Semantic 3D segmentation of 3D Gaussian Splats

Karol Jurski
K.K.Jurski@student.tudelft.nl
Author

Xucong Zhang
xucong.zhang@tudelft.nl
Responsible Professor

TU Delft

1 Background

What is semantic 3D segmentation?

Process of dividing a 3D mesh or point cloud into semantically meaningful regions, e.g. different classes of objects.

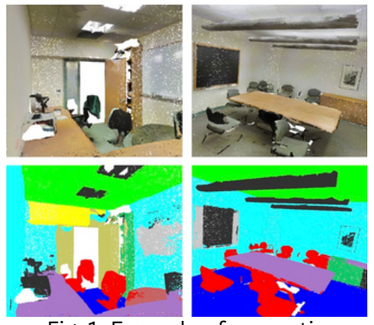

Fig. 1: Example of semantic segmentation [1]

What is 3D Gaussian Splatting (3DGS) [2]?

A recently introduced technique of 3D reconstruction from images using a large number of tiny, Gaussian-shaped splats. It offers photorealistic novel view synthesis and fast rendering times.

Fig. 2: Example of a 3DGS scene [2]

2 Research question

How can semantic 3D segmentation be performed directly on a 3D Gaussian Splats representation?

Sub-questions:

1. Which existing deep learning architecture could be adapted for segmenting 3D Gaussian Splats?
2. Which of the possible representations of 3DGS data is optimal for the chosen deep learning architecture?
3. What is the performance of the method?

3 Method

No existing 3DGS datasets → need to create a custom dataset

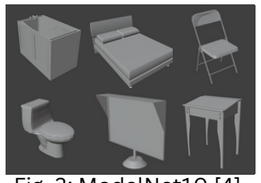

Fig. 3: ModelNet10 [4], a dataset of 4909 objects from 10 categories

Fig. 4: Rendering each object from multiple cameras in Blender

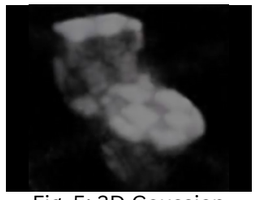

Fig. 5: 3D Gaussian Splats generation

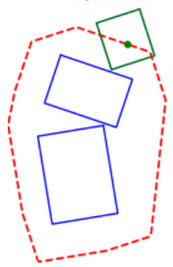

Fig. 6: Composing synthetic scenes from multiple 3DGS objects

Training

+ hyperparameter tuning
+ data augmentation
+ uniform sampling (4096 per scene)

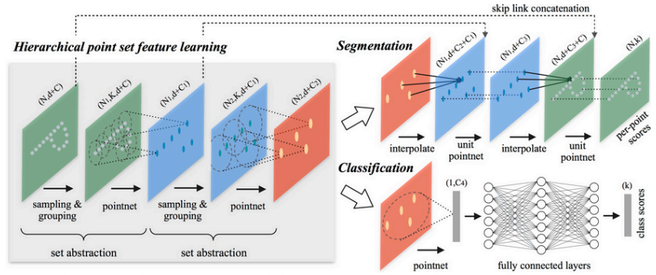

Fig. 7: PointNet++ architecture [3]. PointNet [1] and its successor, PointNet++ [3] are pioneering deep learning architectures for point cloud classification and segmentation.

4 Results

Baseline: PointNet++ segm. on point clouds sampled from meshes from the synthetic scenes.
Experiments: PointNet++ segmentation on the 3DGS synthetic scenes with various 3DGS features used.

	Accuracy	mIoU
Baseline (point cloud sampled from meshes)	84.6%	71%
3DGS: Position	81.2%	68.9%
3DGS: Pos. + opacity	85%	75%
3DGS: Pos. + op. + covariance matrix	81.8%	70.1%
3DGS: Pos. + op. + scale + rotation (matrix)	85.8%	76%
3DGS: Pos. + op. + scale + rotation (quaternion)	**86.8%**	**77.4%**

Table 1: PointNet++ performance on different 3D representations

Fig. 8: Example of a point cloud (baseline) scene segmented by PointNet++

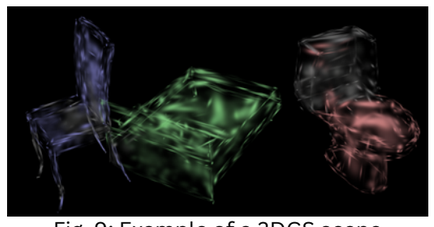

Fig. 9: Example of a 3DGS scene segmented by PointNet++

5 Conclusions

- created a dataset for 3DGS segmentation evaluation
- direct **segmentation of 3DGS** is **possible**
- **PointNet++** is a **suitable** architecture for performing segmentation of 3DGS
- inclusion of additional 3DGS features improves the performance; combination of **position, opacity, size and rotation** (as quaternion) gives accuracy of **86.8%**
- when including all 3DGS features, PointNet++ segmentation **performance for 3DGS higher than for point clouds** sampled from meshes

6 Limitations & further work

- test on real-world (non-synthetic) scenes
- use a textured dataset and assess influence of color on performance
- try more modern point cloud segmentation ML architectures
- compare to other existing (indirect) 3DGS segmentation methods

7 References

[1] R. Qi Charles, Hao Su, Mo Kaichun, and Leonidas J.Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation.
[2] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering.
[3] Charles R. Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: deep hierarchical feature learning on point sets in a metric space.
[4] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes, 2015.