

SLIM-VDB: A Real-Time 3D Probabilistic Semantic Mapping Framework

Anja Sheppard¹, Graduate Student Member, IEEE, Parker Ewen¹, Member, IEEE, Joey Wilson¹, Member, IEEE, Advait V. Sethuraman¹, Benard Adewole¹, Anran Li¹, Yuzhen Chen, Ram Vasudevan¹, Senior Member, IEEE, and Katherine A. Skinner¹, Member, IEEE

Abstract—This letter introduces SLIM-VDB, a new lightweight semantic mapping system with probabilistic semantic fusion for closed-set or open-set dictionaries. Advances in data structures from the computer graphics community, such as OpenVDB, have demonstrated significantly improved computational and memory efficiency in volumetric scene representation. Although OpenVDB has been used for geometric mapping in robotics applications, semantic mapping for scene understanding with OpenVDB remains unexplored. In addition, existing semantic mapping systems lack support for integrating both fixed-category and open-language label predictions within a single framework. In this letter, we propose a novel 3D semantic mapping system that leverages the OpenVDB data structure and integrates a unified Bayesian update framework for both closed- and open-set semantic fusion. Our proposed framework, SLIM-VDB, achieves significant reduction in both memory and integration times compared to current state-of-the-art semantic mapping approaches, while maintaining comparable mapping accuracy.

Index Terms—Mapping, semantic scene understanding, probabilistic inference.

I. INTRODUCTION

ROBOTS require accurate world maps with scene understanding upon which to condition their actions. Recent mapping methods represent both the geometry [1] and semantics [2], [3] of the robot’s surroundings, enabling scene understanding and complex task completion. A popular approach is to represent a scene using a Truncated Signed Distance Function (TSDF) [4], which represents the contact surface of the environment and results in realistic dense reconstructions.

However, existing approaches that build TSDF maps have large memory requirements and struggle to run in real time [5]. Furthermore, non-probabilistic semantic mapping methods are not robust to inconsistencies or “flickering” in labels output by semantic segmentation networks [6]. Even as closed-set

Received 25 May 2025; accepted 6 December 2025. Date of publication 12 January 2026; date of current version 20 January 2026. This article was recommended for publication by Associate Editor J. Luis Sanchez-Lopez and Editor J. Civera upon evaluation of the reviewers’ comments. This work was supported in part by NSF under Grant DGE 224114 and in part by AFOSR MURI under Grant FA9550-23-1-0400. (Corresponding author: Anja Sheppard.)

The authors are with the University of Michigan Robotics Department, Ann Arbor, MI USA (e-mail: anjashep@umich.edu).

An open-source C++ codebase with a Python interface is available at <https://github.com/umfieldrobotics/slim-vdb>.

This article has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2026.3652875>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2026.3652875

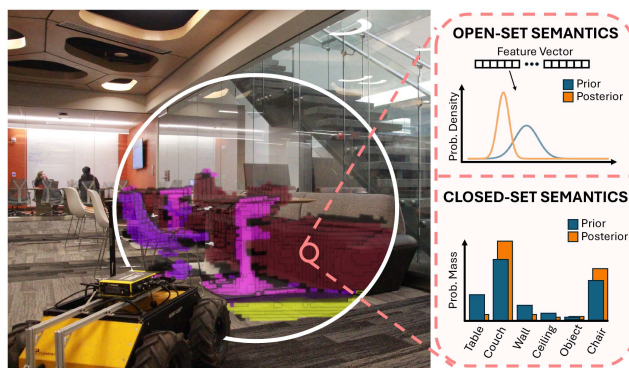


Fig. 1. Example of the real-time semantic mapping method proposed in this letter. SLIM-VDB is capable of using either open- or closed-set semantics and leverages a Bayesian approach for semantic fusion. This approach requires significantly less memory and computation time than existing state-of-the-art semantic mapping methods.

semantics remain useful for simplified world understanding, the growing use of open-language representations [7] highlights the need for mapping frameworks that can handle both paradigms within a unified formulation. These challenges motivate the development of a compute-efficient probabilistic mapping system capable of handling both closed and open-set semantics.

The computer graphics community has developed the OpenVDB [8] data structure for efficient volumetric data storage. VDBFusion [1] has recently adapted OpenVDB into a robotic mapping tool, which demonstrates state-of-the-art performance in computation speed and memory consumption for mapping, while maintaining high geometric accuracy. However, VDBFusion lacks semantic mapping capabilities.

Prior work has used Bayesian updates to deal with either closed-set [9] or open-set [7] semantic uncertainty in order to perform semantic mapping. Unfortunately, these approaches are computationally demanding and are only applicable to either open-set or closed-set semantics, not both.

To address these challenges, we introduce SLIM-VDB, the Semantic Lightweight Implicit Mapping system. SLIM-VDB leverages the OpenVDB data structure and introduces a unified Bayesian semantic fusion update step to enable real-time, memory-efficient semantic mapping for either closed-set or open-set semantics. As far as we are aware, this is the first work to use an OpenVDB backend for semantic mapping and the first work to present a unified probabilistic semantic fusion update

for either open- or closed-set class labels. We also present an open-source C++ codebase with a Python interface to the greater robotics community. A system overview is shown in Fig. 1.

In summary, the contributions of this letter are as follows:

- A novel framework that builds on OpenVDB to enable lightweight, memory-efficient semantic mapping.
- A unified Bayesian inference framework that enables either closed-set or open-set semantic estimation.
- An open-source C++ library with a Python interface for easy integration with robotics applications.

II. RELATED WORKS

SLIM-VDB builds off of advances in real-time TSDF mapping and recurrent Bayesian approaches that fuse neural network predictions into a probabilistic map representation. In this section, we discuss related work on geometric and semantic mapping.

A. Geometric Mapping

OctoMap [10] was the first major step towards a real-time, dynamic, multi-resolution, and probabilistic 3D mapping framework, and it has been the standard in the mapping community for some time. The emergence of implicit surface representations with TSDFs through Voxblox [11] outperformed OctoMap in update speed and reconstruction quality. Voxblox was extended to include instance and semantic information, and Voxfield Panmap [5], [12] has continued improvement in panoptic TSDF voxel mapping. A recent work [1] proposed an OpenVDB-based [8] lightweight TSDF mapping framework that outperforms Voxblox in runtime and memory efficiency. This is a promising new direction for efficient volumetric mapping, but there are no current approaches that allow for the incorporation of semantics.

Neural Radiance Fields and 3D Gaussian Splatting [13] have demonstrated high-quality dense reconstruction but are typically very computationally and memory intensive. Recent work [14] has begun to make these methods operate in real-time, however, voxel-based mapping frameworks remain very useful for mobile systems that have limited computational resources.

B. Semantic Mapping

Maps can store more than just geometric information, such as semantic categories obtained from neural networks. While early approaches to semantic mapping employed voting schemes or heuristics [15], many modern approaches employ learning-based semantic fusion techniques. Semantic Neural Implicit SLAM (SNI-SLAM) [2] proposes a real-time semantic neural implicit representation that correlates appearance, geometry, and semantic features using cross-attention. However, the system is only able to run in real-time on professional server grade GPUs, which makes it ill-suited for mobile robot platforms.

Probabilistic approaches to semantic fusion such as Convolutional Bayesian Kernel Inference (ConvBKI) [16] and SEE-CSOM [17] recurrently incorporate semantic predictions into maps through Bayesian inference. One benefit of Bayesian inference is the quantification of uncertainty, which can be

leveraged for downstream applications such as active perception. By adopting a conjugate prior over a parametric likelihood function that models the distribution of the data, the expectation and variance of voxel-wise semantics can be monitored and updated over time as new measurements are obtained [3], [9], [16], [18]. However, closed-set semantic maps are limited by a fixed number of categories, which may not be suitable for more complex environments or out-of-distribution data.

Recently, the deep learning community has developed Vision-Language Models (VLMs), which generate predictions in a latent space shared with language models, enabling open-dictionary language-based queries [19]. Instead of requiring a fixed set of categories at model selection, open-dictionary segmentation models enable categorization at inference, which allows for reasoning about synonymous categories or obscure objects. LatentBKI [7] extends open-dictionary semantics to mapping by adopting a conjugate prior over a likelihood function that describes the distribution of the feature embeddings. However, it does not run in real-time, and the map size is greatly limited by GPU memory.

C. Openvdb

OpenVDB [8] was developed at DreamWorks Animation to address a critical need for efficient volumetric mapping of animated character assets. The key strengths of OpenVDB lie in its $\mathcal{O}(1)$ lookup, insertion, and delete times, along with a low memory footprint due to its use of a B+ tree to limit the number of levels in the tree. OpenVDB was a pioneer in the use of compact “level set” data structures, which is otherwise known as signed distance fields to the robotics and computer vision community.

OpenVDB has begun to be repurposed for highly efficient volumetric map data storage for mobile robots. The use of OpenVDB led to various efficiency improvements in point insertion and raycasting [22], [23], distributed systems [24], and 3D reconstruction quality [25]. VDBFusion [1] is an open-source library that takes advantage of OpenVDB’s robust C++ implementation and is capable of integrating LiDAR data at 20 frames per second into a map, outperforming Voxblox and OctoMap. However, VDBFusion does not support the incorporation of semantic information into the map, nor does it utilize NanoVDB [21] – a recent GPU-based improvement to OpenVDB – for real-time rendering. The only work that has utilized NanoVDB for map rendering, NanoMap [26], also does not incorporate semantics.

In this work, we develop a GPU-accelerated OpenVDB-based semantic mapping framework for robotics applications that runs in real-time and with a lower memory footprint than existing state-of-the-art approaches. Our framework uses a unified Bayesian probabilistic update paradigm for semantic label fusion that allows for the use of either closed-set [9] or open-set [7] semantic information.

III. METHODOLOGY

This section presents the proposed semantic mapping framework, illustrated in Fig. 2. First, semantic information, either using open- or closed-set semantics, is computed for images collected using a robot’s onboard RGB-D camera or LiDAR. These

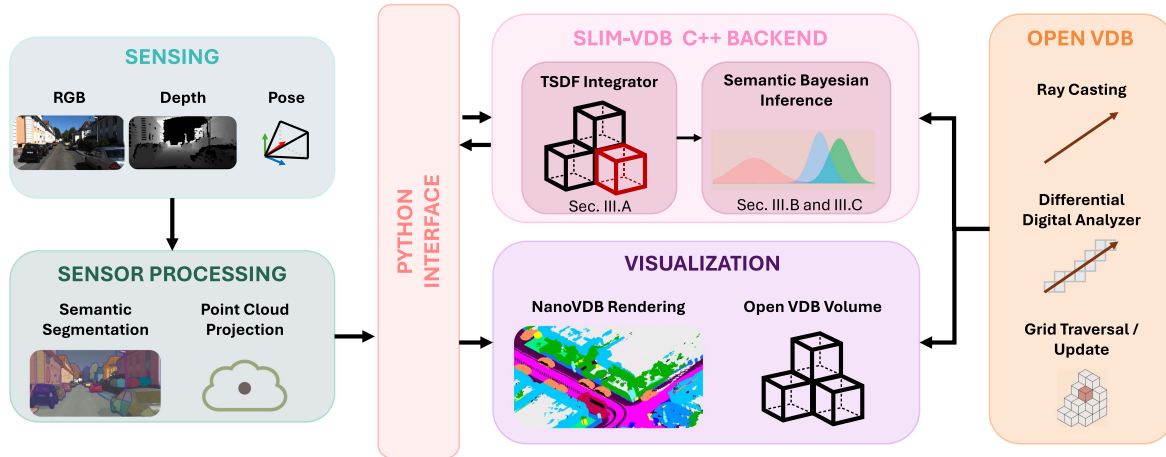


Fig. 2. Overview of the SLIM-VDB mapping system. The input dataset can be either a pointcloud or RGB-D images. Semantic segmentation can be performed with any off-the-shelf network, and we use Semantic Segment Anything [20]. The main integration loop conducts raycasting and TSDF updates (III-A) as well as a probabilistic update for either closed-set (III-C) or open-set (III-D) semantics. Finally, the rendering pipeline (III-E) utilizes NanoVDB [21] for GPU-accelerated viewing of the map.

predictions are projected using the depth image into a semantic point cloud. The geometric information of this point cloud is used to update the TSDF map as described in Section III-A. Next, the global semantic information is updated using Bayesian inference for either closed-set (Section III-C) or open-set (Section III-D) semantics. Lastly, the rendering pipeline is presented to visualize the semantic TSDF in Section III-E.

Vectors, written as columns, are typeset in bold, sets are typeset in uppercase calligraphic font and random variables are uppercase. The element i of a vector \mathbf{x} is denoted as x_i . Let p denote a normalized probability mass or density function. We use $*$ to indicate operations at the voxel level.

A. Surface Reconstruction

The proposed semantic mapping pipeline stores the geometric information of a scene in voxels using the OpenVDB library. This library enables fast and memory-efficient storage and querying via a sparse, hierarchical data structure that is used to store the signed distance, weight, and semantic parameters of voxels close to a surface.

At each frame, a semantically segmented image is projected to 3D using a corresponding depth image. This semantic point cloud is used to sequentially update the voxel map. Raycasting, which determines which voxels correspond to free space and which must be updated, is the critical and most computationally costly operation of this process. We adopt [1]’s implementation of raycasting, which employs OpenVDB’s Differential Digital Analyzer [8]. Motivated by [4], the raycasting method considers only voxels in a configurable “truncated” region around each point, minimizing the integration step computational cost.

Following [4] for volumetric integration with a signed distance function representation, the truncated signed distance and voxel weight values are incrementally updated as new point clouds are integrated into the map. That is, for voxel \mathbf{x}_* at time t , the cumulative truncated signed distance, $D_t(\mathbf{x}_*)$, and cumulative weight, $W_t(\mathbf{x}_*)$, of the voxel are computed using the

signed distance from the latest sensor measurement, $d_t(\mathbf{x}_*)$, and the weighting function, $\tilde{w}(\mathbf{x}_*)$:

$$D_t(\mathbf{x}_*) = \frac{W_{t-1}(\mathbf{x}_*) D_{t-1}(\mathbf{x}_*) + \tilde{w}(\mathbf{x}_*) d_t(\mathbf{x}_*)}{W_t(\mathbf{x}_*)}, \quad (1)$$

$$W_t(\mathbf{x}_*) = W_{t-1}(\mathbf{x}_*) + \tilde{w}(\mathbf{x}_*). \quad (2)$$

A popular choice for a weighting function is $\tilde{w}(\mathbf{x}_*) = 1$.

B. Unified Bayesian Semantic Fusion

In our unified approach to closed or open-set semantic fusion, we utilize a closed-form conjugate pair recursive Bayesian update to estimate the semantic class at each voxel. In probability theory, conjugacy refers to a pair of distributions where, given a prior distribution and a likelihood function, the posterior belongs to the same family of distributions as the prior [27], [28]. This results in a tractable solution for the posterior where only the hyperparameters of the prior need to be updated to reflect the new information from the measurement. Due to the differing likelihood distributions, the Dirichlet-Categorical conjugacy is used for closed-set semantics and the Normal Inverse Gamma-Normal conjugacy is used for open-set semantics.

C. Closed-Set Semantic Fusion

We use the Dirichlet-Categorical conjugate pair for closed-set semantic fusion as also shown in [9], [18]. Let $z \in \{1, \dots, k\}$ denote a discrete Categorical random variable in the predefined sample space with the probability mass function:

$$p(z = i | \boldsymbol{\theta}) = \text{Cat}(z = i | \boldsymbol{\theta}) = \theta_i, \quad (3)$$

where $i \in \{1, 2, \dots, k\}$, $\boldsymbol{\theta} \in [0, 1]^k$, and $\sum_{i=1}^k \theta_i = 1$. Given the parameter $\boldsymbol{\theta}$, this represents the probability that a sample z belongs to semantic class i .

The conjugate prior for the Categorical distribution is the Dirichlet distribution. The Dirichlet distribution is the multivariate generalization of the beta distribution that is parameterized

by a vector $\alpha \in \mathbb{R}_{\geq 0}^k$ and has the probability density:

$$p(\theta | \alpha) = \text{Dir}(\theta | \alpha) = \frac{\Gamma\left(\sum_{j=1}^k \alpha_j\right)}{\prod_{j=1}^k \Gamma(\alpha_j)} \prod_{j=1}^k \theta_j^{\alpha_j - 1}, \quad (4)$$

where $\Gamma(\alpha_j)$ is the Gamma function. The conjugacy of these two distributions is formalized in the following theorem [9]:

Theorem 1: Consider a set of semantic observations $\mathcal{Z} = \{z_1, \dots, z_n\}$, each independently drawn from a Categorical distribution $p(\mathcal{Z} | \theta)$. Let the prior be a Dirichlet distribution, $p(\theta | \alpha)$. Then, the posterior of θ given the observations \mathcal{Z} and parameters α is also a Dirichlet distribution:

$$p(\theta | \mathcal{Z}, \alpha) = \text{Dir}(\theta | \tilde{\alpha}), \quad (5)$$

where the updated Dirichlet parameters are given by:

$$\tilde{\alpha}_i = \alpha_i + \sum_{z_j \in \mathcal{Z}} \mathbb{1}\{z_j = i\}, \quad (6)$$

and $\mathbb{1}\{z_j = i\}$ is an indicator function that equals 1 if observation z_j belongs to class i , and 0 otherwise.

Leveraging Theorem 1 enables us to take a Bayesian approach to integrating multiple semantic predictions at one voxel over time. At each timestep t , a set of 3D points and their corresponding closed-set semantic predictions, $\mathcal{P}_C = \{(\mathbf{x}_j, Z_j)\}_{j=1}^N$ where $\mathbf{x}_j \in \mathbb{R}^3$ and $Z_j \in \{1, \dots, k\}$ is a discrete random variable, is derived from either an RGB-D image or a point cloud. This semantic prediction is typically represented as a one-hot vector output of a segmentation network. At each voxel, Bayes' Rule is recurrently applied to compute the Dirichlet posterior distribution using (6).

In general, computing the predictive posterior involves marginalizing over θ , which can be intractable. However, due to the Dirichlet-Categorical conjugate pair there exists a closed-form solution [9], [27]:

$$p(z = i | \tilde{\alpha}) = \frac{\tilde{\alpha}_i}{\sum_{j=1}^k \tilde{\alpha}_j}. \quad (7)$$

This defines the posterior predictive distribution over the closed-set semantic class. In practice, it is sufficient to track the Dirichlet parameter α to compute these predictions. We will explore in the next section how a similar Bayesian framework can also handle open-set semantics from VLMs.

D. Open-Set Semantic Fusion

Motivated by the use of conjugate pairs for Bayesian inference from the previous section, we use the Normal Inverse Gamma and Normal conjugate pair for open-set semantic fusion as introduced in [7].

Let $\mathbf{Z} \in \mathbb{R}^l$ denote a multivariate, continuous Gaussian random variable with independent elements, where each element has the probability density:

$$p(Z_i | \mu_i, \sigma_i) = \mathcal{N}(Z_i | \mu_i, \sigma_i^2). \quad (8)$$

The Normal Inverse Gamma distribution is a multivariate continuous probability distribution parameterized by four variables: m , λ , ν , and β . Together, these parameters define a probability

density for the mean, μ , and variance, σ^2 , of the Normal distribution such that:

$$p(\mu_i, \sigma_i^2 | m_i, \lambda_i, \nu_i, \beta_i) = \mathcal{N}(\mu_i | m, \sigma_i^2 / \lambda_i) \quad (9)$$

$$\cdot \Gamma^{-1}(\sigma_i^2 | \nu_i, \beta_i), \quad (10)$$

where Γ^{-1} is the inverse gamma function. The Normal Inverse Gamma distribution is the conjugate pair to the Normal distribution, as is formalized in the following theorem:

Theorem 2: Let $\mathcal{Z}_i = \{Z_{i,1}, \dots, Z_{i,n}\}$ be a set of samples drawn from the same Normal distribution, $p(Z_{i,j} | \mu_i, \sigma_i^2)$, with unknown mean and variance and let the prior for μ_i, σ_i^2 be a Normal Inverse Gamma distribution, $p(\mu_i, \sigma_i^2 | m_i, \lambda_i, \nu_i, \beta_i)$. The posterior computed using Bayes' theorem, $p(\mu_i, \sigma_i^2 | \mathcal{Z}_i, m_i, \lambda_i, \nu_i, \beta_i)$, is also a Normal Inverse Gamma distribution. The closed form equations to compute the parameters are as follows, where $|\mathcal{Z}_i|$ is the cardinality of \mathcal{Z}_i and \bar{Z}_i is the mean of \mathcal{Z}_i :

$$\tilde{m}_i = \frac{\lambda_i m_i + |\mathcal{Z}_i| \bar{Z}_i}{\tilde{\lambda}_i}, \quad (11)$$

$$\tilde{\lambda}_i = \lambda_i + |\mathcal{Z}_i|, \quad (12)$$

$$\tilde{\nu}_i = \nu_i + \frac{|\mathcal{Z}_i|}{2}, \quad (13)$$

$$\tilde{\beta}_i = \beta_i + \frac{1}{2} \sum_{z_{i,j} \in \mathcal{Z}_i} (z_{i,j} - \bar{Z}_i)^2 + \frac{\lambda_i |\mathcal{Z}_i| (\bar{Z}_i - m_i)^2}{\tilde{\lambda}_i}. \quad (14)$$

For open-set semantics, we model the measurement likelihood in the VLM feature space as a multivariate Gaussian random variable with i.i.d. elements. Given an RGB-D image and a VLM, we derive a set of 3D points and their corresponding open-set semantic predictions, $\mathcal{P}_O = \{(\mathbf{x}_i, \mathbf{Z}_i)\}_{i=1}^N$ where $\mathbf{x}_i \in \mathbb{R}^3$ is the 3D point and $\mathbf{Z}_i \in \mathbb{R}^l$ is the aforementioned multivariate continuous Gaussian random variable. Notably, the mean and variance of each element of \mathbf{Z}_i is unknown. This enables us to represent the feature embedding as a vector of l 1-dimensional Gaussian random variables, each with an unknown mean and variance.

Following Theorem 2, we utilize the conjugacy between the Normal distribution and the Normal Inverse Gamma to tractably compute the posterior for each element of \mathbf{Z}_i . This is done through the parameter updates defined in (11)–(14). Theorem 2 then enables the estimate of the open-set semantics via the predictive posterior for the Gaussian random variable Z_i given previous semantic measurements:

$$p(Z_i | \mathcal{Z}_i, m_i, \lambda_i, \nu_i, \beta_i) = t_{2\nu_i} \left(Z_i | m_i, \frac{\beta_i(\lambda_i + 1)}{\lambda_i \nu_i} \right) \quad (15)$$

where $t_{2\nu_i}$ is the Student- t distribution [28]. The expectation of the Student- t distribution is $\mathbb{E}(Z_i) = m_i$ when $\lambda_i > 1$, and a covariance of $\text{Cov}(Z_i) = \frac{\lambda_i}{\lambda_i - 2} \left(\frac{\beta_i(\lambda_i + 1)}{\lambda_i \nu_i} \right)$ for $\lambda_i > 2$. This results in a distribution over the feature vectors, and we then apply a softmax in order to get absolute class probabilities.

In the implementation, we store the mean vector $\mathbf{m}_* = (\tilde{m}_1, \tilde{m}_2, \dots, \tilde{m}_l)$ at a voxel level, where l is the size of the feature vector. We simplify the recursive estimation of the

TABLE I
 COMPARISON OF AVERAGE FPS AND CPU+GPU RUNTIME MEMORY USAGE FOR CLOSED- (^C) AND OPEN-SET (^O) SEMANTIC MAPPING, INCLUDING ± ONE STD. DEVIATION. SLIM-VDB* DENOTES A HEURISTIC VERSION OF OUR METHOD WITH A NON-BAYESIAN SEMANTIC UPDATE. **BOLD** IS BEST

Method	SemanticKITTI [29]		SceneNet [30]		
	FPS ↑	Memory (GB) ↓	FPS ↑	Memory (GB) ↓	
Non-Bayesian	SNI SLAM [17]	0.13 ± 0.02	19.43 ± 0.98	0.27 ± 0.01	19.43 ± 0.07
	VoxField [5]	4.43 ± 0.72	3.06 ± 0.63	2.49 ± 1.67	1.05 ± 0.45
	SLIM-VDB* (Ours)	6.25 ± 1.12	0.67 ± 0.05	16.55 ± 2.87	0.51 ± 0.01
Bayesian	ConvBKI [7]	0.53 ± 0.26	27.73 ± 0.20	1.10 ± 1.00	14.32 ± 7.02
	SEE-CSOM [17]	0.10 ± 0.01	10.98 ± 0.35	2.67 ± 8.00	5.19 ± 1.30
	SLIM-VDB ^C (Ours)	1.69 ± 0.69	2.60 ± 0.80	10.84 ± 1.94	1.08 ± 0.16
	LatentBKI [7]	-	-	1.67 ± 0.36	27.68 ± 1.37
	SLIM-VDB ^O (Ours)	-	-	1.76 ± 0.11	3.49 ± 0.68




Fig. 3. Visual runtime comparison on SceneNet and KITTI scenes for both Bayesian and non-Bayesian semantic mapping frameworks.

posterior by using the cumulative TSDF weight $W_t(\mathbf{x}_*)$ for λ_* and $2\nu_*$. Additionally, due to our independence assumption for each element of the feature vector, we only track the diagonal of the covariance matrix β . During inference, the highest class probability at each voxel resulting from the posterior predictive is thresholded against $T_p = 0.1$, in order to filter out uncertain class predictions.

E. Rendering

The truncated signed distance values, weight, and semantic parameters corresponding to either open- or closed-set semantics are stored in voxels contained inside the OpenVDB data structure. Rather than constructing a voxel mesh and publishing it to a visualizer such as Rviz, we use the native VDB rendering tools to provide real-time views of the scene, which render the grid from a given camera viewpoint. Rendering is traditionally a computationally intensive operation, so we utilize NanoVDB [21], a GPU-optimized version of OpenVDB, for faster, parallelized rendering.

IV. EXPERIMENTS

We validate the proposed SLIM-VDB algorithm on both simulation and real world benchmarks, demonstrating a significant reduction in computation time and memory consumption while maintaining comparable accuracy to state-of-the-art semantic mapping approaches. All experiments except those in subsection

IV-E were run on an Intel Xeon W-2245 CPU @ 3.90 GHz with an NVIDIA GeForce 32 GB RTX 3090 GPU. All experiments on KITTI [29] map at a 10 cm per voxel resolution, closed-set experiments on SceneNet [30] map at a 5 cm per voxel, and open-set experiments on SceneNet map at 10 cm per voxel. Evaluation is run in C++.

A. Closed-Set Semantic Mapping

We compare SLIM-VDB against five state-of-the-art semantic mapping approaches: ConvBKI [16], SEE-CSOM [17], Voxfield Panmap [5], SNI-SLAM [2], and LatentBKI [7]. We evaluate runtime and memory consumption given ground truth semantic labels. For semantic prediction accuracy, we use segmentation results from the Semantic Segment Anything [20] network for SceneNet, and PolarSeg [31] for SemanticKITTI. PolarSeg is finetuned on SemanticKITTI scenes 0–5, and we evaluate on the first 100 frames of scenes 6–10. We randomly select 12 full SceneNet scenes for evaluation from two different trajectories.

SemanticKITTI [29] and SceneNet [30] are used to evaluate SLIM-VDB on large-scale scenes and on indoor simulated scenes, respectively. ConvBKI and SEE-CSOM both include a Bayesian update for closed-set semantic Dirichlet parameters (see Section III-C). In order to evaluate ConvBKI we do not use free space sampling, since the global mapping approach does not scale with voxels introduced from free space. We

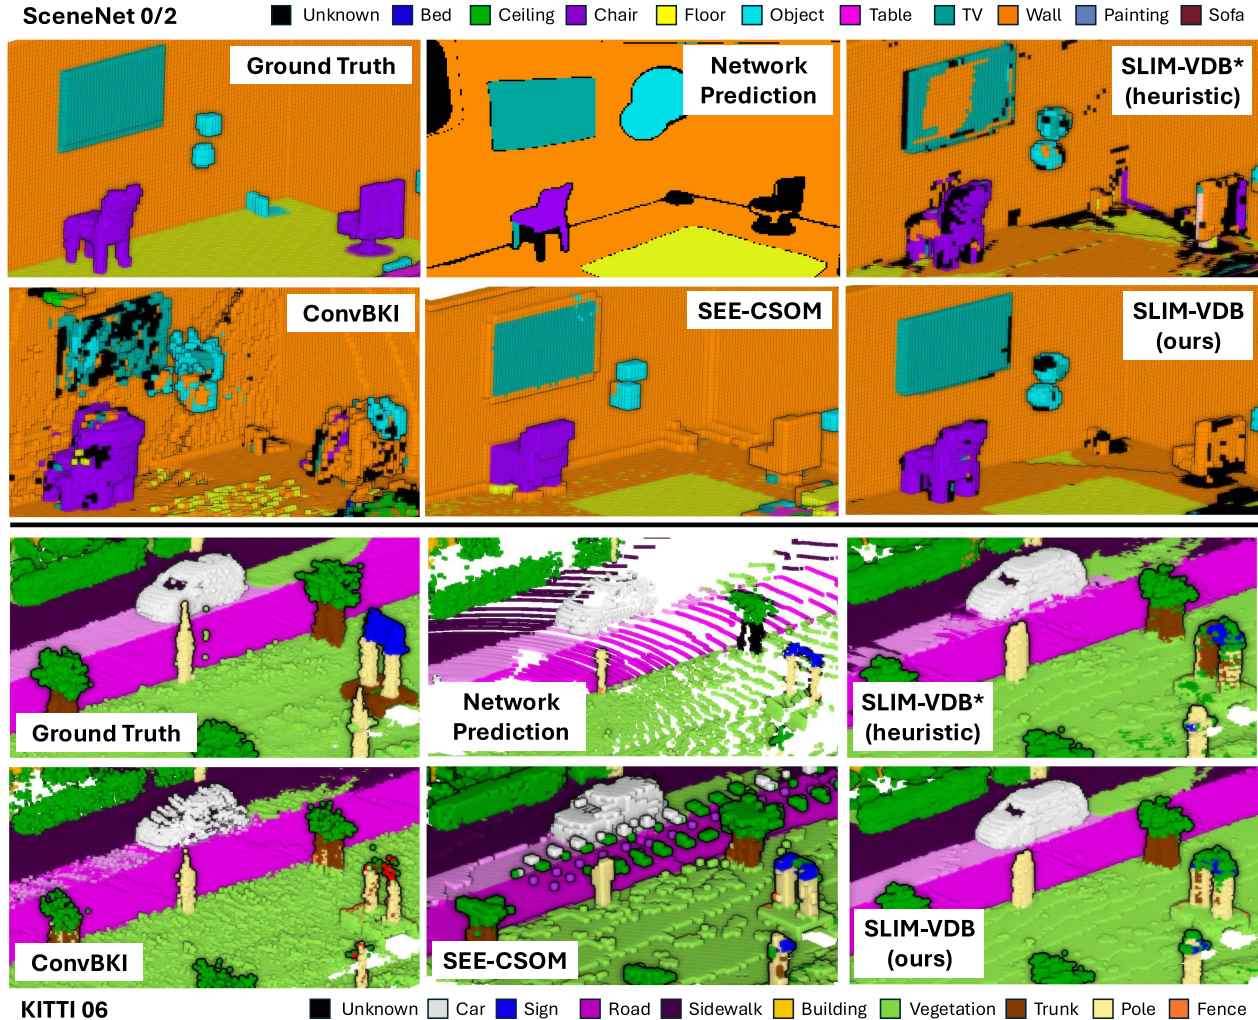


Fig. 4. Qualitative comparison on SceneNet 0/2 and SemanticKITTI 06 between the ground truth, the network output on a single frame, SLIM-VDB, and other probabilistic baselines (SEE-CSOM [17] and ConvBKI [16]).

would like to note that this is a significant advantage of our method SLIM-VDB, as it employs a similar Bayesian update with improved global scalability.

In contrast, Voxfield Panmap and SNI-SLAM assume consistent and error-free semantic predictions. As such, they do not use Bayesian inference for semantic estimation. For a fair comparison, we evaluate our approach without the Bayesian update step – denoted as SLIM-VDB* in the results – against Voxfield Panmap and SNI-SLAM.

1) *Runtime*: Table I presents the average frames per second (FPS) for each method. SLIM-VDB has the shortest runtime on both the large-scale outdoor SemanticKITTI and indoor SceneNet scenes, outperforming all probabilistic baselines with the highest FPS. When the Bayesian update is not used, SLIM-VDB* also outperforms the non-probabilistic baselines Voxfield Panmap and SNI-SLAM. This demonstrates the efficiency improvements of having OpenVDB as a mapping backbone. Fig. 3 visualizes the average runtimes for the pre-processing, integration, and visualization steps of a single frame across all methods. SLIM-VDB outperforms all baseline methods for each

step with notable improvements to the integration. Although SLIM-VDB also greatly improves on baseline visualization cost, these are less comparable because our method only renders a single viewpoint of the scene at every frame rather than publishing a mesh to Rviz.

2) *Memory*: SLIM-VDB also greatly outperforms all baseline methods with respect to combined CPU and GPU memory consumption, as shown in Table I. This reduction in memory consumption is due to the underlying properties of the OpenVDB structure, which does not require initializing all voxels within the map bounds. In contrast, ConvBKI recurrently updates the map through a dense convolutional filter, which requires the entire local map to be stored on GPU memory. Additionally, SLIM-VDB does not require the user to set a static global map size beforehand, as required by ConvBKI and SEE-CSOM.

3) *Semantic Prediction Accuracy*: Finally, we compare the accuracy of SLIM-VDB against the closed-set probabilistic mapping baselines, ConvBKI and SEE-CSOM. In Table II, we present averaged mean Intersection over Union (mIoU)

TABLE II
 AVG SEMANTIC ACCURACY RESULTS ON SCENENET (SN) AND KITTI

	mIoU \uparrow				
	ConvBKI	SEE-CSOM	SLIM-VDB*	SLIM-VDBC	SLIM-VDB ^o
KITTI	0.165	0.128	<u>0.216</u>	0.252	-
SN	0.049	0.128	0.087	<u>0.122</u>	0.116

across 12 random SceneNet scenes and five SemanticKITTI scenes. Qualitative results are shown in Fig. 4. For SceneNet, SLIM-VDB and SEE-CSOM have comparable performance and both significantly outperform ConvBKI. For the KITTI dataset, SLIM-VDB^C outperforms all baselines, including SEE-CSOM. This is likely due to geometric errors in the SEE-CSOM results with larger maps, as discussed in IV-C. We include a non-probabilistic ablation SLIM-VDB*, which simply updates the semantic prediction at each voxel according to the most recent measurement, demonstrating the value of our probabilistic update. Note that a voxel-wise comparison rewards geometric accuracy and penalizes filling in extra voxels.

B. Open-Set Semantic Mapping

In this section, we demonstrate the application of a Bayesian semantic probabilistic update framework to open-set semantics using the Normal Inverse Gamma distribution. For the closed-set semantics evaluation, we use Semantic Segment Anything [20] to generate open-set text labels for each pixel in the SceneNet trajectories and then bin the text labels into the pre-defined SceneNet categories. For fair comparison, we use the same open-set text labels previously generated for the closed-set experiment and then embed them into the feature space of size \mathbb{R}^{512} with Contrastive Language-Image Pre-training (CLIP) [19]. These implicit semantic feature vectors are then integrated into the volume at every time step using the Normal Inverse Gamma update as previously described.

1) *Runtime*: We indicate the runtime performance of SLIM-VDB with open-set labels (denoted by SLIM-VDB^o) in Table I. Note that the runtime reflects the average pre-processing and integration times across the 300 frames of the SceneNet scene, but does not include visualization time. This is because LatentBKI has no real-time visualization.

2) *Memory*: The memory performance (combined CPU and GPU) of SLIM-VDB with open-set labels is also shown in Table I. The use of neural implicit representations in the baseline LatentBKI leads to a very large GPU memory usage. We evaluated the two open-set methods at 10 cm per voxel resolution rather than 5 cm per voxel resolution used for the closed-set experiments as the testbench did not have enough memory to evaluate LatentBKI with a smaller resolution.

3) *Semantic Prediction Accuracy*: To evaluate, we show a comparison against the closed-set (SLIM-VDB^C) mapping results in Table II. The open-set and closed-set SLIM-VDB semantic accuracy is very comparable, reinforcing the unified Bayesian



Fig. 5. Ground truth SceneNet 0/2 RGB image with overlaid red voxels showing regions with SLIM-VDB feature vectors of a high similarity to a text embedding for “sit.” We notice that chairs are highlighted as potential places to take a “sit” action.

TABLE III
 GEOMETRIC ACCURACY OF RESULTING MAPS. **BOLD IS BEST**

	Average Chamfer Distance (\downarrow)				
	VoxField	SNI-SLAM	ConvBKI	SEE-CSOM	SLIM-VDB
KITTI	6.079	-	4.000	504.952	0.079
SN	0.008	-	1.189	0.003	0.004

update framework proposed in this work. Some potential users of this system may prefer the expressiveness of open-set semantics, as the semantic class “chair” could be more directly linked to the downstream action “sit.” See Fig. 5 for an overlay of SLIM-VDB voxels with semantics similar to the CLIP embedding for “sit” in red. Similarity in the feature space is measured using cosine similarity when querying against a single word or phrase.

C. Geometric Mapping Accuracy

In this experiment we demonstrate that the volumetric mapping accuracy of SLIM-VDB is comparable to our baselines (see Table III). We use the L2 Chamfer distance on the final exported map from our selected SceneNet trajectories. SLIM-VDB and VDBFusion have identical geometric mapping accuracy, as they both have the same underlying TSDF update framework. Our method is comparable in accuracy to SEE-CSOM in the indoor SceneNet scenes, although it is slightly outperformed due to the SEE-CSOM method’s approach to removing less certain voxels using an entropy heuristic. However, SLIM-VDB drastically outperforms all methods in the outdoor scene as the baseline methods require more memory to map the entire scene than the testbench can provide, and SEE-CSOM’s entropy heuristic results in “boxy” objects not representative of the real scene. ConvBKI populates extra voxels into the map due to noisy depth measurements without free space sampling. We are unable to include geometric accuracy results from SNI-SLAM as the NeRF sampling step is too computationally expensive for our machines.

D. Rendering

We compare rendering from the VDB map with both CPU and GPU. Over 10 images of size 691 x 256, the average GPU render speed is 2.1 ms and the average CPU render speed is 181.8 ms. For a larger image of 1209 x 448 px, the GPU average render

speed is 3.9 ms and the CPU average render speed is 527.7 ms. If the user does not have access to a GPU, the mapping program could be run headless, as the map pre-processing and integration is run only on the CPU.

E. Mobile Robot Hardware

We present experiments on an NVIDIA Jetson Orin Development Kit with just 4 GB of GPU memory to demonstrate the usability of SLIM-VDB on real mobile robot hardware. We run SLIM-VDB with closed-set semantics, visualization, and a resolution of 5 cm per voxel on the selected SceneNet scenes, and the resulting average FPS is 6.12. This demonstrates that running SLIM-VDB on mobile robot hardware on indoor scenes is feasible.

V. CONCLUSIONS

In this work, we present SLIM-VDB: a 3D probabilistic semantic mapping framework that leverages OpenVDB's efficient volumetric data storage. Our proposed system supports closed-or open-set semantics in a single parallel Bayesian update framework, which to the best of our knowledge is the first single mapping package to do so. In our experiments, we find that SLIM-VDB is able to perform integration of new data points into the map at higher speeds and with a lower memory consumption than the baselines, while retaining comparable semantic and geometric accuracy. Additionally, we integrate NanoVDB for GPU-accelerated rendering. One limitation of our approach is that the real-time viewing of the 3D map is limited to whichever camera rendering viewpoints are previously chosen. We hope to address this in the future by adapting the existing NanoVDB Viewer tool to update dynamically during mapping. Additionally, future work will include integration with fVDB [32], a new deep learning toolset that enables convolution and splatting on VDB grids. In order to propel the usability of SLIM-VDB, an open-source codebase will be released.

REFERENCES

- [1] I. Vizzo, T. Guadagnino, J. Behley, and C. Stachniss, "VDBFusion: Flexible and efficient TSDF integration of range sensor data," *Sensors*, vol. 22, no. 3, 2022, Art. no. 1296.
- [2] S. Zhu et al., "SNI-SLAM: Semantic neural implicit SLAM," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2024, pp. 21167–21177.
- [3] P. Ewen et al., "You've got to feel it to believe it: Multi-modal Bayesian inference for semantic and property prediction," in *Proc. Robot., Sci. Syst.*, 2024, pp. 83–97.
- [4] B. Curless and M. Levoy, "A volumetric method for building complex models from range images," in *Proc. 23rd Annu. Conf. Comput. Graph. Interact. Techn.*, 1996, pp. 303–312.
- [5] L. Schmid et al., "Panoptic Multi-TSDFs: A flexible representation for online multi-resolution volumetric mapping and long-term dynamic scene consistency," in *Proc. 2022 Int. Conf. Robot. Automat.*, 2022, pp. 8018–8024.
- [6] H. Park et al., "Real-Time, accurate, and consistent video semantic segmentation via unsupervised adaptation and cross-unit deployment on mobile device," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 21431–21438.
- [7] J. Wilson et al., "LatentBKI: Open-dictionary continuous mapping in visual-language latent spaces with quantifiable uncertainty," *IEEE Robot. Automat. Lett.*, vol. 10, no. 4, pp. 3102–3109, Apr. 2025.
- [8] K. Museth, "VDB: High-resolution sparse volumes with dynamic topology," *ACM Trans. Graph.*, vol. 32, pp. 1–22, 2013.
- [9] P. Ewen, A. Li, Y. Chen, S. Hong, and R. Vasudevan, "These maps are made for walking: Real-time terrain property estimation for mobile robots," *IEEE Robot. Automat. Lett.*, vol. 7, no. 3, pp. 7083–7090, Jul. 2022.
- [10] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "OctoMap: An efficient probabilistic 3D mapping framework based on octrees," *Auton. Robots*, vol. 34, pp. 189–206, 2013.
- [11] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3D euclidean signed distance fields for on-board MAV planning," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2017, pp. 1366–1373.
- [12] 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 *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2022, pp. 5331–5338.
- [13] F. Tosi et al., "How NeRFs and 3D Gaussian splatting are reshaping SLAM: A survey," 2024, *arXiv:2402.13255*.
- [14] S. Isaacson, P.-C. Kung, M. Ramanagopal, R. Vasudevan, and K. A. Skinner, "LONER: LiDAR only neural representations for real-time SLAM," *IEEE Robot. Automat. Lett.*, vol. 8, no. 12, pp. 8042–8049, Dec. 2023.
- [15] H. He and B. Upcroft, "Nonparametric semantic segmentation for 3D street scenes," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2013, pp. 3697–3703.
- [16] J. Wilson et al., "ConvBKI: Real-time probabilistic semantic mapping network with quantifiable uncertainty," *IEEE Trans. Robot.*, vol. 40, pp. 4648–4667, 2024.
- [17] Y. Deng, M. Wang, Y. Yang, D. Wang, and Y. Yue, "SEE-CSOM: Sharp-edged and efficient continuous semantic occupancy mapping for mobile robots," *IEEE Trans. Ind. Electron.*, vol. 71, no. 2, pp. 1718–1728, Feb. 2024.
- [18] L. Gan, R. Zhang, J. W. Grizzle, R. M. Eustice, and M. Ghaffari, "Bayesian spatial kernel smoothing for scalable dense semantic mapping," *IEEE Robot. Automat. Lett.*, vol. 5, no. 2, pp. 790–797, Apr. 2020.
- [19] A. Radford et al., "Learning transferable visual models from natural language supervision," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 8748–8763.
- [20] J. Chen, Z. Yang, and L. Zhang, "Semantic segment anything," 2023.
- [21] K. Museth, "NanoVDB: A GPU-friendly and portable VDB data structure for real-time rendering and simulation," in *Proc. ACM SIGGRAPH 2021 Talks*, 2021, pp. 1–2.
- [22] M. G. Besselmann, L. Puck, L. Steffen, A. Roennau, and R. Dillmann, "VDB-Mapping: A high resolution and real-time capable 3D mapping framework for versatile mobile robots," in *Proc. IEEE 17th Int. Conf. Automat. Sci. Eng.*, 2021, pp. 448–454.
- [23] R. Hagmanns, T. Emter, M. Grosse Besselmann, and J. Beyerer, "Efficient global occupancy mapping for mobile robots using OpenVDB," in *Proc. IROS Agile Robot. Workshop*, 2022.
- [24] M. G. Besselmann, A. Rönnau, and R. Dillmann, "Remote VDB-mapping: A level-based data reduction framework for distributed mapping," in *Proc. Robot. Natural Settings*, Berlin, Germany, 2022, pp. 448–459.
- [25] Y. Bai, Z. Miao, X. Wang, Y. Liu, H. Wang, and Y. Wang, "VDBblox: Accurate and efficient distance fields for path planning and mesh reconstruction," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 7187–7194.
- [26] V. Walker, F. Vanegas, and F. Gonzalez, "NanoMap: A GPU-accelerated OpenVDB-based mapping and simulation package for robotic agents," *Remote Sens.*, vol. 14, no. 21, 2022, Art. no. 5463.
- [27] S. Tu, *The Dirichlet-Multinomial and Dirichlet-Categorical Models for Bayesian Inference*. Berkeley, CA, USA: UC Berkeley, 2014.
- [28] K. P. Murphy, "Conjugate Bayesian analysis of the Gaussian distribution," 2007.
- [29] J. Behley et al., "SemanticKITTI: A dataset for semantic scene understanding of LiDAR sequences," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 9297–9307.
- [30] J. McCormac, A. Handa, S. Leutenegger, and A. J. Davison, "SceneNet RGB-D: Can 5M synthetic images beat generic ImageNet pre-training on indoor segmentation?," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2678–2687.
- [31] Z. Zhou, Y. Zhang, and H. Foroosh, "Panoptic-polarnet: Proposal-free LiDAR point cloud panoptic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 13194–13203.
- [32] F. Williams et al., "fVDB: A deep-learning framework for sparse, large scale, and high performance spatial intelligence," *ACM Trans. Graph.*, vol. 43, no. 4, pp. 1–15, 2024.