

SCP: SCENE COMPLETION PRE-TRAINING FOR 3D OBJECT DETECTION

Yiming Shan^{1†}, Yan Xia^{1,2†*}, Yuhong Chen¹, Daniel Cremers^{1,2}

¹ Technical University of Munich, Germany

² Munich Center for Machine Learning (MCML), Germany
(yiming.shan, yan.xia, yuhong.chen, cremers)@tum.de

KEY WORDS: LiDAR Point Clouds, 3D Object Detection, Pre-training, Scene Completion, Autonomous Driving

ABSTRACT:

3D object detection using LiDAR point clouds is a fundamental task in the fields of computer vision, robotics, and autonomous driving. However, existing 3D detectors heavily rely on annotated datasets, which are both time-consuming and prone to errors during the process of labeling 3D bounding boxes. In this paper, we propose a Scene Completion Pre-training (SCP) method to enhance the performance of 3D object detectors with less labeled data. SCP offers three key advantages: (1) Improved initialization of the point cloud model. By completing the scene point clouds, SCP effectively captures the spatial and semantic relationships among objects within urban environments. (2) Elimination of the need for additional datasets. SCP serves as a valuable auxiliary network that does not impose any additional efforts or data requirements on the 3D detectors. (3) Reduction of the amount of labeled data for detection. With the help of SCP, the existing state-of-the-art 3D detectors can achieve comparable performance while only relying on 20% labeled data.

1. INTRODUCTION

3D object detection using LiDAR point clouds is a key task in the domains of computer vision (Mao et al., 2021), robotics (Zhou and Tuzel, 2018), and autonomous driving (Yan et al., 2018). In contrast to 2D images, point clouds obtained through mobile laser scanning (MLS) offer accurate 3D geometric properties and depth insights (Xia et al., 2021c), endowing them with superior resilience for object detection under diverse illumination conditions (Xia et al., 2023b). Over the past few years, numerous learning-based 3D detection techniques have exhibited remarkable performance by leveraging extensive supervised training on large annotated datasets. Nevertheless, the annotation of point clouds presents a substantial challenge due to (1) inherent incompleteness and occlusion, which render the identification of points ambiguous (Li et al., 2023); and (2) the time-consuming and error-prone nature of labeling individual points or delineating 3D bounding boxes (Wang et al., 2021).

A possible solution is to learn 3D model initialization using an unsupervised pre-training way and then fine-tune the models with small labeled data. A recent line of pre-training works based on generative adversarial networks (GANs) is proposed. (Sauder and Sievers, 2019) proposes to learn the rearrangement of the point clouds by predicting the original voxel location of each point. However, this approach is incapable of handling rotated and translated point clouds effectively due to the permutation variability exhibited in their voxel representations. Furthermore, (Sharma and Kaul, 2020) explores to classify each point into the assigned partitions based on cover trees. However, it ignores the semantically contiguous regions (e.g., airplane wings, car tires). In addition, PointContrast (Xie et al., 2020) leverages established point-wise correspondences between various views of a complete 3D scene to pre-train weights for point clouds. However, this approach may not be suitable for dynamic urban environments. Recently, a novel approach has been proposed to

learn model initialization by completing the 3D shape of single objects (Wang et al., 2021). This method demonstrates notable advancements in various downstream tasks, such as object classification and segmentation, by effectively completing individual objects. However, it overlooks the significance of both spatial and semantic relationships among objects, which are crucial considerations for successful 3D detection tasks in complex urban environments.

To tackle this problem, we propose a novel Scene Completion Pre-training network, named SCP, aiming to learn a robust model initialization for 3D object detection from single LiDAR scans. Our SCP involves training a voxel-based scene completion network consisting of a feature encoder and a decoder. The encoder utilizes a Transformer-based 3D backbone (Mao et al., 2021) to efficiently extract informative features from the raw point clouds. Concurrently, the decoder incorporates an anisotropic convolution (AIC) module (Li et al., 2020), which dynamically adapts the receptive fields for different voxels. By completing the scene point cloud, our method accurately learns the spatial and semantic relationships among the objects. This enables the pre-training model to serve as an effective model initialization when employed as the 3D backbone in a 3D detection network. Furthermore, a small labeled data is introduced to fine-tune the 3D detection network.

To summarize, the main contributions of this work are:

- We propose a voxel-based Scene Completion Pre-training network, called SCP, which purely applies to LiDAR point clouds. With the help of the carefully designed encoder and decoder, SCP provides a robust model initialization for the next detection network, encoding the spatial and semantic relationships within urban environments.
- We conduct extensive experiments on the KITTI 3D detection benchmark (Geiger et al., 2013) to demonstrate the effectiveness of our SCP. Notably, the existing state-of-the-art methods with SCP yield comparable performance while relying on 20% labeled data.

* Corresponding author

† Equal contribution




Figure 1. Overview of our SCP. Step 1: An encoder-decoder scene completion model is designed to complete the raw point cloud. Step 2: Pre-training model weights from step 1 are used for the following 3D detection task.

2. RELATED WORK

In this section, we give a brief literature review of 3D object detection and scene completion, respectively.

3D object detection. Early 3D object detection works (Shi et al., 2019, Yang et al., 2020, Mao et al., 2021, Zhou and Tuzel, 2018, Yan et al., 2018, Shi et al., 2020, Xia et al., 2023b) can be broadly categorized into two main methods: the point-based and the voxel-based detectors. The point-based approaches focus on directly capturing features and predicting 3D bounding boxes from the raw points. PointRCNN (Shi et al., 2019) extracts features from the foreground points and derives the corresponding 3D bounding box. 3DSSD (Yang et al., 2020) removes the FP layer and refinement module to reduce computational complexity and proposes a new fusion sampling strategy that yields improved results using fewer representative points. VoTr (Mao et al., 2021) introduces a voxel transformer-based 3D detection backbone, presenting an alternative solution to the task of 3D object detection. The voxel-based approach in 3D object detection involves transforming the large and non-structured point cloud data into voxels, which enables efficient feature extraction and saves computational time. VoxelNet (Zhou and Tuzel, 2018) and SECOND (Yan et al., 2018), for instance, partition the points into voxels and utilize 3D sparse convolution to extract features. Subsequently, they employ the Region Proposal Network (RPN) to obtain 3D bounding boxes. On the other hand, PV-RCNN (Shi et al., 2020) combines the strengths of both approaches. It leverages multi-scale techniques to generate high-quality proposals from voxel-based methods while also incorporating fine-grained local information from point-based methods. Recently, DMT (Xia et al., 2023b) explores motion prior knowledge to generate accurate 3D positions and rotation.

3D scene completion. Early works on 3D completion mainly focus on single objects (Xia et al., 2021a, Wang et al., 2022, Xia et al., 2021c). Comparably, completing the whole scene poses greater challenges since the scene point cloud is large-scale and has many objects with various densities. The pioneering work by Song (Song et al., 2017) explores the depth maps for 3D scene completion and leverages the scene information derived from the depth map for semantic segmentation. Scene completion and semantic segmentation are closely intertwined tasks, and jointly processing them can yield mutual performance improvements. JS3C (Yan et al., 2021) introduces the Point-Voxel Interaction (PVI) module to enhance knowledge fusion between the semantic segmentation and semantic scene completion tasks. This module facilitates interaction between incomplete local geometries in point clouds and complete global structures in voxels, enabling a more comprehensive understanding of the scene. AICNet (Li et al., 2020) proposes a

novel anisotropic convolution, which decomposes a 3D convolution into three consecutive 1D convolutions. S3CNet (Cheng et al., 2021) tackles the challenge of large-scale environments by incorporating sparsity considerations and leveraging a sparse convolution-based neural network. Recently, SCPNet (Xia et al., 2023c) introduces a novel knowledge distillation objective termed as Dense-to-Sparse Knowledge Distillation (DSKD).

3. METHOD

The overview of our SCP for 3D object detection can be divided into two stages, as illustrated in Fig. 1. In the first step (scene completion pre-training), we employ an encode-decode model to effectively complete the partial scene point clouds. This involves leveraging available data to predict and generate the missing parts of the point cloud, resulting in a more comprehensive representation of the scene. In the second step (3D object detection), we utilize the learned weights from the scene completion pre-training model as an initialization for the 3D detectors. By transferring the knowledge acquired during the scene completion, we establish a strong spatial and semantic relationship of objects, leading to improved detection performance and efficiency, especially in the case of smaller labeled data.

Scene completion pre-training Next, we provide a detailed pipeline of SCP in Fig. 2. This section is divided into four stages: voxelization, encoder, decoder, and prediction.




Figure 2. The pipeline of scene completion network. The encoder is a Transformer-based backbone and the decoder includes an AIC module and three convolutional layers.

a) Voxelization. The raw point cloud is initially transformed into structured voxels, which serve to facilitate the subsequent feature extraction process.

b) Encoder. We use a transformer-based 3D backbone network to extract features from voxels. The architecture of the 3D backbone is the same as VoTr (Mao et al., 2021), as illustrated in Fig. 3. The voxel undergoes a sequence of three "VoTr Block" layers. Each block layer consists of one sparse voxel module and two submanifold voxel modules. As the voxel passes through each block layer, its features are effectively extracted,

and the voxel is downsampled three times. In these voxels, both non-empty voxels and empty voxels are present. The submanifold voxel modules are designed to handle the non-empty voxels and utilize self-attention mechanisms (Xia et al., 2021b, Xia et al., 2023a) to effectively extract features from them. On the other hand, the sparse voxel modules are specifically designed for empty voxels, allowing them to perform feature extraction on these regions.

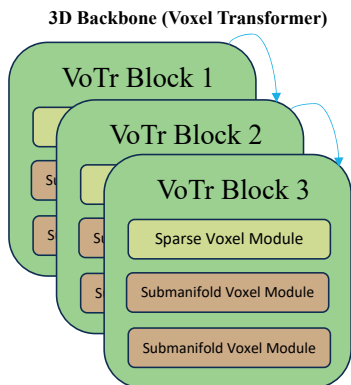


Figure 3. The architecture of the 3D backbone network. It consists of three VoTr blocks, each layer containing two different self-attention modules.

c) Decoder. Subsequently, the features obtained from the encoder are fed into the decoder, which includes Anisotropic Convolutional Networks (Li et al., 2020) and three layers of 3D convolution. Instead of employing traditional 3D convolution directly, the AIC model decomposes the 3D convolution into three separate operations, each corresponding to one dimension (x, y, z). Within each dimension, three 1D convolutions are inserted based on modulation factors, represented by a 3×3 matrix. The modulation factors matrix contains information about nine different weights in X, Y, and Z three different dimensions ($K_{x1}, K_{x2}, \dots, K_{z3}$). These modulation factors are all positive, and the sum of each row (dimension) equals one. By utilizing distinct convolution kernels and sizes for each 1D convolution, the AIC model enables improved incorporation of geometric information and provides enhanced flexibility. Ultimately enhancing the performance of the model.

d) Prediction. The completion predictions obtained from this pipeline represent a comprehensive and coherent depiction of the entire scene, effectively filling in the gaps and providing a more holistic representation of the raw data.

3D object detection. In the first step, we acquired well-performing pre-training weights, which are then utilized to aid in 3D detector (Mao et al., 2021). The pipeline of the detector VoTr is shown in Fig. 4. This knowledge transfer enables the utilization of valuable insights gained from scene completion to improve the accuracy and reliability of 3D detection results. In the initial step, the raw point cloud data undergoes voxelization, converting them into the structured voxel representation. These voxelized data are then fed into the 3D backbone, which is the same in the scene completion. The extracted features are then projected onto a bird's-eye view (BEV) map to generate 3D proposals, followed by the utilization of a 2D backbone and a detection head for further processing. By leveraging the learned knowledge from scene completion, the pre-training weights gain a deeper understanding of the scene, resulting in improved accuracy and reliability in the 3D detection results.




Figure 4. The pipeline of the 3D detector VoTr. VoTr mainly contains voxelization, 3D backbone, BEV map, 2D backbone, and detection head modules.

4. EXPERIMENTS

This section first provides a comprehensive description of the datasets used for both the pre-training and fine-tuning detection stages. Subsequently, we present and discuss the detection results, shedding light on the performance and effectiveness of our SCP.

4.1 Datasets

SemanticKITTI dataset. We utilize SemanticKITTI (Behley et al., 2019) as the pre-training dataset. SemanticKITTI is a pioneering real-world outdoor benchmark specifically designed for 3D semantic scene completion. The dataset consists of volumetric representations of the 3D scenes, with dimensions of $256 \times 256 \times 32$ voxels. Empty voxels, which do not contain any points, are labeled accordingly. For our training phase, we utilize a total of 19,130 pairs of input and target voxel grids, while reserving 815 pairs for validation purposes.

KITTI detection benchmark. To evaluate the effectiveness of our SCP in 3D object detection, we use the KITTI benchmark (Geiger et al., 2013), which comprises 7,481 samples for training and 7,518 for testing. Following VoTr (Mao et al., 2021), the training dataset is further subdivided into 3,712 samples for training data and 3,769 samples for validation data (Chen et al., 2015).

4.2 Evaluation metric

Following (Hossin and Sulaiman, 2015), we first employ the Intersection over Union (IoU) metric for the 3D scene completion. In 3D object detection, we then use the Average Precision (AP) with 11 recall points (AP11) for performance assessment.

Intersection over union (IoU). IoU is a metric used to measure the overlap between a model's predictions and the ground truth. It quantifies the degree of alignment between the two sets. An IoU of 0 indicates no overlap, meaning there is no intersection between the predicted regions and the ground truth. Conversely, an IoU of 1 represents a complete overlap, where the predicted regions perfectly match the ground truth. In practice, a larger IoU indicates better algorithm performance. The IoU is calculated as follows:

$$IoU = \frac{TP}{TP + FP + FN}, \quad (1)$$

where TP (True Positive) denotes the correct prediction of the ground truth. FP (False Positive) denotes incorrectly predicting an object that is not in the ground truth, and FN (False Negative) denotes the ground truth not being predicted.

Average precision 11 (AP11). The Average Precision (AP) is computed by integrating the Precision-Recall (PR) curve. To approximate the PR curve, the 11-point interpolation method

is commonly employed. In the 11-point interpolation, the Precision-Recall curve is sampled at 11 equally spaced recall levels between 0 and 1. The formula for AP₁₁ is as follows:

$$AP_{11} = \frac{1}{11} \sum_{R \in \{0, 0.1, \dots, 0.9, 1\}} P_{\text{interp}}(R) \quad (2)$$

where $P_{\text{interp}}(R)$ is the maximum precision for recall greater than R (Recall).

4.3 Scene completion pre-training

The raw point cloud serves as input and is fed into our SCP. Notably, only 20% of the training data is sampled from the overall training set for training the SCP. Through this step, we obtain the pre-training model for scene completion. As depicted in Fig. 5, it is evident that after pre-training the complementary network, the original data is effectively completed and generates a more holistic representation of the scene.




Figure 5. The comparison before and after the scene completion.

4.4 3D object detection

During this stage, the objective is to train and optimize the components of the detection network that complement the fully trained 3D backbone. The pre-training weights in the scene completion provide a solid starting point and enable efficient transfer of prior knowledge. By building upon the already well-performing 3D backbone part, we can concentrate on refining and fine-tuning the other components of the detection network. To evaluate the effectiveness of our SCP, we employ the entire validation dataset and compare it with the fully trained 3D detectors.

4.5 Implementation details

The pre-training scene completion module employed the Adam optimizer with a batch size of 3 and an initial learning rate of 0.001. The learning rate dynamically changed during training based on the number of training rounds. The training process utilized 20% of the SemanticKITTI dataset, with approximately 50 epochs of training conducted. Regarding the 3D detection module, the Adam-one cycle optimizer was employed with a batch size of 6 for SCP-SSD and 3 for SCP-TSD. The learning rate was set to 0.003. The SCP-SSD model was trained for 100 epochs, while the SCP-TSD model was trained for 80 epochs. All experiments and training processes were conducted on GPUs, specifically using hardware such as the Quadro RTX 8000 and NVIDIA A40. These GPUs offer high-performance computing capabilities and possess a substantial memory capacity of 48GB.

4.6 Results

We conduct a comparative analysis by assembling our SCP and the state-of-the-art 3D detector VoTr (Mao et al., 2021). To

ensure a fair comparison, we adopt the same evaluation metrics, following (Mao et al., 2021). Table. 2 presents the AP₁₁ achieved by each method on the KITTI validation split for the car category. Table. 1 illustrates the results of scene completion. Despite the utilization of a relatively small amount (20%) of training data, our SCP exhibits remarkable performance in object detection tasks. It is worth highlighting that our method achieves detection results that are remarkably close to those obtained through full data training, even surpassing the performance of full-data training.

3D scene completion. In this study, we utilized only 20% of the SemanticKITTI dataset for training the completion network. The completion results obtained in our experiments ranged from 47.39% to 54.58%. These results reflect the effectiveness of our completion network in generating accurate and reliable scene completions. Notably, our pre-training model, which underwent rigorous training and optimization, yielded the highest completion result of **54.58%**. The training results of the scene completion network are shown in Fig. 1.

Table 1. The IoU results of 3D scene completion in different training epochs.

Training rounds	IoU(%)
epochs=10	51.67
epochs=20	52.11
epochs=30	53.79
epochs=40	54.14
epochs=50	54.48

3D object detection. In this experiment, we focus on training the VoTr-SCP model using a significantly reduced amount of training data. Specifically, we utilized only 20% of the available data for training purposes. This deliberate reduction in training data allowed us to investigate the model's performance under resource-constrained scenarios and assess its ability to leverage limited data effectively. The results are presented in Table. 2. Remarkably, the reduction in training data does not lead to a significant degradation in the detection performance of the VoTr-SCP model. Despite utilizing a smaller amount of training data, our approach maintains a highly effective detection performance that is comparable to, and in some cases even surpasses, the results obtained through full data training. Here, we have only calculated the detection results on the car category.

Specifically, our SCP demonstrates notable improvements in the easy car class, elevating the performance from 89.04 to **89.10** compared to VoTr-TSD. This finding highlights the effectiveness of our approach in enhancing the detection results. Additionally, we provide a visualization of a scene completion result in Fig. 6.

5. DISCUSSIONS

Different labeled data volumes. To comprehensively assess the influence of the pre-training model on the detection performance under different amounts of labeled data, we carried out experiments utilizing varying fractions of the training set. Specifically, we conducted experiments using four different fractions: 10%, 20%, 30%, and 40% of the total training set.

The results of these experiments are presented in Table. 3 and Table. 4. It is evident from the tables that as the volume of data

Table 2. Comparisons on the KITTI with AP11 for the car category (20% training data).

Methods	AP11(%)		
	Easy	Mod	Hard
VoTr-SSD (Mao et al., 2021)	87.86	78.27	76.93
VoTr-SSD-SCP (Ours)	86.86	75.75	68.79
VoTr-TSD (Mao et al., 2021)	89.04	84.04	78.68
VoTr-TSD-SCP (Ours)	89.10	78.78	78.04

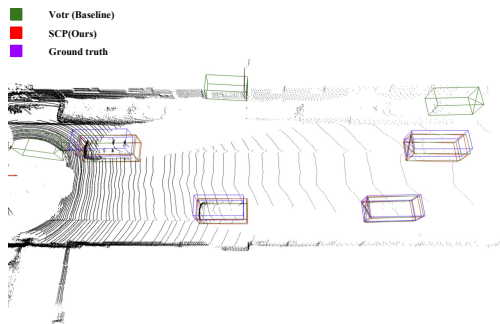


Figure 6. Visualization of the example results. Due to the existence of some points similar to cars in the target area, VoTr incorrectly identified the wall as a car (the green bounding box), whereas our SCP can help to reduce these errors.

increases, the performance of the models improves. Notably, our experimental results highlight the robustness and efficacy of our model, even when trained with very small amounts of data. At data volumes as low as 10% and 20% of the training set, our model consistently achieved commendable results.

Table 3. Comparisons on different data volume with VoTr-SSD-SCP (IoU=51.6) for the car category.

Data Volume	AP11(%)		
	Easy	Mod	Hard
10%	85.34	71.89	67.85
20%	85.91	72.36	68.25
30%	86.43	74.43	68.28
40%	86.79	76.02	72.01

Table 4. Comparisons on different data volume with VoTr-TSD-SCP (IoU=54.1) for the car category.

Data Volume	AP11(%)		
	Easy	Mod	Hard
10%	88.11	77.90	76.84
20%	89.10	78.78	78.04
30%	88.50	78.43	77.67
40%	89.06	78.95	78.44

Different completion effects. To investigate the influence of scene completions on the 3D detection performance, we conducted an in-depth analysis by selecting two different completion results with varying intersection over union (IoU). Specifically, we chose one completion result with an IoU of 51.6% and another with an IoU of 54.1%. We then compared the performance of these completion results across different volume datasets. Table.5 and Table.6 demonstrate that different scene

completion results yield varying effects on 3D detection for the VoTr-SSD-SCP and VoTr-TSD-SCP frameworks, respectively. Notably, when utilizing the completion results with better performance (IoU=54.1%), the detection results show significant improvements on the 10% and 20% datasets. We thus get two conclusions: (1) Accurate and reliable scene completions play a critical role in improving the overall detection performance, particularly in scenarios with very limited data. (2) As the completion achieves better results, it becomes increasingly beneficial for 3D object detection.

Table 5. Comparisons on different IoU with VoTr-SSD-SCP for the car category.

Completion Results	AP11(%)		
	Easy	Mod	Hard
IoU=51.6(10%)	85.34	71.89	67.85
IoU=54.1(10%)	85.73	72.63	67.38
IoU=51.6(20%)	85.91	72.36	68.25
IoU=54.1(20%)	86.86	75.75	68.79

Table 6. Comparisons on different IoU with VoTr-TSD-SCP for the car category.

Completion Results	AP11(%)		
	Easy	Mod	Hard
IoU=51.6(10%)	88.23	77.95	76.81
IoU=54.1(10%)	88.11	77.90	76.84
IoU=51.6(20%)	88.43	78.43	77.57
IoU=54.1(20%)	89.10	78.78	78.04

6. CONCLUSIONS

In this paper, we propose SCP, a novel scene completion pre-training network for 3D object detection. We have demonstrated that scene completion can learn a model initialization to help the 3D detectors only trained on a small amount of dataset. Experiments indicate that the quality of the scene completion has a positive correlation with the effectiveness of object detection. In the future, we hope to explore the scene completion pre-training into more downstream tasks, for example, 3D semantic segmentations and 3D forecasting.

REFERENCES

- Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., Gall, J., 2019. SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences. *Proc. of the IEEE/CVF International Conf. on Computer Vision (ICCV)*.
- Chen, X., Kundu, K., Zhu, Y., Berneshawi, A. G., Ma, H., Fidler, S., Urtasun, R., 2015. 3d object proposals for accurate object class detection. *Advances in neural information processing systems*, 28.
- Cheng, R., Agia, C., Ren, Y., Li, X., Bingbing, L., 2021. S3cnet: A sparse semantic scene completion network for lidar point clouds. *Conference on Robot Learning*, PMLR, 2148–2161.
- Geiger, A., Lenz, P., Stiller, C., Urtasun, R., 2013. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11), 1231–1237.

- Hossin, M., Sulaiman, M. N., 2015. A review on evaluation metrics for data classification evaluations. *International journal of data mining & knowledge management process*, 5(2), 1.
- Li, J., Han, K., Wang, P., Liu, Y., Yuan, X., 2020. Anisotropic convolutional networks for 3d semantic scene completion. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3351–3359.
- Li, X., Liu, Y., Xia, Y., Lakshminarasimhan, V., Cao, H., Zhang, F., Stilla, U., Knoll, A., 2023. Fast and deterministic (3+1) DOF point set registration with gravity prior. *ISPRS Journal of Photogrammetry and Remote Sensing*, 199, 118–132.
- Mao, J., Xue, Y., Niu, M., Bai, H., Feng, J., Liang, X., Xu, H., Xu, C., 2021. Voxel transformer for 3d object detection. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3164–3173.
- Sauder, J., Sievers, B., 2019. Self-supervised deep learning on point clouds by reconstructing space. *Advances in Neural Information Processing Systems*, 32.
- Sharma, C., Kaul, M., 2020. Self-supervised few-shot learning on point clouds. *Advances in Neural Information Processing Systems*, 33, 7212–7221.
- Shi, S., Guo, C., Jiang, L., Wang, Z., Shi, J., Wang, X., Li, H., 2020. Pv-rnn: Point-voxel feature set abstraction for 3d object detection. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10529–10538.
- Shi, S., Wang, X., Li, H., 2019. Pointcnn: 3d object proposal generation and detection from point cloud. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 770–779.
- Song, S., Yu, F., Zeng, A., Chang, A. X., Savva, M., Funkhouser, T., 2017. Semantic scene completion from a single depth image. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1746–1754.
- Wang, H., Liu, Q., Yue, X., Lasenby, J., Kusner, M. J., 2021. Unsupervised point cloud pre-training via occlusion completion. *Proceedings of the IEEE/CVF international conference on computer vision*, 9782–9792.
- Wang, Y., Tan, D. J., Navab, N., Tombari, F., 2022. Learning local displacements for point cloud completion. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1568–1577.
- Xia, Y., Gladkova, M., Wang, R., Li, Q., Stilla, U., Henriques, J. F., Cremers, D., 2023a. Casspr: Cross attention single scan place recognition. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, IEEE.
- Xia, Y., Wu, Q., Li, W., Chan, A. B., Stilla, U., 2023b. A Lightweight and Detector-Free 3D Single Object Tracker on Point Clouds. *IEEE Transactions on Intelligent Transportation Systems*.
- Xia, Y., Xia, Y., Li, W., Song, R., Cao, K., Stilla, U., 2021a. Asfm-net: Asymmetrical siamese feature matching network for point completion. *Proceedings of the 29th ACM international conference on multimedia*, 1938–1947.
- Xia, Y., Xu, Y., Li, S., Wang, R., Du, J., Cremers, D., Stilla, U., 2021b. Soe-net: A self-attention and orientation encoding network for point cloud based place recognition. *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 11348–11357.
- Xia, Y., Xu, Y., Wang, C., Stilla, U., 2021c. VPC-Net: Completion of 3D vehicles from MLS point clouds. *ISPRS Journal of Photogrammetry and Remote Sensing*, 174, 166–181.
- Xia, Z., Liu, Y., Li, X., Zhu, X., Ma, Y., Li, Y., Hou, Y., Qiao, Y., 2023c. Scpnet: Semantic scene completion on point cloud. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17642–17651.
- Xie, S., Gu, J., Guo, D., Qi, C. R., Guibas, L., Litany, O., 2020. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, Springer, 574–591.
- Yan, X., Gao, J., Li, J., Zhang, R., Li, Z., Huang, R., Cui, S., 2021. Sparse single sweep lidar point cloud segmentation via learning contextual shape priors from scene completion. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35number 4, 3101–3109.
- Yan, Y., Mao, Y., Li, B., 2018. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10), 3337.
- Yang, Z., Sun, Y., Liu, S., Jia, J., 2020. 3dssd: Point-based 3d single stage object detector. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11040–11048.
- Zhou, Y., Tuzel, O., 2018. Voxelnet: End-to-end learning for point cloud based 3d object detection. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4490–4499.