

HIO-SDF: Hierarchical Incremental Online Signed Distance Fields

Vasileios Vasilopoulos, Suveer Garg, Jinwook Huh, Borham Lee and Volkan Isler

Abstract—A good representation of a large, complex mobile robot workspace must be space-efficient yet capable of encoding relevant geometric details. When exploring unknown environments, it needs to be updatable incrementally in an online fashion. We introduce HIO-SDF, a new method that represents the environment as a Signed Distance Field (SDF). State of the art representations of SDFs are based on either neural networks or voxel grids. Neural networks are capable of representing the SDF continuously. However, they are hard to update incrementally as neural networks tend to forget previously observed parts of the environment unless an extensive sensor history is stored for training. Voxel-based representations do not have this problem but they are not space-efficient especially in large environments with fine details. HIO-SDF combines the advantages of these representations using a hierarchical approach which employs a coarse voxel grid that captures the observed parts of the environment together with high-resolution local information to train a neural network. HIO-SDF achieves a 46% lower mean global SDF error across all test scenes than a state of the art continuous representation, and a 30% lower error than a discrete representation at the same resolution as our coarse global SDF grid. Videos and code are available at: <https://samsunglabs.github.io/HIO-SDF-project-page/>

I. INTRODUCTION

Constructing a geometric representation of a robot’s environment which is suitable for real-time planning and control is a fundamental robotics challenge. A good representation must be accurate, support fast inference and, for exploration tasks, incremental update. Signed Distance Field (SDF) [1] is a scalar field that maps each point in the 3D workspace to its signed distance to the closest obstacle. Representing the robot’s environment as a SDF can be an efficient method for performing collision cost queries during reactive planning [2], [3]. However, the underlying field must be incrementally updated online.

Discrete online SDF representations such as Voxblox [4], FIESTA [5] or Voxfield [6] have become popular due to their accuracy, relatively low memory footprint and fast update rates. However, discrete methods have several disadvantages. Interpolation is necessary to obtain continuous SDF values and their real-time performance degrades significantly as their voxel resolution increases. Also, because they rely on voxel hashing, they cannot be easily parallelized and therefore are not directly applicable within modern reactive planning schemes such as Model Predictive Path Integral (MPPI) [7], that typically need to make thousands of queries to the SDF module at each iteration. Additionally, such methods only construct the SDF in observed areas and cannot make predictions in occluded regions.

Real-time neural implicit representations such as iSDF [8] offer a more compact description of the underlying SDF, which can be queried for both values and gradients at an

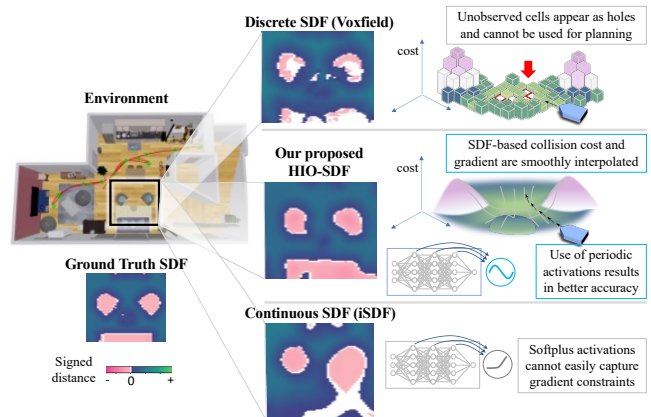


Fig. 1: **Overview of approach:** Online discrete Signed Distance Field (SDF) representations can be very accurate but do not include any information in unobserved areas. Existing continuous SDF representations solve that problem but cannot easily capture fine details in the environment. Our proposed method (HIO-SDF) uses a hierarchy of discrete and continuous algorithms for accurate online 3D SDF reconstruction by a global neural network using periodic activation functions.

arbitrary resolution. However, neural representations also raise technical difficulties. For example, in order to avoid forgetting they typically need to save and replay data from previously visited locations and the learned model tends to over-smooth values resulting in a very coarse representation of the environment. Additionally, due to the absence of a model that can intrinsically capture fine details in the environment, iSDF needs to use heuristics in the definition of the loss function for solving the challenging problem of accurately estimating global Euclidean SDF values from local projective distance estimates for training. Finally, training on a fixed number of sparse data points from the past can lose its efficiency as the space becomes larger.

This work seeks to reconcile these two real-time approaches: we propose a two-level hierarchy which contains a global SDF and a local SDF. At the fine level, the local SDF works by performing direct distance computations of query points to the observed point cloud on a GPU, and informs the global SDF updates. The global SDF is built by combining a coarse voxel-based representation (which prevents forgetting) with the local SDF, in order to fit a frequency-based neural network (SIREN [9]) to the whole environment.

A. Outline of Contributions

This paper proposes Hierarchical Incremental Online SDF (HIO-SDF), a method for constructing a differentiable global Signed Distance Field suitable for real-time application in unknown environments (Fig. 1). The main contributions of our work are as follows:

- 1) Our method uses a mixture of approaches to extract data for the self-supervised training of the global SDF:

Work performed while at Samsung AI Center NY.




Fig. 2: **Overview of our hierarchical incremental online global SDF architecture (HIO-SDF)**: The input to the algorithm is a stream of posed point cloud measurements (dense or sparse) and the output is a global implicit 3D SDF map of the environment that can be used directly for downstream tasks. At each time step, the input is consumed by two different modules that produce SDF data (i.e., tuples of 3D coordinates and associated SDF values) for the supervision of a SIREN [9] representing our global SDF: 1) a discrete global SDF module fusing measurements in the background at a coarse resolution, and 2) a continuous local SDF module that does not need training and performs brute-force distance computations on the GPU between stratified raycasted samples and the currently observed point cloud. Training of the global SIREN is incremental and online, as new sensor measurements are received.

- a) a discrete model for coarse global SDF value lookup without replaying past data, and b) a sensor-based local SDF representation that captures more fine details around the collision boundaries.
- 2) The deep neural network representing the global SDF is constructed using periodic activation functions (SIREN [9]), in order to enforce better accuracy and satisfaction of SDF-related constraints. Unlike past work, the use of SIREN allows us to achieve superior performance with a very simple definition of the loss function.
- 3) Unlike prior work [8], our method offers more flexibility, as it works both with depth sensors and sparse point cloud inputs (e.g., from a 3D LiDAR sensor).

We compare our method against both a discrete SDF representation and a state-of-the-art differentiable SDF model on scenes from two different open-source datasets and a real dataset collected in an office environment. The experimental results show that our method constructs a better underlying SDF representation, even in the presence of sensor noise and odometry uncertainties, both with dense depth and sparse 3D LiDAR point cloud data.

II. RELATED WORK

Although occupancy-based representations have been the mainstream in robotics for decades, SDF-based representations have also been widely employed across different domains. SDF was initially suggested for its accuracy and robustness in surface modeling [1]. Since then, use of SDF for scene reconstruction appear in many previous works [10], [11], [12], [13], [14]. Besides its high accuracy as a surface representation, SDF provides rich spatial information of traversable space and thus it is very useful for collision checks and trajectory optimization [15], [16].

Some of the most successful SDF mapping methods in robotics such as Voxblox [4], FIESTA [5], and Voxfield [6] are based on voxel grids. Voxblox incrementally builds Euclidean Signed Distance Fields (ESDF) from Truncated Signed Distance Fields (TSDF) integrated over time. FIESTA offers improved speed over Voxblox by utilizing occupancy-based wavefront propagation instead of TSDF’s projective distance, but at the cost of the sign of the distance field; it

builds unsigned distance field maps. Voxfield has improved accuracy over both methods by using non-projective TSDF and surface normal estimates from the incoming point cloud. All of the aforementioned methods run in real-time on a single CPU core and have been released as open source, contributing much to the robotics community. However, although conceptually serving as scalar functions, discrete methods are essentially lookup data structures with limited interpolation and extrapolation capabilities, and make it hard to exploit GPU acceleration for massive queries.

Alternative studies have focused on continuous SDF representations. Whereas non-parametric regression methods [17], [18] suffer from increased computational complexity, parametric methods using Deep Neural Networks (DNN) provide a more practical solution in the form of neural implicit functions. One of the earliest methods following this approach is DeepSDF [19], which inspired other studies followed [20], [21], [9]. While most NN-based SDFs are suggested for offline training schemes, there has been prior work on SDF models for incremental training from online observations. The method in [22] was the first incremental SDF NN algorithm using a relay buffer but did not achieve real-time performance. iSDF [8] is the first real-time incremental continuous SDF. DI-Fusion [23] suggests an online method to train SDF using partitioning. Their SDF implementation consists of multiple independent NNs representing local cells, which improves scalability yet sacrifices continuity. Note that these online methods are developed for dense depth sensing and demonstrated in room-scale surface reconstruction applications.

In this study, we seek a solution to incrementally build a continuous SDF for downstream robotic tasks. Prior work has already used SDF representations for offline path refinement [15], learning models for task planning [24], predicting collision clearance in high-dimensional configuration spaces [25], or querying the robot’s distance to the environment during reactive manipulation planning [2], [26], [27]. Our prior work [3] uses an SDF module that computes raw distances on the GPU between query points and the scene’s point cloud for reactive manipulation planning in tabletop scenarios. However, this brute-force approach becomes memory inefficient outside of tabletop scenarios as the point cloud of the scene grows.

III. PROBLEM STATEMENT

A 3D SDF parameterized by θ is defined as an implicit scalar field $f_\theta : \mathbb{R}^3 \mapsto \mathbb{R}$ mapping each 3D location $\mathbf{x} \in \mathbb{R}^3$ to its signed distance $s \in \mathbb{R}$ to the closest surface in the environment, i.e., $f_\theta(\mathbf{x}) = s$. The surface of all objects in the environment is implicitly described by the zero-level set of the SDF, $\mathcal{S} := \{\mathbf{x} \in \mathbb{R}^3 \mid f_\theta(\mathbf{x}) = 0\}$.

One important property of a valid SDF is that its gradient at every location $\mathbf{x} \in \mathbb{R}^3$ needs to have unit norm: $\|\nabla f_\theta(\mathbf{x})\| = 1$. Intuitively, this property gives the SDF distance-like properties; it implies that approaching or moving away from the closest surface point of \mathbf{x} by dx in the direction of $\nabla f_\theta(\mathbf{x})$ should decrease or increase the SDF value by the same amount, respectively. This constraint is known as the *Eikonal equation*. The gradient $\nabla f_\theta(\mathbf{x})$ of the

SDF at a surface point $\mathbf{x} \in \mathcal{S}$ gives the surface normal at this location.

The goal of this work is to construct a continuous and differentiable 3D SDF representation of the environment in real-time from a stream of posed point cloud inputs which is not restricted to dense depth streams unlike prior work [8].

IV. METHOD AND SYSTEM ARCHITECTURE

This section presents HIO-SDF, our method for estimating global SDF values (Fig. 2). The global SDF is represented as a neural network trained in an incremental and online fashion, using data from two separate sources: 1) a discrete global representation running in the background at a coarse resolution, and 2) a brute-force SDF model that computes signed distances of points along rays in the sensor’s local field of view to the currently observed point cloud.

A. Global SDF Network

The global SDF network is a SIREN [9] fully-connected neural network with 4 hidden layers of 256 units each, periodic (sinusoidal) activations and a linear output layer. We choose to use a SIREN because it naturally captures the underlying signal from the environment than networks with traditional activation functions (e.g., ReLUs [8]); the learned weights and biases of the SIREN essentially represent angular frequencies and phase offsets, respectively. An additional property of such networks that make them particularly well-suited for SDF constructions is that the derivative of a SIREN is itself a phase-shifted, well-behaved SIREN, which enables more accurate supervision of the Eikonal constraint that involves the gradient of the network.

Unlike iSDF [8] and other prior work [28], [29], we do not need to apply positional embedding to the 3D input coordinates before feeding them to the network, even though we do not know the scale of the scene ahead of time. The frequency properties of the SIREN directly handle this issue, as discussed above. Additionally, we find that we do not need to feedforward the input to downstream layers [8], or introduce any other heuristics to the network for accelerated learning. These are examples that demonstrate the representation power of the SIREN, which allows us to keep a simple architecture. We initialize each layer of the SIREN by drawing weights from the uniform distribution $\mathcal{U}(-\sqrt{6/n}/\omega_0, \sqrt{6/n}/\omega_0)$ with n the dimension of the input, as described in the original paper [9]. We found that a nominal frequency of $\omega_0 = 10$ works well for our experiments.

B. Coarse Global SDF Data Generation

Even though any discrete global SDF representation from the literature could be used, we choose Voxfield [6] as the backbone method for generating online coarse global SDF data. Voxfield is shown to outperform other prior work in terms of both accuracy and update time.

At each iteration, Voxfield consumes the most recent posed point cloud input and updates an SDF map of the environment, saved in the form of a voxel grid of a user-specified resolution. Considering that Voxfield’s update time and memory consumption increase as the voxel size becomes

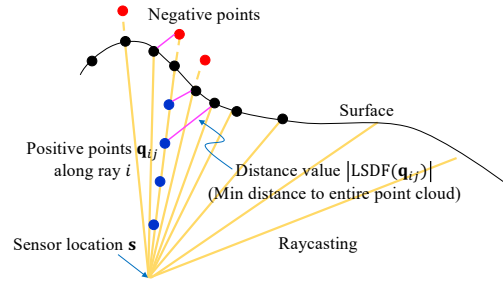


Fig. 3: **Local SDF data generation method:** The local SDF module uses raycasting from the sensor location to find points with positive (blue) and negative (red) SDF values before and after the currently observed point cloud (black), respectively. The absolute local SDF value of each point is the minimum distance (magenta) to the entire local point cloud.

smaller and/or the environment grows, we use a coarse voxel size (10cm). After each Voxfield update, we sample points and their associated SDF values from the voxel grid; specifically, we sample uniformly at random $N = 10000$ points close to the estimated surface (i.e., points with SDF values less than a fixed threshold $\varepsilon = 5\text{cm}$), and $M = 30000$ points away from the surface.

C. Local SDF Data Generation

We complement coarse global SDF data with data generated by a continuous local SDF framework. The local SDF performs direct computation of distances between query points and the currently observed point cloud, in a parallelized fashion on a GPU. This is illustrated in Fig. 3. Specifically, given a local point cloud with P points, described as $\mathcal{P} := \{\mathbf{p}_i \in \mathbb{R}^3\}_{i=1,\dots,P}$, we first sample S points to construct a subset $\mathcal{P}_S \subseteq \mathcal{P}$. Given the current sensor location $\mathbf{s} \in \mathbb{R}^3$, we perform raycasting to find Q stratified samples $\{\mathbf{q}_{ij} \in \mathbb{R}^3\}_{j=1,\dots,Q}$ along the rays connecting \mathbf{s} to each point $\mathbf{p}_i \in \mathcal{P}_S$ up to a maximum distance of $(\|\mathbf{p}_i - \mathbf{s}\| + t)$ behind the point cloud, with $t \in \mathbb{R}$ the *truncation distance*. To each such point \mathbf{q}_{ij} , we assign a sign based on its distance to the sensor \mathbf{s} as follows:

$$\text{sign}(\mathbf{q}_{ij}) := \begin{cases} 1, & \text{if } \|\mathbf{q}_{ij} - \mathbf{s}\| < \|\mathbf{p}_i - \mathbf{s}\| \\ -1, & \text{otherwise} \end{cases} \quad (1)$$

After the computation of these stratified samples with raycasting, we compute the *local SDF value* (LSDF) of each sample \mathbf{q}_{ij} as the direct signed distance to the entire local point cloud \mathcal{P} :

$$\text{LSDF}(\mathbf{q}_{ij}) := \text{sign}(\mathbf{q}_{ij}) \cdot \min_{k \in \{1,\dots,P\}} \|\mathbf{q}_{ij} - \mathbf{p}_k\| \quad (2)$$

and add all pairs $\{(\mathbf{q}_{ij}, \text{LSDF}(\mathbf{q}_{ij})), \forall i, j\}$ to the dataset for the next global SDF training iteration. Even though this direct computation seems expensive, such distance queries between points in \mathbb{R}^3 and the point cloud can be easily parallelized on a GPU, as we demonstrated in our prior work [3]. In our experiments, we choose $S = 1000$, $Q = 20$ (for a maximum of 20000 local samples) and $t = 0.2m$. We typically perform distance queries between stratified samples and points in the local point cloud in less than 5ms. Using a single point cloud measurement, local SDF values of points away from the surface might not accurately reflect their global SDF values. Therefore, we only keep samples with SDF values less than

the truncation distance t . This captures fine details around the surface, i.e., the zero level set of the underlying SDF.

D. Loss Function and Training

Each time a new posed point cloud measurement is received, we run a Voxfield update step and concatenate the coarse global SDF data with the local SDF data. Using this data, consisting of pairs (\mathbf{x}, s) of 3D coordinates and their associated SDF values, we run a training iteration of the global SDF network to update its weights θ . For supervision, we use the following loss function

$$\mathcal{L}_{\text{total}}(\theta) := \lambda_{\text{SDF}} \mathcal{L}_{\text{SDF}}(\theta) + \lambda_{\text{Eikonal}} \mathcal{L}_{\text{Eikonal}}(\theta) \quad (3)$$

with the SDF loss defined as the L_1 norm

$$\mathcal{L}_{\text{SDF}}(\theta) := \mathcal{L}_{\text{SDF}}(f_\theta(\mathbf{x}), s) = |f_\theta(\mathbf{x}) - s| \quad (4)$$

and the Eikonal loss defined as

$$\mathcal{L}_{\text{Eikonal}}(\theta) := \mathcal{L}_{\text{Eikonal}}(\nabla f_\theta(\mathbf{x})) = |||\nabla f_\theta(\mathbf{x})|| - 1| \quad (5)$$

For our experiments, we set $\lambda_{\text{SDF}} = 5.0$, $\lambda_{\text{Eikonal}} = 2.0$, and train for 10 epochs per incremental update, using batch optimization and the Adam optimizer with learning rate 0.0004 and weight decay 0.012.

Our loss function is straightforward to implement and does not require an elaborate SDF loss definition, which often depends on the distance to the surface, the truncation of the Eikonal loss for samples close to the surface, or direct supervision on the direction of SDF gradients. These introduce either additional hyper-parameters or additional pre-processing of the data (e.g., the requirement for surface normals estimation [8], which might not be an option with sparse point cloud data from a 3D LiDAR).

V. RESULTS

In this section, we compare HIO-SDF against two state-of-the-art approaches: iSDF and Voxfield. Note that Voxfield also provides our coarse global SDF estimates, but we will see that our method outperforms Voxfield at 10cm resolution when using a 10cm resolution grid and is even competitive with Voxfield at 5cm resolution. In our experiments, three datasets are used: ReplicaCAD [30], ScanNet [31], and data collected from our office which we call ‘SAIC-NY’. For ReplicaCAD and ScanNet, we use similar scenes as those used in iSDF (`apt_2_nav`, `apt_2_obj`, `apt_2_mnp`, `apt_3_nav` from ReplicaCAD and `scene0030_00`, `scene0004_00` from ScanNet) and follow the same evaluation method. The SAIC-NY dataset includes depth images from an Azure Kinect camera and Velodyne 3D LiDAR scans collected by a Clearpath Ridgeback mobile robot. We use this dataset to compare all methods in a large unmapped space with uncertain odometry (computed online using gMapping [32]), as well as demonstrate HIO-SDF’s flexibility with different sensor inputs.

For the evaluation of HIO-SDF, we play back data in ROS and let the whole pipeline run online, training for 10 epochs per incremental update of the global SDF. The data playback is not paused when evaluating our network, in order to also test the online capabilities of our approach. The reader is referred to the accompanying video submission for more experiments and a demonstration of our algorithm.

A. Qualitative Evaluation

Fig. 4 compares global SDF slices of the three methods at a particular height against ground truth slices in six different environments. For a fair comparison, we evaluate the SDF values from all three approaches using a 3D mesh grid of the same resolution and interpolate Voxfield’s coarse discrete values at a finer resolution. The used height for all environments is close to the floor where most of the geometric details in the environment lie (tables, chairs, etc.). As expected and shown in Fig. 4, HIO-SDF and iSDF can continuously represent the entire space with SDF values, while Voxfield includes areas with nonexistent SDF values. In addition, HIO-SDF is able to better capture detailed information around furniture compared to iSDF, as highlighted in yellow. Fig 5 includes a qualitative comparison of HIO-SDF’s reconstructed 3D meshes against iSDF and Voxfield. The meshes in the figure are generated using the zero level set of SDF values with marching cubes on uniform 3D grids using iSDF’s visualization methods. HIO-SDF generates smooth surfaces that also include finer geometric details over the entire environment compared to iSDF, while the discrete Voxfield cannot generate a complete mesh. Overall, our approach qualitatively outperforms the other two approaches as it has less holes and more detail.

B. Quantitative Evaluation

Fig. 7 includes our quantitative comparisons. We compare HIO-SDF, iSDF and Voxfield against ground truth data over time (as each sequence progresses), by evaluating SDF values and gradients on 25000 points sampled uniformly at random from the visible region at each time instant. For iSDF, we use their released evaluation method; for our method and Voxfield we evaluate the error each time a new measurement is received. Two resolutions of Voxfield grid are evaluated: one at 10cm, which also serves as the backbone coarse global SDF for HIO-SDF, and the other at 5cm that is more accurate yet approaches the limits of Voxfield’s real-time capabilities. For each scene, we run our method and iSDF 10 times to capture the variance but only once for Voxfield since it is a deterministic algorithm. We evaluate all methods using 4 different metrics against ground truth data: 1) the SDF error defined in Eq. 4 using sampled points from the entire visible region (global SDF error), 2) the same SDF error using only local points within a radius of 3m from the sensor position at each time instant (local SDF error), 3) the Eikonal loss in Eq. 5, and 4) the cosine similarity between the predicted gradient $\nabla f_\theta(\mathbf{x})$ and the ground truth SDF gradient \mathbf{g} , defined as

$$\mathcal{L}_{\text{sim}}(\mathbf{x}) := 1 - \frac{\nabla f_\theta(\mathbf{x}) \cdot \mathbf{g}}{||\nabla f_\theta(\mathbf{x})|| \cdot ||\mathbf{g}||}$$

Our quantitative comparisons present some interesting results. First, looking at the global and local SDF errors, our method is more accurate than iSDF and Voxfield at 10cm resolution and is competitive with Voxfield at 5cm resolution, even though it just uses a 10cm-resolution voxel grid as its discrete global SDF backbone. We achieve a mean global SDF error of 5.57cm across all scenes, compared to iSDF’s 10.27cm error and Voxfield’s 7.98cm and 5.48cm

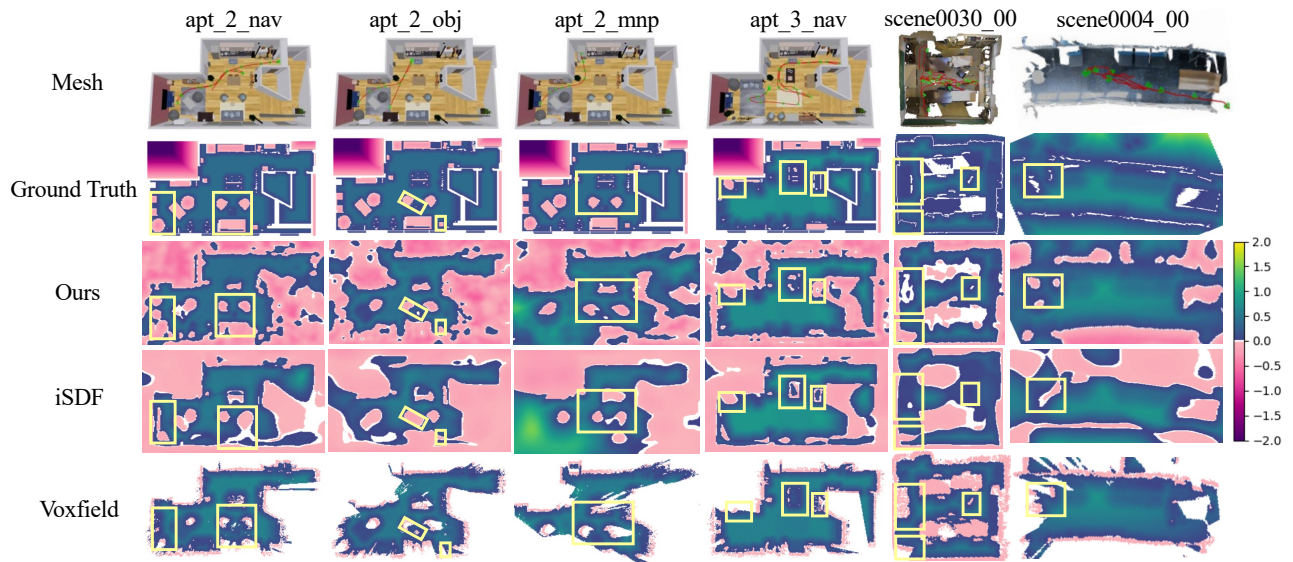


Fig. 4: **Global SDF slice comparison:** The sensor trajectories for each scene are shown in red at the top row with ground truth meshes. SDF slices are computed using the final SDF at the end of each sequence. HIO-SDF is able to represent more detailed SDF information around furniture, compared to both iSDF and Voxfield. For Voxfield, invisible regions by the sensor are shown in white. Details captured by our method even though not captured by the underlying coarse-SDF Voxfield backbone are highlighted in yellow.




Fig. 5: **Mesh reconstruction comparison:** By virtue of using SIREN and utilizing both continuous local and coarse global information, HIO-SDF generates complete and smooth surfaces that also include finer geometric details over the entire environment compared to iSDF and Voxfield.

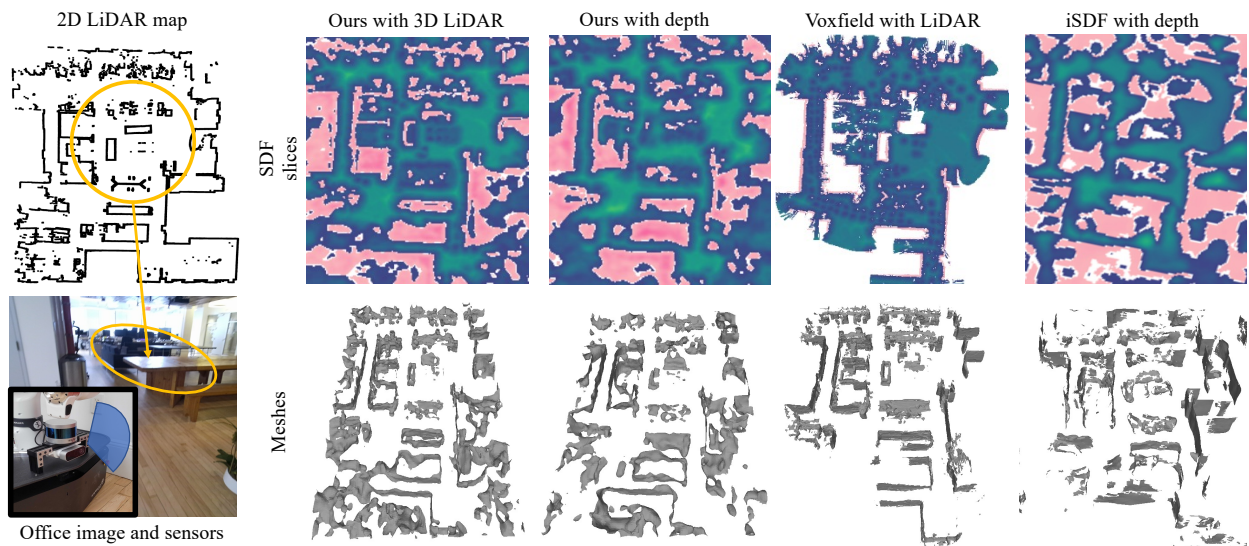


Fig. 6: **Qualitative results in the SAIC-NY office dataset collected with a Clearpath Ridgeback:** Unlike iSDF, HIO-SDF works with both sparse 3D LiDAR data from a Velodyne sensor and dense depth information. Unlike Voxfield, it also attempts to plausibly fill out missing information.




Fig. 7: Quantitative comparison of all methods across 6 different scenes (4 scenes from ReplicaCAD and 2 scenes from ScanNet). The x -axis represents the sequence running time (wall clock) as the simulated sensor navigates each environment.

error at 10cm and 5cm resolution, respectively. Additionally, being trained from this discrete representation, our initial errors are much smaller than iSDF and rapidly converge to their final values, and our measurements have much lower variance across all metrics/scenes. Voxfield errors at both resolutions are much higher than both our method and iSDF in the ScanNet scenes. We believe this happens because ScanNet includes scenes where the measurements exceed the scale of the room, introducing additional noise to the point cloud measurements; in contrast, HIO-SDF and iSDF are capable of filtering out such noise with neural network interpolation. The Eikonal loss error for each method is fairly constant across different sequences, with iSDF performing slightly better than HIO-SDF. Lastly, the value of the synergy between the continuous local and coarse global data is captured in our cosine similarity plots: even though iSDF explicitly supervises this loss to achieve good performance, our method achieves better results (sometimes by a wide margin) in all ReplicaCAD scenes and the loss still decreases or stays constant in ScanNet scenes, implying that the global SIREN SDF is also capable of accurately capturing the direction of the SDF gradient without direct supervision.

C. Evaluation in SAIC-NY Office Dataset

Fig. 6 includes a comparison of SDF reconstruction results in our independently-collected SAIC-NY office dataset. HIO-SDF is able to construct high-quality SDF maps, even in the presence of uncertain odometry measurements. We use this dataset, which represents a large $25m \times 20m$ office space, to also keep track of timings and memory consumption. By

the end of the run, HIO-SDF and Voxfield needed to store a 42.6MB coarse global grid. By comparison, iSDF needed to store 178 depth keyframes of 1.8MB each, accounting for 320.4MB in total. The memory needed for storing the weights of both HIO-SDF’s and iSDF’s network was similar (approximately 1MB). In terms of update times, by the end of the run our method requires around 160ms to perform a Voxfield update, 270ms to extract global SDF data, 38ms to construct the local SDF data, and 29ms per training epoch of the global SIREN, for a total update time of 1sec for 10 training epochs. A forward pass and a backward pass through the SIREN take around 1ms and 9ms respectively on a NVIDIA GeForce RTX 2080 Ti GPU, while querying the final trained model with 100,000 3D coordinates requires only 9ms compared to Voxfield’s 40ms. Overall, evaluations in this dataset demonstrate HIO-SDF’s time- and memory-efficiency, as well as its ability to construct accurate SDFs in large environments extending the room-scale.

VI. CONCLUSION

This paper presents HIO-SDF, a hierarchical, incremental and online global Signed Distance Field model, represented by a SIREN neural network with periodic activation functions. Our approach combines the benefits of coarse, discrete global SDF representations and locally-accurate, continuous SDF models. In several runs through a variety of datasets and scenes, we demonstrate that HIO-SDF is more accurate and efficient than other state-of-the-art representations. Future work will seek to include applications of HIO-SDF in mobile manipulation for reactive 3D planning and motion control.

REFERENCES

- [1] B. Curless and M. Levoy, "A volumetric method for building complex models from range images," in *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*. ACM, 1996, pp. 303–312.
- [2] M. Koptev, N. Figueroa, and A. Billard, "Neural joint space implicit signed distance functions for reactive robot manipulator control," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 480–487, 2023.
- [3] V. Vasilopoulos, S. Garg, P. Piacenza, J. Huh, and V. Isler, "RAMP: Hierarchical Reactive Motion Planning for Manipulation Tasks Using Implicit Signed Distance Functions," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2023.
- [4] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3D Euclidean Signed Distance Fields for on-board MAV planning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 1366–1373.
- [5] L. Han, F. Gao, B. Zhou, and S. Shen, "FIESTA: Fast incremental euclidean distance fields for online motion planning of aerial robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019, pp. 4423–4430.
- [6] Y. Pan, Y. Kompis, L. Bartolomei, R. Mascaro, C. Stachniss, and M. Chli, "Voxfield: Non-Projective Signed Distance Fields for Online Planning and 3D Reconstruction," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022, pp. 5331–5338.
- [7] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Information-theoretic model predictive control: Theory and applications to autonomous driving," *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1603–1622, 2018.
- [8] 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," in *Robotics: Science and Systems*, 2022.
- [9] V. Sitzmann, J. N. Martel, A. W. Bergman, D. B. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [10] J. C. Carr, R. K. Beatson, J. B. Cherrie, T. J. Mitchell, W. R. Fright, B. C. McCallum, and T. R. Evans, "Reconstruction and representation of 3D objects with radial basis functions," in *SIGGRAPH*, 2001, pp. 67–76.
- [11] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon, "KinectFusion: Real-time dense surface mapping and tracking," in *IEEE International Symposium on Mixed and Augmented Reality*, 2011, pp. 127–136.
- [12] E. Bylow, J. Sturm, C. Kerl, F. Kahl, and D. Cremers, "Real-time camera tracking and 3D reconstruction using signed distance functions," in *Robotics: Science and Systems*, vol. 2, 2013.
- [13] A. Zeng, S. Song, M. Nießner, M. Fisher, J. Xiao, and T. Funkhouser, "3DMatch: Learning local geometric descriptors from RGB-D reconstructions," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017.
- [14] L. Yariv, P. Hedman, C. Reiser, D. Verbin, P. P. Srinivasan, R. Szeliski, J. T. Barron, and B. Mildenhall, "BakedSDF: Meshing Neural SDFs for Real-Time View Synthesis," *arXiv pre-print: 2302.14859*, 2023.
- [15] N. Ratliff *et al.*, "CHOMP: Gradient optimization techniques for efficient motion planning," in *IEEE International Conference on Robotics and Automation*, 2009, pp. 489–494.
- [16] H. Oleynikova, A. Millane, Z. Taylor, E. Galceran, J. Nieto, and R. Siegwart, "Signed distance fields: A natural representation for both mapping and planning," in *RSS 2016 Workshop: Geometry and Beyond-Representations, Physics, and Scene Understanding for Robotics*, 2016.
- [17] B. Lee, C. Zhang, Z. Huang, and D. D. Lee, "Online continuous mapping using gaussian process implicit surfaces," in *IEEE International Conference on Robotics and Automation*, 2019.
- [18] L. Wu, K. M. B. Lee, C. Le Gentil, and T. Vidal-Calleja, "Log-GPIS-MOP: A unified representation for mapping, odometry, and planning," *IEEE Transactions on Robotics*, 2023.
- [19] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "DeepSDF: Learning continuous signed distance functions for shape representation," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2019, pp. 165–174.
- [20] A. Gropp, L. Yariv, N. Haim, M. Atzmon, and Y. Lipman, "Implicit geometric regularization for learning shapes," in *International Conference on Machine Learning*, 2020, pp. 3569–3579.
- [21] J. Chibane, A. Mir, and G. Pons-Moll, "Neural unsigned distance fields for implicit function learning," in *Advances in Neural Information Processing Systems (NeurIPS)*, December 2020.
- [22] Z. Yan, Y. Tian, X. Shi, P. Guo, P. Wang, and H. Zha, "Continual neural mapping: Learning an implicit scene representation from sequential observations," in *IEEE/CVF International Conference on Computer Vision*, 2021.
- [23] J. Huang, S.-S. Huang, H. Song, and S.-M. Hu, "Di-fusion: Online implicit 3D reconstruction with deep priors," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [24] D. Driess, J.-S. Ha, M. Toussaint, and R. Tedrake, "Learning models as functionals of signed-distance fields for manipulation planning," in *Conference on Robot Learning*, 2021.
- [25] J. Chase Kew, B. Ichter, M. Bandari, T.-W. E. Lee, and A. Faust, "Neural collision clearance estimator for batched motion planning," in *Algorithmic Foundations of Robotics XIV*. Springer International Publishing, 2021, pp. 73–89.
- [26] P. Liu, K. Zhang, D. Tateo, S. Jauhari, J. Peters, and G. Chalvatzaki, "Regularized deep signed distance fields for reactive motion generation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022, pp. 6673–6680.
- [27] Y. Li, Y. Zhang, A. Razmjoo, and S. Calinon, "Learning robot geometry as distance fields: Applications to whole-body manipulation," *arXiv pre-print: 2307.00533*, 2023.
- [28] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing scenes as neural radiance fields for view synthesis," in *European Conference on Computer Vision*, 2020.
- [29] J. T. Barron, B. Mildenhall, D. Verbin, P. P. Srinivasan, and P. Hedman, "Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields," *arXiv pre-print: 2111.12077*, 2021.
- [30] A. Szot *et al.*, "Habitat 2.0: Training Home Assistants to Rearrange their Habitat," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [31] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, "ScanNet: Richly-annotated 3D Reconstructions of Indoor Scenes," in *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [32] G. Grisetti, C. Stachniss, , and W. Burgard, "Improved techniques for grid mapping with rao-blackwellized particle filters," *IEEE Transactions on Robotics*, vol. 23, pp. 34–46, 2007.