Faithful Euclidean Distance Field from Log-Gaussian Process Implicit Surfaces

Lan Wu[1], Ki Myung Brian Lee[1], Liyang Liu[1], and Teresa Vidal-Calleja[1]

Abstract—In this letter, we introduce the Log-Gaussian Process Implicit Surface (Log-GPIS), a novel continuous and probabilistic mapping representation suitable for surface reconstruction and local navigation. Our key contribution is the realisation that the regularised Eikonal equation can be simply solved by applying the logarithmic transformation to a GPIS formulation to recover the accurate Euclidean distance field (EDF) and, at the same time, the implicit surface. To derive the proposed representation, *Varadhan's formula* is exploited to approximate the non-linear Eikonal partial differential equation (PDE) of the EDF by the logarithm of a linear PDE. We show that members of the Matérn covariance family directly satisfy this linear PDE. The proposed approach does not require post-processing steps to recover the EDF. Moreover, unlike sampling-based methods, Log-GPIS does not use sample points inside and outside the surface as the derivative of the covariance allow direct estimation of the surface normals and distance gradients. We benchmarked the proposed method on simulated and real data against state-of-the-art mapping frameworks that also aim at recovering both the surface and a distance field. Our experiments show that Log-GPIS produces the most accurate results for the EDF and comparable results for surface reconstruction and its computation time still allows online operations.

Index Terms—Gaussian Process Implicit Surfaces, Euclidean Distance Field, Mapping.
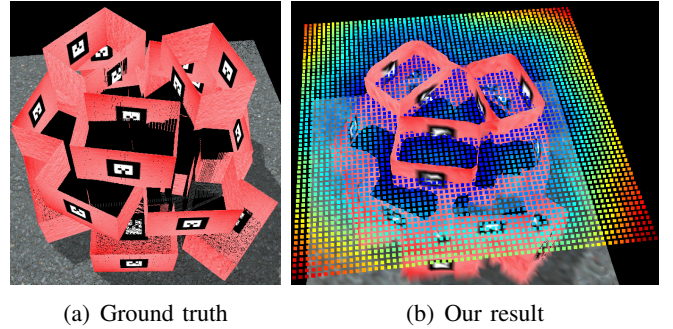

(a) Ground truth   (b) Our result

Fig. 1. 3D reconstruction of a pile of boxes in a simulated environment with (a) the ground truth and (b) the surface reconstruction using Log-GPIS (mesh) and a slice of the Euclidean distance field (point cloud).

I. INTRODUCTION

AUTONOMOUS robots navigating in unknown and unstructured environments require accurate representations of the world. These representations can be used for different purposes such as localisation, mapping, path planning, manipulation, amongst others. Usually, each application requires its dedicated map representation. For instance, navigation requires the robot to build an efficient representation with distance and direction to collision information that are suitable for local or global path planning. On the other hand, recovering a high-quality surface reconstruction of the scene for mapping purposes, requires a dense, accurate and high-resolution representation. In this work, we are interested in world representations that can be used concurrently for accurate reconstruction of the scene and for local navigation.

A wide variety of mapping representations have been proposed in the literature that are suitable for navigation [1], [2], [3], [4] or for surface reconstruction [5], [6], [7], [8]. Notably, occupancy maps for navigation and path planning have received plenty of attention over the years. Occupancy maps vary from discrete and probabilistic 2D [1] and 3D [2] to continuous and also probabilistic [3], [8] representations. On the other hand, a widely adopted mapping representation for surface reconstruction in robotics is the truncated signed distance field (TSDF) used in Kinect fusion [5] to recover a 3D mesh of the scene. Although some authors have exploited it for navigation [9] [10], TSDF is more suitable as an implicit representation of the surface than as a navigation distance field, given that it is a non-accurate approximation of the Euclidean signed distance field (ESDF) [11].

Distance fields are a natural representation for navigation and planning and have been proven advantageous in the past [4]. In fact, Oleynikova *et al.* in [11] show that distance fields can be used for both surface reconstruction and planning seamlessly. In [12] the same authors proposed a discrete mapping framework called Voxblox based on distance fields, which recovers the surface reconstruction of a scene with TSDFs and uses an approximation of the ESDF for local path planning. Despite the fact that Voxblox requires conversion from TSDF to ESDF, the framework runs efficiently online.

Inspired by the idea of distance fields as a unified representation for surface reconstruction and navigation, our work here proposes a single representation that has the property of directly inferring both, the implicit surface and the accurate Euclidean distance field of the scene (see Fig. 1b)). Moreover, unlike most TSDF-based approaches, our proposed representation is continuous and probabilistic.

Gaussian process implicit surface (GPIS) [13], [14] is a recently proposed probabilistic and continuous representation of implicit surfaces that exploits distance information. However, with GPIS, the distance information is not accurate far away from the surface due to the lack of training points. Our work addresses this issue by faithfully modelling a fundamental property of Euclidean distance fields (EDFs) inspired by the heat method [15].

We anticipate that Log-GPIS will be pivotal in the adoption of GPIS in practical robotic mapping and planning applications. As we show, its mathematical foundation provides significant improvement in accuracy and fidelity, while requiring only minor modifications to standard GPIS [14]. We believe that the combination of smooth, probabilistic nature of GPIS and the improved accuracy afforded by our method will unlock the potential of GPIS as a unified representation for robotic mapping and planning. The contributions of this paper are as follow:

- A novel continuous and probabilistic representation that directly infers the EDF and the implicit surface with gradients. Given any query point in space, our representation provides the Euclidean distance and direction to the nearest surface.
- The derivation of the Log-Gaussian Process Implicit Surface (Log-GPIS) by exploiting Varadhan's approach [16] to approximately enforce the Eikonal equation via a linear partial differential equation (PDE). We show that by simply applying the logarithmic function to a GPIS with a specific covariance of the Matérn family, the regularised Eikonal equation is satisfied to accurately infer the EDF and concurrently the implicit surface.
- An implementation suitable for online inference following [14] that can be used for local path planning, which employs the marching cubes algorithm in [17] to recover the surface mesh for visualisation.

The remainder of the paper presents a brief background on EDF and GPIS in Section II followed by the derivation of the Log-GPIS in Section III. The evaluation using simulated and real data is presented in Section IV and the conclusions are in Section V.

II. BACKGROUND

A. Euclidean Distance Field

Consider a manifold $S \subset \mathbb{R}^D$ with boundary ∂S, of which the orientation is given by the normal $\mathbf{n}$. For all $\mathbf{x} \in \mathbb{R}^D$, the EDF $d(\mathbf{x})$ is the closest distance to the boundary ∂S:

$$d(\mathbf{x}) = \min_{\mathbf{y} \in \partial S} |\mathbf{x} - \mathbf{y}| \qquad (1)$$

Consider the scenario where the boundary ∂S dilates along its normal direction with unit speed. Then, the arrival time will be equal to the distance $d(\mathbf{x})$. This intuition is captured by the Eikonal equation [18], which is given by

$$|\nabla d(\mathbf{x})| = 1 \quad \mathbf{x} \in \mathbb{R}^D \qquad (2)$$

with boundary constraints,

$$d(\mathbf{x}) = 0 \quad \text{and} \quad \partial d(\mathbf{x})/\partial \mathbf{n} = 1 \quad \text{on } \mathbf{x} \in \partial S. \qquad (3)$$

Suppose we are given sparse measurements of the points on surface, $\mathcal{X} = \{\mathbf{x}_i\} \subset \partial S$, $i = 1 \ldots N$. The aim of this paper is to estimate $d(\mathbf{x})$ given $\mathcal{X}$, thereby reconstructing S.

B. Gaussian Process Implicit Surfaces

Gaussian process (GP) regression [19] is a well-established non-parametric approach to solving non-linear regression problems. It allows estimating the value of an unknown function at an arbitrary query point given noisy and sparse measurements at other points. In this work, we consider zero-mean GPs:

$$f \sim \mathcal{GP}(0, k(\mathbf{x}, \mathbf{x}')), \qquad (4)$$

where $k(\mathbf{x}, \mathbf{x}')$ is the covariance function that characterises the properties of f.

Suppose we are given a training dataset of measurements $\mathbf{y} = \{y_i = f(\mathbf{x}_i) + \epsilon_i\}_{i=1}^N \subset \mathbb{R}$, taken at locations $\mathbf{x}_i \in \mathbb{R}^D$ and corrupted by additive Gaussian noise $\epsilon_i \sim \mathcal{N}(0, \sigma_f^2)$. The posterior distribution of the value of f at an arbitrary testing point $\mathbf{x}_*$ is given by $f(\mathbf{x}_*) \sim \mathcal{N}(\bar{f}_*, \mathbb{V}[f_*])$, where the predictive mean and variance are given by:

$$\bar{f}_* = \mathbf{k}_*^\top \left(K + \sigma_f^2 I \right)^{-1} \mathbf{y}, \qquad (5)$$

$$\mathbb{V}[f_*] = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top \left(K + \sigma_f^2 I \right)^{-1} \mathbf{k}_*. \qquad (6)$$

Here, I is the identity matrix, and $\mathbf{y}$ represents the input values. $\mathbf{k}_*$ is the vector of covariances between the n_x training points and the testing point. K is the $n_x \times n_x$ covariance matrix of the training points. $k(\mathbf{x}_*, \mathbf{x}_*)$ is the covariance function value of the testing point.

Furthermore, as mentioned above, GP is capable of predicting not only the function values, but also the gradients. This is because differentiation is a linear operator on the space of functions, and consequently the derivative of a GP is another GP [20]. For $f \sim \mathcal{GP}(0, k(\mathbf{x}, \mathbf{x}'))$, the gradient GP ∇f is given by applying the gradient operator ∇ on the covariance function,

$$\begin{bmatrix} f \\ \nabla f \end{bmatrix} \sim \mathcal{GP}(\mathbf{0}, \tilde{k}(\mathbf{x}, \mathbf{x}')), \qquad (7)$$

where $\tilde{k}(\mathbf{x}, \mathbf{x}')$ is the joint covariance matrix of $k(\mathbf{x}, \mathbf{x}')$ and the partial derivatives of $k(\mathbf{x}, \mathbf{x}')$ at $\mathbf{x}$ and $\mathbf{x}'$ [7],

$$\tilde{k}(\mathbf{x}, \mathbf{x}') = \begin{bmatrix} k(\mathbf{x}, \mathbf{x}') & k(\mathbf{x}, \mathbf{x}') \nabla_{\mathbf{x}'}^\top \\ \nabla_{\mathbf{x}} k(\mathbf{x}, \mathbf{x}') & \nabla_{\mathbf{x}} k(\mathbf{x}, \mathbf{x}') \nabla_{\mathbf{x}'}^\top \end{bmatrix}. \qquad (8)$$

Correspondingly, given the training dataset including the observations $\mathbf{y}$ and its gradients $\nabla \mathbf{y}$, where the gradients are computed in the same way as in [14], the following equations express the mean and variance at any testing point with gradient inference,

$$\begin{bmatrix} \bar{f}_* \\ \nabla \bar{f}_* \end{bmatrix} = \tilde{\mathbf{k}}_*^\top (\tilde{K} + \Sigma_f^2 I)^{-1} \begin{bmatrix} \mathbf{y} \\ \nabla \mathbf{y} \end{bmatrix} \qquad (9)$$

$$\begin{bmatrix} \mathbb{V}[f_*] \\ \mathbb{V}[\nabla f_*] \end{bmatrix} = \tilde{k}(\mathbf{x}_*, \mathbf{x}_*) - \tilde{\mathbf{k}}_*^\top (\tilde{K} + \Sigma_f^2 I)^{-1} \tilde{\mathbf{k}}_*. \qquad (10)$$

GP implicit surface (GPIS) techniques [13], [14], [21] [6] use GP regression to estimate the distance field of the surface. Consequently, the surface is given by the zero-level set of

the distance field. Most standard GPIS techniques directly model the distance field as a GP [13], and often use normal measurements to ensure correct results [14]. However, while this allows continuous and probabilistic representation of the surface with high precision and continuity of the distance field near the surface, it is often highly inaccurate further from the surface. For example, using a zero-mean GP implies that the prediction will include false-positive artifacts as the predicted function value falls back to zero. Our method, *Log-GPIS*, alleviates this issue, by elegantly enforcing the Eikonal equation (2) into GP regression.

III. LOG-GAUSSIAN PROCESS IMPLICIT SURFACE

A. Method

As the name suggests, Log-GPIS simply requires that one takes the logarithm of the standard GP regression in (9) to produce accurate EDF estimates near and far away from a surface sampled with a 2D or 3D depth sensor.

We consider GPIS as described in Sec. II-B, equipped with the Matérn family of covariance functions:

$$k(\mathbf{x}, \mathbf{x}') = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}}{\lambda} |\mathbf{x} - \mathbf{x}'| \right)^\nu K_\nu \left(\frac{\sqrt{2\nu}}{\lambda} |\mathbf{x} - \mathbf{x}'| \right),$$

(11)

where the characteristic length scale λ and variance σ^2 are hyperparameters of the GP model. K_ν is the modified Bessel function of the second kind of order ν.

Given a set of points on the surface, $\mathcal{X} = \{\mathbf{x}_i\}$, the distance $\bar{d}_*$ at a query point $\mathbf{x}^*$ is obtained by taking the logarithm of the GP inference result with target measurement set as $y_i = 1$,

$$\bar{d}_* = -\frac{\ln \bar{f}_*}{\lambda},$$

(12)

where $\bar{f}_*$ is the predictive mean from GP inference (9). An immediate benefit of the log-transformation is that the predicted distance approaches infinity as we query points further away from the surface, while the standard GPIS [14] will predict a distance value around zero, resulting in undesirable artifacts. Further, note that for Log-GPIS distance value of $\log(1) = 0$ only occurs on the surface boundary. Thus, Log-GPIS solves the undesirable artifacts in the predicted surface geometry issue.

The log-transformation also affects the gradient. Taking the gradient of both sides of (12) with respect to the distance,

$$\nabla \bar{d}_* = -\frac{1}{\lambda \bar{f}_*} \nabla \bar{f}_*,$$

(13)

where it can be seen that the gradient of $\bar{d}_*$ is in the opposite direction as the gradient of $\bar{f}_*$, subject to a scaling factor $1/\lambda \bar{f}_*$. However, because the Eikonal equation (2) requires that the magnitude of the gradient is normalised to 1, the scaling factor is unimportant, and we simply need to normalise the gradient of $\bar{f}_*$, and invert the sign.

Based on the gradient (13), the predictive variance can be calculated using a first-order approximation:

$$\mathbb{V}[d_*] = \frac{1}{\lambda \bar{f}_*} \mathbb{V}[f_*] \frac{1}{\lambda \bar{f}_*}^\top.$$

(14)

Despite its simplicity, a surprising benefit is that (12) approximately satisfies the Eikonal equation (2) for all query points. In what follows, we present the derivation.

B. Derivation

A major challenge in exploiting the Eikonal equation (2) to recover the EDF is that it is non-linear and hyperbolic. It is therefore difficult to solve directly, and special numerical approaches are often required, including discrete and combinatorial methods such as fast marching or label-correcting to propagate the distance field through a grid (*e.g.* [12], [22]).

A work in the computer graphics literature [15] presents a smooth alternative based on Varadhan's distance formula [16] [23], which approximates the EDF $d(\mathbf{x})$ using the heat kernel on the manifold S. The physical intuition behind Varadhan's formula is as follows. Imagine that the target surface S is hot and emanates heat. The conduction of heat can be viewed as particles taking a random walk starting from the boundary ∂S. If we restrict the duration of the random walks to be short (*i.e.* as time $t \to 0$), the paths taken will be close to the shortest possible one.

Heat conduction on S is modelled by the heat kernel, denoted by $v(\mathbf{x})$. The heat kernel is the solution of the heat equation[1]:

$$(1 - t\Delta)v = (1/t - \Delta)v = 0 \quad \text{in } S,$$
$$v = 1 \quad \text{on } \partial S,$$

(15)

where $\Delta = (\sum_i \frac{\partial^2}{\partial x_i^2})$ is the Laplace operator and t represents time. It is important to note that the heat equation (15) is linear, unlike the Eikonal equation (2).

The celebrated result of Varadhan (Theorem 2.3, [16]) is that the heat kernel $v(\mathbf{x})$ and the EDF $d(\mathbf{x})$ are related as follows[2]:

$$d(\mathbf{x}) = \lim_{t \to 0} \{-\sqrt{t} \ln[v(\mathbf{x})]\},$$

(16)

where it can be seen that the approximation gets better as t approaches zero (*i.e.* after a short duration). Consequently, we define the approximation of $d(\mathbf{x})$ as $u(\mathbf{x})$:

$$u(\mathbf{x}) = -\sqrt{t} \ln v(\mathbf{x}).$$

(17)

In this case, t serves as a smoothing factor of $u(\mathbf{x})$. To check the effect of the approximation, we substitute $v(\mathbf{x}) = \exp(-u(\mathbf{x})/\sqrt{t})$ into (15):

$$v - t\Delta v = v \left[\left(1 - |\nabla u|^2\right) + \sqrt{t}\Delta u \right] = 0,$$

(18)

which shows that we effectively solve the regularised form of the Eikonal equation.

While the theory described thus far only applies to the case of $\mathbf{x} \in S$ (due to the domain constraint in (15)), we can easily extend it to $\mathbf{x} \in \mathbb{R}^D$. To do so, consider the EDF of the complement, S^c. The corresponding heat equation (15) is then

---

[1]Technically, this is a screened Poisson equation. Nonetheless, it exhibits similar behaviours to the conventional heat equation. We use the term heat equation in the interest of physical intuition.

[2]Note that an analogous relationship also holds for the conventional heat equation (Theorem 2.2, [16]), albeit with a more cumbersome counterpart of the log-transform.

defined in S^c. Meanwhile, the boundary condition remains the same, because $\partial S^c = \partial S$. Combining the two, we have:

$$
\begin{aligned}
v - t\Delta v &= 0 \quad \text{in } \mathbb{R}^D, \\
v &= 1 \quad \text{on } \partial S,
\end{aligned}
\tag{19}
$$

in place of (15). Note that, in doing so, we have lost the information on whether $\mathbf{x}$ is inside or outside S (*i.e.* the sign is lost), because we have 'stitched' the EDFs on S and S^c. In our view, the loss of sign is a small price to pay compared to the convenience of using (19) to approximate the (unsigned) distance field everywhere on $\mathbb{R}^D$.

Varadhan's formula (16) replaces the *nonlinear* PDE constraint (2) on $d(\mathbf{x})$ with a *linear* one (19) on $v(\mathbf{x})$, which we can easily incorporate into GP regression. Inspired by [24], we do so by choosing a suitable covariance function using the Wiener-Khinchin theorem [19]. For brevity, we only discuss the 2D case and re-write (15) as:

$$
\frac{\partial^2 v\,(x_1, x_2)}{\partial x_1^2} + \frac{\partial^2 v\,(x_1, x_2)}{\partial x_2^2} - \lambda^2 v\,(x_1, x_2) = w\,(x_1, x_2)
\tag{20}
$$

where $w(x, y)$ is white noise and $\lambda = 1/\sqrt{t}$. The spectral density is the square of Fourier transform of (20) as:

$$
S\,(\omega_1, \omega_2) = \frac{1}{\left(\omega_1^2 + \omega_2^2 + \lambda^2\right)^2} \, .
\tag{21}
$$

The Wiener-Khinchin theorem [19] implies that the covariance function of GP satisfying (19) is given by the inverse Fourier transform of spectral density (21). This is precisely the Whittle covariance function [25], a special case of the Matérn family (11) with $\nu = 1$:

$$
k\,(\mathbf{x}, \mathbf{x}') = \frac{|\mathbf{x} - \mathbf{x}'|}{2\lambda} K_\nu\,(\lambda\,|\mathbf{x} - \mathbf{x}'|) \, .
\tag{22}
$$

Finally, the boundary conditions in (19) translate to the choice of target measurements $y_i = 1$ at locations $\mathbf{x}_i$ on the surface.

C. Practical Considerations

Note that when λ is set to a large value, t in (16) gets closer to zero, thus producing a better distance approximation. As shown in Fig. 2, relatively large λ values will produce a more accurate EDF. However, t cannot be exactly zero as it is in the denominator in $v(\mathbf{x}) = \exp(-u(\mathbf{x})/\sqrt{t})$. As a trade-off, we use $\lambda = 40$ in our experiments. In this paper, hyperparameter

Fig. 3. RMSE varying λ and covariances for the 2D example in Fig. 2.

optimisation is not considered. Nevertheless, λ is exactly the hyperparameter of the Whittle covariance, and (16) provides the necessary intuition to choose a good λ value for the EDF inference.

As sample functions from Matérn family are $|\nu - 1|$ times differentiable, the Whittle covariance is $|\nu - 1| = 0$ times differentiable. This means it can not be used for GPIS gradient inference. Given that the Whittle covariance is a special case of the Matérn covariance family, and Matérn $\nu = 3/2$ is once differentiable, the Matérn covariances are more suitable for distance estimation with gradients.

We empirically observed that a modified version of Matérn $3/2$ with $l = \sqrt{2\nu}/\lambda$ is a good approximation of Whittle covariance with the advantage that it is once differentiable. As an example of this observation Fig. 3 presents the root mean squared error in the case of different λ values of both covariances for the example of Fig. 2 (see demo code[3]). Given that the modified Matérn $3/2$ allows surface normal and distance gradient prediction and Whittle does not, and the approximation is acceptable, we propose to use this Matérn version as a trade-off for extra information.

To summarise, Log-GPIS is capable of producing accurate EDF estimates by simply applying a log-transform to a GPIS with a covariance of the Matérn family. While the resulting EDF now loses the sign, it is possible to recover the sign using the surface normal and sensor's location, as we show in Fig. 4d). Moreover, there is no need to provide sample points inside and outside the surface. Our algorithm only requires the measurements around the surface, and then the EDF will be predicted through the queried points in the field.

IV. EVALUATION

A. Implementation

We implemented the proposed Log-GPIS based on the open source implementation[4] of Lee *et al.* in [14]. Our implementation maintains the online continuous mapping setup as the original code, with modifications that implement the Log-GPIS method described in the previous section to produce a faithful EDF of the scene.
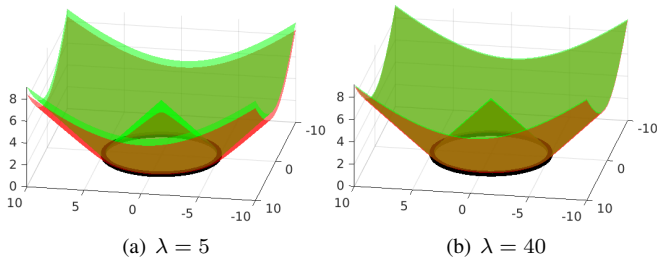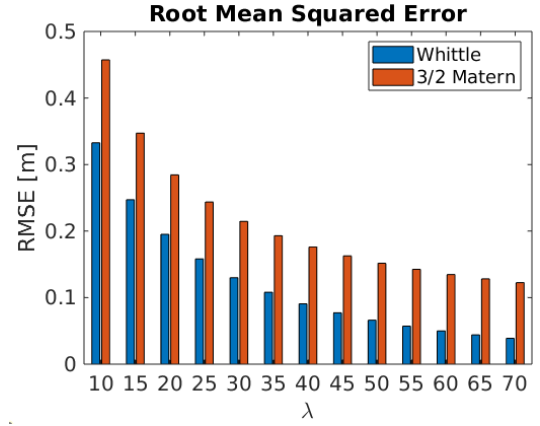

(a) $\lambda = 5$ (b) $\lambda = 40$

Fig. 2. Whittle covariance approximation of the EDF in a 2D example. The measurements are the black points in a circle of 5 m radius. The testing area is a square of 20×20 m. The difference of distance field between the prediction (shown as red color) and the ground truth (shown as green color) for two different λ values. The larger the λ the closer to the true distance.

---

[3]https://github.com/LanWu076/Log-GPIS-demo
[4]https://github.com/leebhoram/GPisMap

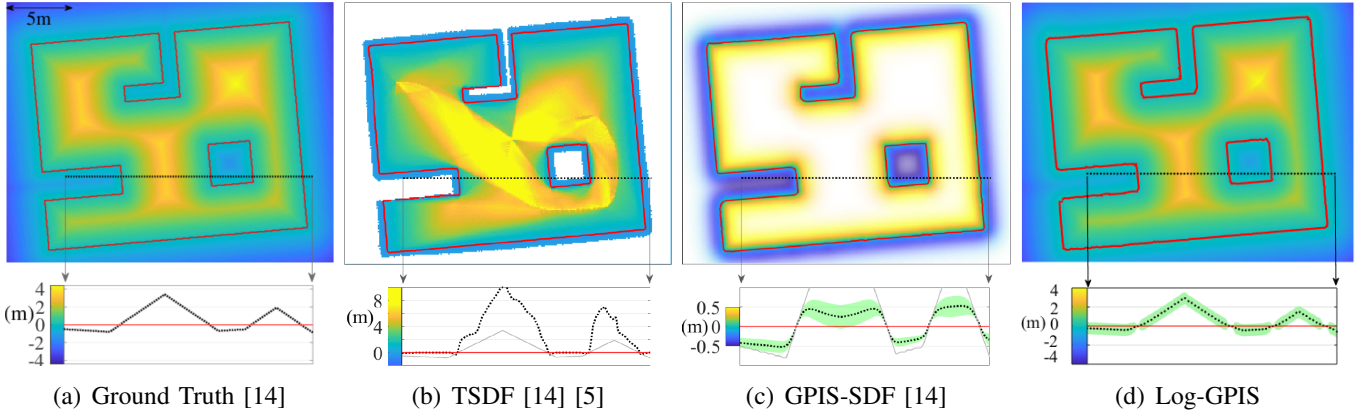(a) Ground Truth [14]　　(b) TSDF [14] [5]　　(c) GPIS-SDF [14]　　(d) Log-GPIS

Fig. 4. Distance field comparison of 2D map dataset. The 1D surface illustrates is shown in solid red line in all maps. The ground truth surface and ESDF map shown in a), then followed by the TSDF map in b) from the results presented in [14]. In c) we present GPIS-SDF [14], which is also based on a GPIS method. d) shows the results of our Log-GPIS method. All methods closely follow the surface, but only Log-GPIS infers the EDF accurately over the whole field. The bottom subplot show the distance value across the black line. Note that the colormap is different for the different maps visualisation purposes, except for the GT and ours. The light green region in c) and d) is variance.

It is well-known that GP-based methods are computationally expensive due to matrix inversion. The computational complexity scales as $\mathcal{O}\left(N^3 D^3\right)$ with gradients, where N is the number of training points and D is the dimension of gradient. Moreover, the computation of predictive mean and variance costs $\mathcal{O}\left(N_x D\right)$ and $\mathcal{O}\left(N_x^2 D^2\right)$ respectively per query point. Therefore, it is challenging to model large maps using most GP approaches. To address this practical issue, we adopt the incremental approach of [14], where QuadTrees [26] and Octree [27] data structures are used to maintain sub-GPIS of clusters, and online Bayesian update of surface points [14] is used to deal with multiple overlapping observations. Furthermore, we modified the original code to render the final mesh using the marching cubes method [28]. To increase the speed and accuracy of the extraction of marching cube vertices and faces, we generate the isosuface with Hermite data [17]. As discussed in Section III, we use Matérn $3/2$ covariance. The code is implemented in C++ and Matlab and runs on an Octa-core i7 CPU at 2.5GHz for all our experimental validations.

In the following sections, we evaluate the performance of Log-GPIS in three scenarios. First, we show the comparison of the true distance value with ours and other approaches on a simulated 2D dataset from [14] with 2D laser observations. Subsequently, we quantitatively evaluate the accuracy of the predicted distance field in a 3D structure simulated in a Gazebo simulation environment with RGB-D image observations. Finally, in order to demonstrate the performance of our algorithm on a real dataset, we provide a quantitative and qualitative evaluation on a publicly available dataset that contains RGB-D images, ground truth sensor poses and structure ground truth. For all experiments, we represent 2D surfaces with continuous lines and 3D surfaces as meshes[5].

B. 2D Results in Simulated Dataset

To ensure fair comparison between the results published in [14], we first verify the accuracy of the Euclidean distance value by reproducing the same 2D simulation environment as

[5]A video is available at https://youtu.be/bSGx_WNdQvo

in [14]. The dataset consists of 28 Lidar scans from different robot poses with angular range of $-135°$ to $+135°$, the sensing resolution of $1°$, and noise $\sigma = 0.01m$.

The results of TSDF [5], GPIS-SDF [14], and our method Log-GPIS against the ground truth are shown in Fig. 4. The sub-figures at the bottom show the distance field along the solid black line. The ground truth in Fig. 4a) shows the linear behaviour of the distance to the surface. Given that this simulation shows the signed distance field, the surface is exactly at the zero crossing. Fig. 4b) shows the distance field estimated by the TSDF method of [5] with constant weights. It turns out that the distance value of TSDF is overestimated within the limited viewpoints. It matches the true Euclidean distance only if the sensor ray is in the direction of surface normal. The benefit of TSDF is that it still produces an accurate continuous implicit surface. Fig. 4c) is the result of what the authors in [14] called GPIS-SDF, which is an online continuous and probabilistic method to recover the implicit surface and the distance field. Note that in this case although the distance value within a small region around the surface is accurate, if the testing points are far away from the measurements, the uncertainty increases and the distance values are not properly estimated. As we can see, the region without measurements nearby is covered by the white color with brighter being more uncertain. Fig. 4d) is the surface and distance field implemented by our method Log-GPIS. As we have mentioned previously, our approach does not contain signed distance information. However, we recover the sign of the Log-GPIS by comparing the predicted gradient of each testing point with sensor position. If the gradient is in the opposite direction of the sensor, we flip the sign of distance value of the testing point. In this figure, it is evident that the EDF is linear and accurately predicted even far from the surface measurements, and the surface is accurately predicted. This shows that our Log-GPIS has the advantage of both mapping and distance accurate estimation.
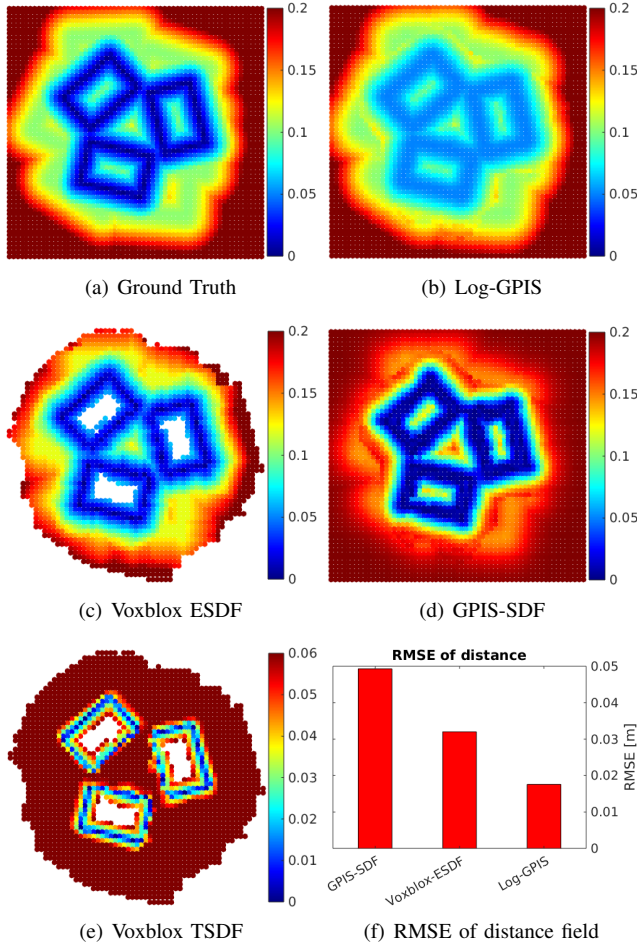
Fig. 5. Quantitative comparison of the distance field on a horizontal slice (size $1.2m \times 1.2m$) of boxes dataset. The slice is 0.5 meter above the ground. All figures are from top view. It can be seen that Log-GPIS estimates the EDF with the lowest root mean square error.

Fig. 6. Same quantitative evaluation as in Fig. 5. The slice (size $3.75m \times 3.2m$) is 1 meter above the ground. All figure are from top view.

C. 3D Results in Simulated and Real datasets

In order to benchmark our method on 3D datasets, we present results from both simulation and real world data.

The simulated dataset consists of a pile of boxes simulated in Gazebo. The simulation environment contains multiple boxes arrange in a three layer pile. The size of each box is $0.3 \times 0.2 \times 0.2m$. The raw measurements are 52 depth images captured by a robot with RGB-D camera moving around the boxes with accurate poses.

The real world experiment is based on the publicly available dataset *cow and lady dataset* [6]. The sensor measurements were obtained from a Kinect 2 RGB-D camera moving in 3D in a small room, with ground-truth poses measured by a Vicon motion capture system. The ground truth structure is a PLY Pointcloud, created from multiple scans with a Leica MS50 professional laser scanner. Since there is minor misalignment between each frame given the ground-truth poses, we use Iterative Closest Point algorithm [29] to correct the poses using the point clouds from consecutive frames.
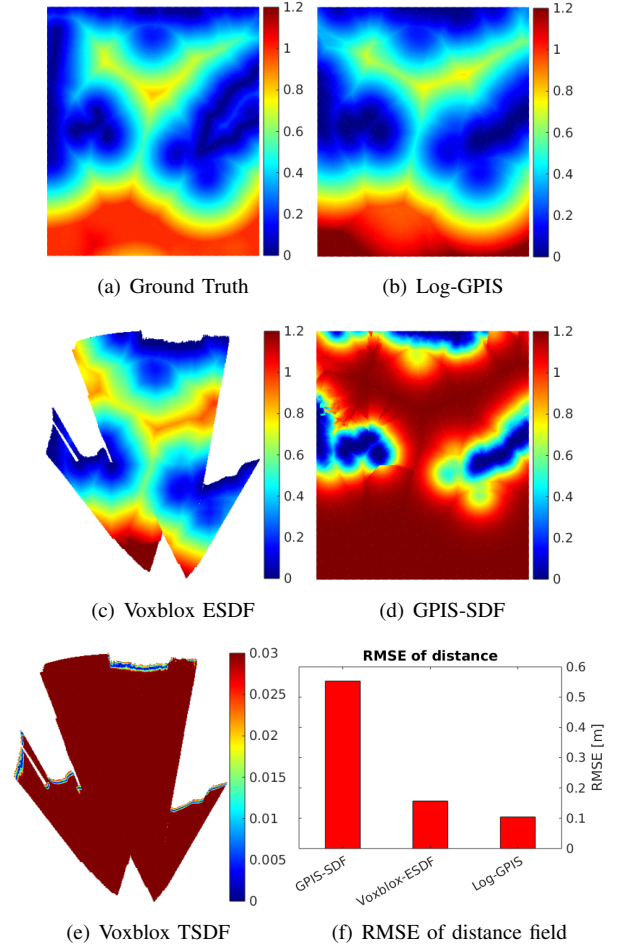
[6]https://projects.asl.ethz.ch/datasets/doku.php?id=iros2017

We benchmark Log-GPIS with two state-of-the-art mapping frameworks; the recently proposed Voxblox [12] and the continuous mapping framework used in the previous section GPIS-SDF from [14]. These two frameworks are particularly relevant to our work as both produce a mesh reconstruction for visualisation purposes and generate concurrently the ESDF for navigation purposes. Given that [12] has validated the accuracy and efficiency of Voxblox against an occupancy representation, we directly compare Log-GPIS with Voxblox. Due to our implementation is based on [14], we simply use the same mapping parameters for GPIS-SDF and Log-GPIS. The parameters used for Voxblox were chosen to make the comparison as fair as possible. We chose a voxel size for simulation and real world datasets as 0.02 and 0.01 respectively, and the truncation distance as 3 voxels. We also chose the merged method as the TSDF Integrator to achieve the best TSDF and ESDF. Log-GPIS and GPIS-SDF testing points, and the Voxblox voxels are at the same resolution.

1) Accuracy of the Euclidean distance field: To study the accuracy of distance field for the different methods, a cross-section of the simulated environment illustrating distance values is shown in Fig. 5 (see also Fig. 1b)). Except for the colormap of the TSDF, the other sufigures are in the same colormap scale. To show clearly the edge of the brick in the
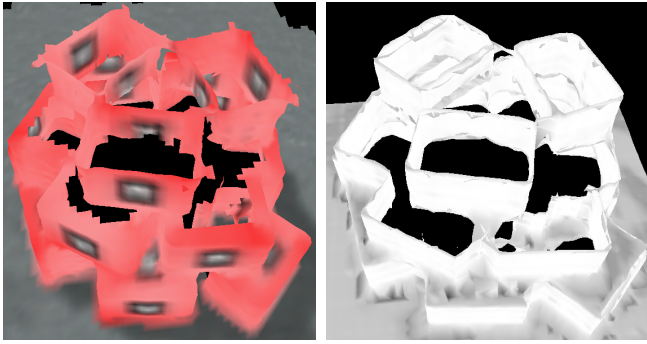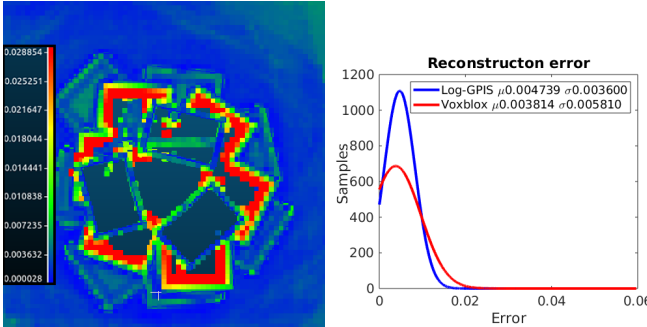
(a) Voxblox mesh   (b) Log-GPIS mesh
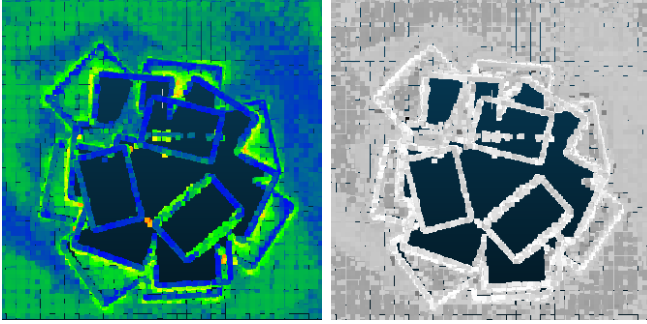
(c) Log-GPIS error   (d) Log-GPIS variance

(e) Voxblox error   (f) Error of vertices

Fig. 7. 3D Reconstruction evaluation in boxes dataset. a) and b) show the meshes for Voxblox and Log-GPIS respectively with their associated reconstruction error in c) and e) and the error distribution in f).

boundary, we limit the colormap to a maximum of 0.2 m for display purposes. In this case, we opted to compare the absolute distances, since the Log-GPIS explicity produces the unsigned distance. Fig. 5a) presents the ground truth Euclidean distance computed from ground truth point cloud. The true surface is shown in black in the centre of the dark blue area. The ESDF and TSDF results from Voxblox are shown in Fig. 5c) and e) respectively. Note that the field is only estimated by Voxblox in the areas within the field of view of the camera for the different poses. Voxblox constructs the TSDF first by grouped raycasting, and then propagating the ESDF based on the relatively accurate TSDF within truncation distance. By considering multiple observations of same region, the ESDF becomes more accurate. As we can see from the results, given the limited number of frames (52 depth images), the ESDF of Voxblox do not produce proper boundaries,
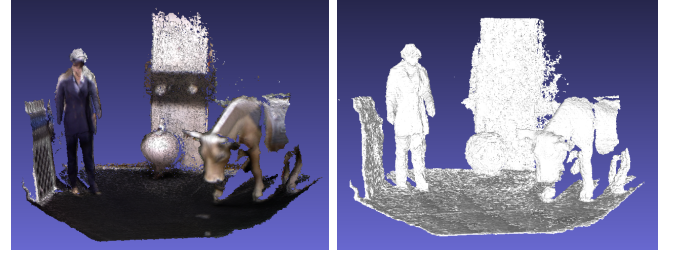
(a) Voxblox Mesh   (b) Log-GPIS Mesh

Fig. 8. 3D reconstruction of the Lady and Cow dataset. a) Voxblox mesh with colour and b) Log-GPIS mesh coloured by variance.

especially the second layer of the boxes (shown in green colour) is not present. As for the TSDF, accurate distance values only within the vicinity of the hit surface. Fig. 5d) presents the results of the distance field from GPIS-SDF. The distance follows a truncated distance field due to the lack of information far away from the surface. Fig. 5a) presents the results of the proposed Log-GPIS EDF. In this case the distance follows very similar trend to the ground truth. Note that Log-GPIS instead of having zero-crossing on the surface, it is an inflection point (minimum value) at the surface. Our algorithm then finds the global minimum to recover the surface, which should be exactly in the middle of the light blue region. The accurate mesh estimation from Log-GPIS is shown in Fig. 1b). As a quantitative measure we compute the root mean squared error for all the approaches. This is shown in Fig. 5f). Note that for Voxblox, we only use the distance in the mapped space. The TSDF is not used in the RMSE evaluation since the distance is only valid inside the truncated region. While Voxblox-ESDF has a relatively good performance, our method produces the most accurate distance estimation.

We use the same benchmark in a more challenging scenario with the real data from the cow and lady dataset. Despite sensor noise and relatively inaccurate poses, the observations of the simulated dataset are similar in the real dataset as Fig. 6 shows. Again the proposed Log-GPIS produces the most accurate distance field with the lowest RMSE amongst Voxblox ESDF and GPIS-SDF.

2) Surface reconstruction: Here, we aim to evaluate the quality of the surface mesh using the same benchmarks as for the distance field. The ground-truth structure for the box dataset is shown in Fig. 1a) and Fig. 7 shows quantitative and qualitative results for Voxblox and Log-GPIS. Note that in this case GPIS-SDF and Log-GPIS produce nearly identical results for the reconstruction of the surface, thus we opted to omit the GPIS-SDF result.

Fig. 7a) and b) present the 3D meshes of Voxblox and Log-GPIS respectively. In order to show the probabilistic information of Log-GPIS, we colored the surface with uncertainty, where darker areas relate to higher variance. Fig. 1b) shows the colored version of the Log-GPIS mesh. Voxblox and Log-GPIS produce high-quality reconstructions. We note that the edges of Log-GPIS mesh are less noisy than Voxblox. Fig. 7c) and e) display the error evaluation of each mesh with respect the ground truth. The points in these two plots are coloured with same scale of colormap shown in Fig. 7e). Interestingly,

the Log-GPIS uncertainty is well correlated with the error as shown Fig. 7d) and c) respectively. The overall error distribution displayed in Fig. 7f) shows that the mesh reconstruction performance for Voxblox and Log-GPIS is equivalent. The median computation time per frame of Voxblox including TSDF, ESDF and meshing is 1.54 s with the selected parameters. Our Log-GPIS consumes a median of 4.56 s including marching cubes for mesh generation, which allows online operations even with surface reconstruction. Note that our method also allows partial evaluation at a faster computation time. Evaluating a horizontal line as a representative use-case for, *e.g.* collision checking for navigation purposes, the median time is reduced to 3.21 s.

Finally, we present qualitative 3D reconstruction results for the real dataset in Fig. 8. Voxblox reconstruction results are shown in Fig. 8a) and the Log-GPIS colored by variance is shown in Fig. 8b), where darker gray relates to high variance. We can see that both algorithms produce similar reconstruction results.

V. CONCLUSION

We proposed a unified representation well-suited for robotic mapping and planning. The novel continuous and probabilistic Log-GPIS provides both, accurate EDF and implicit surface estimation with gradients. The derivation of this representation follows an elegant mathematical formulation exploiting the heat method. The simple yet relevant modification of adding a logarithm transform to the well-known GPIS has the potential to enable wide adoption in practical mapping and planning applications. The results presented here validate the accurate performance of the Log-GPIS with respect to the state-of-the-art method in terms of EDF estimation. In future work, we consider to use our representation for collision check within a local planner. Further development aims at exploiting the provided continuous surface in an ICP-like framework to solve the simultaneous pose estimation and mapping problem.

REFERENCES

[1] H. Moravec and A. Elfes, "High resolution maps from wide angle sonar," in *Proceedings. 1985 IEEE international conference on robotics and automation*. IEEE, 1985, pp. 116–121.

[2] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "Octomap: an efficient probabilistic 3d mapping framework based on octrees," *Autonomous Robots*, pp. 189–206, 2013.

[3] S. T. O'Callaghan and F. T. Ramos, "Gaussian process occupancy maps," *The International Journal of Robotics Research*, 2012.

[4] M. Zucker, N. Ratliff, A. D. Dragan, M. Pivtoraiko, M. Klingensmith, C. M. Dellin, J. A. Bagnell, and S. S. Srinivasa, "Chomp: Covariant hamiltonian optimization for motion planning," *The International Journal of Robotics Research*, pp. 1164–1193, 2013.

[5] S. Izadi, D. Kim, O. Hilliges, D. Molyneaux, R. Newcombe, P. Kohli, J. Shotton, S. Hodges, D. Freeman, A. Davison *et al.*, "Kinectfusion: real-time 3d reconstruction and interaction using a moving depth camera," in *Proceedings of the 24th annual ACM symposium on User interface software and technology*, 2011, pp. 559–568.

[6] L. Wu, R. Falque, V. Perez-Puchalt, L. Liu, N. Pietroni, and T. Vidal-Calleja, "Skeleton-based conditionally independent gaussian process implicit surfaces for fusion in sparse to dense 3d reconstruction," *IEEE Robotics and Automation Letters*, no. 2, pp. 1532–1539, 2020.

[7] W. Martens, Y. Poffet, P. R. Soria, R. Fitch, and S. Sukkarieh, "Geometric priors for gaussian process implicit surfaces," *IEEE Robotics and Automation Letters (RA-L)*, pp. 373–380, 2017.

[8] S. Kim and J. Kim, *GPmap: A Unified Framework for Robotic Mapping Based on Sparse Gaussian Processes*. Springer International Publishing, 2015, pp. 319–332.

[9] R. Wagner, U. Frese, and B. Bäuml, "3d modeling, distance and gradient computation for motion planning: A direct gpgpu approach," in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 3586–3592.

[10] R. Wagner, U. Frese, and B. Bauml, "Real-time dense multi-scale workspace modeling on a humanoid robot," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013.

[11] H. Oleynikova, A. Millane, Z. Taylor, E. Galceran, J. Nieto, and R. Siegwart, "Signed distance fields: A natural representation for both mapping and planning," in *RSS 2016 Workshop: Geometry and Beyond-Representations, Physics, and Scene Understanding for Robotics*. University of Michigan, 2016.

[12] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 1366–1373.

[13] O. Williams and A. Fitzgibbon, "Gaussian process implicit surfaces," 2007.

[14] B. Lee, C. Zhang, Z. Huang, and D. D. Lee, "Online continuous mapping using gaussian process implicit surfaces," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019.

[15] K. Crane, C. Weischedel, and M. Wardetzky, "Geodesics in heat," *ACM Transactions on Graphics*, 2012.

[16] S. R. S. Varadhan, "On the behavior of the fundamental solution of the heat equation with variable coefficients," *Communications on pure and applied mathematics*, pp. 431–455, 1967.

[17] S. Fuhrmann, M. Kazhdan, and M. Goesele, "Accurate isosurface interpolation with hermite data," in *2015 International Conference on 3D Vision*. IEEE, 2015, pp. 256–263.

[18] S. Mauch, "A fast algorithm for computing the closest point and distance transform," *Technical Report caltechASCI/2000.077, Applied and Computational Math., Calif. Inst. of Technology*, 2000.

[19] C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*. Cambridge, Mass.: MIT Press, 2006.

[20] E. Solak, R. Murray-Smith, W. Leithead, D. Leith, and C. Rasmussen, "Derivative observations in gaussian process models of dynamic systems." in *Advances in Neural Information Processing Systems (NIPS)*, 01 2002, pp. 1033–1040.

[21] J. A. Stork and T. Stoyanov, "Ensemble of sparse gaussian process experts for implicit surface mapping with streaming data," *2020 IEEE International Conference on Robotics and Automation(ICRA)*, 2020.

[22] B. Lau, C. Sprunk, and W. Burgard, "Improved updating of euclidean distance maps and voronoi diagrams," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010.

[23] A. G. Belyaev and P. A. Fayolle, "On variational and PDE-based distance function approximations," *COMPUTER GRAPHICS Forum*, pp. 1–16, 2015.

[24] S. Särkkä, "Linear operators and stochastic partial differential equations in gaussian process regression," in *International Conference on Artificial Neural Networks and Machine Learning*, T. Honkela, W. Duch, M. Girolami, and S. Kaski, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 151–158.

[25] P. Whittle, "On stationary processes in the plane," *Biometrika*, 1954.

[26] R. A. Finkel and J. L. Bentley, "Quad trees a data structure for retrieval on composite keys," *Acta informatica*, no. 1, pp. 1–9, 1974.

[27] D. Meagher, "Geometric modeling using octree encoding," *Computer graphics and image processing*, no. 2, pp. 129–147, 1982.

[28] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3d surface construction algorithm," *ACM siggraph computer graphics*, no. 4, pp. 163–169, 1987.

[29] P. J. Besl and N. D. McKay, "Method for registration of 3-d shapes," in *Sensor fusion IV: control paradigms and data structures*. International Society for Optics and Photonics, 1992, pp. 586–606.