

Differentiable SLAM Helps Deep Learning-based LiDAR Perception Tasks

Prashant Kumar¹

prahantk.nan@gmail.com

Dheeraj Vattikonda²

dheeraj.46329@gmail.com

Vedang Bhupesh Shenvi Nadkarni³

vedang.nadkarni@gmail.com

Erqun Dong²

doneq13@gmail.com

Sabyasachi Sahoo⁴

sabyasachi.sahoo.1@ulaval.ca

¹ IIT Delhi, India

² McGill University, MILA

³ BITS Pilani, India

⁴ Université Laval, MILA

Abstract

We investigate a new paradigm that uses differentiable SLAM architectures in a self-supervised manner to train end-to-end deep learning models in various LiDAR based applications. To the best of our knowledge there does not exist any work that leverages SLAM as a training signal for deep learning based models. We explore new ways to improve the efficiency, robustness, and adaptability of LiDAR systems with deep learning techniques. We focus on the potential benefits of differentiable SLAM architectures for improving performance of deep learning tasks such as classification, regression as well as SLAM. Our experimental results demonstrate a non-trivial increase in the performance of two deep learning applications - Ground Level Estimation and Dynamic to Static LiDAR Translation, when used with differentiable SLAM architectures. Overall, our findings provide important insights that enhance the performance of LiDAR based navigation systems. We demonstrate that this new paradigm of using SLAM Loss signal while training LiDAR based models can be easily adopted by the community.

1 Introduction

We investigate the impact of differentiable SLAM on training better deep learning-based machine perception (DLMP) models through the methodology of fully differentiable back-propagation. SLAM is a foundational component in mobile robotics. Robots build a map of their environment while simultaneously determining their locations within the map. The potential of using SLAM to help DLMP has been used on tasks such as learning observation models for new modalities [65], object localization and tracking [21, 23], etc. State-of-the-art SLAM systems are often not differentiable, presenting a challenge in integrating them with deep learning approaches. While recent works such as GradSLAM [24] and GradLidarSLAM [8] have addressed this issue by proposing differentiable SLAM architectures,

there has been little investigation into how these architectures impact the performance of deep learning models. It is an open question whether differentiable SLAM architectures can be effectively utilized to enhance the performance of deep learning models in various LiDAR based applications. We propose a self-supervised framework that leverages a differentiable SLAM architecture. It enable fully differentiable training of deep learning models with SLAM error for various LiDAR applications. Our method is based on the principle of minimizing the discrepancy between the output of the deep learning model and the ground truth, including the trajectory error obtained from ground truth and predicted LiDAR scans. Through extensive experimentation, we demonstrate that our approach outperforms existing methods and achieves improvements in deep learning tasks. Our results highlight the potential of utilizing differentiable SLAM architectures to enhance the performance of deep learning models. Our main contributions are:

- We propose a new framework to train differentiable LiDAR-based SLAM using deep learning-based machine perception (DLMP) tasks in a self-supervised manner.
- We demonstrate its effectiveness by applying it on two tasks - (1) Ground Plane Estimation and Ground Point Segmentation (2) Dynamic to static LiDAR translation for improved SLAM.
- Our experiments show that our proposed framework significantly improves the performance of deep learning models on DLMP tasks.

1.1 Related Work

1.1.1 Differentiable SLAM

The idea of making SLAM differentiable has been investigated in some previous works [8, 14, 40]. Incorporating differentiable SLAM modules to help deep learning training has huge potential. However, to the best of our knowledge it has not been implemented and well studied. Several works on SLAM that integrate deep learning-based techniques have been introduced in recent years - learning observation model for new modalities [35], learning object pose tracking [21, 23], learning a compact scene representation [8, 43], learning a CNN-based depth predictor as the front-end of a monocular SLAM system [36], etc. While these works leverage learning techniques, they often focus only on specific modules within the SLAM system. Furthermore, these methods are typically limited to visual odometry. Sodhi et.a l. [35] attempt to optimize end-to-end tracking performance by learning observation models using energy-based methods for SLAM on novel modalities like tactile sensors. They do not use trajectory error directly to optimize the observation model. They do not conduct perception tasks explicitly (i.e. no perception based results are available). Different from existing literature, we investigate the use of SLAM trajectory error in a fully differentiable fashion to help LiDAR based deep learning tasks.

1.1.2 LiDAR based Deep Learning

Several works [9, 11, 10, 15, 26, 27, 37, 45] have explored generative modelling for LiDAR. LiDAR based generative modelling was first introduced by Caccia et. al. [4]. They use deep generative models - VAE, GANs to reconstruct as well as generate high quality LiDAR samples. Another work, DSLR [16], extended this idea to generate static structures occluded by

dynamic objects for 3D LiDAR scene reconstruction in an adversarial setting. The work also aims to improve SLAM performance with these static reconstructions. We therefore consider DSLR as a suitable test bed for our work. Another work [26] focuses on alleviating the problem of dropped points on the LiDAR depth map by introducing measurement uncertainty in the generative models. LiDAR data is a rich source of information for the 3D world of vital use for autonomous navigation systems. There exists a good body of work on LiDAR based segmentation [0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12], object detection [13, 14, 15, 16, 17, 18, 19, 20, 21, 22], ground elevation estimation [23, 24, 25]. For tasks other than generative modelling, one of the requirements for utilizing differentiable SLAM based error is that the output must be in the form of a per point prediction/regression so that it can be mapped to the original LiDAR points. A subset of points is then selected for SLAM. SLAM can then be performed between the mapped predicted LiDAR and the input LiDAR for SLAM Loss back-propagation. However, multi-class (except binary) per-point classification based tasks (e.g. segmentation) require certain non differentiable operations (e.g. `torch.isin()`, `torch.argmax()`, etc) to map the predictions to the original LiDAR based on a given criteria (e.g. only static object classes) and cannot be integrated with differentiable SLAM. This is a limitation of the differentiable SLAM module.

We choose a task that unifies binary classification as well as regression to show the effect of differentiable SLAM on the selected task. Ground plane estimation and ground point segmentation unifies both these modalities. We use this task to show the benefit of differentiable SLAM. We use a well-known baseline GndNet [27] that has shown impressive performance on the above mentioned task.

Our primary focus is to propose a highly accurate SLAM solution that can provide more effective supervisory signals. Current DL-based supervised pose estimation methods may introduce errors into deep learning perception models, while differentiable SLAM has shown promising results and offers better accuracy for our task. Therefore we do not currently use DL-based pose estimation methods for our work.

2 Problem Setup

2.1 Model

Our framework is composed of two primary components: a generic deep learning module and a differentiable SLAM module. These are coupled together to allow training of the entire architecture in an end-to-end fashion. We focus on optimizing the overall loss function - the sum of the loss for the deep learning model and the SLAM module, denoted as

$$Loss = L_{model} + \gamma L_{slam}, \quad (1)$$

where γ is a coefficient balancing the impact of the SLAM on deep learning. The loss function for the deep learning model L_{model} is defined as

$$L_{model} = \frac{1}{n} \sum_{i=1}^n l(f(X_i; \theta), Y_i) + \beta R(\theta), \quad (2)$$

where n is the number of training samples, f is the deep learning model with parameters θ , l is a loss function, X_i is the input to the model, Y_i is the ground truth output, β is a regularization parameter, and $R(\theta)$ is the regularization term.

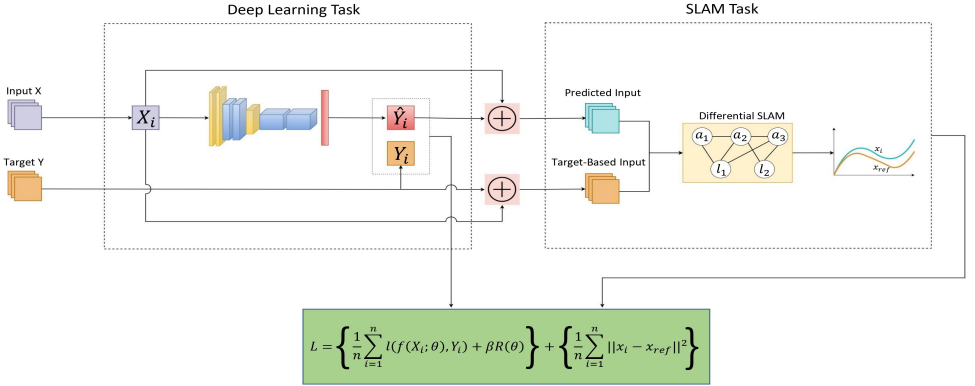


Figure 1: Our proposed framework for integrating differentiable SLAM into deep learning task training. The framework includes a Deep Learning Task which takes the input scans (X_i) and outputs the predicted scans (\hat{Y}_i). In the SLAM task, the trajectory error is calculated between trajectories x_i (calculated from predicted scans \hat{Y}_i) and x_{ref} (calculated from target scans Y_i). The framework aims to optimize the combined loss of both Deep Learning loss and SLAM (Absolute Trajectory Loss between trajectories) Loss.

The SLAM loss L_{slam} includes both translational and orientational errors. These are used to calculate the SLAM loss between the trajectories, which is defined as

$$L_{slam} = \frac{1}{n} \sum_{i=1}^n \|x_i - x_{ref}\|^2, \quad (3)$$

x_i is the estimated way-point of the trajectory generated by the SLAM algorithm using the predicted output of the deep learning model. x_{ref} is the estimated ground truth way-point, generated by estimating the trajectory from the input LiDAR sequence instead of using actual ground-truth pose estimates that are available in the dataset. This enables the differentiable SLAM framework to work in a self-supervised fashion. It makes our method practically viable and efficient.

2.2 Learning with SLAM

To enhance the performance of deep learning by optimizing SLAM errors, it is necessary for the SLAM to be fully differentiable. Our differentiable SLAM module takes two branches of input and predicts one trajectory for each of them. First - the outputs of the deep learning model (e.g. generated LiDAR scan, predicted segmentation mask) are input to the differentiable SLAM as input and the differentiable SLAM predicts a trajectory based on deep learning's output. Second - ground truth LiDAR information (e.g. LiDAR scan with only static points annotated, ground truth segmentation masks) are input to the differentiable SLAM, and the differentiable SLAM predicts a ground truth trajectory using them.

Classical SLAM systems [25] are non-differentiable. A common technique in these systems - the non-linear optimization is based on the Levenberg-Marquardt algorithm [27]. It switches the damping factor discretely at each iteration of the optimization process. This stops the gradient from backpropagating to the nodes when we unroll the optimization iter-




Figure 2: Left: Setup of DSLR with Differentiable SLAM Loss. Right: Setup of GndNet with Differentiable SLAM Loss.

ations to build the computational graph [14]. We use the generalized logistic function [62] for soft switching of the damping factor as well as the optimization update [14]

$$\lambda = \lambda_{\min} + \frac{\lambda_{\max} - \lambda_{\min}}{1 + De^{-\sigma(r_1 - r_0)}}, \quad (4)$$

$$x_{t+1} = x_t + \frac{\delta_t}{1 + e^{-(r_1 - r_0)}}, \quad (5)$$

λ_{\max} and λ_{\min} are damping coefficient bounds in Levenberg-Marquardt solvers. r_0 and r_1 represent error norms at the current and lookahead iterates. D and σ are tunable parameters.

3 Differentiable SLAM Integration

We now discuss the methodology of integrating SLAM into DLMP training tasks. In general, differentiable SLAM can be backpropagated as additional information in a deep learning model to help train the model for better performance. To this end, we study three LiDAR-related tasks - Ground v/s non-Ground Segmentation, Ground Elevation Estimation, the Dynamic to Static LiDAR translation, and Generative modelling for LiDAR.

3.1 Ground Elevation Estimation and Ground Segmentation

3.1.1 GndNet

Ground Elevation Estimation for LiDAR scans is crucial for tasks like navigable space detection, registration, to name a few. GndNet [28] estimates the ground elevation information as well as segments the LiDAR points into ground and non-ground (object/obstacle) points. We adapt their models to train with differentiable SLAM error along with the existing loss function. The goal is to achieve better estimates of ground plane elevation and classification into ground v/s non-ground points, using differentiable SLAM error. GndNet discretizes the raw point cloud into a evenly spaced $x - y$ grid, without binning the z -dimension (here the x , y , z direction refer to orientation of the LiDAR points in 3D coordinates), thereby creating a set of pillars [18]. Next, PointNet [29] is used to generate features for every non-empty pillar. Then, these pillar features are placed on the $x - y$ grid leading to a psuedo-image. Finally, a convolutional encoder-decoder network learns features from this image and regresses the ground elevation per cell in the grid. *This regression output is compared against the ground-truth elevation to compute L_{model} .* Further based on the elevation, points above a threshold

are classified as obstacle/object/above ground, while points below threshold are classified as ground points, thereby segmenting into ground and obstacle class.

3.1.2 Differentiable SLAM based GndNet

We insert our differentiable SLAM module after the regression of ground elevation per cell. Using the elevation output, we extract the corresponding points in the LiDAR scan that are classified as above ground (via a threshold parameter used by GndNet). We consider this as a predicted LiDAR scan (\tilde{L}) generated based on thresholding of the elevation output of the model. Given that the dataset also has groundtruth elevation information, we generate a groundtruth LiDAR scan (L), by thresholding the LiDAR points with groundtruth elevation information (using the same threshold parameter). Given we have a batch of contiguous groundtruth and predicted LiDAR using the above strategy, we can use the differentiable SLAM error module to estimate the trajectory for both the batches. Thereafter, we evaluate the rotational and translation trajectory error between both the trajectories (L_{slam}). This is our SLAM error that can be successfully backpropagated through the network owing to the differentiable SLAM module. For a visual description, refer Figure 2.

3.2 Dynamic to Static LiDAR Translation

3.2.1 DSLR

We choose a generative modelling application to show the effect of the differentiable SLAM on a generative modelling task. Dynamic to Static Translation of LiDAR point cloud [16] translates a LiDAR scan with occlusions due to dynamic objects, to a fully static scan with all dynamic occlusions replaced by static background. *DSLR* uses a 3-module based model - an Autoencoder, Pair Discriminator and an Adversarial Module to achieve the translation. Given a set of dynamic scans $X = (X_1, X_2 \dots X_n)$ and corresponding static scans $Y = (Y_1, Y_2 \dots Y_n)$, the autoencoder simply learns to reconstruct a LiDAR scan. The pair discriminator module classifies a LiDAR scan pair into 2 classes based on the below equation.

$$DI((l_1), (l_2)) = \begin{cases} 1 & l_1 \in X, l_2 \in X \\ 0 & l_1 \in X, l_2 \in Y \end{cases} \quad (6)$$

The adversarial modules trick the discriminator to predict 1 for a pair that should be labelled as 0, thus generating the adversarial loss (L_{model}), which helps to achieve static translation for a dynamic LiDAR scan.

3.2.2 Differentiable SLAM based DSLR

We modify the adversarial module of DSLR in order to plug differentiable SLAM. Given a dynamic scan (d_i) as input, the output of the adversarial module is a reconstructed static scan (\bar{s}_i) with the dynamic occlusions replaced by the actual static background. We also have the groundtruth static scans (s_i) to compare the generated scans against. Given that we have a batch of contiguous reconstructed static scans ((\bar{s}_i)), as well as the groundtruth static scans (s_i), we can use the differentiable SLAM error module to calculate the trajectories for both the sets and calculate the error between the two (L_{slam}), that can be backpropagated using the deep learning model. For a visual description, refer Figure 2.

3.3 Conditional LiDAR Generation (AE_{lidar})

We demonstrate the benefit of our differentiable SLAM module on a standard LiDAR autoencoder which serves as backbone for multiple downstream tasks. We use Caccia et. al. [9] to reconstruct LiDAR scan using range-image based LiDAR representation. The encoder and decoder architectures are adapted from Radford et. al. [5]. We train AE_{lidar} with and without differentiable SLAM. For the differentiable SLAM variant, we calculate the SLAM loss between the reconstructed output LiDAR and the input LIDAR scan. The SLAM loss along with the reconstruction loss is back-propagated to ensure that the model learns from the SLAM error as well. The pipeline for this task is similar to DSLR (Figure 2(left)).

4 Experiments

4.1 Experimental Setup

For all the deep learning models used in our paper: DSLR[16], GndNet [28] and AE_{lidar} [9] we follow the experimental setting of the respective models as used in their work, except a minor change for GndNet. GndNet uses every fourth contiguous LiDAR scan for training. However, we require finer contiguity as we compute SLAM error between contiguous scans. Therefore, we use every second contiguous scan for training GndNet and run the experiments with and without the SLAM module using this setting to report the results.

SLAM error module is time consuming. Thus, we do not calculate SLAM error for every epoch - we calculate SLAM error after every k^{th} epoch, where k is a hyperparameter. More details in the Appendix (Section 1.2).

4.2 Datasets

CARLA-64: CARLA-64 [16] is an extensive simulated LiDAR dataset. It mimics the exact settings of a VLP-64 LiDAR sensor. The dataset consists of 8 sequences for training and 6 for testing. It has 4 sequences for testing on SLAM.

ARD-16: ARD-16 [16] is a real-world sparse industrial dataset collected using a VLP-Puck LiDAR sensor. It is $4\times$ sparse compared to CARLA-64 and KITTI.

SemanticKITTI: SemanticKITTI [10][9] is a well known LiDAR dataset with semantic segmentation labels. It has 11 sequences(00-10) for which semantic labels are available.

For more details on the datasets, please refer to the Appendix (Section 1.1).

4.3 Results

Ground Elevation Estimation and Segmentation We, for the first time show the possibility of integration of a differentiable SLAM module into a regression task (ground elevation estimation) and segmentation (into ground v/s non-ground points) based downstream task. We compare the results of the GndNet with and without differentiable SLAM error as explained in Section 3. The model performs 2 task - regressing the ground elevation of the scan points, and segmenting the points into ground and above ground points. The performance of the model is shown in the Table 1. Our variant comfortably surpasses the MSE (Mean-Squared Error) estimate while regressing the elevation of the points with an *improvement of 0.04* on MSE. Also our model fares better than GndNet on recall - increase of 3% and makes




Figure 3: SLAM using (1) static output scans of DSLR(left) and (2) with static output from DSLR+Differentiable SLAM Error model (right). Red boxes indicate regions where the trajectory generated from our framework is better than the baseline. Differentiable SLAM helps to reduce the drift accumulated while navigation.

less false positive mistakes. This improvement is gained by using the differentiable SLAM module for only 27 epochs out of the total 150 epochs. (Refer 4.1) .

Method	Frames	MSE	mIOU	Prec	Recall
GndNet	6554	0.76	0.81	0.85	0.94
GndNet+Diff SLAM	6554	0.72	0.81	0.83	0.97

Table 1: Comparison of the Differentiable SLAM module on ground elevation estimation task and segmentation. For MSE, lower the better, for the rest, opposite.

4.3.1 Dynamic to Static LiDAR Translation for SLAM

In this section we for the first time show the application of differentiable SLAM for a generative modelling task - Dynamic to Static Translation for LiDAR scan for effective SLAM [17]. In this task, we evaluate the relative benefit of using differentiable SLAM over plain DSLR. We also compare the reconstruction quality of the static translations with the ground truth static scan using Chamfer’s Distance[18]. As we see in Table 2, with differentiable SLAM the Chamfer Distance is always better. Here we would like to point that we add SLAM error as a loss term only for 7 interleaved epochs (Section 4.1) which gives a meaningful reduction in the error. For results on all the six CARLA test sequences, please refer to Table 1 in the Appendix.

We further investigate the effect of integrating the differentiable SLAM module in DSLR on downstream SLAM. As we see in Table 3 and Figure 3, using static reconstructions obtained from differentiable SLAM integrated DSLR gives reduced navigation error on all the four LIDAR SLAM sequences for CARLA-64.

Dataset	Run	DSLRL with Diff. SLAM	DSLRL without Diff. SLAM
CARLA-64	Avg [9..14]	6.96	7.85
ARD-16	3	0.31	0.34
KITTI	8	5.00	5.23

Table 2: Relative Comparison of Static Translation(using Chamfer Distance) for DSLR with and without Differentiable SLAM. Lower the better.

Run	With Diff SLAM			Without Diff SLAM		
	ATE	RPE		ATE	RPE	
		Trans	Rot		Trans	Rot
CARLA-64 Dataset						
0	2.37	0.440	0.09	4.73	0.440	0.11
1	1.3	0.400	0.070	2.9	0.400	0.070
2	0.76	0.567	0.07	1.36	0.571	0.15
3	4.09	0.399	0.081	4.4	0.395	0.104
ARD-16 Dataset						
3	1.94	4.81	0.186	2.05	4.81	0.188

Table 3: Relative Comparison of Static Translation for DSLR with and without Differentiable SLAM on CARLA-64 and ARD-16 dataset. We use Chamfer Distance metric.

CARLA-64	Chamfer Distance with SLAM	Chamfer Distance without SLAM
Avg[9..14]	2.83	3.03

Table 4: Comparison of a general purpose generative autoencoder, AE_{lidar} with and without differentiable SLAM using Chamfer’s Distance metric.

4.3.2 LiDAR Reconstruction using AE_{lidar}

We demonstrate the effect of our differentiable SLAM on a general purpose simple generative model in Table 4. Our result demonstrates that such general purpose model that is used in several complex models as a *backbone* can benefit from differentiable SLAM. For detailed results on all the CARLA test sequences, please refer to Table 2 in the Appendix.

5 Discussion and Limitations

Certain limitations of Differentiable SLAM are discussed in the Appendix (Section 1.4).

LiDAR based applications has seen significant progress with the development of new techniques and technologies that have revolutionized the field. One such technique is SLAM, which is a popular approach to map an unknown environment and localize a robot within it. In this paper, we propose a novel method that uses differentiable SLAM to improve the performance of deep learning tasks such as binary segmentation, generative modeling, and regression. Our core idea lies in the fact that SLAM prefers certain properties over others, such as static structures/non-ground points over dynamic/ground ones. We assume that the reference trajectory provided to SLAM is close to the ground truth, which enables us to minimize SLAM loss in a way that is equivalent to minimizing with ground truth poses. By doing so, we encourage DSLR to give more static-like predictions, and segmentation models to make clear distinctions between ground and non ground objects. Additionally, we use SLAM to improve elevation regression so that ground points can be deleted, which improves SLAM performance. Our approach is based on a two-step reasoning process - we first assume that the reference trajectory provided to SLAM is close to the ground truth, and then we exploit the properties that SLAM prefers to improve deep learning tasks. We argue that SLAM preferences can be used to improve the performance of deep learning tasks. We present empirical results that demonstrate the effectiveness of our approach. Overall, we believe that our approach has the potential to significantly advance the field of robotics and open up new avenues for research in this exciting area.

References

- [1] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9297–9307, 2019.
- [2] Daan Bloembergen and Chris Eijgenstein. Automatic labelling of urban point clouds using data fusion. *arXiv preprint arXiv:2108.13757*, 2021.
- [3] Michael Bloesch, Jan Czarnowski, Ronald Clark, Stefan Leutenegger, and Andrew J. Davison. Codeslam - learning a compact, optimisable representation for dense visual slam. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2560–2568, 2018. doi: 10.1109/CVPR.2018.00271.
- [4] L. Caccia, H. v. Hoof, A. Courville, and J. Pineau. Deep generative modeling of lidar data. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5034–5040, 2019.
- [5] Kang Chen, Yu-Kun Lai, Yu-Xin Wu, Ralph Martin, and Shi-Min Hu. Automatic semantic modeling of indoor scenes from low-quality rgb-d data using contextual information. *ACM Transactions on Graphics*, 33(6), 2014.
- [6] Xieyuanli Chen, Shijie Li, Benedikt Mersch, Louis Wiesmann, Jürgen Gall, Jens Behley, and Cyrill Stachniss. Moving object segmentation in 3d lidar data: A learning-based approach exploiting sequential data. *IEEE Robotics and Automation Letters*, 6(4):6529–6536, 2021.
- [7] George Eskandar, Janaranjani Palaniswamy, Karim Guirguis, Barath Somashekar, and Bin Yang. Glpu: A geometric approach for lidar pointcloud upsampling. *arXiv preprint arXiv:2202.03901*, 2022.
- [8] Aryan FNU, Dheeraj Vattikonda, Erqun Dong, and Sabyasachi Sahoo. Grad-lidar-SLAM: Fully differentiable global SLAM for lidar with pose-graph optimization. In *IROS 2022 Workshop Probabilistic Robotics in the Age of Deep Learning*, 2022.
- [9] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012.
- [10] Benoît Guillard, Sai Vemprala, Jayesh K Gupta, Ondrej Miksik, Vibhav Vineet, Pascal Fua, and Ashish Kapoor. Learning to simulate realistic lidars. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8173–8180. IEEE, 2022.
- [11] Hao-Su. 3d deep learning on point cloud representation (analysis), 2017. URL [http://graphics.stanford.edu/courses/cs468-17-spring/LectureSlides/L14%20-%203d%20deep%20learning%20on%20point%20cloud%20representation%20\(analysis\).pdf](http://graphics.stanford.edu/courses/cs468-17-spring/LectureSlides/L14%20-%203d%20deep%20learning%20on%20point%20cloud%20representation%20(analysis).pdf).

- [12] Chenhong He, Hui Zeng, Jianqiang Huang, Xian-Sheng Hua, and Lei Zhang. Structure aware single-stage 3d object detection from point cloud. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11873–11882, 2020.
- [13] Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and Andrew Markham. Randla-net: Efficient semantic segmentation of large-scale point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11108–11117, 2020.
- [14] Krishna Murthy Jatavallabhula, Ganesh Iyer, and Liam Paull. Gradslam: Dense slam meets automatic differentiation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2130–2137, 2020. doi: 10.1109/ICRA40945.2020.9197519.
- [15] Hyun-Koo Kim, Kook-Yeol Yoo, and Ho-Youl Jung. Color image generation from range and reflection data of lidar. *Sensors*, 20(18):5414, 2020.
- [16] Prashant Kumar, Sabyasachi Sahoo, Vanshil Shah, Vineetha Kondameedi, Abhinav Jain, Akshaj Verma, Chiranjib Bhattacharyya, and Vinay Vishwanath. Dynamic to static lidar scan reconstruction using adversarially trained auto encoder. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 1836–1844, 2021.
- [17] Loic Landrieu and Martin Simonovsky. Large-scale point cloud semantic segmentation with superpoint graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4558–4567, 2018.
- [18] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12697–12705, 2019.
- [19] Seungjae Lee, Hyungtae Lim, and Hyun Myung. Patchwork++: Fast and robust ground segmentation solving partial under-segmentation using 3d point cloud. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 13276–13283. IEEE, 2022.
- [20] Hyungtae Lim, Minh Oh, and Hyun Myung. Patchwork: Concentric zone-based region-wise ground segmentation with ground likelihood estimation using a 3d lidar sensor. *IEEE Robotics and Automation Letters*, 6(4):6458–6465, 2021.
- [21] Ziqi Lu, Yihao Zhang, Kevin Doherty, Odin Severinsen, Ethan Yang, and John Leonard. Slam-supported self-training for 6d object pose estimation. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2833–2840. IEEE, 2022.
- [22] Kaj Madsen, Hans Bruun Nielsen, and Ole Tingleff. *Methods for non-linear least squares problems*. 2004.
- [23] Nathaniel Merrill, Yuliang Guo, Xingxing Zuo, Xinyu Huang, Stefan Leutenegger, Xi Peng, Liu Ren, and Guoquan Huang. Symmetry and uncertainty-aware object slam for 6dof object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14901–14910, 2022.

- [24] Andres Milioto, Ignacio Vizzo, Jens Behley, and Cyrill Stachniss. Rangenet++: Fast and accurate lidar semantic segmentation. In *2019 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 4213–4220. IEEE, 2019.
- [25] Raúl Mur-Artal, J. M. M. Montiel, and Juan D. Tardós. ORB-SLAM: a versatile and accurate monocular SLAM system. *IEEE Transactions on Robotics*, 31(5):1147–1163, 2015. doi: 10.1109/TRO.2015.2463671.
- [26] Kazuto Nakashima and Ryo Kurazume. Learning to drop points for lidar scan synthesis. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 222–229. IEEE, 2021.
- [27] Kazuto Nakashima, Yumi Iwashita, and Ryo Kurazume. Generative range imaging for learning scene priors of 3d lidar data. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1256–1266, 2023.
- [28] Anshul Paigwar, Özgür Ercent, David Sierra-Gonzalez, and Christian Laugier. Gndnet: Fast ground plane estimation and point cloud segmentation for autonomous vehicles. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2150–2156, 2020. doi: 10.1109/IROS45743.2020.9340979.
- [29] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
- [30] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 918–927, 2018.
- [31] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- [32] Francis J Richards. A flexible growth function for empirical use. *Journal of experimental Botany*, 10(2):290–301, 1959.
- [33] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointcnn: 3d object proposal generation and detection from point cloud. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 770–779, 2019.
- [34] Shaoshuai Shi, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network. *IEEE transactions on pattern analysis and machine intelligence*, 43(8):2647–2664, 2020.
- [35] Paloma Sodhi, Eric Dexheimer, Mustafa Mukadam, Stuart Anderson, and Michael Kaess. Leo: Learning energy-based models in factor graph optimization. In *Conference on Robot Learning*, pages 234–244. PMLR, 2022.
- [36] Keisuke Tateno, Federico Tombari, Iro Laina, and Nassir Navab. Cnn-slam: Real-time dense monocular slam with learned depth prediction. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6565–6574, 2017.

- [37] Larissa T Triess, Christoph B Rist, David Peter, and J Marius Zöllner. A realism metric for generated lidar point clouds. *International Journal of Computer Vision*, 130(12): 2962–2979, 2022.
- [38] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018.
- [39] Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 3dssd: Point-based 3d single stage object detector. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11040–11048, 2020.
- [40] Brent Yi, Michelle A Lee, Alina Kloss, Roberto Martín-Martín, and Jeannette Bohg. Differentiable factor graph optimization for learning smoothers. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1339–1345. IEEE, 2021.
- [41] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11784–11793, 2021.
- [42] Yang Zhang, Zixiang Zhou, Philip David, Xiangyu Yue, Zerong Xi, Boqing Gong, and Hassan Foroosh. Polarnet: An improved grid representation for online lidar point clouds semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9601–9610, 2020.
- [43] Shuaifeng Zhi, Michael Bloesch, Stefan Leutenegger, and Andrew J. Davison. Scenecode: Monocular dense semantic reconstruction using learned encoded scene representations. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11768–11777, 2019.
- [44] Xinge Zhu, Yuexin Ma, Tai Wang, Yan Xu, Jianping Shi, and Dahua Lin. Ssn: Shape signature networks for multi-class object detection from point clouds. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXV 16*, pages 581–597. Springer, 2020.
- [45] Vlas Zyrianov, Xiyue Zhu, and Shenlong Wang. Learning to generate realistic lidar point clouds. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIII*, pages 17–35. Springer, 2022.