PhD position in Computer Vision and geometry processing

Object vectorization from remote sensing data

Context

Monitoring and anticipating natural disasters and emerging risks using remote sensing data is a key scientific challenge. One of the underlying problems is to describe objects of interest contained in the remote sensing data, typically large-scale images, with vectorized representations such as polygons, planar graphs or networks of parametric curves. These compact and parametrizable representations are important for understanding, analyzing and simulating natural risks.

Object vectorization has been a long-standing problem in Computer Vision. To capture the silhouette of a solid object in an image for instance, many works have considered polygons as a relevant representation choice: they are compact, editable and can basically approximate well any close chain of pixels. Traditionally, algorithms such as Douglas-Peucker [1] were used to simplify such a chain of pixels extracted after a semantic segmentation of an image at the pixel level. More recently, deep Learning methods were proposed to extract polygons directly from input images using typically recurrent neural networks to predict the next node of the polygon [2,3] or graph neural networks to link a list of predicted nodes [4]. Detecting and vectorizing objects simultaneously in an end-to-end fashion is however very ambitious as neural networks, while impressive for predicting, are still not armed for exploring the complex solution spaces of geometric objects and structures. Recent works, e.g., for line-segments [5] or low complexity polygons [6], suggest that hybrid methods remain the best option with typically prediction of raster maps by deep learning and vectorization of elements of these raster maps by geometric approaches.

One of the main difficulties in our context is that objects of interest can have various appearances and geometric specificities. For instance, buildings are surface objects whose boundaries are piecewise-lineic while roads or faults are curved line networks. In the literature methods are usually specific to only one type of vector representation.

Objectives

The goal of this PhD is to investigate new geometric models to vectorize a large family of objects contained in remote sensing data. These models should take as input the remote sensing data and some prediction maps that estimate the semantic class of objects at the pixel level. These maps will be typically provided by the last generation of deep learning architectures for classification and semantic segmentation tasks.

The candidate will design a generic method able to vectorize different types of geometric structures, going from polygons, e.g. [6], to planar graphs, e.g. [7], through networks of parametric curved, e.g. [8]. To offer such a high genericity on the structure types, a strategy could consist in fitting an underlying space partitioning data structure to input images and prediction maps. This idea was explored in [9] with promising results to extract polygons and planar graphs using Delaunay triangulations. This preliminary work however suffers from some weaknesses, including a lack of genericity of Delaunay triangulations and a slow inference based on Monte Carlo sampling. The candidate will analyze the possible space partitioning data structures that are relevant in terms of (i) representation genericity and (ii) capacity to operate efficiently geometric and semantic modifications to better fit input data. The candidate will also investigate efficient optimization mechanisms to explore the large solution space of such problems. In particular, deterministic processes that order local modifications, e.g. [6], and that reach solutions of a similar quality of those obtained by slow stochastic simulations will be an direction to study.

The candidate will also investigate the extension of the 2D vectorization models to 3D, with a focus on the vectorization of 3D line networks from 3D data measurements such as airborne Lidar scans and Multiview stereo imagery.

Keywords

computer vision, geometry processing, remote sensing, object vectorization, geometric data structures

Candidate profile

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

How to candidate

Candidates must send their CV, academic transcript for their Bachelor and Master, a motivation letter and optionally some recommendation letters to Florent.Lafarge@inria.fr before April 8, 2024

References

[1] Wu and Marquez. A non-self-intersection Douglas-Peucker algorithm. In IEEE Symposium on Computer Graphics and Image Processing, 2003

[2] Li, Wegner and Lucchi. Topological Map Extraction From Overhead Images. Proc. of the IEEE International Conference on Computer Vision (ICCV) 2019

[3] Acuna, Ling, Kar and Fidler. Efficient interactive annotation of segmentation datasets with polygon-rnn++. Proc. of the IEEE conference on Computer Vision and Pattern Recognition (CVPR), 2018

[4] Zorzi, Bazrafkan, Habenschuss, Fraundorfer. PolyWorld: Polygonal Building Extraction With Graph Neural Networks in Satellite Images. Proc. of the IEEE conference on Computer Vision and Pattern Recognition (CVPR), 2022

[5] Pautrat, Barath, Larsson, Oswald, Pollefeys. DeepLSD: Line Segment Detection and Refinement with Deep Image Gradients. Proc. of the IEEE conference on Computer Vision and Pattern Recognition (CVPR), 2023

[6] Li M., Lafarge and Marlet. Approximating shapes in images with low-complexity polygons. Proc. of the IEEE conference on Computer Vision and Pattern Recognition (CVPR), 2020

[7] Bahl, Bahri and Lafarge. Single-Shot End-to-end Road Graph Extraction. Proc. of the CVPR EarthVision Workshop, 2022

[8] Favreau, Lafarge and Bousseau. Fidelity vs. Simplicity: a Global Approach to Line Drawing Vectorization. roc. of SIGGRAPH conference, 2016

[9] Favreau, Lafarge, Bousseau and Auvolat. Extracting Geometric Structures in Images with Delaunay Point Processes. IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI), Vol. 42(4), 2020