

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, 2] and agriculture [3], and in general tasks such as robot navigation [4, 5]. 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 (RKHSs).

GP regression [6] 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 [4, 5], infrastructure [7], and agricultural [3] or oceanographic data [1, 8]. 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 [9–12]. One such approximation is the *inducing points* formulation [9], 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 [3, 7, 13–16].

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 [17, 18] 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 [19]. These classify as a submodular maximisation problem, which is NP-hard [20]. When the measurement locations are subject to a dynamics constraint, a non-myopic search is necessary [21], which can be achieved by, e.g., growing a search tree with an appropriate pruning condition [5, 22, 23].

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 [24, 25]. 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 [26–28]. 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 [27] 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

Consider a sensor-equipped, mobile robot that operates in environment $\mathcal{X} \subseteq \mathbb{R}^D$. The robot's motion is described by a discrete-time non-linear dynamic model

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

where $\mathbf{x}_t \in \mathcal{X}$ is the state of the robot in the environment and $\mathbf{u}_t \in \mathcal{U} \subseteq \mathbb{R}^{D_u}$ is the control input at time t . A sequence of N robot states and control actions are denoted as $\mathbf{X} = \{\mathbf{x}_0 \dots \mathbf{x}_{N-1}\}$ and $\mathbf{U} = \{\mathbf{u}_0 \dots \mathbf{u}_{N-1}\}$ respectively.

The onboard sensor takes online measurements $y_t \in \mathbb{R}$ of a scalar spatial phenomenon $s(\mathbf{x}_t)$ with additive noise ϵ_t according to the following measurement model:

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

Given measurement set corresponding to a sequence of robot states \mathbf{X} , denoted as $\mathbf{y}_{\mathbf{X}}$, the estimate of spatial phenomenon s over \mathcal{X} is denoted as $\hat{s}(\mathbf{x} \mid \mathbf{y}_{\mathbf{X}})$. The deterministic error between true spatial phenomenon s and estimate \hat{s} is the point-wise absolute difference, defined as

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

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}^* that minimises the total deterministic error (3) after time-step N over the domain \mathcal{X} :*

$$\min_{\mathbf{U} \in \mathcal{U}^N} \int_{\mathcal{X}} E(\mathbf{x} \mid \mathbf{y}_{\mathbf{X}}) d\mathbf{x} \quad (4)$$

Difficulty arises in solving Problem 1 in practice, as full knowledge of the ground truth $s(\mathbf{x})$ is naturally unavailable. 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. We use the tools from interpolation theory on RKHSs to find such bounds, 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 $\mathbf{y}_{\mathbf{X}}$, we generate an estimate $\hat{s}(\mathbf{x} \mid \mathbf{y}_{\mathbf{X}})$ of the spatial phenomenon $s(\mathbf{x})$ using GP regression with sparse approximation. A GP is a generalisation of multivariate Gaussian random variables (RVs) to *random functions*. A GP $s(\mathbf{x}) \sim GP(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$ is characterised by a mean function $m(\mathbf{x})$ and a *covariance* function $k(\mathbf{x}, \mathbf{x}')$, which specifies the covariance between function values at different points \mathbf{x} and \mathbf{x}' [6]:

$$\mathbb{E}[s(\mathbf{x})] = m(\mathbf{x}), \quad \text{Cov}[s(\mathbf{x}), s(\mathbf{x}')] = k(\mathbf{x}, \mathbf{x}'). \quad (5)$$

We impose a zero-mean GP prior on the scalar field of interest, $s(\mathbf{x}) \sim GP(0, k(\mathbf{x}, \mathbf{x}'))$, with $k(\mathbf{x}, \mathbf{x}')$ specified by the user. Let $\mathbf{y}_{\mathbf{X}}$ be a vector containing noisy measurements up to time-step N as per measurement model (2), i.e. $[\mathbf{y}_{\mathbf{X}}]_i = y_i$. With the zero-mean prior, the estimate $\hat{s}(\mathbf{x} \mid \mathbf{y}_{\mathbf{X}})$ given the measurements $\mathbf{y}_{\mathbf{X}}$ is given by another GP [29]:

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

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.

Inversion of the matrix $K_{\mathbf{X}}$ in (6) is computationally taxing at $\mathcal{O}(N^3)$. We use the *inducing point*-based approximation of the regression (6) introduced in [9] for its reduced computational complexity. Intuitively, this formulation introduces a small set of inducing points $\mathbf{Z} = \{\mathbf{z}_i\}_{i=1}^M$ ($M \ll N$) whose function values $[\mathbf{y}_{\mathbf{Z}}]_i = s(\mathbf{z}_i)$ ‘summarise’ the entire set of measurements $\mathbf{y}_{\mathbf{X}}$. One way to assert this is to impose that $\mathbf{y}_{\mathbf{X}}$ are *conditionally independent* (CI) given $\mathbf{y}_{\mathbf{Z}}$. Mathematically, the CI property holds if and only if:

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

which follows from asserting that the conditional cross-covariance vanishes given inducing measurements $\mathbf{y}_{\mathbf{Z}}$, i.e. $\sigma^2(\mathbf{x}, \mathbf{x}' \mid \mathbf{y}_{\mathbf{Z}}) = 0$. In other words, the correlation between any two measurements is indirect and is limited by their correlation to the inducing measurements, $\mathbf{y}_{\mathbf{Z}}$.

We consider two popular inducing point-based approximations that satisfy CI: the subset of regressors (SoR) and fully independent conditional (FIC) approximations. As noted in [9], the SoR and FIC approximations are equivalent to replacing the kernel $k(\mathbf{x}, \mathbf{x}')$ with approximate ones as follows:

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

where $\delta(\cdot)$ is the Kronecker delta function. Then, complexity reduces to $\mathcal{O}(M^2N)$ where $M \ll N$, a significant reduction compared to full GP regression (6).

B. Worst-case Error Bounds

We use the tools from interpolation theory on RKHSs [25] to derive worst-case bounds on the deterministic error (3) of sparse GP regression. The connection between GPs and RKHSs derive from an alternate interpretation of the covariance function in a GP (5) as a *positive-definite kernel*¹. Any positive-definite kernel k uniquely defines an RKHS \mathcal{H}_k . An RKHS \mathcal{H}_k is a space of real-valued functions equipped with an inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}_k}$ such that: 1) $k(\cdot, \mathbf{x}) \in \mathcal{H}_k \forall \mathbf{x} \in \mathcal{X}$, and 2) $\langle f, k(\cdot, \mathbf{x}) \rangle_{\mathcal{H}_k} = f(\mathbf{x}) \forall \mathbf{x} \in \mathcal{X}, \forall f \in \mathcal{H}_k$. That is, the kernel function $k(\cdot, \mathbf{x})$ is itself an element of \mathcal{H}_k , and ‘reproduces’ all other functions in \mathcal{H}_k at \mathbf{x} . The inner product induces a norm $\|f\|_{\mathcal{H}_k} = \sqrt{\langle f, f \rangle_{\mathcal{H}_k}}$. In the

¹A function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is positive definite if, for any choice of $\mathbf{X} \subset \mathcal{X}$, the matrix $[K_{\mathbf{X}}]_{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)$ is positive definite.

absence of noise, it can be shown that the predictive mean of GP regression (6) is exactly the minimum-norm interpolant, i.e., the function in the RKHS \mathcal{H}_k that agrees with all measurements $\mathbf{y}_\mathbf{X}$ and has the minimum norm [29, 30].

This RKHS view of GPs is beneficial because its deterministic nature allows bounding the deterministic error (3), as was done in [25] for interpolation on RKHSs. Inspired by such work, we establish the following worst-case error bound for GP regression with measurements containing bounded noise, which is representative of robotics applications:

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

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

where $P_\mathbf{X}(\mathbf{x}) = \sqrt{\sigma^2(\mathbf{x}, \mathbf{x} | \mathbf{y}_\mathbf{X})}$ 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 [25]:

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

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 holds true for 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., [19]. 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 [10], the interpolation of $s(\mathbf{x})$ given $\mathbf{y}_\mathbf{X}$ can be decomposed into two stages: 1) the interpolation of inducing measurements $\mathbf{y}_\mathbf{Z}$ given $\mathbf{y}_\mathbf{X}$, 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:*

$$P_\mathbf{Z}(\mathbf{x}) \leq P_\mathbf{X}(\mathbf{x}) \leq P_\mathbf{Z}(\mathbf{x}) \exp H(\mathbf{y}_\mathbf{Z} | \mathbf{y}_\mathbf{X}), \quad (11)$$

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

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

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

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

Using Theorem 2, we approximately solve Problem 1 via the following surrogate problem, which minimises the control-dependent terms in the new upper bound (13):

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

$$\min_{\mathbf{U} \in \mathcal{U}^N} H(\mathbf{y}_\mathbf{Z} | \mathbf{y}_\mathbf{X}). \quad (14)$$

The merit of the reformulation in Problem 2 is that the worst-case total deterministic error (9) can be approximately minimised without explicitly integrating over the operating region in the noise-free case. This is because the posterior entropy $H(\mathbf{y}_\mathbf{Z} | \mathbf{y}_\mathbf{X})$ is the only term dependent on the control actions, but is independent of the integrand \mathbf{x} . Further, Problem 2 serves as a closer proxy to Problem 1 as $H(\mathbf{y}_\mathbf{Z} | \mathbf{y}_\mathbf{X}) \rightarrow 0$ (i.e., as more measurements are added, or with better solution quality). This is owing to the ‘tightness’ of the inequality (11) in that $P_\mathbf{X}(\mathbf{x}) \rightarrow P_\mathbf{Z}(\mathbf{x})$ as $H(\mathbf{y}_\mathbf{Z} | \mathbf{y}_\mathbf{X}) \rightarrow 0$. The fact that $P_\mathbf{X}(\mathbf{x})$ approaches $P_\mathbf{Z}(\mathbf{x})$ illustrates the importance of selecting good inducing points \mathbf{Z} with low $P_\mathbf{Z}(\mathbf{x})$.

V. PLANNING FRAMEWORK

A further benefit of the reformulation in Problem 2 is that we need only to compute and minimise the posterior entropy of a fixed number of variables (i.e., the inducing measurements $\mathbf{y}_\mathbf{Z}$), unlike previous formulations [8, 19] where the dimensionality grows. We exploit this benefit by using an efficient recursive sparse GP algorithm presented in [13] that maintains a Gaussian belief over the inducing measurements $\mathbf{y}_\mathbf{Z}$. In turn, the Gaussian representation allows the use of the reduced value iteration (RVI) algorithm [27] for minimising the posterior entropy while retaining its strong guarantees.

A. Recursive Sparse GP Regression

Given measurements up to time t , the recursive sparse GP regression algorithm [13] permits equivalent calculation to GP regression (6), while only storing the posterior mean and covariance of the inducing measurements $\mathbf{y}_\mathbf{Z}$:

$$\boldsymbol{\mu}_t = \mathbb{E}[\mathbf{y}_\mathbf{Z} | \mathbf{y}_t], \quad \Sigma_t = \text{Cov}[\mathbf{y}_\mathbf{Z} | \mathbf{y}_t]. \quad (15)$$

Importantly, the posterior entropy can be calculated as a function of Σ_t :

$$H(\mathbf{y}_\mathbf{Z} | \mathbf{y}_\mathbf{X}) \equiv c(\Sigma_t) = \frac{1}{2} \log \det 2\pi e \Sigma_t. \quad (16)$$

This is more efficient than direct computation using (12), and has additional benefit of fitting directly into the sparse recursive GP regression algorithm as follows.

Letting $\mathbf{q}(\mathbf{x}) = K_\mathbf{Z}^{-1} \mathbf{k}_\mathbf{Z}(\mathbf{x})$, we can recover the posterior GP (6) (i.e., perform regression) given the belief $\boldsymbol{\mu}_t$ and Σ_t :

$$\begin{aligned} \boldsymbol{\mu}(\mathbf{x} | \mathbf{y}_t) &= \mathbf{q}^T(\mathbf{x}) \boldsymbol{\mu}_t, \\ \sigma_{SoR}^2(\mathbf{x}, \mathbf{x}' | \mathbf{y}_t) &= \mathbf{q}^T(\mathbf{x}) \Sigma_t \mathbf{q}(\mathbf{x}'), \\ \sigma_{FIC}^2(\mathbf{x}, \mathbf{x}' | \mathbf{y}_t) &= \sigma_{SoR}^2(\mathbf{x} \mathbf{x}' | \mathbf{y}_t), \\ &+ \delta(\mathbf{x}, \mathbf{x}') (k(\mathbf{x}, \mathbf{x}') - \hat{k}_{SoR}(\mathbf{x}, \mathbf{x}')). \end{aligned} \quad (17)$$

The update procedure is analogous to a Kalman filter. Initially, the belief is set to $\boldsymbol{\mu}_0 = 0$ and $\Sigma_0 = K_Z$. Given measurement y_t at \mathbf{x}_t , we generate the predictive mean, variance and cross-covariance at future time-step:

$$\begin{aligned} \hat{y}_{t+1} &= \mu(\mathbf{x}_t | \mathbf{y}_t), \\ \Sigma_{t+1}^{yy} &= \sigma_*^2(\mathbf{x}_t, \mathbf{x}_t) + \sigma_e^2, \\ \Sigma_{t+1}^{yZ} &= \mathbf{q}(\mathbf{x}_t) \Sigma_t. \end{aligned} \quad (18)$$

Using the predictions (18), and new measurement y_{t+1} at \mathbf{x}_{t+1} we perform a Kalman update:

$$\begin{aligned} \boldsymbol{\mu}_{t+1} &= \boldsymbol{\mu}_t + \Sigma_{t+1}^{yZ} (\Sigma_{t+1}^{yy})^{-1} (y_{t+1} - \hat{y}_{t+1}), \\ \Sigma_{t+1} &= \Sigma_t - \Sigma_{t+1}^{yZ} (\Sigma_{t+1}^{yy})^{-1} (\Sigma_{t+1}^{yZ})^T. \end{aligned} \quad (19)$$

As with the Kalman filter, here the covariance prediction and update (18)-(19) are measurement independent, enabling posterior entropy (16) minimisation over non-myopic horizons.

B. Receding Horizon Planning

Given the belief maintained by recursive GP, we plan a path that minimises the posterior entropy (16) using an adapted version of the RVI algorithm [27]. 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 ($\mathbf{x}_t^v, \Sigma_t^v$).

Alg. 1 shows an operational overview. The tree is initialised with the robot's initial position and belief (\mathbf{x}_0, Σ_0) (line 1). Initially, we perform an offline search over horizon N by repeatedly iterating the RVI algorithm (Alg. 2) in 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 extracts the optimal control from the tree by searching for the lowest-cost leaf node (lines 5-6). 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} (line 7). The belief ($\boldsymbol{\mu}_{t+1}, \Sigma_{t+1}$) is updated using the recursive GP equation (19) (line 8). 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 9). 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 Σ_{N-1} forward. Note that the propagation of posterior covariance Σ_{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 $\{\alpha_q\}_{q \in Q}^{|Q|}$ 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 $c(\Sigma_N^{RVI})$ returned by RVI and the

Algorithm 1 Receding horizon planning for worst-case error minimisation

- 1: $\mathcal{T} \leftarrow \{(\mathbf{x}_0, \Sigma_0)\}$
 - 2: **for** $t = 1, \dots, N$
 - 3: $\mathcal{T} \leftarrow \text{RVI}(\mathcal{T})$
 - 4: **while** robot is operational
 - 5: $l^* \leftarrow \arg \min_{l \in \mathcal{L}(\mathcal{T})} c(\Sigma^l)$
 - 6: $\mathbf{u}^* \leftarrow \text{backtrace}(l^*)$
 - 7: $\mathbf{x}_{t+1}, y_{t+1} \leftarrow \text{execute } \mathbf{u}^*$ and sample measurement
 - 8: Update $\boldsymbol{\mu}_{t+1}, \Sigma_{t+1}$ with y_{t+1} using (19)
 - 9: $\mathcal{T} \leftarrow \text{RVI}(\mathcal{T}_{l^*})$
-

Algorithm 2 Reduced value iteration

- 1: **for** $\forall l \in \mathcal{L}(\mathcal{T})$
 - 2: **for** $\forall \mathbf{u}_{N-1} \in \mathcal{U}$
 - 3: $\mathbf{x}_N^u \leftarrow \mathbf{f}(\mathbf{x}_{N-1}^l, \mathbf{u}_{N-1})$
 - 4: $\Sigma_N^u \leftarrow \text{update } \Sigma_{N-1}^l \text{ with (19)}$
 - 5: $\mathcal{N} \leftarrow \mathcal{T} \cup \{(\mathbf{x}_N^u, \Sigma_N^u)\}$
 - 6: $S_{\min} \leftarrow \{l \in \mathcal{L}(\mathcal{T}) \mid \Sigma_N^l = \arg \min c(\Sigma_N)\}$
 - 7: **for** $l \in \mathcal{L}(\mathcal{T}) \setminus S_{\min}$ in ascending order of $c(\Sigma_N)$
 - 8: $Q \leftarrow \{\Sigma_N^u \mid u \in \mathcal{L}(\mathcal{T}), d(\mathbf{x}_N^u, \mathbf{x}_N^l) \leq \delta\}$
 - 9: **if** Q is not empty and Σ_N is ϵ -alg. red. w.r.t. Q
 - 10: $\mathcal{T} \leftarrow \mathcal{T} \setminus l$
 - 11: **Return** \mathcal{T}
-

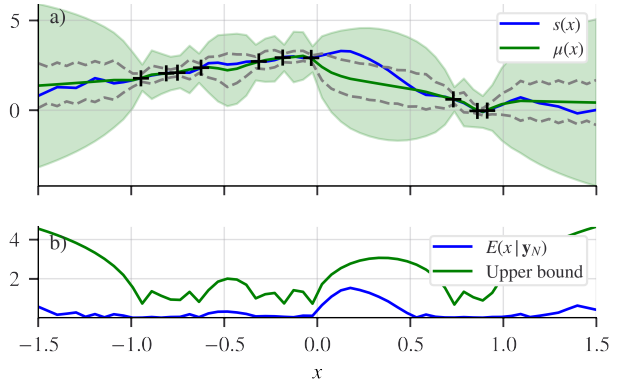


Fig. 1. a) One-dimensional sparse Gaussian process regression of target function in the RKHS. Measurement locations are indicated with black crosses. The proposed bound on deterministic error is shown in the green shaded area. The area enclosed by the grey dotted lines is the standard 1σ -confidence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

optimal cost $c(\Sigma_N^*)$ satisfy $0 < c(\Sigma_N^*) - c(\Sigma_N^{RVI}) < C(\epsilon, \delta)$, where C is a problem-specific function. In particular, with $\epsilon, \delta = 0$ the result is optimal [27]. Note that the same bound holds for Problem 2, because it is of the same form as in [27].

As noted in [31], 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.

VI. EXPERIMENTAL RESULTS

A. Characterisation of the Error Bound

We first corroborate the error bound proposed in Theorem 1 with an example. As the RKHS norm required for computing the error bound may be hard to compute, particularly in higher dimensional settings, we consider a one-dimensional regression problem where the target function is designed with 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 demonstrate the tightness and behaviour of the proposed bound.

Fig. 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:

$$\begin{bmatrix} x_{t+1} \\ y_{t+1} \end{bmatrix} = \begin{bmatrix} x_t \\ y_t \end{bmatrix} + \Delta t (V_g \begin{bmatrix} -\sin(\pi x) \cos(\pi y) \\ \cos(\pi x) \sin(\pi y) \end{bmatrix}) + V \begin{bmatrix} \cos u_t \\ \sin u_t \end{bmatrix}. \quad (20)$$

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 verify that the algorithm solves the deterministic error minimisation problem (Problem 1), we evaluate the average absolute error sampled over a 30×30 grid. We vary the horizon N between $\{1, 5, 10\}$ and examine the average absolute error over time. Fig. 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




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




Fig. 3. Example trajectories after 100 steps with varying search horizon. Green: executed trajectory. Left column: mean ($\mu(\mathbf{x} | \mathbf{y}_{\mathbf{x}})$). Right column: variance, or square of power function ($\sigma^2(\mathbf{x} | \mathbf{y}_{\mathbf{x}}) = P_{\mathbf{x}}^2(\mathbf{x})$). Red: current plan. Black: search tree. Brighter colour means higher value. Non-myopic 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.

of the spatial domain. Additionally, the myopic nature of this planner results in trajectories that often get “stuck” in an attracting region of a gyre. The robot is unable to use




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.




Fig. 6. An example trajectory from measurement entropy maximisation. (a): mean ($\mu(\mathbf{x} \mid \mathbf{y}_{\mathbf{X}})$). (b): variance, or square of power function ($\sigma^2(\mathbf{x} \mid \mathbf{y}_{\mathbf{X}}) = P_{\mathbf{X}}^2(\mathbf{x})$). 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.

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 successfully manoeuvres 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 $N = 10$. Fig. 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 [19]. 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 [8]. 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 and multi-dimensional spatial fields [14, 32–35].

APPENDIX

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

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

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

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

$$H(s(\mathbf{x}) \mid \mathbf{y}_{\mathbf{X}}) = \frac{1}{2}(\log P_{\mathbf{X}}(\mathbf{x}) + \log(2\pi e)).$$

Using the CI assumption, it can be shown that:

$$H(s(\mathbf{x}) \mid \mathbf{y}_{\mathbf{X}}) = H(s(\mathbf{x}) \mid \mathbf{y}_{\mathbf{Z}}) + \mathcal{I}(\mathbf{y}_{\mathbf{Z}}; s(\mathbf{x}) \mid \mathbf{y}_{\mathbf{X}}). \quad (21)$$

It then follows that:

$$H(s(\mathbf{x}) \mid \mathbf{y}_{\mathbf{Z}}) \leq H(s(\mathbf{x}) \mid \mathbf{y}_{\mathbf{X}}) \leq H(s(\mathbf{x}) \mid \mathbf{y}_{\mathbf{Z}}) + H(\mathbf{y}_{\mathbf{Z}} \mid \mathbf{y}_{\mathbf{X}}), \quad (22)$$

and the claimed bound is recovered by taking the exponential of (22). \square

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] G. D'urso, J. J. H. Lee, O. Pizarro, C. Yoo, and R. Fitch, "Hierarchical mcts for scalable multi-vessel multi-float systems," in *Proc. of IEEE ICRA*. IEEE, 2021, pp. 8664–8670.
- [3] 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.
- [4] L. Wu, K. M. B. Lee, L. Liu, and T. Vidal-Calleja, "Faithful Euclidean distance field from log-Gaussian process implicit surfaces," *Rob. and Autom. Lett.*, vol. 6, pp. 2461–2468, 2021.
- [5] 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.
- [6] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.
- [7] 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.
- [8] K.-C. Ma, L. Liu, H. K. Heidarrsson, and G. S. Sukhatme, "Data-driven learning and planning for environmental sampling," *J. Field Robot.*, vol. 35, pp. 643–661, 2018.
- [9] J. Quinonero-Candela and C. E. Rasmussen, "A unifying view of sparse approximate Gaussian process regression," *J. Mach. Learn. Res.*, vol. 6, pp. 1939–1959, 2005.
- [10] 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.
- [11] M. Titsias, "Variational learning of inducing variables in sparse Gaussian processes," in *Proc. of Artificial intelligence and statistics*. PMLR, 2009, pp. 567–574.
- [12] K. M. B. Lee, W. Martens, J. Khatkar, R. Fitch, and R. Mettu, "Efficient updates for data association with mixtures of Gaussian processes," IEEE, 2020, pp. 335–341.
- [13] M. Schürch, D. Azzimonti, A. Benavoli, and M. Zaffalon, "Recursive estimation for sparse Gaussian process regression," *Automatica*, vol. 120, p. 109127, 10 2020.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] K. C. To, J. J. H. Lee, C. Yoo, S. Anstee, and R. Fitch, "Streamline-based control of underwater gliders in 3d environments," in *Proc. of IEEE CDC*. IEEE, 2019, pp. 8303–8310.
- [18] K. C. To, C. Yoo, S. Anstee, and R. Fitch, "Distance and steering heuristics for streamline-based flow field planning," in *Proc. of IEEE ICRA*. IEEE, 2020, pp. 1867–1873.
- [19] 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.
- [20] A. Krause and D. Golovin, "Submodular function maximization." *Tractability*, vol. 3, pp. 71–104, 2014.
- [21] G. Best and R. Fitch, "Probabilistic maximum set cover with path constraints for informative path planning," in *Proc. of ARAA ACRA*, 2016.
- [22] 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.
- [23] T. Patten and R. Fitch, "Monte Carlo planning for active object classification," *Auton. Robot.*, vol. 42, pp. 391 – 421, 2018.
- [24] T. Karvonen and S. Särkkä, "Worst-case optimal approximation with increasingly flat Gaussian kernels," *Adv. Comput. Math.*, vol. 46, 03 2020.
- [25] T. Karvonen, S. Särkkä, and K. Tanaka, "Kernel-based interpolation at approximate Fekete points," *Numer. Algorithms*, vol. 87, no. 1, pp. 445–468, 2021.
- [26] R. Platt, R. Tedrake, L. P. Kaelbling, and T. Lozano-Perez, "Belief space planning assuming maximum likelihood," in *Proc. of RSS*, 2010.
- [27] 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.
- [28] J. Wakulicz, H. Kong, and S. Sukkarieh, "Active information acquisition under arbitrary unknown disturbances," in *Proc. of IEEE ICRA*, 2021.
- [29] M. Kanagawa, P. Hennig, D. Sejdinovic, and B. K. Sriperumbudur, "Gaussian processes and kernel methods: A review on connections and equivalences," 2018.
- [30] G. Wahba, *Spline models for observational data*. SIAM, 1990.
- [31] N. Atanasov, "Active information acquisition with mobile robots," Ph.D. dissertation, University of Pennsylvania, 2015.
- [32] K. M. B. Lee, C. Yoo, B. Hollings, S. Anstee, S. Huang, and R. Fitch, "Online estimation of ocean current from sparse GPS data for underwater vehicles," in *Proc. of IEEE ICRA*. IEEE, 2019, pp. 3443–3449.
- [33] C. Yoo, S. Anstee, and R. Fitch, "Stochastic path planning for autonomous underwater gliders with safety constraints," in *Proc. of IEEE/RSJ IROS*. IEEE, 2019, pp. 3725–3732.
- [34] C. Yoo, J. J. H. Lee, S. Anstee, and R. Fitch, "Path planning in uncertain ocean currents using ensemble forecasts," in *Proc. of IEEE ICRA*. IEEE, 2021, pp. 8323–8329.
- [35] J. J. H. Lee, C. Yoo, S. Anstee, and R. Fitch, "Hierarchical planning in time-dependent flow fields for marine robots," in *Proc. of IEEE ICRA*. IEEE, 2020, pp. 885–891.