

A LiDAR Point Cloud Generator: from a Virtual World to Autonomous Driving

Xiangyu Yue, Bichen Wu, Sanjit A. Seshia, Kurt Keutzer and Alberto L. Sangiovanni-Vincentelli

EECS, University of California, Berkeley

{xyyue,bichen,sseshia,keutzer,alberto}@eecs.berkeley.edu

ABSTRACT

3D LiDAR scanners are playing an increasingly important role in autonomous driving as they can generate depth information of the environment. However, creating large 3D LiDAR point cloud datasets with point-level labels requires a significant amount of manual annotation. This jeopardizes the efficient development of supervised deep learning algorithms which are often data-hungry. We present a framework to rapidly create point clouds with accurate point-level labels from a computer game. To our best knowledge, this is the first publication on LiDAR point cloud simulation framework for autonomous driving. The framework supports data collection from both auto-driving scenes and user-configured scenes. Point clouds from auto-driving scenes can be used as training data for deep learning algorithms, while point clouds from user-configured scenes can be used to systematically test the vulnerability of a neural network, and use the falsifying examples to make the neural network more robust through retraining. In addition, the scene images can be captured simultaneously in order for sensor fusion tasks, with a method proposed to do automatic registration between the point clouds and captured scene images. We show a significant improvement in accuracy (+9%) in point cloud segmentation by augmenting the training dataset with the generated synthesized data. Our experiments also show by testing and retraining the network using point clouds from user-configured scenes, the weakness/blind spots of the neural network can be fixed.

CCS CONCEPTS

• **Computing methodologies** → **Simulation environments**; • **Information systems** → *Extraction, transformation and loading*; • **Software and its engineering** → *Software infrastructure*; *Simulator / interpreter*; Software verification;

KEYWORDS

LiDAR Point Cloud, Simulation Environment, Autonomous Driving, Neural Network Analysis, Neural Network Retraining

ACM Reference Format:

Xiangyu Yue, Bichen Wu, Sanjit A. Seshia, Kurt Keutzer and Alberto L. Sangiovanni-Vincentelli. 2018. A LiDAR Point Cloud Generator: from a

Virtual World to Autonomous Driving. In *ICMR '18: 2018 International Conference on Multimedia Retrieval, June 11–14, 2018, Yokohama, Japan*, Jennifer B. Sartor, Theo D'Hondt, and Wolfgang De Meuter (Eds.). ACM, New York, NY, USA, Article 4, 7 pages. <https://doi.org/10.1145/3206025.3206080>

1 INTRODUCTION

Autonomous driving requires accurate and reliable perception of the environment. As 3D data has some superiority in information representation over 2D data [9, 12, 27, 28], 3D LiDAR (Light Detection And Ranging) sensors are playing an increasingly important role of all the environment sensors for autonomous driving. On the one hand, 3D LiDAR sensors can provide direct distance measurements that allow detection of all kinds of obstacles, and their resolution and field of view exceed radar and ultrasonic sensors [17]. On the other hand, LiDAR sensors are robust under a variety of conditions: day or night, with or without glare and shadows [26]. While LiDAR point clouds contain accurate depth measurement of the environment, navigation of autonomous vehicles also relies on correct understanding of the semantics of the environment. Most of the LiDAR-based perception tasks, such as semantic segmentation [3, 4, 26] and drivable area detection [8, 16], require significant amount of point-level labels for training and/or validation. Such annotation, however, is usually very expensive.

To facilitate the manual annotation process, much work has been done on interactive annotation. Annotation methods have been proposed for labeling 3D point clouds of both indoor scenes [21] and outdoor driving scenes [10]. These methods utilize little computer assistance during the annotation process and thus need a significant amount of human effort. In [15, 23], approaches have been proposed to enhance the human-machine interaction to improve annotation efficiency. In [20, 25], annotation suggestions for indoor RGBD scenes are proposed by the system that are interactively corrected or refined by the user. In order to provide faster interactive labeling rates, Boyko *et al.* [2] proposed a group annotation approach for labeling objects in 3D LiDAR scans. Active learning has also been introduced in the annotation process to train a classifier with fewer interactions [14, 22], yet it requires users to interact with examples one-by-one. Other frameworks further take into account the risk of mislabeling and cost of annotation. In [24], Welinder *et al.* proposed a model of the labeling process and dynamically chooses which images will be labeled next in order to achieve a desired level of confidence.

Recently, video games have been used for creating large-scale ground truth data for training purposes. In [19], a video game is used to generate ground truth semantic segmentation for the synthesized in-game images. However, human effort is still required in the annotation process. In [13], the same game engine is used to generate ground truth 2D bounding boxes of objects in the images.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICMR '18, June 11–14, 2018, Yokohama, Japan

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5046-4/18/06...\$15.00

<https://doi.org/10.1145/3206025.3206080>




Figure 1: Sample data extracted from an in-game scene. (a): Image of the scene; (b): Extracted point cloud from the same scene; (c): Point cloud of car mapped to image after registration (Blue dots) matches car in image.

Richter *et al.* [18] further extended the work of [19] so that various ground truth information (e.g. semantic segmentation, semantic instance segmentation, and optical flow) can be extracted from the game engine. In addition, many driving simulation environments [1, 5, 11] have been built in order to obtain various kinds of labeled data for autonomous driving purposes. Many of these work [5, 13, 18, 19] show the effectiveness of synthetic data in image-based learning tasks by showing improved performance after training with additional synthetic data. However, little work has been done on extracting annotated 3D LiDAR point clouds from simulators, not to mention showing the efficacy of the synthetic point clouds during the training process of neural networks.

Note that even if we could provide large amounts of training data, it is still almost impossible for any algorithms to achieve 100% accuracy. For Cyber-Physical Systems used for safety critical purposes, such as autonomous driving, verifying neural networks is of extreme importance [6]. In [7], a framework is proposed to systematically analyze Convolutional Neural Networks (CNNs) used in objection detection in autonomous driving systems. However, the framework only takes into account cars from direct front/back view and thus has a very limited modification space. In addition, each background image needs to be manually annotated, making it expensive to generate a dataset with large diversity. To the best of our knowledge, no similar work has been done on LiDAR point clouds. In this paper, we propose an extraction-annotation-CNN testing framework based on a popular video game. The main contributions of this paper are as follows:

- This is the first publication on LiDAR point cloud simulation framework for autonomous driving.
- The framework can automatically extract point-cloud data with ground truth labels together with the corresponding image frame of the in-game scene, as shown in Figure 1.
- The framework can do automatic registration between collected point clouds and images which can then be used together for sensor fusion tasks, e.g. inferring depth information from RGB images.
- Users can construct specified scenarios in the framework interactively and the collected data (point clouds and images) can then be used to systematically test, analyze and improve LiDAR-based and/or image-based learning algorithms for autonomous driving.

We conduct experiments on a Convolutional Neural Network (CNN)-based model for 3D LiDAR point cloud segmentation using the data collected from the proposed framework. The experiments show 1) significantly improved performance on KITTI dataset [10] after retraining with additional synthetic LiDAR point clouds, and 2) efficacy of using the data collected from user-configured scenes in the framework to test, analyze and improve the performance of the neural network. The performance improvements come from the fact that the data collected in the rich virtual world contains a lot of information that the neural network failed to learn from the limited amount of original training samples.

2 TECHNICAL APPROACH

2.1 In-Game Simulation Setup and Method for Data Collection

We choose to utilize the rich virtual world in Grand Theft Auto V (GTA-V), a popular video game, to obtain simulated point clouds as well as captured in-game images with high fidelity¹. Our framework is based on DeepGTAV², which uses Script Hook V³ as a plugin.

In order to simulate realistic driving scenes, an ego car is used in the game with a virtual LiDAR scanner mounted atop, and it is set to drive autonomously in the virtual world with the AI interface provided in Script Hook V. While the car drives on a street, the system collects LiDAR point clouds and captures the game screen, simultaneously. We place the virtual LiDAR scanner and the game camera at the same position in the virtual 3D space. This set-up offers two advantages: 1) a sanity check can be easily done on the collected data, since point clouds and corresponding images must be consistent; 2) registration between the game camera and the virtual LiDAR scanner can be done automatically, and then collected point clouds and scene images can be combined together as training dataset for neural networks for sensor fusion tasks. Details of the proposed registration method will be described in Section 2.2.

Ray casting is used to simulate each laser ray emitted by the virtual LiDAR scanner. The ray casting API takes as input the 3D coordinates of the starting and ending point of the ray, and returns the 3D coordinates of the first point the ray hits. This point is used, with another series of API function calls, to calculate, among other

¹The publisher of GTA-V allows non-commercial use of footage of gameplay [19].

²<https://github.com/aitorzip/DeepGTAV>

³<http://www.dev-c.com/gtav/scripthookv/>




Figure 2: Sample configurable parameters of the virtual LiDAR. (a) shows front view of the virtual LiDAR: black dotted line is the horizontal line, α is the vertical field of view (FOV), θ is the vertical resolution, σ is the pitch angle; (b) shows top view of the virtual LiDAR, β is the horizontal FOV, and ϕ is the horizontal resolution.

data, the distance of the point, the category and instance ID of the object hit by the ray, thus allowing automatic annotation on the collected data.

In our framework, users can provide configurations of the LiDAR scanner including vertical field of view (FOV), vertical resolution, horizontal FOV, horizontal resolution, pitch angle, maximum range of laser rays, and scanning frequency. Some of the configurable parameters are shown in Figure 2.

2.2 Automatic Registration Method

The goal of the registration process is to find the corresponding pixel in the image for each LiDAR point. In our framework, the registration process can be done automatically by the system based on the parameters of the camera and LiDAR scanner. In addition, the centers of the camera and LiDAR scanner are set to the same position in the virtual world, making the registration projection similar to the camera perspective projection model, as shown in Figure 3.

The problem is formulated as follows: for a certain laser ray with *azimuth* angle ϕ and *zenith* angle θ , calculate the index (i, j) of the corresponding pixel on image. \mathcal{F}_c , \mathcal{F}_o , P , P' and P_{far} are 3D coordinates of a) center of camera/LiDAR scanner, b) center of camera near clipping plane, c) point first hit by the virtual laser ray (in red), d) pixel on image corresponding to P , and e) a point far away in the laser direction (required as an input argument to the ray casting API), respectively. m and n are the width and height of the near clipping plane. γ is 1/2 vertical FOV of camera while ψ is 1/2 vertical FOV of the LiDAR scanner. Note that LiDAR scanner FOV is usually smaller than camera FOV, since there is usually no object in the top part of the image, and thus emitting laser to open space is not necessary. After a series of 3D geometry calculation, we can get:

$$\begin{aligned} i &= \frac{R_m}{m} \cdot (f \cdot \tan \gamma \cdot \frac{m}{n} - \frac{f}{\cos \theta} \cdot \tan \phi), \\ j &= \frac{R_n}{n} \cdot (f \cdot \tan \gamma + f \cdot \tan \theta), \end{aligned} \quad (1)$$

where $f = \|\vec{\mathcal{F}_c \mathcal{F}_o}\|$, and (R_m, R_n) is the pixel resolution of the image/near clipping plane.




Figure 3: Projection for Registration. \mathcal{F}_o is the center of the near clipping plane of the camera; \mathcal{F}_c is the center of the camera and of the LiDAR scanner; the red line is the laser ray and P is the point hit by the ray; the calibrated on-image point has pixel index (i, j) and 3D coordinates P' ; γ is 1/2 the camera vertical FOV and ψ is 1/2 the LiDAR vertical FOV; ϕ and θ are the *azimuth* and *zenith* angles of the laser ray.

Further, as an input argument to the ray casting API, the 3D coordinates of P_{far} are also required. Using similar 3D geometry calculations, we obtain:

$$\begin{aligned} P' &= \mathcal{F}_c + f \cdot \vec{x}_c - \frac{f}{\cos \theta} \cdot \tan \phi \cdot \vec{y}_c - f \cdot \tan \theta \cdot \vec{z}_c, \\ P_{far} &= \mathcal{F}_c + k \cdot (P' - \mathcal{F}_c), \end{aligned} \quad (2)$$

where k is a large coefficient, and $\vec{x}_c, \vec{y}_c, \vec{z}_c$ are unit vectors of the camera axis in the world coordinate system.

An example of the registration result is shown in Figure 1. After simulation, both image and point cloud of the specified in-game scene are collected by the framework (Figure 1 (a, b)). Then with the proposed registration method, we map all the points with category "Car" to the corresponding image. As shown in Figure 1 (c), the mapped car point cloud (blue dots) matches the car in the image fairly accurately.

2.3 Configurable In-game Scene

Besides the auto-driving mode for large-scale data collection, our framework offers a configurable mode, where the user can configure desired in-game scenes and collect data from them. One advantage of configurable scenes is generating training data of driving scenes that are dangerous or rare in real world. Another advantage is that we can systematically sample the modification space (e.g. number of cars, position and orientation of a car) of an in-game scene. The data can then be used to test a neural network, expose its vulnerabilities and improve its performance through retraining. Our framework offers a large modification space of the in-game scene. As shown in Figure 4, the user can specify and change 8 dimensions of in-game scene: car model, car location, car orientation, number of cars, scene background, color of car, weather, and time of day. The first 5 dimensions affect both LiDAR point cloud and scene image, while the last three dimensions affect only the scene image. An example of sampling is shown in Figure 5, where the scenes are only sampled from the spatial dimensions (X, Y) with only one car in each scene. X and Y are the location offset of the car relative to the camera/LiDAR location in the left-right and forward-backward



Figure 4: Modification dimensions of the framework with image in center showing the reference scene.

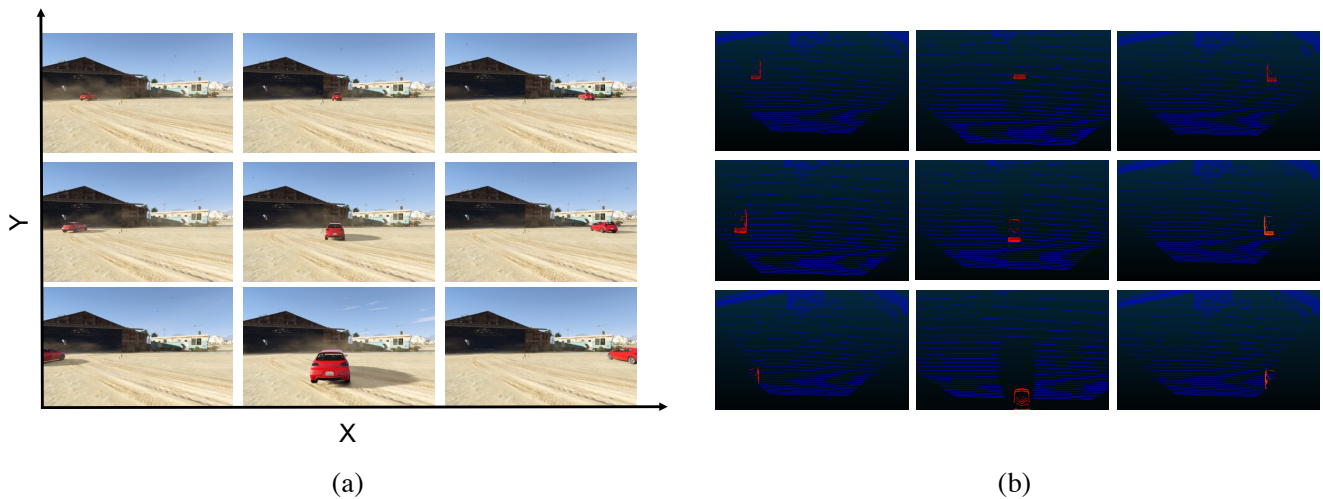


Figure 5: Scenes with one car sampled from spatial dimensions and corresponding point cloud. (a) shows the scene image while changing the location of the car on X (left-right) and Y (forward-backward) directions; (b) shows point clouds (red for car and blue for background) of scenes in (a).

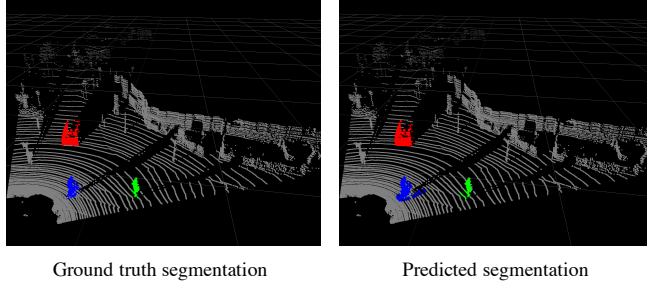


Figure 6: LiDAR point cloud segmentation

directions. Figure 5 (b) shows collected point cloud of the samples shown in Figure 5 (a). The red points represent car points while the blue points represent the scene background. The collected point clouds match the scenes well thus allowing the use of the data to test neural networks systematically.

3 EXPERIMENTS AND RESULTS

We performed experiments to show the efficacy of our data synthesis framework: 1) Data collected by the framework can be used in the training phase and help improve the validation accuracy; 2) Collected data can be used to systematically test a neural network and improve its performance via retraining.

3.1 Evaluation Metrics

Our experiments are performed on the task of LiDAR point cloud segmentation; specifically, given a point cloud detected by a LiDAR sensor, we wish to perform point-wise classification, as shown in Figure 6. This task is an essential step for autonomous vehicles to perceive and understand the environment, and navigate accordingly.

To evaluate the accuracy of the point cloud segmentation algorithm, we compute *Intersection-over-Union* (IoU), *Precision* and *Recall* as:

$$IoU_c = \frac{|\mathcal{P}_c \cap \mathcal{G}_c|}{|\mathcal{P}_c \cup \mathcal{G}_c|}, Pr_c = \frac{|\mathcal{P}_c \cap \mathcal{G}_c|}{|\mathcal{P}_c|}, Recall_c = \frac{|\mathcal{P}_c \cap \mathcal{G}_c|}{|\mathcal{G}_c|}.$$

Here, \mathcal{P}_c denotes the set of points that our model predicted to be of class- c , \mathcal{G}_c denotes the ground-truth set of points belonging to class- c , and $|\cdot|$ denotes the cardinality of a set. *Precision* and *Recall* measures accuracy with regard to false positives and false negatives, respectively; while *IoU* takes both into account. For this, *IoU* is used as the primary accuracy metric in our experiments.

3.2 Experimental Setup

Our analysis is based on SqueezeSeg [26], a convolutional neural network based model for point cloud segmentation. To collect the real-world dataset, we used LiDAR point cloud data from the KITTI dataset and converted its 3D bounding box labels to point-wise labels. Since KITTI dataset only provides reliable 3D bounding boxes for front-view LiDAR point clouds, we limit the horizontal field of view (FOV) to the forward-facing 90°. This way, we obtained 10,848 LiDAR scans with manual labels. We used 8,057 scans for training and 2,791 scans for validation. Each point in a KITTI LiDAR scan has 3 cartesian coordinates (x, y, z) and an intensity

Table 1: Segmentation Performance Comparison on the Car Category. Only data used in the first row has Intensity channel. All numbers are in percentage.

	Precision	Recall	IoU
KITTI w/ Intensity	66.7	95.4	64.6
KITTI w/o Intensity	58.9	95.0	57.1
GTA-V only	30.4	86.6	29.0
KITTI w/o Intensity + GTA-V	69.6	92.8	66.0

value, which measures the amplitude of the laser signal returned. Although the intensity measurement as an extra input feature is beneficial to improve the segmentation accuracy, simulating the intensity measurement is very difficult and not supported in our current framework. Therefore we excluded intensity as an input feature to the neural network for GTA-V synthetic LiDAR data. We use NVIDIA TITAN X GPUs for the experiments during both the training and validation phases.

3.3 Experimental Results

Synthetic Data for Training: For the first set of experiments, we used our data synthesis framework to generate 8,585 LiDAR point cloud scans in autonomous-driving scenes. The generated data contain (x, y, z) measurements but do not contain intensity. The horizontal FOV of the collected point clouds are set to be 90° to match the setting of KITTI point clouds described in Section 3.2.

To quantify the effect of training the model with synthetic data, we first trained two models on the KITTI training set with intensity included and excluded, and validated on the KITTI validation set. The performance is shown in the first 2 rows of Table 1 as the baseline. The model with intensity achieved better result. Then we trained another model with only GTA-V synthetic data. As shown in the third row of Table 1, the performance drops a lot. This is mostly because the distributions of the synthetic dataset and KITTI dataset are quite different. Therefore, through training purely on synthetic dataset, it is hard for the neural network to learn all the required details for the KITTI dataset, which might be missing or insufficient in the synthetic training dataset. Finally, we combined the KITTI data and GTA-V data together as the training set and train another model. As shown in the last row of Table 1, the performance gets improved significantly, almost 9% better than the accuracy achieved only using real-world data. Despite the loss of the intensity channel, the GTA+KITTI dataset gives better accuracy (66.0%) than if intensity is included (64.6%). This demonstrates the efficacy of the synthetic data extracted in our framework.

Neural Network Testing and Robustness Enhancement: For the second set of experiments, we first used our framework to systematically test SqueezeSeg. As an illustrative experiment, we only performed sampling in the car location X-Y dimensions as in Figure 5, rather than the whole modification space. 555 scenes were sampled to test SqueezeSeg, with the IoU results shown in Figure 7. The blue and green dots show the car locations resulting in low IoU. Most of the "blind spot" are locations far from the LiDAR scanner, but there are also closer locations that result in low IoU scores. Close locations with low IoUs are dangerous in autonomous




Figure 7: IoU scatter with the change of car location




Figure 8: mIoU map of the validation set before retraining

driving, since they can mislead the decision-making system of the autonomous vehicles and result in immediate accident.

Further experiments are then done to show the efficacy of using generated synthetic data to possibly improve performance and robustness of the network in the modification space. We synthesized totally 2,250 LiDAR point cloud scans in 15 different scene backgrounds. In each scene background, only one car is placed with the same orientation as the camera view. We obtained 150 point cloud scans in each scene background by changing the position of the car (X, Y) in the sampled space: $S = \{(x, y) \mid x \in \{-5, \dots, 4\}, y \in \{5, \dots, 19\}\}$, where X, Y are respectively the left-right and forward-backward offset relative to the position of the camera. For each scene background, the position and orientation of the camera were fixed.

We split the collected point cloud scans based on the scene background. 1200 point cloud scans in the first 8 backgrounds are used as validation set \mathcal{V} , and the rest 1050 scans from the other 7 backgrounds, which we call retraining set \mathcal{R} , are used for retraining purpose. First, we train a neural network with purely KITTI data and do evaluation on the synthetic 1200 scans in the validation set \mathcal{V} . We define *mean IoU (mIoU)* for each point in the 15×10 X-Y modification space as averaging IoUs over all the scene backgrounds in \mathcal{V} :

$$mIoU(i, j) = \frac{1}{n} \sum_{k=1}^n IoU(i, j, k),$$




Figure 9: mIoU map of the validation set after selection with mIoU less than 0.65 set to 0. All the point clouds in the retraining set \mathcal{R} corresponding to the blue positions in the new mIoU map will be added to the original training set.

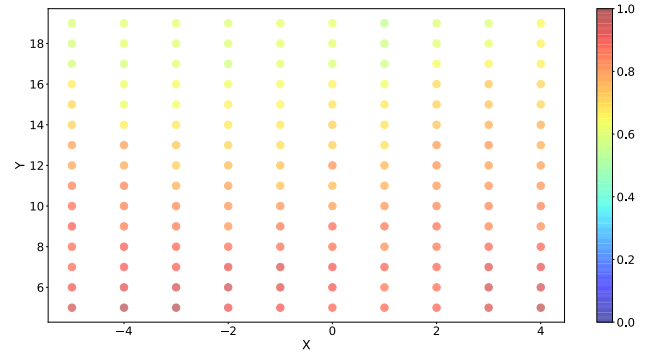


Figure 10: mIoU map of the validation set after retraining

where n is the number of scene backgrounds ($n = 8$ in this experiment), (i, j) is in $\{(i, j) \mid i \in [-5, 4], j \in [5, 19], i, j \in \mathbb{Z}\}$ and $IoU(i, j, k)$ refers to the IoU of the point cloud scan sampled at (i, j) in the X-Y modification with the k_{th} scene background.

The mIoU map of the validation set is computed, as shown in Figure 8. We can see that the pre-trained network performs poorly on positions that are far away, at the boundary of the FOV. But more surprisingly, we also observed that on some positions that are fairly close to the ego-vehicle, e.g. $(-3, 5)$, the mIoU scores are also very low. Detection errors at such near distance can be very dangerous.

Based on the mIoU map, we choose positions with an mIoU smaller than a threshold to form a retraining set, as shown in Figure 9. Then all the point clouds in the retraining set \mathcal{R} with a selected position are added to the original training set. After the retraining process, we re-evaluate the validation set \mathcal{V} , with the new mIoU map shown in Figure 10. As the figure shows, at almost all the close-to-center positions originally with low mIoU, the neural network performs much better than before the retraining. In order to visualize the performance improvements better, we plot the mIoU improvement after the retraining process for each position. The mIoU improvements are sorted and plotted in Figure 11. We see that after retraining, performance on point clouds at most of the




Figure 11: mIoU improvements in ascending order for all 150 positions

positions gets much better, with slightly degraded performance at only a small fraction of positions. Meanwhile, the performance on KITTI dataset remained almost the same with IoU changing from 60.8% to 60.6%. These experiments show the efficacy of using synthetic data from user-configured scenes of the proposed framework to test, analyze and improve the performance of neural networks through retraining.

4 CONCLUSIONS AND FUTURE WORK

In this paper, we proposed a framework that synthesizes annotated LiDAR point clouds from a virtual world in a game, with a method to automatically calibrate the point cloud and scene image. Our framework can be used to: 1) obtain a large amount of annotated point cloud data, which can then be used to help neural network training; 2) systematically test, analyze and improve performance of neural networks for tasks such as point cloud segmentation. Experiments show that for a point cloud segmentation task, synthesized data help improve the validation accuracy (IoU) by 9% on a real-world benchmark. Furthermore, the systematical sampling and testing framework can help us to identify potential weakness/blind spots of our neural network model and fix them. Our future works will focus on two directions: 1) Experiments show that the intensity channel in LiDAR point clouds is crucial for LiDAR-based perception tasks, but current framework does not support intensity synthesis. We plan to explore techniques to simulate realistic LiDAR intensity. 2) The effectiveness of using synthetic data to train perception algorithms is usually limited by domain shift. We plan to explore domain adaptation methods to diminish the gap between virtual world and real world.

ACKNOWLEDGEMENT

This work was partially supported by NSF, Award 1645964, together with Berkeley Deep Drive. We thank Kostadin Ilov for providing system assistance.

REFERENCES

- [1] Daniel Biedermann, Matthias Ochs, and Rudolf Mester. 2016. Evaluating visual ADAS components on the CONGRATS dataset. In *IEEE Intelligent Vehicles Symposium*. 986–991.
- [2] Aleksey Boyko and Thomas Funkhouser. 2014. Cheaper by the Dozen: Group Annotation of 3D Data. In *ACM Symposium on User Interface Software and Technology*. 33–42.
- [3] Ayush Dewan, Gabriel L Oliveira, and Wolfram Burgard. 2017. Deep Semantic Classification for 3D LiDAR Data. *arXiv preprint arXiv:1706.08355* (2017).
- [4] David Dohan, Brian Matejek, and Thomas Funkhouser. 2015. Learning hierarchical semantic segmentations of lidar data. In *International Conference on 3D Vision*. 273–281.
- [5] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. 2017. CARLA: An Open Urban Driving Simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*. 1–16.
- [6] Tommaso Dreossi, Alexandre Donzé, and Sanjit A Seshia. 2017. Compositional Falsification of Cyber-Physical Systems with Machine Learning Components. In *NASA Formal Methods Symposium*. 357–372.
- [7] Tommaso Dreossi, Shromona Ghosh, Alberto Sangiovanni-Vincentelli, and Sanjit A. Seshia. 2017. Systematic testing of convolutional neural networks for autonomous driving. *arXiv preprint arXiv:1708.03309* (2017).
- [8] R. Fernandes, C. Prenebida, P. Peixoto, D. Wolf, and U. Nunes. 2014. Road Detection Using High Resolution LIDAR. In *IEEE Vehicle Power and Propulsion Conference*. 1–6.
- [9] Yue Gao, Meng Wang, Dacheng Tao, Rongrong Ji, and Qionghai Dai. 2012. 3-D object retrieval and recognition with hypergraph analysis. *IEEE Transactions on Image Processing* 21, 9 (2012), 4290–4303.
- [10] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. 2013. Vision meets Robotics: The KITTI Dataset. *International Journal of Robotics Research* 32, 11 (2013), 1231–1237.
- [11] Vladimir Haltakov, Christian Unger, and Slobodan Ilic. 2013. Framework for generation of synthetic ground truth data for driver assistance applications. In *German Conference on Pattern Recognition*. 323–332.
- [12] Rongrong Ji, Ling-Yu Duan, Jie Chen, Tiejun Huang, and Wen Gao. 2014. Mining compact bag-of-patterns for low bit rate mobile visual search. *IEEE Transactions on Image Processing* 23, 7 (2014), 3099–3113.
- [13] Matthew Johnson-Roberson, Charles Barto, Rounak Mehta, Sharath Nittur Sridhar, Karl Rosaen, and Ram Vasudevan. 2017. Driving in the matrix: Can virtual worlds replace human-generated annotations for real world tasks?. In *IEEE International Conference on Robotics and Automation*. 746–753.
- [14] Ashish Kapoor, Kristen Grauman, Raquel Urtasun, and Trevor Darrell. 2007. Active learning with gaussian processes for object categorization. In *IEEE International Conference on Computer Vision*. 1–8.
- [15] Regis Kopper, Felipe Bacim, and Doug A Bowman. 2011. Rapid and accurate 3D selection by progressive refinement. In *IEEE Symposium on 3D User Interfaces*. 67–74.
- [16] Ziyi Liu, Siyu Yu, Xiao Wang, and Nanning Zheng. 2017. Detecting Drivable Area for Self-driving Cars: An Unsupervised Approach. *arXiv preprint arXiv:1705.00451* (2017).
- [17] Frank Moosmann, Oliver Pink, and Christoph Stiller. 2009. Segmentation of 3D lidar data in non-flat urban environments using a local convexity criterion. In *IEEE Intelligent Vehicles Symposium*. 215–220.
- [18] Stephan R Richter, Zeeshan Hayder, and Vladlen Koltun. 2017. Playing for benchmarks. In *IEEE International Conference on Computer Vision*. 2232–2241.
- [19] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. 2016. Playing for data: Ground truth from computer games. In *European Conference on Computer Vision*. 102–118.
- [20] Tianjia Shao, Weiwei Xu, Kun Zhou, Jingdong Wang, Dongping Li, and Baining Guo. 2012. An Interactive Approach to Semantic Modeling of Indoor Scenes with an RGBD Camera. *ACM Transactions on Graphics* 31, 6 (2012), 136:1–136:11.
- [21] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. 2012. Indoor segmentation and support inference from rgb-d images. In *European Conference on Computer Vision*. 746–760.
- [22] Andrew Top, Ghassan Hamarneh, and Rafeef Abugharbieh. 2011. Active learning for interactive 3D image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. 603–610.
- [23] Manuel Veit and Antonio Capobianco. 2014. GoThenTag: A 3-D point cloud annotation technique. In *IEEE Symposium on 3D User Interfaces*. 193–194.
- [24] Peter Welinder and Pietro Perona. 2010. Online crowdsourcing: rating annotators and obtaining cost-effective labels. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. 25–32.
- [25] Yu-Shiang Wong, Hung-Kuo Chu, and Niloy J Mitra. 2014. Smartannotator: An interactive tool for annotating RGBD indoor images. *arXiv preprint arXiv:1403.5718* (2014).
- [26] Bichen Wu, Alvin Wan, Xiangyu Yue, and Kurt Keutzer. 2018. SqueezeSeg: Convolutional Neural Nets with Recurrent CRF for Real-Time Road-Object Segmentation from 3D LiDAR Point Cloud. In *IEEE International Conference on Robotics and Automation*.
- [27] Sicheng Zhao, Lujun Chen, Hongxun Yao, Yanhao Zhang, and Xiaoshuai Sun. 2015. Strategy for dynamic 3D depth data matching towards robust action retrieval. *Neurocomputing* 151 (2015), 533–543.
- [28] Sicheng Zhao, Hongxun Yao, Yanhao Zhang, Yasi Wang, and Shaohui Liu. 2015. View-based 3D object retrieval via multi-modal graph learning. *Signal Processing* 112 (2015), 110–118.