Uncertainty-aware geometry processing on Gaussian Process Implicit Surfaces
Supplementary material

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

ACM Reference Format:
Baptiste Genest and David Coeurjolly. 2026. Uncertainty-aware geometry processing on Gaussian Process Implicit Surfaces Supplementary material. *ACM Trans. Graph.* 45, 4 (July 2026), 3 pages. https://doi.org/10.1145/3811280

1 Barnes-Hut scheme [Barnes and Hut 1986]

Following the Barnes-Hut implementation of Chen et al. [2024], the β parameter controls the criterion for approximating a subtree: the larger the β, the more accurate (but more expensive) the approximation. Using this definition, we first validate that the Barnes–Hut approximation for jointly querying the field $\varphi(x)$ and its gradient $\nabla\varphi(x)$ converges as β increases (Fig. 1-top). To assess the suitability of this scheme for joint GPIS queries, we study the approximation error of the resulting surface density μ.

In Fig. 1-bottom, show the query time on a grid for the Barnes-Hut scheme with respect to input size and the β parameter (set to 3 in our experiments).

2 Quasi-Monte Carlo estimation

We study the Monte Carlo error when approximating the expected value of non-linear functions, as a function of the number of low-discrepancy samples S. In Fig. 2, we compare the μ density estimated using different S from a reference value (constructed with a large S value). In our experiments, 512 samples were used.

3 Convergence of the numerical scheme for the Poisson problem

In order to avoid memory issues at very high resolutions, we study the convergence of the error for Poisson problems on the following simple 2D GPIS model:

$$\varphi(x) = \|x\| - 1 + \sum_{i}^{2} x_i \epsilon_i$$

where $\epsilon_i \sim \mathcal{N}(0, \sigma^2 I_2)$, for varying resolutions and η parameter. Fig. 3 shows the impact of the parameters on the order of convergence. We observe that η has a limited impact while higher σ values decreases slightly the order of convergence (convergence of order ≈ 1.3 for highest σ).

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.

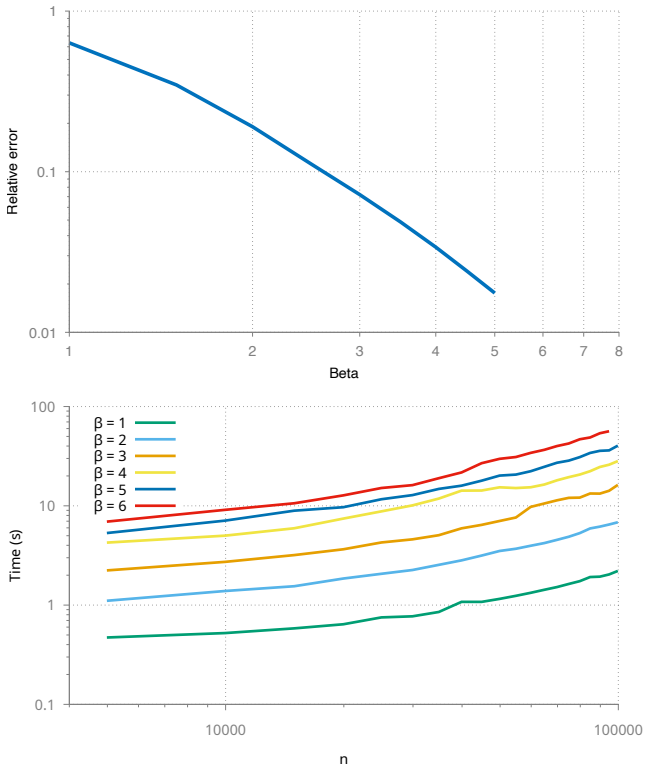

Fig. 1. (top) On the spot object (see main paper), we plot the relative error on the μ field for increasing values of β in the Barnes-Hut scheme, on a fixed grid with $h = \frac{1}{100}$ and $\sigma = 1.0$. (Bottom) Query time as a function of the input size n and β on a fixed grid with $h = \frac{1}{100}$ and $\sigma = 1.0$ (result on samples of the bunny shape.)

4 Eigenfunctions of the Δ_φ operator for varying σ

In Fig. 5, we validate the stability of the eigenfunctions associated to the smallest eigenvalues of the cotan Laplacian of the mean surface, and of our $M_\varphi^{-1} L_\varphi$ operator on the mean surface for various sigma σ values and $h = 1/200$.

5 Behavior of the Heat Method

The heat method of Crane et al. [2013] is a standard algorithm to compute geodesic distances on surfaces from simple differential operators. However, it is based on Varadhan's formula which does not hold in our setting to the best of our knowledge, and amounts to solving a different equation than the one we derived in Eq. 29. We nonetheless show results in Fig. 4 where we replace each operator

ACM 1557-7368/2026/7-ART
https://doi.org/10.1145/3811280

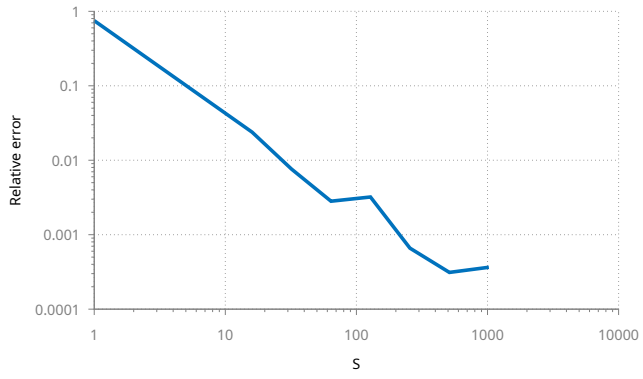

Fig. 2. Relative error on the μ field for increasing values of S on a fixed grid with $h = \frac{1}{100}$ and $\sigma = 1$, reference value computed with $S^\star = 8192$. Result on the Spot shape.

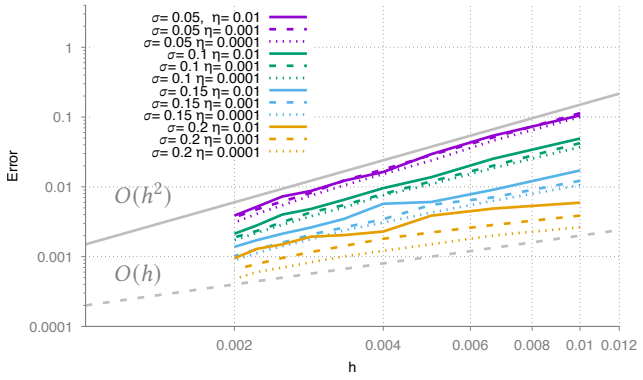

Fig. 3. Convergence of the distance, for various σ and η values between the solution of a Poisson problem when the resolution diminishes, for multiple threshold values η. Computed as $\|\tilde{u}_h - u^\star\|_\mu/\mu(\Omega)$ where $u_h = \Delta_\varphi^{-1} f_h$ and $u^\star = \Delta_\varphi^{-1} f_{h^\star}$, were $h^\star = 10^{-4}$ and $\tilde{u}_h$ is an upsampled version of u_h to $h^\star$, $f(x, y) = \sin(10x)\sin(10y)$.

used in the heat method by ours and use $\|\cdot\|_T$ as norm. Compared to the fast marching method we use, the uncertainty also affects the results but shows a less interesting behavior (over-smoothed).

6 Selling's algoritm

A superbase of $\mathbb{Z}^d$ is a tuple $(\mathbf{b}_0, \dots, \mathbf{b}_d) \in (\mathbb{Z}^d)^{d+1}$ such that

$$\mathbf{b}_0 + \cdots + \mathbf{b}_d = 0, \qquad |\det(\mathbf{b}_1, \dots, \mathbf{b}_d)| = 1\,.$$

From the superbase obtained by Alg. 1, the decomposition *Eq.* **??** is given by:

$$T = -\sum_{0 \le i < j \le d} \underbrace{\langle \mathbf{b}_i, T\mathbf{b}_j \rangle}_{\lambda_{i'}} \underbrace{\mathbf{v}_{ij}\mathbf{v}_{ij}^T}_{\mathbf{e}_{i'}},$$

with $\mathbf{v}_{ij} = \mathbf{b}_k \times \mathbf{b}_l$ where $\{i, j, k, l\} = \{0, 1, 2, 3\}$ and i' goes from 0 to $d' = \frac{d(d+1)}{2}$. Selling's algoritms runs in $O(\log(\kappa(T)))$ iterations where $\kappa(T) = \frac{\lambda^{\max}(T)}{\lambda^{\min}(T)}$ is the anisotropy ratio of T [Fehrenbach and Mirebeau 2014]. In our experiments the run time of the decomposition is negligible compared to the rest of the pipeline.

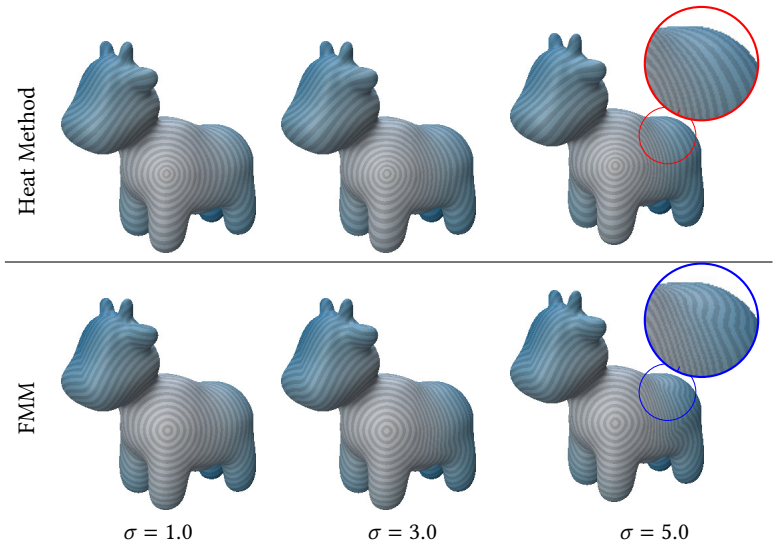

Fig. 4. Comparison between an analog of the Heat Method (top) using our operators (to be compared to Fig. 8 in the main paper), and our fast-marching approach (bottom).

Algorithm 1: Selling's algorithm in 3D

Data: Matrix $T \in S_3^{++}$
// Initialize superbase
$\mathbf{b}_0 = (-1, -1, -1), \mathbf{b}_1 = (1, 0, 0), \mathbf{b}_2 = (0, 1, 0), \mathbf{b}_3 = (0, 0, 1)$;
$B = (\mathbf{b}_0, \mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3)$;
while $\exists i, j$ *s.t.* $\langle \mathbf{b}_i, T\mathbf{b}_j \rangle > 0$ **do**
  $B \leftarrow (\mathbf{b}_i, -\mathbf{b}_j, \mathbf{b}_k + \mathbf{b}_j, \mathbf{b}_l + \mathbf{b}_j)$; // $\{i, j, k, l\} = \{0, 1, 2, 3\}$
return B;

References

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

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).

Keenan Crane, Clarisse Weischedel, and Max Wardetzky. 2013. Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Transactions on Graphics (ToG)* 32, 5 (2013), 1–11.

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.

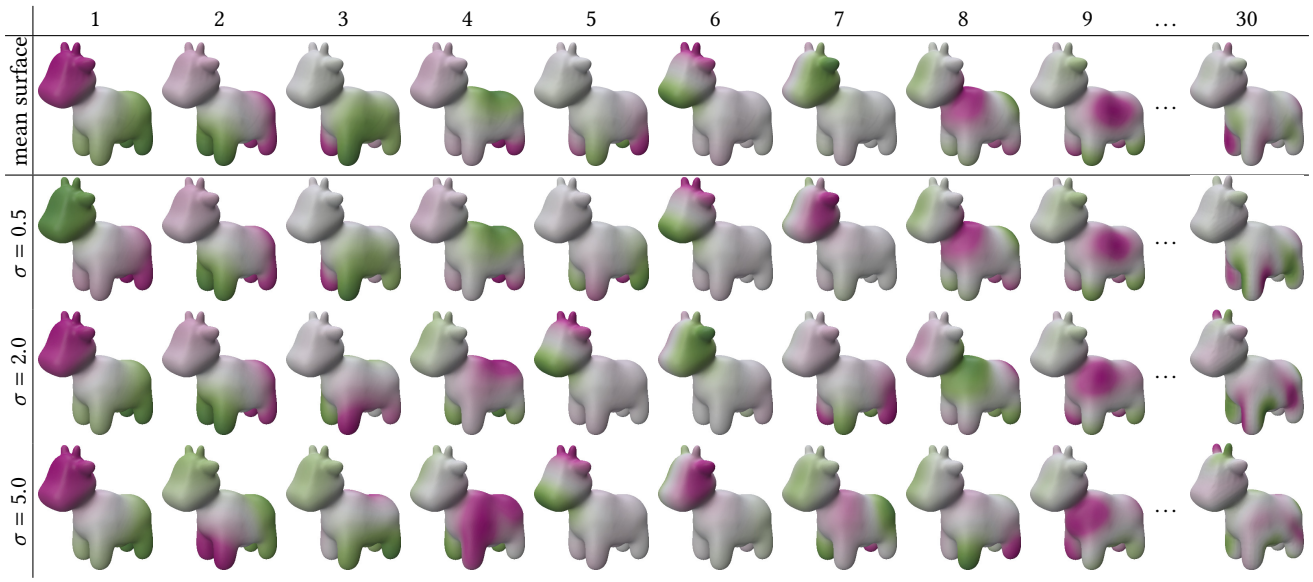

Fig. 5. We show the eigenfunctions associated to the smallest 9 eigenvalues of the cotan Laplacian of the mean surface (first row) and of our $M_\varphi^{-1}L_\varphi$ operator (also the 30th) on the mean surface for various sigma σ values and $h = 1/200$.