Internship proposal: How Can Neural Networks Express 3D Shapes?

M1/M2 Internship, 4 months; INRIA Paris-Saclay; GEOMERIX Team, in Turing building
Supervision: Mathieu Desbrun (Researcher & Professor), Nissim Maruani (PhD Student)
Keywords: Machine Learning, Geometry Processing
Contact: nissim.maruani@inria.fr & mathieu.desbrun@inria.fr

Context

Recent advancements in 3D geometry processing have leveraged machine learning to tackle ill-posed problems such as shape creation from text prompts [6] or reconstruction from sparse inputs [1, 5, 3]. However, unlike 2D images which are commonly represented as fixed $N \times M$ arrays, there is no single standard *learnable representation for 3D shapes*. Although continuous fields have been used extensively in learning-based tasks [4], they (a) suffer from a distinctive "blobby" aspect, (b) are unable to fit sharp features and small details, and (c) require supervision on whole volumes rather than surfaces. On the other hand, polygon meshes are efficient and compact, but their combinatorial nature makes them challenging to use in a learning-based context. Learnable point-based representations [1, 5, 3] have circumvented that issue by tailoring classical geometry processing methods to stochastic gradient descent (SGD) based optimizations tasks with great success.

Objectives

Depending on the interests of the intern, we propose two research directions:

1. Explore novel ideas to represent 3D shapes with neural networks. As mentioned earlier, the goal would be to bridge the gap between continuous functions and discrete 3D shapes. One possible first step would be to investigate classical geometry processing techniques suitable for SGD optimization, such as APSS [2].

2. Extend existing point-based representations to novel applications. The goal would then be to adapt existing methods to new tasks, such as learning from images, simulation, texture generation, etc.

These two subjects allow for an in-depth exploration of both recent state-of-the-art methods and classical geometry processing techniques. They have a strong publication potential, and could lead to a PhD position.

Expected Skills

The applicant must have a solid background in computer science (and ideally, applied mathematics), and a strong interest in machine learning and geometry processing. Proficiency in *Python* and *PyTorch* would be best.

References

[1] Zhiqin Chen and Hao Zhang. "Neural marching cubes". en. In: *ACM Transactions on Graphics* 40.6 (Dec. 2021), pp. 1–15. ISSN: 0730-0301, 1557-7368. DOI: 10.1145/3478513.3480518. URL: https://dl.acm.org/doi/10.1145/3478513.3480518.

[2] Gaël Guennebaud and Markus Gross. "Algebraic point set surfaces". en. In: *ACM SIGGRAPH 2007 papers*. San Diego California: ACM, July 2007, p. 23. ISBN: 978-1-4503-7836-9. DOI: 10.1145/1275808.1276406. URL: https://dl.acm.org/doi/10.1145/1275808.1276406 (visited on 11/10/2023).

[3] Nissim Maruani, Roman Klokov, Maks Ovsjanikov, Pierre Alliez, and Mathieu Desbrun. "VoroMesh: Learning Watertight Surface Meshes with Voronoi Diagrams". en. In: *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*. 2023, pp. 14565–14574. URL: https://openaccess.thecvf.com/content/ICCV2023/html/Maruani_VoroMesh_Learning_Watertight_Surface_Meshes_with_Voronoi_Diagrams_ICCV_2023_paper.html.

[4] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. "DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation". en. In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Long Beach, CA, USA: IEEE, June 2019, pp. 165–174. ISBN: 978-1-72813-293-8. DOI: 10.1109/CVPR.2019.00025. URL: https://ieeexplore.ieee.org/document/8954065/ (visited on 11/10/2023).

[5] Songyou Peng, Chiyu Jiang, Yiyi Liao, Michael Niemeyer, Marc Pollefeys, and Andreas Geiger. "Shape As Points: A Differentiable Poisson Solver". In: *Advances in Neural Information Processing Systems*. Vol. 34. Curran Associates, Inc., 2021, pp. 13032–13044. URL: https://proceedings.neurips.cc/paper_files/paper/2021/hash/6cd9313ed34ef58bad3fdd504355e72c-Abstract.html.

[6] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. *DreamFusion: Text-to-3D using 2D Diffusion*. en. arXiv:2209.14988 [cs, stat]. Sept. 2022.