

# PCGen: Point Cloud Generator for LiDAR Simulation

Chenqi Li, Yuan Ren and Bingbing Liu

Huawei Noah's Ark Lab, Toronto, Canada

{chenqi.li, yuan.ren3, liu.bingbing}@huawei.com

**Abstract**—Data is a fundamental building block for LiDAR perception systems. Unfortunately, real-world data collection and annotation is extremely costly & laborious. Recently, real data based LiDAR simulators have shown tremendous potential to complement real data, due to their scalability and high-fidelity compared to graphics engine based methods. Before simulation can be deployed in the real-world, two shortcomings need to be addressed. First, existing methods usually generate data which are more noisy and complete than the real point clouds, due to 3D reconstruction error and pure geometry-based raycasting method. Second, prior works on simulation for object detection focus solely on rigid objects, like cars, but Vulnerable Road User (VRU)s, like pedestrians, are important road participants. To tackle the first challenge, we propose First Peak Averaging (FPA) raycasting and surrogate model raydrop. FPA enables the simulation of both point cloud coordinates and sensor features, while taking into account reconstruction noise. The ray-wise surrogate raydrop model mimics the physical properties of LiDAR's laser receiver to determine whether a simulated point would be recorded by a real LiDAR. With minimal training data, the surrogate model can generalize to different geographies and scenes, closing the domain gap between raycasted and real point clouds. To tackle the simulation of deformable VRU simulation, we employ Skinned Multi-Person Linear model (SMPL) dataset to provide a pedestrian simulation baseline and compare the domain gap between CAD and reconstructed objects. Applying our pipeline to perform novel sensor synthesis, results show that object detection models trained by simulation data can achieve similar result as the real data trained model.

## I. INTRODUCTION

The success of deep learning is deeply rooted in the availability of large-scale, high-fidelity datasets. Pioneering datasets [1]–[13] facilitate development of cutting edge visual recognition systems, providing challenging benchmarks for the community. However, collection and annotation of data in the real world is very inefficient, slow and uneconomical. Simulation, on the other hand, give users the flexibility to generate diverse scenarios with ease, as well as providing automatically generated ground truth annotations. For LiDAR simulation, two distinct approaches have been explored: graphics engine based and real data based. Results show that real data based methods produces simulation data with lower domain gap compared to graphics engine based methods. However, real data based methods suffer from noisy reconstruction and sim-to-real domain gap. This paper tackles these challenges and makes the following contributions:

- We present FPA raycasting to simulate LiDAR point clouds and sensor features, accounting for noise in the reconstructed scenes.
- We develop the surrogate model of a single laser head and use it for raydrop. Comparing to the UNet

based raydrop method, the proposed method is scene-independent. The surrogate model of a specific LiDAR can be trained once and used in different scenes.

- We perform novel sensor synthesis with our simulation pipeline. The test results show that it provides high-fidelity data for the new sensor configuration, achieving similar result as the model trained on the real data.
- We propose Learned Point Cloud Similarity (LPCS) metric to measure domain gap between real and simulation point clouds, from the perspective of perception models
- We provide a baseline pedestrian simulation result, using SMPL and reconstructed human models

## II. RELATED WORK

### A. Graphics Engine Based LiDAR Simulator

Initial attempts to LiDAR simulation were spearheaded by Car Learning to Act (CARLA) [14]. Building on top of Unreal Engine 4 (UE4) [15], CARLA's simulation platform allows the user to customize scenarios, including agent model, density, interaction with the world, weather conditions, and sensor suite. Similarly, Yue et al. [16] leveraged the popular, high-fidelity simulation of Grand Theft Auto V (GTA V) to automatically extract point cloud with ground truth labels. The framework also enables users to construct diverse, customized scenarios interactively to test neural network performance in corner cases. Experiments have shown that retraining with additional synthetic point cloud significantly improves model's performance on KITTI dataset [17], [18]. This work is further extended by the Precise Synthetic Image and LiDAR (PreSIL) dataset [19], which improves the raycasting functionality within GTA V to address the issues of approximating human with cylinders and missed ray-scene collisions. PreSIL provides a large simulated dataset in KITTI format, and demonstrate that it can boost the performance on state-of-art model in KITTI 3D Object Detection benchmark.

### B. Real Data Based LiDAR Simulator

Generating CAD model assets and complex scenarios are labor-intensive, making simulation difficult and costly to scale. Furthermore, domain gap between noiseless simulation and real-world data leads to poor model performance if only trained on simulation data, prompting the development of domain adaptation techniques [18], [20], [21]. Fang et al. [22] investigated a hybrid, data-driven approach to point cloud generation framework, which combines real world scanned background point cloud and synthetic foreground objects. They show that by augmenting the real dataset

with synthetic frames, instance segmentation and object detection performance is improved. Concurrently, LiDARsim [23] employed a similar approach, leveraging real data to reconstruct both background and foreground objects. LiDARsim further extended the simulation with the addition of a learning system to model the physics of LiDAR raydrop, closing the gap between real and simulation point clouds. They show that the simulation data trained models can obtain similar performance in object detection and semantic segmentation as models trained using real data, without domain adaptation techniques. Similar to LiDARsim, Langer et al. [24] developed a simulation pipeline for domain transfer in the context of semantic segmentation. Their results show that closest point raycasting, along with geodesic correlation alignment, successfully generated simulation data to adjust a model trained on the source domain (Velodyne-64) to the target domain (Velodyne-32).

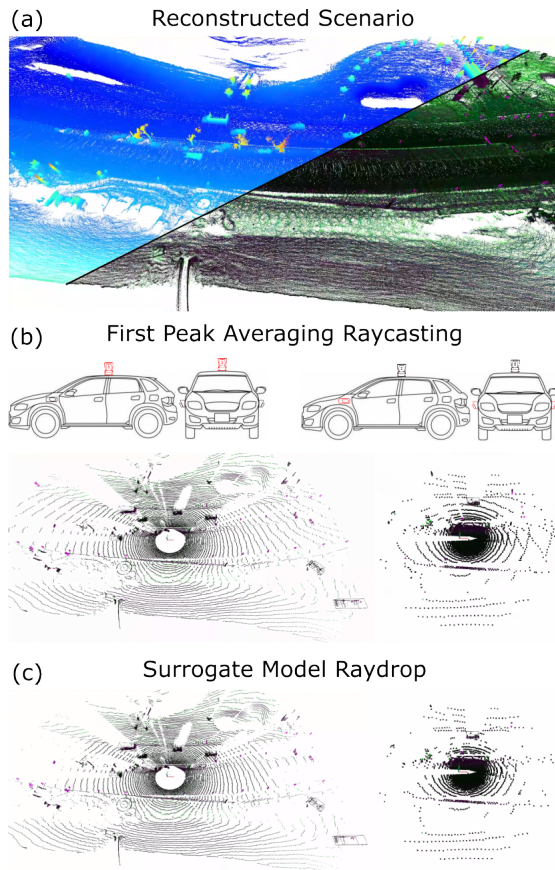


Fig. 1: Overview of Simulation Pipeline: (a) A dense point cloud representation of the map is constructed using Iterative Closest Point (ICP), from sparse individual frames whose objects have been cropped out. Road participants are also reconstructed with ICP and re-inserted into the dense point cloud map to simulate the desired scenarios. (b) LiDAR with a predefined configuration is placed into the scene. FPA raycasting simulates point cloud collected by real LiDAR. (c) Point cloud post-processed by surrogate model raydrop algorithm to remove data point that are unlikely to be observed by a real LiDAR.

### III. METHODOLOGY

#### A. 3D Scene Reconstruction

Given a sequence of single frame point clouds, 3D scene reconstruction aims to generate a dense point cloud mapping of the background. Ground truth bounding box annotations are used to crop out foreground object points in each frame. Since annotation for dynamic objects are less accurate, the bounding box dimensions are slightly enlarged before cropping, to ensure complete removal of all foreground points. Using odometry, cropped single frame point cloud are transformed into the global coordinate frame and subsequently accumulated to obtain a dense 3D reconstruction of the sequence. Voxel downsampling and point cloud radius outlier removal are performed as post-processing steps, in order to reduce memory requirements and remove noisy points.

For object reconstruction, point cloud clusters belonging to the same object instance are cropped from individual frames. Each cluster is returned to x-axis aligned origin using bounding box location and orientation. To reduce the impact of human annotation and odometry imperfections, a generalized ICP [25] can be used to improve the alignment. However, if the source and/or target point cloud contain very few points, minimizing the point cloud distance using ICP often leads to unreasonable reconstructions. Thus, we adaptively employ ICP if both source and target point cloud contains greater than threshold number of points, otherwise, object clusters are simply accumulated. The foreground objects can then be inserted into the background to simulate a variety of scenarios.

#### B. Raycasting method

Raycasting emits rays based on predefined LiDAR configuration and perform collision checks with the reconstructed scenario, in order to simulate single frame point clouds observed by real LiDARs. Theoretically, the reconstructed 3D scene is a 2D surface embedded in the 3D space, represented in the form of a dense point cloud. LiDARsim rendered dense point cloud as surfels [26] and used Intel Embree Engine to compute ray-disc intersections [27]. Langer et al. used Closest Point (CP) raycasting, which projects dense point cloud into a range image, and the closest point from each pixel is extracted to render the raycasted point cloud [24]. However, localization, calibration and sensor synchronization are subject to error, points in the reconstructed point cloud do not strictly lie on 2D surface of the scene. Both raycasting methods suffer from noisy reconstructions, leading to noisy raycasted point clouds. Moreover, CP raycasting only works with spinning scan LiDAR. For LiDARs with irregular scan pattern, such as the DJI Livox, multiple rays might land in the same pixel, leading to redundant points.

To solve this problem, we propose First Peak Averaging (FPA) raycasting. First, the reconstructed point cloud is projected into a range image. Each pixel of the range image corresponds to a frustum in the 3D space, as shown in Figure 2(a). Each point within the frustum can be defined using the Spherical coordinates  $(d, \lambda, \phi)$  or the Cartesian coordinates

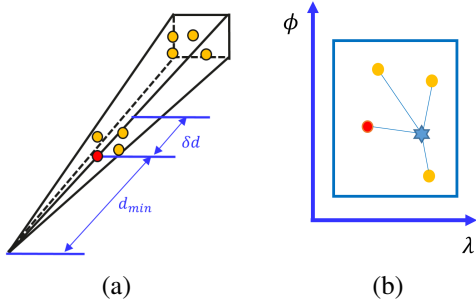


Fig. 2: First peak averaging raycasting

$(x, y, z)$ . Typically, the depth distribution of all points within the frustum forms multiple peaks, due to occlusion and observation of the scene from multiple angles. The intuition behind FPA raycasting is to average points from the closest peak, in order to estimate the true 2D surface from the noisy 3D point cloud. We are interested in the closest peak, because it corresponds to the first intersection between the ray and the 3D scene. To speed up the simulation, a fixed peak width  $\delta d$  is used to screen the first peak points. In Figure 2(b),  $\lambda_l$  and  $\phi_l$  are the azimuth and elevation of the casted laser beam (indicated by the blue star),  $\lambda_i$  and  $\phi_i$ ,  $i = 1, 2, \dots, N$  are the azimuth and elevation of each point within the closest peak (indicated by the red and yellow circles, with the red circle being the closest point within the frustum). The intersection between the ray and the 3D scene is estimated with the Inverse Distance Weighted (IDW) averaging.

$$f = \sum_{i=0}^N w_i f_i \quad (1)$$

where  $f$  can be the  $x$ ,  $y$ ,  $z$  coordinates or the feature, such as the intensity, of a point. The inverse distance weights  $w_i$  and inverse distance  $d_i$  are defined as:

$$w_i = \frac{\tilde{d}_i}{\sum_{j=1}^N \tilde{d}_j} \quad (2)$$

$$\tilde{d}_i = \frac{1}{\sqrt{(\lambda_i - \lambda_l)^2 + (\phi_i - \phi_l)^2}} \quad (3)$$

### C. Raydrop

The FPA raycasting algorithm records every ray-scene intersection, without taking into account the aforementioned factors. However, real laser returns are affected by many factors, such as the distance, incidence angle, material reflectivity and atmospheric composition. Raydrop aim to reduce the existing domain gap, where the simulated point clouds are more geometrically complete than real point clouds. LiDARsim developed a 2D UNet for raydrop [28], trained using pairs of simulation and real point cloud range images. This requires strict pixel-wise correspondence between the two range images, and it cannot be guaranteed due to odometry and calibration error. Furthermore, deep learning models require large amounts of data to generalize. With convolution layers, the encoder-decoder network learns to drop rays based on context. Generalization to different

geographies will be limited. A real LiDAR, however, does not use the wide perceptive field of the UNet to determine laser return. To address these problems, we propose a surrogate model (MLP) for the laser head.

$$r = MLP(d, \theta, i) \quad (4)$$

where  $r$  is a ray's return probability,  $d$ ,  $\theta$ ,  $i$  are the distance, incidence angle and simulation intensity of the ray. By learning ray-wise drop/return probability, the surrogate model does not require pixel-wise range image correspondence between the real and the simulation point clouds, nor large amounts of data across geographic regions.

Given a pair of raycasted and real dataset, the  $d$ ,  $\theta$  and  $i$  of each point can be computed and projected into the parameter space, as shown in Fig. 3(a).

$$d = \sqrt{x^2 + y^2 + z^2} \quad (5)$$

where  $x, y, z$  is Cartesian coordinate of a raycasted point in the LiDAR frame

$$\theta = \arccos \frac{\vec{R} \cdot \vec{N}}{\|\vec{R}\| \|\vec{N}\|} \quad (6)$$

where  $\vec{R}$  is the ray vector,  $\vec{N}$  is the point's normal vector, computed from single frame point cloud using Open3D [29]. Intensity,  $i$ , is taken directly from the real point cloud. For the FPA simulation point cloud, the raycasted intensity is the average reflection intensity of laser beams from different directions, which abstractly describes the material's reflectivity.

To train the model, we pair the real dataset with a simulation dataset, which is raycasted using the real dataset's LiDAR poses. With sufficient data, the parameter space will be filled, which can then be voxelized. The ratio between the number of real points and the number of simulation points,  $r$ , can be computed for each voxel. This represents the probability the ray would be returned, at the given  $d$ ,  $\theta$ ,  $i$ . The parameter space can be converted directly into a lookup table for inference. We approximate the parameter space with an MLP, to enable GPU inference acceleration and infer values for voxels without data.

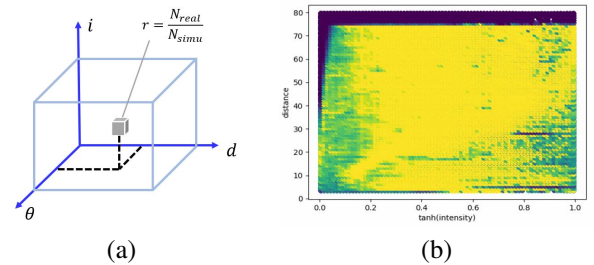


Fig. 3: Input parameter space of the surrogate model

Figure 3(b) shows an example 2D projection of the voxelized input parameter space. Each pixel shows the ray's return probability, with yellow and purple representing a probability of 1 and 0 respectively. A LiDAR point cloud cutoff at 75 meters is observed, and the rays that intersect

farther away from the LiDAR require greater intensity to be returned. The parameter space is not very smooth, due to voxelization. In order to reduce the influence of the unsmoothness, the data of a voxel will not be used for training if the number of simulation points in the voxel is less than the given threshold or the number of the real points is larger than that of the simulation points.

#### D. Learned Point Cloud Similarity (LPCS)

Raycasting algorithm requires the selection of three hyperparameters, the range image width & height and averaging peak width. To tune hyperparameters, a point cloud similarity metric is required. By computing the sum of squared distances between nearest neighbors, Chamfer Distance (CD) [30], [31] does not take into account mismatching local density. Such low-level features, however, will be extracted by backbones of the perception architectures and influence model prediction. Inspired by Learned Perceptual Image Patch Similarity (LPIPS) [32], we propose LPCS, which compares the backbone feature distance between pairs of point clouds from a model pretrained on real data. Given a pair of corresponding point cloud, all objects of interest can be cropped out using bounding box annotations, yielding two subset point clouds for simulation and real,  $P_s, P_r \in \mathbb{R}^3$ , respectively. LPCS can be computed via the absolute element-wise distance between the feature of both subsets:

$$LPCS(P_s, P_r) = \sum_{i=1}^N |F(P_s)_i - F(P_r)_i| \quad (7)$$

where  $F$ , the model backbone, outputs a high-dimensional vector of  $N$  terms

#### E. Novel Sensor Synthesis

With reconstructed maps and objects, a dense representation of any single frame can be obtained by inserting the objects using their poses from the real frame. By providing the raycasting algorithm with 1. the elevation and azimuth angle of each ray with respect to the LiDAR frame 2. the pose of the LiDAR with respect to the global frame, we can simulate a corresponding point cloud, under a new LiDAR configuration. Note that if the new LiDAR's Field of View (FOV) dramatically exceed that of the original dataset, part of the new dataset might be missing. This can be circumvented by collecting the original dataset using the LiDAR with largest FOV, and optionally, highest density. Our simulation pipeline greatly reduces the data collection and annotation costs to train a model that can generalize to any combination of car and LiDAR models.

#### F. Pedestrian Simulation

Previous works on simulation for object detection focused primarily on rigid object classes, such as car. We attempt to expand simulation classes to VRU, in particular, pedestrians. Following the same reconstruction procedure from Section III-A, a pedestrian object library is generated. However, body movements lead to deformed reconstructions. To address this problem, we employ the SMPL [33] dataset, which provides

realistic 3D human CAD models. Unlike cars, pedestrians share similar geometry and mainly differ in their postures. Thus, we believe that the SMPL models provide sufficient degrees of freedom to capture the diverse set of pedestrian poses that can be found on the road. During raycasting, CAD models are converted to point clouds by sampling a large number of points on the surface of triangle meshes. A tight bounding box can also be generated to enclose all points. Using statistics, dimensions of the real bounding boxes and its enclosed point cloud cluster can be computed. The distributions are used to filter out unrealistic CAD poses, as well as loosening the bounding box annotations to match human annotations.

## IV. EXPERIMENTAL EVALUATION

Through the following experiments, we intend to demonstrate the fidelity and potential applications of our pipeline.

- 1) To verify the fidelity of our raycasting and raydrop algorithms, we show that object detection models trained using real data or simulation data, achieves similar performance when evaluated on the real validation data.
- 2) To prove LPCS as an useful point cloud similarity metric, we show negative correlation between LPCS and object detection Mean Average Precision (mAP)
- 3) To demonstrate sensor type conversion, we show that with only source sensor data and 10% annotated target sensor data, we can simulate target sensor and achieve similar object detection performance as 100% annotated target sensor data.
- 4) To demonstrate VRU simulation, we provide a baseline for pedestrian simulation and compare CAD models against reconstructed pedestrian.

#### A. Experimental Setup

1) *Dataset*: We evaluate our simulation pipeline using Waymo Open Dataset, which consists of 800 training and 200 validation segments spanning over different geographical locations [3]. Each segment provides synchronized sensor data over 20 seconds. Each LiDAR frame is composed of 5 LiDARs: top, left, right, front and rear, captured at 10Hz. For all of our experiments, we perform reconstruction and mapping using top LiDAR, due to its high density and long range of 75 meters.

2) *Object Detection Model*: OpenPCDet [34] implementation of Centerpoint [35] will be used to evaluate the quality of our simulation dataset. During training, no data augmentation is performed, in order to minimize stochasticity among trials. All models are trained over 80 epochs.

#### B. First Peak Averaging Raycasting Ablation

With 2D and 3D object detection, we quantitatively benchmark CP against FPA. Given a real LiDAR frame, a dense representation can be obtained via reconstruction. Two simulation datasets, which are identical to the real dataset, can be raycasted using CP and FPA. Both datasets employ the same LiDAR pose and ray configurations from Waymo Open Dataset's calibration. Centerpoint is trained using the

TABLE I: Top LiDAR Raycasting (Real Top Validation)

(a) XYZ			
	3D mAP IoU 0.5/0.7	2D BEV mAP IoU 0.5/0.7	
Simulation CP	74.06/41.93	77.93/61.30	
Simulation FPA	74.41/41.98	78.84/62.61	
Real	77.59/47.27	81.19/66.01	
(b) XYZ + Intensity + Elongation			
	3D mAP IoU 0.5/0.7	2D BEV mAP IoU 0.5/0.7	
Simulation CP	74.44/40.33	79.46/61.86	
Simulation FPA	76.16/42.18	80.56/63.75	
Real	79.23/48.21	82.57/67.24	

CP, the FPA and the real dataset. The trained models are evaluated on the real validation dataset, which is summarized in Table I. When trained using only XYZ, FPA method provides a small performance gain of +0.35/+0.05% and +0.91/+1.31% for 3D and 2D  $mAP_{0.5/0.7}$ . This shows that the FPA, compared to CP, reduces noise in the simulated point clouds. When point cloud features (intensity, elongation) are included during training, FPA provides +1.72/+1.85%, +1.10/+1.89% improvement for 3D and 2D  $mAP_{0.5/0.7}$ . More importantly, the inclusion of point cloud features lead to improved model performance for FPA, but worsened 3D  $mAP_{0.7}$  model performance for CP. This indicates that the simulated point cloud feature by CP compared to FPA is also more noisy and can cause confusion for the model.

### C. Raydrop Ablation

To show the generalization capability of the proposed raydrop model, we will use a surrogate model trained on the validation set to drop points in the training set. Following IV-B, FPA simulation of the real validation dataset is first generated. Following Section III-C, simulation and real validation dataset is used to train the surrogate model. For each point in the raycasted simulation training set, the surrogate model checkpoint provides a probability the point should be kept. By choosing a threshold, probabilities can be converted to binary drop/keep masks. The higher the threshold, the more points will be dropped. The trained model’s performance on the real validation set is summarized in Table II. "No raydrop" indicates original raycasting results, and "Real" indicates real data. Since XYZIE is used for training, these two rows are identical to "Simulation FPA" and "Real" rows of Table I(b). At threshold 0.28, 3D  $mAP_{0.5/0.7}$  is maximized at 77.56/44.93%, a +1.40/+2.75% improvement compared to without raydrop. 2D  $mAP_{0.5/0.7}$  also increases +1.51/+1.49% and +0.65/+2.00% at 0.30 and 0.28 threshold, respectively. Thus, the surrogate model is an extremely light-weight and generalizable method to effectively reduce the sim2real domain gap. Figure 4(a) provides a visual comparison of point cloud with and without raydrop. After raydrop, simulation point cloud more closely matches that of the real point cloud.

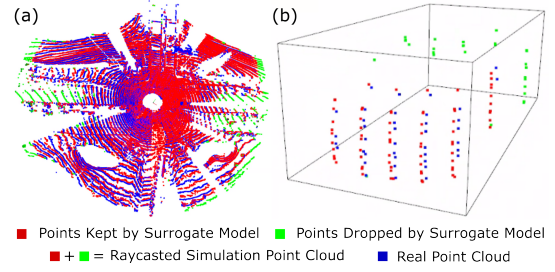


Fig. 4: Surrogate Model Raydrop Visualization (a) Top LiDAR raydrop (b) Side LiDAR raydrop

TABLE II: Top LiDAR Raydrop (Real Top Validation)

Threshold	3D mAP IoU 0.5/0.7	2D BEV mAP IoU 0.5/0.7
No raydrop	76.16/42.18	80.56/63.75
0.28	77.56/44.93	81.21/65.75
0.30	77.36/42.80	82.07/65.24
0.32	77.22/43.63	81.30/64.80
0.34	76.63/41.24	80.51/63.02
Real	79.23/48.21	82.57/67.24

### D. Learned Point Cloud Similarity

We perform a small hyperparameter search using 10% of the validation set. Using different combinations of hyperparameters (averaging peak width, range image width, range image height), new simulation validation datasets are raycasted. With a real data trained model, we can obtain the mAP of each raycasted dataset, as well as computing the LPCS between the raycasted and the real dataset.

Figure 5 shows a negative correlation between LPCS and mAP, which implies that LPCS provides a good estimate of the domain gap between the simulation and the real dataset. Using LPCS to guide gridsearch of raycasting hyperparameters, we conclude that the optimal range image dimension is  $2560 \times 128$ , along with an averaging peak width of 20cm. These hyperparameters are used in all of the top LiDAR experiments.

### E. Novel Sensor Synthesis

In this section, we showcase an example novel sensor synthesis, from top to side LiDARs. Following Section III-E, side LiDAR dataset can be raycasted using top LiDAR reconstructions. First 10% of the real and raycasted side LiDAR training set is used to train a surrogate raydrop model, to raydrop the entire raycasted training set. Using the first 10% and random 10% of the real side LiDAR training set, the

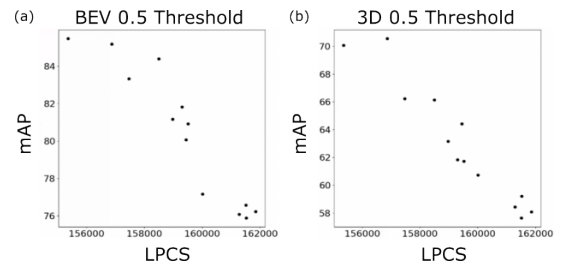


Fig. 5: Correlation between LPCS and Object Detection mAP

TABLE III: Side LiDAR Simulation (Real Side Validation)

(a) Simulation Ablation				
Raycast	Raydrop	Finetune	3D mAP IoU 0.5/0.7	2D BEV mAP IoU 0.5/0.7
✓	✗	✗	61.98/28.25	68.98/51.00
✓	✓	✗	77.00/36.05	81.99/63.19
✓	✓	First 10%	82.09/45.86	85.79/66.46
✓	✓	Random 10%	82.99/49.90	85.77/69.45

(b) Real Data Baselines				
Dataset Sensor	Volume	3D mAP IoU 0.5/0.7	2D BEV mAP IoU 0.5/0.7	
Top	All 100%	62.31/32.39	65.02/51.70	
Side	First 10%	75.22/35.06	79.63/56.95	
Side	Random 10%	83.77/49.57	86.96/68.39	
Side	All 100%	85.28/52.84	88.02/70.78	

simulation data trained models are further finetuned. Object detection performance, evaluated on the real side LiDAR validation set, is summarized under Table III(a). Some real data trained baselines are also provided in Table III(b). If a real top LiDAR trained model (Top 100%) is evaluated on the real side validation set, we observe around 20% domain gap compared to the real side LiDAR trained model (Side 100%). A similar performance is observed for raycasted simulation dataset (Raycast only). However, raydrop (Raycast + Raydrop) is extremely useful. Compared to (Raycast only), we observe +15.02/+7.80% and +13.01/+12.19% improvement. Compared to (Side First 10%), we observe +1.78/+1.01% and +2.36/+6.24% improvement. Figure 4b shows that after dropping the scanlines around the top and rear sections of the car, the raydropped point cloud is a lot more similar compared to the real point cloud than the raycasted point cloud. The reduced domain gap is likely the main contributing factor to the improved performance. Finetuning with first 10% further boosts model performance (Raycast + Raydrop + First 10%). Compared to (Side First 10%), we observe +6.87/+10.8% and +6.16/+9.51% improvement. Compared to (Side All 100%), we observe a gap of -3.19/-6.98% and -1.17/-1.93%. This suggests that given a large scale primary dataset of an old sensor and a small scale secondary dataset of a new sensor, our simulation pipeline can leverage both datasets and train a model that is capable of achieving similar performance compared to a model trained on the large scale dataset of the new sensor. Alternatively, the simulation pipeline can be thought of as a data augmentation generator. The large scale primary dataset can be converted to augmentation frames for the secondary dataset, to improve the model's generalization capability on the secondary dataset.

#### F. Deformable Object Simulation

Table IV compares the pedestrian simulation using reconstructed and CAD pedestrians, when evaluated on the real top validation set. "Real" represents real data trained model. "CAD Naive" represents randomly sampling from the CAD library and replacing real pedestrians. "CAD

TABLE IV: Pedestrian Simulation (Real Top Validation)

Training Dataset	3D mAP IoU 0.25/0.50	2D BEV mAP IoU 0.25/0.50
CAD Naive	55.62/06.56	59.31/11.92
Reconstruction	69.76/47.14	70.49/55.44
CAD Modified	74.56/54.31	75.19/62.47
Real	78.31/61.62	79.24/68.31

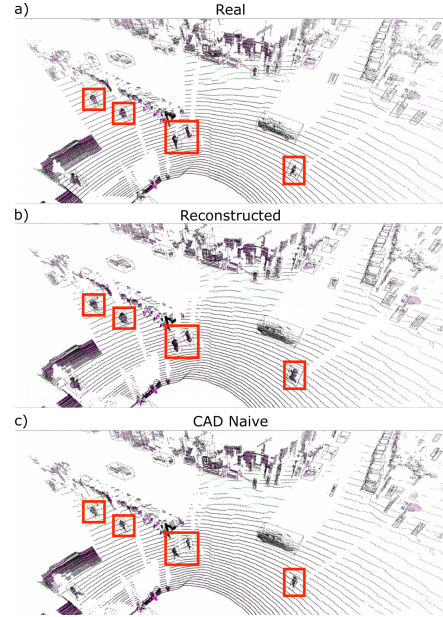


Fig. 6: Pedestrian Simulation Visualization

Modified" improves upon "CAD Naive" with the inclusion of pose filtering and bounding box adjustment as described in Section III-F. "Reconstructed" replaces real with reconstructed pedestrians. Figure 6 provides a visual comparison of the above experiments. The reconstructed pedestrians leads to noticeably thicker outline compared to the real and CAD point clouds.

Compared to car classes, reconstructed pedestrians suffer from increased sim-real domain gap of -8.55/-14.48% and -8.75/-12.87% for 3D and 2D  $mAP_{0.5/0.7}$ . "CAD Modified" achieves the lowest domain gap of -3.75/-7.31% and -4.05/-5.84%. Our results provides a baseline for VRU simulation and show that SMPL is a viable alternative to replace reconstruction for pedestrians.

#### V. CONCLUSION

In this work, we propose a point cloud based simulation pipeline. Experiments show that the pipeline is capable of transferring a dataset collected by an old sensor to recreate the LiDAR stream that would have been collected by a new sensor configuration of different density, placement and scanning mechanism. It greatly reduces the cost of data collection and annotation to generalize model performance from a particular vehicle-LiDAR setup to any desired combination. To harness the full potential of our pipeline, we look to close the domain gap between simulated and real pedestrians by augmenting the SMPL dataset with accessories, such as backpacks and handbags, as well as expanding VRU simulation to cyclists.

## REFERENCES

- [1] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [2] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631.
- [3] P. Sun, H. Kretschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, *et al.*, "Scalability in perception for autonomous driving: Waymo open dataset," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 2446–2454.
- [4] M.-F. Chang, J. Lambert, P. Sangkloy, J. Singh, S. Bak, A. Hartnett, D. Wang, P. Carr, S. Lucey, D. Ramanan, *et al.*, "Argoverse: 3d tracking and forecasting with rich maps," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8748–8757.
- [5] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. K. Pontes, *et al.*, "Argoverse 2: Next generation datasets for self-driving perception and forecasting," in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [6] X. Huang, X. Cheng, Q. Geng, B. Cao, D. Zhou, P. Wang, Y. Lin, and R. Yang, "The apolloscope dataset for autonomous driving," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2018, pp. 954–960.
- [7] J. Houston, G. Zuidhof, L. Bergamini, Y. Ye, L. Chen, A. Jain, S. Omari, V. Iglovikov, and P. Ondruska, "One thousand and one hours: Self-driving motion prediction dataset," *arXiv preprint arXiv:2006.14480*, 2020.
- [8] J. Mao, M. Niu, C. Jiang, H. Liang, J. Chen, X. Liang, Y. Li, C. Ye, W. Zhang, Z. Li, *et al.*, "One million scenes for autonomous driving: Once dataset," *arXiv preprint arXiv:2106.11037*, 2021.
- [9] Y. Choi, N. Kim, S. Hwang, K. Park, J. S. Yoon, K. An, and I. S. Kweon, "Kaist multi-spectral day/night data set for autonomous and assisted driving," *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 3, pp. 934–948, 2018.
- [10] A. Patil, S. Malla, H. Gang, and Y.-T. Chen, "The h3d dataset for full-surround 3d multi-object detection and tracking in crowded urban scenes," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 9552–9557.
- [11] J. Xue, J. Fang, T. Li, B. Zhang, P. Zhang, Z. Ye, and J. Dou, "Blvd: Building a large-scale 5d semantics benchmark for autonomous driving," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 6685–6691.
- [12] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 year, 1000 km: The oxford robotcar dataset," *The International Journal of Robotics Research*, vol. 36, no. 1, pp. 3–15, 2017.
- [13] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 9297–9307.
- [14] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning*. PMLR, 2017, pp. 1–16.
- [15] A. Sanders, *An introduction to Unreal engine 4*. AK Peters/CRC Press, 2016.
- [16] X. Yue, B. Wu, S. A. Seshia, K. Keutzer, and A. L. Sangiovanni-Vincentelli, "A lidar point cloud generator: from a virtual world to autonomous driving," in *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval*, 2018, pp. 458–464.
- [17] B. Wu, A. Wan, X. Yue, and K. Keutzer, "Squeezeseg: Convolutional neural nets with recurrent crf for real-time road-object segmentation from 3d lidar point cloud," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1887–1893.
- [18] B. Wu, X. Zhou, S. Zhao, X. Yue, and K. 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)*. IEEE, 2019, pp. 4376–4382.
- [19] B. Hurl, K. Czarniecki, and S. Waslander, "Precise synthetic image and lidar (presil) dataset for autonomous vehicle perception," in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 2522–2529.
- [20] B. Sun, J. Feng, and K. Saenko, "Correlation alignment for unsupervised domain adaptation," in *Domain Adaptation in Computer Vision Applications*. Springer, 2017, pp. 153–171.
- [21] P. Morerio, J. Cavazza, and V. Murino, "Minimal-entropy correlation alignment for unsupervised deep domain adaptation," *arXiv preprint arXiv:1711.10288*, 2017.
- [22] J. Fang, D. Zhou, F. Yan, T. Zhao, F. Zhang, Y. Ma, L. Wang, and R. Yang, "Augmented lidar simulator for autonomous driving," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1931–1938, 2020.
- [23] S. Manivasagam, S. Wang, K. Wong, W. Zeng, M. Sazanovich, S. Tan, B. Yang, W.-C. Ma, and R. Urtasun, "Lidarsim: Realistic lidar simulation by leveraging the real world," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [24] F. Langer, A. Milioto, A. Haag, J. Behley, and C. Stachniss, "Domain transfer for semantic segmentation of lidar data using deep neural networks," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 8263–8270.
- [25] P. J. Besl and N. D. McKay, "Method for registration of 3-d shapes," in *Sensor fusion IV: control paradigms and data structures*, vol. 1611. Spie, 1992, pp. 586–606.
- [26] H. Pfister, M. Zwicker, J. Van Baar, and M. Gross, "Surfels: Surface elements as rendering primitives," in *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 2000, pp. 335–342.
- [27] T. Möller and B. Trumbore, "Fast, minimum storage ray/triangle intersection," in *ACM SIGGRAPH 2005 Courses*, 2005, pp. 7–es.
- [28] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [29] Q.-Y. Zhou, J. Park, and V. Koltun, "Open3d: A modern library for 3d data processing," *arXiv preprint arXiv:1801.09847*, 2018.
- [30] G. Borgefors, "Distance transformations in arbitrary dimensions," *Computer vision, graphics, and image processing*, vol. 27, no. 3, pp. 321–345, 1984.
- [31] —, "Distance transformations in digital images," *Computer vision, graphics, and image processing*, vol. 34, no. 3, pp. 344–371, 1986.
- [32] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *CVPR*, 2018.
- [33] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black, "Smpl: A skinned multi-person linear model," *ACM transactions on graphics (TOG)*, vol. 34, no. 6, pp. 1–16, 2015.
- [34] O. D. Team, "Openpcdet: An open-source toolbox for 3d object detection from point clouds," <https://github.com/open-mmlab/OpenPCDet>, 2020.
- [35] T. Yin, X. Zhou, and P. Krahenbuhl, "Center-based 3d object detection and tracking," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 11 784–11 793.