OSM-BKI: OpenStreetMap-Guided Bayesian Kernel Inference for Domain-Robust LiDAR Semantic Mapping

Doncey Albin[1], Sandilya Sai Garimella[2], Dreycey Albin[3], Xujia Zhang[1], Lu Gan[2], Christoffer Heckman[1]

***Abstract*— LiDAR semantic segmentation models degrade sharply under domain shift, and continuous mapping frameworks such as Bayesian Kernel Inference (BKI) smooth these errors but cannot correct them. OpenStreetMap (OSM) provides a globally available, domain-independent source of semantic structure that has been widely exploited for place recognition, SLAM, and navigation, yet to our knowledge no prior work has integrated OSM into ground-based robotic semantic mapping. We present a method that fuses OSM semantic priors directly into the BKI map update as a spatially varying Dirichlet prior. For each incoming LiDAR scan, nearby OSM features are rasterized into a soft semantic class vector, projected onto each point within the scan using a common label taxonomy, and modulated by a per-class height filter requiring only a known sensor mounting height. Evaluations are made on the KITTI-360 and Multi-Campus (MCD) datasets using cross-domain and in-domain experiments. We find that our approach improves cross-domain mIoU by up to 87% relative over S-BKI (9.49 to 17.72 on KITTI-360) and in-domain mIoU by up to 6% relative over S-BKI (55.23 to 58.55 on KITTI-360).**

I. INTRODUCTION

Modern robotic perception demands spatial representations that are both geometrically accurate and semantically rich. Dense semantic mapping is a fundamental prerequisite for high-level autonomy tasks, ranging from traversability analysis and path planning to long-term scene understanding. While traditional occupancy grids provide essential structural data [1], they lack the categorical context required for task-level decision-making. Consequently, significant research effort has been directed toward fusing noisy, sparse measurements from semantic segmentation networks into coherent, dense 3D maps [2].

A primary challenge in this domain is the inherent uncertainty of sensor data. Deep learning-based segmentation models, such as RangeNet++ [5], while powerful, frequently exhibit residual noise, flickering predictions, and confusion between geometrically similar classes (e.g., *Road* vs. *Sidewalk*). Furthermore, these models often suffer from performance degradation when deployed in environments that differ from their training data [6].

To address sparsity and noise, continuous mapping frameworks such as Semantic Bayesian Kernel Inference (S-BKI) have been proposed [7]. BKI generalizes the traditional occupancy grid update by applying a spatial kernel to smooth sparse observations into a continuous probability distribution [8]. This approach effectively handles sensor sparsity and allows for querying the map at arbitrary resolutions. However, existing BKI formulations typically operate in a "sensor-only" regime, relying exclusively on online perception. This creates a significant vulnerability: systematic sensor errors, such as consistent misclassifications due to domain shift, cannot be corrected by spatial smoothing alone.

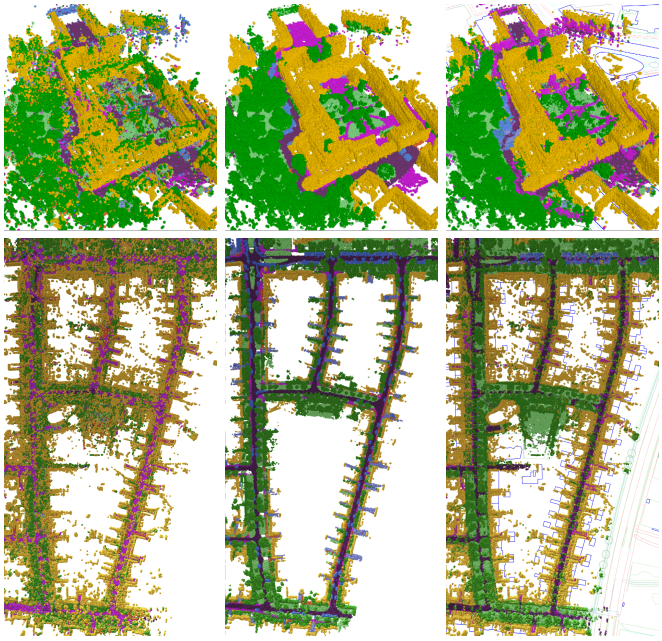

Fig. 1: Resulting semantic maps under cross-domain evaluation. **Top:** *kth_day_09* from MCD [3]. **Bottom:** sequence *0000* from KITTI-360 [4]. **Left to right:** S-BKI, ground truth, and OSM-BKI (ours).

We address this limitation by bridging the gap between global semantic priors and local sensor perception. We propose a framework that unifies these disparate information sources into a single Bayesian update rule. We demonstrate that map priors, such as those from OpenStreetMap (OSM), can function as probabilistic filters that regulate the confidence of incoming perception data. We extend the standard generalized kernel inference [9] with a *semantic compatibility kernel* that weighs observations based on their alignment with a continuous, OSM-derived soft prior. The specific contributions of this paper are as follows:

- **Dual-Injection BKI Formulation:** OSM information enters the Bayesian update at two stages: as a spatially varying Dirichlet prior over voxel classes, and as a per-class observation weight that amplifies OSM-consistent sensor measurements during kernel inference.

[1]Authors are with the Department of Computer Science, University of Colorado Boulder, Boulder, CO 80309, USA

[2]Authors are with the School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

[3]Dreycey Albin is with Microsoft, Redmond, WA, USA. This work was performed independently and does not reflect the views of Microsoft.

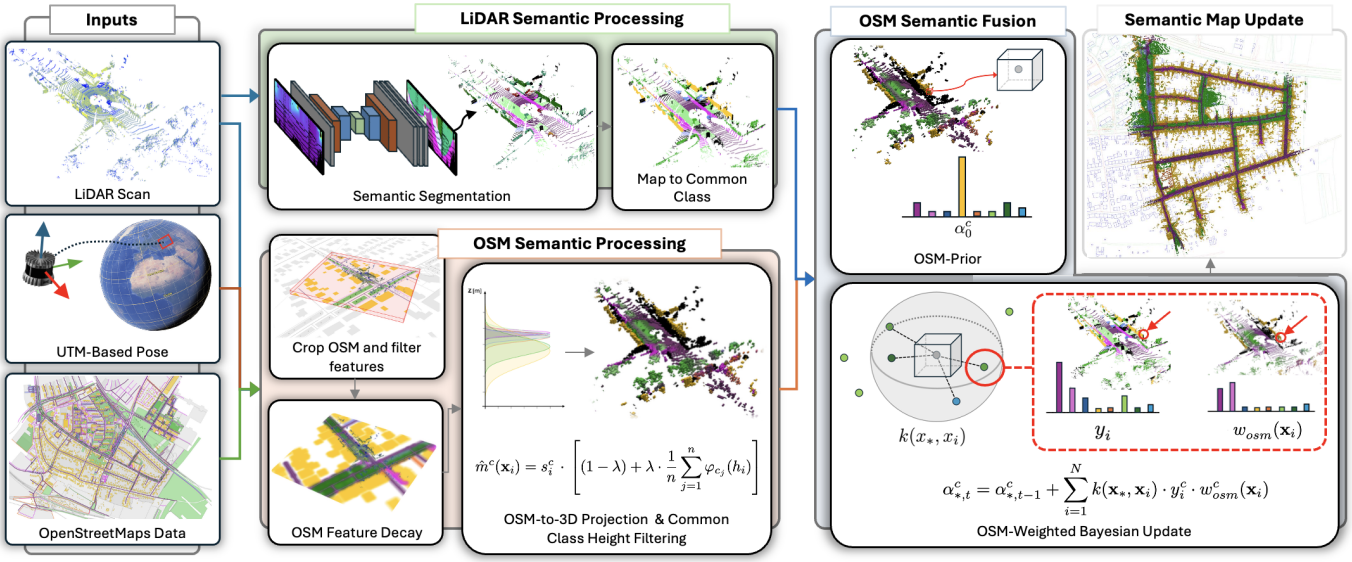

Fig. 2: Pipeline of our OpenStreetMap-fused semantic BKI framework.

A conversion matrix bridges the OSM and sensor label taxonomies, encoding known inter-class confusions so that plausible misclassifications are tolerated.

- **Height-Aware Semantic Filtering:** A per-class Gaussian height weight modulates each semantic channel as a function of measured height, lifting inherently 2D OSM priors into 3D using only a known sensor mounting height and no external elevation data.
- **Cross-Domain and In-Domain Experiments:** We demonstrate on a train-on-one, test-on-the-other protocol across KITTI-360 and MCD that OSM priors recover up to 76% relative mIoU under large domain shift, with the largest gains in classes most affected by cross-dataset confusion.

II. METHOD

To isolate the contribution of the OSM signal to semantic mapping, we assume accurate georegistration of all sensor poses. Prior work has addressed LiDAR-to-OSM localization [10], [11], as well as the integration of OSM in SLAM [12]. While these works have inspired our approach, integration of such upstream methods lies beyond the scope of this study.

We use Bayesian Kernel Inference (BKI) as the substrate for continuous LiDAR semantic mapping and incorporate OpenStreetMap (OSM) at two stages of the Bayesian update: as a spatially varying Dirichlet prior over voxel classes, and as a per-class likelihood weight that modulates how strongly each sensor measurement updates the posterior. Together, these biases each voxel toward classes consistent with surrounding OSM features while remaining adaptable where OSM is sparse, misaligned, or ambiguous.

A. Preliminaries

Semantic Bayesian kernel inference (S-BKI) [7] models each voxel $*$ as a Dirichlet distribution over C semantic classes with concentration parameters $\boldsymbol{\alpha}_* \in \mathbb{R}^C$. Given semantically labeled points $\{(x_i, y_i)\}_{i=1}^N$, where y_i^c is the categorical likelihood for class c, the closed-form update for voxel $*$ at time t is

$$\alpha_{*,t}^c = \alpha_{*,t-1}^c + \sum_{i=1}^{N} k(x_*, x_i)\, y_i^c, \tag{1}$$

where $k(x_*, x_i) \in [0, 1]$ is a sparse kernel [13] that decays with Euclidean distance and vanishes beyond a length scale l. We refer the reader to [7], [9], [14]–[17] for a more complete introduction.

B. OSM Feature Cropping and Distance Weighting

OpenStreetMap (OSM) encodes geographic features as three primitive geometries: polygons, polylines, and points. To impose a consistent semantic footprint across primitives, we inflate non-areal features to polygons at runtime: polylines are buffered to rectangles using a class-dependent width (e.g., lane width for roads), and points are dilated to disks of class-dependent radius. These initial widths and radii are updated at runtime when the available feature metadata is provided. For both computational efficiency and experimental tractability, we restrict our analysis to a curated subset of C_o OSM feature classes. Throughout, $o \in C_o$ indexes OSM classes and $c \in C$ indexes the common semantic classes shared between datasets.

For each incoming LiDAR scan, the resulting polygonal layers are cropped to a rectangular region centered on and bounded by the scan extent. Each point i at $x_i = [x_i^1, x_i^2, x_i^3]$ is then assigned a soft class vector $\mathbf{m}_i \in \mathbb{R}^{C_o}$ over the C_o OSM channels. Interior points receive unit support, while exterior support decays linearly with the signed distance d_s to the polygon boundary,

$$w(d_s) = \max\left(0,\ 1 - \frac{\max(d_s, 0)}{\delta}\right), \quad \delta \geq d_s \tag{2}$$

where the falloff scale δ sets the maximum distance beyond a polygon boundary at which it can still contribute support,

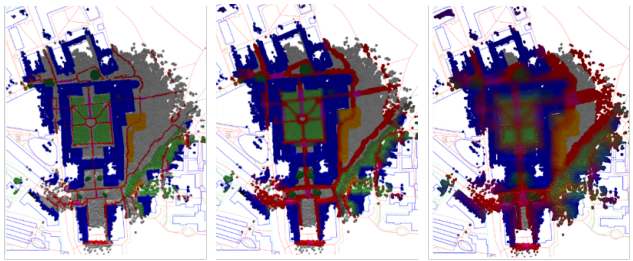

$\delta = 1.0\,\mathrm{m}$ $\delta = 3.0\,\mathrm{m}$ $\delta = 10.0\,\mathrm{m}$

Fig. 3: Illustration of increasing δ for OSM prior decay on MCD's *kth_day_09* sequence.

providing a buffer for spatial misalignment between OSM and the robot's local map.

C. OSM-to-Semantic Label Mapping

Since OSM layers do not correspond one-to-one with common semantic labels, we define a conversion matrix $\mathbf{M} \in \mathbb{R}^{C \times C_o}$, where each entry $M_{c,o} = P(c \mid o)$ maps OSM class probabilities into a common set of C semantic labels shared between KITTI-360 and MCD taxonomies. We adopt the convention that row $c \in C$ indexes common semantic classes and column $o \in C_o$ indexes OSM classes, and each column $\mathbf{M}_{:,o}$ is normalized to sum to one, distributing each OSM class's support across the semantic classes. For each point i, the projected class vector $\mathbf{s}_i \in \mathbb{R}^C$ over common semantic classes is given by:

$$\mathbf{s}_i = \mathbf{M}\,\mathbf{m}_i, \tag{3}$$

whose c-th entry is:

$$\mathbf{s}_i^c = \sum_{o=1}^{C_o} \mathbf{M}_{c,o}\,\mathbf{m}_i^o. \tag{4}$$

Off-diagonal entries of $\mathbf{M}$ encode known LiDAR confusions so that plausible misclassifications are not penalized, while implausible ones receive negligible weight.

D. Height Filtering

Because OSM polygons reside strictly in 2D, $\mathbf{s}_i$ alone does not distinguish classes that overlap in the ground plane but differ in height (*Road* at ground level and a *Building* facade above it). We address this using a per-class height filter that modulates each channel of $\mathbf{s}_i$ by a Gaussian weight on measured height. For each measurement $\mathbf{x}_i$ with vertical coordinate x_i^3, we estimate the height above ground using the known sensor mounting height z_{sensor} and the scan origin elevation z_{origin}:

$$h_i = x_i^3 - (z_{\text{origin}} - z_{\text{sensor}}), \tag{5}$$

where z_{origin} is the vertical coordinate of the sensor pose for the current scan and z_{sensor} is the known mounting height of the sensor above ground. For each common semantic class c, we define a Gaussian height weight centered on the expected height-above-ground μ_c with tolerance τ_c:

$$\varphi_c(h) = \exp\left(-\frac{\left[\max(0,\ |h - \mu_c| - d_c)\right]^2}{2\tau_c^2}\right), \tag{6}$$

where $d_c \geq 0$ is a per-class dead zone that maintains full weight ($\varphi_c = 1$) within a band $[\mu_c - d_c,\ \mu_c + d_c]$ before Gaussian decay begins. This is particularly important for classes like *Building*, whose points span a wide range of heights: the dead zone preserves full weight across the expected vertical extent of the structure, while the Gaussian tails attenuate the weight at heights inconsistent with μ_c.

Each row of $\mathbf{s}_i$ (Eq. 3) corresponds to a common semantic class c and is modulated by the per-class Gaussian weight $\varphi_c(h_i)$. We blend this weight with the unmodulated channel using a scalar $\lambda \in [0, 1]$:

$$\hat{m}^c(\mathbf{x}_i) = s_i^c \cdot \left[(1 - \lambda) + \lambda \cdot \frac{1}{n}\sum_{j=1}^{n} \varphi_{c_j}(h_i)\right]. \tag{7}$$

The height-modulated vector $\hat{\mathbf{m}}_i$ attenuates classes whose expected height-above-ground is inconsistent with h_i (e.g., *Road* at $h_i = 10\,\mathrm{m}$) and preserves those that are consistent (e.g., *Building* at $h_i = 10\,\mathrm{m}$). When $\lambda = 0$, height filtering is disabled and OSM polygons are treated as infinitely tall, applying equal weight at all elevations. Because the ground estimate depends only on the sensor pose and a single calibration constant, the height filter requires no external elevation data and generalizes to any platform with a known sensor mounting height.

E. Prior Initialization

For each voxel we convert the height-filtered OSM projection into a per-class boost on the BKI Dirichlet concentration. The initial Dirichlet parameter for semantic class c at voxel $\mathbf{x}_*$ is

$$\alpha_0^c(\mathbf{x}_*) = \bar{\alpha} + \beta_p \cdot \hat{m}^c(\mathbf{x}_*), \tag{8}$$

where $\bar{\alpha}$ is the uniform BKI base count and β_p is a global scalar controlling how strongly OSM-derived term shapes the Dirichlet prior. Voxels with no OSM support fall back to the uniform prior $\bar{\alpha}$.

F. OSM-Weighted BKI Update

During sensor data integration, we introduce a class-dependent multiplicative weight for each training point $(\mathbf{x}_i, \mathbf{y}_i)$:

$$w_{osm}^c(\mathbf{x}_i) = 1 + \beta_k \cdot \hat{m}^c(\mathbf{x}_i), \tag{9}$$

where β_k is the OSM kernel strength and $\hat{m}^c$ is the height-filtered projection (Eq. 7). Classes not supported by OSM at location $\mathbf{x}_i$ retain $w_{osm}^c = 1$. The Dirichlet parameter update for voxel $\mathbf{x}_*$ at time t is then:

$$\alpha_{*,t}^c = \alpha_{*,t-1}^c + \sum_{i=1}^{N} k(\mathbf{x}_*, \mathbf{x}_i) \cdot y_i^c \cdot w_{osm}^c(\mathbf{x}_i), \tag{10}$$

where $k(\mathbf{x}_*, \mathbf{x}_i)$ is the sparse kernel and y_i^c is the one-hot class likelihood for point i. The sparse kernel ensures only points within distance l of the voxel center contribute, while w_{osm}^c boosts the contribution of points at locations that OSM associates with class c.

III. EXPERIMENTS

We evaluate on two LiDAR semantic mapping datasets, KITTI-360 [4] and the Multi-Campus Dataset (MCD) [3], under cross-domain and in-domain settings. The upstream segmentation model, CENet [18], is held fixed across all mapping methods so that differences are attributable solely to the training and test data distributions. We compare against S-BKI as the most direct baseline sharing the same sparse kernel-weighted Bayesian update formulation. We additionally compare our results with the baseline semantic segmentation results.

For each dataset, we hold out two sequences for evaluation and train on the remainder. We evaluate under two protocols. In the *cross-domain* setting, CENet trained on MCD is evaluated on the KITTI-360 holdout sequences, and vice versa. Because the two datasets use distinct semantic taxonomies, we define a common label space and discard classes without a correspondence. In the *in-domain* setting, the model is evaluated on the holdout sequences of the dataset it was trained on. We report Intersection-over-Union (IoU) for each semantic class, as well as the mean Intersection-over-Union (mIoU) and mean accuracy (Acc) for each sequence.

For all experiments, a new keyframe is extracted when the Euclidean distance from the previous keyframe exceeds a threshold (Δ_{kf}). For all experiments, we maintain the same hyperparameters besides β_p and β_k. For cross-domain experiments, we use $\beta_p = 0.5$ and $\beta_k = 1.0$. For in-domain, we select $\beta_p = 0.1$ and $\beta_k = 0.1$. For all other hyperparameters, please reference Table I.

TABLE I: Hyperparameters used across all experiments.

Dataset	Resolution	α_0^c	l	β_p	β_k	δ	Δ_{kf}
KITTI-360	0.5m	1.0	1.0	0.5	1.0	2.0m	5.0m
MCD	0.3m	1.0	1.0	0.5	1.0	2.0m	2.5m

TABLE II: Quantitative results on *cross-domain* experiments.

		Per-class IoU (%)									
Sequence	**Method**	road	sidewalk	parking	building	fence	terrain	vegetation	vehicle	mIoU [%]	Acc [%]
KITTI-360 Dataset											
0000	*Baseline*	1.95	12.58	0.00	48.39	**1.90**	9.21	18.72	**0.48**	11.65	19.87
	S-BKI [7]	0.51	7.88	0.00	48.86	0.40	3.98	14.26	0.06	9.49	16.94
	Ours	**3.08**	**23.70**	**8.24**	**52.67**	0.56	**34.16**	**19.09**	0.25	**17.72**	**34.24**
0009	*Baseline*	1.20	15.92	0.00	40.10	**1.40**	14.10	**26.78**	**0.25**	12.47	21.71
	S-BKI [7]	0.17	11.02	0.00	40.27	0.34	7.38	24.67	0.02	10.48	18.97
	Ours	**1.35**	**30.29**	**1.03**	**41.92**	0.51	**22.84**	22.64	0.05	**15.08**	**31.31**
Multi-Campus Dataset											
kth_day_09	*Baseline*	24.27	3.65	0.04	74.90	3.28	28.49	47.86	19.07	25.20	36.64
	S-BKI [7]	24.98	2.27	0.01	81.04	5.18	30.46	56.18	**29.12**	28.66	38.10
	Ours	**34.25**	**20.30**	**9.87**	**83.90**	**7.54**	**57.95**	**68.86**	18.57	**37.66**	**51.65**
kth_night_05	*Baseline*	12.64	2.05	0.00	78.01	1.01	18.04	40.59	8.51	20.11	32.66
	S-BKI [7]	**13.02**	1.01	0.00	83.09	0.84	17.02	50.61	**14.10**	22.46	33.90
	Ours	13.01	**16.89**	**1.19**	**84.78**	**2.36**	**32.54**	**66.96**	6.67	**28.05**	**42.44**

TABLE III: Quantitative results on *in-domain* experiments.

		Per-class IoU (%)									
Sequence	**Method**	road	sidewalk	parking	building	fence	terrain	vegetation	vehicle	mIoU [%]	Acc [%]
KITTI-360 Dataset											
0000	*Baseline*	72.47	63.28	24.47	70.38	23.99	55.18	56.75	56.34	52.86	69.14
	S-BKI [7]	69.46	59.23	22.58	**77.53**	**28.95**	**58.78**	62.13	63.19	55.23	69.16
	Ours	**73.97**	**64.07**	**40.66**	71.98	27.59	48.13	**65.82**	**76.17**	**58.55**	**76.60**
0009	*Baseline*	84.51	67.13	36.46	67.76	23.68	54.69	65.07	55.81	56.89	73.62
	S-BKI [7]	83.47	63.01	38.53	**75.89**	29.16	**57.34**	69.44	62.87	59.96	74.05
	Ours	**86.01**	**68.70**	**46.13**	69.83	**29.27**	53.68	**72.63**	**74.77**	**62.63**	**79.50**
Multi-Campus Dataset											
kth_day_09	*Baseline*	64.72	73.39	29.15	91.97	12.49	73.85	83.04	**49.65**	59.78	70.78
	S-BKI [7]	64.63	74.23	28.87	94.04	**34.99**	74.37	**86.31**	47.41	63.11	68.25
	Ours	**74.35**	**79.66**	**33.13**	**94.08**	19.51	**76.35**	84.21	44.58	**63.23**	**76.65**
kth_night_05	*Baseline*	67.60	68.41	33.78	90.34	6.68	64.25	78.91	25.40	54.42	64.21
	S-BKI [7]	**70.33**	67.71	**40.02**	92.44	8.44	65.61	**86.88**	**28.38**	57.48	63.87
	Ours	59.88	**76.57**	39.98	**92.73**	**13.57**	**67.76**	86.15	27.43	**58.01**	**71.18**

IV. RESULTS AND DISCUSSION

Table II summarizes per-class IoU, mIoU, and accuracy metrics across all sequences. OSM-BKI consistently outperforms all baselines in mIoU and accuracy under both cross-domain directions. The largest gains occur in ground classes such as road, terrain, and sidewalk, where OSM polygons help disambiguate structurally similar surfaces that domain shift causes the network to confuse. However, classes with sparse OSM coverage see limited benefit: parking remains near zero on KITTI-360 due to poor annotation density, and fence is unaffected since fences are rarely mapped in OSM. Overall, improvements are larger when the domain gap is more severe (up to 76% relative mIoU gain on KITTI-360), suggesting the OSM prior is most valuable precisely when the learned model transfers worst.

While we do not report runtime metrics here, the nearest-feature distance queries required for OSM support introduces additional computational overhead relative to S-BKI. A detailed runtime analysis is left to future work.

V. CONCLUSION AND FUTURE WORK

The utilization of OSM as an auxiliary semantic signal within the BKI framework can be highly effective at correcting noisy predictions. Despite OSM data being constrained to a 2D plane, projections into 3D space are still viable through per-class height consistency kernels that require only a known sensor mounting height to estimate ground level, with no external elevation data needed. In future work, we plan to refine the static per-class height priors using Bayesian inference over observed class-height distributions, and to incorporate publicly accessible geospatial data such as DEMs and DSMs as additional priors for semantic occupancy estimation.

REFERENCES

[1] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "OctoMap: An efficient probabilistic 3D mapping framework based on octrees," *Autonomous Robots*, 2013, software available at https://octomap.github.io. [Online]. Available: https://octomap.github.io

[2] B. Gao, Y. Pan, C. Li, S. Geng, and H. Zhao, "Are we hungry for 3d lidar data for semantic segmentation? a survey of datasets and methods," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 6063–6081, 2022.

[3] T.-M. Nguyen, S. Yuan, T. H. Nguyen, P. Yin, H. Cao, L. Xie, M. Wozniak, P. Jensfelt, M. Thiel, J. Ziegenbein, and N. Blunder, "Mcd: Diverse large-scale multi-campus dataset for robot perception," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6 2024, pp. 22 304–22 313. [Online]. Available: https://mcdviral.github.io/

[4] Y. Liao, J. Xie, and A. Geiger, "Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 3, pp. 3292–3310, 2023.

[5] A. Milioto, I. Vizzo, J. Behley, and C. Stachniss, "RangeNet++: Fast and Accurate LiDAR Semantic Segmentation," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2019.

[6] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall, "Semantickitti: A dataset for semantic scene understanding of lidar sequences," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 9296–9306.

[7] L. Gan, R. Zhang, J. W. Grizzle, R. M. Eustice, and M. Ghaffari, "Bayesian spatial kernel smoothing for scalable dense semantic mapping," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 790–797, 2020.

[8] K. Doherty, T. Shan, J. Wang, and B. Englot, "Learning-aided 3-d occupancy mapping with bayesian generalized kernel inference," *IEEE Transactions on Robotics*, vol. 35, no. 4, pp. 953–966, 2019.

[9] K. Doherty, J. Wang, and B. Englot, "Bayesian generalized kernel inference for occupancy map prediction," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 3118–3124.

[10] S. Kang, M. Y. Liao, Y. Xia, O. Wysocki, B. Jutzi, and D. Cremers, "Opal: Visibility-aware lidar-to-openstreetmap place recognition via adaptive radial fusion," 2025. [Online]. Available: https://arxiv.org/abs/2504.19258

[11] F. Yan, O. Vysotska, and C. Stachniss, "Global localization on openstreetmap using 4-bit semantic descriptors," in *2019 European Conference on Mobile Robots (ECMR)*, 2019, pp. 1–7.

[12] M. Frosi, V. Gobbi, and M. Matteucci, "Osm-slam: Aiding slam with openstreetmaps priors," *Frontiers in Robotics and AI*, vol. Volume 10 - 2023, 2023. [Online]. Available: https://www.frontiersin.org/journals/robotics-and-ai/articles/10.3389/frobt.2023.1064934

[13] W. Vega-Brown, M. Doniec, and N. Roy, "Nonparametric bayesian inference on multivariate exponential families," in *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, Eds., vol. 27, 2014.

[14] J. Wilson, Y. Fu, J. Friesen, P. Ewen, A. Capodieci, P. Jayakumar, K. Barton, and M. Ghaffari, "Convbki: Real-time probabilistic semantic mapping network with quantifiable uncertainty," *IEEE Transactions on Robotics*, vol. 40, pp. 4648–4667, 2024.

[15] J. Kim, J. Seo, and J. Min, "Evidential semantic mapping in off-road environments with uncertainty-aware bayesian kernel inference," 2024. [Online]. Available: https://arxiv.org/abs/2403.14138

[16] 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 Transactions on Industrial Electronics*, vol. 71, no. 2, pp. 1718–1728, 2024.

[17] J. Kim, M. Jeon, J. Min, K. Kwak, and J. Seo, "E2-bki: Evidential ellipsoidal bayesian kernel inference for uncertainty-aware gaussian semantic mapping," 2026. [Online]. Available: https://arxiv.org/abs/2509.11964

[18] H. Cheng, X. Han, and G. Xiao, "Cenet: Toward concise and efficient lidar semantic segmentation for autonomous driving," in *2022 IEEE International Conference on Multimedia and Expo (ICME)*, 2022, pp. 01–06.