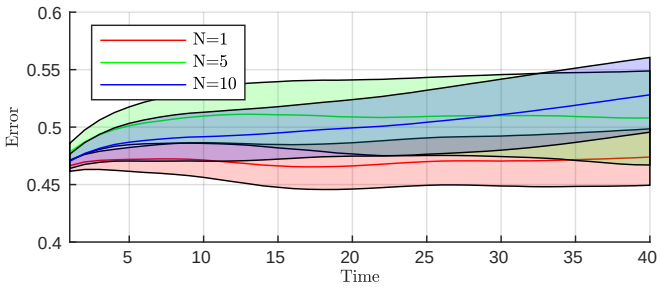

Fig. 5. Average absolute reconstruction error with varying search horizon for the measurement entropy maximisation formulation. Shaded areas represent 95% confidence interval over 20 trials.

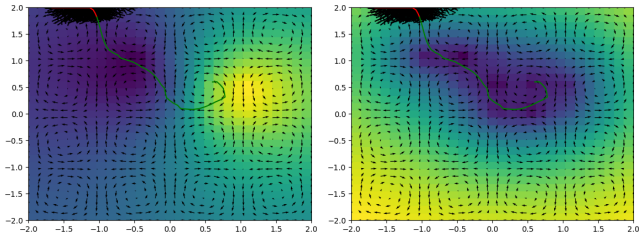

(a) (b)

Fig. 6. An example trajectory from measurement entropy maximisation. (a): mean ($\mu(\mathbf{x} \mid \mathbf{y}_N)$). (b): variance, or square of power function ($\sigma^2(\mathbf{x} \mid \mathbf{y}_N) = P^2_{\mathbf{X}}(\mathbf{x})$) Green: executed trajectory. Red: current plan. Black: search tree. Measurement entropy maximisation compels the robot to always explore outward. However, further away from the inducing points, the variance never decrease regardless of measurements.

an attracting region of a gyre. The robot is unable to use the flow field dynamics to its advantage to best explore, and reconstruction of the spatial field is thus poor.

With longer planning horizons the robot is able to successfully manoeuvre through the flow field to increase coverage of the domain, giving improved estimation of the spatial field as seen in Fig. 3c and 3e, with best estimation and coverage given by the longest horizon $N = 10$. Figure 3 further demonstrates the influence of the objective (2) on trajectories. For all horizon lengths, the robot preferentially takes measurements near inducing points over exploring regions further away, such that variance is minimised at inducing points. This is exemplified in Fig. 3 where with increasing horizon, broader coverage of the region around inducing points and greater minimisation of variance is achieved.

C. Comparison to Measurement Entropy Maximisation

Under the same experimental scenario (Fig. 2), we examine the behaviour of the solution of entropy maximisation, which is the usual approach to information gathering in GPs [15]. This was implemented by setting the cost in Alg. 2 as $-\log\det K_{\mathbf{X}}$ with $\epsilon = \infty$.

Fig. 5 shows the average absolute error over the environment with 20 random initial positions and varying search horizon. Surprisingly, the error does not decrease over time, and actually increases with larger search horizons ($N = 10$).

To understand the finding, we examine an example trajectory after 40 time-steps in Fig. 5, where the robot expands outwards making use of the ambient flow field. This is expected because the measurement entropy maximisation formulation demands the robot to simply move as far away from its previous trajectory as possible. However, the expansive behaviour is problematic when using inducing point-based GPs. As the robot gets further from the inducing points, the measurements do not make a significant contribution. Estimates further away from inducing points remain to be of poor quality regardless of measurements taken, as seen near the end of the green trajectory in Fig. 6b. Coupled with measurement entropy maximisation, the robot expands outwards in a positive feedback loop. A possible solution is to re-adjust the inducing points in an online manner via gradient updates [7]. We defer online update of inducing points to future work.

VII. CONCLUSION

We derived an upper bound to worst-case deterministic error for sparse GP regression with bounded noise. We proved this upper bound naturally gives way to an information-theoretic analogue to minimisation of deterministic error. Thus, one may minimise deterministic error via an information-theoretic proxy. Our results demonstrate the proposed approach outperforms conventional methods in reducing deterministic error. Results illuminated clear limitations involving placement of sparse GP inducing points which will be addressed in future work via online updates to inducing point locations. Further work also lies in extensions for reconstruction of time-varying spatial fields.

APPENDIX

Proof of Theorem 1. Letting $[\mathbf{s}(\mathbf{x}_N)]_i = s(\mathbf{x}_i,\, i = 1, \ldots, N$ be the vector of function evaluations at $\mathbf{x}_N$, and similarly $[\boldsymbol{\epsilon}_N]_i = \epsilon_i$, we have:

$$
\begin{aligned}
&E(\mathbf{x} \mid \mathbf{y}_N) = |s(\mathbf{x}) - \mu(\mathbf{x} \mid \mathbf{y}_N)| \\
&\leq |s(\mathbf{x}) - \mathbf{k}_{\mathbf{X}}^{\mathrm{T}}(\mathbf{x}) K_{\mathbf{X}}^{-1} \mathbf{s}(\mathbf{x}_N)| + \\
&\qquad |\mathbf{k}_{\mathbf{X}}^{\mathrm{T}}(\mathbf{x}) K_{\mathbf{X}}^{-1} \mathbf{s}(\mathbf{x}_N) - \mathbf{k}_{\mathbf{X}}^{\mathrm{T}}(\mathbf{x}) K_{\mathbf{X}}^{-1} [\mathbf{s}(\mathbf{x}_N) + \boldsymbol{\epsilon}_N]| \\
&= |\langle s, k(\cdot, \mathbf{x}) - \mathbf{k}_{\mathbf{X}}^{\mathrm{T}}(\mathbf{x}) K_{\mathbf{X}}^{-1} \mathbf{k}_{\mathbf{X}}(\cdot) \rangle_{\mathcal{H}_k}| + |\mathbf{k}_{\mathbf{X}}^{\mathrm{T}}(\mathbf{x}) K_{\mathbf{X}}^{-1} \boldsymbol{\epsilon}_N| \\
&\leq ||s||_{\mathcal{H}_k} ||k(\cdot, \mathbf{x}) - \mathbf{k}_{\mathbf{X}}^{\mathrm{T}}(\mathbf{x}) K_{\mathbf{X}}^{-1} \mathbf{k}_{\mathbf{X}}(\cdot)||_{\mathcal{H}_k} + |\mathbf{k}_{\mathbf{X}}^{\mathrm{T}}(\mathbf{x}) K_{\mathbf{X}}^{-1} \boldsymbol{\epsilon}_N| \\
&\leq ||s||_{\mathcal{H}_k} P_{\mathbf{X}}(\mathbf{x}) + \sqrt{\sigma_\epsilon^2 N} ||\mathbf{k}_{\mathbf{X}}^{\mathrm{T}}(\mathbf{x}) K_{\mathbf{X}}^{-1}||
\end{aligned}
$$

where the final inequality follows from $|\cdot| = \sqrt{\langle \cdot, \cdot \rangle}$, Cauchy-Schwarz inequality and the assumption $\epsilon_i^2 < \sigma_\epsilon^2$ $\forall i = 1, \ldots, N$. □

Proof of Theorem 2. We exploit the fact that the power function $P_N(\mathbf{x})$ can be linked to conditional entropy as:

$$H(s(\mathbf{x})|\mathbf{y}_N) = \frac{1}{2}(\log P_N(\mathbf{x}) + \log(2\pi e)).$$

Using the CI assumption, it can be shown that:

$$H(s(\mathbf{x})|\mathbf{y}_N) = H(s(\mathbf{x})|\mathbf{y}_{\mathbf{Z}}) + \mathcal{I}(\mathbf{y}_{\mathbf{Z}}; s(\mathbf{x})|\mathbf{y}_N). \tag{20}$$

It then follows that:

$$H(s(\mathbf{x})|\mathbf{y}_{\mathbf{Z}}) \leq H(s(\mathbf{x})|\mathbf{y}_N) \leq H(s(\mathbf{x})|\mathbf{y}_{\mathbf{Z}}) + H(\mathbf{y}_{\mathbf{Z}}|\mathbf{y}_N), \tag{21}$$

and the claimed bound is recovered by taking the exponential of (21). □

REFERENCES

[1] K. M. B. Lee, J. J. H. Lee, C. Yoo, B. Hollings, and R. Fitch, "Active perception for plume source localisation with underwater gliders," in *Proc. of ARAA ACRA*, 2018.

[2] M. Popovic, T. Vidal-Calleja, G. Hiltz, J. J. Chung, I. Sa, R. Siegwart, and J. Nieto, "An informative path planning framework for UAV-based terrain monitoring," *Auton. Rob.*, vol. 44, pp. 889 – 911, 2020.

[3] L. Wu, K. M. B. Lee, L. Liu, and T. Vidal-Calleja, "Faithful Euclidean distance field from log-Gaussian process implicit surfaces," *Rob. and Auotom. Lett.*, vol. 6, pp. 2461–2468, 2021.

[4] M. G. Jadidi, J. V. Miro, and G. Dissanayake, "Gaussian processes autonomous mapping and exploration for range-sensing mobile robots," *Auton. Rob.*, vol. 42, pp. 273–290, 2020.

[5] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.

[6] L. Sun, T. Vidal-Calleja, and J. V. Miro, "Bayesian fusion using conditionally independent submaps for high resolution 2.5D mapping," in *Proc. of IEEE ICRA*, 2015.

[7] K.-C. Ma, L. Liu, H. K. Heidarsson, and G. S. Sukhatme, "Data-driven learning and planning for environmental sampling," *J. Field Robot.*, vol. 35, pp. 643–661, 2018.

[8] J. Quinonero-Candela and C. E. Rasumussen, "A unifying view of sparse approximate Gaussian process regression," *J. Mach. Learn. Res.*, vol. 6, pp. 1939–1959, 2005.

[9] A. G. Wilson and H. Nickisch, "Kernel interpolation for scalable structured Gaussian processes (KISS-GP)," *32nd International Conference on Machine Learning, ICML 2015*, vol. 3, pp. 1775–1784, 2015.

[10] M. Titsias, "Variational learning of inducing variables in sparse Gaussian processes," in *Proc. of Artificial intelligence and statistics*. PMLR, 2009, pp. 567–574.

[11] M. Schürch, D. Azzimonti, A. Benavoli, and M. Zaffalon, "Recursive estimation for sparse Gaussian process regression," *Automatica*, vol. 120, p. 109127, 10 2020.

[12] K. Y. C. To, F. H. Kong, K. M. B. Lee, C. Yoo, S. Anstee, and R. Fitch, "Estimation of spatially correlated ocean currents from ensemble forecasts and online measurements," in *Proc. of IEEE ICRA*, 2021.

[13] H. A. Kingravi, H. Maske, and G. Chowdhary, "Kernel observers: Systems-theoretic modeling and inference of spatiotemporally evolving processes," in *Proc. of Adv. Neural Inf. Process. Syst.* Curran Associates, Inc., 2016, pp. 3990–3998.

[14] T. Vidal-Calleja, D. Su, F. De Bruijn, and J. V. Miro, "Learning spatial correlations for Bayesian fusion in pipe thickness mapping," in *Proc. of IEEE ICRA*. IEEE, 2014, pp. 683–690.

[15] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies," *J. Mach. Learn. Res*, vol. 9, pp. 235–284, 02 2008.

[16] A. Krause and D. Golovin, "Submodular function maximization." *Tractability*, vol. 3, pp. 71–104, 2014.

[17] G. Best and R. Fitch, "Probabilistic maximum set cover with path constraints for informative path planning," in *Proc. of ARAA ACRA*, 2016.

[18] G. A. Hollinger and G. S. Sukhatme, "Sampling-based robotic information gathering algorithms," *The Int. Journ. of Rob. Res.*, vol. 33, no. 9, pp. 1271–1287, 2014.

[19] T. Patten and R. Fitch, "Monte Carlo planning for active object classification," *Auton. Robot*, vol. 42, pp. 391 – 421, 2018.

[20] T. Karvonen and S. Särkkä, "Worst-case optimal approximation with increasingly flat Gaussian kernels," *Adv. Comput. Math.*, vol. 46, 03 2020.

[21] T. Karvonen, S. Särkkä, and K. Tanaka, "Kernel-based interpolation at approximate Fekete points," *Numer. Algorithms*, vol. 87, no. 1, pp. 445–468, 2021.

[22] R. Platt, R. Tedrake, L. P. Kaelbling, and T. Lozano-Perez, "Belief space planning assuming maximum likelihood," in *Proc. of RSS*, 2010.

[23] N. Atanasov, J. Le Ny, K. Daniilidis, and G. J. Pappas, "Information acquisition with sensing robots: Algorithms and error bounds," in *Proc. of IEEE ICRA*, 2014, pp. 6447–6454.

[24] J. Wakulicz, H. Kong, and S. Sukkarieh, "Active information acquisition under arbitrary unknown disturbances," in *Proc. of IEEE ICRA*, 2021.

[25] M. Kanagawa, P. Hennig, D. Sejdinovic, and B. K. Sriperumbudur, "Gaussian processes and kernel methods: A review on connections and equivalences," 2018.

[26] G. Wahba, *Spline models for observational data*. SIAM, 1990.

[27] N. Atanasov, "Active information acquisition with mobile robots," Ph.D. dissertation, University of Pennsylvenia, 2015.