Supplementary Material

Towards Generating Realistic Autonomous Driving 3D Semantic Training Data

This supplementary material provides further information regarding the method presented in the main article. Sec. 1 provides further details about the network architectures used as the variational auto-encoder (VAE) and as the denoising diffusion probabilistic model (DDPM). Sec. 2 provides further details regarding the binary cross-entropy (BCE) and dice losses used in the VAE training. Sec. 3 shows class IoU from the experiments done in Sec. 4.2 in the main article. Sec. 4 shows further qualitative results from conditional and unconditional generation from our method.

1 Architectures

This section presents the network architectures used in our approach. The VAE is trained to encode the scene $\mathcal{P}$ in the latent $\mathcal{Z}$, which is then padded to a dense grid $\boldsymbol{Z}$ such that the VAE decoder can reconstruct $\hat{\mathcal{P}} \approx \mathcal{P}$ from it. The VAE architecture layers employ sparse operations to enable the processing of large-scale data and avoid exponential memory growth. The DDPM is trained with the VAE learned latent $\boldsymbol{Z}$ by receiving as input the noisy latent $\boldsymbol{Z}^t$ at step t, predicting $\mathbf{v}^t$ following the $\mathbf{v}$-parameterization formulation [9], where $\mathbf{v}$ is parameterized in terms of the sampled noise $\boldsymbol{\epsilon}$ and the uncorrupted data $\boldsymbol{Z}^0$. Given the dense latent $\boldsymbol{Z}$, the DDPM uses dense convolutional operations. Both model architectures are depicted in Fig. 9.

2 Pruning Losses

This section details the $\mathcal{L}_{\text{bce}}$ and $\mathcal{L}_{\text{dice}}$ losses used to learn the pruning mask at each VAE upsampling layer. As mentioned in Sec. 3.1 in the main article, the VAE is trained with a cross-entropy loss as the semantic loss and the KL divergence as the latent loss. To learn the pruning masks, we follow recent mask-based segmentation approaches [1], [2], [6], [7], supervising the mask prediction with the $\mathcal{L}_{\text{bce}}$ and $\mathcal{L}_{\text{dice}}$ losses. The $\mathcal{L}_{\text{bce}}$ aims at predicting whether an individual voxel is occupied or not at a given upsampling layer l and is computed as:

$$\mathcal{L}^l_{\text{bce}} = -(\boldsymbol{m}_l \log(\hat{\boldsymbol{m}}_l) + (1 - \boldsymbol{m}_l) \log(1 - \hat{\boldsymbol{m}}_l)). \tag{11}$$

where $\boldsymbol{m}_l$ is the target pruning mask and $\hat{\boldsymbol{m}}_l$ is the predicted mask. At the same time, the dice loss aims at predicting the whole scene layout by computing the dice coefficient, which approximates the IoU computation and maximizing this computation as:

$$\text{dice}^l = \frac{2\,|\hat{\boldsymbol{m}}_l \cap \boldsymbol{m}_l|}{|\hat{\boldsymbol{m}}_l| + |\boldsymbol{m}_l|}, \tag{12}$$

$$\mathcal{L}^l_{\text{dice}} = 1 - \text{dice}^l, \tag{13}$$

where as in Eq. (11), $\boldsymbol{m}_l$ is the target pruning mask and $\hat{\boldsymbol{m}}_l$ the predicted mask. Therefore, both losses complement each other, predicting both the individual voxel occupancy and the layout of the entire scene.

3 Quantitative Results

In this section we provide further quantitative results regarding the class IoU from the scenes generated by our method and the baselines. In Tab. 11, we report the class IoU from the model trained with real data evaluated on the real validation set and the scenes generated by the different methods. Similar to Tab. 2 from the main article, we report the numbers at two resolutions, $0.2\,\text{m}$, and $0.1\,\text{m}$, since PDD [5] and SemCity [3] can only generate scenes up to $0.2\,\text{m}$ resolution. Similar to what was presented in Tab. 2, at $0.2\,\text{m}$ resolution, all the methods achieve similar performance still our method achieves the best mIoU, as shown in Tab. 2. At $0.1\,\text{m}$ resolution, our method surpasses the method able to generate scenes at this resolution, achieving the best performance in all classes compared to the scene generated by the baseline. Tab. 12 depicts the class IoU results from the experiment presented in Tab. 9 from the main article. As seen, with the curation process described in Sec. 4.4 the model performance improved in almost all the classes, even using less synthetic training data. This suggests that the curation process can improve the model performance when using synthetic scenes as training data by selecting only the most realistic scenes. This curation helps to bridge the gap between real and synthetic data by removing less realistic samples.

4 Qualitative Results

In this section we show further examples of scenes generated by our method. Figs. 10 to 12 shows the unconditionally generated scene examples. As shown, the randomly generated point clouds present different scenarios. This variability agrees with the discussion in Sec. 4.2 where the training with a mixed set of samples from real and synthetic scenes performed better than using only real scenes. The synthetic scenes add variability to the data compared to the real training set since the real scenes were collected sequentially, where consecutive scenes may have few changes. In contrast, by using randomly generated scenes, each individual sample will be different from each other, adding variability to the data, thus improving the performance of the model when trained with this synthetic training set. Regarding the conditioned scene generation, Figs. 13 and 14 show examples of scenes generated conditioned to LiDAR scans from KiTTI-360 [4] and to our own collected data, respectively. In those examples, we can notice that the model is able to generate reasonable scenes according to the scan used as condition. Such conditioning also works for point clouds collected with a LiDAR sensor different from the one used to train the model, as seen in Fig. 14 from the scenes generated conditioned to Ouster LiDAR collected by us.

References

[1] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in *Proc. of the Europ. Conf. on Computer Vision (ECCV)*, 2020. 1

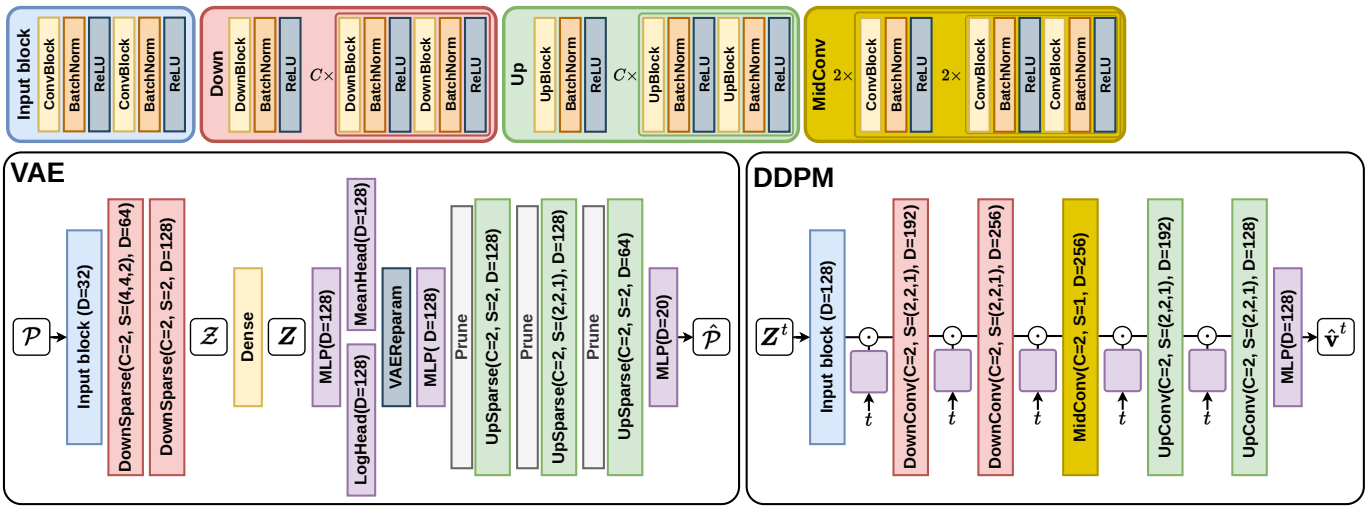

Fig. 9. VAE and DDPM model architectures. The VAE receives the voxelized point cloud $\mathcal{P}$, encode it to the latent $\mathcal{Z}$ which is densified to $\boldsymbol{Z}$ and decoded to $\hat{\mathcal{P}}$. The DDPM receives the noisy latent $\boldsymbol{Z}^t$ at step t and predicts $\hat{\mathbf{v}}^t$ following the $\mathbf{v}$-parameterization formulation [9].

		IoU ↑												
	Res.	car	truck	OV	road	park.	sidewalk	build.	fence	veg.	trunk	terrain	pole	sign
Val. scenes	0.2	91.10	45.90	31.73	87.76	34.87	58.18	86.33	25.32	86.64	42.79	64.11	54.75	47.72
PDD [5]	0.2	80.24	**11.81**	0.29	81.14	8.99	57.00	79.19	34.82	78.71	7.96	**54.65**	15.67	0.81
SemCity [3]	0.2	80.46	03.72	0.71	85.81	**24.22**	66.61	78.40	**36.32**	**82.85**	21.50	51.95	24.18	1.78
XCube [8]	0.2	47.94	11.62	**13.28**	82.75	18.19	54.70	67.86	36.08	76.91	**25.57**	51.01	**35.47**	**17.74**
Ours	0.2	**88.67**	5.54	03.41	**87.62**	13.95	**68.26**	**82.47**	27.19	80.16	22.94	53.59	34.22	13.71
Val. scenes	0.1	93.77	42.24	49.00	92.16	46.52	69.24	89.28	41.08	88.07	51.88	65.81	61.12	58.36
XCube [8]	0.1	15.42	1.68	**03.33**	71.26	2.84	42.39	45.69	23.69	68.30	22.04	39.39	24.22	9.70
Ours	0.1	**89.53**	**10.55**	2.41	**87.30**	**14.51**	**70.88**	**82.32**	**25.04**	**81.27**	**26.29**	**53.84**	**33.38**	**14.06**

TABLE 11. Class-wise IoU evaluated on real data validation set and synthetic data generated by the different methods with semantic segmentation model trained on real data.

	IoU ↑												
	car	truck	OV	road	park.	sidewalk	build.	fence	veg.	trunk	terrain	pole	sign
Real only	93.77	42.24	49.00	92.16	46.52	69.24	89.28	41.08	88.07	51.88	65.81	**61.12**	58.36
Real + 75% Synth.	93.97	78.95	57.44	**92.55**	45.86	**71.67**	89.59	**42.67**	88.01	**52.03**	65.44	60.44	58.49
Real + 25% Synth.†	**94.60**	**80.55**	**64.48**	92.15	**49.47**	71.01	**89.76**	42.43	**88.14**	50.74	**66.02**	59.95	**58.82**

TABLE 12. Dense scenes class-wise IoU evaluated on the real data validation set comparing the network trained only with the full real training set with the training with the full training set with additional 75% synthetic data generated with our method. OV refers to other-vehicle. † refers to the curated conditional generated scenes

[2] B. Cheng, I. Misra, A. G. Schwing, A. Kirillov, and R. Girdhar, "Masked-attention mask transformer for universal image segmentation," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1

[3] J. Lee, S. Lee, C. Jo, W. Im, J. Seon, and S.-E. Yoon, "Semcity: Semantic scene generation with triplane diffusion," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1, 2

[4] Y. Liao, J. Xie, and A. Geiger, "KITTI-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d," *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 45, no. 3, pp. 3292–3310, 2022. 1, 6

[5] Y. Liu, X. Li, X. Li, L. Qi, C. Li, and M.-H. Yang, "Pyramid diffusion for fine 3d large scene generation," in *Proc. of the Europ. Conf. on Computer Vision (ECCV)*, 2024. 1, 2

[6] R. Marcuzzi, L. Nunes, L. Wiesmann, J. Behley, and C. Stachniss, "Mask-Based Panoptic LiDAR Segmentation for Autonomous Driving," *IEEE Robotics and Automation Letters (RA-L)*, vol. 8, no. 2, pp. 1141–1148, 2023. 1

[7] R. Marcuzzi, L. Nunes, L. Wiesmann, E. Marks, J. Behley, and C. Stachniss, "Mask4D: End-to-End Mask-Based 4D Panoptic Segmentation for LiDAR Sequences," *IEEE Robotics and Automation Letters (RA-L)*, vol. 8, no. 11, pp. 7487–7494, 2023. 1

[8] X. Ren, J. Huang, X. Zeng, K. Museth, S. Fidler, and F. Williams, "Xcube: Large-scale 3d generative modeling using sparse voxel hierarchies," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2

[9] T. Salimans and J. Ho, "Progressive distillation for fast sampling of diffusion models," in *Proc. of the Intl. Conf. on Learning Representations (ICLR)*, 2022. 1, 2

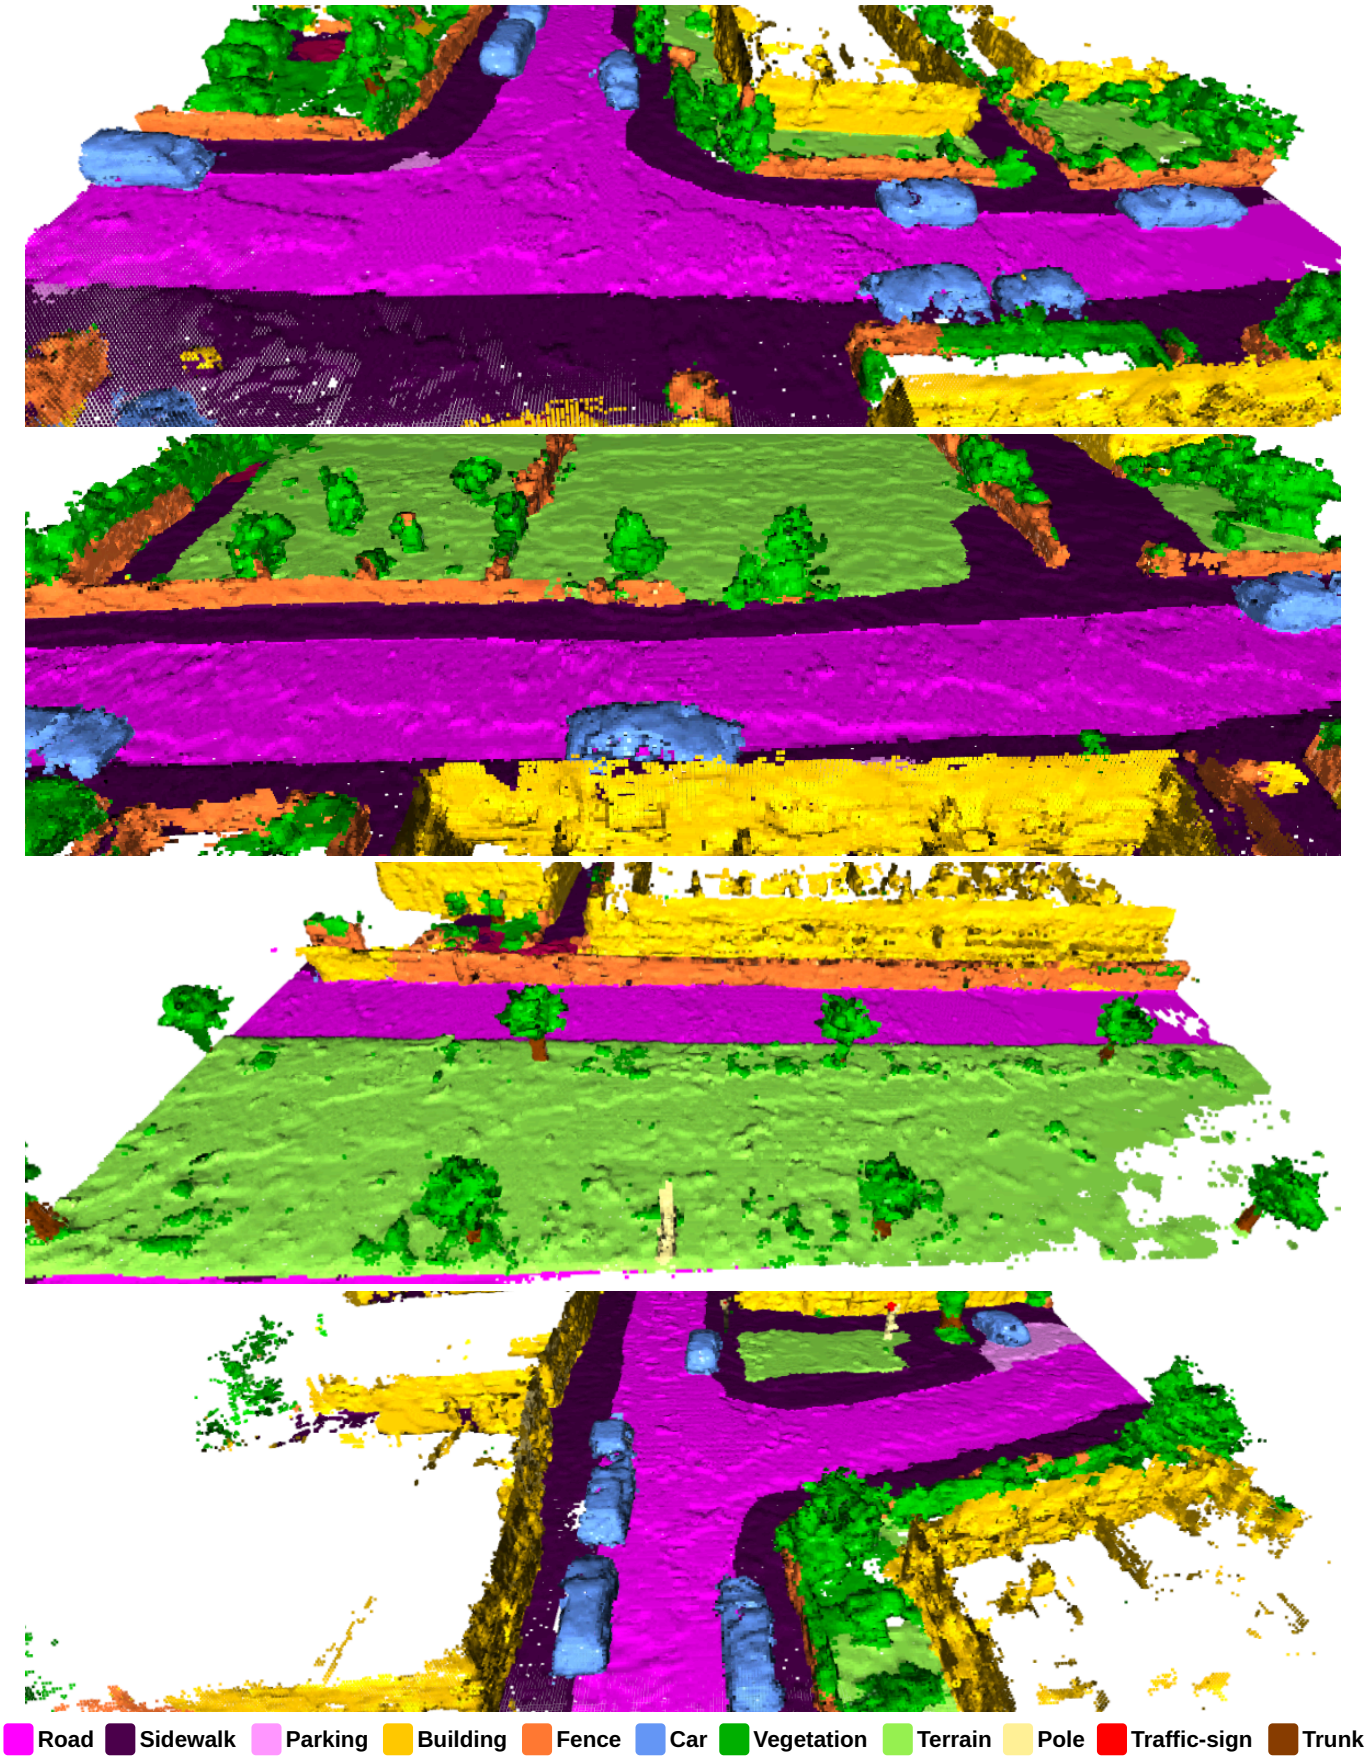

Fig. 10. Unconditional scenes generated by our method.

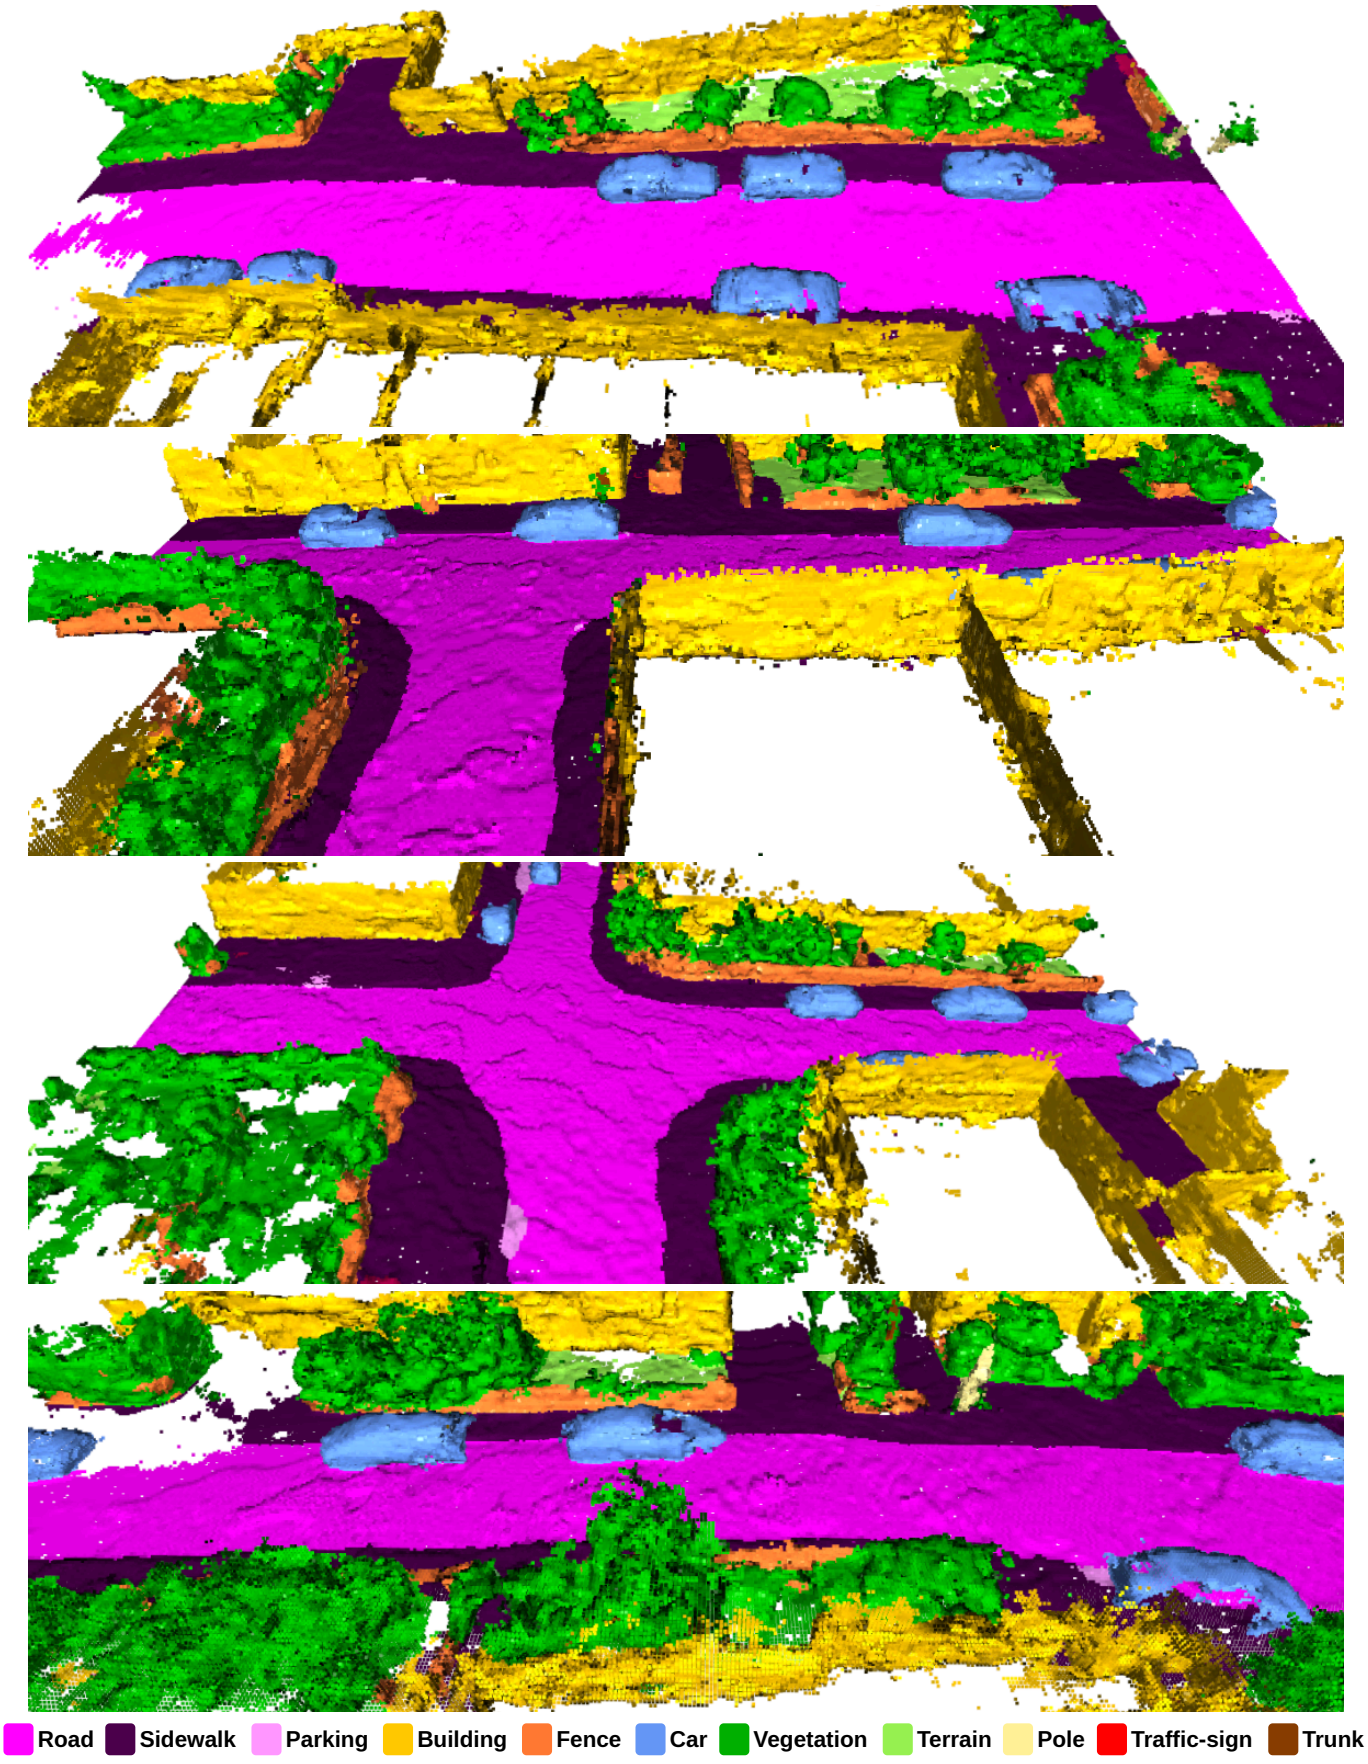

Fig. 11. Unconditional scenes generated by our method.

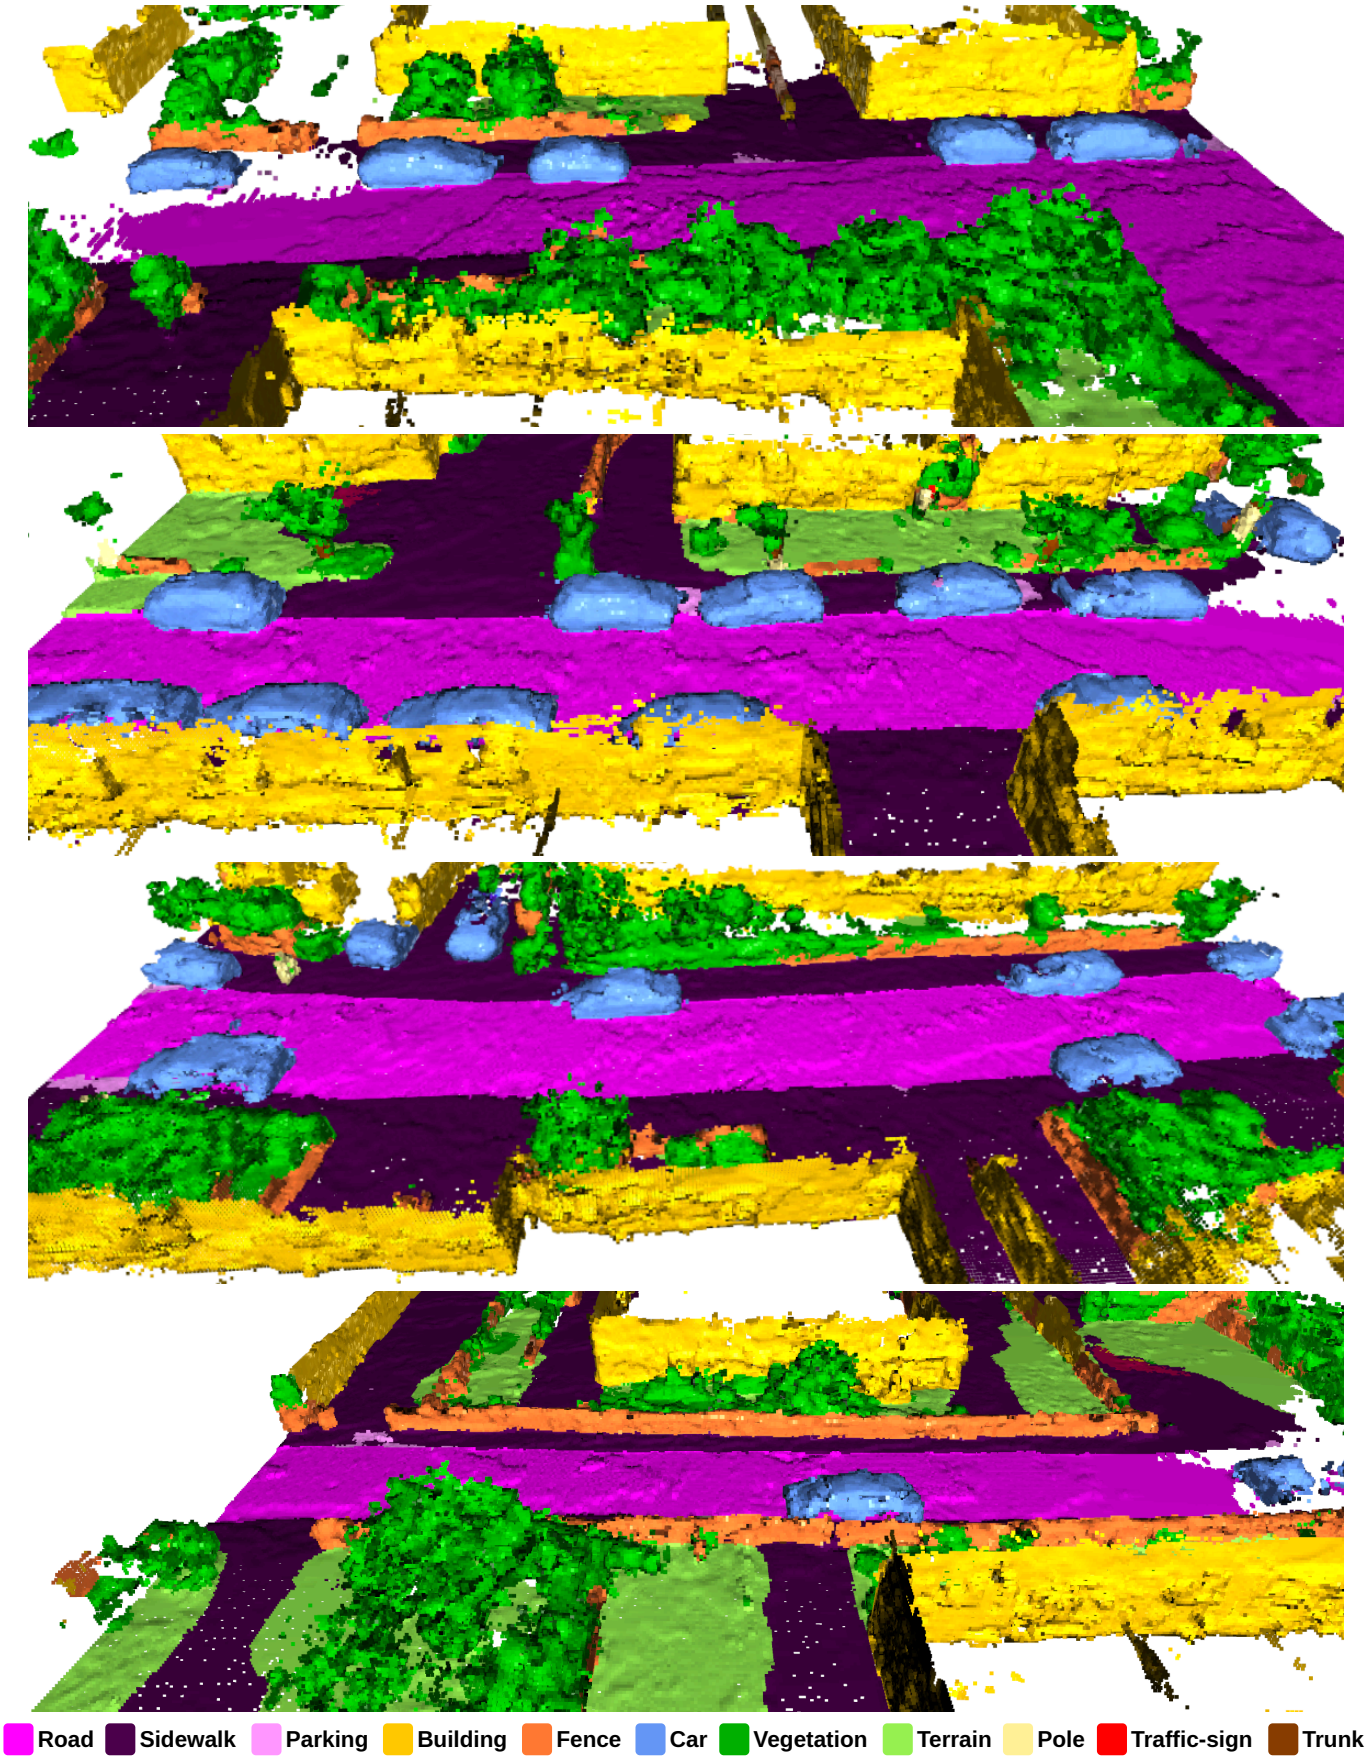

Fig. 12. Unconditional scenes generated by our method.

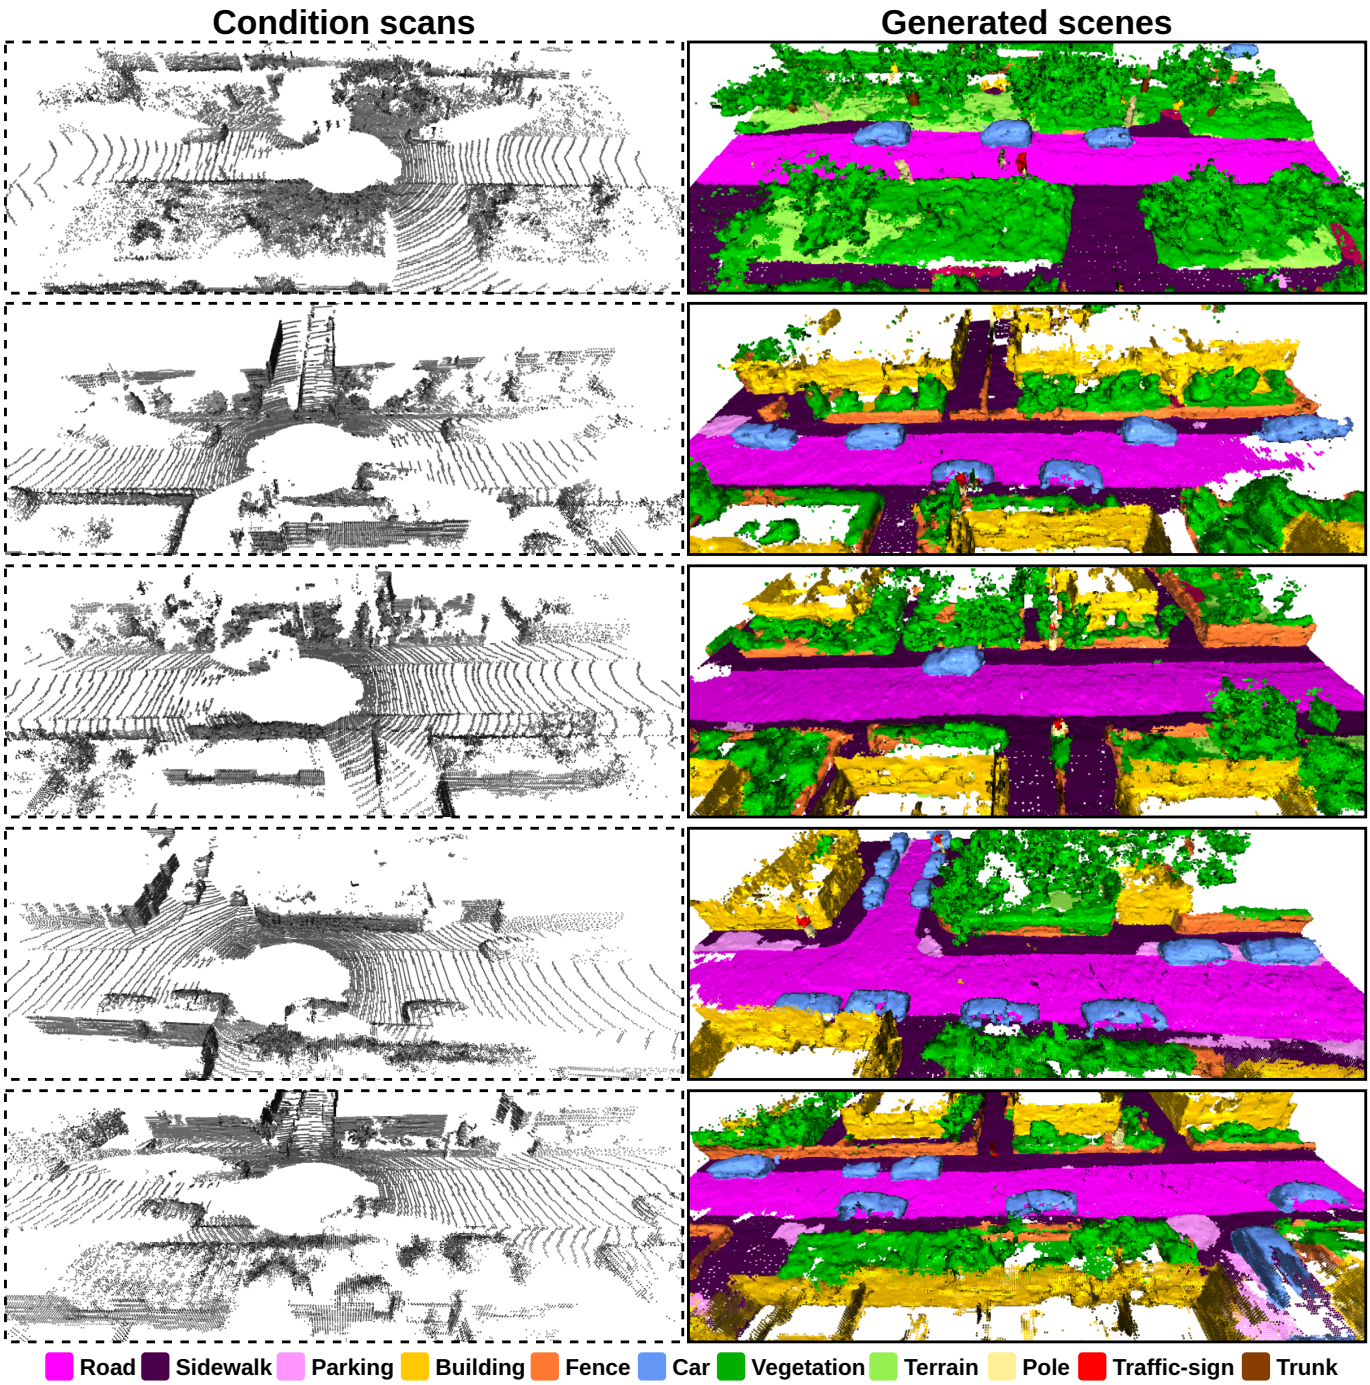

Fig. 13. Generated scenes conditioned to KITTI-360 dataset [4].

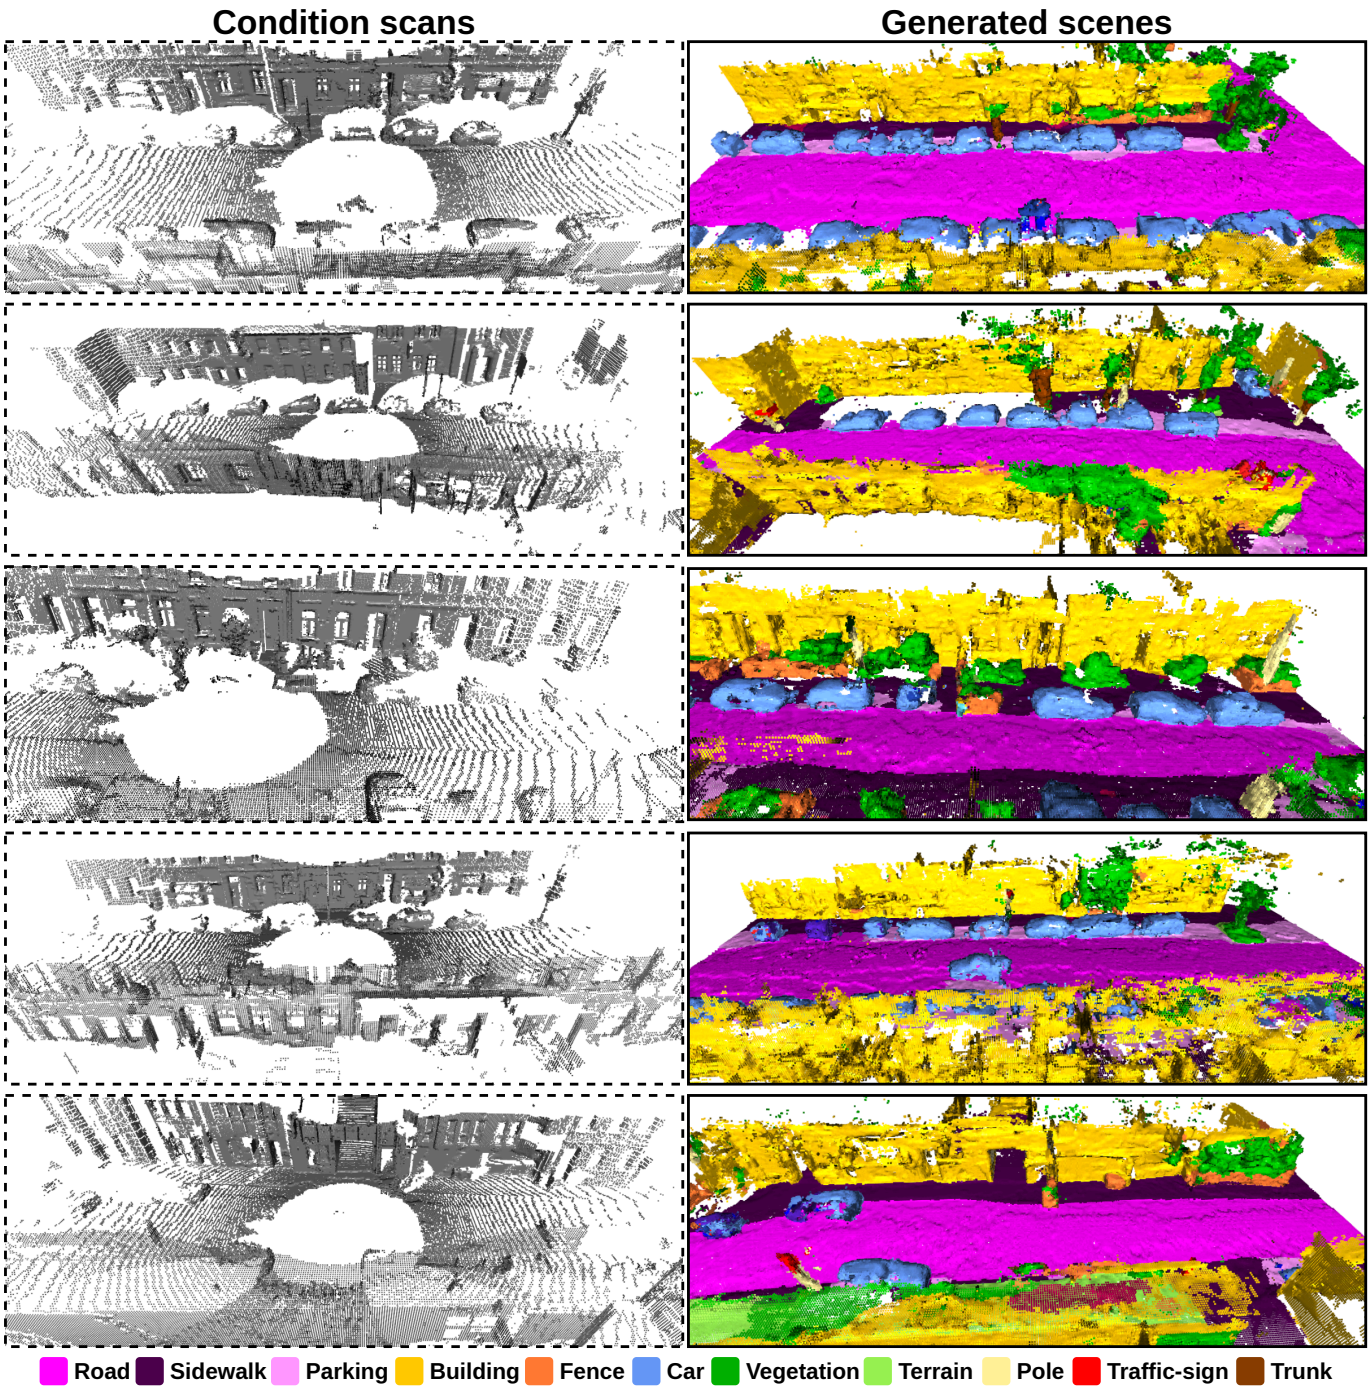

Fig. 14. Generated scenes conditioned to our data collected with ouster OS-1 128 beams LiDAR.