

“© 2017 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

Warped Gaussian Processes Occupancy Mapping with Uncertain Inputs

Maani Ghaffari Jadidi, Jaime Valls Miro, and Gamini Dissanayake

Abstract—In this paper, we study extensions to the Gaussian Processes (GPs) continuous occupancy mapping problem. There are two classes of occupancy mapping problems that we particularly investigate. The first problem is related to mapping under pose uncertainty and how to propagate pose estimation uncertainty into the map inference. We develop expected kernel and expected sub-map notions to deal with uncertain inputs. In the second problem, we account for the complication of the robot's perception noise using Warped Gaussian Processes (WGP). This approach allows for non-Gaussian noise in the observation space and captures the possible nonlinearity in that space better than standard GPs. The developed techniques can be applied separately or concurrently to a standard GP occupancy mapping problem. According to our experimental results, although taking into account pose uncertainty leads, as expected, to more uncertain maps, by modeling the nonlinearities present in the observation space WGP can improve the map quality.

Index Terms—Mapping, SLAM, Probability and Statistical Methods, Range Sensing.

I. INTRODUCTION

IN many scenarios such as robotic navigation, the robot pose is partially observable; and we have access only to an estimate (noise-corrupted version) of the robot pose, as depicted in Figure 1. Under these circumstances, the robot requires to navigate in an uncertain environment [1]–[4], and the probability distribution of the robot pose will be the input for the mapping problem. In practice, and based on the application, most of the occupancy mapping techniques ignore robot pose uncertainty for map representation; either for efficiency or as the resultant map is not suitable for navigation. Furthermore, dense representation of the state often makes uncertainty propagation intractable. This problem is not unique to *Gaussian Processes Occupancy Maps* (GPOMs), but it is also present in *Occupancy Grid Maps*. With this motivation, we study the problem of GP occupancy mapping under pose uncertainty. The first solution is uncertainty propagation through *kernel* functions. The second solution we propose uses the *expected sub-map* notion to incorporate pose uncertainties into the map building process.

The second problem studied is motivated by the fact that due to the smoothness common in the resultant regressed




Fig. 1: The pose-graph estimation of the robot trajectory from the Intel research lab. dataset [8], solved using the techniques in [3], [9]. The large uncertainties from pose estimation are often ignored in the dense occupancy map representation. For the sake of clarity, loop-closures are omitted, and only one-sixth of the pose covariances are illustrated.

maps, inferring a high-quality map compatible with the actual shape of the environment is non-trivial [5]. Furthermore, for a complicated task such as robotic mapping [6], the additive white Gaussian noise assumption in standard GPs can be simplistic. To account for these problems, we improve the incremental GPOM technique using *Warped Gaussian Processes* (WGP) [7]. The core idea is to map the target values through a warping (transforming) function to capture the nonlinear behavior of the observations.

We tackle the mentioned problems to improve map quality and provide results using incremental WGP Occupancy Maps (WGPOMs) under pose uncertainty.

Notation

Probabilities and probability densities are not distinguished in general. Matrices are capitalized in bold, such as in \mathbf{X} , and vectors are in lower case bold type, such as in \mathbf{x} . Vectors are column-wise and $1:n$ means integers from 1 to n . The Euclidean norm is shown by $\|\cdot\|$. $\text{tr}(\mathbf{X})$ and $|\mathbf{X}|$ denote the trace and the determinant of matrix \mathbf{X} , respectively. Random variables, such as X , and their realizations, x , are sometimes denoted interchangeably. $x^{[i]}$ denotes a reference to the i -th element of the variable. An alphabet such as \mathcal{X} denotes a set. A reference to a test set quantity is shown by \mathbf{x}_* . The n -by- n identity matrix is denoted by \mathbf{I}_n . The sign function, and the

Manuscript received: August, 31, 2016; Revised November, 22, 2016; Accepted December, 19, 2016.

This paper was recommended for publication by Editor Cyrill Stachniss upon evaluation of the Associate Editor and Reviewers' comments.

The authors are with the Centre for Autonomous System, Faculty of Engineering and IT, University of Technology Sydney, Ultimo, NSW 2007, Australia {maani.ghaffarijadidi, jaime.vallsmiro, gamini.dissanayake}@uts.edu.au

Digital Object Identifier (DOI): see top of this page.

absolute value are denoted by $\text{sgn}(x)$, and $|x|$, respectively. Finally, $\mathbb{E}[\cdot]$ and $\mathbb{V}[\cdot]$ denote the expected value and variance of a random variable, respectively.

Outline

A review of related works is given in the following section. We first discuss the problem of Gaussian processes occupancy mapping under pose uncertainty in Section III; followed by presenting the warped Gaussian processes occupancy mapping in Section IV. Robotic mapping experiments are presented in Section V, and finally, Section VI concludes the paper.

II. RELATED WORK

Current robotic navigation algorithms rely on a dense environment representation for safe navigation. Occupancy grid maps are the standard way of environment representation in robotics and are the natural choice for online implementations [10]–[14]. However, they simplify the mapping problem to a set of independent cells and estimate the probability of occupancy for each cell by ignoring the correlation between them. The map posterior is then approximated as the product of its marginals.

The FastSLAM algorithm [15] exploits *Rao-Blackwellized particle filters* [16], and relies on the fact that, given the robot trajectory, landmarks are conditionally independent. In practice, FastSLAM can generate accurate maps; however, it suffers from degeneracy and depletion problems [17]. Full SLAM methods estimate the *smoothing distribution* of the robot trajectory and the map, given all available data. In robotics, full SLAM using probabilistic graphical models representation and the *maximum a posteriori* estimate derived by solving a nonlinear least squares problem [18]–[22] can be considered the state-of-the-art and we thus focus our attention on using a pose-graph estimation of the robot trajectory as in [3], [9].

Kernel methods in the form of Gaussian processes (GPs) framework [23] are non-parametric regression and classification techniques that have been used to model spatial phenomena [24]. The development of GPOM is discussed in [25], [26]. The incremental GP map building using the Bayesian Committee Machine (BCM) technique [27] is developed in [28]–[31] and for online applications in [32]. In [33], the Hilbert maps technique is proposed and is more scalable. However, it is an approximation for continuous occupancy mapping and produces maps with less accuracy than GPOM.

Approximate methods for uncertainty propagation into GPs models through kernel functions are proposed in [34], and developed for GPOM in [35]. Generally speaking, training and query points can both be noisy. In [34], the problem of prediction at an uncertain input is discussed while it is assumed the input in training data is noise free. In [35], using a similar approach, the idea is extended to account for noisy training input. Similarly, we assume the input in training data is uncertain and query points are deterministic. In this paper, we employ this technique and develop the expected sub-map technique for approximate uncertainty propagation, then incorporate both techniques into the WGPOM framework.

III. MAPPING UNDER POSE UNCERTAINTY

The main challenge in building occupancy maps under robot pose uncertainty is the dense representation of map belief which makes uncertainty propagation computationally expensive. Maximum likelihood dense map representations are currently the common practice which does not necessarily produce correct maps, especially if pose estimation uncertainties are significant. This popularity can be understood from the fact that employing an environment representation constructed with significant uncertainties results in vague obstacles and free space and is not suitable for robotic motion planning and navigation. However, accounting for pose uncertainties in mapping is not only important for correct map representations, but also for motion planning (prediction) tasks.

A. Problem Statement and Formulation

Let \mathcal{M} be the set of possible static occupancy maps. We consider the map of the environment as an n_m -tuple random variable $(M^{[1]}, \dots, M^{[n_m]})$ whose elements are described by a normal distribution $M^{[i]} \sim \mathcal{N}(\mu^{[i]}, \sigma^{[i]})$, $\forall i \in \{1: n_m\}$. Let $\mathcal{X} \subset \mathbb{R}^2$ be the set of spatial coordinates to build a map on. Let $y = \tilde{y} + \epsilon_y$ be a noisy measurement (class label; -1 and 1 for unoccupied and occupied, respectively) at a noisy sample $\mathbf{x} = \tilde{\mathbf{x}} + \epsilon_x$, where $\epsilon_y \sim \mathcal{N}(0, \sigma_n^2)$ and $\epsilon_x \sim \mathcal{N}(\mathbf{0}, \Sigma_x)$. Define a training set $\mathcal{D} = \{(\mathbf{x}^{[i]}, y^{[i]}) \mid i = 1: n_t\}$ which consists of noisy measurements at noisy locations. Let function $f: \mathcal{X} \rightarrow \mathcal{M}$, i.e. $y = f(\tilde{\mathbf{x}} + \epsilon_x) + \epsilon_y$, be the real underlying process that we model as a Gaussian process $f(\mathbf{x}) \sim \mathcal{GP}(0, k(\mathbf{x}, \mathbf{x}'))$, where $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is the *covariance function* or *kernel*; and \mathbf{x} and \mathbf{x}' are either in the training or the test (query) sets. Estimate $p(M = m \mid \mathcal{D})$, i.e. the map posterior probability given a noisy training set. For a given query point in the map, \mathbf{x}_* , GP predicts a mean, μ_* , and an associated variance, σ_* . We can write $m^{[i]} = y(\mathbf{x}_*^{[i]}) \sim \mathcal{N}(\mu^{[i]}, \sigma^{[i]})$. To show a valid probabilistic representation of the map $p(m^{[i]} \mid \mathcal{D})$, the classification step squashes data into the range $[0, 1]$.

We propose two methods to solve the defined problem. The first approach is based on the *expected kernel*. The alternative approach, *expected sub-map*, treats all inputs deterministically and propagates pose uncertainties through uncertain map fusion. The following assumptions are made in the present work:

Assumption 1 (Deterministic query points). *In the problem of Gaussian processes occupancy mapping under robot pose uncertainty, query points are deterministic.*

Assumption 2. *The covariance function in the expected sub-map method is stationary, $k(\mathbf{x}, \mathbf{x}') = k(\|\mathbf{x} - \mathbf{x}'\|)$.*

Remark 1. *Using Assumption 2, map inference in the local coordinates of the robot (local map) can be done using deterministic inputs.*

B. System Dynamics

The equation of motion of the robot is governed by the nonlinear partially observable equation as follows.

$$\mathbf{x}_{t+1}^- = f(\mathbf{x}_t, \mathbf{u}_t, \mathbf{w}_t) \quad \mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_t) \quad (1)$$

moreover, with appropriate linearization at the current state estimate, we can predict the state covariance matrix as

$$\Sigma_{t+1}^- = \mathbf{F}_t \Sigma_t \mathbf{F}_t^T + \mathbf{W}_t \mathbf{Q}_t \mathbf{W}_t^T \quad (2)$$

where $\mathbf{F}_t = \frac{\partial f}{\partial \mathbf{x}}|_{\mathbf{x}_t, \mathbf{u}_t}$ and $\mathbf{W}_t = \frac{\partial f}{\partial \mathbf{w}}|_{\mathbf{x}_t, \mathbf{u}_t}$ are the Jacobian matrices calculated with respect to \mathbf{x} and \mathbf{w} , respectively.

C. Expected Kernel

The core idea in the expected kernel approach is taking an expectation of the covariance function over uncertain inputs. Let \mathbf{x} be distributed according to a probability distribution $p(\mathbf{x})$. The expected covariance function can be computed as

$$\tilde{k} = \mathbb{E}[k] = \int_{\Omega} k d\mathbf{p} \quad (3)$$

In general, this integral is analytically intractable; therefore we employ two numerical approximations to solve (3). However, for the case of the squared exponential (SE) kernel, a closed-form solution exists [36]. Hence, once the expected covariance matrix is calculated, we can compute the predictive conditional distribution for a single query point similar to standard GPs.

Monte Carlo Integration: Since we assume the distribution of the uncertain input is known, by drawing independent samples, $\mathbf{x}^{[i]}$, from $p(\mathbf{x})$ and using a Monte-Carlo technique, we can approximate Equation (3) by

$$\tilde{k} = \frac{1}{n} \sum_{i=1}^n k_i \quad (4)$$

where k_i is the covariance function computed at $\mathbf{x}^{[i]}$.

Remark 2. In Equation (4), the covariance $k(\mathbf{x}, \mathbf{x})$ and cross-covariance $k(\mathbf{x}, \mathbf{x}_*)$ are both denoted as k_i . Depending on the input, the integration is only performed on training points as it is assumed query points are deterministic.

Gauss-Hermite Quadrature: Gauss-Hermite quadrature [37] of integrals of the kind $\int_{-\infty}^{\infty} \exp(-x^2) f(x) dx$ are given by $\int_{-\infty}^{\infty} \exp(-x^2) f(x) dx \approx \sum_{j=1}^n w^{[j]} f(x^{[j]})$. The multi-variate normal distribution of noisy input is given by $\mathcal{N}(\tilde{\mathbf{x}}, \Sigma_x)$. Through a change of variable such that $\mathbf{L}\mathbf{L}^T = 2\Sigma_x$ and $\mathbf{u} = \mathbf{L}^{-1}(\mathbf{x} - \tilde{\mathbf{x}})$, where \mathbf{L} is a lower triangular matrix that can be calculated using a Cholesky factorization, the Equation (3) can be approximated as

$$\tilde{k} = (2\pi)^{-\frac{d}{2}} \sum_{i_1=1}^n \dots \sum_{i_d=1}^n \bar{w} k_{i_{1:d}} \quad (5)$$

where $\bar{w} \triangleq \prod_{j=1}^d w^{[i_j]}$, $u^{[i_j]}$ are the roots of the Hermite polynomial H_n , $\mathbf{u}_{i_{1:d}} \triangleq [u^{[i_1]}, \dots, u^{[i_d]}]^T$, and $k_{i_{1:d}}$ is the covariance function computed at $\mathbf{x}_{i_{1:d}} = \mathbf{L}\mathbf{u}_{i_{1:d}} + \tilde{\mathbf{x}}$. When $d = 2$, we can simplify Equation (5) and write $\tilde{k} = \frac{1}{2\pi} \sum_{i_1=1}^n \sum_{i_2=1}^n \bar{w} k_{i_{1:2}}$.

Remark 3. We assumed points from the map spatial support, $\mathbf{x} \in \mathcal{X}$, are global coordinates. In practice, to transform local points, an unscented transform [38] is used to reduce linearization errors [35].

An illustrative example of GP regression where inputs are uncertain is shown in Figure 2. By propagating the input




Fig. 2: The plot shows an example of GP regression with uncertain inputs. The GP-EK shows GP regression by incorporating the input uncertainty using the expected kernel. The standard GP results are generated by ignoring the input uncertainty which cannot provide a consistent solution. The groundtruth function is $f(x) = \frac{1}{5} \cos(x^2)e^{-x} + \frac{3}{20} \sin(x)$ and training points are corrupted by $\epsilon_x \sim \mathcal{N}(0, 0.6^2)$.

uncertainty using the expected kernel, the output does not follow the observations exactly yet remains consistent as the underlying function is within the estimated uncertainty bounds.

D. Expected Sub-map

We exploit the fact that a stationary covariance function does not depend on the selected coordinates, i.e. the global or a local sub-map frame. Therefore, we treat all training inputs as noise free and conduct map inference using deterministic inputs in the local (sensor/robot) frame. To fuse the inferred sub-map into the global map, we draw independent samples from $p(\mathbf{x}_t)$. In other words, by taking the expectation over the location of the sub-map, we propagate uncertainty of each map point to its neighborhood. Thus, we have $p(m|\tilde{\mathbf{x}}, y) = \int p(m|\mathbf{x}, y)p(\mathbf{x}_t)d\mathbf{x}_t$, and by drawing independent samples from $p(\mathbf{x})$ and using a Monte-Carlo approximation it follows that $p(m|\tilde{\mathbf{x}}, y) \approx \frac{1}{n} \sum_{j=1}^n p(m_j|\mathbf{x}_j, y)$. Note that any sub-map $p(m_j|\mathbf{x}_j, y)$ can be fused into the global map using the algorithms in [5]. However, as a result of sampling, the expected map, $p(m|\tilde{\mathbf{x}}, y)$, is similar to a mixture distribution; therefore, the mean and variance calculations need to be addressed accordingly. We present the following proposition to calculate the first two moments of $p(m|\tilde{\mathbf{x}}, y)$.

Lemma 1. Let X_1, X_2, \dots, X_n be random variables that are distributed according to probability densities $p(x_1), p(x_2), \dots, p(x_n)$, with constant weights w_1, w_2, \dots, w_n , where $\sum_{i=1}^n w_i = 1$. The probability density function of the mixture is $p(x) = \sum_{i=1}^n w_i p(x_i)$. Given $\mu_i = \mathbb{E}[X_i]$ and $\sigma_i^2 = \mathbb{V}[X_i]$, the mean and variance of the mixture density is given by $\mu = \mathbb{E}[X] = \sum_{i=1}^n w_i \mathbb{E}[X_i]$ and $\sigma^2 = \mathbb{V}[X] = \sum_{i=1}^n w_i (\sigma_i^2 + \mu_i^2) - (\sum_{i=1}^n w_i \mu_i)^2$.

Proof. The proof follows from the fact that for the k -th moment of the mixture, we can write $\mathbb{E}[X^{(k)}] = \sum_{i=1}^n w_i \mathbb{E}[X_i^{(k)}]$ and define the variance accordingly. \square

Proposition 2 (Expected sub-map fusion). In incremental map building, to compute $p(m|\tilde{\mathbf{x}}, y)$, sampled sub-maps can be fused into the global map using the following equations

$$\mathbb{E}[M] = \frac{1}{n} \sum_{j=1}^n \mathbb{E}[M_j] \quad (6)$$

$$\mathbb{V}[M] = \frac{1}{n} \left(\sum_{j=1}^n (\mathbb{V}[M_j] + \mathbb{E}[M_j]^2) - \frac{1}{n} \left(\sum_{j=1}^n \mathbb{E}[M_j] \right)^2 \right) \quad (7)$$

where $\mathbb{E}[M_j]$ is the updated global map built using the j -th independently drawn robot pose sample, and (6) and (7) can be computed point-wise for every map point $m^{[i]}$.

Proof. The proof directly follows from Lemma 1. \square

Therefore, we can perform incremental map fusion by considering the robot pose uncertainty without modifying the GP framework.

IV. WARPED GP OCCUPANCY MAPPING

The primary challenge in modeling the environment “accurately” is the different nature of free and occupied classes. Free space tends to span vast areas while occupied space often represents the structural shape of the environment. In addition, the assumption of additive Gaussian noise in the observations in standard GPs is unable to capture complexity in observations appropriately. We propose to employ Warped Gaussian Processes to account for the nonlinear behavior of observations. This method is appealing as it allows for non-Gaussian noise in the observation space. However, exact inference is not possible anymore, and approximate inference algorithms such as *expectation propagation* [39] or *variational Bayes* [40] are required.

The idea to accommodate non-Gaussian distributions and noise is to use a nonlinear monotonic function for warping (transforming) the observation space [7]. Let $g_w(\cdot)$ be a transformation from the observation space to a latent space as

$$t^{[i]} = g_w(y^{[i]}; \psi) \quad i = 1 : n \quad (8)$$

where ψ denotes the vector of warping function hyperparameters and $\mathbf{t} = [t^{[1]}, \dots, t^{[n]}]^T$ is the vector of latent targets. Now we can re-write the GP formulation for the latent target and by accounting for the transformation between a true observation and the latent target, the negative log of the marginal likelihood (NLML) can be written as

$$\begin{aligned} \log p(\mathbf{y}|\mathbf{X}, \boldsymbol{\theta}, \psi) = & -\frac{1}{2} g_w(\mathbf{y})^T [\mathbf{K}(\mathbf{X}, \mathbf{X}) + \sigma_n^2 \mathbf{I}_n]^{-1} g_w(\mathbf{y}) \\ & - \frac{1}{2} \log |\mathbf{K}(\mathbf{X}, \mathbf{X}) + \sigma_n^2 \mathbf{I}_n| - \frac{n}{2} \log 2\pi + \sum_{i=1}^n \log \left| \frac{\partial g_w(y)}{\partial y} \right|_{y^{[i]}} \end{aligned} \quad (9)$$

in which the last term is the Jacobian of the defined transformation. To compute the mean at a new test point, it is possible to calculate the expectation of the inverse warping function over the latent target predictive density, therefore

$$\mathbb{E}[y^{[n+1]}] = \int g_w^{-1}(t) \mathcal{N}(\hat{t}^{[n+1]}, \sigma^{[n+1]}) dt = \mathbb{E}[g_w^{-1}] \quad (10)$$

This integral can be computed numerically using Gauss-Hermite quadrature with a weighted sum of the inverse warping function g_w^{-1} .

Inspired by the neural network transfer functions, a sum of hyperbolic tangent functions satisfies the requirements for the transformation to be monotonic and at the same time




Fig. 3: A challenging example of regression using standard and Warped GPs. The measurements are sampled from the true output values by adding noise, i.e. $\epsilon_y \sim \mathcal{N}(0, 0.05^2)$. The warping functions are $\tanh(\ell = 2)$ and polynomial of degree five. The \tanh demonstrate a better extrapolating behavior, while the polynomial function can follow the underlying function more closely. However, polynomials are prone to over-fitting as it can be seen that the estimated uncertainties (blue shaded region) are significantly narrower.

allowing for complicated mappings. With hyperparameters vector $\psi = [\mathbf{a}, \mathbf{b}, \mathbf{c}]^T$, the function can be defined as

$$g_w(y; \psi) = y + \sum_{i=1}^{\ell} a^{[i]} \tanh(b^{[i]}(y + c^{[i]})) \quad a^{[i]}, b^{[i]} \geq 0 \quad \forall i \in \{1 : \ell\} \quad (11)$$

where the parameter ℓ is the number of steps and has to be set depending on the complexity of observations, $\mathbf{a} = [a^{[1]}, \dots, a^{[\ell]}]$, $\mathbf{b} = [b^{[1]}, \dots, b^{[\ell]}]$, and $\mathbf{c} = [c^{[1]}, \dots, c^{[\ell]}]$. Alternative warping functions can be polynomials ($\psi = \mathbf{c}$):

$$g_w(y; \psi) = y + \sum_{i=2}^{\ell} c^{[i-1]} \text{sgn}(y) |y|^{i-1} \quad c^{[i]} \geq 0 \quad \forall i \in \{2 : \ell\} \quad (12)$$

Figure 3 shows a simple yet challenging example for regression using standard and Warped GPs. The measurements are corrupted by an additive Gaussian noise. Even though the noise is still Gaussian, the complicated structure of the underlying function makes modeling it non-trivial. Note that the inputs are deterministic.

Remark 4. By increasing the number of training points, it is possible to generate more accurate results using standard GPs. However, given the cubic time complexity of GPs, dense training datasets reduce the scalability of the algorithms significantly.

V. RESULTS AND DISCUSSION

We now present results from experiments using a synthetic dataset and a real publicly available pose-graph dataset. The synthetic dataset, Figure 4, is built in such a way as to highlight the strength of WGPOM to model complicated structural shapes and to better appreciate the mapping performance of the incremental GPOM and WGPOM under the expected kernel (EK) and expected sub-map (ESM) uncertainty propagation techniques. On the other hand, the Intel dataset [8], as shown in Figure 1, exposes an extreme real-world example of the problem at hand where highly uncertain robot poses along the estimated trajectory are present.

We compare the overall taken time to build the entire map (using all the available data) as well as the map accuracy using the Area Under the receiving operating characteristic Curve




Fig. 4: The synthetic dataset used for comparison of GPOM and WGPOM under various uncertainty propagation conditions. The figure shows collected observations along the robot trajectory (orange dots) and the robot position uncertainty ellipse at each corresponding pose (Q_3 scenario). The robot starting position is shown in green. Map dimensions are in meters.

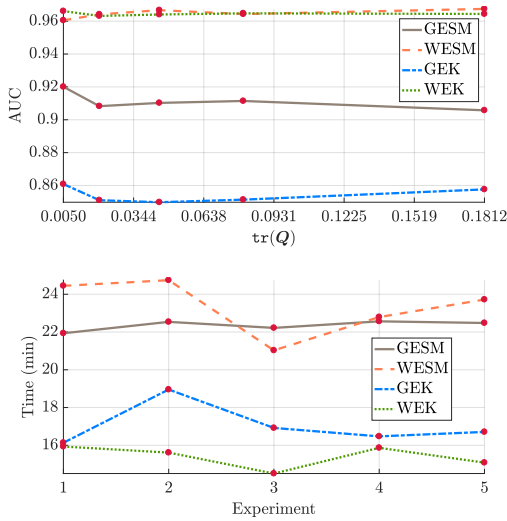


Fig. 5: The AUC and runtime for incremental GPOM and WGPOM using the synthetic dataset and proposed uncertainty propagation methods. The runtimes are in minutes. All maps are computed using 0.5 m resolution and a k d-tree data structure for the storage and nearest neighbor queries. As it is expected, generally, by propagating pose uncertainties into the map inference, the map quality degrades (GESM and GEK). However, applying WGs can alleviate this effect and produce maps with improved accuracies (WESM and WEK).

(AUC) [41]. For each model, we learn the hyperparameters at the first increment of map building by minimization of the NLML using the first set of training data with manual supervision to ensure the best possible outcome for all models. The data processing and computations for the incremental map building are implemented using MATLAB.

A. First Experiment: Motion Uncertainty Effect

The map of the environment, the robot trajectory, the observations collected at each pose using a simulated rangefinder sensor, together with the evolution of the robot pose uncertainty due to its motion noise, are illustrated in Figure 4. The uncertainty ellipsoids show the worst-case covariance of the robot position, and there is no uncertainty reduction along the path by closing loops.

Details from the model selection, compared techniques, and conducted experiments are collected in Table I. We use Matérn

TABLE I: Details of the experiments for the motion uncertainty effect in a synthetic dataset using the expected kernel and the expected sub-map uncertainty propagation schemes. The experiment is repeated five times by increasing the robot motion noise covariance, i.e. Q_1, \dots, Q_5 . Note that all mapping techniques are incremental.

Model selection:		
Technique	Cov. func.	Warping func.
GPOM	Matérn ($\nu = 5/2$)	n/a
WGPOM	SE (ARD)	tanh ($\ell = 2$)
Compared mapping techniques:		
Technique	Uncertainty propagation	Abbreviation
GPOM	EK	GEK
GPOM	ESM	GESM
WGPOM	EK	WEK
WGPOM	ESM	WESM
Numerical integration:		
Uncertainty propagation	Technique	Samples
EK	Gauss-Hermite	9
ESM	Monte-Carlo	10
Training and test (query) points size:		
Parameter	Symbol	Value
Tot. training points	n_t	87833
Ave. training points	\bar{n}_t	829
Test points	n_q	6561
Robot motion model noise covariances:		
$Q_1 = \text{diag}(0.05 \text{ m}, 0.05 \text{ m}, 0.25 \text{ rad})^2$, $Q_2 = \text{diag}(0.1 \text{ m}, 0.1 \text{ m}, 0.5 \text{ rad})^2$		
$Q_3 = \text{diag}(0.15 \text{ m}, 0.15 \text{ m}, 0.75 \text{ rad})^2$, $Q_4 = \text{diag}(0.2 \text{ m}, 0.2 \text{ m}, 1.00 \text{ rad})^2$		
$Q_5 = \text{diag}(0.3 \text{ m}, 0.3 \text{ m}, 2.00 \text{ rad})^2$		

($\nu = 5/2$) covariance function for GPOM as its performance has been shown fitting in earlier works [5]. For WGPOM, we use SE covariance function with Automatic Relevance Determination (ARD) [42]. Since the observations do not cover the entire map, we use tanh with $\ell = 2$ as the warping function to improve the extrapolation ability of GPs (Figure 3). The experiment for each mapping technique using EK and ESM uncertainty propagation is repeated by increasing the robot motion noise covariance in five steps. In Figure 5, the map accuracy and runtime comparisons of all methods using AUC are shown. In the bottom plot, the runtime for the ESM is higher than the EK for both mapping techniques. From the top plot, we can see that applying WGs improves the map quality regardless of the uncertainty propagation choice. The expected kernel is computed using Gauss-Hermite quadrature with 9 sample points, and the expected sub-map computations are performed using Monte-Carlo approximations with 10 samples. When increasing the number of samples from 9 to 18 we did not observe any improvement in the map accuracy.

Figure 6 illustrates the results of all combinations of the proposed techniques. GPOM and WGPOM by ignoring the robot pose uncertainty are shown in Figures 6a and 6d, respectively. WGPOM demonstrates a better discrimination performance between class labels in the absence of measurements. This effect can be seen in the middle of the smaller star in both maps. Note that further optimization of hyperparameters can lead to over-fitting instead of solving the discussed problem.

Even though ESM and EK try to achieve the same goal, they demonstrate different behaviors. Generally speaking, integration over the covariance function (EK) has a smoothing effect that sometimes can be desirable. For example, in all the




Fig. 6: Illustrative examples of the occupancy map from the first experiment (for Q_3). The top row corresponds to the GPOM, and the bottom row shows WGPOM results. The maps in (a) and (d) show GPOM and WGPOM results by ignoring the robot pose uncertainties. In (b) and (e) the robot pose uncertainty is incorporated using the expected sub-map method. In (c) and (f) the robot pose uncertainty is incorporated using the expected kernel method. The WGPOM-based maps can deal with input uncertainty better and provide maps with higher quality, shown in (e) and (f). All maps are computed using 0.5 m resolution.

presented maps the central part of the map due to the lack of observation is partially complete. As a result of this smoothing effect of the EK, this part is correctly classified to be closer to the free space class. However, the probabilities are closer to 0.5, and the robot cannot be completely confident about the status of the area. Alternatively, ESM leads to relatively more confident maps with smaller smoothing effects. The resultant maps are safer for navigation as the gap between occupied and unoccupied areas is classified as an unknown region. While this behavior can be an appealing property from a motion planning point of view, the occupied areas are faded; and we cannot see the structural shape of the environment accurately. Overall, by propagating more uncertainty into the map inference process, we expect less accuracy in the outcome and more realistic estimation of the belief. However, WGPOMs by modeling nonlinearity in the observation space improve the map quality.

B. Experimental Results

The experimental results of occupancy mapping using the Intel dataset are shown in Table II and Figure 7. In the absence of a complete groundtruth map, the groundtruth map for this dataset is generated using the estimated robot trajectory and rangefinder measurements. In this way, the groundtruth map has the same orientation which makes the comparison convenient. The covariance function used is an intrinsically sparse kernel [43]. The logic behind this choice is that the structural shape of the environment is complex, and it is cluttered with random people and typical office furniture; therefore, using a covariance function that correlates map points over a long range is not suitable. We use the Sparse

TABLE II: The AUC and runtime for incremental GPOM and WGPOM using the Intel dataset and proposed uncertainty propagation methods. The runtimes are in minutes. Maps are computed using 0.2 m resolution and a kd -tree data structure for the storage and nearest neighbor queries. The Sparse covariance function is used for all techniques and Polynomials of degree seven as the warping function. Total and average number of training points, and the number of test points are $n_t = 137979$, $\bar{n}_t = 186$, and $n_q = 40401$, respectively.

Method	No uncertainty		EK		ESM	
	AUC	Time	AUC	Time	AUC	Time
GPOM	0.8499	72	0.7323	532	0.7026	657
WGPOM	0.8463	125	0.7887	577	0.7436	766

covariance function for both GPOM and WGPOM and their corresponding uncertainty propagation experiments.

Table II shows the runtime and accuracy comparison of the mapping techniques. In this scenario, where the pose uncertainties are ignored, GPOM performs marginally better than WGPOM. This result, shown in Figures 7a and 7d, can be understood from the fact that WGPOM has covered more partially observed areas in the map. While this is desirable, in the absence of a complete groundtruth map, it leads to a lower AUC. However, WGPOM maintains more accurate maps by incorporating the pose uncertainties which is the actual problem to be solved.

In this experiment, the uncertainty propagation using the ESM method leads to poor map qualities and the maps, shown in Figures 7b and 7e, are almost entirely faded due to the significant pose uncertainties resulting from the robot repeatedly traveling through the same areas. This fading effect could be partially mitigated by increasing the number of samples, but based on the runtimes in Table II, it is not justifiable. Nevertheless, WGPOM using the EK demonstrates a better performance in comparison to all methods that incorporated




Fig. 7: Occupancy mapping results using the Intel dataset. The top row corresponds to the GPOM, and the bottom row shows WGPOM results. The maps built by ignoring pose uncertainties are shown in (a) and (d) using GPOM and WGPOM, respectively. In (b) and (e) the robot pose uncertainty is incorporated using the expected sub-map method. In (c) and (f) the robot pose uncertainty is incorporated using the expected kernel method.

pose uncertainties; moreover, it produces a map that is also usable for navigation tasks, i.e. obstacle avoidance, and path planning.

C. Discussion and Limitations

Uncertainty propagation through the developed methods can provide safer maps for robotic navigation. However, long-term uncertainty propagation leads to highly faded maps. If the uncertainty at each step is large and the robot cannot improve the localization confidence through loop-closures, this fading effect is more severe. Also, the uncertainty of the robot orientation has a large impact on the map quality. As we have seen in the presented results, in the expected sub-map approach, if the orientation uncertainty is significant, integration using a small number of samples leads to poor map accuracies. The expected kernel technique has an advantage from this perspective, as the unscented transform maps the measurement into the map space, and samples are related to the map spatial dimensions.

It is also demonstrated that the concept of accounting for possible nonlinearities in the observation space, here through the warping function, has desirable effects on the map quality. Our results showed that regardless of the uncertainty propagation technique, applying WGP provide better maps than standard GPs. We reiterate that by ignoring the robot pose uncertainty, the map is a potentially incorrect representation of the environment.

Finding an appropriate warping function that is compatible with non-Gaussianity in the observation space can be time-

consuming unless the model selection is performed in a more systematic way. In this work, we tried several functions and chose the best one. Moreover, since the exact inference is not possible, approximate methods may not always converge. Although the upper-bound time complexity of WGPOM is similar to that of GPOM, in practice for larger datasets the inference takes longer. Improving the computational efficiency of the proposed methods is an interesting direction to follow.

D. Computational Complexity

For both GPOM and WGPOM, the worst-case time complexity is cubic in the number of training data, $\mathcal{O}(\bar{n}_t^3)$. For ESM, the number of sub-map fusion into the global map scales linearly with the number of samples, n_s , and the sub-map fusion involves a nearest neighbor query for each test point resulting in $\mathcal{O}(\bar{n}_t^3 + n_s n_q \log n_q)$. In the case of two-dimensional mapping, using Gauss-Hermite quadrature, the time complexity of EK computation is quadratic in the number of sample points and together with sub-map fusion leads to $\mathcal{O}(\bar{n}_t^3 + n_s^2 + n_q \log n_q)$. In EK, applying the unscented transform to all training points involves a Cholesky factorization of the input covariance matrix which is ignored for the two-dimensional case.

VI. CONCLUSION

In this paper, we studied incremental GP occupancy mapping extensions through warped GPs. Since occupancy maps have dense belief representations, the robot pose uncertainty

is often ignored. We proposed two methods to incorporate robot pose uncertainty into the map inference, the expected kernel and the expected sub-map. While the expected kernel handles the input uncertainty within the GPs framework, the expected sub-map exploits the inherent property of stationary covariance functions for map inference in the local frame with deterministic inputs. The proposed methods can also be useful if the belief representation is not dense (as opposed to occupancy mapping). Furthermore, the WGPM technique can deal with the nonlinear behavior of measurements through a nonlinear transformation which improves the ability of GPs to learn complex structural shapes more accurately, especially, under uncertain inputs.

Future work includes further examinations of the proposed methods in practical robotic exploration and obstacle avoidance scenarios, for example, using currently popular visual-inertial odometry systems.

ACKNOWLEDGMENT

The authors would like to thank three anonymous reviewers for their valuable feedback and Kasra Khosoussi for the helpful discussions.

REFERENCES

- [1] N. Roy, W. Burgard, D. Fox, and S. Thrun, "Coastal navigation-mobile robot navigation with uncertainty in dynamic environments," in *Proc. IEEE Int. Conf. Robot Automat.*, vol. 1. IEEE, 1999, pp. 35–40.
- [2] S. Prentice and N. Roy, "The belief roadmap: Efficient planning in belief space by factoring the covariance," *The Int. J. Robot. Res.*, vol. 28, no. 11-12, pp. 1448–1465, 2009.
- [3] R. Valencia, M. Morta, J. Andrade-Cetto, and J. M. Porta, "Planning reliable paths with Pose SLAM," *IEEE Trans. Robot.*, vol. 29, no. 4, pp. 1050–1059, 2013.
- [4] J. Vallvé and J. Andrade-Cetto, "Potential information fields for mobile robot exploration," *Robot. Auton. Syst.*, vol. 69, pp. 68–79, 2015.
- [5] M. Ghaffari Jadidi, J. Valls Miro, and G. Dissanayake, "Gaussian process autonomous mapping and exploration for range sensing mobile robots," *arXiv preprint arXiv:1605.00335*, 2016.
- [6] S. Thrun, "Exploring artificial intelligence in the new millennium," G. Lakemeyer and B. Nebel, Eds. Morgan Kaufmann Publishers Inc., 2003, ch. Robotic Mapping: A Survey, pp. 1–35.
- [7] E. Snelson, C. E. Rasmussen, and Z. Ghahramani, "Warped Gaussian processes," *NIPS*, vol. 16, pp. 337–344, 2004.
- [8] A. Howard and N. Roy, "The robotics data set repository (Radish)," 2003. [Online]. Available: <http://radish.sourceforge.net>
- [9] V. Ila, J. Porta, and J. Andrade-Cetto, "Information-based compact Pose SLAM," *IEEE Trans. Robot.*, vol. 26, no. 1, pp. 78–93, 2010.
- [10] H. P. Moravec and A. Elfes, "High resolution maps from wide angle sonar," in *Proc. IEEE Int. Conf. Robot Automat.*, vol. 2. IEEE, 1985, pp. 116–121.
- [11] A. Elfes, "Sonar-based real-world mapping and navigation," *Robotics and Automation, IEEE Journal of*, vol. 3, no. 3, pp. 249–265, 1987.
- [12] S. Thrun, "Learning occupancy grid maps with forward sensor models," *Autonomous robots*, vol. 15, no. 2, pp. 111–127, 2003.
- [13] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "OctoMap: An efficient probabilistic 3D mapping framework based on octrees," *Auton. Robot.*, vol. 34, no. 3, pp. 189–206, 2013.
- [14] R. S. Merali and T. D. Barfoot, "Optimizing online occupancy grid mapping to capture the residual uncertainty," in *Proc. IEEE Int. Conf. Robot Automat.* IEEE, 2014, pp. 6070–6076.
- [15] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit, "FastSLAM 2.0: An improved particle filtering algorithm for simultaneous localization and mapping that provably converges," in *In Proc. of the Int. Conf. on Artificial Intell. (IJCAI)*, 2003, pp. 1151–1156.
- [16] A. Doucet, N. De Freitas, K. Murphy, and S. Russell, " Rao-Blackwellised particle filtering for dynamic Bayesian networks," in *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc., 2000, pp. 176–183.
- [17] T. Bailey, J. Nieto, and E. Nebot, "Consistency of the FastSLAM algorithm," in *Proc. IEEE Int. Conf. Robot Automat.* IEEE, 2006, pp. 424–429.
- [18] F. Dellaert and M. Kaess, "Square root SAM: Simultaneous localization and mapping via square root information smoothing," *The Int. J. Robot. Res.*, vol. 25, no. 12, pp. 1181–1203, 2006.
- [19] G. Grisetti, D. L. Rizzini, C. Stachniss, E. Olson, and W. Burgard, "On-line constraint network optimization for efficient maximum likelihood map learning," in *Proc. IEEE Int. Conf. Robot Automat.* IEEE, 2008, pp. 1880–1885.
- [20] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert, "iSAM2: Incremental smoothing and mapping with fluid relinearization and incremental variable reordering," in *Proc. IEEE Int. Conf. Robot Automat.* IEEE, 2011, pp. 3281–3288.
- [21] K. Konolige, G. Grisetti, R. Kümmerle, W. Burgard, B. Limketkai, and R. Vincent, "Efficient sparse pose adjustment for 2D mapping," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.* IEEE, 2010, pp. 22–29.
- [22] S. Thrun and M. Montemerlo, "The graph SLAM algorithm with applications to large-scale mapping of urban structures," *The Int. J. Robot. Res.*, vol. 25, no. 5-6, pp. 403–429, 2006.
- [23] C. Rasmussen and C. Williams, *Gaussian processes for machine learning*. MIT press, 2006, vol. 1.
- [24] S. Vasudevan, F. Ramos, E. Nettleton, and H. Durrant-Whyte, "Gaussian process modeling of large-scale terrain," *Journal of Field Robotics*, vol. 26, no. 10, pp. 812–840, 2009.
- [25] S. O'Callaghan, F. T. Ramos, and H. Durrant-Whyte, "Contextual occupancy maps using Gaussian processes," in *Proc. IEEE Int. Conf. Robot Automat.* IEEE, 2009, pp. 1054–1060.
- [26] S. T. O'Callaghan and F. Ramos, "Gaussian process occupancy maps," *The Int. J. Robot. Res.*, vol. 31, no. 1, pp. 42–62, 2012.
- [27] V. Tresp, "A Bayesian committee machine," *Neural Computation*, vol. 12, no. 11, pp. 2719–2741, 2000.
- [28] S. Kim and J. Kim, "Building occupancy maps with a mixture of Gaussian processes," in *Proc. IEEE Int. Conf. Robot Automat.* IEEE, 2012, pp. 4756–4761.
- [29] M. Ghaffari Jadidi, J. Valls Miro, R. Valencia Carreño, J. Andrade-Cetto, and G. Dissanayake, "Exploration in information distribution maps," in *RSS Workshop on Robotic Exploration, Monitoring, and Information Content*, 2013.
- [30] —, "Exploration using an information-based reaction-diffusion process," in *Australasian Conf. Robot. Automat.*, 2013.
- [31] M. Ghaffari Jadidi, J. V. Miro, R. Valencia, and J. Andrade-Cetto, "Exploration on continuous Gaussian process frontier maps," in *Proc. IEEE Int. Conf. Robot Automat.*, 2014, pp. 6077–6082.
- [32] J. Wang and B. Englot, "Fast, accurate Gaussian process occupancy maps via test-data octrees and nested Bayesian fusion," in *Proc. IEEE Int. Conf. Robot Automat.*, 2016, pp. 1003–1010.
- [33] F. Ramos and L. Ott, "Hilbert maps: scalable continuous occupancy mapping with stochastic gradient descent," in *Robotics: Science and Systems*, 2015.
- [34] A. Girard, "Approximate methods for propagation of uncertainty with Gaussian process models," Ph.D. dissertation, University of Glasgow, 2004.
- [35] S. T. O'Callaghan, F. T. Ramos, and H. Durrant-Whyte, "Contextual occupancy maps incorporating sensor and location uncertainty," in *Proc. IEEE Int. Conf. Robot Automat.* IEEE, 2010, pp. 3478–3485.
- [36] A. Girard, C. E. Rasmussen, J. Q. nonero Candela, and R. Murray-Smith, "Gaussian process priors with uncertain inputs application to multiple-step ahead time series forecasting," in *NIPS*, 2003, pp. 545–552.
- [37] P. J. Davis and P. Rabinowitz, *Methods of numerical integration*. Academic Press, 1984.
- [38] S. J. Julier and J. K. Uhlmann, "New extension of the Kalman filter to nonlinear systems," in *AeroSense'97*. Int. Society for Optics and Photonics, 1997, pp. 182–193.
- [39] T. P. Minka, "A family of algorithms for approximate Bayesian inference," Ph.D. dissertation, Massachusetts Institute of Technology, 2001.
- [40] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul, "An introduction to variational methods for graphical models," *Machine learning*, vol. 37, no. 2, pp. 183–233, 1999.
- [41] T. Fawcett, "An introduction to ROC analysis," *Pattern recognition letters*, vol. 27, no. 8, pp. 861–874, 2006.
- [42] R. M. Neal, *Bayesian learning for neural networks*. Springer New York, 1996, vol. 118.
- [43] A. Melkumyan and F. Ramos, "A sparse covariance function for exact Gaussian process inference in large datasets," in *IJCAI*, vol. 9, 2009, pp. 1936–1942.