PhD position in 3D Computer Vision and geometry processing

Efficient data structures and algorithms for processing massive point clouds

Context

Analyzing 3D point clouds captured from real-world environments is a core component of Geometry Processing and 3D Computer Vision. Processing tasks include, for instance, the estimation of local geometric properties, semantic segmentation, extraction of geometric primitives or reconstruction into surface meshes. Algorithms that perform these tasks are typically designed to handle up to a few million points efficiently [1,2]. With the technological advances on sensors and storage capacity, new acquisition protocols generate more and more massive point clouds that contain billions of points, as illustrated in Figure 1. The naive solution then consists in decomposing the space into blocks of reasonable number of points before performing parallel computing. This solution is however prone to border effect errors and does not allow the analyze of point clouds at global scales. Moreover, it requires high computing resources and storage capacity.

Fig.1 – Illustration of a massive point cloud at three scales with, from left to right, the city of Amsterdam and its surroundings, the downtown and a few building blocks. Image from [6].

Scaling point cloud processing algorithms to billion points without naïve block decomposition is a challenging scientific problem. Among existing works, streaming methods that process data on the fly have been designed towards this goal. They however are tailored made for specific applications [3,4] and cannot be generalized easily to a generic toolbox. Other methods, e.g. [5], operate block decomposition by focusing on border effect reduction. Besides these strategies, the nature of the data structure that encodes input points is also a central question. For visualization applications for instance, octrees constitute a popular choice as levels of details for rendering points can be easily defined by this hierarchical structure [6,7].

Objectives

The goal of this PhD is to (i) investigate new data structures to read, compress and store the information contained in massive point clouds efficiently, and (ii) to rethink popular processing tasks so that they can operate at multiple scales directly from such data structures.

The candidate will study the potential of different space partitioning data structures that can be built efficiently in a hierarchical way and from which information can be stored and requested easily. He/she will also propose compression operations to convert clusters of input points into lightweight geometric objects, and clusters of these geometric objects into single one. The choice of geometric objects will have to account for representation genericity, compactness and efficiency to connect and aggregate them. Prior work shows, for example, that planar components (which are frequent in urban environments) can be turned into a hierarchy of

floating polygons with a limited loss of information. Similarly, the notion of "superpoints" introduced in [9] could also be a solution for compressing non-planar components.

The candidate will also revisit some traditional point cloud processing tasks such as estimation of local geometric properties, surface reconstruction or primitive detection under the idea that the atomic geometric element is not a 3D point anymore, but geometric object living at a given scale of the data structure. Continuing on the previous example with polygons and superpoints, planar shape detection could be simply addressed by selecting polygons in the hierarchy of the data structure, and surface reconstruction, by assembling the geometric objects with a space partition.

The candidate will also investigate the potential of the proposed data structures in recent 3D deep learning architectures which still largely suffer from scalability issues. In particular, the proposed data structures could be an effective alternative to the very coarse simplification of input point clouds [10].

Keywords

Geometry processing, 3D computer vision, massive point clouds, point set processing, geometric data structures

Candidate profile

The ideal candidate should have good knowledge in 3D geometry and applied mathematics, be able to program in C/C++, be fluent in English, and be creative and rigorous.

Application deadline: April, 28th 2024

Contact : Florent.Lafarge@inria.fr

References
[1] The CGAL Project. *CGAL User and Reference Manual*. CGAL Editorial Board, 5.5.1 edition, 2022.
[2] CloudCompare, version 2.10.3, 2022.
[3] Pajarola. Stream-Processing Points. IEEE Visualization 2005
[4] Zhou and Neumann. A streaming framework for seamless building reconstruction from large-scale aerial lidar data. CVPR 2009
[5] Mostegel, Prettenthaler, Fraundorfer and Bischof. Scalable Surface Reconstruction from Point Clouds with Extreme Scale and Density Diversity. CVPR 2017
[6] Schütz, Ohrhallinger, Wimmer. Fast Out-of-Core Octree Generation for Massive Point Clouds. Computer Graphics Forum, vol 39(7), 2020
[7] Elseberg, borrmann and Nuchter. One billion points in the cloud – an octree for efficient processing of 3D laser scans. ISPRS Journal of Photogrammetry and Remote Sensing, vol 76, 2013
[8] Fang, Lafarge, and Desbrun. Shape detection at structural scales. CVPR 2018
[9] Landrieu and Simonovsky. Large-scale Point Cloud Semantic Segmentation with Superpoint Graphs. CVPR 2018
[10] Potamias, Bouritsas and Zafeiriou. Revisiting Point Cloud Simplification : A Learnable Feature Preserving Approach. ECCV 2022