

Complete & Label: A Domain Adaptation Approach to Semantic Segmentation of LiDAR Point Clouds

Li Yi Boqing Gong Thomas Funkhouser
Google Research

{ericyi, bgong, tfunkhouser}@google.com

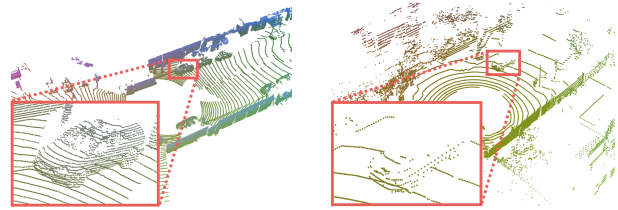
Abstract

We study an unsupervised domain adaptation problem for the semantic labeling of 3D point clouds, with a particular focus on domain discrepancies induced by different LiDAR sensors. Based on the observation that sparse 3D point clouds are sampled from 3D surfaces, we take a Complete and Label approach to recover the underlying surfaces before passing them to a segmentation network. Specifically, we design a Sparse Voxel Completion Network (SVCN) to complete the 3D surfaces of a sparse point cloud. Unlike semantic labels, to obtain training pairs for SVCN requires no manual labeling. We also introduce local adversarial learning to model the surface prior. The recovered 3D surfaces serve as a canonical domain, from which semantic labels can transfer across different LiDAR sensors. Experiments and ablation studies with our new benchmark for cross-domain semantic labeling of LiDAR data show that the proposed approach provides 6.3-37.6% better performance than previous domain adaptation methods.

1. Introduction

Semantic segmentation of LiDAR point clouds is important for many applications, including autonomous driving, semantic mapping, and construction site monitoring to name a few. Given a LiDAR sweep (frame), the goal is to produce a semantic label for each point.

Although there is a great potential for deep neural networks on this semantic segmentation task, their performance is limited by the availability of labeled training data. Acquiring manual labels for 3D points is very expensive. Several datasets have recently been released by autonomous driving companies [1, 4, 5, 13, 14, 25, 27, 51]. However, each has a different configuration of LiDAR sensors, which produce different 3D sampling patterns (Figure 1), and each covers distinct geographic regions with distinct distributions of scene contents. As a result, deep networks trained on one dataset do not perform well on others.



(a) captured by a 64-beam LiDAR (b) captured by a 32-beam LiDAR

Figure 1. The sampling discrepancy between point clouds captured by two LiDAR sensors. All figures are best viewed in color.

There is a domain adaptation problem. While the mismatch of scene contents is similar to those studied in 2D visual domain adaptation [36, 8], the sampling mismatch is unique to 3D point clouds. Each time a new LiDAR sensor configuration is selected, data is acquired with a different 3D sampling pattern, so models trained on the old data are no longer effective, and new labeled data must be acquired for supervised training in the conventional machine learning paradigm. In contrast, domain adaptation aims to take better advantage of the old labeled data by revealing unlabeled data of the new LiDAR configuration to a machine learner so that it can account for the new scenarios.

To address the sampling caused domain gap, we observe that LiDAR samples have an underlying geometric structure, and domain adaptation can be performed more effectively with a 3D model leveraging that structure. Specifically, assuming the physical world is composed of 3D surfaces, and that LiDAR sensor samples come from those surfaces, we address the domain adaption problem by transforming it into a 3D surface completion task. That is, if we can recover the underlying complete 3D surfaces from sparse LiDAR point samples, and train networks that operate on the completed surfaces, then we can leverage the labeled data from any LiDAR scanner to work on the data from any other.

The motivation for this approach is that surface completion is an easier task than semantic segmentation. First, there are strong priors on the shapes of 3D surfaces encountered in the real world, and thus a network trained to densify a point cloud can learn and leverage those priors with relatively little

training data. Second, surface completion can be learned from self-supervision (e.g., from multi-view observations) and/or from synthetic datasets (e.g., from sampled computer graphics models). Unlike semantic segmentation, no manual labels are required. We train our completion network with supervision from complete surfaces reconstructed from multiple frames of LiDAR data.

Our network architecture is composed of two phases: surface completion and semantic labeling. In the first phase, we use a sparse voxel completion network (SVCN) to recover the 3D surface from a LiDAR point cloud. In the second phase, we use a sparse convolutional U-Net to predict a semantic label for each voxel on the completed surface.

Extensive experiments with different autonomous vehicle driving datasets verify the effectiveness of our domain adaptation approach to the semantic segmentation of 3D point clouds. For example, using a network trained on the Waymo open dataset [51] to perform semantic segmentation on the nuScenes dataset [4] provides an absolute mIoU improvement of 6.0 over state-of-the-art domain adaptation methods. Similarly, training on nuScenes and testing on Waymo provides an absolute mIoU improvement of 10.4 over prior arts.

Our contributions are three-fold. First and foremost, we identify the cross-sensor domain gap for LiDAR point clouds caused by sampling differences, and we propose to recover complete 3D surfaces from the point clouds to eliminate the discrepancies in sampling patterns. Second, we present a novel sparse voxel completion network, which efficiently processes sparse and incomplete LiDAR point clouds and completes the underlying 3D surfaces with high resolution. Third, we provide thorough quantitative evaluations to validate our design choices on three datasets.

2. Related Work

Unsupervised domain adaptation. Conventional machine learning relies on the assumption that training and test sets share the same underlying distribution, but the practice often violates the assumption. Unsupervised domain adaptation (UDA) [8, 36] handles the mismatch by revealing some test examples to the machine learner such that it can account for the test-time scenarios while learning from the training set. Early work on UDA mainly reweights [45, 49, 67] or re-samples [15, 16] the source-domain examples to match the target distribution. Besides, there is a fruitful line of works on learning domain-invariant representations, such as subspace alignment [11] and interpolation [17, 18], adversarial training [12, 55, 2, 46, 23], maximum mean discrepancy [33], maximum classifier discrepancy [43], correlation alignment [50], etc. As noted in [40], these methods by design align two domains in a holistic view and fail to capture the idiosyncratic geometric properties in point clouds.

Domain adaptation for 3D point clouds. Relatively little

work has been done to study domain adaptation for 3D point clouds. Rist et al. [42] propose that dense 3D voxels are preferable to point clouds for sensor-invariant processing of LiDAR point clouds. Salah et al. [44] propose a CycleGAN approach to the adaptation of 2D bird’s eye view images of LiDAR between synthetic and real domains. Wu et al. [59] compensate for differences in missing points and intensities between real and synthetic data using geodesic correlation alignment. Qin et al. [40] and Wang et al. [58] propose multi-scale feature matching methods for object detection from 3D point clouds. None of these methods explicitly account for differences in point sampling patterns in the 3D domain.

Deep 3D semantic segmentation. We target at deep 3D semantic segmentation in this paper, which associates semantic labels to 3D data via deep learning approaches. Different from 2D images, 3D data can be represented in various forms, introducing extra challenges for deep learning methods design. Early works use dense voxel grid to represent 3D objects and leverage dense 3D convolution to predict semantic labels [60, 38], with usually a limited resolution due to the heavy computation cost. To reduce the computation load, point cloud based methods are proposed which directly operate on point sets [37, 39, 31, 48, 53]. To further leverage the relationship among 3D points, deep neural networks working on graphs [62, 57] and meshes [3, 24, 22] are used. Recently, sparse convolution based methods [20, 19, 7] have been very popular, achieving superior performance on various indoor and outdoor semantic segmentation benchmarks. They treat 3D data as a set of sparse voxels and restrict 3D convolution to these voxels. Our segmentation backbone is based upon SparseConvNet [19] but we focus on improving its domain transfer ability to 3D data with different sampling patterns.

Deep 3D shape completion. Deep 3D shape completion aims at completing missing geometry pieces of some partial 3D observation using deep learning methods. Dense voxel representation has been explored to complete single 3D objects [9, 61, 21] and indoor scenes [47]. The heavy computation cost is a big issue for these methods, making them not scale well to large LiDAR point clouds. To improve the computation efficiency, octree-based methods have been proposed [41, 52, 66] which are able to produce high resolution 3D outputs. We present a sparse voxel completion network sharing similar flavors to [41, 52, 66] with an improved network architecture and loss function. We demonstrate how we could complete sparse LiDAR point clouds with high resolution using sparse convolution when the output structure is unknown and also one main difference is that we consider the application of shape completion to 3D domain adaptation. Another relevant track of works study point cloud upsampling using deep learning methods [64, 63, 30]. They usually require an upsampling factor and have no control on the sampling patterns of the results.

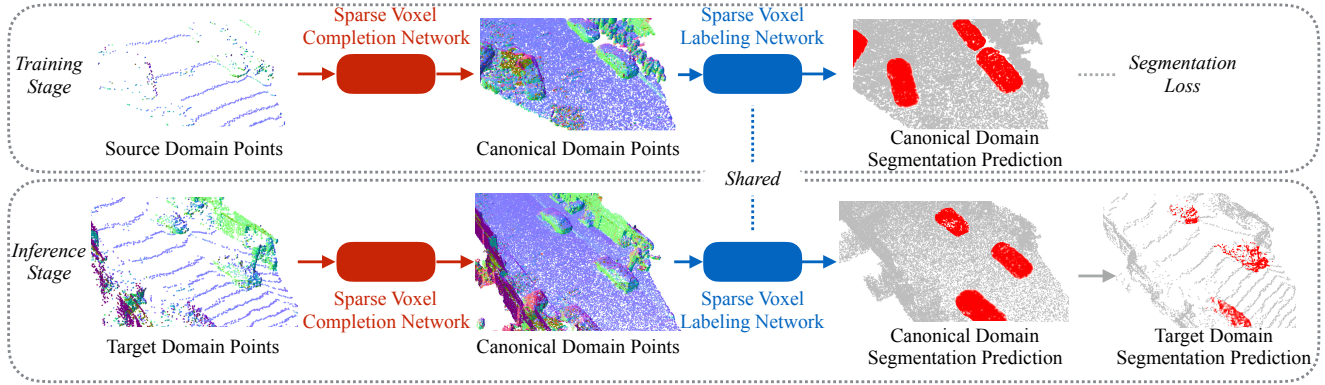


Figure 2. The overall pipeline of our “complete and label” approach.

3. Method

This paper proposes a method to overcome the domain gap caused by different LiDAR sensors’ 3D point sampling. Observing that all the sensors acquire samples of 3D surfaces, we propose a two-stage approach, where a sensor-specific surface completion neural network first recovers the underlying 3D surfaces from the sparse LiDAR point samples, and then a sensor-agnostic semantic segmentation network assigns labels to the recovered 3D surfaces. This two-phase approach focuses the domain adaptation problem on surface completion, which can be learned with self-supervision.

3.1. Overview and notations

Figure 2 illustrates the overall workflow of our approach. We consider an unsupervised domain adaptation (UDA) setting, but our approach is readily applicable to other settings such as multi-domain adaptation [10] and open domain adaptation [35, 32]. In UDA, we have access to a set of labeled LiDAR point clouds, $\{\mathbf{x}_i^s, \mathbf{y}_i^s\}_{i=1}^{N_s}$, from the source domain and a set of unlabeled LiDAR point clouds $\{\mathbf{x}_j^t\}_{j=1}^{N_t}$ in a target domain, where $\mathbf{x}_i^s \in \mathbb{R}^{T_i^s \times 3}$ and $\mathbf{x}_j^t \in \mathbb{R}^{T_j^t \times 3}$ represent two sets of T_i^s and T_j^t 3D points, respectively, and $\mathbf{y}_i^s \in \mathcal{Y} = \{1, \dots, Y\}^{T_i^s}$ corresponds to a per-point semantic label ranging within Y different classes. The two sets of point clouds are captured with different LiDAR sensors, which have their unique sampling patterns. Our goal is to learn a segmentation model that achieves high performance on the target-domain LiDAR points.

To cope with the domain gap caused by different LiDAR sensors, we learn neural surface completion networks to recover the 3D surfaces underlying incomplete 3D point clouds. Denote by $\psi^s(\mathbf{x}_i^s) \in \mathbb{R}^{M_i^s \times 3}$ and $\psi^t(\mathbf{x}_j^t) \in \mathbb{R}^{M_j^t \times 3}$ the surface completion networks for the two domains, respectively, where M_i^s and M_j^t are the numbers of dense points used to represent the completed surfaces. We say the 3D surfaces reside in a *canonical domain*.

We train a semantic segmentation network, $\phi(\psi^s(\mathbf{x}_i^s))$, over this canonical domain by using the labeled training

set of the source domain, and then apply it to the densified point clouds of the target domain, i.e., $\phi(\psi^t(\mathbf{x}_j^t))$. The per-point labels of the original target-domain point cloud \mathbf{x}_j^t are obtained by projecting the segmentation results back to the target domain.

3.2. SVCN for Surface Completion

This section describes the sparse voxel completion network (SVCN), which recovers the underlying 3D surfaces from a sparse, incomplete LiDAR point cloud and is the core of our approach.

3.2.1 Architecture

Figure 3 shows the architecture of SVCN, which comprises a structure generation sub-net and a structure refinement sub-net. The former consumes a set of sparse voxels obtained by voxelizing an input point cloud, and it outputs denser voxels representing the underlying 3D surfaces from which the input points are sampled. The structure refinement network then prunes out redundant voxels.

Both sub-nets are highly relevant to the sparse convolutional U-Net [19], which is an encoder-decoder architecture involving a series of sparse conv/deconv operations. Multi-scale features can be integrated, and skip connections provide additional information pathways in the network. However, the sparse convolutional U-Net is not directly applicable to our setting since it applies all convolutional operations only to active sites without changing the voxel structure, while we need extrapolation.

Structure generation network. In order to generate new structures for completion purposes, we replace sparse deconvolutions with dense upsampling and voxel pruning operations. Specifically, in the decoder, each voxel in the lower resolution level l will generate 2^3 voxels in the higher resolution level $l - 1$ after a dense upsampling operation (see the layers in purple in Figure 3). Low-resolution voxel features are also duplicated to the corresponding positions of high-resolution voxels.

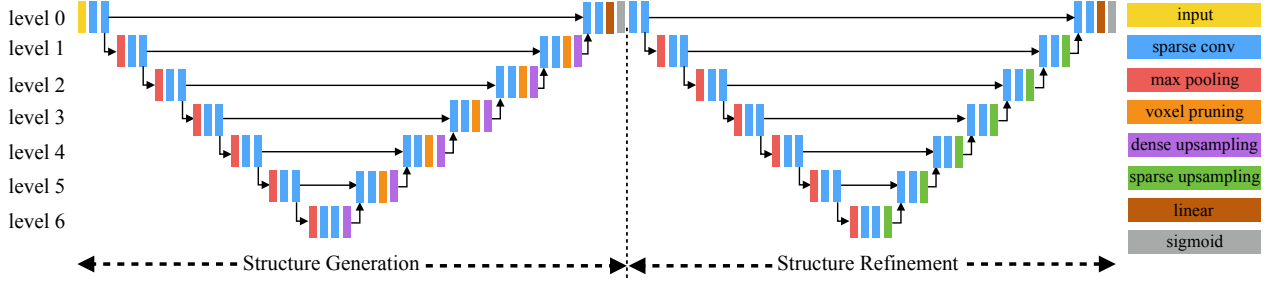


Figure 3. The architecture of the sparse voxel completion network. We use a sparse convolution with a kernel size 3 and a stride 1, max pooling with a kernel size 2 and a stride 2. Both dense and sparse upsampling are done with a factor of 2. We supervise the structure generation module and the structure refinement module with separate losses after each module as described in Section 3.2.3

The above procedure allows generating new structures, but it could easily break the inherent sparsity of voxelized 3D surfaces. Similar to [41, 52, 66], we introduce a voxel pruning module to trim voxels and avoid expanding too many in the decoder. Given a set of voxels equipped with features, the voxel pruning module applies a linear layer together with a sigmoid function on each voxel, and outputs a probability score indicating the existence of each voxel. At training time, only the ground truth voxels are kept. At test time, we prune the voxels whose existence probabilities are lower than 0.5.

To maintain the faithfulness of the generated shape to input voxels, for each resolution level l , we single out intersections between the densely upsampled voxels in the decoder and the sparse voxels in the corresponding encoder level, excluding these voxels from pruning. The skip connections need special care. Through them, we pass encoder features to the upsampled voxels, and zeros for newly generated voxels because they have no counterparts in the encoder.

Structure refinement network. The structure generation network is able to generate new structures for shape completion purposes. However, since we prune voxels using the ground truth existence probabilities on each level during training, the network could be sensitive to noisy outlier points (e.g., an outlier input voxel could possibly add a big chunk of wrong voxels to the final prediction). To cope with this issue, we introduce a structure refinement network, which is essentially a sparse convolutional U-Net adding no new voxels any more. Instead, it predicts an existence confidence score for each voxel. This is achieved by replacing the dense upsampling and voxel pruning modules in the structure generation network with sparse upsampling operations, which unpool voxel features only to the voxels that exist in the higher encoder level (see the layers in green in Figure 3). This way, the network is able to reevaluate the structure generation outputs across the whole scene in a spirit similar to stacked hourglass networks [34].

For more details of the SVCN network architecture, please refer to the supplementary materials.




Figure 4. An example of the generated training data for SVCN.

3.2.2 Training Data

We need to prepare training data, $\{(z_i^s, z_i^c)\}$ and $\{(z_i^t, z_i^c)\}$, for the surface completion networks ψ^s and ψ^t of the source and target domains, respectively, where z_i^c stands for a dense surface point cloud in the canonical domain from which we can sample both a source-domain point cloud z_i^s and a target-domain point clouds z_i^t . It is important to note that the training data for the surface completion network SVCN could be different from that for semantic segmentation, so we use different notations here. Indeed, one advantage of surface completion is that it can be learned from self-supervision which does not require manual labels. Exemplar supervisions include dense surface points via simulation, multi-view registration, and high-resolution LiDAR point clouds, to name a few. We first describe how we obtain the dense surface point clouds $\{z_i^c\}$, followed by how to sample domain-specific incomplete point clouds $\{z_i^s\}$ or $\{z_i^t\}$ for constructing the training pairs for SVCN.

Dense surface point clouds. To obtain the dense point clouds of complete 3D surfaces, we leverage the LiDAR sequences in existing autonomous driving datasets, for example [51]. Specifically, we aggregate multiple LiDAR frames within a sequence to generate a denser and more complete point cloud. Poisson surface reconstruction [26] is then applied to recover the underlying mesh surfaces with a SurfaceTrimmer step removing parts with low sampling density. We discretize a surface by uniformly sampling points on it, ensuring the point resolution is higher than the resolutions in the source or target domain. An example is shown in Figure 4(a). The complete scene point clouds act as a canonical domain with uniform sampling patterns.

Domain-specific incomplete point clouds. Given the dense, complete surface point clouds $\{z_i^c\}$, we simulate a

“virtual LiDAR” to generate incomplete point clouds for the source (target) domain such that the virtual LiDAR point clouds $\{\mathbf{z}_i^s\}$ ($\{\mathbf{z}_i^t\}$) share the same distribution as the real point clouds in that domain. In particular, we propose a polar sampling scheme to implement the “virtual LiDAR”. First, we randomly pick up a reference point cloud from a domain and compute the polar coordinate (r, θ, ϕ) for each point (x, y, z) , where $r = \sqrt{x^2 + y^2 + z^2}$, $\theta = \text{atan2}(\sqrt{x^2 + y^2}, z)$, $\phi = \text{atan2}(y, x)$. We argue that (θ, ϕ) reveals the sampling pattern in this point cloud without much scene-specific information and can be used to re-sample a different complete scene point cloud to simulate the corresponding sampling pattern. Second, we select a sensor location in the complete scene point cloud, remove occluded points, and convert the rest points into their polar coordinates. Finally, we search for the nearest neighbor point in the complete scene point cloud for each point in a reference frame and sample these points to imitate the reference sampling pattern. Notice this is done in the (θ, ϕ) space to transfer the sampling pattern only. In Figure 4 (b) and (c), we shown the simulated incomplete point clouds with sampling patterns transferred from reference point clouds in the Waymo [51] and nuScenes [4] datasets, respectively.

3.2.3 Training Algorithm

Given the paired training data, we convert them to voxels and employ a voxel-wise binary cross-entropy loss to first pre-train the structure generation sub-net. We then fix it and switch it to the inference mode, using the predicted voxel existence probability as the input to train the structure refinement sub-net with another voxel-wise binary cross-entropy loss.

Local Adversarial Learning. Since we have a strong prior that voxels densified by SVCN should lie on 3D surfaces, we propose an adversarial loss to capture this prior, in a similar spirit to [28, 56, 30]. This loss can be added to the training of either the structure generation sub-net or the refinement sub-net. A notable property of our adversarial loss is that it is imposed over local surface patches, as opposed to the global output of SVCN. Please refer to the supplementary materials for more details.

3.3. Segmentation in the Canonical Domain

We train a semantic segmentation network $\phi(\cdot)$ over the canonical domain using the labeled data $\{(\mathbf{x}_i^s, \mathbf{y}_i^s)\}$ in the source domain. We leverage a state-of-the-art 3D semantic segmentation method, MinkowskiNet [7], as our segmentation network. Given a test point cloud \mathbf{x}_i^t in the target domain, we first map it to the canonical domain by the surface completion network SVCN $\psi^t(\mathbf{x}_i^t)$, apply the segmentation network over it $\phi(\psi^t(\mathbf{x}_i^t))$, and finally project the segmentation results back to the original target-domain point cloud

\mathbf{x}_i^t . Please refer to the supplementary material for how to propagate the source-domain labels to the dense, complete point clouds in the canonical domain and how to project segmentation results to the target domain. Both depend on nearest neighbor search and majority voting.

4. Experiments

We experiment with three autonomous driving datasets captured by different LiDAR configurations.

- Waymo open dataset [51]: It contains LiDAR point cloud sequences from 1K scenes, each sequence containing about 200 frames. There are five LiDAR sensors. We use the top 64-beam LiDAR in our experiments. The LiDAR frames are labeled with 3D object bounding boxes in four categories, from which we crop the LiDAR point clouds to obtain per-point semantic labels. The data is officially split into 798 training scenes and 202 validation scenes. Following this split, we have $\sim 160\text{K}$ training frames and $\sim 40\text{K}$ validation frames.
- nuScenes-lidarseg dataset [4]: It contains $\sim 40\text{K}$ LiDAR frames annotated with per-point semantic labels from 1K scenes. Officially these points are cast into 16 categories for the semantic segmentation task with one additional “ignored” class excluded from evaluations. Different from the Waymo open dataset, it adopts a 32-beam LiDAR sensor with different configurations, causing a sampling gap from the Waymo point clouds. Following the dataset’s recommendation, we train our models using $\sim 28\text{K}$ frames from 700 training scenes and evaluate on $\sim 6\text{K}$ frames from 150 validation scenes.
- SemanticKITTI dataset [1, 13]: It adopts a Velodyne 64-beam LiDAR similar to Waymo but with a different sensor configuration. Points are classified into 19 categories with one additional “ignored” class excluded from evaluations. Following the official recommendation, we use sequence 00-07 and 09-10 for training and evaluate on sequence 08, resulting in $\sim 19\text{K}$ training frames and $\sim 4\text{K}$ frames for evaluation.

While transferring semantic segmentation from Waymo to nuScenes-lidarseg and SemanticKITTI or the inverse, we consider the only two overlapping categories in all these three datasets: vehicles and pedestrians. The two classes are both common and safety-critical in self-driving scenes. For the domain transfer between nuScenes-lidarseg and SemanticKITTI, we consider all the 10 overlapping categories between the two datasets: car, bicycle, motorcycle, truck, other vehicle, pedestrian, drivable surface, sidewalk, terrain, and vegetation. We carefully remap the semantic categories to guarantee the class definitions in different datasets are consistent. Please refer to the supplementary material for the remapping process. The three datasets provide an organic, large-scale testbed to study domain adaptation methods for

3D point clouds. By design, our approach copes with the domain discrepancy among the three datasets caused by different configurations of LiDAR sensors.

4.1. Sparse LiDAR Point Cloud Completion

We first evaluate our sparse voxel completion network (SVCN) in this section.

Training data. SVCN takes as input an incomplete point cloud and predicts its underlying complete 3D surfaces in a dense volumetric form. To generate data pairs for training and evaluation, we aggregate multiple frames within each sequence from the Waymo open dataset, resulting in 2400 complete scene point clouds for training and 200 for test. We then sample incomplete point clouds via the “virtual LiDAR” described in Section 3.2. We voxelize the complete point clouds using a voxel size of 20cm to provide ground truth supervision for various learning methods.

Evaluation. We use two evaluation metrics. One is voxel-level intersection over union (IoU). The other is the Chamfer Distance (CD) between the predicted voxel set and the ground truth voxel set. To compute the CD between two voxel sets, we convert each voxel set into a point cloud by keeping the center of each voxel and then compute CD between the two point clouds.

SVCN vs. Baselines. There is not much prior work on using sparse convolution to complete sparse LiDAR inputs. The closest baseline is ESSCNet [66] which achieves state-of-the-art results for semantic indoor scene completion. Our structure generation network (the first half of SVCN) is a variant of ESSCNet with two improvements. It leverages only one group in the spatial group convolution for higher generation quality, and we densely supervise voxel pruning at each resolution with resolution balancing weights. So we will also refer to the ablation baseline of our SVCN without the structure refinement as ESSCNet⁺⁺. A second ablation baseline is our full SVCN network trained without the local adversarial loss.

Table 1. Comparison for sparse LiDAR point cloud completion (W: Waymo, N: nuScenes-lidarseg, K: SemanticKITTI).

Test domain	Metric	ESSCNet ⁺⁺ [66]	SVCN w/o adv.	SVCN
W	IoU(%) ↑	44.1	46.3	47.5
	CD(m) ↓	1.070	1.013	0.968
N	IoU(%) ↑	24.9	26.7	28.8
	CD(m) ↓	1.745	1.730	1.610
K	IoU(%) ↑	40.9	42.9	44.3
	CD(m) ↓	1.147	1.122	1.052

Results. Table 1 shows the comparison results. Our full network with local adversarial learning outperforms all competing methods for the inputs with the sampling patterns of all three datasets. Comparing ESSCNet⁺⁺ and SVCN without the local adversarial loss, we can see that the structure refine-

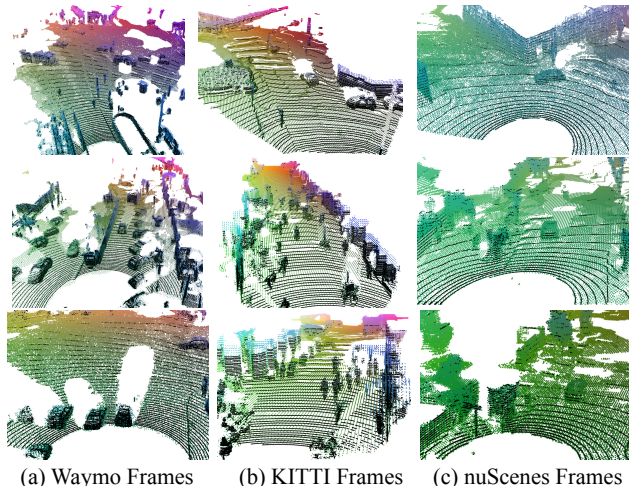


Figure 5. Visualizations of the surface completion results from different datasets. Black points indicate the original sparse incomplete LiDAR points, and we use colored points to represent the output of our sparse voxel completion network.

ment network does improve the scene completion quality. Finally, the local adversarial loss, which accounts for surface priors, results in better completions than SVCN without it. Notice that the LiDAR point clouds from nuScenes-lidarseg hold a much sparser sampling pattern compared with those from Waymo or SemanticKITTI, and are thus more challenging for the completion task. This is revealed by the relatively low IoU and high CD scores when the inputs hold nuScenes-lidarseg sampling patterns.

To better understand how our surface completion network SVCN could canonicalize different sampling patterns and therefore mitigate the corresponding domain gap, we visualize the surface completion results from different datasets in Figure 5. We use black points to represent the incomplete LiDAR inputs and colored points for the outputs of our SVCN. It is clear that SVCN is able to recover the underlying surfaces regardless of the input sampling patterns. SVCN is also able to fill small holes to make the geometry more complete. Comparing the vehicles from Waymo and nuScenes-lidarseg datasets respectively, we observe a clear domain gap in the inputs. In contrast, after surface completion they share more similar sampling patterns and geometry.

4.2. Unsupervised Domain Adaptation Results

In this section, we study the domain transfer ability of our approach among the Waymo, nuScenes-lidarseg and SemanticKITTI datasets. We compare our method with the previous state-of-the-art method of SqueezeSegV2 (SQSGV2) on this topic [59], which projects LiDAR point clouds to form 2D range images and uses 2D convolutional neural networks as the segmentation backbone. Since domain adaptation methods for 3D LiDAR point cloud segmentation has not been studied much previously, we also compare with




Figure 6. A comparison of different domain adaptation methods on an example Waymo frame. We consider the domain adaptation direction from nuScenes-lidarseg to Waymo dataset. Different colors indicate different semantic classes. FeaDA and OutDA represent feature space domain and output space domain adaptation respectively. We use green circles to highlight the prediction errors.

state-of-the-art adaptation methods for 2D semantic segmentation, including feature space adversarial domain adaptation (FeaDA) [6], output space adversarial domain adaptation (OutDA) [54], and Sliced Wasserstein Discrepancy-based domain adaptation (SWD) [29]. In addition, we incorporate the geodesic correlation alignment technique used in [59] with our 3D segmentation backbone, forming another baseline named 3DGCA. We report the results in Table 2.

From Table 2, we can see an obvious performance drop when transferring segmentation networks from one domain to another. For example, compared with both training and testing on nuScenes-lidarseg ($N \rightarrow N$), training on the Waymo dataset while evaluating on the nuScenes-lidarseg dataset ($W \rightarrow N$) would cause the mean IoU (mIoU) to drop from 69.9% to 42.9%. This shows the importance of studying the domain adaptation problem. Our method successfully brings the mIoU to 50.2% and outperforms the prior arts. We draw the same observation on the other pairs of domains tested.

The 2D domain adaptation methods do not work well on the 3D point clouds. FeaDA tries to bring close two domains in a global feature space, but it fails to model rich local cues in 3D point clouds, such as sampling patterns, surfaces, and contexts. OutDA fails in most cases with no surprise because it assumes that the segmentation masks of two domains are indistinguishable. While this assumption works for 2D scenes, it breaks given different 3D sampling patterns in two domains. SWD in general does a better job than the previous two methods. However, since it does not explicitly model the cause of the domain gap, we do not observe a huge performance improvement either. The performance of SQSGV2 is significantly lower than the other methods when evaluated on 3D point clouds. It requires projecting LiDAR point clouds onto range images to exploit 2D convolutional methods. When projecting the predicted labels back to 3D, we observe huge errors especially along object boundaries, leading to the subpar performance. The performance improvement by 3DGCA is also restricted.

Qualitative Results. Figure 6 shows some qualitative results of both surface completion and semantic segmentation when adapting from nuScenes-lidarseg to Waymo. We can see that the baseline methods mislabel objects that are both close to and distant from the sensor location. The sparsity of distant objects is a great challenge for all methods. Our

approach completes the underlying 3D surfaces from only sparse observations, making it easier for the segmentation network to handle those distant objects.

Correlation between scene completion and domain adaptation. We provide additional ablation studies about the correlation between the quality of scene completion and the performance of domain adaptation. We replace SVCN in our method with its variants described in Section 4.1 and report the resulting segmentation results in Table 3. It can be clearly seen that better scene completion qualities lead to better domain transfer performances.

Comparison with handcrafted baselines. The sampling patterns differ across the autonomous driving datasets due to complex interactions of various factors, such as the heights, poses, and types of LiDAR sensors, making it hard to design analytical solutions to align them. To show the difficulty of the problem and demonstrate the necessity of our approach, we also study some analytical methods aiming to align different sampling patterns. Consider the Waymo open dataset captured by 64-beam LiDARs and nuScenes consisting of 32-beam LiDAR data. We design two handcrafted baselines to align them: B1) analytically downsampling or upsampling the LiDAR beams and B2) piecewise linear interpolation. For both baselines, we first convert a LiDAR point cloud into a $H \times W$ range image, where H equals the beam number of the LiDAR and W is set to 2048. We say two points are adjacent if they have the same column index and their row indices are off by one. In B1, we simply remove the points of every other row in the range image to downsample Waymo point clouds. To upsample nuScenes, we add the midpoint of every pair of adjacent points. In B2, we linearly interpolate every pair of adjacent points, densely adding points to the line segment and making their spacing follow a hyperparameter δ . We use the same δ while densifying Waymo and nuScenes to bring their sampling patterns close.

We compare our approach with the above two baselines in Table 4. Our method outperforms both by large margins. B1 fails because the sampling pattern difference is on more than just the number of LiDAR beams. B2 densifies the point cloud by interpolation, but it produces “phantom” points along back-projections of silhouette boundaries and does not complete occluded regions. This study shows the importance of our learning-based approach and the SVCN network.

Table 2. Unsupervised domain adaptation for 3D semantic segmentation among the Waymo, nuScenes-lidarseg and SemanticKITTI datasets. N denotes nuScenes-lidarseg, W denotes Waymo and K denotes SemanticKITTI. We report the mIoU in each cell.

Source→Target	#classes	No DA	FeaDA [6]	OutDA [54]	SWD [29]	SQSGV2 [59]	3DGCA	Ours
N→N	2	69.9	-	-	-	-	-	-
W→N	2	42.9	44.2	40.2	42.4	14.2	43.85	50.2
N→N	10	54.4	-	-	-	-	-	-
K→N	10	27.9	27.2	26.5	27.7	10.1	27.4	31.6
W→W	2	86.3	-	-	-	-	-	-
N→W	2	46.2	48.7	43.6	49.3	30.4	48.7	59.7
K→W	2	46.3	43.9	43.4	47.2	34.2	46.1	52.0
K→K	10	50.2	-	-	-	-	-	-
N→K	10	23.5	21.4	22.7	24.5	13.4	23.9	33.7
K→K	2	61.0	-	-	-	-	-	-
W→K	2	55.0	56.4	54.1	56.8	36.8	56.2	60.4

Table 3. The segmentation mIoU of our approach when using different scene completion methods.

Source→Target	Ours w/o refinement	Ours w/o adv.	Ours-full
W→N	46.6	49.1	50.2
N→W	58.2	59.0	59.8

Table 4. Comparison with handcrafted sampling aligning baselines.

Src→Tgt	No DA	B1	B2	Ours
W→N	42.9	46.0	45.1	50.2
N→W	46.2	55.6	55.2	59.8

4.3. Domain Generalization Results

Getting rid of the dependency on target domain data during training was argued to be an important feature in real applications [65], which allows domain generalization to multiple unseen target domains. In this section, we demonstrate our approach performs well for domain generalization.

We choose the Waymo dataset as our source domain and aim at generalizing a segmentation neural network to nuScenes-lidarseg and SemanticKITTI without accessing them during training. Our strategy is to train a generic surface completion network. For this purpose, we introduce data augmentation while generating virtual LiDAR point clouds $\{\mathbf{z}_i^s\}$ from complete surface point clouds $\{\mathbf{z}_i^c\}$. More concretely, we use (θ, ϕ) to denote the polar coordinates of a Waymo point cloud and we evenly quantize θ into 64 bins. We randomly select 30% to 70% of the bins and use the corresponding points as our augmented reference sampling pattern. This augmentation strategy enforces SVCN to handle various sampling patterns, therefore generalizing to different target domains. After training our SVCN using only the reference point clouds from Waymo (with the data augmentation strategy), we evaluate its surface completion quality for the nuScenes-lidarseg and SemanticKITTI sampling patterns. The evaluation metric is the same as in Section 4.1. Furthermore, we show how the surface completion quality

contributes to the domain generalization performance of the semantic segmentation task using mIoU as the evaluation metric. Table 5 shows the results, where we additionally report the domain adaptation results for comparison (i.e., by training a target-domain-specific SVCN).

Table 5. Domain generalization from Waymo to nuScenes-lidarseg and SemanticKITTI by training a generic SVCN.

Source→Target	Method	Surface Completion IoU(%) / CD(m)	Semantic Segmentation mIoU(%)
W→N	No Adaptation	-/-	42.9
	Adaptation	28.8/1.610	50.2
	Generalization	25.7/1.800	49.8
W→K	No Adaptation	-/-	55.0
	Adaptation	44.3/1.052	60.4
	Generalization	42.8/1.115	59.6

We can see that generic SVCN trained under the domain generalization setting performs a little worse on surface completion and slightly degrades the semantic segmentation results on nuScenes-lidarseg and SemanticKITTI, compared with the target-domain-specific SVCNs learned under the domain adaptation setting. However, both adaptation and generalization are better than the no adaptation baselines by large margins, indicating the efficacy of our method.

5. Conclusion

In this paper, we present “complete and label”, a novel domain adaptation approach designed to overcome the domain gap in 3D point clouds acquired with different LiDAR sensors. We show that by leveraging geometric priors, we can transform this domain adaptation problem into a 3D surface completion task, and then perform downstream tasks such as semantic segmentation on the completed 3D surfaces with sensor-agnostic networks. Extensive experiments with multiple autonomous driving datasets demonstrate the significant improvement of our approach over prior arts.

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 International Conference on Computer Vision*, pages 9297–9307, 2019.
- [2] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3722–3731, 2017.
- [3] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.
- [4] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. *arXiv preprint arXiv:1903.11027*, 2019.
- [5] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, et al. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8748–8757, 2019.
- [6] Yi-Hsin Chen, Wei-Yu Chen, Yu-Ting Chen, Bo-Cheng Tsai, Yu-Chiang Frank Wang, and Min Sun. No more discrimination: Cross city adaptation of road scene segmenters. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1992–2001, 2017.
- [7] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3075–3084, 2019.
- [8] Gabriela Csurka. *Domain adaptation in computer vision applications*, volume 2. Springer, 2017.
- [9] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5868–5877, 2017.
- [10] Lixin Duan, Ivor W Tsang, Dong Xu, and Tat-Seng Chua. Domain adaptation from multiple sources via auxiliary classifiers. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 289–296, 2009.
- [11] Basura Fernando, Amaury Habrard, Marc Sebban, and Tinne Tuytelaars. Unsupervised visual domain adaptation using subspace alignment. In *Proceedings of the IEEE international conference on computer vision*, pages 2960–2967, 2013.
- [12] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.
- [13] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.
- [14] Jakob Geyer, Yohannes Kassahun, Mentar Mahmudi, Xavier Ricou, Rupesh Durgesh, Andrew S Chung, Lorenz Hauswald, Viet Hoang Pham, Maximilian Mühlegg, Sebastian Dorn, et al. A2d2: Audi autonomous driving dataset. *arXiv preprint arXiv:2004.06320*, 2020.
- [15] Boqing Gong, Kristen Grauman, and Fei Sha. Connecting the dots with landmarks: Discriminatively learning domain-invariant features for unsupervised domain adaptation. In *International Conference on Machine Learning*, pages 222–230, 2013.
- [16] Boqing Gong, Kristen Grauman, and Fei Sha. Reshaping visual datasets for domain adaptation. In *Advances in Neural Information Processing Systems*, pages 1286–1294, 2013.
- [17] Boqing Gong, Kristen Grauman, and Fei Sha. Geodesic flow kernel and landmarks: Kernel methods for unsupervised domain adaptation. In *Domain Adaptation in Computer Vision Applications*, pages 59–79. Springer, 2017.
- [18] Raghuraman Gopalan, Ruonan Li, and Rama Chellappa. Domain adaptation for object recognition: An unsupervised approach. In *2011 international conference on computer vision*, pages 999–1006. IEEE, 2011.
- [19] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9224–9232, 2018.
- [20] Benjamin Graham and Laurens van der Maaten. Submanifold sparse convolutional networks. *arXiv preprint arXiv:1706.01307*, 2017.
- [21] Xiaoguang Han, Zhen Li, Haibin Huang, Evangelos Kalogerakis, and Yizhou Yu. High-resolution shape completion using deep neural networks for global structure and local geometry inference. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 85–93, 2017.
- [22] Rana Hanocka, Amir Hertz, Noa Fish, Raja Giryes, Shachar Fleishman, and Daniel Cohen-Or. Meshcnn: a network with an edge. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019.
- [23] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. *arXiv preprint arXiv:1711.03213*, 2017.
- [24] Jingwei Huang, Haotian Zhang, Li Yi, Thomas Funkhouser, Matthias Nießner, and Leonidas J Guibas. TextureNet: Consistent local parametrizations for learning from high-resolution signals on meshes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4440–4449, 2019.
- [25] Xinyu Huang, Xinjing Cheng, Qichuan Geng, Binbin Cao, Dingfu Zhou, Peng Wang, Yuanqing Lin, and Ruigang Yang. The apolloscape dataset for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 954–960, 2018.

- [26] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, volume 7, 2006.
- [27] R. Kesten, M. Usman, J. Houston, T. Pandya, K. Nadhamuni, A. Ferreira, M. Yuan, B. Low, A. Jain, P. Ondruska, S. Omari, S. Shah, A. Kulkarni, A. Kazakova, C. Tao, L. Platinsky, W. Jiang, and V. Shet. Lyft level 5 av dataset 2019. url: <https://level5.lyft.com/dataset/>, 2019.
- [28] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.
- [29] Chen-Yu Lee, Tanmay Batra, Mohammad Haris Baig, and Daniel Ulbricht. Sliced wasserstein discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10285–10295, 2019.
- [30] Ruihui Li, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng. Pu-gan: a point cloud upsampling adversarial network. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7203–7212, 2019.
- [31] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. In *Advances in neural information processing systems*, pages 820–830, 2018.
- [32] Ziwei Liu, Zhongqi Miao, Xingang Pan, Xiaohang Zhan, Stella X Yu, Dahua Lin, and Boqing Gong. Compound domain adaptation in an open world. *arXiv preprint arXiv:1909.03403*, 2019.
- [33] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I Jordan. Learning transferable features with deep adaptation networks. *arXiv preprint arXiv:1502.02791*, 2015.
- [34] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hour-glass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016.
- [35] Pau Panareda Busto and Juergen Gall. Open set domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 754–763, 2017.
- [36] Vishal M Patel, Raghuraman Gopalan, Ruonan Li, and Rama Chellappa. Visual domain adaptation: A survey of recent advances. *IEEE signal processing magazine*, 32(3):53–69, 2015.
- [37] 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.
- [38] Charles R Qi, Hao Su, Matthias Nießner, Angela Dai, Mengyuan Yan, and Leonidas J Guibas. Volumetric and multi-view cnns for object classification on 3d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5648–5656, 2016.
- [39] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in neural information processing systems*, pages 5099–5108, 2017.
- [40] Can Qin, Haoxuan You, Lichen Wang, C-C Jay Kuo, and Yun Fu. Pointdan: A multi-scale 3d domain adaption network for point cloud representation. In *Advances in Neural Information Processing Systems*, pages 7190–7201, 2019.
- [41] Gernot Riegler, Ali Osman Ulusoy, Horst Bischof, and Andreas Geiger. Octnetfusion: Learning depth fusion from data. In *2017 International Conference on 3D Vision (3DV)*, pages 57–66. IEEE, 2017.
- [42] Christoph B Rist, Markus Enzweiler, and Dariu M Gavrilă. Cross-sensor deep domain adaptation for lidar detection and segmentation. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 1535–1542. IEEE, 2019.
- [43] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3723–3732, 2018.
- [44] Khaled Saleh, Ahmed Abobakr, Mohammed Attia, Julie Iskander, Darius Nahavandi, Mohammed Hossny, and Saeid Nahvandi. Domain adaptation for vehicle detection from bird’s eye view lidar point cloud data. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019.
- [45] Hidetoshi Shimodaira. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of statistical planning and inference*, 90(2):227–244, 2000.
- [46] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell Webb. Learning from simulated and unsupervised images through adversarial training. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2107–2116, 2017.
- [47] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1746–1754, 2017.
- [48] Hang Su, Varun Jampani, Deqing Sun, Subhansu Maji, Evangelos Kalogerakis, Ming-Hsuan Yang, and Jan Kautz. Splatnet: Sparse lattice networks for point cloud processing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2530–2539, 2018.
- [49] Masashi Sugiyama, Taiji Suzuki, Shinichi Nakajima, Hisashi Kashima, Paul von Büna, and Motoaki Kawanabe. Direct importance estimation for covariate shift adaptation. *Annals of the Institute of Statistical Mathematics*, 60(4):699–746, 2008.
- [50] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *European conference on computer vision*, pages 443–450. Springer, 2016.
- [51] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. *arXiv*, pages arXiv–1912, 2019.

- [52] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2088–2096, 2017.
- [53] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6411–6420, 2019.
- [54] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7472–7481, 2018.
- [55] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7167–7176, 2017.
- [56] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018.
- [57] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (TOG)*, 38(5):1–12, 2019.
- [58] Ze Wang, Sihao Ding, Ying Li, Minming Zhao, Sohini Roychowdhury, Andreas Wallin, Guillermo Sapiro, and Qiang Qiu. Range adaptation for 3d object detection in lidar. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019.
- [59] Bichen Wu, Xuanyu Zhou, Sicheng Zhao, Xiangyu Yue, and Kurt Keutzer. Squeezesegv2: Improved model structure and unsupervised domain adaptation for road-object segmentation from a lidar point cloud. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 4376–4382. IEEE, 2019.
- [60] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015.
- [61] Bo Yang, Hongkai Wen, Sen Wang, Ronald Clark, Andrew Markham, and Niki Trigoni. 3d object reconstruction from a single depth view with adversarial learning. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 679–688, 2017.
- [62] Li Yi, Hao Su, Xingwen Guo, and Leonidas J Guibas. Syncspecnn: Synchronized spectral cnn for 3d shape segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2282–2290, 2017.
- [63] Lequan Yu, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng. Ec-net: an edge-aware point set consolidation network. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 386–402, 2018.
- [64] Lequan Yu, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng. Pu-net: Point cloud upsampling network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2790–2799, 2018.
- [65] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2100–2110, 2019.
- [66] Jiahui Zhang, Hao Zhao, Anbang Yao, Yurong Chen, Li Zhang, and Hongen Liao. Efficient semantic scene completion network with spatial group convolution. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 733–749, 2018.
- [67] Kun Zhang, Bernhard Schölkopf, Krikamol Muandet, and Zhikun Wang. Domain adaptation under target and conditional shift. In *International Conference on Machine Learning*, pages 819–827, 2013.