Generative Models for non Euclidean Data

Internship Offer - LIRIS - Origami team, (Univ Lyon 1, France)

Advisor: Julie Digne (julie.digne@liris.cnrs.fr)

Realistic images synthesized by StyleGAN3 [5] (a) and comparison with a state-of-the-art shape synthesis method [6] (b). (c) shows a real world shape with details and geometric content.

Context:

Synthesizing data is a big challenge of today's computer vision and computer graphics research. From initial VAE[1] to GAN [2] to diffusion generative models [3] image synthesis has reached a high level of realism, with globally coherent images and realistic high frequency details. Several reasons explain this discrepancy, the most important one being the fact that shapes are non-euclidean data, for which defining an equivariant rotation is an open problem.

While some methods reach some level of realism using either explicit structure retrieval [4], continuous normalizing flows [7] or more recently denoising diffusion models [6], no (or only few) high frequency geometric details are synthesized and the level of realism is not comparable to the one reached for image generation. High resolution detail synthesis is often tackled from a different perspective, by adding some local details to existing shapes [8].Furthermore, the synthesis remains completely shape topology agnostic.

Goal

In this internship, we will focus first on a single of these issues: re-introducing some topological knowledge into shape synthesis. The idea is to synthesize first a shape of the desired topology, and then to deform it into a detailed geometric shape. This point of view will allow for topology-consistent shape interpolation. The internship will start by an extensive review of various recent shape synthesis methods and how shape topology can be introduced in these models.

Required skills: geometry processing, machine learning, optimization

Languages: Python/pytorch

Advisor: Julie Digne

Location: LIRIS, Origami team, Nautibus building, Univ Lyon 1

Salary: 600€ per month

Duration: 4 to 6 months (starting date anytime between Feb 2022 to April 2022)

References

[1] Kingma, D. P., & Welling, M. (2014, April). Stochastic gradient VB and the variational auto-encoder. In Second International Conference on Learning Representations, ICLR (Vol. 19, p. 121).

[2] Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative Adversarial Nets, Proc. NeurIPS 2014

[3] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 10684-10695).

[4] Yang, J., Mo, K., Lai, Y. K., Guibas, L. J., & Gao, L. (2022). DSG-Net: Learning disentangled structure and geometry for 3D shape generation. ACM Transactions on Graphics (TOG).

[5] Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., & Aila, T. (2021). Alias-free generative adversarial networks. Advances in Neural Information Processing Systems, 34, 852-863.

[6] Zeng, X., Vahdat, A., Williams, F., Gojcic, Z., Litany, O., Fidler, S., & Kreis, K. (2022). LION: Latent Point Diffusion Models for 3D Shape Generation, NeurIPS 2022

[7] Yang, G., Huang, X., Hao, Z., Liu, M. Y., Belongie, S., & Hariharan, B. (2019). Pointflow: 3d point cloud generation with continuous normalizing flows. In Proceedings of the IEEE/CVF International Conference on Computer Vision (pp. 4541-4550).

[8] Hamdi-Cherif, A., Digne, J., & Chaine, R. (2018, February). Super-Resolution of Point Set Surfaces Using Local Similarities. In Computer Graphics Forum (Vol. 37, No. 1, pp. 60-70).