Uncertainty-aware geometry processing on Gaussian Process Implicit Surfaces

BAPTISTE GENEST, Université Lyon 1, CNRS, INSA Lyon, France
DAVID COEURJOLLY, CNRS, Université Lyon 1, INSA Lyon, France

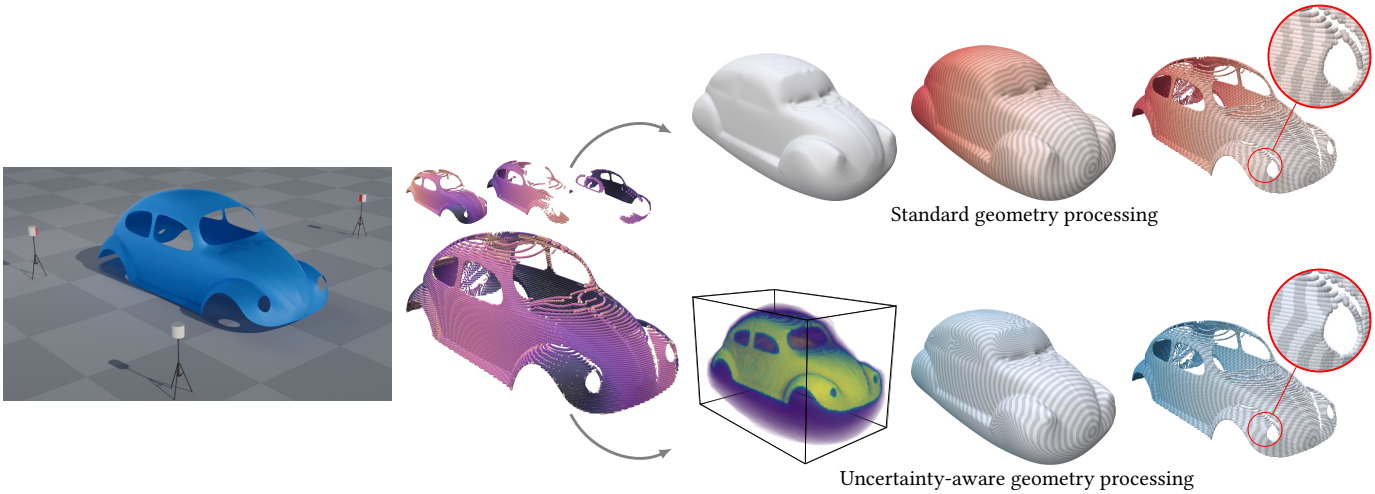
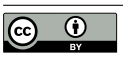


Fig. 1. Starting from a LiDAR simulation producing samples with uncertainty values (left, RIEGL-VZ400 TLS model from `helios++` [Winiwarter et al. 2022]), instead of performing downstream geometry processing tasks (*e.g.,* geodesic computation) on a (Poisson) reconstructed surface and projecting back the results to the point cloud (top row), we first express uncertainty through a simple Gaussian Process Implicit Surface (GPIS) model, which can be efficiently evaluated pointwise. From this GPIS model, we consider its surface density μ and define uncertainty-aware differential operators (bottom row). Where the confidence is high (*e.g.,* on the car doors), our operators (illustrated here with the geodesic distance) behave as in the standard calculus on the mean surface (which coincides with the PSR output). Where confidence is low, *e.g.,* where data is missing, the distribution of plausible geometries is more diffuse, which results in a transition toward a volumetric representation of the shape (PDE solutions are volumetric but plotted on the mean surface and the input for the sake of clarity).

We present a framework for uncertainty-aware geometry processing on Gaussian Process Implicit Surfaces (GPIS), enabling computations directly on such probabilistic representations of shapes. In contrast to classical geometry processing pipelines that assume deterministic surface meshes or point clouds, our approach considers uncertainty in the input data and defines analogs of fundamental differential operators–gradient, divergence, and Laplacian– that account for the distribution of plausible geometries encoded by the GPIS. Leveraging the Kac-Rice formula, we embed computations from random surfaces into a volumetric Cartesian domain, enabling efficient evaluation of expected integrals and differential operators. The proposed approach bridges classical surface PDE-based geometry processing and volumetric representations, enabling a principled handling of noise and ambiguity for various downstream geometry processing tasks.

CCS Concepts: • **Computing methodologies** → **Shape modeling**.

Authors' Contact Information: Baptiste Genest, Université Lyon 1, CNRS, INSA Lyon, France, baptiste.genest@liris.cnrs.fr; David Coeurjolly, CNRS, Université Lyon 1, INSA Lyon, France, david.coeurjolly@cnrs.fr.



This work is licensed under a Creative Commons Attribution 4.0 International License.
© 2026 Copyright held by the owner/author(s).
ACM 1557-7368/2026/7-ART
https://doi.org/10.1145/3811280

ACM Reference Format:
Baptiste Genest and David Coeurjolly. 2026. Uncertainty-aware geometry processing on Gaussian Process Implicit Surfaces. *ACM Trans. Graph.* 45, 4 (July 2026), 14 pages. https://doi.org/10.1145/3811280

1 Introduction

When processing data from complex simulations or captured measurements, uncertainty in the quality of the data can critically impact the reliability of downstream applications. Modeling input uncertainty and propagating it through processing pipelines are key challenges for ensuring computational reliability. In the context of geometry processing, our goal is to compute geometric quantities or solve numerical partial differential equations on the probability distribution of the possible geometries of a (partially known) shape.

In particular, Gaussian processes (GPs) provide a simple yet powerful framework to model uncertainty. In short, a Gaussian process is a continuous collection of random variables (*i.e.,* a random *field*), where any finite subset of these variables follows a multivariate Gaussian distribution. Combined with the theory of implicit surfaces, GPs have also been used recently to directly encode a shape. Gaussian Process Implicit Surfaces (GPIS) enrich the usual implicit description of surfaces by providing a principled way to quantify the likelihood that a point belongs to the shape (*i.e.*, an iso-surface of the field) or the distribution of its tangent planes at each point.

This probabilistic representation of surfaces raises novel challenges for geometry processing. Rather than reconstructing a deterministic surface mesh and applying classical processing pipelines, we aim to perform processing *directly* on the distribution of shapes, preserving the influence of uncertainty throughout the pipeline.

To address these challenges, our contributions are threefold. First, we avoid the computational burden of Gaussian process regression by introducing a closed-form GPIS model that can be evaluated very efficiently by simply modeling how the uncertainty, induced by noise, on the input is transferred to the output reconstructed field. Secondly, we leverage tools from stochastic geometric measure theory to derive probabilistic analogs of fundamental differential operators such as gradient, divergence, and Laplacian on such probabilistic surfaces. Finally, we demonstrate that the proposed differential operators enable the extension of many classic geometry processing tools to GPIS, explicitly accounting for shape uncertainty.

2 Related works

We review previous works that address uncertainty in surface representation and computation, with applications in computer graphics and applied mathematics. For each section, we specify whether the inputs and outputs are deterministic or stochastic.

Probabilistic surface representation (deterministic input, stochastic output). From a set of points, Pauly et al. [2004] explore the distribution of shapes that passes through them. While this approach is purely surface-based, alternative approaches focus on implicit representation, assuming that the shape is the iso-surface of a scalar field. Williams and Fitzgibbon [2006] use Gaussian process (GP for short) regression to build an implicit surface by inferring the indicatrix of the shape from a set of samples, labeled as inside or outside of the shape. This representation has been widely used in the robotics community [Dragiev et al. 2011; Wu et al. 2024]. Sellán and Jacobson [2022] interpret the classic Poisson surface reconstruction algorithm of Kazhdan et al. [2006] as a GP regression process, which augments the PSR reconstruction with a Gaussian distribution over the output, at the cost of having to solve many large linear systems to obtain the variance. Follow-up work by Holalkere et al. [2025] proposes a way to significantly speed-up the regression process by using recent tools from the Gaussian process community by replacing the multiple linear solves by a single one. Pujol and Chica [2025] use a Bayesian approach to build a Gaussian representation of a shape from a noisy point cloud by first estimating normals using Bayesian PCA and then formulating a reconstruction scheme that is able to propagate uncertainty to the output. This work also requires solving large linear systems.

While the term "Gaussian Process Implicit Surface" tends to refer to the techniques that use Gaussian process *regression* directly to obtain the scalar field inducing the iso-surface, we use it to designate any iso-surface of a Gaussian random field φ, *e.g.,* at each point of the domain Ω, the field value follows a Gaussian distribution:

$$\forall x \in \Omega, \varphi(x) \sim \mathcal{N}\left(\gamma(x), \sigma(x)\right). \tag{1}$$

for some real valued mean function $\gamma(x)$ and variance $\sigma(x)$ we clarify later. The differential operators we propose can be built independently of how the field is obtained.

Probabilistic numerics and computing (stochastic input, stochastic output). The field of probabilistic numerics aims to quantify how the uncertainty in the data defining a PDE (*e.g.,* coefficients or initial value) transfers to the solution [Cockayne et al. 2017; Wang et al. 2021]; see Hennig et al. [2022] for an overview. In contrast, our approach defines deterministic operators constructed on stochastic input geometry.

PDE on unknown manifolds (deterministic input, stochastic output). Closer to our work, several papers show how to derive computations when the geometry of the domain is only known at a few locations. Computations can be performed either on the input points directly by inferring tangent spaces [Li et al. 2025], or through parameterization using a Riemannian metric [Niu et al. 2023; Peoples and Harlim 2024; Tosi et al. 2014]. Our approach works differently, as we consider the complete implicit field.

Noise-aware geometry processing (deterministic input, deterministic output). Several works in geometry processing do not operate on uncertain data directly but consider the hypothetical impact of randomness (typically noise) in their derivations, making algorithms more robust for surface reconstruction [Pöthkow et al. 2011], or mesh simplification [Trettner and Kobbelt 2020]. Our approach is similar in the sense that the output of the algorithm is deterministic and is governed by a probabilistic approach. In contrast, we consider the probability distribution itself as the input.

Embedding methods (deterministic input, deterministic output). As we carry out computation on the Euclidean domain surrounding the shapes, our method can be seen as related to *embedding methods*. For example, the Closest Point Method [Ruuth and Merriman 2008] solve PDEs on surfaces by alternating between computation in the ambient space and a re-projection step on the surface. For a comprehensive review of such techniques in computer graphics and in applied mathematics (for deterministic surfaces), we refer to King et al. [2024]. However, a a key distinction is that the operators we propose are fundamentally volumetric, while embedding methods use the 3D space to solve PDEs on 2D surfaces.

Stochastic geometry in rendering and visualization (stochastic input, deterministic output). The field of uncertainty visualization also processes random fields and the geometry of their level-sets (see Brodlie et al. [2012] for a review). For example, Pfaffelmoser et al. [2011] use the level-set crossing probability to visualize the likelihood that a random surface reaches a region of space. In rendering, random geometry has been widely used to emulate the behavior of light on highly detailed surfaces, for example on the micro-facet model. Recently, Miller et al. [2024] and Seyb et al. [2024] propose unified models for light scattering for the rendering of stochastic volumes and surfaces, mostly defined by Gaussian processes. Seyb et al. [2024] extend light transport to GPIS by considering averages over realizations of the entire Gaussian field, while we only consider realizations of the level sets. In a follow-up work, Xu et al. [2025], synthesize GPIS for stochastic geometries in a fashion similar to Perlin noise. They achieve constant time evaluation by using compact support kernels from Gaussians defined on a (multiscale) grid. Such an approach is not applicable for surface reconstruction, since the kernels we consider typically require very long-range interactions

(*e.g.;* the PSR kernel decreases quadratically with the distance) from unstructured data points. Furthermore, this approach still requires solving, possibly large, sparse linear systems to condition the GPIS to fit to a given geometry, something that we avoid completely.

3 Background

3.1 Gaussian variables and Gaussian processes

First, a d dimensional Gaussian vector $X = (X_1, \dots, X_d)$ follows a multivariate normal distribution characterized by its mean vector $\mathbb{E}[X] \in \mathbb{R}^d$ and covariance matrix $\Sigma_X \in \mathbb{R}^{d\times d}$:

$$X \sim \mathcal{N}(\mathbb{E}[X], \Sigma_X). \tag{2}$$

A Gaussian Process (GP) φ over a domain Ω is a random field such that, for any finite set of points $\{x_i\} \in \Omega$, $(\varphi(x_1), \dots, \varphi(x_n))$ is a Gaussian vector. We recall two properties of Gaussian vectors that will be used later (see Seeger [2004] for further results). First, Gaussian vectors are stable by affine maps, given $A \in \mathbb{R}^{d'\times d}, b \in \mathbb{R}^{d'}$:

$$AX + b \sim \mathcal{N}\left(A\mathbb{E}[X] + b, A\Sigma_X A^\top\right). \tag{3}$$

Then, given two Gaussian vectors X and Y such that the pair (X, Y) is Gaussian, the conditioning of X under the event that Y takes the values c is also Gaussian and expressed as:

$$X|Y = c \sim \mathcal{N}\Big(\mathbb{E}[X] + \Sigma_{X,Y}\Sigma_{Y,Y}^{-1}(c - \mathbb{E}[Y]), \Sigma_{X,X} - \Sigma_{X,Y}\Sigma_{Y,Y}^{-1}\Sigma_{Y,X}\Big), \tag{4}$$

where $\Sigma_{X,Y}$ is the matrix of cross-covariances $(\mathrm{Cov}(x_i, y_j))_{i,j}$.

3.2 Gaussian Process regression

Gaussian processes are widely used in Machine Learning, as the closed form conditioning of Gaussian variables can be used to perform regression. Given a set of n observations $\{y_i\}$ at locations $\mathbf{x} \coloneqq \{x_i\} \in \Omega$, and an explicit covariance function $k(\cdot,\cdot)$ (for example, a kernel $k(u,v) = e^{-\kappa\|u-v\|^2}$ for some $\kappa > 0$), we aim to infer the value y_* at an unknown location x_*, assuming that the data is generated from a Gaussian process φ, i.e. $Y = (\varphi(x_1), \dots, \varphi(x_n))$:

$$\begin{pmatrix} Y \\ Y_* \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} f(\mathbf{x}) \\ f(x_*) \end{pmatrix}, \begin{pmatrix} K & K_* \\ K_*^\top & k(x_*, x_*) \end{pmatrix}\right),$$

where $(K)_{i,j} = k(x_i, x_j) \in \mathbb{R}^{n\times n}$, $(K_*)_i = k(x_i, x_*) \in \mathbb{R}^{n\times 1}$, and f is a prior function, usually set to 0. Using Eq. 4, the Gaussian distribution of the prediction is given by:

$$Y_*|Y_1 = y_1, \dots, Y_n = y_n.$$

In practice, the conditioning requires computing the inverse of the matrix K ($\Sigma_{Y,Y}^{-1}$ term in Eq. 4), which is often dense. This leads to an $O(n^3)$ complexity for pre-computation and a linear complexity at inference at a given point x_* for computing K_*. This process can be made more scalable using stochastic gradient descent as discussed by Lin et al. [2023]).

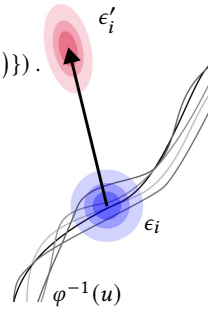
4 A linearized perturbation GPIS reconstruction model

Whereas Gaussian process *regression* is computationally expensive, we propose a simple and efficient way to define a Gaussian process implicit surface (GPIS). Instead of regression, our approach focuses on expressing how the uncertainty on the input data impacts the output of a surface reconstruction scheme (see App. A for further discussion on the difference of interpretation with previous GPIS models). Given pairs $\{(p_i, \omega_i)\}$ of positions and (weighted) normal vectors in $\mathbb{R}^d$, pointwise implicit surface reconstruction schemes (*e.g.;* [Guennebaud and Gross 2007]) can be generically expressed as:

$$\gamma(x) = \lambda(x, \{(p_i, \omega_i)\}).$$

for some real-valued function λ. Let us consider random centered Gaussian perturbations of the position and normals of each data point, *i.e.* with $\epsilon_i \sim \mathcal{N}(0, \Sigma_{\epsilon_i})$ and $\epsilon_i' \sim \mathcal{N}(0, \Sigma_{\epsilon_i'})$:

$$\tilde{\gamma}(x) \coloneqq \lambda(x, \{(p_i + \epsilon_i, \omega_i + \epsilon_i')\}).$$



The output distribution of $\tilde{\gamma}$ is not Gaussian in general. To obtain a Gaussian output, the expression must be affine with respect to ϵ and ϵ' [Nabeya and Kariya 1986]. Hence, we linearize the perturbation with respect to the input (positions and normals):

$$\varphi(x) \coloneqq \gamma(x) + \sum_{i=1}^{n} \nabla_{p_i}\lambda(x, \{(p_i, \omega_i)\}) \cdot \epsilon_i + \nabla_{\omega_i}\lambda(x, \{(p_i, \omega_i)\}) \cdot \epsilon_i'. \tag{5}$$

We thus retrieve a pointwise view of a Gaussian process, instead of a regression one (see Xu et al. [2025] for a discussion between these representations). While this linearization process from Gaussian noise on the input is a standard trick to study the robustness of an estimation algorithm [Lachaud et al. 2023], to the best of our knowledge, no prior work uses it to define probabilistic surface models. In the specific case of triangle soups, we use a simple geometric model to define both ϵ_i and ϵ_i' values from a single parameter σ to easily control the variance of the GPIS (see App. B).

We further assume that the reconstruction scheme is derived from a potential ρ:

$$\lambda(x, \{(p_i, \omega_i)\}) = \sum_{i=1}^{n} \nabla\rho(x - p_i) \cdot \omega_i.$$

In dimension 3, it is motivated by the regularized winding number, introduced by Chen et al. [2024] as a closed-form solution to the Poisson surface reconstruction, with:

$$\rho(x) = \frac{\mathrm{erf}_s(\|x\|)}{\|x\|}, \tag{6}$$

where erf_s is the CDF of a centered Gaussian variable of standard deviation s (where s intuitively controls the level of smoothness of the reconstruction).

Our closed form estimate (Eq. 5) provides direct access to the joint distribution of the field values *and* its gradients by decomposing

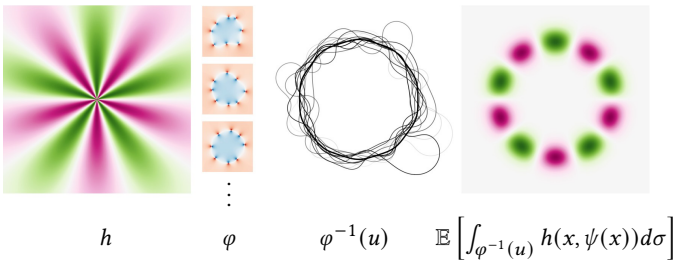


Fig. 2. The (generalized) Kac-Rice formula allows to compute the expected value of the integral of an ambient field h, over the random iso-surfaces $\varphi^{-1}(u)$ of a Gaussian process φ. The expected integral weights $h(x)$ by the likelihood that the $\varphi^{-1}(u)$pass through x. (Eq. (9)). Here, h does not depend on the random field ψ for the right figure.

the contribution of each random vector:

$$\begin{pmatrix}\varphi(x)\\ \nabla\varphi(x)\end{pmatrix} = \begin{pmatrix}\gamma(x)\\ \nabla\gamma(x)\end{pmatrix} + \sum_{i=1}^{n}\begin{pmatrix}(D^2\rho(x-p_i)\omega_i)^\top\\ D^3\rho(x-p_i):\omega_i\end{pmatrix}\epsilon_i + \begin{pmatrix}\nabla\rho(x-p_i)^\top\\ D^2\rho(x-p_i)\end{pmatrix}\epsilon_i' \tag{7}$$

where $D^2\rho, D^3\rho$ are respectively the Hessian and the third-order differential tensor of ρ, and " : " is tensor contraction. Then, the joint distribution can be computed using Eq. 3 as everything is expressed as sums of 4×3 matrix-(Gaussian) vector products. Finally, the main advantage of using such a pointwise scheme is to enable the use of accelerating schemes *à-la* Barnes and Hut [1986] to bring the evaluation of Eq. 7 from $O(n)$ down to $O(\log(n))$ time (see Sec. 9.1).

5 The Kac-Rice formula

From this point on, our formulation applies to any GPIS model, whether constructed as in Sec. 4 or following prior work of Sellán and Jacobson [2022] (see App. F for discussion). We consider random fields defined over Ω, a compact domain in $\mathbb{R}^d$.

Our core contributions rely on a specific variant of the Kac-Rice formula. For a comprehensive review of the related field and its standard usage, we refer to Azäis and Wschebor [2009]. In geometric measure theory, the co-area formula states that the integral over the level sets of a smooth field γ of a function $h \in L^1(\Omega)$ can be expressed as an integral over the domain:

$$\int_{-\infty}^{\infty}\int_{\gamma^{-1}(u)} h(x)dH_{d-1}(x)du = \int_\Omega h(x)\|\nabla\gamma(x)\|dx\,, \tag{8}$$

where dH_{d-1} is the $(d-1)$-Hausdorff measure on the iso-surface $\gamma^{-1}(u)$. Intuitively, the integral of a field restricted to a level-set must be weighted by the norm of the gradient of the field to compensate the speed at which the left-hand side integral will sweep over each level-set.

The Kac-Rice theory extends this formula to the stochastic case: given a Gaussian random field φ, $\varphi^{-1}(u)$ is a random surface (Fig.2-third column), and the integral over such random level-set is also a random variable. We aim to compute the expected value:

$$\mathbb{E}\left[\int_{\varphi^{-1}(u)\cap\Omega} h(x)\sigma_u(\varphi,dx)\right],$$

where σ_u is a generalization of the Hausdorff measure to a random level-set [Azäis and Wschebor 2009], simply denoted dA hereafter.

Under mild regularity conditions on φ and with the integrand h that might depend on another (possibly vector valued) random field ψ such that $(\varphi(x),\psi(x))$ is Gaussian for all $x\in\Omega$, the Kac-Rice formula (in the scalar-field case) provides a closed-form expression for the expected value:

$$\mathbb{E}\left[\int_{\varphi^{-1}(u)\cap\Omega} h(x,\psi(x))dA\right] = \int_\Omega p_{\varphi(x)}(u)\mathbb{E}\left[\|\nabla\varphi(x)\|h(x,\psi(x))|\varphi(x)=u\right]dx\,, \tag{9}$$

where $p_{\varphi(x)}$ is the probability density function (PDF) of the 1D Gaussian variable $\varphi(x)$.

While this formula is more involved than its deterministic counterpart, its intuition remains clear: since φ is a random Gaussian field, the level-set $\varphi^{-1}(u)$ has a non-zero probability of passing through any point $x\in\Omega$. The expectation therefore aggregates contributions of the co-area formula applied to all possible realizations of the level-set $\varphi^{-1}(u)$ conditioned to run through each point x of the domain, weighted by the likelihood $p_{\varphi(x)}(u)$ of this event.

The Kac-Rice formula variants have been used in pure mathematics to estimate geometric and topological features the level sets of simple random models (for example the expected Euler characteristic). We refer to Berzin et al. [2022] for a review of some applications. To the best of our knowledge, the Kac-Rice formula has not been used to embed computations and differential operators from random surfaces to the surrounding Euclidean domain Ω (see Sec. 7).

6 A geometrically coherent notion of surface density for GPIS

Previous works on GPIS [Pujol and Chica 2025; Sellán and Jacobson 2022; Williams and Fitzgibbon 2006] define the likelihood that point x belongs to the iso-surface $\mathcal{M}=\varphi^{-1}(u)$ as being proportional to the likelihood that the field at x equals u, i.e.:

$$\Pr(x\in\mathcal{M}) \propto p_{\varphi(x)}(u)\,.$$

While this definition is intuitive in the context of Gaussian fields, as it mimics a Dirac delta distribution over the shape, the Kac-Rice formula hints at a slight modification to enrich it with a meaningful notion of mass. Considering the expected uniform sampling (non-normalized) distribution ($h(x,y)=1$ in Eq. 9) over the random iso-surfaces yields:

$$\mathbb{E}\left[\int_{\varphi^{-1}(u)} 1dA\right] = \int_\Omega p_{\varphi(x)}(u)\mathbb{E}_{\varphi=u}\left[\|\nabla\varphi(x)\|\right]dx\,,$$

with the notation $\mathbb{E}_{\varphi=u}[X] \coloneqq \mathbb{E}[X|\varphi(x)=u]$. Hence, we consider:

$$\Pr(x\in\mathcal{M}) \propto \mu(x)\,, \tag{10}$$

with μ as a surface density for the levet-set u:

$$\boxed{\mu(x) \coloneqq p_{\varphi(x)}(u)\mathbb{E}_{\varphi=u}\left[\|\nabla\varphi(x)\|\right]\,.} \tag{11}$$

One can check that the total mass of μ recovers the mean area of the GPIS:

$$\int_\Omega \mu(x)dx = \mathbb{E}\left[\mathrm{Area}(\varphi^{-1}(u))\right],$$

which ensures that it defines a coherent integration measure.



Fig. 3. From the input geometry on the left, we first build a Gaussian field φ such that its mean coincides with the scalar field obtained by the surface reconstruction scheme, here the regularized winding number [Chen et al. 2024] (center left). Our calculus framework relies mostly on the iso-surface density μ (center right) and the tangent plane distribution tensor field T (right). Where data is missing, here around the hole, both objects capture that the distribution of plausible geometries is more spread out than near data points.

Finally, one can check that μ is invariant to any transformation $g \in C^1(\mathbb{R})$ of the field φ that preserves the level-set, *i.e.*, $g^{-1}(u) = \{u\}$. While $g \circ \varphi$ is not Gaussian in general, it guarantees that this definition is intrinsic to the distribution of the random surfaces themselves and not attached to a given representation.

7 Probabilistic calculus over GPIS

In this section, we use the Kac-Rice formula to transfer the expected values of standard integrals on random surfaces (inner products of functions, Dirichlet energy,...) to the volumetric Cartesian domain Ω, to define differential operators over general GPIS.

7.1 Functional spaces and inner products

As we integrate functions both on Ω and their restriction to level sets, we first characterize the integrands such that all integrals are well defined. By considering fields in $H_0^1(\Omega, dx) = \{u \in L^2(\Omega, dx) \text{ s.t. } \nabla u \in L^2(\Omega, dx), u|_{\partial\Omega} = 0\}$, the Trace theorem [Evans 2022] ensures that such functions restricted to smooth level sets $\varphi^{-1}(u)$ are in $L^2(\varphi^{-1}(u), dA)$.

In this space, we can use the Kac-Rice formula (with $h(x, y) = f(x)g(x)$, for $f, g \in H_0^1(\Omega)$) to express how the expected L^2 inner product between functions on the random manifold is transferred on Ω:

$$\mathbb{E}\left[\langle f, g\rangle_{L^2(\varphi^{-1}(u), dA)}\right] = \int_\Omega f(x)g(x)d\mu(x). \tag{12}$$

Hence, the expected L^2 inner product on $\varphi^{-1}(u)$ induces integration with respect to the surface density measure μ we introduced in Eq. 11. The inner product $\langle\cdot,\cdot\rangle_{L^2(\Omega,\mu)}$ is then the natural inner product of interest, since it captures information at a point x proportionally to the likelihood of the presence of the surface at x (Sec. 6).

7.2 The Kac-Rice Laplace-Beltrami operator

On a manifold implicitly expressed as $\mathcal{M} = \varphi^{-1}(u)$, the Dirichlet energy of a scalar field defined in the ambient space restricted to $\mathcal{M}$ can be written as:

$$E_D(f) = \int_{\varphi^{-1}(u)} \|\Pi_{\nabla\varphi^\perp}\nabla f\|^2 dA, \tag{13}$$

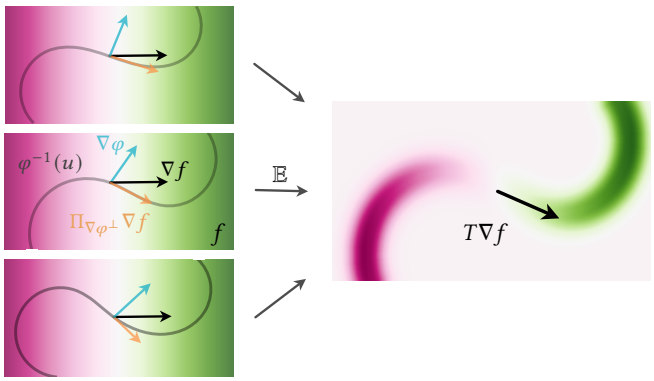


Fig. 4. At each point $x \in \Omega$, the expected Dirichlet energy of an ambient scalar field f over a GPIS isosurface $\varphi^{-1}(u)$, captures the tangential variations of f, *i.e.* $\Pi_{\nabla\varphi^\perp}\nabla f$, along the distribution of possible tangent planes, conditioned under the event that the surface runs through x. This distribution of tangent planes is encoded in the tensor field T that averages the orthogonal projectors weighted by the likelihood that the shape lies in x, *i.e.* $\|\nabla\varphi\|p_{\varphi(x)}(u)$. On the figure on the right, the field μf is plotted to illustrate what part of f is actually captured by the Dirichlet energy.

where ∇ is the gradient operator on $\mathbb{R}^d$ and $\Pi_{\nabla\varphi^\perp}$ is the orthogonal projector onto the tangent plane of $\varphi^{-1}(u)$. Now, over random surfaces, for a fixed scalar field f, we consider its expected value:

$$\mathbb{E}[E_D(f)] = \mathbb{E}\left[\int_{\varphi^{-1}(u)} \|\Pi_{\nabla\varphi^\perp}\nabla f\|^2 dA\right].$$

Our key idea is to use the Kac-Rice formula to incorporate the distribution of the gradients into the orthogonal projectors onto the tangent planes. By setting $h(x, y) = \|\Pi_{y^\perp}\nabla f(x)\|^2$ and $\psi(x) = \nabla\varphi(x)$ in Eq. 9,the expected Dirichlet energy becomes:

$$\mathbb{E}[E_D(f)] = \int_\Omega \mathbb{E}_{\varphi=u}\left[\|\Pi_{\nabla\varphi^\perp}\nabla f\|^2\|\nabla\varphi(x)\|\right] p_{\varphi(x)}(u)dx. \tag{14}$$

From this expression, we define the *Kac-Rice Laplace-Beltrami operator* Δ_φ as the gradient of the above energy with respect to f, and the $L^2(\Omega, \mu)$ inner product, i.e.:

$$\boxed{\Delta_\varphi f(x) := \frac{1}{\mu(x)}\nabla\cdot(T(x)\nabla f(x)),} \tag{15}$$

with

$$\boxed{T(x) := p_{\varphi(x)}(u)\mathbb{E}_{\varphi=u}\left[\|\nabla\varphi\|\Pi_{\nabla\varphi^\perp}\right].} \tag{16}$$

(see App. C for details). The matrix field T is critical in the following derivations, as it encodes the distribution of the tangent planes everywhere in space (see Fig. 4). One can check that Δ_φ is the quadratic form associated with the expected Dirichlet energy, i.e.: $\langle f, \Delta_\varphi f\rangle_{L^2(\Omega,\mu)} = \mathbb{E}[E_D(f)]$.

This differential operator, defined on the ambient space, can be seen as a probabilistic generalization of the Laplacian on implicit surfaces defined by Bertalmio et al. [2001], which reads, on $\varphi^{-1}(u)$:

$$\frac{1}{\|\nabla\varphi\|}\nabla\cdot(\|\nabla\varphi\|\Pi_{\nabla\varphi^\perp}\nabla f). \tag{17}$$

As this operator is derived as the gradient of the Dirichlet energy in the deterministic case Eq. 13, we recover this operator (in the weak

sense) when the variance σ tends to zero. We provide the derivation of Δ_φ and its weak convergence to Eq. 17 in App. C.

Interpretation. The Laplacian Δ_φ can be interpreted as a standard anisotropic diffusion operator in $\mathbb{R}^d$ that allows heat (and more generally, information) to spread in the directions given by the matrix T, as illustrated in Fig. 3 . We recall that T is given by $T(x) = p_{\varphi(x)}\mathbb{E}_{\varphi=u}\left[\|\nabla\varphi\|\Pi_{\nabla\varphi^\perp}\right]$. Hence, when the variance is small, it boils down to a projection onto the tangent plane of the mean surface, with the consequence that heat spreads only tangentially. When the variance is high, the diffusion becomes more uniform in space. The speed at which the scalar field spreads is roughly proportional to μ itself: the less likely the shape lies in x, the slower information spreads.

7.3 First order operators

First, the divergence of a vector field V is defined as the variation of the volume of an object when advected infinitesimally by V, i.e., as the Lie derivative of the volume form [Bogdanskii and Shram 2022]:

$$\mathrm{div}(V) = \mathcal{L}_V(\mathrm{vol}).$$

With our non-uniform volume form, $\mathrm{vol} = d\mu = \mu\mathrm{dx}$, the divergence operator is given by (see App. D):

$$\boxed{\mathrm{div}_\varphi(V)(x) \coloneqq \frac{1}{\mu(x)}\nabla\cdot(\mu(x)V(x))\,.} \tag{18}$$

As a consistency check, we verify in App. D that if a set A_t is advected by a vector field V, then $d\mu(A_t)$ is preserved if and only if $\mathrm{div}_\varphi(V) = 0$.

Since (scalar) Laplacians are formally defined as $\Delta = \mathrm{div}\circ\mathrm{grad}$, Eq. 15 and Eq. 18 lead to:

$$\Delta_\varphi f(x) = \frac{1}{\mu(x)}\nabla\cdot(T(x)\nabla f(x)) = \frac{1}{\mu(x)}\nabla\cdot\left(\mu(x)\frac{T(x)}{\mu(x)}\nabla f(x)\right).$$

By identification, we define

$$\boxed{\mathrm{grad}_\varphi f(x) \coloneqq \frac{T(x)}{\mu(x)}\nabla f(x)\,.} \tag{19}$$

In App. E we show that grad_φ and div_φ operators are dual with respect to an inner product that converges weakly to the metric of the mean surface when the variance tends to 0.

Interpretation. From Eq. 11 and the definition of $T(x)$, the scaling by $p_{\varphi(x)}(u)$ cancels out in the gradient operator expression, leading to:

$$\mathrm{grad}_\varphi f(x) = \frac{T(x)}{\mu(x)}\nabla f(x) = \frac{\mathbb{E}_{\varphi=u}\left[\|\nabla\varphi\|\Pi_{\nabla\varphi^\perp}\right]}{\mathbb{E}_{\varphi=u}\left[\|\nabla\varphi\|\right]}\nabla f(x)\,.$$

Hence, the linear operator applied to the ambient gradient is a weighted average of the orthogonal projectors over the distribution of tangent planes. The strength of the projection is reduced with uncertainty and we recover the projected gradient in the deterministic case.

7.4 Domain restriction and boundary conditions

Since μ decreases exponentially with the distance to the mean surface, we restrict computations to a narrow band to avoid conditioning problems for Δ_φ:

$$\Omega_{\varphi,\eta} \coloneqq \Omega \cap \mu^{-1}\left([\eta,+\infty)\right).$$

Over $\Omega_{\varphi,\eta}$, various homogeneous boundary conditions could be considered. When solving Poisson like problems, we choose to impose homogeneous Dirichlet conditions:

$$\begin{cases}\Delta_\varphi u(x) = f(x), & x\in\Omega_{\varphi,\eta}\\ u(x) = 0, & x\in\partial\Omega_{\varphi,\eta}\,.\end{cases} \tag{20}$$

From a theoretical point of view, restriction to a compact domain is necessary, since, with $\sigma > 0$, one can bound the eigenvalues of T uniformly from above and below. In other words, there exist $m, M > 0$, such that

$$\forall x\in\Omega_{\varphi,\epsilon}, m \le \lambda^x_{\min} \le \lambda^x_{\max} \le M\,, \tag{21}$$

where $\lambda^x_{\min}, \lambda^x_{\max}$ are the smallest and largest eigenvalues of $T(x)$ respectively. This condition, from the standard elliptic PDE theory [Evans 2022], ensures the existence and smoothness of the solutions of Eq. 20. From this perspective, Δ_φ is a regularized version of the Laplacian on implicit surfaces Eq. 17: as the latter involves $\Pi_{\nabla\varphi^\perp}$ which is always rank deficient, Eq. 17 does not satisfy the ellipticity condition Eq. 21.

8 Discretization

Our differential operators rely on standard gradient and divergence operators of $\Omega\subset\mathbb{R}^d$. Among alternative discretization schemes, for the sake of simplicity, we use finite differences on a regular grid with spacing h:

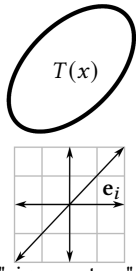
$$\Omega^h_{\varphi,\eta} = h\mathbb{Z}^d\cap\Omega_{\varphi,\eta}\,.$$

We denote $N = |\Omega^h_{\varphi,\eta}|$ the number of variables in the domain. The divergence and gradient operators are then discretized as:

$$D_\varphi \coloneqq M_\varphi^{-1}G^\top\left(M_\varphi\otimes I_d\right)\qquad G_\varphi \coloneqq \mathrm{BlockDiag}\left(\frac{T(x_i)}{\mu(x_i)}\right)G \tag{22}$$

where $M_\varphi \coloneqq \mathrm{Diag}\left(\mu(x_i)\right)$, for $x_i\in\Omega^h_{\varphi,\eta}$, $\otimes$ is the Kronecker product, and $G\in\mathbb{R}^{dN\times N}$ is the standard forward finite differences scheme for the gradient.

8.1 Adaptive stencil finite difference operator



While directly composing $D_\varphi G_\varphi$ provides a valid discretization of Δ_φ, in the sense that it converges when grid resolution increases, its accuracy suffers when the anisotropy of $T(x)$ is high. Dedicated finite difference schemes exist for such highly anisotropic diffusion operators. In particular, we use the adaptive stencil scheme developed by Fehrenbach and Mirebeau [2014]. This scheme uses tools from discrete geometry that give an analog of the eigen decomposition of a symmetric matrix where the "eigenvectors" are constrained to have integer coordinates only (see inset). Any SPD matrix T, can be decomposed in the following way: there exists

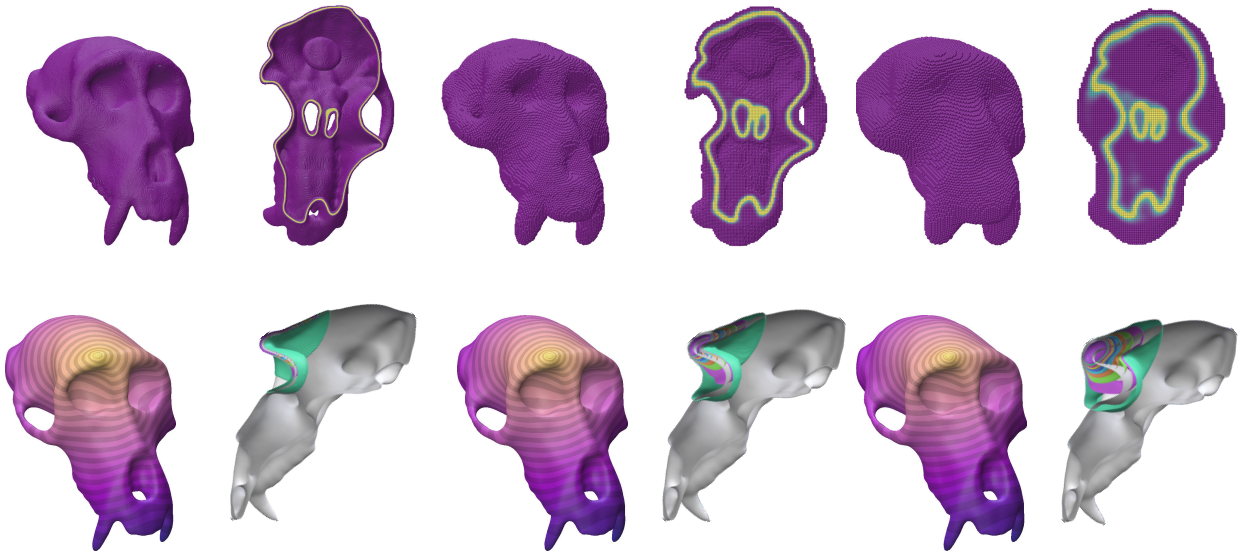


Fig. 5. We use the width of the adaptive stencil (Eq. 25) to automatically adapt the resolution of the grid $h^\star$ to the anisotropy of the field T (top row). We also use a BFS approach to avoid querying the GPIS in the entire domain by starting the exploration from input points. In the bottow row, we illustrate the Green's function (logscale) obtained from the associated Laplacian operator (Eq. 24), *i.e.,* we solve $L_\varphi f_{\text{green}} = M_\varphi \delta_y$ for some Dirac function at $y \in \Omega^{h^\star}_{\varphi,\eta}$. As our operators are volumetric, we display values on the mean surface and visualize some isosurfaces of f_{green}, highlighting that the anisotropy of the diffusion is affected by the σ value. Timings are given in Sec. 9

a family of coefficients (unique for $d \leq 4$) $\{(\lambda_i, \mathbf{e}_i)\}_{1\leq i\leq d'} \in \mathbb{R}^+ \times \mathbb{Z}^d$, where $d' \leq \frac{d(d+1)}{2}$, such that:

$$T = \sum_{1\leq i\leq d'} \lambda_i \mathbf{e}_i \mathbf{e}_i^\top. \tag{23}$$

This decomposition is efficiently computed in practice using Selling's algorithm [Fehrenbach and Mirebeau 2014] (see supplementary material Sec. 6). We construct the anisotropic diffusion operator defined by the tensor field $T(x)$ (with associated $\{(\lambda_i(x), \mathbf{e}_i(x))\}$) as:

$$L_\varphi u_h = \frac{1}{h} \sum_{\mathbf{e}\in\mathbb{Z}^d, s=\pm 1} \frac{\lambda^{\mathbf{e}}(x) + \lambda^{\mathbf{e}}(x + h\, s\, \mathbf{e})}{2} \frac{u_h(x + h\, s\, \mathbf{e}) - u_h(x)}{h}, \tag{24}$$

where $\lambda^{\mathbf{e}}(x)$ equals $\lambda_i(x)$ if $\mathbf{e} = \mathbf{e}_i(x)$ for any $\mathbf{e}_i(x)$ in the decomposition in Eq. 23 of $T(x)$, 0 if not. Intuitively, the stencil now looks for information aligned with the local anisotropy of the tensor $T(x)$. The term $\frac{1}{2}\left(\lambda^{\mathbf{e}}(x) + \lambda^{\mathbf{e}}(x + h\, s\, \mathbf{e})\right)$ ensures that the resulting matrix is symmetric.

The discrete Δ_φ is then finally $M_\varphi^{-1} L_\varphi$ with a sparse symmetric positive definite matrix L_φ, and a (diagonal) mass matrix M_φ. Compared to $D_\varphi G_\varphi$, $M_\varphi^{-1} L_\varphi$ is guaranteed to be non-negative, ensuring compatibility with fast SPD solvers. In Fig. 6-(top) we confirm the numerical stability of $M_\varphi^{-1} L_\varphi$ compared to $D_\varphi G_\varphi$ when the anisotropy is high (small σ value).

8.2 Adaptive resolution and efficient narrow band computation

Given an input set of n samples $\{p_i, \omega_i\}$, we can use the geometry of the stencils of $T(p_i)$ to construct a heuristic adapting the grid resolution h to a GPIS. Indeed, the width of the stencil encodes the speed at which information propagates in the numerical schemes, a larger stencil should be compensated with a finer resolution. In our experiments, we have considered:

$$h^* \coloneqq \frac{\lambda}{\max_{i\in\{1..n\}} \max_{1\leq j\leq d'} \|\mathbf{e}_j(p_i)\|_1}, \tag{25}$$

where the $\mathbf{e}_j(p_i)$ are obtained from the Selling's algorithm on $T(p_i)$ (see Sec. 8.1), and λ is a scaling parameter (*e.g.;* the average edge length for triangle soup inputs). As illustrated in Fig. 7, for fixed η, changing the uncertainty value σ on the input as a direct impact on the width of the stencil, and thus the h^* resolution. The higher the uncertainty, the coarser the grid.

Similarly to Closest-Point Methods [King et al. 2024], to efficiently build the narrow band and avoid querying φ on the entire domain, we start from the points in the input and use a breadth-first search (BFS) strategy on Ω_h to extend the narrow band only to the points x such that $\mu(x) > \eta$. As demonstrated in Fig. 7, this approach is an order of magnitude faster than a naive global query for fine resolutions in our context.

8.3 Expectations of non-linear functions of Gaussian variables

Many terms involve the expected value of a non-linear function of the Gaussian vector $\nabla\varphi(x)|\varphi(x) = u$ (see for example Eq. 11). For a general non-linear function f, no closed formula exists to estimate:

$$\mathbb{E}_{\varphi=u}\left[f(\nabla\varphi(x))\right].$$

To evaluate such expressions, we use a Quasi Monte-Carlo (QMC) approach [Soboĺ 1990]: we fix a given set of S low discrepancy

samples S_i, of $\mathcal{N}(0, I_d)$ and evaluate the expectation using:

$$\mathbb{E}_{\varphi=u}\left[f(\nabla\varphi(x))\right] \approx \frac{1}{S}\sum_{i=1}^{S} f(C_x S_i + \mu_x),$$

where $\nabla\varphi(x)|\varphi(x) = u \sim \mathcal{N}(\mu_x, \Sigma_x)$ is computed using Eq. 4 and Eq. 7 and C_x is the Cholesky factor of $\Sigma_x = C_x^\top C_x$. In all our experiments, we set $S = 512$ (see supplementary for further discussions).

9 Numerics & Implementation details

All experiments were performed on an Intel(R) Xeon(R) W-2245 3.90Hz, with 8 cores. Code is provided in supplementary materials. All σ values used in our experiments are expressed as normalized factors relative to the length scale of the input model, *i.e.*, mean edge length of the input triangle soup. The s value used for the regularized winding number Eq. 6 is set to 0.4.

In Alg. 1, we provide a complete pipeline for the Green's function example of Fig. 5 and Fig. 8. Other geometry processing tasks discussed in this section can be obtained adapting the input and last lines. Note that in practice we store the $\mu(x_i)$ that are computed during the narrow band construction. We recall that a query of the joint value/gradient of a GPIS model (line 2) yields a $(d+1)-$dimensional Gaussian vector with mean in $\mathbb{R}^{d+1}$ and covariance matrix $\mathbb{R}^{(d+1)\times(d+1)}$. When conditioned on the isovalue (line 3, Eq. 4), we use the block decomposition of the covariance matrix (*e.g.*, $\Sigma_{\varphi,\nabla\varphi} \in \mathbb{R}^{d\times 1}$, $\Sigma_{\nabla\varphi,\nabla\varphi} \in \mathbb{R}^{d\times d}$).

Algorithm 1: Overall pipeline (Green's function)

Data: A GPIS model: $x \mapsto (\nabla\varphi(x), \varphi(x))$, resolution $h > 0$, narrow-band width $\eta > 0$, iso-value $u \in \mathbb{R}$, Green source $x_s \in \mathbb{R}^d$

1 **Function** EvalMu(x):
  // Get mean and covariance for the value/gradient Gaussian pair (e.g., Barnes-Hut)
2 $(\nabla\varphi(x), \varphi(x)) \leftarrow$ GPIS(x) // Sec. 9.1, $O(\log n)$
  // Setting $X = \nabla\varphi(x)$, $Y = \varphi(x)$, $c = u$ in Eq. 4
3 $\nabla\varphi_{\text{cond}} \leftarrow \nabla\varphi(x)|\varphi(x) = u$ // Sec. 9.1, $O(d)$
  // Approximate Eq. 11 with Quasi-Monte-Carlo
4 **return** MuQMC $\left(\nabla\varphi_{\text{cond}}, p_{\varphi(x)}(u)\right)$ // Sec. 8.3, $O(S)$
5 **Function** EvalT(x):
  // Same as EvalMu but for Eq. 16
// Construct the narrow-band variables ($N = |\Omega^h_{\varphi,\eta}|$)
6 $\Omega^h_{\varphi,\eta} \leftarrow \{x_i \in h\mathbb{Z}^d, \text{ s.t. EvalMu}(x_i) > \eta\}$ // Sec. 8.2
// Evaluate the quantities
7 $(\mu(x_i), T(x_i)) \leftarrow \{(\text{EvalMu}(x_i), \text{EvalT}(x_i)), x_i \in \Omega^h_{\varphi,\eta}\}$ // $O(N(\log n + d + S))$
// Build operators using the adaptive stencil
8 $M_\varphi \leftarrow Sec.\ 8$, $L_\varphi \leftarrow Sec.\ 8.1$ // $O(N)$
// Geometry processing task (e.g., Green's function)
9 $f_{\text{green}} \leftarrow \text{SPDSolve}(L_\varphi, M_\varphi \delta_y)$ // closest $y \in \Omega^h_{\varphi,\eta}$ to x_s

Alg. 1 also summarizes the main computational costs as functions of the input (n and d) or method parameters (S). As discussed in App. **??**, we do not include the computational cost of Selling's algorithm. For representative timings on the Green's problem of Fig. 5 ($n = 7000$ input position/normal pairs): for $\sigma \in \{0.5, 0.5, 1.0\}$, the computation of $h^\star$ in $O(n(\log n + S + d)$ (Eq. 25)), the narrow band construction and the operators (lines 6-8) are obtained in $\{21, 4.5, 3.2\}$ seconds respectively, for $\Omega^h_{\varphi,\eta}$ bounding boxes in $\{299 \times 477 \times 282, 106 \times 165 \times 101, 83 \times 128 \times 81\}$. Solving the Green's problem using a conjugate gradient (line 9) is obtained in $\{21, 6.7, 6.2\}$ seconds.

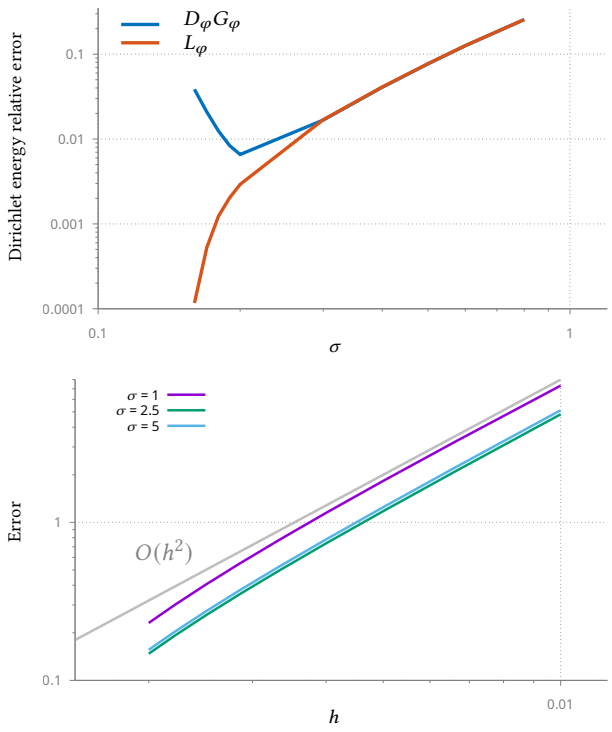


Fig. 6. Convergence of the Dirichlet energy (top): on the spot GPIS model of Fig. 8, we plot the relative error between the expected Dirichlet energies $\langle f, D_\varphi G_\varphi f\rangle$ or $\langle f, L_\varphi f\rangle$, and the cotan Laplacian Dirichlet energy $\langle f|_{\mathcal{M}}, L_{\text{cotan}} f|_{\mathcal{M}}\rangle$ on the mean surface mesh (for $f(x) = \sin(10x_1)\sin(10x_2)\sin(10x_3)$). Using the same setting, convergence of the discrete operator $M_\varphi^{-1}L_\varphi$ when resolution diminishes, for various σ values (bottom). Computed as $\|M_\varphi^{-1}L_\varphi f_h - M_\varphi^{-1}L_\varphi f_{h^\star}\|_{L^2(\Omega,\mu)}$, were $h^\star = 10^{-3}$.

9.1 Accelerating scheme

A key feature of the GPIS model described in Sec. 4 based on the regularized winding number potential [Chen et al. 2024] is its compatibility with acceleration schemes *à-la* Barnes-Hut [Barnes and Hut 1986], to approximate Eq. 7 with a $O(n\log(n))$ pre-computation and $O(\log(n))$ for a single evaluation of $\varphi(x)$ and $\nabla\varphi(x)$. The code is provided in supplementary materials. Compared to a classical implementation of the regularized winding number potential, each node in the hierarchy accumulates the covariance matrices of the points *below* it. We refer to the supplementary for details on the Barnes-Hut scheme. Note that one could easily adapt Madan et al. [2025] for a fast GPU implementation.

9.2 Numerical analysis

Dirichlet energy convergence. We first numerically validate in Fig. 6-(top) that the closed form of the expected energy we derive from the Kac-Rice formula converges to the Dirichlet energy computed on the mean surface when the variance σ tends to 0 when using Eq. 24, while the naive discretization fails at high uncertainty.

Convergence of the discretization scheme. In Fig. 6-(bottom), we validate that our discretization scheme for $M_\varphi^{-1}L_\varphi$ converges as a function of the grid resolution h and for various σ. We validate the order 2 convergence of the scheme of Fehrenbach and Mirebeau [2014].

Number of variables in the narrow band. As all computations are performed on the narrow band $\Omega^h_{\varphi,\eta}$, we study the number of variables N in it in Fig. 7-(top) depending on the variance, controlled by σ, and the density threshold η.

Speed-up provided by the adaptive narrow band exploration. In Fig. 7-(bottom) we study the impact of the BFS exploration strategy described Sec. 8.2 compared to a naive volumetric query (note that the computation includes both $\mu(x_i)$ and $T(x_i)$). We show that an order of magnitude speedup can be obtained using the BFS approach.

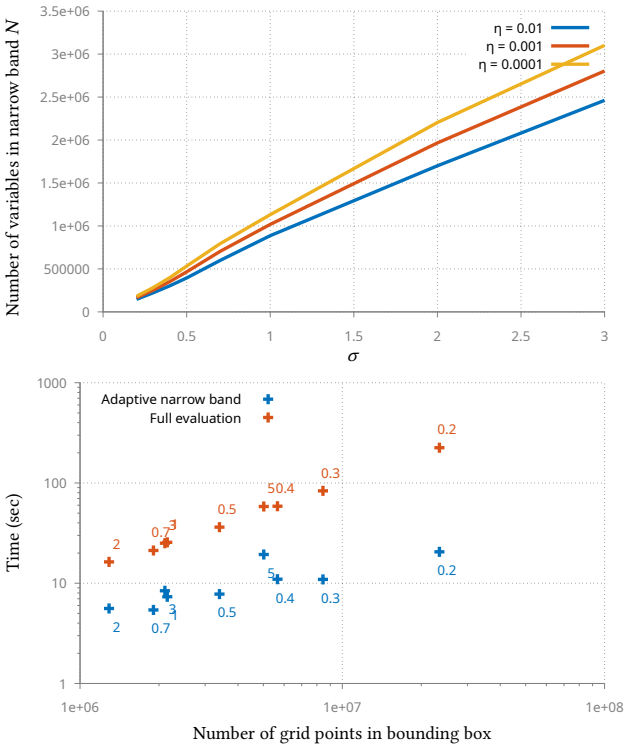


Fig. 7. Adaptive grid performances: on the macaca GPIS model of Fig. 5, we first plot the number of variables in the narrow band (N) as a function of σ for various η values (top). For various σ values, on a grid resolution given by Eq. 25, we compare the efficiency of the BFS approach ($O(N)$ queries) to the naive one ($O\left(h^{\star -3}\right)$ queries) by plotting the respective compute time of $\{\mu(x_i)\}$ (in sec) as a function of the number of points in the model bounding box.

10 Geometry processing on GPIS

In this section, we review several standard geometry processing applications that can be extended to GPIS using our framework. We show in Fig. 8 that considering the uncertainty of the GPIS has a deep impact on the downstream applications that is ignored by the calculus on the mean surface (*e.g.*, PSR output).

10.1 Spectral analysis

Computing the eigenfunctions of a manifold is a widely used operation in geometry processing as they are intrinsic invariants of the shape and generalize Fourier analysis to curved domains [Lévy 2006]. As for standard Laplacians, we compute them by solving the following generalized eigenvalue problem:

$$\begin{cases} L_\varphi u = \lambda M_\varphi u, & x \in \Omega_{\varphi,\epsilon} \\ u = 0, & x \in \partial\Omega_{\varphi,\epsilon}. \end{cases}$$

In Fig. 8 (third row), we illustrate the impact of the uncertainty parameter σ to the eigenvalues/eigenvectors decomposition. Additional results are provided in the supplemental.

Application: Uncertainty aware kernels for scalar inference on GPIS. Our approach provides deterministic operators, but the stochastic nature of Gaussian process regression can be retrieved for GPIS using an extension of the Mattérn family of kernels to manifolds by using the eigenfunctions $\{u_i\}$ of the Laplacian [Borovitskiy et al. 2020]:

$$k_{(\cdot)}(x,x') = \frac{1}{C_{(\cdot)}} \sum_{i=0}^{\infty} \rho_{(\cdot)}(\lambda_i) u_i(x) u_i(x'), \tag{26}$$

where $\rho_\nu(\lambda) = \left(\frac{2\nu}{\kappa^2} + \lambda\right)^{-\nu-\frac{d}{2}}$ for $\nu < \infty$ and $\rho_\infty(\lambda) = e^{-\frac{\kappa^2}{2}\lambda}$, $C_{(\cdot)} = \sum_i \rho_{(\cdot)}(\lambda_i)$, for $\kappa > 0$. In practice, the sums are truncated (*e.g.*, to the first 200 eigenvectors in Fig. 9). We show that simply plugging the $L^2(\Omega,\mu)$-orthonormal eigenfunctions of the Δ_φ operators incorporates the *geometric* uncertainty into the regression: the variance of the extrapolated field is small near the data points, grows with the distance and is higher where the shape itself is uncertain.

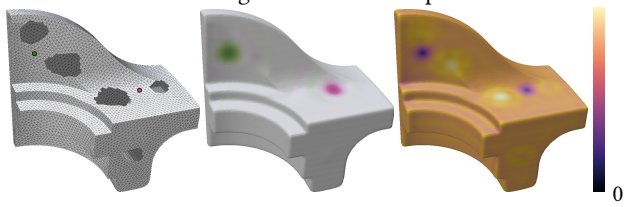


Fig. 9. Gaussian process scalar field inference using the kernel derived from the eigenfunctions of the Δ_φ operator. Regression is performed by building a Gaussian process using the kernel Eq. 26 from scalar data (left) as in Sec. 3.2. The mean field of the regression is plotted in the middle and the variance field on the right. One can observe that the variance of the prediction is lower near the data points and is highest at locations where the uncertainty of the underlying GPIS is maximum.

10.2 Geodesic distance

A meaningful notion of geodesic distance can be obtained on GPIS with a kinetic energy formulation. The kinetic energy $\mathcal{H}$ –or

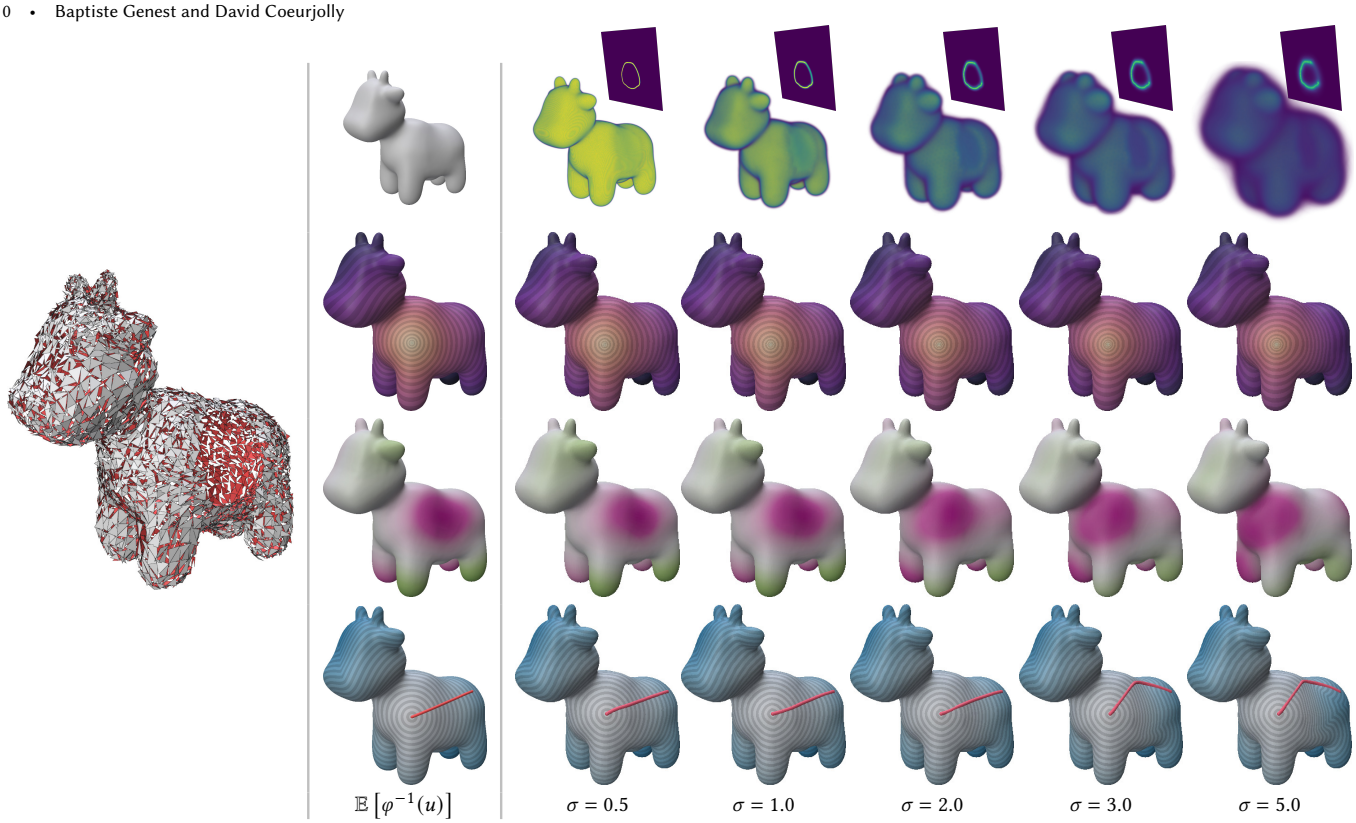


Fig. 8. The differential operators we propose depend continuously on the level of uncertainty of the shape. As the σ parameter on the input increases (left, see App. B), the region without data, on the side of the input, is more and more uncertain (density μ top row vanishing in this region). While we recover results performed on the mean surface directly (first column) at low noise level, at high uncertainty the results evolve where data is missing: the Green's function (second row) decreases faster in it (log scale colormap), a nodal domain of the 9-nth eigenfunction of $M_\varphi^{-1}L_\varphi$ (Sec. 10.1), drifts away from it (third row), and the geodesic distance (Sec. 10.2) progressively considers it as a hole (fourth row). The geodesic from a fixed point is also shown. Remarkably, we observe that the shortest path changes abruptly above a certain uncertainty level, i.e., when the geodesic front that gets around the hole is faster than the one that goes through it.

Hamiltonian without potential energy– of an ambient velocity field $V : \Omega \to \mathbb{R}^d$ over an embedded manifold $\mathcal{M} = \varphi^{-1}(u)$ is defined as

$$\mathcal{H}(V) \coloneqq \int_{\varphi^{-1}(u)} \|\Pi_{\nabla\varphi(x)^\top} V(x)\|^2 dA\,. \tag{27}$$

Similar to Sec. 7.2 for the Dirichlet energy, we can use the Kac-Rice formula to obtain the pointwise expression $\tilde{\mathcal{H}}$ of the expected kinetic energy.

$$\tilde{\mathcal{H}}(V(x), x) = \|V(x)\|^2_{T(x)}\,,$$

with $T(x)$ defined in Eq. 16. We consider the geodesic distance U from a source x_0 as the scalar map whose gradient has unitary kinetic energy, leading to the generalized Eikonal equation:

$$\|\nabla U(x)\|_{T(x)} = 1\,, \tag{28}$$

where $\|\nabla U(x)\|_{T(x)} = \sqrt{\nabla U(x) \cdot T(x) \nabla U(x)}$, for all $x \in \Omega$, except for source points x_s where $U(x_s) = 0$.

Numerically, we solve Eq. 28 using a Hamiltonian Fast Marching method also based on the integer decomposition of T (Eq. 23), as proposed by Mirebeau and Portegies [2019], in $O(N \log(N))$ time.

Once field U is computed, the geodesics can be extracted by integrating ∇U. The influence of the uncertainty on the obtained distance is shown in Fig. 8. The shortest path between two points can drastically change with the uncertainty, even though the mean surfaces are the same. Indeed, Eq. 28 imposes a constant tangential –weighted by likelihood– velocity at each point x. Hence the geodesic distance must increase faster in zones with high uncertainty to match the velocity constraint.

10.3 Hodge decomposition

In order to validate the proposed gradient and divergence operators, we show that they generalize the behavior of the Hodge decomposition of vector fields. The usual projection on the set of divergence-free vector fields can trivially be extended to GPIS using the proposed operators:

$$V_{\mathrm{div}_\mu\text{-free}} = (I - \mathrm{grad}_\varphi \Delta_\varphi^{-1} \mathrm{div}_\varphi)\, V\,. \tag{29}$$

We also refer to App. E for a discussion about the inner product implicitly used here. In the discrete setting, divergence and gradient

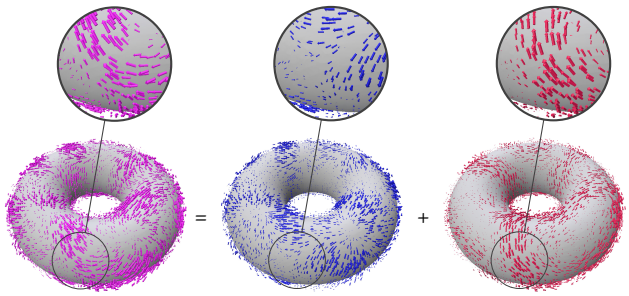


Fig. 10. From div_φ and grad_φ operators, we decompose a (volumetric) vector field (left) into a sum of a vector field in $\mathrm{Im}(\mathrm{grad}_\varphi)$ (middle) and a μ-div-free vector field (right).

operators are given by D_φ and G_φ, and exact projections onto the kernel of D_φ is ensured using the $D_\phi G_\phi$ composition for Δ_φ.

As illustrated in Fig. 10, we can recover the components of a vector field that is a sum of a purely μ-divergence-free one and a μ-curl-free one (*i.e.* in $\mathrm{Im}(\mathrm{grad}_\varphi)$). In the standard Riemannian setting, the Hodge decomposition has a deep connection with the topology of the manifold with the harmonic component of a vector field. Studying how this property extends to the volumetric nature of a GPIS is an interesting axis for future research.

10.4 Smoothed distance to a GPIS

The distance to a surface can be expressed as $d(x, \mathcal{M}) = \min_{p\in\mathcal{M}} \|x - p\|$. A smoothed formulation has been proposed by Sethi et al. [2012] and extended by Madan and Levin [2022] following a standard LogSumExp approach:

$$\tilde{d}_\alpha(x, \mathcal{M}) = -\frac{1}{\alpha} \log\left(\int_{\mathcal{M}} \exp\left(-\alpha\|x - p\|\right) dA(p)\right), \qquad (30)$$

with a smoothness parameter $\alpha > 0$. We can once again extend this formulation to a GPIS using the Kac-Rice formula. As the integral involved in Eq. 30 does not depend on $\nabla\varphi$, it simply reduces to integration with respect to μ:

$$\tilde{d}_\alpha(x, \varphi) = -\frac{1}{\alpha} \log\left(\int_{\Omega} \exp\left(-\alpha\|x - p\|\right) d\mu(p)\right). \qquad (31)$$

While the mean surface of the GPIS has no boundaries, we show in Fig. 11 that the geometry of the level-sets of this smoothed distance may greatly vary with uncertainty for various downstream distance field applications (*e.g.,* sphere tracing).

11 Conclusion

In this paper, we propose a general framework to define a meaningful notion of calculus on Gaussian process implicit surfaces. As everything follows from the *exact* expected value of the Dirichlet energy on GPIS, we argue that the operators obtained are a natural generalization of the differential toolbox commonly used in Geometry Processing.

Generally speaking, the Kac-Rice formula appears as a very useful tool when dealing with GPIS, which deserves more attention from the geometry processing community and beyond. We believe that our derivation opens the door to many interesting axes of research for the processing of such stochastic surfaces.

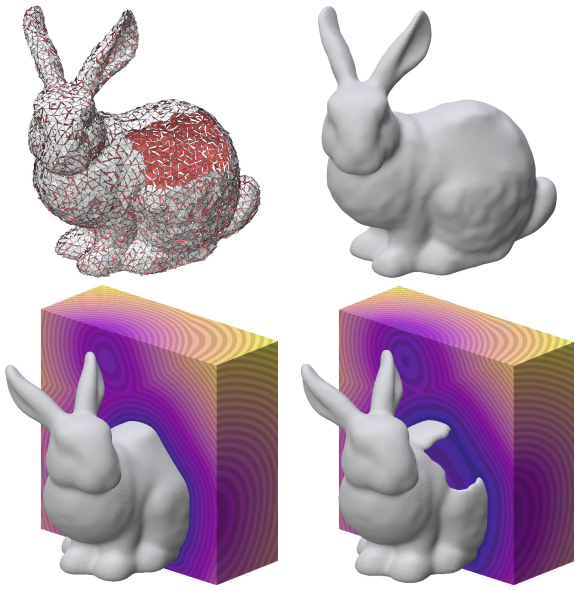


Fig. 11. From an input triangle soup (top left), despite having the same mean surface of the GPIS model (top-right), various uncertainty parameters ($\sigma = 1.0$ bottom left, $\sigma = 5.0$ bottom right), lead to different level-sets (here 0.08) of the smoothed distance, allowing an implicit representation of the input shape.with various geometrical charasteristics (hollow bunny as illustrated here).

While the simple structure of the Dirichlet energy allows to express everything in closed form, we believe that our approach could easily be extended to the very wide range of PDEs that can be expressed as gradient flows of geometric functionals, such as the Willmore flow or vector operators.

Finally, while finite differences are trivial to implement, they impose a strong structure on the discretization of the domain. Exploring other numerical schemes for discretizing the operators might be beneficial. For example, a finite volume approach on the Voronoi diagram of samples of μ could provide more flexibility on the number of variables and more easily fit highly curved geometries.

Acknowledgments

This work is partially supported by the French National Research Agency within the StableProxies project (ANR-22-CE46-0006) and donations from Adobe Inc. Figures were made using the Polyscope library. Bunny model courtesy of the Stanford 3D scanning repository, Spot model from Keenan Crane, Nefertiti model from Nora Al-Badri and Jan Nikolai Nelles and the Beetle and Fandisk models from the Aim@Shape repository.

References

Jean-Marc Azäis and Mario Wschebor. 2009. *Level Sets and Extrema of Random Processes and Fields* (1 ed.). Wiley. doi:10.1002/9780470434642

Josh Barnes and Piet Hut. 1986. A hierarchical O (N log N) force-calculation algorithm. *nature* 324, 6096 (1986), 446–449.

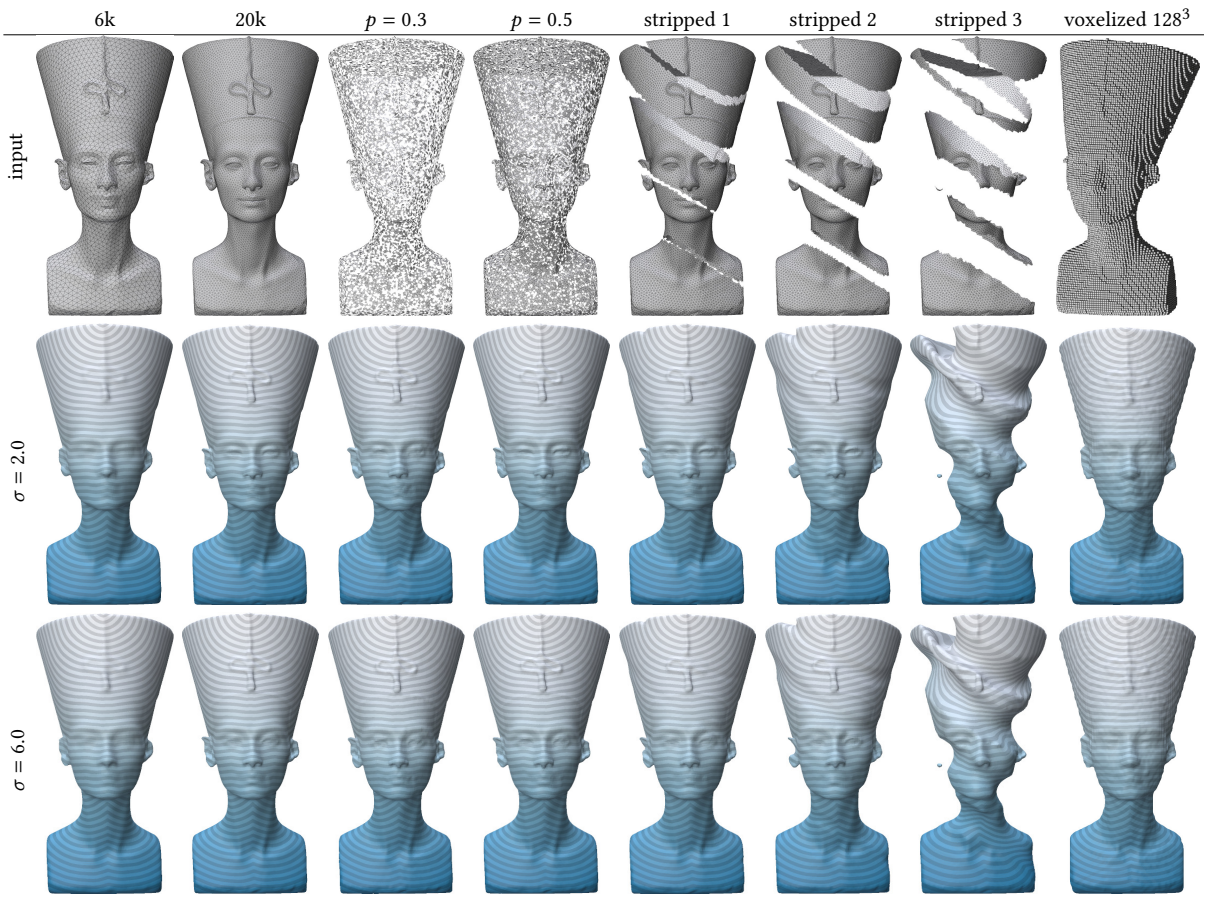


Fig. 12. We analyze how the quality of the input data used in our GPIS model (Sec. 4, App. B) affects the rest of the framework, illustrated here through geodesic distance computation. We vary geometric resolution (first two columns), subsampling levels (next two), and the presence of large missing regions and voxelized output (last four), each shown for two confidence values σ. Even when the mean surface is heavily degraded—and may break into multiple connected components (e.g., stripped 3)—the resulting geodesic distance field remains stable.

Marcelo Bertalmio, Guillermo Sapiro, Li-Tien Cheng, and Stanley Osher. 2001. Variational problems and PDEs on implicit surfaces. In *Proceedings IEEE Workshop on Variational and Level Set Methods in Computer Vision*. 186–193. doi:10.1109/VLSM.2001.938899

Corinne Berzin, Alain Latour, and José León. 2022. Kac-Rice Formula: A contemporary overview of the main results and applications. *arXiv preprint arXiv:2205.08742* (2022).

Yuri Bogdanskii and Vladyslav Shram. 2022. Divergence of multivector fields on infinite-dimensional manifolds. *Ukrains' kyi Matematychnyi Zhurnal* 74, 12 (2022), 1640–1653.

Viacheslav Borovitskiy, Alexander Terenin, Peter Mostowsky, and Marc Deisenroth (he/him). 2020. Matérn Gaussian Processes on Riemannian Manifolds. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 12426–12437.

Ken Brodlie, Rodolfo Allendes Osorio, and Adriano Lopes. 2012. A review of uncertainty in data visualization. In *Expanding the frontiers of visual analytics and visualization*. Springer, 81–109.

Hanyu Chen, Bailey Miller, and Ioannis Gkioulekas. 2024. 3D Reconstruction with Fast Dipole Sums. *ACM Transactions on Graphics* 43, 6 (Dec. 2024), 1–19. doi:10.1145/3687914 Publisher: Association for Computing Machinery (ACM).

Jon Cockayne, Chris Oates, Tim Sullivan, and Mark Girolami. 2017. Probabilistic numerical methods for PDE-constrained Bayesian inverse problems. In *AIP Conference Proceedings*, Vol. 1853. AIP Publishing LLC, 060001.

Stanimir Dragiev, Marc Toussaint, and Michael Gienger. 2011. Gaussian process implicit surfaces for shape estimation and grasping. In *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2845–2850.

Lawrence C Evans. 2022. *Partial differential equations*. Vol. 19. American mathematical society.

Jérôme Fehrenbach and Jean-Marie Mirebeau. 2014. Sparse non-negative stencils for anisotropic diffusion. *Journal of Mathematical Imaging and Vision* 49, 1 (2014), 123–147.

Gaël Guennebaud and Markus Gross. 2007. Algebraic point set surfaces. In *ACM SIGGRAPH 2007 papers (SIGGRAPH '07)*. Association for Computing Machinery, New York, NY, USA, 23–es. doi:10.1145/1275808.1276406

Philipp Hennig, Michael A Osborne, and Hans P Kersting. 2022. *Probabilistic Numerics: Computation as Machine Learning*. Cambridge University Press.

Sidhanth Holalkere, David S. Bindel, Silvia Sellán, and Alexander Terenin. 2025. Stochastic Poisson Surface Reconstruction with One Solve using Geometric Gaussian Processes. doi:10.48550/arXiv.2503.19136

Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. 2006. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, Vol. 7.

Nathan King, Haozhe Su, Mridul Aanjaneya, Steven Ruuth, and Christopher Batty. 2024. A Closest Point Method for PDEs on manifolds with interior boundary conditions for geometry processing. *ACM Transactions on Graphics* 43, 5 (2024), 1–26.

Jacques-Olivier Lachaud, David Coeurjolly, Céline Labart, Pascal Romon, and Boris Thibert. 2023. Lightweight Curvature Estimation on Point Clouds with Randomized Corrected Curvature Measures. *Computer Graphics Forum* 42, 5 (Aug. 2023), e14910. doi:10.1111/cgf.14910

Bruno Lévy. 2006. Laplace-beltrami eigenfunctions towards an algorithm that" understands" geometry. In *IEEE International Conference on Shape Modeling and Applications 2006 (SMI'06)*. IEEE, 13–13.

Rongji Li, Qile Yan, and Shixiao W. Jiang. 2025. Generalized Moving Least-Squares Methods for Solving Vector-valued PDEs on Unknown Manifolds. doi:10.48550/arXiv.2406.12210

Jihao Andreas Lin, Javier Antorán, Shreyas Padhy, David Janz, José Miguel Hernández-Lobato, and Alexander Terenin. 2023. Sampling from Gaussian process posteriors using stochastic gradient descent. *Advances in Neural Information Processing Systems* 36 (2023), 36886–36912.

Abhishek Madan and David IW Levin. 2022. Fast evaluation of smooth distance constraints on co-dimensional geometry. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–17.

Abhishek Madan, Nicholas Sharp, Francis Williams, Ken Museth, and David I.W. Levin. 2025. Stochastic Barnes-Hut Approximation for Fast Summation on the GPU. In *ACM SIGGRAPH 2025 Conference Papers*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3721238.3730725

Bailey Miller, Hanyu Chen, Alice Lai, and Ioannis Gkioulekas. 2024. Objects as volumes: A stochastic geometry view of opaque solids. (2024).

Jean-Marie Mirebeau and Jorg Portegies. 2019. Hamiltonian fast marching: A numerical solver for anisotropic and non-holonomic eikonal PDEs. *Image processing on line* 9 (2019), 47–93.

Seiji Nabeya and Takeaki Kariya. 1986. Transformations preserving normality and Wishart-ness. *Journal of Multivariate Analysis* 20, 2 (1986), 251–264. doi:10.1016/0047-259X(86)90081-3

Mu Niu, Zhenwen Dai, Pokman Cheung, and Yizhu Wang. 2023. Intrinsic Gaussian Process on Unknown Manifolds with Probabilistic Metrics. (2023).

Mark Pauly, Niloy J Mitra, and Leonidas J Guibas. 2004. Uncertainty and Variability in Point Cloud Surface Data. (2004).

John Wilson Peoples and John Harlim. 2024. A Higher Order Local Mesh Method for Approximating Laplacians on Unknown Manifolds. doi:10.48550/arXiv.2405.15735

Tobias Pfaffelmoser, Matthias Reitinger, and Rüdiger Westermann. 2011. Visualizing the positional and geometrical variability of isosurfaces in uncertain scalar fields. In *Computer graphics forum*, Vol. 30. Wiley Online Library, 951–960.

Eduard Pujol and Antonio Chica. 2025. Bayesian 3D Shape Reconstruction from Noisy Points and Normals. (2025). https://doi.org/10.1111/cgf.70201 Publisher: The Eurographics Association and John Wiley & Sons Ltd..

Kai Pöthkow, Britta Weber, and Hans-Christian Hege. 2011. Probabilistic Marching Cubes. *Computer Graphics Forum* 30, 3 (June 2011), 931–940. doi:10.1111/j.1467-8659.2011.01942.x Publisher: Wiley.

Steven J Ruuth and Barry Merriman. 2008. A simple embedding method for solving partial differential equations on surfaces. *J. Comput. Phys.* 227, 3 (2008), 1943–1961.

Matthias Seeger. 2004. Gaussian processes for machine learning. *International journal of neural systems* 14, 02 (2004), 69–106.

Silvia Sellán and Alec Jacobson. 2022. Stochastic Poisson Surface Reconstruction. *ACM Transactions on Graphics* 41, 6 (Dec. 2022), 1–12. doi:10.1145/3550454.3555441

Manu Sethi, Anand Rangarajan, and Karthik Gurumoorthy. 2012. The Schrödinger distance transform (SDT) for point-sets and curves. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 198–205.

Dario Seyb, Eugene d'Eon, Benedikt Bitterli, and Wojciech Jarosz. 2024. From microfacets to participating media: A unified theory of light transport with stochastic geometry. *ACM Transactions on Graphics* 43, 4 (July 2024), 1–17. doi:10.1145/3658121

IM Soboĺ. 1990. Quasi-monte carlo methods. *Progress in Nuclear Energy* 24, 1-3 (1990), 55–61.

Alessandra Tosi, Søren Hauberg, Alfredo Vellido, and Neil D. Lawrence. 2014. Metrics for Probabilistic Geometries. doi:10.48550/arXiv.1411.7432 arXiv:1411.7432 [stat].

Philipp Trettner and Leif Kobbelt. 2020. Fast and Robust QEF Minimization using Probabilistic Quadrics. *Computer Graphics Forum* 39, 2 (May 2020), 325–334. doi:10.1111/cgf.13933

Junyang Wang, Jon Cockayne, Oksana Chkrebtii, T. J. Sullivan, and Chris J. Oates. 2021. Bayesian Numerical Methods for Nonlinear Partial Differential Equations. *Statistics and Computing* 31, 5 (Sept. 2021). doi:10.1007/s11222-021-10030-w

Oliver Williams and Andrew Fitzgibbon. 2006. Gaussian process implicit surfaces. In *Gaussian Processes in Practice*. https://www.microsoft.com/en-us/research/wp-content/uploads/2006/06/gpc.pdf

Lukas Winiwarter, Alberto Manuel Esmorís Pena, Hannah Weiser, Katharina Anders, Jorge Martínez Sánchez, Mark Searle, and Bernhard Höfle. 2022. Virtual laser scanning with HELIOS++: A novel take on ray tracing-based simulation of topographic full-waveform 3D laser scanning. *Remote Sensing of Environment* 269 (2022). doi:10.1016/j.rse.2021.112772

Lan Wu, Ki Myung Brian Lee, Cedric Le Gentil, and Teresa Vidal-Calleja. 2024. Log-GPIS-MOP: A Unified Representation for Mapping, Odometry and Planning. doi:10.48550/arXiv.2206.09506

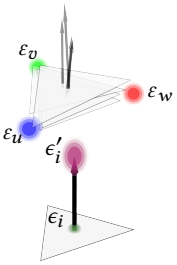
Kehan Xu, Benedikt Bitterli, Eugene d'Eon, and Wojciech Jarosz. 2025. Practical Gaussian Process Implicit Surfaces with Sparse Convolutions. *ACM Transactions on Graphics (TOG)* 44, 6 (2025), 1–18.

A Difference of interpretation with regression based GPIS

Compared to the Gaussian process regression approach of previous works [Sellán and Jacobson 2022; Williams and Fitzgibbon 2006], this approach has a different probabilistic focus. While GPR focuses on quantifying the uncertainty arising when extrapolating data, hence with a variance increasing away from known regression points, this model has the opposite behavior, as the output is more uncertain close to the input points and averages out away from the data. However, since the Kac-Rice formula Eq. 9 takes into account the norm of the gradient of the field φ, the actual surface density rapidly decreases in regions where data is missing, as illustrated in Fig. 3.

B Geometric noise model from triangle soups

The joint distribution of the noise on the inputs we assume in Sec. 4 can be derived explicitly from a noisy triangle soup (*i.e.*; a set of triangles without connectivity information): for each triangle $T_i = \{x_u, x_v, x_w\}$, we assume that each vertex has been perturbed by centered Gaussian noise of standard deviation σ (see inset). This noise impacts the estimation of the triangle's barycenter p_i, and of its vector-area ω_i. Since each ε_i follows a multivariate Gaussian distribution $\varepsilon_i \sim \mathcal{N}(0, \sigma I_3)$, ϵ_i and ϵ_i' of Eq. 5 form a $2d$ dimensional Gaussian vector given by:



$$\begin{pmatrix}\epsilon_i\\ \epsilon_i'\end{pmatrix} = \begin{pmatrix} I_d/3 & I_d/3 & I_d/3 \\ [x_w - x_v]_\times & [x_u - x_w]_\times & [x_v - x_u]_\times \end{pmatrix} \begin{pmatrix}\varepsilon_u\\ \varepsilon_v\\ \varepsilon_w\end{pmatrix},$$

where $[v]_\times$ is the matrix of the cross product with v (here divided by two in order to match the area formula). The cross correlations between positions and normals cancel out but the scale of the noise on the weighted normal is set consistently. This model is illustrated in Fig. 8, Fig. 11 and Fig. 12.

C Derivation of the Δ_φ operator and weak convergence

If φ is a smooth Gaussian random field and $f \in H_0^1(\Omega)$, from Eq. 14, we have:

$$\begin{aligned}
\mathbb{E}\left[E_D(f)\right] &= \int_\Omega \mathbb{E}_{\varphi=u}\left[\|\Pi_{\nabla\varphi^\perp}\nabla f\|^2 \|\nabla\varphi(x)\|\right] p_{\varphi(x)}(u)dx \\
&= \int_\Omega \mathbb{E}_{\varphi=u}\left[\langle \nabla f, \Pi_{\nabla\varphi^\perp}\nabla f\rangle \|\nabla\varphi(x)\|\right] p_{\varphi(x)}(u)dx \qquad \text{(i)} \\
&= \int_\Omega \langle \nabla f, p_{\varphi(x)}(u)\mathbb{E}_{\varphi=u}\left[\|\nabla\varphi(x)\|\Pi_{\nabla\varphi^\perp}\right]\nabla f\rangle dx \qquad \text{(ii)} \\
&= -\int_\Omega f\nabla\cdot\left(p_{\varphi(x)}(u)\mathbb{E}_{\varphi=u}\left[\|\nabla\varphi(x)\|\Pi_{\nabla\varphi^\perp}\right]\nabla f\right)dx \qquad \text{(iii)} \\
&= -\int_\Omega f\frac{1}{\mu}\nabla\cdot\left(\underbrace{p_{\varphi(x)}(u)\mathbb{E}_{\varphi=u}\left[\|\nabla\varphi(x)\|\Pi_{\nabla\varphi^\perp}\right]}_{T}\nabla f\right)d\mu \\
&= -\langle f, \frac{1}{\mu}\nabla\cdot(T\nabla f)\rangle_{L^2(\Omega,\mu)} = -\langle f, \Delta_\varphi f\rangle_{L^2(\Omega,\mu)}\,.
\end{aligned}$$

Hence, up to a sign, $\Delta_\varphi f$ is the functional gradient of the expected Dirichlet energy with respect to $L^2(\Omega,\mu)$. To obtain (i), we used $\Pi^\top_{\nabla\varphi^\perp}\Pi_{\nabla\varphi^\perp} = \Pi_{\nabla\varphi^\perp}$, between (i) and (ii) the linearity of the expected value, between (ii) and (iii) by integration by parts where the boundary term vanishes since $f \in H^1_0$. Remaining lines consist in considering $dx = \frac{1}{\mu}d\mu$ and using Eq. 16 to obtain Δ_φ.

The weak convergence to the Laplacian on implicit surfaces follows the convergence of the expected Dirichlet energy to its mean, *i.e.* the Dirichlet energy of the mean surface:

$$\mathbb{E}\left[E_D(f)\right] = \int_\Omega \langle \nabla f, p_{\varphi(x)}(u)\mathbb{E}_{\varphi=u}\left[\|\nabla\varphi(x)\|\Pi_{\nabla\varphi^\perp}\right]\nabla f\rangle dx\,,$$

which converges, when the variance σ of φ tends to 0, to:

$$\int_\Omega \langle \nabla f, \Pi_{\nabla\varphi^\perp}\nabla f\rangle \|\nabla\varphi(x)\|\delta_{\varphi=u}dx\,. \tag{32}$$

This last expression is the Dirichlet energy of the mean surface by the co-area formula and is the expression used by Bertalmio et al. [2001] to derive Eq. 17 as the gradient of Eq. 32 with respect to $L^2(\Omega, \|\nabla\varphi\|dx)$.

D Derivations of the divergence operator

In order to be self-contained, we provide the derivation leading to our divergence operator from the Lie derivative operator:

$$\begin{aligned}
\mathcal{L}_V(d\mu) = \mathcal{L}_V(\mu dx) &= \nabla\mu\cdot V dx + \mu\mathcal{L}_V(dx)\\
&= \nabla\mu\cdot V dx + \mu\nabla\cdot V dx\\
&= \frac{\nabla\mu}{\mu}\cdot V d\mu + \frac{\mu}{\mu}\nabla\cdot V d\mu\\
&= \frac{1}{\mu}\left(\nabla\mu\cdot V + \mu\nabla\cdot V\right)d\mu\\
&= \frac{1}{\mu}\nabla\cdot(\mu V)\,d\mu = \mathrm{div}_\varphi(V)d\mu\,.
\end{aligned}$$

The μ-measure preserving property of a div_μ-free flow can be proved using the Reynolds theorem. Given a set A_t, advected by V, i.e., $A_t = \phi^*_V(A_0)$, we have, $\forall t \geq 0$:

$$\begin{aligned}
\frac{d}{dt}d\mu(A_t) &= \frac{d}{dt}\int_{A_t}\mu(x)dx\\
&= \underbrace{\int_{A_t}\frac{\partial}{\partial t}\mu(x)dx}_{0} + \int_{\partial A_t}\mu(x)V(x)\cdot n(x)d\sigma(x)\\
&= \int_{A_t}\nabla\cdot(\mu(x)V(x))\,dx = \int_{A_t}\frac{1}{\mu(x)}\nabla\cdot(\mu(x)V(x))\,d\mu\,,
\end{aligned}$$

where the third equality is obtained by the divergence theorem. Hence, the measure is indeed preserved if and only if $\mathrm{div}_\mu(V) = 0$.

E Structure of the first order operators

Over a Riemannian manifold, the divergence and the gradient operators are adjoint with respect to the volume form and the metric tensor. Such structure is preserved when considering the following pair of inner products:

$$\langle \mathrm{div}_\varphi V, f\rangle_{L^2(\Omega,\mu)} = -\langle V, \mathrm{grad}_\varphi f\rangle_{L^2(\Omega,S)}\,,$$

where $S := \mu^2 T^{-1}$. This inner product also arises in the Hodge decomposition. Indeed, the projection Eq. 29 does not correspond to an orthogonal projection with respect to the canonical inner product of $\mathbb{R}^{dV}$. The orthogonal projection onto the kernel of an operator A with respect to an inner product of the form

$$\langle x, y\rangle_Q := x\cdot Qy\,,$$

expressed as a solution to:

$$\Pi_{\ker(A),Q}(x) = \underset{y\in\mathrm{Ker}(A)}{\arg\min}\ \|x-y\|^2_Q$$

is given by:

$$\Pi_{\ker(A),Q} = I - Q^{-1}A^\top(AQ^{-1}A)^{-1}A\,.$$

From the definition of the operators, one can check that the projection Eq. 29 on $\ker(\mathrm{div}_\varphi)$ is orthogonal when $Q = S$.

Remarkably, this inner product converges in the weak sense to the metric tensor on $\gamma^{-1}(u) = \mathcal{M}$: for any vector fields $u, v \in H^1_0(\Omega)$ we have

$$\langle u, v\rangle_{L^2(S)} \xrightarrow{\sigma\to 0} \int_\Omega u(x)\cdot y(x)\|\nabla\gamma(x)\|\delta_{\gamma^{-1}(u)}(x)dx = \langle u|_M, v|_\mathcal{M}\rangle_{L^2(\mathcal{M})}$$

if $u|_\mathcal{M}, v|_\mathcal{M} \in TM$, $+\infty$ else (as S involves the inverse of the expected orthogonal projection). In summary, the first order operators we propose are adjoint with respect to an inner product that converges weakly to the metric tensor of the mean surface when the variance tends to 0.

F Behavior of the operators on regression based GPIS

We verify that our approach is agnostic to the method used to construct the GPIS. For this experiment, we infer the field φ via standard GP regression [Williams and Fitzgibbon 2006] using the same kernel, with training samples set to 1 inside the sphere, 0 on the sphere, and -1 outside. More precisely, from a Gaussian process defined by a kernel $k(x,x')$, the joint distribution of the value/gradient is given by another GP with block kernel:

$$k^\nabla := \begin{pmatrix} k & \nabla_x k^\top \\ \nabla_{x'}k & \nabla_x\nabla_{x'}k\end{pmatrix}.$$

The resulting behavior matches that of our model in Eq. 5: geodesic distances slow down in regions with higher uncertainty.

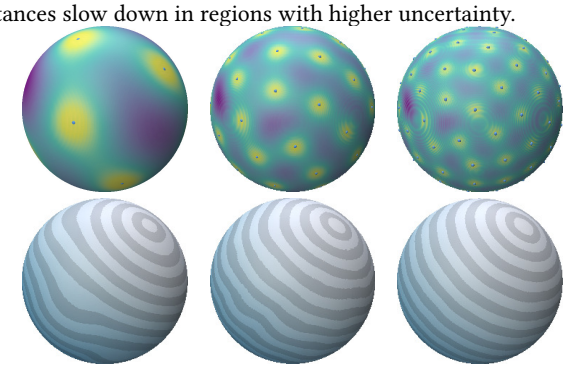


Fig. 13. Convergence of the geodesic distance when the variance of a regression-based GPIS model diminishes. Here, the GPIS is built from a regression process by learning the indicator function of a sphere for more and more samples.