

Geometric Priors for Gaussian Process Implicit Surfaces

Wolfram Martens, Yannick Poffet, Pablo Ramón Soria, Robert Fitch, and Salah Sukkarieh

Abstract—This paper presents an extension of Gaussian process implicit surfaces (GPIS) by the introduction of geometric object priors. The proposed method enhances the probabilistic reconstruction of objects from three-dimensional (3-D) pointcloud data, providing a rigorous way of incorporating prior knowledge about objects expected in a scene. The key ideas, including the systematic use of surface normal information, are illustrated with one-dimensional and two-dimensional examples, and then applied to simulated and real pointcloud data for 3-D objects. The performance of our method is demonstrated in two different application scenarios, using complete and partial surface observations. Qualitative and quantitative analysis of the results reveals the superiority of the proposed approach over existing GPIS configurations that do not exploit prior knowledge.

Index Terms—Agricultural automation, categorization, probability and statistical methods, object detection, rgb-d perception, segmentation.

I. INTRODUCTION

GAUSSIAN Process Implicit Surfaces (GPIS) [1] are a tool for non-parametric probabilistic reconstruction of object surfaces from partial and/or noisy 3D data. Interpreting an object's surface as the level-set of an underlying Gaussian process (GP) in 3D allows for typical GP operations, such as probabilistic inference. In particular, given a set of data points on the surface as obtained with a 3D range sensor, GPIS yield the

Manuscript received June 16, 2016; revised September 5, 2016; accepted October 30, 2016. Date of publication November 21, 2016; date of current version December 26, 2016. This paper was recommended for publication by Associate Editor Dr. B. Upcroft and Editor Prof. J. Kosecka upon evaluation of the reviewers' comments. This work was supported in part by the Australian Centre for Field Robotics, in part by the NSW Government, in part by the Australian Research Council's Discovery Project funding scheme under Grant DP140104203, and in part by the Faculty of Engineering and Information Technologies, The University of Sydney under the Faculty Research Cluster Program.

W. Martens and S. Sukkarieh are with the Australian Centre for Field Robotics, The University of Sydney, Camperdown, NSW 2006, Australia (e-mail: w.martens@acfr.usyd.edu.au; salah@acfr.usyd.edu.au).

Y. Poffet is with the Australian Centre for Field Robotics, The University of Sydney, Camperdown, NSW 2006, Australia, and also with the Swiss Federal Institute of Technology Lausanne, École Polytechnique Fédérale de Lausanne, Lausanne 1015, Switzerland (e-mail: yannick.poffet@epfl.ch).

P. R. Soria is with the Australian Centre for Field Robotics, The University of Sydney, Camperdown, NSW 2006, Australia, and also with the Robotics, Vision, and Control Group, University of Seville, Sevilla 41004, Spain (e-mail: prs@us.es).

R. Fitch is with the Australian Centre for Field Robotics, The University of Sydney, Camperdown, NSW 2006, Australia, and also with the Centre for Autonomous Systems, University of Technology Sydney, Ultimo, NSW 2007, Australia (e-mail: rfitch@uts.edu.au).

Color versions of one or more of the figures in this letter are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/LRA.2016.2631260

expectation and variance of the surface contours of an object at arbitrary query points. GPIS are related to GP occupancy maps [2], which construct a spatial GP to probabilistically map the occupancy of a scene, but are more directly targeted at reconstructing compact objects, exploiting the topology of bodies with non-zero volume.

GPIS have strong potential for autonomous perception of objects because they permit rigorous probabilistic reasoning about the geometry of closed volumes. Object perception is a fundamental problem in robotics that underlies a wide variety of tasks ranging from information gathering to manipulation. Probabilistic reasoning about objects is essential for robust approaches to these tasks such as active perception and planning under uncertainty. Further, in many instances, strong *prior knowledge* is available that describes the geometric properties of objects of interest, such as in agricultural robotics where the ability to perceive fruit and other natural objects is of immense value [3], [4]. This paper presents methods for exploiting such prior knowledge during reconstruction of object surfaces and occupancy from range data using GPIS, dramatically expanding their capabilities for real-world applications.

While the concept of using geometric prior knowledge is novel for GPIS, the use of prior shapes has already been explored in the literature on reconstruction and segmentation of 2D and 3D data. [5] and [6] employ elementary shape primitives, whereas [7] and [8] use more complex prior shapes. The methods we propose combine the well-established concept of using prior object shapes with the flexibility and probabilistic capabilities of GPIS.

Previous work on GPIS addresses the probabilistic reconstruction of objects, for example in grasping under uncertainty [9], [10]. In [11], GPIS are used for informative path planning, whereas [12] proposes an active grasp planner based on GP implicit shape potentials, which are closely related to GPIS. Complimentary to visual data, [13] uses a tactile sensor to enhance shape perception.

The linearity of GPs implies that gradient information can be integrated in the regression equations equivalently to function value observations [14]–[16]. As a result, [17] enhances GP reconstruction by integrating surface normal information computed from pointclouds, whereas [9] uses tactile sensors for this purpose. In [18] and [19], surface normal information is introduced by adding points with opposite signs to the interior and exterior of objects.

The main contribution of this paper lies in the introduction of geometric object priors for GPIS which modify the mean function to encode shape primitives in the prior distribution. We propose a library of object shape primitives consisting of spheres, ellipsoids, cylinders and infinite planes that are use-

ful in representing natural objects. By choosing a type of object prior with corresponding geometric parameters and GP covariance function, essential characteristics of an expected object can be exploited to enhance probabilistic operations such as surface reconstruction under uncertainty. While [17] is related to our work in that it also employs a non-zero mean function, their mean function is represented by an additional GP, which is inferred from the observations, instead of using a priori information about the object shape.

In addition to the introduction of geometric priors for GPIS, we also restate the fundamental mathematical concepts of GPIS inference for surface observations and surface normal data in concise form that resolves previous inconsistencies in the literature. We provide the explicit relations for two common GPIS covariance functions, namely the Squared-Exponential and Thin-Plate kernels. We also address the ubiquitous problem of hyperparameter learning, and present an efficient method for rendering object surfaces.

The paper is structured as follows. Section II provides background information on GPIS. In Section III we present the concept of geometric object priors and our library of generic object shapes. Section IV addresses two key components of our implementation, and results are presented in Section V. In Section VI we discuss the implications of our approach and comment on future prospects building on this work.

II. GAUSSIAN PROCESS IMPLICIT SURFACES

A. Gaussian Processes

We consider scalar Gaussian processes (GPs) over d -dimensional spaces, $\mathbf{x} \in \mathbb{R}^d$. A GP can be interpreted as a distribution over functions $f(\mathbf{x})$,

$$f(\mathbf{x}) \sim \mathcal{GP}(m_f(\cdot), k_f(\cdot, \cdot)), \quad (1)$$

and is fully characterised by its mean $m_f(\mathbf{x})$ and covariance (or kernel) function $k_f(\mathbf{x}, \mathbf{x}') = \text{cov}(f(\mathbf{x}), f(\mathbf{x}'))$.

Let $\mathcal{D} = \{y_n\}$ for $n = 1, \dots, N$, be a set of noisy observations $y_n = f(\mathbf{x}_n) + \epsilon_n$, where ϵ_n denotes independent zero-mean Gaussian noise with intensity σ_N . The distribution of $f^* = f(\mathbf{x}^*)$ at some query location \mathbf{x}^* follows from the joint distribution of \mathcal{D} and f^* , and is given by

$$f^* \sim \mathcal{N}(\langle f^* \rangle, \sigma^{*2}), \quad (2)$$

with mean and variance

$$\langle f^* \rangle = m_f(\mathbf{x}^*) + \mathbf{k}^{*T} \mathbf{K}^{-1} (\mathbf{y}_{\mathcal{D}} - \mathbf{m}_{f, \mathcal{D}}), \quad (3)$$

$$\sigma^{*2} = k^{**} - \mathbf{k}^{*T} \mathbf{K}^{-1} \mathbf{k}^*. \quad (4)$$

$\mathbf{y}_{\mathcal{D}} \in \mathbb{R}^{N \times 1}$ and $\mathbf{m}_{f, \mathcal{D}} \in \mathbb{R}^{N \times 1}$ denote the concatenation of the observations and mean values $m_f(\mathbf{x}_n)$, respectively, for all data points, and the covariance matrices are given by

$$\mathbf{K} \in \mathbb{R}^{N \times N}, \quad (\mathbf{K})_{mn} = k_f(\mathbf{x}_m, \mathbf{x}_n), \quad (5)$$

$$\mathbf{k}^* \in \mathbb{R}^{N \times 1}, \quad (\mathbf{k}^*)_n = k_f(\mathbf{x}_n, \mathbf{x}^*), \quad (6)$$

$$k^{**} \in \mathbb{R}, \quad k^{**} = k_f(\mathbf{x}^*, \mathbf{x}^*). \quad (7)$$

For noisy observations, the modified covariance function is given by $k_N(\mathbf{x}_m, \mathbf{x}_n) = k_f(\mathbf{x}_m, \mathbf{x}_n) + \delta_{mn} \sigma_N^2$, where $\delta_{mn} = 1$ if $m = n$ and $= 0$ otherwise.

B. Gaussian Process Implicit Surfaces

Implicit surfaces in Euclidean space $\mathbf{x} \in \mathbb{R}^d$ are induced by the \mathbb{R}^{d-1} -dimensional level sets of continuous functions $f(\mathbf{x})$, $f(\mathbf{x}) = c$, where we set $c = 0$, without loss of generality. The interior and exterior of a surface are interpreted according to

$$f(\mathbf{x}) \begin{cases} > 0, & \text{if } \mathbf{x} \text{ is outside the surface} \\ = 0, & \text{if } \mathbf{x} \text{ is on the surface} \\ < 0, & \text{if } \mathbf{x} \text{ is inside the surface,} \end{cases}$$

with different conventions about the sign in the literature.

In a physical interpretation, the level-sets implied by functions in \mathbb{R}^3 represent the surfaces of objects with non-zero volume. The key idea of GPIS is to let $f(\mathbf{x})$ be distributed according to a GP, so that the surface contour is also random and characterised by the underlying process.

C. Surface normal observations

Additional to the fact that an observed point must lie on the surface of an object, the geometry of neighbouring surface points usually yields the direction of surface normals, which immediately discriminate the interior from the exterior of an object in that region. In the literature there exist two ways of incorporating this topological information into the construction of GPIS. In [1] and [18], the sign of the GP at the inside and the outside of an object is triggered by adding artificial data points with the corresponding sign. A somewhat more elegant way of incorporating the local surface geometry is by use of the surface normals explicitly [9], [17], as will be elaborated in the following.

Derivatives of a GP are also Gaussian, and the covariance between data points and derivatives is readily obtained through differentiation of the covariance function [15],

$$\text{cov} \left(\frac{\partial f(\mathbf{x})}{\partial x_i}, f(\mathbf{x}') \right) = \frac{\partial k(\mathbf{x}, \mathbf{x}')}{\partial x_i} \quad (8)$$

$$\text{and } \text{cov} \left(\frac{\partial f(\mathbf{x})}{\partial x_i}, \frac{\partial f(\mathbf{x}')}{\partial x'_j} \right) = \frac{\partial^2 k(\mathbf{x}, \mathbf{x}')}{\partial x_i \partial x'_j}. \quad (9)$$

Note that for stationary covariance functions, $k(\mathbf{x}, \mathbf{x}') = k(\mathbf{x} - \mathbf{x}')$, as considered in this work,

$$\text{cov} \left(f(\mathbf{x}), \frac{\partial f(\mathbf{x}')}{\partial x'_i} \right) = -\text{cov} \left(\frac{\partial f(\mathbf{x})}{\partial x_i}, f(\mathbf{x}') \right). \quad (10)$$

As a consequence, observations as well as queries of derivatives are incorporated into the general regression equation (2), similar to additional data or query points. If each data point is equipped with full derivative information, the resulting system dimension is $N^+ = (d+1)N$. The covariance sub-matrices (7) are extended by the corresponding covariance terms according to equations (8) through (10). Similar to function value observations y_n , gradient observations are assumed to be noisy, $\nabla y_n = \nabla f(\mathbf{x}_n) + \epsilon_{n, \text{Grad}}$, with noise intensity $\sigma_{N, \text{Grad}}$.

Let $\mathbf{y}_n^+ \in \mathbb{R}^{(d+1) \times 1}$ denote the concatenation of y_n and its gradient ∇y_n . Similarly, assuming sufficient differentiability of $m_f(\mathbf{x})$, $\mathbf{m}_{f, n}^+ \in \mathbb{R}^{(d+1) \times 1}$ denotes the concatenation of the mean function and its gradient at \mathbf{x}_n . Lastly, let $\mathbf{y}_{\mathcal{D}}^+ \in \mathbb{R}^{N^+ \times 1}$ and $\mathbf{m}_{f, \mathcal{D}}^+ \in \mathbb{R}^{N^+ \times 1}$ be the concatenated vectors for all data

points. The joint distribution for function value and gradient at \mathbf{x}^* , similar to equation (2), is then given by

$$\mathbf{f}^{+*} \sim \mathcal{N}(\langle \mathbf{f}^{+*} \rangle, \Sigma^{+*}), \quad (11)$$

with mean and covariance matrix

$$\langle \mathbf{f}^{+*} \rangle = \mathbf{m}^{+*} + \mathbf{K}^{+*T} (\mathbf{K}^+)^{-1} (\mathbf{y}_D^+ - \mathbf{m}_{f,D}^+), \quad (12)$$

$$\Sigma^{+*} = \mathbf{K}^{+**} - \mathbf{K}^{+*T} (\mathbf{K}^+)^{-1} \mathbf{K}^{+*}. \quad (13)$$

The covariance matrices $\mathbf{K}^{+**} \in \mathbb{R}^{(d+1) \times (d+1)}$, $\mathbf{K}^{+*} \in \mathbb{R}^{N^+ \times (d+1)}$ and $\mathbf{K}^+ \in \mathbb{R}^{N^+ \times N^+}$ are computed in analogy to k^{**} , \mathbf{k}^* , \mathbf{K} in (7), making use of the value-derivative and derivative-derivative covariance functions (8) and (9).

According to the definition of interior and exterior points, the gradient at the surface $\nabla f(\mathbf{x}_{\text{surf}})$ must point outwards; however, for general GPs and corresponding implicit surfaces, the magnitude of the gradients at the surface is not fixed. This can be illustrated by the fact that multiplication of $f(\mathbf{x})$ by a scalar changes the magnitude of the gradient field, leaving the zero-level set invariant. Per definition, we will assume a fixed magnitude of the gradient $|\nabla f(\mathbf{x}_{\text{surf}})|_2 = 1$ at surface points, such that the gradient observations ∇y_n are directly given by the surface normals with $|\mathbf{n}_{\text{surf}}|_2 = 1$.

D. Covariance functions

The covariance function $k(\mathbf{x}, \mathbf{x}')$ governs the correlation between data points drawn from a GP and hence determines the smoothness properties of the process. In terms of GPIS, the covariance function can thus be set to control the smoothness properties of the surface of an object.

1) *Squared-exponential covariance function*: Among various applications of GPs, the squared-exponential (SE) covariance function experiences high popularity, partly due to its comparably easy to interpret form, but also because it models well the smoothness characteristics of many random processes. Its isotropic form is given by

$$k_{\text{SE}}(\mathbf{x}, \mathbf{x}') = \sigma^2 \exp\left(-\frac{1}{2} \gamma (\mathbf{x} - \mathbf{x}')^T (\mathbf{x} - \mathbf{x}')\right), \quad (14)$$

with intensity σ and inverse-squared length scale $\gamma = \frac{1}{L^2}$.

Differentiation yields the value-derivative and derivative-derivative covariance functions according to (8) and (9),

$$\frac{\partial k_{\text{SE}}(\mathbf{x}, \mathbf{x}')}{\partial x_i} = -\gamma (x_i - x'_i) k_{\text{SE}}(\mathbf{x}, \mathbf{x}'), \quad (15)$$

$$\begin{aligned} \frac{\partial^2 k_{\text{SE}}(\mathbf{x}, \mathbf{x}')}{\partial x_i \partial x'_j} &= (\gamma \delta_{ij} - \gamma^2 (x_i - x'_i)(x_j - x'_j)) \\ &\times k_{\text{SE}}(\mathbf{x}, \mathbf{x}'). \end{aligned} \quad (16)$$

Note that these relations appear with minor errors in [9], and are shown in corrected form here.

2) *Thin-plate covariance function*: A second popular kernel function for GPIS is given by the *thin-plate (TP)* covariance function which, for the 3D-case, is

$$k_{\text{TP}}(\mathbf{x}, \mathbf{x}') = 2d^3 - 3Rd^2 + R^3, \quad (17)$$

with $d = \|\mathbf{x} - \mathbf{x}'\|_2$ and hyperparameter R . Its name derives from the mechanical properties of a thin plate, whose deforma-

tion energy is lowest for zero second-order derivatives, i.e. in its undeformed state. Note that (17) was first proposed in [1] with an incorrect sign, and corrected in [13].

The value-derivative and derivative-derivative covariance functions for the TP kernel are given by

$$\frac{\partial k_{\text{TP}}(\mathbf{x}, \mathbf{x}')}{\partial x_i} = 6(x_i - x'_i)(d - R), \quad (18)$$

$$\frac{\partial^2 k_{\text{TP}}(\mathbf{x}, \mathbf{x}')}{\partial x_i \partial x'_j} = -6 \left(\frac{(x_i - x'_i)(x_j - x'_j)}{d} + \delta_{ij}(d - R) \right). \quad (19)$$

A drawback of the TP kernel is that these properties can be realised only for a limited domain of a GP (characterised by R) [1]. It also provides no design parameters (like the SE kernel does) to adapt it to desired surface properties.

III. GEOMETRIC OBJECT PRIORS

In most contexts it is convenient to remove the mean of a GP by subtraction of m_f , so that $g(\mathbf{x}) = f(\mathbf{x}) - m_f(\mathbf{x})$ is drawn from a zero-mean GP with identical covariance properties as f . We will show how non-zero mean functions can be used explicitly to define prior functions that model the expected geometry of object surfaces.

A. Library of prior object shapes

1) *Spherical prior*: The most generic object prior for unknown objects is a sphere of radius r . For simplicity, a spherical prior located at $\boldsymbol{\mu} = \mathbf{0}$ is considered first. A suitable mean function is then given by

$$m_S(\mathbf{x}) = \frac{r}{2} (\mathbf{x}^T \mathbf{A}_S \mathbf{x} - 1), \quad (20)$$

with scale matrix $\mathbf{A}_S = r^{-2} \mathbf{I}$. Note that $m_S(\mathbf{x}) = 0$ for all \mathbf{x} with $\|\mathbf{x}\|_2 = r$, satisfying the zero-level set assumption for points on the surface of a sphere centred at $\mathbf{0}$. Next, consider the gradient of the mean function, which is given by

$$\nabla m_S(\mathbf{x}) = r \mathbf{A}_S \mathbf{x}, \quad (21)$$

and satisfies the previous assumption that the gradient should point outwards with magnitude 1 at all points on the surface.

Fig. 1 shows the 1D-prior and characteristics of the resulting GP for an object with radius r , without (Fig. 1(a)) and after two surface observations (Fig. 1(b)). Fig. 2 shows the reconstructed shape of a 2D-object with spherical (circular) prior for increasing numbers of observations. The inferred contour tends towards the circular prior in regions without observations and reproduces the correct shape with increasing accuracy as observations are added.

2) *Ellipsoidal prior*: Generalisation of the spherical prior to object shapes with non-isotropic dimensions is straightforward by modification of \mathbf{A}_S in the spherical prior (20). Let a, b, c be the semi-major axes of an ellipsoid in 3D, and $\mathbf{A}_E = \text{diag}[a^{-2}, b^{-2}, c^{-2}]$. For any $h \in \mathbb{R}$, the mean function

$$m_S(\mathbf{x}) = \frac{h}{2} (\mathbf{x}^T \mathbf{A}_E \mathbf{x} - 1) \quad (22)$$




Fig. 1. 1D-GP for spherical prior with radius r and SE-kernel ($\sigma = 0.5$ and $\gamma = 0.72$). Thick blue lines represent the mean, dashed blue lines represent 3σ -confidence bounds and black lines represent samples from each GP. (a) Prior object shape described by equation (20). (b) Inferred object shape after observing two surface points




Fig. 2. Reconstructed object shape for increasing number of viewpoints (red dashed line). The black line shows the true object shape as sampled from the underlying GP, the shaded regions illustrate the variance in the resulting GP (dark areas indicating high variance). Red circles mark the viewpoint locations and red arrows represent surface normal observations. (a) Prior without observations. (b) Single viewpoint. (c) Two viewpoints. (d) Three viewpoints

satisfies the zero-level set assumption for the surface of an ellipsoid defined by a, b, c . However, unless $a = b = c$, the magnitude-one assumption for the surface gradients will be violated at a non-empty set of points, regardless of how h is selected. Possible choices are, for example, setting h equal to either of a, b, c , resulting in correct gradient magnitudes in the corresponding direction, or setting h equal to their mean.

Fig. 3 shows the reconstruction of the surface of a car from simulated surface observations using an ellipsoidal object prior. Note how the non-isotropic prior helps reconstruct the car's surface in regions without observations, as will be discussed in more detail in Sec. V.

3) *Cylindrical prior*: Cylindrical object priors are interpreted as the limit of an ellipsoid with one of the major axes going to infinity, resulting in a scale matrix $\mathbf{A}_C = \text{diag}[a^{-2}, b^{-2}, 0]$ for $c \rightarrow \infty$, for instance. Fig. 4 shows the reconstruction of a tree trunk using a cylindrical prior. In this example the data points belonging to the tree trunk were manually segmented from the remaining data points.




Fig. 3. Reconstruction of a car from simulated surface observations. (a) CAD model of a car and ellipsoidal prior imposed. (b) Surface reconstruction using ellipsoidal prior




Fig. 4. Reconstruction of a tree trunk using a cylindrical prior.

4) *Infinite plane prior*: While the introduction of implicit surfaces was motivated to model closed surfaces, it may also be useful to use planar GPIS priors. The mean function

$$m_P(\mathbf{x}) = \mathbf{n}^T \mathbf{x}, \quad (23)$$

where $\|\mathbf{n}\|_2 = 1$, describes a linear prior with planar zero-level set. The resulting surface intersects with the origin, and its gradient magnitude is 1 everywhere, thus satisfying the gradient magnitude assumption. In particular, let $\mathbf{n} = [0, 0, 1]^T$, so that $m_P(\mathbf{x}) = x_3$ describes a linear function in 3D, with constant gradient pointing in positive x_3 -direction. The zero-level set is then given by the x_1, x_2 -plane, with the infinite half-volumes for $x_3 < 0$ and $x_3 > 0$ interpreted as the “inside” and the “outside” of the object. A natural application of this prior is for representing a ground surface in 3D-pointclouds, with positive x_3 pointing upwards.

B. General object poses

General poses of observed objects are immediately integrated in the prior functions, as outlined in the following.

1) *Translation*: A translation of the object by a vector $\boldsymbol{\mu}$ is realised by shifting the mean function by an offset, $m^\mu(\mathbf{x}) = m(\mathbf{x} - \boldsymbol{\mu})$. For example, the mean function of a spherical object prior, translated by $\boldsymbol{\mu}$ yields

$$m_S^\mu(\mathbf{x}) = \frac{r}{2} \left((\mathbf{x} - \boldsymbol{\mu})^T \mathbf{A}_S (\mathbf{x} - \boldsymbol{\mu}) - 1 \right). \quad (24)$$

2) *Rotation*: Furthermore, anisotropic objects like ellipsoids, cylinders or planes exhibit rotational degrees of freedom in $SO(3)$ additionally to translation. For a mean function $m(\mathbf{x})$, the result of translation $\boldsymbol{\mu} \in \mathbb{R}^3$ and rotation $\boldsymbol{\theta} \in SO(3)$ is given by the transformed mean function $m^{\boldsymbol{\mu}, \boldsymbol{\theta}}(\mathbf{x}) = m(\mathbf{R}_\theta(\mathbf{x} - \boldsymbol{\mu}))$, where \mathbf{R}_θ denotes the rotation matrix induced by $\boldsymbol{\theta}$. For example, the mean for a general ellipsoid reads

$$m_E^{\boldsymbol{\mu}, \boldsymbol{\theta}}(\mathbf{x}) = \frac{h}{2} \left((\mathbf{x} - \boldsymbol{\mu})^T \mathbf{A}_E^\theta (\mathbf{x} - \boldsymbol{\mu}) - 1 \right), \quad (25)$$

where $\mathbf{A}_E^\theta = \mathbf{R}_\theta^T \mathbf{A}_E \mathbf{R}_\theta$. In the case of a newly observed set of surface points and given an object prior function, $\boldsymbol{\mu}$ and $\boldsymbol{\theta}$ represent a set of unknown parameters that need to be inferred,

for example via maximising the log-likelihood of the data, as discussed in the next section.

IV. IMPLEMENTATION

This section elaborates on two key components needed for the generation of the results presented later in Section V.

A. Learning the hyperparameters

Learning GP hyperparameters is a common task in the literature [14], [18], and we follow the usual approach of maximising the log-likelihood of a training dataset. In the case of zero-mean GPs, the properties of a GP are fully characterised by the parameters of its covariance function. Introducing object priors leads to higher-dimensional optimisation problems, as the geometrical parameters that define the prior also need to be learned.

Optimising the kernel hyperparameters using iterative methods is computationally expensive because it involves re-computation and inversion of the covariance matrix \mathbf{K}^+ in (12). However, for isotropic and stationary kernels, where $k(\mathbf{x}, \mathbf{x}') = k(|\mathbf{x} - \mathbf{x}'|)$, the covariance matrix is invariant to the object prior and depends only on the covariance between observed data points. If the prior parameters of an object, i.e. prior shape and GP kernel, are known, determining the most likely pose reduces to a low-dimensional, inexpensive optimisation problem. Similarly, if any shape parameters are unknown (e.g., the radius of a sphere) the dimension of the optimisation problem stays small, and \mathbf{K}^+ is unaffected.

B. Rendering of GPIS

We provide an outline of how the mean surfaces in our results were generated by efficient region-growing across the surface of an object, allowing us to create smooth surface plots quickly (less than a second). While tools are available for constructing level surfaces, such as `isosurface.m` in MATLAB or implementations of the *marching cube algorithm* [20], these were observed to be inconveniently slow. Standard implementations of these methods take as input a volume of sampled function values, which implies many GP evaluations for points that do not lie on the object surface.

1) *Initialisation*: The proposed algorithm is initialised by choosing a point on or near the zero-level surface. Conveniently, such points are always available as all observed points lie on the surface. It then generates two candidate vertices in a plane perpendicular to the local gradient, with edge lengths of the resulting triangle small enough that the change in function value and gradient is small.

2) *Newton update*: Equation (12) allows the computation of gradients at each candidate point, and a Newton-Raphson method can be applied to find points closer to the level surface, up to any desired accuracy. In our implementation, a single Newton step was sufficient to approximate the closest point on the zero-level. The updated vertex points are added to the mesh, creating an initial triangular face.

3) *Expanding the mesh*: The algorithm then chooses one of the edges of the initial triangle to expand the mesh outwards, perpendicular to the local gradient direction. Similar to the initial two candidate points, the newly created point is then updated by a single Newton step and added to the mesh, creating a new




Fig. 5. Stages of GPIS rendering for an eggplant with frontier “colliding”.

triangular face. The current mesh is thus bounded by the set of its outer edges, its *frontier*, and consecutively spreads out across the object surface.

4) *Frontier splitting and merging*: As the frontier gradually explores the entire surface, it eventually “collides” with itself, leading to termination if the entire surface has been rendered, or a *split* into two frontiers, with each continuing to explore the remaining surface. With multiple frontiers expanding across the surface we also need a *merge* action, which takes place if two separate frontiers “collide”, merging them back into a single one. Fig. 5 shows an example of frontier splitting during surface construction of an eggplant.

V. RESULTS

In this section we evaluate the benefits of object priors in two relevant applications. The performance of our approach is compared with two state-of-the-art GPIS configurations without priors. In particular, we will refer to:

- 1) NZM-SE: Non-zero constant mean > 0 with squared-exponential kernel, similar to [9],
- 2) ZM-TP: Zero-mean with thin-plate kernel,
- 3) SP/EP-SE: Spherical/Ellipsoidal object prior with squared-exponential kernel.

Note that ZM-TP corresponds to the approach proposed in [1], although we take into account surface normal information (Sec. II-C) for this and the other techniques, instead of adding artificial data points. All GPIS hyperparameters as well as prior shape and pose parameters for the selected prior were learned automatically for each object as discussed in Sec. IV-A unless stated otherwise.

In our unoptimised implementation, finding the pose parameters (typically required online), given kernel function and prior shape, took less than a second. Learning the kernel hyperparameters is by far the most expensive operation, as it involves the repeated inversion of \mathbf{K}^+ , and took up to 10 minutes in the considered cases, using a standard optimisation tool (`fminsearch.m`) in MATLAB. If strong prior knowledge is available, as assumed in this paper, this is cast as an offline problem, but it could easily be sped up, for example by using gradient-based optimisers [14].

A. Reconstruction of object surfaces

The first scenario considers GPIS reconstruction from partial as well as complete surface observations of an object, such as obtained from single or multiple viewpoints. Unlike deterministic surface reconstruction methods, our results imply a distribution over surfaces instead of a single solution. For illustration, we render the zero-level mean of the inferred GPIS and indicate standard deviation in colour, where blue indicates small and yellow large values.




Fig. 6. Surface reconstruction of an apple from partial surface observations, colour represents uncertainty (yellow for large standard deviations). (a) Scanned surface. (b) NZM-SE. (c) ZM-TP. (d) SP-SE (our method)




Fig. 7. Surface reconstruction of simulated car from partial surface observations, colour represents uncertainty (yellow for large standard deviations). (a) Poisson surface reconstruction. (b) NZM-SE. (c) ZM-TP. (d) EP-SE (our method).

The first data set is a 3D-scan of an apple, acquired with the Artec Eva 3D-scanner. The second data set is generated in BlenSor [21] from a CAD model of a car. We use a spherical prior for the apple, and an ellipsoidal prior for the car.

Both data sets are downsampled using the method proposed in [22], which also provides surface normals (244 points with surface normals for the apple and 152 for the car). For very noisy or sparse pointclouds, surface normal computation is generally difficult. A benefit of downsampling is to increase the relative voxel size used in surface normal computation, which increases the stability of the surface normals (and hence the quality of surface reconstruction). Further, this paper does not focus on optimised GP inference for large data sets, which involves the inversion of the covariance matrix ($\mathbf{K}^+ \in \mathbb{R}^{4N \times 4N}$ for 3D-pointclouds). Hence, we chose smaller data sets simply for speed, accounting for the ubiquitous trade-off between speed and accuracy using GPs.

The computational complexity of GP inference and the reliability of surface normal computation are general challenges to GPIS using surface normals, and not specific to the use of geometric priors. We hence do not discuss these effects in detail, and observe that all our data sets provide reliable surface normal information for the chosen resolutions.

While the test data sets consist of observations across the entire surface of the objects, we are particularly interested in surface reconstruction for partial observations. Figs. 6 and 7 present results for partial surface observations of the apple and the car, respectively, using NZM-SE, ZM-TP and SP/EP-SE. In Fig. 6(b), NZM-SE reproduces only a degenerate version of the fruit, with a significant part of its volume missing. This effect is to be expected, as the GP is biased towards positive values in regions without observations. Conversely, reconstruction with ZM-TP results in an almost opposite behaviour (Fig. 6(c)). Instead of reducing the fruit’s volume, the surface expands out, as a direct result of the thin-plate covariance function attempting to minimise the second-order derivatives of the underlying GP. Fig. 6(d) shows how for SP-SE, in regions with no data, the reconstructed surface adapts to the spherical prior shape (similar




Fig. 8. Surface reconstruction of car with full data. (a) EP-SE (our method). (b) ZM-TP

to Fig. 2), yielding a reasonable and usable reconstruction of the fruit. Similar results can be observed for the car, shown in Fig. 7. For reference, Fig. 7(a) shows the Poisson surface reconstruction [23] of a simulated pointcloud of size 16994 for the entire surface of the car. EP-SE generates a reasonable reconstruction from partial observations, whereas the other methods produce degenerate results.

Reconstruction for the entire data sets returned reasonable results for all techniques in both test cases, but close inspection of the car’s surface reconstruction using ZM-TP (Fig. 8(b)) reveals a general limitation of the thin-plate covariance function. Its single hyperparameter R , which is usually fixed by the scene dimensions, allows no variation of the smoothness properties of a surface. As a result, the mean surface in Fig. 8(b) exhibits an undesired edginess with high uncertainty even close to the observations.

B. Probabilistic occupancy maps

The second scenario is to use GPIS to build probabilistic occupancy maps (POMs) [24]. POMs are used in a variety of robotics applications, such as path planning and exploration, and are critical in representing uncertainty about occupied or unoccupied space from incomplete observations. Using the convention where object interiors have function values < 0 , the probability of a point being occupied is given by $P(f^* < 0)$. As illustrated in Fig. 9, it is given by the cumulative distribution func-




Fig. 9. Evaluation of occupancy probabilities at two query locations (black crosses) for a 1D-GP, based on three observations of f (red stars).




Fig. 10. Ground truth eggplant model. (a) Model surface. (b) True occupancy

tion of a Gaussian with mean $\langle f^* \rangle$ and variance σ^{*2} , evaluated at zero.

Fig. 10 shows the surface and corresponding ground truth occupancy map of an eggplant model; Figs. 11 and 12 show the surface and POM resulting from simulated complete and partial surface observations with 187 data points and surface normals in the full data set. The resulting 3D-POMs are evaluated by illustration of the horizontal plane at $x_3 = 0$. We compare results obtained using an elliptical prior with squared exponential kernel (EP-SE) against the ZM-TP configuration. In the TP-kernel, R was set manually to achieve good surface reconstruction results.

Fig. 11 shows results for observations across the entire surface of the object, Figs. 11(a) and 11(b) for ZM-TP, and Figs. 11(c) and 11(d) for EP-SE. In [1] it was illustrated how the TP-kernel is superior to the SE-kernel in terms of surface reconstruction for zero-mean functions. However, in the same reference it can also be observed that the resulting variance increases abruptly, even very close to the observations. As a consequence, the POM in Fig. 11(b) exhibits high uncertainty close to the object surface. Conversely, Fig. 11(d) reveals how using an ellipsoidal prior with SE-kernel returns a POM with much higher confidence, corresponding to the true shape and occupancy of the object. It is worth noting how the ellipsoidal prior creates a minor bulge in a region of few observations at the bottom of the mean surface, indicating that a more complex geometric prior might yield an improved result. However, Fig. 11(d) illustrates the great advantage of using GPIS in this context, as poorly reconstructed regions such as this are highlighted as areas of high uncertainty.

In Fig. 12 we analyse the behaviour of the same methods for partial observations. The elliptical prior function includes prior knowledge about the object's geometry, resulting in the reconstruction of the entire expected object occupancy, even in regions without surface observations (Fig. 12(d)). Fig. 12(b), on the other hand, illustrates how for ZM-TP the true occupancy is poorly reconstructed, with uncertainty increasing sharply in regions without surface observations.




Fig. 11. Reconstruction and POM from complete surface observations. (a) ZM-TP surface reconstruction. (b) ZM-TP POM. (c) EP-SE surface reconstruction. (d) EP-SE POM (our method)

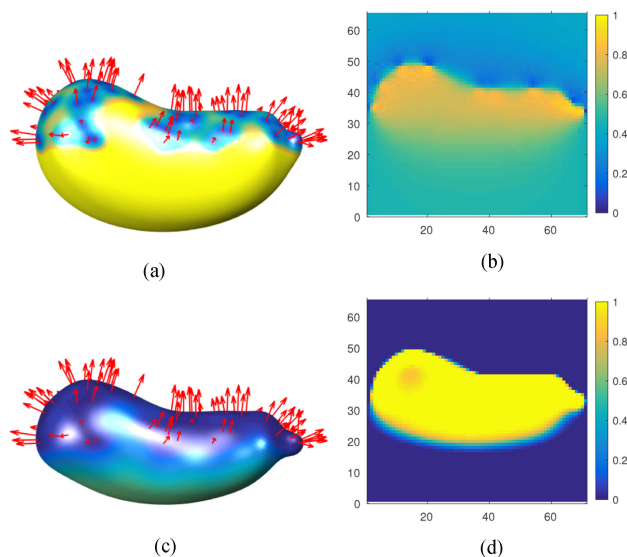


Fig. 12. Reconstruction and POM from partial surface observations. (a) ZM-TP surface reconstruction. (b) ZM-TP POM. (c) EP-SE surface reconstruction. (d) EP-SE POM (our method)

TABLE I
POM EVALUATION USING THE SUM OF SQUARED DIFFERENCES

	Complete observations	Partial observations
NZM-SE	354.2828	310.8350
ZM-TP	407.8959	702.1192
EP-SE	29.4521	191.0672

Table I provides a quantitative analysis of the EP-SE technique against ZM-TP and NZM-SE. As an evaluation metric, we use the sum of squared differences (SSD) between the ground truth occupancy map and the resulting POM. As can be seen in Figs. 11 and 12, the large improvement in the SSD score for non-constant prior functions stems from two reasons: 1) For partial

observations, prior-less methods fail to reproduce significant parts of the object volume; 2) For partial and full observations, the use of a prior decreases uncertainty in regions further away from the observations, exploiting the prior information about the object's shape.

Remark: Determining suitable hyperparameters for the prior-less NZM-SE technique was tedious as there is no established selection method in the literature for the GP mean value. Joint optimisation for the mean and kernel hyperparameters was sensitive to initialisation, and the parameters that globally maximised the data-likelihood yielded a meaningless GPIS reconstruction, with mean close to zero and very short length-scales. We achieved suitable results using a hand-engineered combination of data-likelihood maximisation and manual tuning. In comparison, optimisation for EP-SE and SP-SE was robust and unambiguous.

VI. DISCUSSION AND FUTURE DIRECTIONS

This paper proposes the use of object prior functions to enhance the probabilistic surface reconstruction from noisy or incomplete 3D data using GPIS. The results demonstrate how our approach, in combination with the systematic exploitation of surface normal information, overcomes major limitations of existing GPIS techniques. In addition to the superior results in surface reconstruction and probabilistic occupancy mapping, our approach resulted in stable and unambiguous hyperparameter optimisation. While ZM-TP, depending on only a single hyperparameter R , did not exhibit the same sensitivity as NZM-SE, the existing literature relies on the insertion of artificial data points. It remains unclear how this should be implemented rigorously for general datasets, in particular for incomplete observations.

Besides the applications discussed here we advocate several problems to be investigated in follow-up research.

- 1) *Outlier removal:* Most previous work on GPIS assumed structured scenarios and pre-segmented pointclouds, where the given data belongs only to one object of interest. In more realistic scenarios the pointcloud data is expected to include data belonging to other objects or general noise. Geometric object priors would be a strong tool to discriminate the relevant data from outliers.
- 2) *Multi-object segmentation:* In cluttered scenes with multiple objects, associating data to different objects is inherent to higher-level tasks such as object detection or classification. In our ongoing research we have produced preliminary results for probabilistic segmentation of cluttered scenes using geometric GPIS priors. In this context it is also interesting to investigate inter-object relationships, such as walls being vertical to the ground, similar to [25].
- 3) *Expanding prior libraries:* It is interesting to explore more sophisticated mean functions to produce priors for more complex object shapes. In particular, we plan to investigate the construction of composite priors, built from multiple shape primitives, such as those proposed in this paper. Similar approaches exist in the literature for surface reconstruction [6], however none of these consider GPIS.

REFERENCES

- [1] O. Williams and A. Fitzgibbon, "Gaussian process implicit surfaces," in *Proc. Gaussian Process. Practice Workshop*, 2007.
- [2] S. O'Callaghan and F. Ramos, "Gaussian process occupancy maps," *Int. J. Robot. Res.*, vol. 31, no. 1, pp. 42–62, 2012.
- [3] C. Hung, J. Nieto, Z. Taylor, J. Underwood, and S. Sukkarieh, "Orchard fruit segmentation using multi-spectral feature learning," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2013, pp. 5314–5320.
- [4] S. Nuske, K. Wilshusen, S. Achar, L. Yoder, S. Narasimhan, and S. Singh, "Automated visual yield estimation in vineyards," *J. Field Robot.*, vol. 31, no. 5, pp. 837–860, 2014.
- [5] N. Somani, C. Cai, A. Perzly, M. Rickert, and A. Knoll, "Object recognition using constraints from primitive shape matching," in *Proc. Int. Symp. Visual Comput.*, 2014, pp. 783–792.
- [6] R. Schnabel, R. Wahl, and R. Klein, "Efficient RANSAC for point-cloud shape detection," *Comput. Graph. Forum*, vol. 26, no. 2, pp. 214–226, 2007.
- [7] D. Cremers and S. Soatto, "A pseudo-distance for shape priors in level set segmentation," in *Proc. 2nd IEEE Workshop Variational, Geometric Level Set Methods Comput. Vis.*, 2003, pp. 169–176.
- [8] S. Dambreville, R. Sandhu, A. Yezzi, and A. Tannenbaum, "Robust 3D pose estimation and efficient 2D region-based segmentation from a 3D shape prior," in *Proc. Eur. Conf. Comput. Vis.*, 2008, pp. 169–182.
- [9] S. Dragiev, M. Toussaint, and M. Gienger, "Gaussian process implicit surfaces for shape estimation and grasping," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2011, pp. 2845–2850.
- [10] J. Mahler *et al.*, "GP-GPIS-OPT: Grasp planning with shape uncertainty using Gaussian process implicit surfaces and sequential convex programming," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2015, pp. 4919–4926.
- [11] G. A. Hollinger, B. Englot, F. S. Hover, U. Mitra, and G. S. Sukhatme, "Active planning for underwater inspection and the benefit of adaptivity," *Int. J. Robot. Res.*, vol. 32, no. 1, pp. 3–18, 2013.
- [12] S. Dragiev, M. Toussaint, and M. Gienger, "Uncertainty aware grasping and tactile exploration," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2013, pp. 113–119.
- [13] M. Bjorkman, Y. Bekiroglu, V. Hogman, and D. Kragic, "Enhancing visual perception of shape through tactile glances," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2013, pp. 3180–3186.
- [14] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, 2006.
- [15] E. Solak, R. Murray-Smith, W. E. Leithead, D. J. Leith, and C. E. Rasmussen, "Derivative observations in Gaussian process models of dynamic systems," in *Proc. Neural Inf. Process. Syst.*, 2003, pp. 1057–1064.
- [16] R. Adler, *The Geometry of Random Fields*. Philadelphia, PA, USA: SIAM, 2010.
- [17] S. Kim and J. Kim, "Hierarchical Gaussian processes for robust and accurate map building," in *Proc. Aust. Robot. Autom. Assoc. Australas. Conf. Robot. Autom.*, 2015, pp. 117–124.
- [18] M. Gerardo-Castro, T. Peynot, and F. Ramos, "Laser-radar data fusion with Gaussian process implicit surfaces," in *Proc. Field Serv. Robot.*, 2015, pp. 289–302.
- [19] M. G. López, B. Mederos, and O. Dalmau, "GP-MPU method for implicit surface reconstruction," in *Human-Inspired Computing and Its Applications*, vol. 8856, A. Gelbukh, F. Espinoza, and S. Galicia-Haro, Eds. New York, NY, USA: Springer-Verlag, 2014, pp. 269–280.
- [20] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3D surface construction algorithm," in *Proc. Special Interest Group Comput. Graph.*, 1987, pp. 163–169.
- [21] M. Gschwandtner, R. Kwitt, A. Uhl, and W. Pree, "Blensor: Blender sensor simulation toolbox," in *Proc. Int. Symp. Visual Comput.*, 2011, pp. 199–208.
- [22] J. Papon, A. Abramov, M. Schoeler, and F. Wörgötter, "Voxel cloud connectivity segmentation—Supervoxels for point clouds," in *Proc. Comput. Vis. Pattern Recog.*, 2013, pp. 2027–2034.
- [23] M. B. M. Kazhdan and H. Hoppe, "Poisson surface reconstruction," in *Proc. Eurographics Symp. Geom. Process.*, 2006, pp. 61–70.
- [24] F. Fleuret, J. Berclaz, R. Lengagne, and P. Fua, "Multicamera people tracking with a probabilistic occupancy map," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 2, pp. 267–282, Feb. 2008.
- [25] T. T. Pham, T. J. Chin, K. Schindler, and D. Suter, "Interacting geometric priors for robust multimodel fitting," *IEEE Trans. Image Process.*, vol. 23, no. 10, pp. 4601–4610, Oct. 2014.