Terrain Traversability Prediction for Off-Road Autonomous Driving

Anonymous Author(s)
Affiliation
Address
email

Abstract: In this paper, we focus on the problem of estimating the *traversability* of terrain for autonomous off-road navigation. To produce dense and accurate local traversability maps, a robot must reason about geometric and semantic properties of the environment. To achieve this goal, we develop a novel *Bird's Eye View Network* (BEVNet), a deep neural network that directly predicts dense traversability maps from sparse LiDAR inputs. BEVNet processes both geometric and semantic information in a temporally consistent fashion. More importantly, it uses learned prior and history to *predict* traversability in unseen space and into the future, allowing a robot to better appraise its situation. We quantitatively evaluate BEVNet on both on-road and off-road scenarios, and show that it outperforms a variety of strong baselines.

Keywords: Off-road Driving, Autonomous Driving, Deep Learning, Perception

1 Introduction

While there has been great recent interest in the development of autonomous vehicles, the vast majority of this work has focused on *on-road* and *urban* driving. However, a wide range of application areas including defence, agriculture, conservation, and search and rescue could benefit from autonomous *off-road* vehicles that can operate in complex, natural terrain. In such environments, understanding the *traversability* of terrain surrounding the vehicle is crucial for successful planning and control. Perceiving whether terrain is actually traversable from sparse LiDAR data can be a challenging problem as off-road terrain is often characterized by rapid changes to the ground plane, heavy vegetation, overhanging branches, and negative obstacles. In other words, a successful off-road robot must reason about both the geometric and semantic content of its surroundings in order to determine what terrain is traversable and what is impassible.

An effective terrain traversability prediction system should efficiently 1) aggregate the observations over time [1, 2] with noisy odometry [3, 4], 2) reason about the partially seen or even yet to be seen parts of the environment [5, 6, 7], and 3) detect overhanging structures in the environment, such as tree branches, tunnels, and power-lines [8, 9]. While previous work has addressed each of these issues individually, the aforementioned challenges are related, and solving each one should benefit the others.

In this paper, we propose a *Bird's Eye View Network* (BEVNet), a recurrent neural network that directly predicts the traversability of the terrain in the form of a 2D grid around the robot from LiDAR scans. As shown in Figure 1, our model has three main parts: 1) a 3D sparse convolution subnetwork to process the voxelized point cloud, 2) a Convolutional Gated Recurrent Unit (ConvGRU) which uses convolutional layers in gated recurrent unit [10] to aggregate the 3D information, 3) a 2D convolutional encoder-decoder with efficient backbone [11] that simultaneously inpaints the empty spaces and projects the 3D data into the 2D Bird's Eye View (BEV) map. To train the model, we use both past and future labeled LiDAR scans to build a complete 3D semantic point cloud and build the ground-truth 2D traversability map. Unlike previous work [8], our model does not use any explicit rule to find the overhanging structures when doing the projection. We detect and remove any overhanging structures when *building the ground-truth BEV map* which can be done efficiently

Submitted to the 5th Conference on Robot Learning (CoRL 2021). Do not distribute.

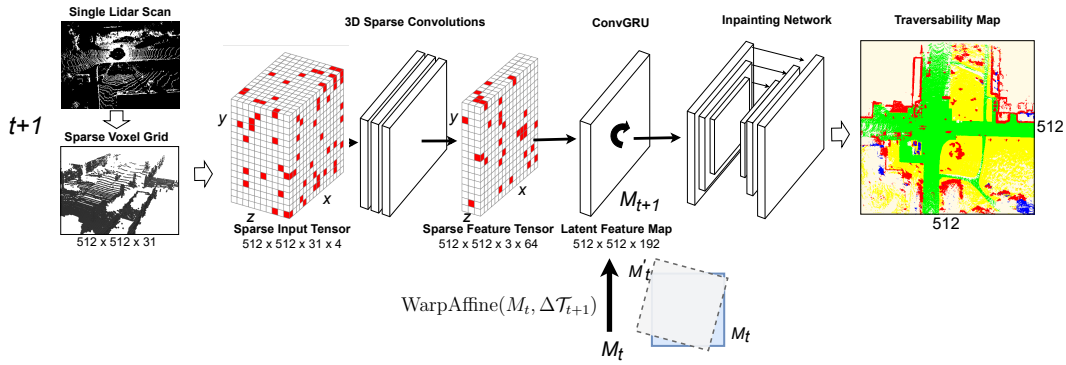

Figure 1: The network architecture of BEVNet. The incoming LiDAR scan is first discretized into a sparse voxel grid, which is then fed into a sequence of sparse convolution layers to compress the z dimension. The compressed sparse feature tensor is aggregated over time via the ConvGRU unit. We use differentiable affine warping to align the latent feature map with the current odometry frame. Finally, the inpainting network "inpaints" the latent map to output a dense traversability map.

as the environment is fully observed and labeled. The network *learns* what should be considered as a overhanging obstacle given sparse LiDAR scans.

We make several contributions and empirical observations. We proposed a novel framework to build the BEV map by simultaneously 1) aggregating observations over time, 2) predicting the unseen areas of the map, and 3) filtering out irrelevant obstacles like overhanging tree branches that do not affect traversability. Experimental results on SemanticKITTI [12] and RELLIS-3D [13] demonstrate the novel framework works well in both on-road and off-road settings, and outperforms the strong baselines in all the aforementioned tasks.

2 Related Work

While there exists a vast literature on perception for autonomous driving [14, 15, 16, 17, 18, 19, 20, 21], most prior work focuses on *urban* environments and leverage large datasets and structure inherent in cities and road networks. Since there are not as many recent papers focused on off-road driving and rural environments, we compare our system to prior works that share the mutual components to ours, namely traversability analysis [22, 23, 24, 25], online semantic mapping [15, 16, 1, 17, 21], recurrency handling [15, 1, 17], bird's eye view semantic segmentation [1, 19, 26, 27], and semantic scene completion [5, 6, 7].

Traversability Analysis Traversability Analysis is less common in the task of urban autonomous driving but crucial for a successful off-road autonomy [23, 28, 29]. Traversability of surrounding terrain may be analyzed based on various criteria, including surface roughness [22], negative obstacles [23], and terrain classification [24, 25]. Our system performs end-to-end semantic mapping by classifying each surrounding location based on traversabilty.

Online Semantic Mapping Semantic mapping provides structured information for an autonomous driving system, with the earliest works in this area dating back to well before deep-learning driven methods gained traction [30]. A core paradigm is the bird's eye view representation, which stores local information in a two-dimensional grid surrounding the vehicle [31]. While prior works such as [14, 15, 16] utilize a high definition map as additional prior knowledge, such a map may not be available for off-road terrain. Our system therefore does not assume a map *a priori*, and instead constructs one online. Similar to our system, Casas et al. [17], Wu et al. [21] learns to produce semantic maps for urban autonomous driving. In comparison, our system focuses on semantic mapping with a broader categorization of semantics based on terrain traversability, which is useful for off-road driving.

Recurrent Representations Temporally consistent accumulation of semantic information is crucial for mapping the environment and providing the information necessary for safe, efficient motion planning. Recent work [15, 17] directly concatenates the past 10 voxelized LiDAR scans as its input representation for memory efficiency. In comparison to our system, the author's recurrent architec-

ture is specifically designed for semantic occupancy forecasting, whereas our recurrent architecture accumulates sensor data to better estimate the traversability of the current surrounding terrain. Maturana et al. [1] utilize a more conventional approach by accumulating information via Bayes filtering, which in our system is replaced with a recurrent neural network. There are various approaches to handling recurrency in neural networks including [2, 32, 33, 34, 35, 36, 37]. Our system utilizes ConvGRU to accumulate 2D representations for BEV mapping.

BEV Semantic Segmentation Philion and Fidler [19] learns to produce a BEV map from RGB camera for on-road driving, and Maturana et al. [1] produces a BEV map for off-road driving by segmenting an image and projecting it using depth information. In comparison, our system learns to directly project a LiDAR scan to produce a BEV map. We also compare our approach to LiDAR segmentation where the segmented point cloud from networks such as [26, 27] can be projected onto a BEV map. We evaluate this comparison in detail in Sec. 5.

Semantic Scene Completion (SSC) The goal of SSC is to generate a complete 3D scene given a single LiDAR scan as input. Existing works such as [5, 6, 7] utilize information from semantic segmentation to complete the scene, whereas our system learns to directly predict the completed scene and therefore does not require segmentation prediction from a secondary network or ground truth labels. In addition, our system performs point cloud projection and scene completion in 2D simultaneously, as the main task in our work is to produce a 2D BEV map. We evaluate our system's ability to complete scenes by comparing our model against [5], the details can be found in Sec. 5.

3 Method

3.1 Overview

We consider a mobile robot with a 360° LiDAR mounted at its top. In order for the robot to navigate efficiently and safely in a new environment (either on-road or off-road), the robot builds an *online traversability map* around itself. The traversability map resembles a conventional occupancy map as well as the semantic map from [1], where each cell stores a probabilitity distribution of traversability labels. In this work, we use four levels of traversability: *free*, *low-cost*, *medium-cost* and *lethal*. The number of traversability levels can be trivially extended if so desired. The traversability map is inside the robot's odometry frame, so that the robot is always at the center, with its heading pointing to the east. The traversability map is converted to a costmap by mapping each traversability level to the corresponding cost value via a lookup table. The converted costmap can be easily interfaced with a local planner [38] or a global planner (e.g., A*) for finding the least-cost path to a goal.

We adopt a supervised-learning approach to predict this traversability map. We start by building a traversability dataset from LiDAR segmentation datasets [12, 13] via a traversability-aware projection procedure. Then, we introduce BEVNet, a recurrent neural network that takes the current LiDAR scan and utilizes its history to build a dense traversability map. In the following sections, we will describe each component in detail.

3.2 Building a Traversability Dataset

Recent work [15, 1] focuses on on-road driving where reasoning about a large number of fine-grained semantic classes is necessary. Here we consider a more general driving paradigm where we simply care about the *traversability* of the surrounding terrain. This makes our model applicable to both on-road and off-road driving. Given a dataset with semantically labeled LiDAR scans, we convert it to a traversability dataset via the following procedure (illustrated in Figure 2).

Scan Aggregation. For each scan, we aggregate it with the past t and the future t scans with stride s to construct a larger point set. We set t to a large enough number (e.g., 71) to obtain dense traversability information for a large area around the robot. These parameters may be tuned depending on the vehicle speed and density of the LiDAR points.

Traversability Mapping. We map the semantic classes into our 4-level ontology. The general principle is to map semantic classes with similar costs to the same traversability label. For example, *car* and *building* are mapped to *lethal*, whereas *mud* and *grass* are mapped to *low-cost*. Detailed mapping can be found in the supplementary.

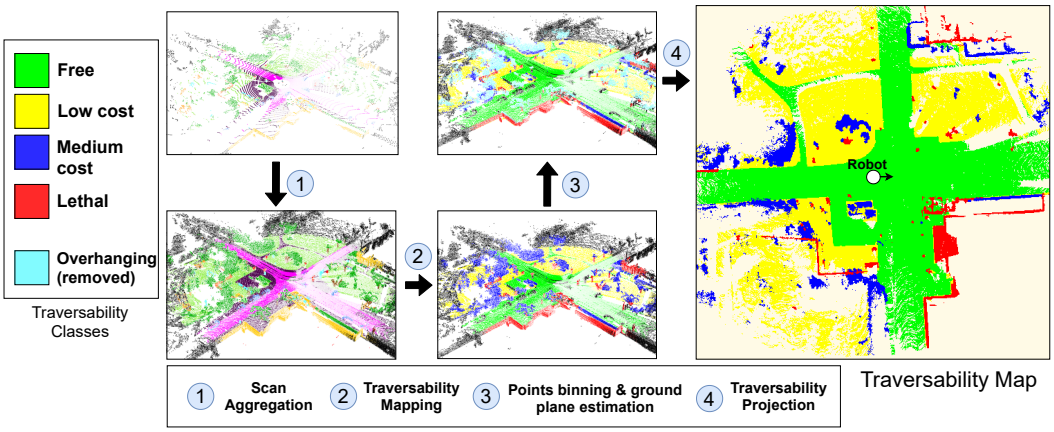

Figure 2: Process for generating the traversability dataset on SemanticKITTI. Single scan labels are aggregated to form complete scenes (scan aggregation), and their semantics are remapped to our 4 class ontology (traversability mapping), legend defined in figure. The remapped scans are then filtered with ground estimation to remove overhanging points (points binning and ground plane estimation). Finally, the filtered points are projected to a traversability map (traversability projection).

Points Binning and Ground Plane Estimation. For each point in the aggregated scan, we do a down projection to find its location x, y on the traversability map. Hence each x, y location of the map contains a pillar of points. We estimate the ground plane (e.g., a height map) by running a mean filter kernel over the lowest z coordinates of the points labeled as *free* and *low-cost* at each x, y location in the map. This ground plane is used as a reference for final traversability projection.
Traversability Projection. For each pillar of points, we filter out overhanging obstacles by removing points that are above the ground plane by a certain threshold because they will not collide with the robot. For the remaining points, we take the point with the lowest traversability at each x, y location as the final traversability label.

3.3 Feature Extraction via Sparse Convolution with Z Compression

The architecture of BEVNet is shown in Figure 1. An input LiDAR scan is first discretized into a $512 \times 512 \times 31$ grid with a resolution of $0.2m$. We perform sparse discretization so that only occupied voxels are preserved. Each voxel contains a 4-dimensional feature $f = \frac{1}{n}\sum_{i=1}^{n}[x_i, y_i, z_i, r_i]$, which is the average of the coordinates and remission values of the points inside the voxel. This sparse voxel grid is fed into a sequence of sparse convolution layers, which compress the z dimension via strided convolutions. We keep x and y dimensions unchanged. The output of the sparse convolution layers is a sparse feature tensor S of size $512 \times 512 \times C$, where C is the feature dimension.

3.4 Temporal Aggregation of Sparse Feature Maps

A single LiDAR scan becomes increasingly sparse as the distance increases, making it difficult to estimate traversability for areas far away from the robot. Contrary to classical SLAM that aggregates LiDAR measurements over time via a hand-engineered Bayesian update rule [39], we let the network learn to aggregate the sparse feature maps from past LiDAR scans via a Convolutional Gated Recurrent Unit (ConvGRU). The ConvGRU maintains a 2D latent feature map M that shares the same coordinate system and dimensions as the final traversability map. The latent feature map M is updated as

$$M_{t+1} = \text{ConvGRU}(\text{WarpAffine}(M_t, \Delta\mathcal{T}_{t+1}), S_{t+1}),$$

where $\Delta\mathcal{T}_{t+1}$ is the relative transform of the robot's odometry frame from t to $t+1$. The WarpAffine operation transforms the latent feature map M_t from the previous odometry frame to the current odometry frame so that the features from M_t and S_{t+1} are spatially aligned. Note that the WarpAffine operation is differentiable to allow the gradients to backpropagate through time.

3.5 Traversability Inpainting

Since the ConvGRU only aggregates sparse feature tensors, M contains little information for areas where there is no LiDAR point. Instead of treating no-hit area as unknown, we let the network fill in the empty space by leveraging the local and global contextual cues via the *Inpainting Network*. The inpainting network is a fully convolutional network inspired by FCHardNet [11] which is originally designed for fast image segmentation. It consists of a sequence of downsampling and upsampling layers with skip connections, making it effective for capturing local and global contextual information for predicting what is missing.

4 Implementation Details

We build the traversability datasets from SemanticKITTI [12] and RELLIS-3D [13] to evaluate BEVNet in both on-road and off-road scenarios. For SemanticKITTI we aggregate 71 frames with stride 2 to generate a single traversability map. For RELLIS-3D we aggregate 141 frames with stride 5. Both datasets provide per-frame odometry, which we use for the differential warping layer in the ConvGRU. The traversability maps have a size of $102.4m \times 102.4m$ with a resolution of $0.2m$. Note that the traversability maps contain an additional "unknown" class marking regions that have never been observed.

Network training. We train our network using the Adam optimizer [40] with an initial learning rate of $3e-4$ and a decay of 0.7 per epoch. We use a weighted Cross-Entropy loss. We start with training a single-frame model without ConvGRU until the model converges. Then we freeze the sparse convolution layers and insert the ConvGRU layer, and then train the ConvGRU and the inpainting network together. While technically we can train the whole network end-to-end , this two-stage training procedure is faster and is more memory-efficient. When training the ConvGRU, we use a sequence length of 5 with a frame stride randomly chosen from [1, 10, 20]. Training takes about 12 hours on a single RTX 3090. The inference time of our network is 6 fps on a RTX 3090.

Data augmentation. During training, we randomly rotate every pair of LiDAR scans and the ground truth traversability map in $\mathcal{U}[-45°, 45°]$, and randomly drop 20% of the points. Furthermore, we perturb the groundtruth odometry with rotation drawn from $\mathcal{N}(0, 0.01^2)$ and translation drawn from $\mathcal{N}(0, 0.1^2)$. Note that the error in odometry will accumulate over time. We evaluate the effect of noisy odometry in Sec. 5.

5 Experiments

We conduct both quantitative and qualitative study on SemanticKITTI (on-road) and RELLIS-3D (off-road) datasets. We trained a separate model for each dataset. We compare with a variety of baselines, ranging from LiDAR segmentation to scene completion on the validation sequences. We also perform an ablation study to better understand the contribution of recurrence, and how our model behaves on the two datasets that have very different characteristics.

5.1 Evaluation Metrics

We use the mean Intersection of Union (mIoU) [11], a widely used metric for image segmentation, as the quantitative measure of the prediction accuracy. Note that our model predicts an additional "unknown" class to improve the visual consistency, and we exclude the "unknown" class in the evaluation. To better understand our model's capability of predicting the future, we report mIoUs in three modes: **seen**, **unseen**, and **all**. In the "seen" mode, we do not include ground truth labels obtained from future frames, effectively excluding any future predictions. For the "unseen" model we only include the future predictions. In the "all" scenario, we evaluate on both.

5.2 Comparison with LiDAR Segmentation with Temporal Aggregation

A strong baseline for building a traversability map is to perform semantic segmentation of the incoming LiDAR scan, project it down to obtain a 2D sparse traversability map, and aggregate the

	SemanticKITTI			RELLIS-3D		
	All	Seen	Unseen	All	Seen	Unseen
BEVNet-S	0.416	0.465	0.308	0.559	0.518	0.545
Clean Odometry						
BEVNet-TA	0.468	0.534	0.335	0.615	**0.605**	0.586
BEVNet-R	**0.480**	0.547	**0.344**	**0.618**	0.577	**0.620**
Cylinder3D-TA	0.465	**0.655**	N/A	0.411	0.568	N/A
Noisy Odometry						
BEVNet-TA	0.379	0.415	0.310	0.452	0.372	0.560
BEVNet-R	**0.468**	**0.529**	**0.343**	**0.614**	**0.572**	**0.616**
Cylinder3D-TA	0.342	0.455	N/A	0.318	0.435	N/A

Table 1: Mean IoU of different methods on SemanticKITTI and RELLIS-3D.

traversability maps over time. To compare with this approach, we choose Cylinder3D [26] (fine-tuned on our 4-class ontology) as the LiDAR segmentation network for its strong performance, and use the same projection procedure in Sec 3 on the input LiDAR scan to obtain the single-frame traversability map. We perform the temporal aggregation by tracking the categorical distribution of traversability via a uniform Dirichlet prior. To do so, we keep a counter map M_C of size $H \times W \times 4$ (initialized to zeros). It is of the same size as the traversability map except that the last dimension counts the traversability labels observed so far. We update M_C incrementally. For each incoming single-frame traversability map, we warp M_C to the current odometry frame via bilinear interpolation, and increment the counts by adding the one-hot version of the incoming single-frame map. The actual traversability label can be obtained by taking the *argmax* of the last dimension of M_C.

Results on SemanticKITTI. In the left half of Table 1 we compare the performance of BEVNet-Recurrent (BEVNet-R) with Cylinder3D+Temporal Aggregation (C3D-TA) on the SemanticKITTI validation set. When only considering what has been observed so far ("seen") and clean odometry, C3D-TA is better than BEVNet-R. This shows that LiDAR segmentation with accurate temporal aggregation can work very well in structured environments such as on-road driving. When evaluated on the full groundtruth ("full"), BEVNet-R outperforms C3D-TA because C3D-TA cannot predict the future traversability. When evaluating on noisy odometry, BEVNet-R surpasses C3D-TA for both "seen" and "full" test scenarios. BEVNet-R uses learned recurrency to "fix" the errors in odometry and to adaptively forget history in case the error in odometry is too large. In comparison, C3D-TA solely uses the provided odometry to aggregate information, which may result in large misalignment as errors accumulate over time.

Results on RELLIS-3D. The results on RELLIS-3D (right half of Table 1) share a similar trend as those in SemanticKITTI, except that BEVNet-R consistently outperforms C3D-TA with a larger gap. This suggests that off-road environment is more challenging, where accurate LiDAR segmentation is hard to obtain due to the lack of environmental structure. Indeed, Cylinder3D only achieves a 64.1 mIoU on RELLIS-3D for LiDAR segmentation, which is lower than the 87.9 mIoU on SemanticKITTI. Interestingly, noisy odometry has almost no impact on BEVNet-R. We hypothesize that it is because the RELLIS-3D dataset contains less clutter and occlusion so BEVNet-R does not rely heavily on the history for traversability prediction.

Qualitative results In left half of Figure 3, we highlight the fact that BEVNet-R can preserve small dynamic objects such as bicyclists better than C3D-TA. Hand-engineered temporal aggregation is prone to treating small dynamics objects as noise and ignoring them. In comparison, BEVNet can learn to keep small dynamic objects, while preserving smoothness in static regions. The right half shows the impact of noisy odometry. We can see large misalignment and smear artefact for C3D-TA, whereas BEVNet-R produces significantly cleaner output. Finally, in Figure 4 we visualize examples on both SemanticKITTI and RELLIS-3D. In general BEVNet-R shows strong performance in predicting future traversability. It learns to predict whole cars, alley entrances, and trail paths with extremely sparse LiDAR points.

5.2.1 Ablation Study

We conduct our ablation study on three variants of BEVNet: BEVNet-Single (BEVNet-S), BEVNet-Single+Temporal Aggregation (BEVNet-TA), and BEVNet-Recurrent (BEVNet-R). We aim to answer three questions: 1) is learned recurrence better than temporal aggregation? 2) does history help

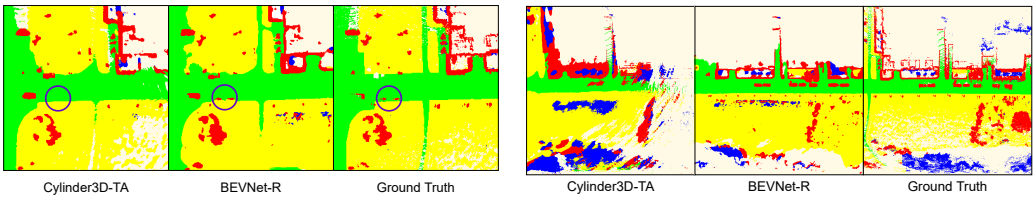

Figure 3: Qualitative comparison of our method and baseline. **Left:** BEVNet is better at preserving small fast-moving objects such as bicyclists (highlighted by the blue circles), which the hand-engineered update rule tends to ignore. (Maps are 50% zoomed in). **Right:** When noise is injected into the odometry, the learned recurrent network is able to fix errors in BEV map caused by the noise, while Cylinder3D+TA fails to do this, resulting in a blurry, inaccurate map.

predict the future? and 3) where should information be aggregated in the network? We answer these questions through a set of experiments on both SemanticKITTI and RELLIS-3D datasets.

Is learned recurrence better than temporal aggregation? When evaluated on the full ground truth, we observe that BEVNet-R consistently outperforms BEVNet and BEVNet-TA on both on-road and off-road scenarios (Table 1). Notably, BEVNet-TA also outperforms BEVNet, which shows that any form of recurrence is beneficial. In particular, we observe that the learned recurrence makes best use of the temporal information in comparison to the hand-engineered TA. When noisy odometry is introduced we observe the same trend as discussed in Sec. 5.2, where BEVNet-R shows robustness to noise and outperforms BEVNet-TA.

Does history help predict the future? In Table 1, we can see that any form of recurrence that accumulates history helps with predicting the unseen area. When evaluated on unseen ground truth, BEVNet-TA and BEVNet-R both consistently outperform BEVNet-S, even when noisy odometry is introduced. This yields an interesting conclusion that any form of memory acquired through recurrence provides more information for forecasting the future. We emphasize that learned recurrence especially proves strong improvement in this aspect, as it consistently outperforms all other approaches on the unseen ground truth.

Where to put ConvGRU? Recurrence may be applied right after the sparse convolutions (*early aggregation*) or may be applied after the 2D inpainting network (*late aggregation*). We compare the two approaches on the SemanticKITTI dataset with clean odometry and including the unseen area for evaluation. Note that here our model is trained with clean odometry. As we emphasize with Table 2, our experiments show that early aggregation yields better results than late aggregation. This is because when early aggregation is applied the inpainting network has access to temporally fused information, and therefore is given more information to complete the scene and maintain temporal consistency across scans. Furthermore, we may infer that if late aggregation is applied, it is more difficult for the recurrent network to learn to correct the odometry as it is given completed scenes with potentially noisy information instead of the sparse feature maps.

	mIoU
Early Aggregation	**0.535**
Late Aggregation	0.479

Table 2: Effect of GRU location in the network.

5.3 Comparison and Discussion with other Related Work

Several recent papers have focused on semantic understanding of scenes from sparse LiDAR scans, which we discuss below. These papers solve related, but slightly different problems. They produce smaller maps, and do not perform temporal aggregation.

Semantic Scene Completion. The scene completion task aims to predict a dense semantic voxel grid from a single LiDAR scan. To compare to our work, we build a traversability map by converting predicted dense voxel grids to 2D traversability maps using the same method described in Sec. 3.

We compare with JS3C-Net [5] because it is one of the top performing models on the SemanticKITTI scene completion task and the code is publicly available. We map the 19-class predictions of JS3C-Net to our 4-class ontology, and project the voxels down to generate 256×256 traversability maps. We adapt our approach to produce the same output format as JS3C-Net. JS3C-Net achieves an mIoU of 0.549, whereas BEVNet-S and BEVNet-R achieve an mIoU of 0.592 and 0.608, respectively. This shows that learning to project traversability and inpainting the map simultaneously can work better than first reconstructing the scene, followed by a rule-based projection.

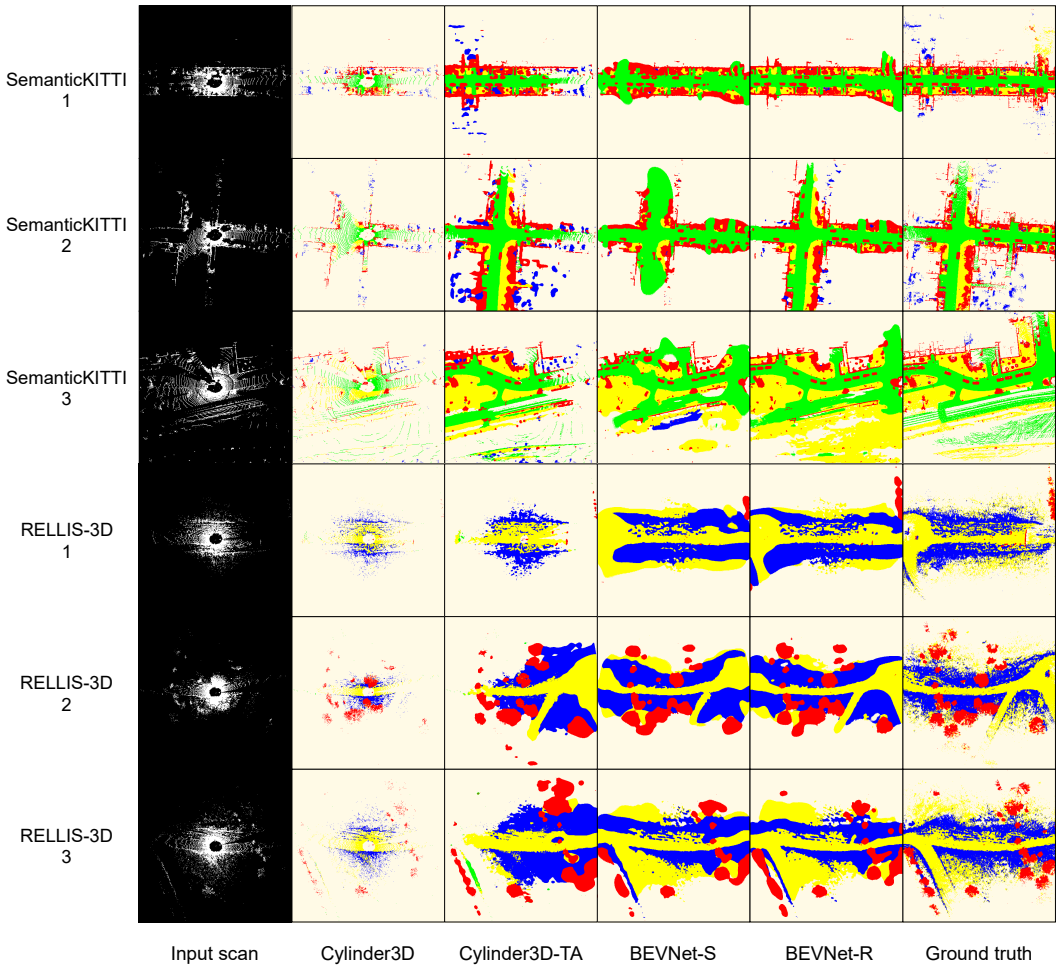

Figure 4: Qualitative comparison of our method on SemanticKITTI [12] and RELLIS-3D [13]. Learned recurrency in our end-to-end network can preserve previously observed information while predicting future observations for the occluded areas.

Inpainting Network. Han et al. [7] recently proposed an approach that uses GANs to inpaint a sparsely segmented BEV image. We trained our model using the same SemanticKITTI dataset with the groundtruth 19-class BEV images as supervision. Note that this task is different from traversability estimation because it does a simple topdown projection. Our single-frame model achieves 0.253 mIoU on this task, which is significantly higher than 0.131 reported in [7]. Note that [7] assumes the LiDAR scan has been already segmented, and it does inpainting in 2D. This makes it unsuitable for our traversability projection due to the lack of 3D reasoning.

6 Conclusion

We propose BEVNet, a framework that predicts dense terrain traversability in a local region around a mobile robot with the aim of helping the robot navigate in a novel on-road or off-road environment. BEVNet addresses a number of challenges in a unified architecture, namely: 1) it learns to aggregate sparse LiDAR information over time, 2) it learns to reason about traversability that involves both geometric and semantic understanding of the environment, and 3) it learns to fill in the unknown space where there are no LiDAR hits, and thus provides the robot with a more complete understanding of its surroundings. Most notably, BEVNet can leverage past information to *better predict the future*. We believe BEVNet provides an important step towards robot autonomy on complex terrains where a prior map is unavailable. In the future, we would like to test BEVNet in more challenging off-road scenarios and predict traversability from cheap, weakly-labeled data.

References

[1] D. Maturana, P.-W. Chou, M. Uenoyama, and S. Scherer. Real-time semantic mapping for autonomous off-road navigation. In *Proceedings of 11th International Conference on Field and Service Robotics (FSR '17)*, pages 335 – 350, September 2017.

[2] Y. Xiang and D. Fox. DA-RNN: semantic mapping with data associated recurrent neural networks. In *Proceedings of Robotics: Science and Systems*, Cambridge, Massachusetts, July 2017. doi:10.15607/RSS.2017.XIII.013.

[3] D. S. Chaplot, D. Gandhi, S. Gupta, A. Gupta, and R. Salakhutdinov. Learning to explore using active neural slam. *International Conference on Learning Representations*, 2020.

[4] S. A. Mohamed, M.-H. Haghbayan, T. Westerlund, J. Heikkonen, H. Tenhunen, and J. Plosila. A survey on odometry for autonomous navigation systems. *IEEE Access*, 7:97466–97486, 2019.

[5] X. Yan, J. Gao, J. Li, R. Zhang, Z. Li, R. Huang, and S. Cui. Sparse single sweep lidar point cloud segmentation via learning contextual shape priors from scene completion. In *AAAI*, 2021.

[6] R. Cheng, C. Agia, Y. Ren, X. Li, and B. Liu. S3cnet: A sparse semantic scene completion network for lidar point clouds. In *CoRL*, 2020.

[7] Y. Han, J. Banfi, and M. Campbell. Planning paths through unknown space by imagining what lies therein. In *CoRL*, 2020.

[8] A. J. Reina, J. L. Martínez, A. Mandow, J. Morales, and A. García-Cerezo. Collapsible cubes: Removing overhangs from 3d point clouds to build local navigable elevation maps. In *2014 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, pages 1012–1017. IEEE, 2014.

[9] J. Larson, M. Trivedi, and M. Bruch. Off-road terrain traversability analysis and hazard avoidance for ugvs. Technical report, CALIFORNIA UNIV SAN DIEGO DEPT OF ELECTRICAL ENGINEERING, 2011.

[10] K. Cho, B. van Merrienboer, C. Gulcehre, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*, 2014.

[11] P. Chao, C.-Y. Kao, Y. Ruan, C.-H. Huang, and Y.-L. Lin. Hardnet: A low memory traffic network. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3551–3560, 2019. doi:10.1109/ICCV.2019.00365.

[12] 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 *Proc. of the IEEE/CVF International Conf. on Computer Vision (ICCV)*, 2019.

[13] P. Jiang, P. Osteen, M. Wigness, and S. Saripalli. Rellis-3d dataset: Data, benchmarks and analysis, 2020.

[14] M. Bansal, A. Krizhevsky, and A. Ogale. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. In *Proceedings of Robotics: Science and Systems*, Freiburgim-Breisgau, Germany, June 2019. doi:10.15607/RSS.2019.XV.031.

[15] A. Sadat, S. Casas, M. Ren, X. Wu, P. Dhawan, and R. Urtasun. Perceive, predict, and plan: Safe motion planning through interpretable semantic representations. In *ECCV*, 2020.

[16] W. Zeng, W. Luo, S. Suo, A. Sadat, B. Yang, S. Casas, and R. Urtasun. End-to-end interpretable neural motion planner. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8652–8661, 2019.

[17] S. Casas, A. Sadat, and R. Urtasun. Mp3: A unified model to map, perceive, predict and plan, 2021.

[18] M. Bojarski, D. Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. Jackel, M. Monfort, U. Muller, J. Zhang, X. Zhang, J. Zhao, and K. Zieba. End to end learning for self-driving cars. 04 2016.

[19] J. Philion and S. Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *Proceedings of the European Conference on Computer Vision*, 2020.

[20] S. Hoermann, M. Bach, and K. Dietmayer. Dynamic occupancy grid prediction for urban autonomous driving: A deep learning approach with fully automatic labeling. *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2056–2063, 2018.

[21] P. Wu, S. Chen, and D. N. Metaxas. Motionnet: Joint perception and motion prediction for autonomous driving based on bird's eye view maps. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11382–11392, 2020.

[22] D. Stavens and S. Thrun. A self-supervised terrain roughness estimator for off-road autonomous driving. UAI'06, page 469–476. AUAI Press, 2006. ISBN 0974903922.

[23] J. Larson and M. Trivedi. Lidar based off-road negative obstacle detection and analysis. *2011 14th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, pages 192–197, 2011.

[24] B. Suger, B. Steder, and W. Burgard. Traversability analysis for mobile robots in outdoor environments: A semi-supervised learning approach based on 3d-lidar data. *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3941–3946, 2015.

[25] D. Barnes, W. P. Maddern, and I. Posner. Find your own way: Weakly-supervised segmentation of path proposals for urban autonomy. *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 203–210, 2017.

[26] X. Zhu, H. Zhou, T. Wang, F. Hong, Y. Ma, W. Li, H. Li, and D. Lin. Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. *arXiv preprint arXiv:2011.10033*, 2020.

[27] T. Cortinhal, G. Tzelepis, and E. Aksoy. Salsanext: Fast, uncertainty-aware semantic segmentation of lidar point clouds. In *ISVC*, 2020.

[28] G. Reina, A. Milella, and R. Rouveure. Traversability analysis for off-road vehicles using stereo and radar data. *2015 IEEE International Conference on Industrial Technology (ICIT)*, pages 540–546, 2015.

[29] J. Larson, M. Trivedi, and M. Bruch. Off-road terrain traversability analysis and hazard avoidance for ugvs. 2011.

[30] C. Thorpe, M. Hebert, T. Kanade, and S. Shafer. Vision and navigation for the carnegie-mellon navlab. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(3):362–373, 1988. doi:10.1109/34.3900.

[31] A. Elfes. Using occupancy grids for mobile robot perception and navigation. *Computer*, 22: 46–57, 1989.

[32] C. Choy, J. Gwak, and S. Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3070–3079, 2019.

[33] H. Fan and Y. Yang. Pointrnn: Point recurrent neural network for moving point cloud processing. *arXiv*, 1910.08287, 2019.

[34] X. Shi, Z. Chen, H. Wang, D. Yeung, W. Wong, and W. Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *NIPS*, 2015.

[35] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9:1735–1780, 1997.

[36] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.

[37] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4724–4733, 2017.

[38] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou. Information theoretic MPC for model-based reinforcement learning. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1714–1721. IEEE, 2017.

[39] S. Thrun, W. Burgard, and D. Fox. Probabilistic robotics (intelligent robotics and autonomous agents). 2005.

[40] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2015.