Explicit Geometry for CPU-Efficient Mapping: Bridging Discrete Grids and Continuous Fields

José E. Maese[1]

***Abstract*— While recent trends in continuous volumetric mapping heavily favor GPU-intensive neural representations, this computational reliance often limits their deployment on resource-constrained robotic platforms. To address this, we present two efficient, CPU-optimized mapping paradigms based on deterministic geometric formulations. First, we detail a Directional Bitmask-based Truncated Signed Distance Field (DB-TSDF) that utilizes precomputed kernels for fast, $O(1)$ discrete spatial integration. Second, we present a Gaussian Euclidean Distance Field (G-EDF), a block-sparse Gaussian Mixture Model that explicitly regresses the distance field to provide a C^1 continuous, analytically differentiable representation. Finally, we discuss our ongoing research to unify these distinct approaches into a highly asynchronous SLAM backend, where the discrete DB-TSDF method is leveraged to rapidly bootstrap the continuous G-EDF blocks. Ultimately, this work highlights how analytically grounded geometric models offer a scalable and computationally efficient alternative to data-intensive neural approaches for autonomous navigation.**

I. INTRODUCTION

Volumetric mapping remains a fundamental prerequisite for autonomous navigation, providing the spatial context required for collision avoidance and motion planning under real-time constraints [1], [2]. Traditional CPU-based platforms frequently rely on discrete structures such as point clouds or occupancy grids due to their simplified data management [3]. However, these representations are susceptible to aliasing and discretization artifacts at high resolutions [4]. To mitigate these issues, Truncated Signed Distance Fields (TSDFs) have been widely adopted to provide smooth proximity information [1], [2], [5]. Despite their benefits, many state-of-the-art TSDF pipelines exhibit computational costs on the CPU that scale unfavorably with resolution or rely heavily on GPU acceleration [6], [7].

In this context, we previously introduced DB-TSDF, a discrete mapping framework optimized for high-resolution CPU execution [8]. By employing a directional bitmask-based integration scheme, each voxel encodes a compact 32-bit distance mask and a hit counter updated via bitwise AND operations [9]. This geometric structure utilizes precomputed directional kernels to model the anisotropic footprint of LiDAR return. Crucially, the system achieves a constant per-scan computational cost that is independent of the global grid dimensions, ensuring real-time performance through integer-only parallelization.

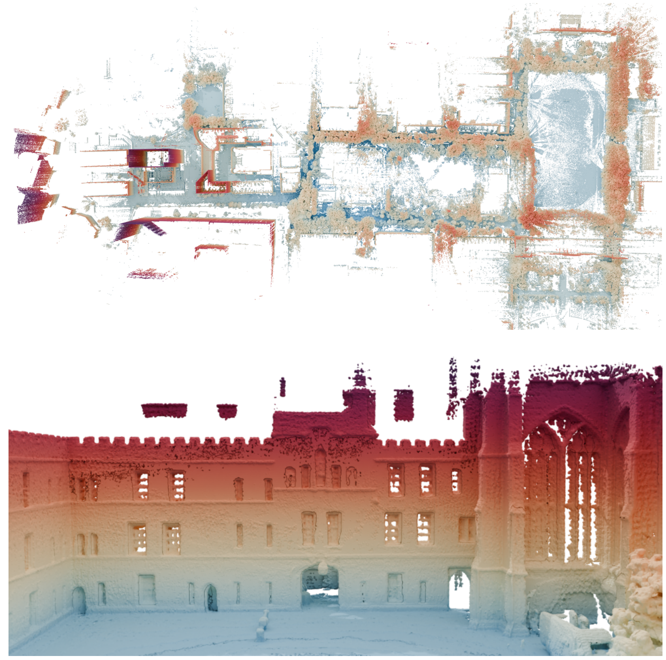

Fig. 1: Visual comparison of the proposed mapping methods. Top: Full-scale G-EDF reconstruction of the $900 \times 610 \times 120$ m *Snail* dataset, demonstrating scalability in unbounded outdoor environments. Bottom: detail of the reconstructed scene from the *Newer College* dataset using DB-TSDF.

Parallel to these discrete advancements, the need for C^1 continuity and smooth gradients for robust localization has driven research toward continuous field representations [4]. While neural approaches like Implicit Neural Representations (INRs) [10], [7] offer resolution-independent queries, their reliance on intensive GPU inference limits their deployment on resource-constrained robots. To address this, we developed G-EDF [11], a continuous 3D distance field based on a Block-Sparse Gaussian Mixture Model [12], [13]. Unlike probabilistic occupancy frameworks [3], [14], G-EDF utilizes anisotropic Gaussians as universal geometric function approximators to explicitly regress the Euclidean Distance Field (EDF) [13], [15]. This mathematical structure provides closed-form analytical gradients that empirically satisfy the Eikonal equation, enabling robust gradient-based localization on the CPU without the artifacts of trilinear interpolation.

Currently, we are working on a unified system that integrates these two distinct geometric paradigms into a high-

*This work was supported by the grants PICRA 4.0 (PLEC2023-010353), funded by the Spanish Ministry of Science and Innovation and the Spanish Research Agency (MCIN/AEI/10.13039/501100011033); and INSERTION (PID2021-127648OB-C31), funded by the "Agencia Estatal de Investigación – Ministerio de Ciencia, Innovación y Universidades" and the "European Union NextGenerationEU/PRTR".

[1]The author is with the Service Robotics Laboratory, Universidad Pablo de Olavide, Seville, Spain. jemaealv@upo.es

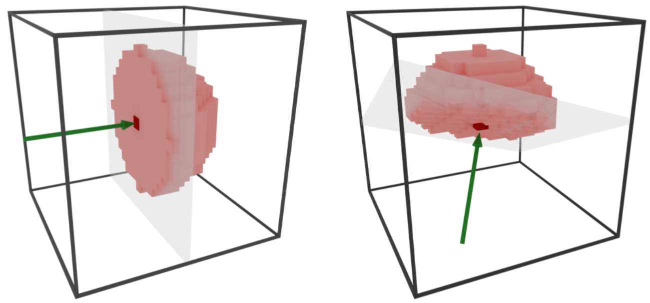

Fig. 2: Examples of the hemispherical shadow region (red) inside the $21 \times 21 \times 21$ voxel kernel for different ray directions (green arrows). The mask consistently aligns with the incoming LiDAR beam while preserving an exclusion pattern around the contact voxel.

efficiency CPU-SLAM framework. By leveraging a simplified version of DB-TSDF as a rapid spatial method, the system can generate a discrete occupancy grid to bootstrap the initialization of the continuous G-EDF blocks. This unified strategy enables a non-linear optimization process to focus on structural refinement rather than global topological discovery, achieving near real-time training of the Gaussian kernels. This evolution demonstrates that efficiency in robotics can be achieved by prioritizing geometric and mathematical rigor over data-intensive neural models, providing a deterministic and interpretable representation for large-scale environmental modeling.

II. Methodology Overview

This framework prioritizes explicit geometric representations over data-intensive neural models, enabling highly efficient, deterministic, and real-time execution entirely on the CPU. The proposed architecture bridges the gap between high-speed discrete mapping and mathematically continuous fields. It first employs a discrete directional grid to instantly capture the environmental topology, which then serves as the exact geometric foundation to bootstrap a continuous, analytically differentiable Gaussian field.

A. Truncated Signed Distance Field: DB-TSDF

To handle unstructured LiDAR data efficiently without GPU acceleration, the system relies on a Directional Bitmask-based Truncated Signed Distance Field (DB-TSDF). This discrete grid maintains an extremely low memory footprint of just 8 bytes per voxel by encapsulating a 32-bit distance mask, an 8-bit hit counter, and an 8-bit occupancy sign flag. The open-source implementation of this discrete mapping framework is publicly available at GitHub[1].

The integration bypasses traditional ray-casting by utilizing precomputed directional kernels. The directional space is discretized into angular bins, and for a given point $\mathbf{p} = (x, y, z)$ in the sensor frame, its corresponding azimuth and elevation bin indices are directly computed to select the appropriate kernel.

Within each precomputed $21 \times 21 \times 21$ cubic kernel, the 32-bit mask encodes the truncated distance to the kernel center. This binary distance mask can be computed using a standard L_1 norm or a more accurate L_2 norm via a precomputed look-up table. Concurrently, the hemispherical shadow region, used to determine voxel occupancy, is explicitly modeled using an L_2 norm to prevent discretization artifacts and maintain a consistent exclusion pattern (Fig. 2). The kernel thus natively models the anisotropic footprint of the LiDAR beam:

$$m(x,y,z) = \begin{cases} 0, & \text{if } r = 0, \\ (2^{32} - 1) \gg (32 - \lceil r \rceil), & \text{if } r > 0 \end{cases}$$

For an incoming point, the corresponding voxel's distance mask is updated via a single bitwise AND operation. This guarantees that the stored distance only decreases as new evidence is accumulated, maintaining a constant computational cost per scan independent of the global grid size:

$$m_{\text{grid}} \leftarrow \begin{cases} m_{\text{grid}} \& m_{\text{kernel}}, & \text{if } m_{\text{grid}} \neq m_{\text{grid}} \& m_{\text{kernel}}, \\ m_{\text{grid}}, & \text{otherwise} \end{cases}$$

B. Continuous Distance Field: G-EDF

While the DB-TSDF provides rapid spatial context, gradient-based localization requires C^1 continuity to prevent the local minima caused by trilinear interpolation. To achieve this, the system models the Euclidean Distance Field (EDF) as a continuous function approximated by a Block-Sparse Gaussian Mixture Model (G-EDF)[2].

Using anisotropic Gaussians as universal geometric function approximators, the continuous distance field $\hat{d}(\mathbf{x})$ at a query point $\mathbf{x} \in \mathbb{R}^3$ is expressed as a weighted sum of K kernels:

$$\hat{d}(\mathbf{x}) = \sum_{k=1}^{K} w_k \cdot \exp\left(-\frac{1}{2} \sum_{j \in \{x,y,z\}} \frac{(x_j - \mu_{kj})^2}{l_{kj}^2} \right)$$

Crucially, the weights w_k are not constrained to be positive. Negative weights act as subtractive terms that carve out sharp minima at the surface manifold, ensuring a tight fit to the underlying geometry.

This mathematical structure yields closed-form analytical gradients, enabling exact evaluation without numerical differentiation errors:

$$\nabla \hat{d}(\mathbf{x}) = -\sum_{k=1}^{K} g_k(\mathbf{x}) \cdot \mathbf{\Sigma}_k^{-1}(\mathbf{x} - \boldsymbol{\mu}_k)$$

By explicitly regressing the distance field, the continuous representation naturally approximates the Eikonal equation ($\|\nabla \hat{d}\| \approx 1$), providing consistent and reliable gradients for robust localization (Fig. 3).

[1] https://github.com/robotics-upo/DB-TSDF

[2] https://github.com/robotics-upo/G-EDF

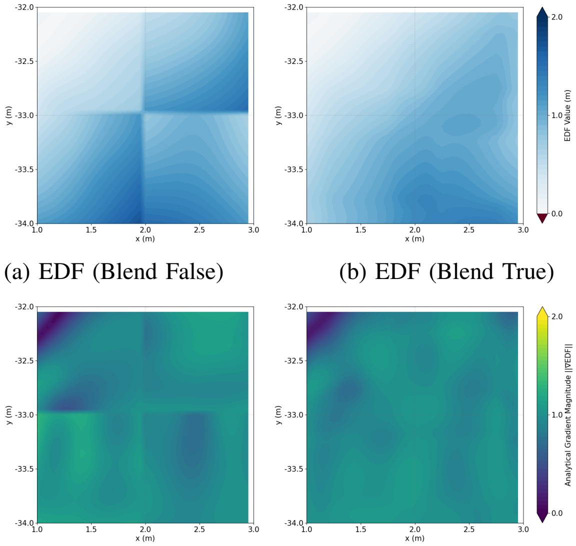

(a) EDF (Blend False) (b) EDF (Blend True)

(c) Gradient (Blend False) (d) Gradient (Blend True)

Fig. 3: Cross-section analysis on a local region of the New College Dataset [16] ($z = 3.0$ m) using 1.0 m^3 blocks with $\delta = 0.25$ m overlap. Left (a, c): Discontinuities in the distance field and gradient without blending. Right (b, d): The proposed G-EDF representation with overlapping blending enabled, ensuring global C^1 continuity.

III. Recent Advancements: From Static Maps to Real-Time SLAM

Building upon these foundational representations, our current ongoing research focuses on integrating both methodologies into a cohesive, real-time spatial estimation framework. A central tenet of this active development is the commitment to explicit geometric representations and analytical optimization. By avoiding the computational bottlenecks associated with neural representations, the unified framework is explicitly designed to execute efficiently on standard CPU architectures.

To transition from static mapping to an unbounded SLAM system, the architecture is being developed as a highly asynchronous framework. The discrete DB-TSDF handles the instantaneous, $O(1)$ integration of incoming LiDAR scans, providing an immediate structural context. In parallel, background CPU threads progressively train and refine the continuous G-EDF blocks. To manage large-scale exploration without exhausting system memory, the framework employs a dynamic sliding window mechanism. This strategy ensures that the dense discrete grid is only maintained for the robot's immediate surroundings. Once a region falls outside this local radius, its lightweight, optimized Gaussian representation is archived and periodically published as a highly compressed, self-contained local map for downstream tasks.

Furthermore, this continuous geometric approach opens new, efficient avenues for place recognition. As part of this ongoing project, we are developing an analytical loop closure module that directly exploits the mathematical continuity of the G-EDF. Instead of relying on computationally heavy point-cloud matching, the system evaluates the explicit distance field across spherical shells to extract rotation-invariant geometric signatures. This allows for rapid, deterministic candidate matching. Ultimately, this ongoing development demonstrates the immense potential of utilizing mathematically rigorous, continuous geometric representations to drive end-to-end SLAM pipelines on resource-constrained platforms.

IV. Conclusions

This work demonstrates that deterministic, analytically grounded geometric methods can effectively rival modern neural approaches for real-time volumetric mapping on resource-constrained platforms. By introducing the fast, discrete DB-TSDF and bridging it with the continuous, analytically differentiable G-EDF, we have established a highly efficient foundation for CPU-based spatial modeling.

Our ongoing efforts to unify these structures into an asynchronous SLAM backend highlight the practical advantages of this unified approach: $O(1)$ data ingestion, C^1 continuous gradients for robust trajectory optimization, and mathematical signatures for instantaneous place recognition. Therefore, this framework stands as a highly deployable and interpretable solution for autonomous navigation, proving that rigorous mathematical structures can drive state-of-the-art robotics without the overhead of GPU acceleration.

References

[1] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, sep 2017, p. 1366–1373. [Online]. Available: http://dx.doi.org/10.1109/IROS.2017.8202315

[2] L. Han, F. Gao, B. Zhou, and S. Shen, "Fiesta: Fast incremental euclidean distance fields for online motion planning of aerial robots," 2019. [Online]. Available: https://arxiv.org/abs/1903.02144

[3] K. Wurm, A. Hornung, M. Bennewitz, C. Stachniss, and W. Burgard, "Octomap: A probabilistic, flexible, and compact 3d map representation for robotic systems," vol. 2, 01 2010.

[4] F. Caballero and L. Merino, "DLL: Direct LIDAR Localization. A map-based localization approach for aerial robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021.

[5] K. Museth, "Vdb: High-resolution sparse volumes with dynamic topology," *ACM Trans. Graph.*, vol. 32, pp. 27:1–27:22, 07 2013.

[6] L. Wu, C. L. Gentil, and T. Vidal-Calleja, "Vdb-gpdf: Online gaussian process distance field with vdb structure," 2024. [Online]. Available: https://arxiv.org/abs/2407.09649

[7] X. Zhong, Y. Pan, J. Behley, and C. Stachniss, "Shine-mapping: Large-scale 3d mapping using sparse hierarchical implicit neural representations," 2023. [Online]. Available: https://arxiv.org/abs/2210.02299

[8] J. E. Maese, F. Caballero, and L. Merino, "Db-tsdf: Directional bitmask-based truncated signed distance fields for efficient volumetric mapping," in *2026 IEEE International Conference on Robotics and Automation (ICRA)*, 2026.

[9] L. Coto-Elena, J. E. Maese, L. Merino, and F. Caballero, "D-lio: 6dof direct lidar-inertial odometry based on simultaneous truncated distance field mapping," 2025. [Online]. Available: https://arxiv.org/abs/2505.16726

[10] J. Ortiz, A. Clegg, J. Dong, E. Sucar, D. Novotny, M. Zollhoefer, and M. Mukadam, "isdf: Real-time neural signed distance fields for robot perception," 2022. [Online]. Available: https://arxiv.org/abs/2204.02296

[11] J. E. Maese, L. Coto-Elena, L. Merino, and F. Caballero, "G-edf-loc: 3d continuous gaussian distance field for robust gradient-based 6dof localization," 2026. [Online]. Available: https://arxiv.org/abs/2604.04525

[12] P. Biber and W. Strasser, "The normal distributions transform: a new approach to laser scan matching," in *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003) (Cat. No.03CH37453)*, vol. 3, 2003, pp. 2743–2748 vol.3.

[13] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3d gaussian splatting for real-time radiance field rendering," 2023. [Online]. Available: https://arxiv.org/abs/2308.04079

[14] F. Ramos and L. Ott, "Hilbert maps: Scalable continuous occupancy mapping with stochastic gradient descent," *The International Journal of Robotics Research*, vol. 35, pp. 1717–1730, 12 2016.

[15] A. Guédon and V. Lepetit, "Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering," in *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024, pp. 5354–5363.

[16] L. Zhang, M. Camurri, D. Wisth, and M. Fallon, "Multi-camera lidar inertial extension to the newer college dataset," 2022. [Online]. Available: https://arxiv.org/abs/2112.08854