Activity Report 2019

Project-Team TITANE

Geometric Modeling of 3D Environments

RESEARCH CENTER
Sophia Antipolis - Méditerranée

THEME
Interaction and visualization

Table of contents

Project-Team TITANE

Creation of the Team: 2013 January 01, updated into Project-Team: 2014 January 01

Keywords:

Computer Science and Digital Science:

A5. - Interaction, multimedia and robotics
A5.3. - Image processing and analysis
A5.3.2. - Sparse modeling and image representation
A5.3.3. - Pattern recognition
A5.5.1. - Geometrical modeling
A5.6. - Virtual reality, augmented reality
A5.6.1. - Virtual reality
A5.6.2. - Augmented reality
A8.3. - Geometry, Topology
A8.12. - Optimal transport
A9.2. - Machine learning

Other Research Topics and Application Domains:

B2.5. - Handicap and personal assistances
B3.3. - Geosciences
B5.1. - Factory of the future
B5.6. - Robotic systems
B5.7. - 3D printing
B8.3. - Urbanism and urban planning

1. Team, Visitors, External Collaborators

Research Scientists

Pierre Alliez [Team leader, Inria, Senior Researcher, HDR]
Florent Lafarge [Inria, Researcher, HDR]
Yuliya Tarabalka [Inria, Researcher, until Jan 2019]
Mathieu Desbrun [Caltech, International Chair, Advanced Research position until Mar 2019]

PhD Students

Gaetan Bahl [IRT Saint-Exupéry, PhD Student, from Mar 2019]
Jean-Philippe Bauchet [Luxcarta, PhD Student]
Oussama Ennafii [IGN, PhD Student, until Sep 2019]
Hao Fang [Inria, PhD Student, until Jan 2019]
Nicolas Girard [Inria, PhD Student]
Muxingzi Li [Inria, PhD Student]
Lionel Matteo [Thales, PhD Student]
Flora Quilichini [Inria, PhD Student, until Jun 2019]
Onur Tasar [Inria, PhD Student]
Vincent Vadez [Dorea, PhD Student]
Julien Vuillamy [Dassault Systèmes, PhD Student]
Mulin Yu [Inria, PhD Student, from Nov 2019]
Tong Zhao [Inria, PhD Student, from Nov 2019]

Technical staff

Dmitry Anisimov [Inria, Engineer, until Apr 2019]
Fernando Ireta Munoz [Inria, Engineer]
Cédric Portaneri [Inria, Engineer]

Interns and Apprentices

Roxane Debord [Inria, from Jun 2019 until Aug 2019]
Mulin Yu [Inria, from Mar 2019 until Aug 2019]
Tong Zhao [Inria, from Apr 2019 until Sep 2019]
Andrii Zhygallo [Inria, until Feb 2019]

Administrative Assistant

Florence Barbara [Inria, Administrative Assistant]

Visiting Scientist

Jorg Peters [Uni. Florida, from Jun 2019 until Jul 2019]

2. Overall Objectives

2.1. General Presentation

Our overall objective is the computerized geometric modeling of complex scenes from physical measurements. On the geometric modeling and processing pipeline, this objective corresponds to steps required for conversion from physical to effective digital representations: *analysis*, *reconstruction* and *approximation*. Another longer term objective is the *synthesis* of complex scenes. This objective is related to analysis as we assume that the main sources of data are measurements, and synthesis is assumed to be carried out from samples.

The related scientific challenges include i) being resilient to defect-laden data due to the uncertainty in the measurement processes and imperfect algorithms along the pipeline, ii) being resilient to heterogeneous data, both in type and in scale, iii) dealing with massive data, and iv) recovering or preserving the structure of complex scenes. We define the quality of a computerized representation by its i) geometric accuracy, or faithfulness to the physical scene, ii) complexity, iii) structure accuracy and control, and iv) amenability to effective processing and high level scene understanding.

3. Research Program

3.1. Context

Geometric modeling and processing revolve around three main end goals: a computerized shape representation that can be visualized (creating a realistic or artistic depiction), simulated (anticipating the real) or realized (manufacturing a conceptual or engineering design). Aside from the mere editing of geometry, central research themes in geometric modeling involve conversions between physical (real), discrete (digital), and mathematical (abstract) representations. Going from physical to digital is referred to as shape acquisition and reconstruction; going from mathematical to discrete is referred to as shape approximation and mesh generation; going from discrete to physical is referred to as shape rationalization.

Geometric modeling has become an indispensable component for computational and reverse engineering. Simulations are now routinely performed on complex shapes issued not only from computer-aided design but also from an increasing amount of available measurements. The scale of acquired data is quickly growing: we no longer deal exclusively with individual shapes, but with entire *scenes*, possibly at the scale of entire cities, with many objects defined as structured shapes. We are witnessing a rapid evolution of the acquisition paradigms with an increasing variety of sensors and the development of community data, as well as disseminated data.

In recent years, the evolution of acquisition technologies and methods has translated in an increasing overlap of algorithms and data in the computer vision, image processing, and computer graphics communities. Beyond the rapid increase of resolution through technological advances of sensors and methods for mosaicing images, the line between laser scan data and photos is getting thinner. Combining, e.g., laser scanners with panoramic cameras leads to massive 3D point sets with color attributes. In addition, it is now possible to generate dense point sets not just from laser scanners but also from photogrammetry techniques when using a well-designed acquisition protocol. Depth cameras are getting increasingly common, and beyond retrieving depth information we can enrich the main acquisition systems with additional hardware to measure geometric information about the sensor and improve data registration: e.g., accelerometers or GPS for geographic location, and compasses or gyrometers for orientation. Finally, complex scenes can be observed at different scales ranging from satellite to pedestrian through aerial levels.

These evolutions allow practitioners to measure urban scenes at resolutions that were until now possible only at the scale of individual shapes. The related scientific challenge is however more than just dealing with massive data sets coming from increase of resolution, as complex scenes are composed of multiple objects with structural relationships. The latter relate i) to the way the individual shapes are grouped to form objects, object classes or hierarchies, ii) to geometry when dealing with similarity, regularity, parallelism or symmetry, and iii) to domain-specific semantic considerations. Beyond reconstruction and approximation, consolidation and synthesis of complex scenes require rich structural relationships.

The problems arising from these evolutions suggest that the strengths of geometry and images may be combined in the form of new methodological solutions such as photo-consistent reconstruction. In addition, the process of measuring the geometry of sensors (through gyrometers and accelerometers) often requires both geometry process and image analysis for improved accuracy and robustness. Modeling urban scenes from measurements illustrates this growing synergy, and it has become a central concern for a variety of applications ranging from urban planning to simulation through rendering and special effects.

3.2. Analysis

Complex scenes are usually composed of a large number of objects which may significantly differ in terms of complexity, diversity, and density. These objects must be identified and their structural relationships must be recovered in order to model the scenes with improved robustness, low complexity, variable levels of details and ultimately, semantization (automated process of increasing degree of semantic content).

Object classification is an ill-posed task in which the objects composing a scene are detected and recognized with respect to predefined classes, the objective going beyond scene segmentation. The high variability in each class may explain the success of the stochastic approach which is able to model widely variable classes. As it requires a priori knowledge this process is often domain-specific such as for urban scenes where we wish to distinguish between instances as ground, vegetation and buildings. Additional challenges arise when each class must be refined, such as roof super-structures for urban reconstruction.

Structure extraction consists in recovering structural relationships between objects or parts of object. The structure may be related to adjacencies between objects, hierarchical decomposition, singularities or canonical geometric relationships. It is crucial for effective geometric modeling through levels of details or hierarchical multiresolution modeling. Ideally we wish to learn the structural rules that govern the physical scene manufacturing. Understanding the main canonical geometric relationships between object parts involves detecting regular structures and equivalences under certain transformations such as parallelism, orthogonality and symmetry. Identifying structural and geometric repetitions or symmetries is relevant for dealing with missing data during data consolidation.

Data consolidation is a problem of growing interest for practitioners, with the increase of heterogeneous and defect-laden data. To be exploitable, such defect-laden data must be consolidated by improving the data sampling quality and by reinforcing the geometrical and structural relations sub-tending the observed scenes. Enforcing canonical geometric relationships such as local coplanarity or orthogonality is relevant for registration of heterogeneous or redundant data, as well as for improving the robustness of the reconstruction process.

3.3. Approximation

Our objective is to explore the approximation of complex shapes and scenes with surface and volume meshes, as well as on surface and domain tiling. A general way to state the shape approximation problem is to say that we search for the shape discretization (possibly with several levels of detail) that realizes the best complexity / distortion trade-off. Such a problem statement requires defining a discretization model, an error metric to measure distortion as well as a way to measure complexity. The latter is most commonly expressed in number of polygon primitives, but other measures closer to information theory lead to measurements such as number of bits or minimum description length.

For surface meshes we intend to conceive methods which provide control and guarantees both over the global approximation error and over the validity of the embedding. In addition, we seek for resilience to heterogeneous data, and robustness to noise and outliers. This would allow repairing and simplifying triangle soups with cracks, self-intersections and gaps. Another exploratory objective is to deal generically with different error metrics such as the symmetric Hausdorff distance, or a Sobolev norm which mixes errors in geometry and normals.

For surface and domain tiling the term meshing is substituted for tiling to stress the fact that tiles may be not just simple elements, but can model complex smooth shapes such as bilinear quadrangles. Quadrangle surface tiling is central for the so-called *resurfacing* problem in reverse engineering: the goal is to tile an input raw surface geometry such that the union of the tiles approximates the input well and such that each tile matches certain properties related to its shape or its size. In addition, we may require parameterization domains with a simple structure. Our goal is to devise surface tiling algorithms that are both reliable and resilient to defect-laden inputs, effective from the shape approximation point of view, and with flexible control upon the structure of the tiling.

3.4. Reconstruction

Assuming a geometric dataset made out of points or slices, the process of shape reconstruction amounts to recovering a surface or a solid that matches these samples. This problem is inherently ill-posed as infinitely-many shapes may fit the data. One must thus regularize the problem and add priors such as simplicity or smoothness of the inferred shape.

The concept of geometric simplicity has led to a number of interpolating techniques commonly based upon the Delaunay triangulation. The concept of smoothness has led to a number of approximating techniques that commonly compute an implicit function such that one of its isosurfaces approximates the inferred surface. Reconstruction algorithms can also use an explicit set of prior shapes for inference by assuming that the observed data can be described by these predefined prior shapes. One key lesson learned in the shape problem is that there is probably not a single solution which can solve all cases, each of them coming with its own distinctive features. In addition, some data sets such as point sets acquired on urban scenes are very domain-specific and require a dedicated line of research.

In recent years the *smooth, closed case* (i.e., shapes without sharp features nor boundaries) has received considerable attention. However, the state-of-the-art methods have several shortcomings: in addition to being in general not robust to outliers and not sufficiently robust to noise, they often require additional attributes as input, such as lines of sight or oriented normals. We wish to devise shape reconstruction methods which are both geometrically and topologically accurate without requiring additional attributes, while exhibiting resilience to defect-laden inputs. Resilience formally translates into stability with respect to noise and outliers. Correctness of the reconstruction translates into convergence in geometry and (stable parts of) topology of the reconstruction with respect to the inferred shape known through measurements.

Moving from the smooth, closed case to the *piecewise smooth case* (possibly with boundaries) is considerably harder as the ill-posedness of the problem applies to each sub-feature of the inferred shape. Further, very few approaches tackle the combined issue of robustness (to sampling defects, noise and outliers) and feature reconstruction.

4. Application Domains

4.1. Domain 1

In addition to tackling enduring scientific challenges, our research on geometric modeling and processingis motivated by applications to computational engineering, reverse engineering, digital mapping and urbanplanning. The main deliverable of our research will be algorithms with theoretical foundations. Ultimatelywe wish to contribute making geometry modeling and processing routine for practitioners who deal with real-world data. Our contributions may also be used as a sound basis for future software and technology developments.

Our first ambition for technology transfer is to consolidate the components of our research experiments in the form of new software components for the CGAL (Computational Geometry Algorithms Library) library. Consolidation being best achieved with the help of an engineer, we will search for additional funding. Through CGAL, we wish to contribute to the “standard geometric toolbox”, so as to provide a generic answer to application needs instead of fragmenting our contributions. We already cooperate with the Inria spin-offcompany Geometry Factory, which commercializes CGAL, maintains it and provide technical support.

Our second ambition is to increase the research momentum of companies through advising Cifre Ph.D. thesesand postdoctoral fellows on topics that match our research program.

5. Highlights of the Year

5.1. Highlights of the Year

Pierre Alliez was program co-chair of the EUROGRAPHICS 2019 conference and of the Symposium on Solid and Physical Modeling (SPM). From February 2019 Yuliya Tarabalka is on leave to the Luxcarta company for two years.

5.1.1. Awards

Cédric Portaneri and Pierre Alliez obtained a best paper award at the ACM Conference on Multimedia Systems for a contribution to the progressive compression of textured meshes, in collaboration with the Draco team from Google. Jean-Philippe Bauchet obtained an award for the best presentation at a national workshop (GMTG 2019). Jean-Dominique Favreau received the best PhD thesis award 2019 (assessit prize) from IG-RV. Onur Tasar was part of the winning team of the tomtom AI summer school challenge organized in the Netherlands.

6. New Software and Platforms

6.1. CGAL Barycentric_coordinates_2

Module CGAL : Barycentric coordinates 2D

KEYWORD: Computational geometry

FUNCTIONAL DESCRIPTION: This package offers an efficient and robust implementation of two-dimensional closed-form generalized barycentric coordinates defined for simple two-dimensional polygons.

- Participants: Dmitry Anisimov and Pierre Alliez
- Contact: Pierre Alliez

6.2. dtk-nurbs-probing

KEYWORDS: Algorithm - CAD - Numerical algorithm - Geometric algorithms

FUNCTIONAL DESCRIPTION: This library offers tools for computing intersection between linear primitives and the constitutive elements of CAD objects (curves and surfaces). It is thus possible to compute intersections between a linear primitive with a trimmed or untrimmed NURBS surface, as well with Bezier surfaces. It is also possible, in the xy plane, to compute the intersections between linear primitives and NURBS curves as well as Bezier curves.

- Participants: Come Le Breton, Laurent Busé and Pierre Alliez
- Contact: Come Le Breton

6.3. Module CGAL : Point Set Processing

KEYWORD: Geometry Processing

FUNCTIONAL DESCRIPTION: This CGAL component implements methods to analyze and process unorganized point sets. The input is an unorganized point set, possibly with normal attributes (unoriented or oriented). The point set can be analyzed to measure its average spacing, and processed through functions devoted to the simplification, outlier removal, smoothing, normal estimation, normal orientation and feature edges estimation.

- Participants: Clément Jamin, Laurent Saboret and Pierre Alliez
- Contact: Pierre Alliez
- URL: http://doc.cgal.org/latest/Point_set_processing_3/index.html#Chapter_Point_Set_Processing

6.4. Module CGAL : Scale space surface reconstruction

KEYWORD: Geometric algorithms

SCIENTIFIC DESCRIPTION: This CGAL package implements a surface reconstruction method which takes as input an unordered point set and computes a triangulated surface mesh interpolating the point set. We assume that the input points were sampled from the surface of an object. The method can also process point sets sampled from the interior of the object, although we cannot provide guarantees on the output. This method can handle a decent amount of noise and outliers. The point set may greatly undersample the object in occluded regions, although no surface will be reconstructed to fill these regions.

FUNCTIONAL DESCRIPTION: This method allows to reconstruct a surface that interpolates a set of 3D points. This method provides an efficient alternative to the Poisson surface reconstruction method. The main difference in output is that this method reconstructs a surface that interpolates the point set (as opposed to approximating the point set). How the surface connects the points depends on a scale variable, which can be estimated semi-automatically.

- Participants: Pierre Alliez and Thijs Van Lankveld
- Contact: Pierre Alliez

6.5. Module Gudhi : Skeleton-Blockers

Skeleton-Blockers data-structure

KEYWORDS: C++ - Mesh - Triangulation - Topology - 3D

FUNCTIONAL DESCRIPTION: Skeleton-Blockers is a compact, efficient and generic data-structure that can represent any simplicial complex. The implementation is in C++11.

- Participant: David Salinas
- Contact: David Salinas
- URL: https://project.inria.fr/gudhi/software/

6.6. DPP

Delaunay Point Process for image analysis

KEYWORDS: Computer vision - Shape recognition - Delaunay triangulation - Stochastic process

FUNCTIONAL DESCRIPTION: The software extract 2D geometric structures (planar graphs, polygons...) from images

- Participants: Jean-Dominique Favreau, Florent Lafarge and Adrien Bousseau
- Contact: Florent Lafarge
- Publication: Extracting Geometric Structures in Images with Delaunay Point Processes

6.7. KIPPI

KInetic Polygonal Partitioning of Images

KEYWORDS: Computer vision - Computational geometry - Image segmentation

SCIENTIFIC DESCRIPTION: The scientific description of the algorithm is detailed in [Bauchet and Lafarge, KIPPI: KInetic Polygonal Partitioning of Images, CVPR 2018]

FUNCTIONAL DESCRIPTION: KIPPI decompose an image, or a bounded 2D space, into convex polygons. The method exploits a kinetic framework for propagating and colliding line-segments until forming convex polygons.

- Participants: Jean-Philippe Bauchet and Florent Lafarge
- Contact: Florent Lafarge

6.8. Module CGAL: 3D Point-Set Shape Detection

KEYWORD: CGAL

FUNCTIONAL DESCRIPTION: This package implements the efficient RANSAC method for shape detection, contributed by Schnabel et al. From an unstructured point set with unoriented normals, the algorithm detects a set of shapes. Five types of primitive shapes are provided by this package: plane, sphere, cylinder, cone and torus. Detecting other types of shapes is possible by implementing a class derived from a base shape.

- Participants: Clément Jamin, Pierre Alliez and Sven Oesau
- Contact: Pierre Alliez

6.9. CGAL module: Classification

KEYWORDS: Classification - Point cloud - Mesh

FUNCTIONAL DESCRIPTION: This CGAL module aims at classifying 3D data, typically point clouds, into arbitrary classes of interest. The module offers the user the possibility to segment data i) locally or globally, and ii) in an supervised or unsupervised way.

- Authors: Florent Lafarge and Simon Giraudot
- Contact: Florent Lafarge

6.10. SMICER

KEYWORDS: Geometric modeling - Computational geometry - Polyhedral meshes

FUNCTIONAL DESCRIPTION: The software allows the decomposition of a 3D domain into a polyhedra from a set of planar shapes

- Participants: Florent Lafarge and Pierre Alliez
- Contact: Florent Lafarge

6.11. Stochastic Vectorization

KEYWORDS: Vector graphics - Stochastic models

FUNCTIONAL DESCRIPTION: The software converts a line-drawing image into Bezier curves.

- Participants: Jean-Dominique Favreau, Florent Lafarge and Adrien Bousseau
- Contact: Florent Lafarge
- Publication: 01309271

7. New Results

7.1. Analysis

7.1.1. *Pyramid scene parsing network in 3D: improving semantic segmentation of point clouds with multi-scale contextual information*

Participants: Hao Fang, Florent Lafarge.

Analyzing and extracting geometric features from 3D data is a fundamental step in 3D scene understanding. Recent works demonstrated that deep learning architectures can operate directly on raw point clouds, i.e. without the use of intermediate grid-like structures. These architectures are however not designed to encode contextual information in-between objects efficiently. Inspired by a global feature aggregation algorithm designed for images, we propose a 3D pyramid module to enrich pointwise features with multi-scale contextual information. Our module can be easily coupled with 3D semantic segmantation methods operating on 3D point clouds. We evaluated our method on three large scale datasets with four baseline models. Experimental results show that the use of enriched features brings significant improvements to the semantic segmentation of indoor and outdoor scenes (See Figure 1). This work was published in the ISPRS journal of Remote Sensing and Photogrammetry [6].

7.1.2. *Low-power neural networks for semantic segmentation of satellite images*

Participants: Gaetan Bahl, Florent Lafarge.

In collaboration with Lionel Daniel and Matthieu Moretti (IRT Saint-Exupéry).

Semantic segmentation methods have made impressive progress with deep learning. However, while achieving higher and higher accuracy, state-of-the-art neural networks overlook the complexity of architectures, which typically feature dozens of millions of trainable parameters. As a result, these networks requires high computational resources and are mostly not suited to perform on edge devices with tight resource constraints, such as phones, drones, or satellites. In this work, we propose two highly-compact neural network architectures for semantic segmentation of images, which are up to 100 000 times less complex than state-of-the-art architectures while approaching their accuracy. To decrease the complexity of existing networks, our main ideas consist in exploiting lightweight encoders and decoders with depth-wise separable convolutions and decreasing memory usage with the removal of skip connections between encoder and decoder. Our architectures are designed to be implemented on a basic FPGA such as the one featured on the Intel Altera Cyclone V family. We demonstrate the potential of our solutions in the case of binary segmentation of remote sensing images, in particular for extracting clouds and trees from RGB satellite images. This work was published in the Low-Power Computer Vision ICCV workshop [13].

7.1.3. *A learning approach to evaluate the quality of 3D city models*

Participants: Oussama Ennafii, Florent Lafarge.

In collaboration with Arnaud Le Bris and Clément Mallet (IGN).

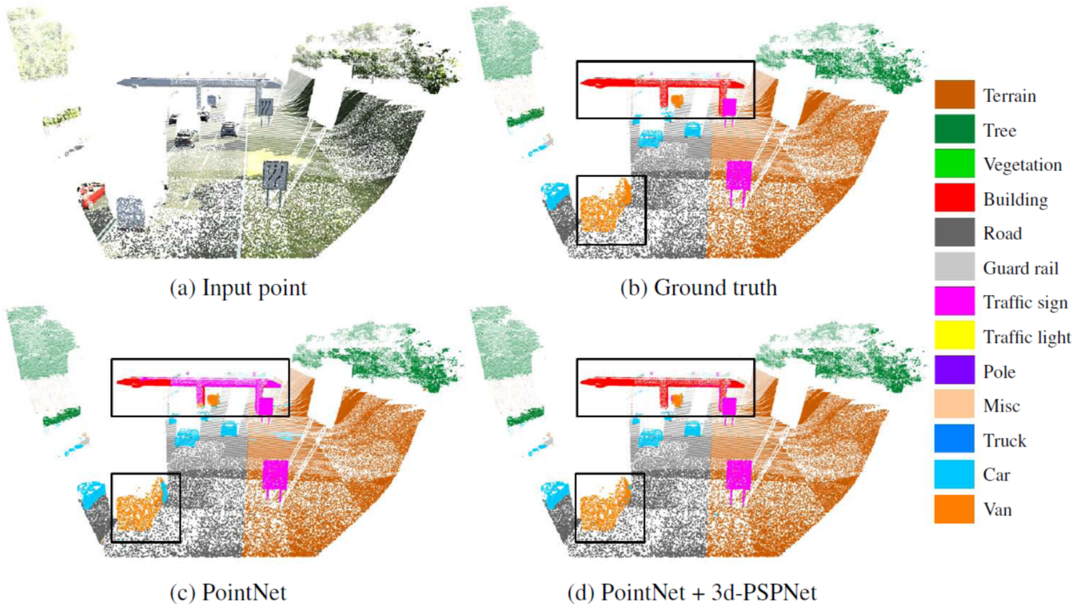

Figure 1. Semantic segmentation of a point cloud with and without our 3d-PSPNet module. Given an input point cloud (a), PointNet fails to predict correct labels for points describing large-scale objects (see rectangles in (c)). PointNet equipped with our 3d-PSPNet module gives better prediction results by enriching global contextual information (d).

The automatic generation of 3D building models from geospatial data is now a standard procedure. An abundant literature covers the last two decades and several softwares are now available. However, urban areas are very complex environments. Inevitably, practitioners still have to visually assess, at city-scale, the correctness of these models and detect frequent reconstruction errors. Such a process relies on experts, and is highly time-consuming with approximately two hours per square kilometer for one expert. This work proposes an approach for automatically evaluating the quality of 3D building models. Potential errors are compiled in a novel hierarchical and versatile taxonomy. This allows, for the first time, to disentangle fidelity and modeling errors, whatever the level of detail of the modeled buildings. The quality of models is predicted using the geometric properties of buildings and, when available, Very High Resolution images and Digital Surface Models. A baseline of handcrafted, yet generic, features is fed into a Random Forest classifier. Both multi-class and multi-label cases are considered: due to the interdependence between classes of errors, it is possible to retrieve all errors at the same time while simply predicting correct and erroneous buildings. The proposed framework was tested on three distinct urban areas in France with more than 3,000 buildings. $80\% - 99\%$ F-score values are attained for the most frequent errors. For scalability purposes, the impact of the urban area composition on the error prediction was also studied, in terms of transferability, generalization and representativeness of the classifiers. It shows the necessity of multimodal remote sensing data and mixing training samples from various cities to ensure stability of the detection ratios, even with very limited training set sizes. This work was presented at the IGARSS conference [16] and published in the PE&RS journal [5].

7.1.4. Robust joint image reconstruction from color and monochrome cameras

Participant: Muxingzi Li.

In collaboration with Peihan Tu (Uni. of Maryland) and Wolfgang Heidrich (KAUST).

Recent years have seen an explosion of the number of camera modules integrated into individual consumer mobile devices, including configurations that contain multiple different types of image sensors. One popular configuration is to combine an RGB camera for color imaging with a monochrome camera that has improved performance in low-light settings, as well as some sensitivity in the infrared. In this work we introduce a method to combine simultaneously captured images from such a two-camera stereo system to generate a high-quality, noise reduced color image. To do so, pixel-to-pixel alignment has to be constructed between the two captured monochrome and color images, which however, is prone to artifacts due to parallax. The joint image reconstruction is made robust by introducing a novel artifact-robust optimization formulation. We provide extensive experimental results based on the two-camera configuration of a commercially available cell phone. This work was presented at the BMVC conference [18].

7.1.5. Noisy supervision for correcting misaligned cadaster maps without perfect Ground Truth data

Participants: Nicolas Girard, Yuliya Tarabalka.

In collaboration with Guillaume Charpiat (Tau Inria project-team).

In machine learning the best performance on a certain task is achieved by fully supervised methods when perfect ground truth labels are available. However, labels are often noisy, especially in remote sensing where manually curated public datasets are rare. We study the multi-modal cadaster map alignment problem for which available annotations are misaligned polygons, resulting in noisy supervision. We subsequently set up a multiple-rounds training scheme which corrects the ground truth annotations at each round to better train the model at the next round. We show that it is possible to reduce the noise of the dataset by iteratively training a better alignment model to correct the annotation alignment. This work was presented at the IGARSS conference [10].

7.1.6. Incremental Learning for Semantic Segmentation of Large-Scale Remote Sensing Data

Participants: Onur Tasar, Pierre Alliez, Yuliya Tarabalka.

In spite of remarkable success of the convolutional neural networks on semantic segmentation, they suffer from catastrophic shortcomings: a significant performance drop for the already learned classes when new classes are added on the data having no annotations for the old classes. We propose an incremental learning methodology, enabling to learn segmenting new classes without hindering dense labeling abilities for the previous classes, although the entire previous data are not accessible. The key points of the proposed approach are adapting the network to learn new as well as old classes on the new training data, and allowing it to remember the previously learned information for the old classes. For adaptation, we keep a frozen copy of the previously trained network, which is used as a memory for the updated network in absence of annotations for the former classes. The updated network minimizes a loss function, which balances the discrepancy between outputs for the previous classes from the memory and updated networks, and the mis-classification rate between outputs for the new classes from the updated network and the new ground-truth. We either regularly feed samples from the stored, small fraction of the previous data or use the memory network, depending on whether the new data are collected from completely different geographic areas or from the same city (see Figure 2). Our experimental results prove that it is possible to add new classes to the network, while maintaining its performance for the previous classes, despite the whole previous training data are not available. This work was published in the IEEE journal of Selected Topics in Applied Earth Observations and Remote Sensing [9].

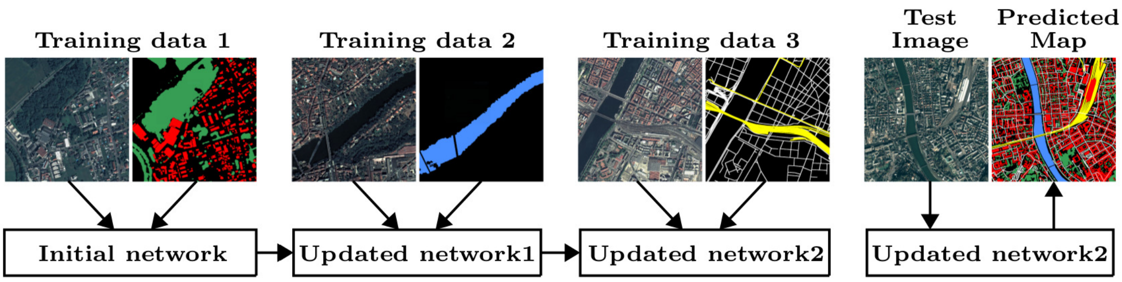

Figure 2. An example of an incremental learning scenario. Firstly, satellite images as well as their label maps for building and high vegetation classes are fed to the network. Then, from the second training data, the network learns the water class without forgetting building and high vegetation classes. Finally, road and railway classes are taught to the network. Whenever new training data are obtained, we store only a small part of the previous ones for the network to remember. When a new test image is provided, the network is able to detect all the classes.

7.1.7. Multi-Task Deep Learning for Satellite Image Pansharpening and Segmentation

Participants: Onur Tasar, Yuliya Tarabalka.

In collaboration with Andrew Khalel (Cairo University), Guillaume Charpiat (Inria, TAU team)

In this work, we propose a novel multi-task framework, to learn satellite image pansharpening and segmentation jointly (Figure 3). Our framework is based on the encoder-decoder architecture, where both tasks share the same encoder but each one has its own decoder. We compare our framework against single-task models with different architectures. Results show that our framework outperforms all other approaches in both tasks. This work was presented at the IGARSS conference [11].

7.1.8. A Generic Framework for Combining Multiple Segmentations in Geographic Object-Based Image Analysis

Participant: Onur Tasar.

In collaboration with Sébastien Lefèvre (Université Bretagne Sud, IRISA) and David Sheeren (DYNAFOR, University of Toulouse, INRA)

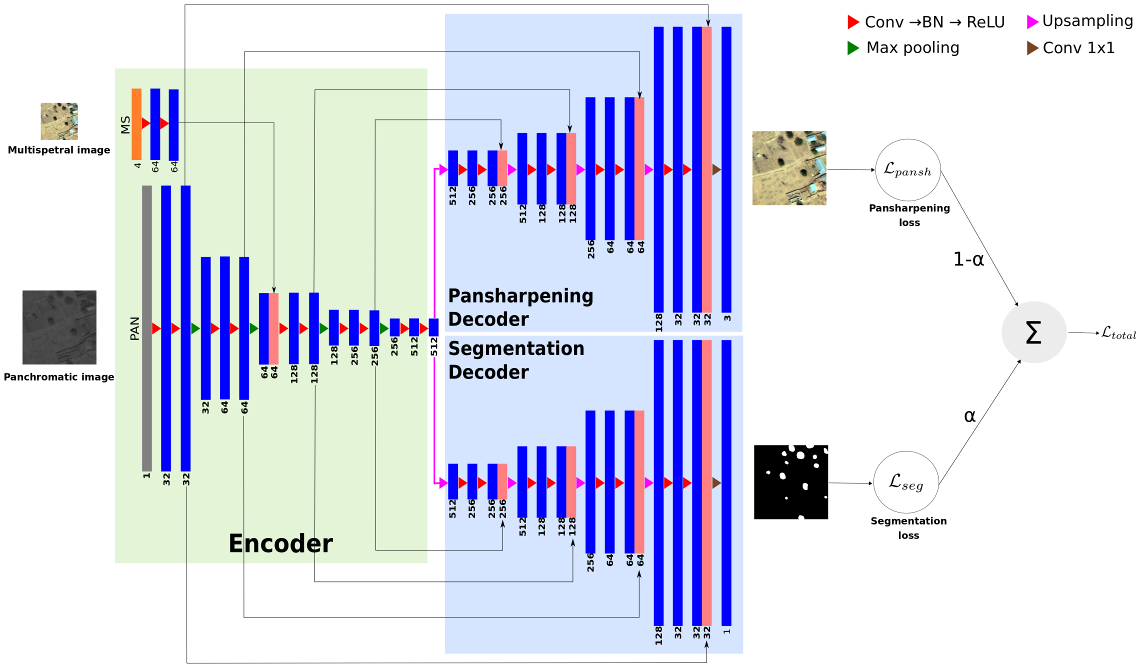

Figure 3. The overall pansharpening and segmentation framework.

The Geographic Object-Based Image Analysis (GEOBIA) paradigm relies strongly on the segmentation concept, i.e., partitioning of an image into regions or objects that are then further analyzed. Segmentation is a critical step, for which a wide range of methods, parameters and input data are available. To reduce the sensitivity of the GEOBIA process to the segmentation step, here we consider that a set of segmentation maps can be derived from remote sensing data. Inspired by the ensemble paradigm that combines multiple weak classifiers to build a strong one, we propose a novel framework for combining multiple segmentation maps (Figure 4). The combination leads to a fine-grained partition of segments (super-pixels) that is built by intersecting individual input partitions, and each segment is assigned a segmentation confidence score that relates directly to the local consensus between the different segmentation maps. Furthermore, each input segmentation can be assigned some local or global quality score based on expert assessment or automatic analysis. These scores are then taken into account when computing the confidence map that results from the combination of the segmentation processes. This means the process is less affected by incorrect segmentation inputs either at the local scale of a region, or at the global scale of a map. In contrast to related works, the proposed framework is fully generic and does not rely on specific input data to drive the combination process. We assess its relevance through experiments conducted on ISPRS 2D Semantic Labeling. Results show that the confidence map provides valuable information that can be produced when combining segmentations, and fusion at the object level is competitive w.r.t. fusion at the pixel or decision level. This work was published in the ISPRS journal of Geo-Information [8].

7.2. Reconstruction

7.2.1. City Reconstruction from Airborne Lidar: A Computational Geometry Approach

Participants: Jean-Philippe Bauchet, Florent Lafarge.

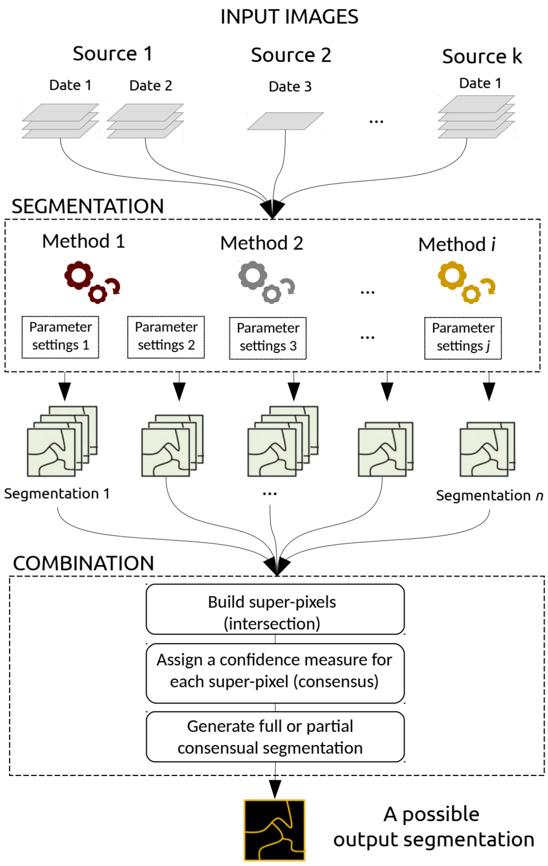

Figure 4. Our generic framework to combine multiple segmentations in the GEOBIA paradigm. Segmentations can come from different data sources (e.g., optical and radar sensors) and be acquired at different dates. They may also be produced using different methods (e.g., region-based or edge-based) relying on different parameter values.

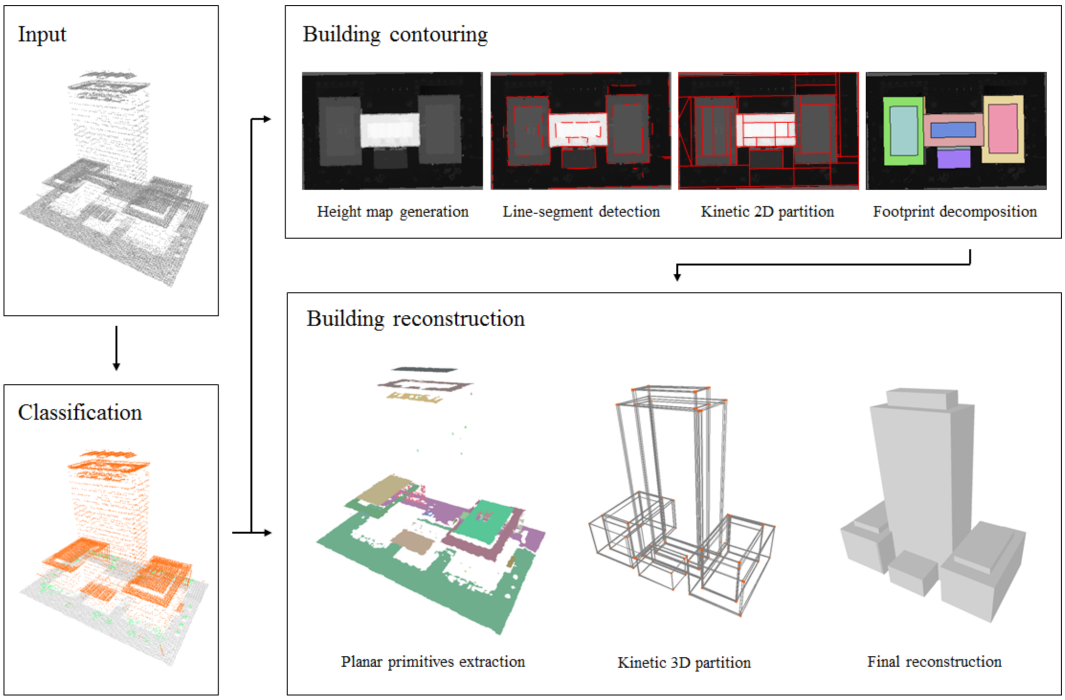

Figure 5. City Reconstruction from Airborne Lidar. Our method consists of three main steps. We first label points of the LiDAR scan as ground, vegetation or roof. Then, we apply a contouring algorithm to the height map, revealing the facades initially absent in the point set. Finally, we extract and propagate planar primitives from the point cloud, dividing the space into polyhedra that are labeled to obtain a 3D reconstruction of buildings.

We introduce a pipeline that reconstructs buildings of urban environments as concise polygonal meshes from airborne LiDAR scans. It consists of three main steps: classification, building contouring, and building reconstruction, the two last steps being achieved using computational geometry tools. Our algorithm demonstrates its robustness, flexibility and scalability by producing accurate and compact 3D models over large and varied urban areas in a few minutes only (See Figure 5). This work was published in the ISPRS international conference 3D GeoInfo [14].

7.2.2. Extracting geometric structures in images with Delaunay point processes

Participant: Florent Lafarge.

In collaboration with Jean-Dominique Favreau (Ekinnox), Adrien Bousseau (GraphDeco Inria team) and Alex Auvolat (Wide Inria team).

We introduce Delaunay Point Processes, a framework for the extraction of geometric structures from images. Our approach simultaneously locates and groups geometric primitives (line segments, triangles) to form extended structures (line networks, polygons) for a variety of image analysis tasks. Similarly to traditional point processes, our approach uses Markov Chain Monte Carlo to minimize an energy that balances fidelity to the input image data with geometric priors on the output structures. However, while existing point processes struggle to model structures composed of interconnected components, we propose to embed the point process into a Delaunay triangulation, which provides high-quality connectivity by construction. We further leverage key properties of the Delaunay triangulation to devise a fast Markov Chain Monte Carlo sampler. We demonstrate the flexibility of our approach on a variety of applications, including line network extraction, object contouring, and mesh-based image compression (See Figure 6). This work was published in the IEEE journal TPAMI [7].

Figure 6. Example applications of Delaunay Point Processes to extract planar graphs representing blood vessels in retina images (left), and complex polygons representing object silhouettes (right). The point distribution creates a dynamic Delaunay triangulation while edge and facet labels specify the geometric structure (see red edges on close-ups).

7.3. Approximation

7.3.1. Cost-driven framework for progressive compression of textured meshes

Participants: Cédric Portaneri, Pierre Alliez.

In collaboration with Michael Hemmer (Google X), Lukas Birklein and Elmar Schoemer (Uni. of Mainz).

Recent advances in digitization of geometry and radiometry generate in routine massive amounts of surface meshes with texture or color attributes. This large amount of data can be compressed using a progressive approach which provides at decoding low complexity levels of details (LoDs) that are continuously refined until retrieving the original model. The goal of such a progressive mesh compression algorithm is to improve the overall quality of the transmission for the user, by optimizing the rate-distortion trade-off. In this paper,

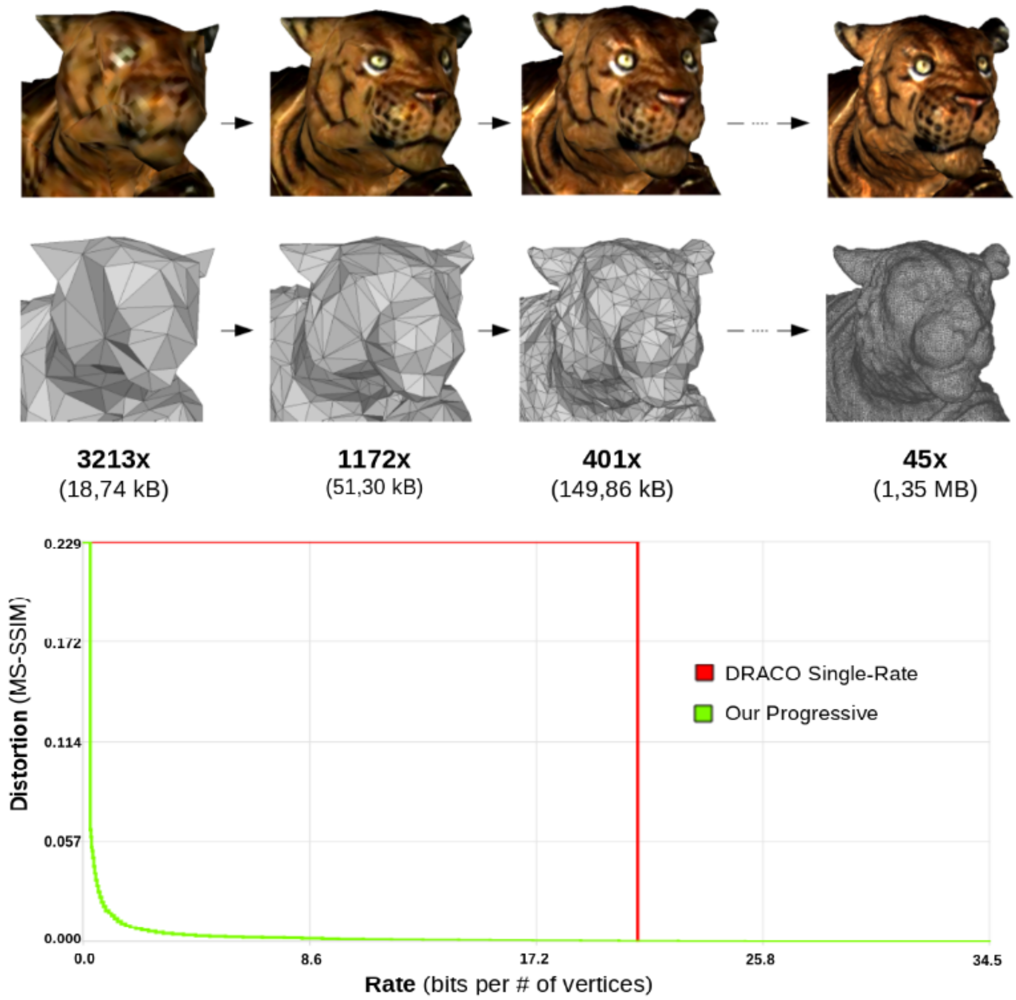

Figure 7. Progressive decompression of a textured surface triangle mesh. Top: key levels of detail with their size and compression rate compared to the raw obj file (texture data not included). Bottom: Distortion against the bit consumption, in bits per vertex, where the number of vertices refers to the input mesh. Green is our progressive approach, red is the state-of-the-art single-rate DRACO encoder.

we introduce a novel meaningful measure for the cost of a progressive transmission of a textured mesh by observing that the rate-distortion curve is in fact a staircase, which enables an effective comparison and optimization of progressive transmissions in the first place. We contribute a novel generic framework which utilizes the cost function to encode triangle surface meshes via multiplexing several geometry reduction steps (mesh decimation via half-edge or full-edge collapse operators, xyz quantization reduction and uv quantization reduction). This framework can also deal with textures by multiplexing an additional texture reduction step. We also design a texture atlas that enables us to preserve texture seams during decimation while not impairing the quality of resulting LODs. For encoding the inverse mesh decimation steps we further contribute a significant improvement over the state-of-the-art in terms of rate-distortion performance and yields a compression-rate of 22:1, on average. Finally, we propose a unique single-rate alternative solution using a selection scheme of a subset among LODs, optimized for our cost function, and provided with our atlas that enables interleaved progressive texture refinements (see Figure 7). This work was presented at the ACM Multimedia Systems conference [19] and obtained the best paper award.

7.3.2. Selective padding for Polycube-based hexahedral meshing

Participant: Pierre Alliez.

In collaboration with Gianmarco Cherchi and Riccardo Scateni from University of Cagliari (Sardinia), Max Lyon from University of Aachen and David Bommes from University of Bern.

Hexahedral meshes generated from polycube mapping often exhibit a low number of singularities but also poor quality elements located near the surface. It is thus necessary to improve the overall mesh quality, in terms of the minimum Scaled Jacobian (MSJ) or average Scaled Jacobian (ASJ). Improving the quality may be obtained via global padding (or pillowing), which pushes the singularities inside by adding an extra layer of hexahedra on the entire domain boundary. Such a global padding operation suffers from a large increase of complexity, with unnecessary hexahedra added. In addition, the quality of elements near the boundary may decrease. We propose a novel optimization method which inserts sheets of hexahedra so as to perform selective padding, where it is most needed for improving the mesh quality. A sheet can pad part of the domain boundary, traverse the domain and form singularities. Our global formulation, based on solving a binary problem, enables us to control the balance between quality improvement, increase of complexity and number of singularities. We show in a series of experiments that our approach increases the MSJ value and preserves (or even improves) the ASJ, while adding fewer hexahedra than global padding. (See Figure 8). This work was published in an international journal and was presented at the EUROGRAPHICS conference [4].

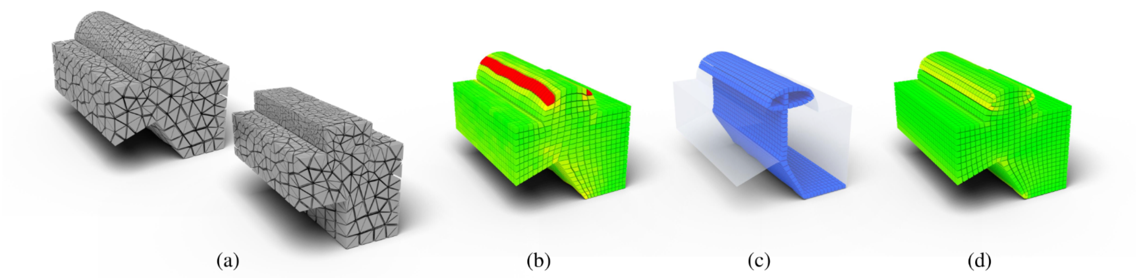

Figure 8. Polycube-based hexahedral meshing. Our pipeline takes as input a model and its polycube mapping (a); we compute the relative hex-mesh and locate the surface areas in need of padding analyzing the mapping quality (b); we set and solve a binary problem to find a set of facets to extrude in order to create a selective padding layer (c); we compute and analyze the mapping with the new hex-mesh structure (d).

8. Bilateral Contracts and Grants with Industry

8.1. Bilateral Contracts with Industry

8.1.1. Google X

Participants: Cédric Portaneri, Pierre Alliez.

We developed a novel approach and software prototype for the compression of 3D models. Our main focus is on progressive compression of surface triangle meshes with color textures, with emphasis on fine grain, genericity and flexible metric. The proposed methodology is to turn the input models into a stream of refinements, in which both mesh and texture refinement details are multiplexed in accordance to rate-distortion principles. Fine grain control is achieved through considering all components, local as well as non-local, from both the mesh and its textures: mesh complexity, vertex accuracy, texture definition and accuracy. We leveraged the recent advances on perceptual metrics to improve the visual appearance, and performed joint consolidation and encoding of the models to further optimize the rate-distortion tradeoffs and visual perception.

8.1.2. Dorea technology

Participants: Vincent Vadez, Pierre Alliez.

In collaboration with SME Dorea Technology, our objective is to advance the knowledge on the radiative thermal simulation of satellites, via geometric model reduction. The survival of a satellite is related to the temperature of its components, the variation of which must be controlled within safety intervals. In this context, the thermal simulation of the satellite for its design is crucial to anticipate the reality of its operation. This CIFRE project started in August 2018, for a total duration of 3 years.

8.1.3. Luxcarta

Participants: Jean-Philippe Bauchet, Florent Lafarge.

The goal of this collaboration is to design automated approaches for producing city models from the last generation of satellites. In particular, this project investigates geometric representations for images and 3D data that are more compact and meaningful than traditional pixel and voxel grids, the intuition being to synthesize massive satellite data to reconstruct objects in 3D in a more scalable manner than existing methods. This CIFRE project started in October 2016, for a total duration of 3 years.

8.1.4. CNES and Acri-ST

Participants: Onur Tasar, Yuliya Tarabalka, Pierre Alliez.

The aim is to devise efficient representations for satellite images. The project started in October 2017, for a total duration of 3 years.

8.1.5. CSTB

Participants: Hao Fang, Mulin Yu, Florent Lafarge.

This collaboration takes the form of two independent contracts. The first project investigated the automatic conversion of raw 3D data to polyhedral surfaces that approximate man-made objects at some key structural representation scales. This project started in March 2016, for a total duration of 3 years. The second project investigates the design of as-automatic-as-possible algorithms for repairing and converting Building Information Modeling (BIM) models of buildings in different urban-specific CAD formats using combinatorial maps. This project started November 2019, for a total duration of 3 years.

8.1.6. IRT Saint-Exupéry

Participants: Gaetan Bahl, Florent Lafarge.

This project investigates low-power deep learning architectures for detecting, localizing and characterizing changes in temporal satellite images. These architectures are designed to be exploited on-board satellites with low computational resources. The project started in March 2019, for a total duration of 3 years.

8.1.7. Dassault Systèmes

Participants: Julien Vuillamy, Pierre Alliez, Florent Lafarge.

This project investigates algorithms for reconstructing city models from multi-sourced data. 3D objects are reconstructed by filtering, parsing and assembling planar shapes. The project started in April 2018, for a total duration of 3 years.

9. Partnerships and Cooperations

9.1. National Initiatives

9.1.1. ANR

9.1.1.1. PISCO: Perceptual Levels of Detail for Interactive and Immersive Remote Visualization of Complex 3D Scenes

Participants: Pierre Alliez [contact], Flora Quilichini, Florent Lafarge.

The way of consuming and visualizing this 3D content is evolving from standard screens to Virtual and Mixed Reality (VR/MR). Our objective is to devise novel algorithms and tools allowing interactive visualization, in these constrained contexts (Virtual and Mixed reality, with local/remote 3D content), with a high quality of user experience. Partners: Inria, LIRIS INSA Lyon Institut National des Sciences Appiquées (coordinator), Laboratoire d'Informatique en Images et Systèmes d'Information LS2N Nantes University. Total budget 550 KE, 121 KE for TITANE. The PhD thesis of Flora Quilichini is funded by this project which started in January 2018, for a total duration of 4 years.

9.1.1.2. LOCA-3D: Localization Orientation and 3D CArtography

Participants: Fernando Ireta Munoz, Florent Lafarge, Pierre Alliez [contact].

This project is part of the ANR Challenge MALIN LOCA-3D (Localization, orientation and 3D cartography). The challenge is to develop and experiment accurate location solutions for emergency intervention officers and security forces. These solutions must be efficient inside buildings and in conditions where satellite positioning systems do not work satisfactorily. Our solution is based on an advanced inertial system, where part of the inertial sensor drift is compensated by a vision system. Partners: SME INNODURA TB (coordinator), IBISC laboratory (Evry university) and Inria. Total budget: 700 KE, 157 KE for TITANE. The engineer position of Fernando Ireta Munoz is funded by this project which started in January 2018, for a total duration of 4 years.

9.1.1.3. EPITOME: efficient representation to structure large-scale satellite images

Participants: Nicolas Girard, Yuliya Tarabalka [PI].

The goal of this young researcher project is to devise an efficient multi-scale vectorial representation, which would structure the content of large-scale satellite images. More specifically, we seek for a novel effective representation for large-scale satellite images, that would be generic, i.e., applicable for images worldwide and for a wide range of applications, and structure-preserving, i.e. best representing the meaningful objects in the image scene. To address this challenge, we plan to bridge the gap between advanced machine learning and geometric modeling tools to devise a multi-resolution vector-based representation, together with the methods for its effective generation and manipulation. Total budget: 225 KE for TITANE. The PhD thesis of Nicolas Girard is funded by this project which started in October 2017, for a total duration of 4 years.

9.1.1.4. Faults_R_GEMS: Properties of FAULTS, a key to Realistic Generic Earthquake Modeling and hazard Simulation

Participants: Lionel Matteo, Yuliya Tarabalka [contact].

The goal of the project is to study the properties of seismic faults, using advanced math tools including learning approaches. The project is in collaboration with Geoazur lab (coordinator), Arizona State University, CALTECH, Ecole Centrale Paris, ENS Paris, ETH Zurich, Geosciences Montpellier, IFSTTAR, IPGP Paris, IRSN Fontenay-aux-Roses, LJAD Nice, UNAVCO Colorado and Pisa University. The PhD thesis of Lionel Matteo is funded by this project which started in October 2017, for a total duration of 4 years.

9.1.1.5. BIOM: Building Indoor and Outdoor Modeling

Participants: Muxingzi Li, Pierre Alliez, Florent Lafarge [contact].

The BIOM project aims at automatic, simultaneous indoor and outdoor modelling of buildings from images and dense point clouds. We want to achieve a complete, geometrically accurate, semantically annotated but nonetheless lean 3D CAD representation of buildings and objects they contain in the form of a Building Information Models (BIM) that will help manage buildings in all their life cycle (renovation, simulation, deconstruction). The project is in collaboration with IGN (coordinator), Ecole des Ponts Paristech, CSTB and INSA-ICube. Total budget: 723 KE, 150 KE for TITANE. The PhD thesis of Muxingzi Li is funded by this project which started in February 2018, for a total duration of 4 years.

9.2. International Initiatives

9.2.1. Inria International Partners

9.2.1.1. Declared Inria International Partners

We collaborated with David Bommes from Bern University (Switzerland), Gianmarco Cherchi and Riccardo Scateni from University of Cagliary (Sardinia), and Elmar Schoemer from Johannes Gutenberg Universität Mainz.

9.3. International Research Visitors

9.3.1. Visits of International Scientists

- Michael Hemmer, research engineer at Google X, visited us in June.
- Jorg Peters, Professor at University of Florida, visited us in June.

9.3.2. Visits to International Teams

- Pierre Alliez visited the Google X team for one week in April.
- Florent Lafarge visited the Institute of Computer Graphics and Vision at TU Graz in March.

10. Dissemination

10.1. Promoting Scientific Activities

10.1.1. Scientific Events: Organisation

10.1.1.1. General Chair, Scientific Chair

Pierre Alliez was program co-chair of the EUROGRAPHICS conference and of the Symposium on Solid and Physical Modeling (SPM).

10.1.2. Scientific Events: Selection

10.1.2.1. Member of the Conference Program Committees

- Pierre Alliez: member of the advisory board for EUROGRAPHICS 2020.
- Florent Lafarge: EUROGRAPHICS
- Yuliya Tarabalka: ICCV

10.1.2.2. Reviewer

- Pierre Alliez: ACM SIGGRAPH, SIGGRAPH Asia, Symposium on Geometry processing.
- Florent Lafarge: CVPR, ICCV, SIGGRAPH.

10.1.3. Journal

10.1.3.1. Member of the Editorial Boards

- Pierre Alliez: associate editor of the Computer Graphics Forum, Computer-Aided Geometric Design and Computer-Aided Design. He is also a member of the editorial board of the CGAL open source project.
- Florent Lafarge: associate editor of The Visual Computer and the Revue Française de Photogrammétrie et de Télédétection, and member of the Editorial Advisory Board of the ISPRS Journal of Photogrammetry and Remote Sensing

10.1.3.2. Reviewer - Reviewing Activities

- Pierre Alliez: reviewer for Computer Graphics Forum and ACM Transactions on Graphics.
- Florent Lafarge: reviewer for the ISPRS Journal of Photogrammetry and Remote Sensing and the International Journal of Computer Vision.

10.1.4. Invited Talks

Pierre Alliez gave invited talks at Cambridge university, Google X and UCLA.

10.1.5. Leadership within the Scientific Community

Pierre Alliez is a member of the Steering Committees of the EUROGRAPHICS Symposium on Geometry Processing, EUROGRAPHICS Workshop on Graphics and Cultural Heritage and Executive Board Member for the Solid Modeling Association.

10.1.6. Scientific Expertise

Pierre Alliez was a reviewer for the European commission and ANR, and is a scientific advisory board member for the Bézout Labex in Paris (Models and algorithms: from the discrete to the continuous).

10.1.7. Research Administration

Pierre Alliez is Head of Science of the Inria Sophia Antipolis center and a member of the scientific committee of the 3IA Côte d'Azur.

10.2. Teaching - Supervision - Juries

10.2.1. Teaching

Master: Pierre Alliez, Florent Lafarge and Gaétan Bahl, advanced machine learning, 21h, M2, university Nice Sophia Antipolis, France.

Master: Pierre Alliez, data visualization, 3h, M2, university Côte d'Azur, France.

Master: Pierre Alliez and Florent Lafarge, 3D Meshes and Applications, 32h, M2, Ecole des Ponts ParisTech, France.

Master: Florent Lafarge, Traitement d'images numériques, 6h, M2, university Nice Sophia Antipolis, France.

Master: Pierre Alliez and Florent Lafarge, Interpolation numérique, 60h, M1, university Nice Sophia Antipolis, France.

Master: Florent Lafarge, Mathématiques pour la géométrie, 20h, M1, EFREI, France.

Master: Yuliya Tarabalka, Discrete inference and learning, 12h, M2 MVA, ENS Paris-Saclay & CentraleSupelec, France.

10.2.2. Supervision

PhD defended January 16th: Hao Fang, Geometric modeling of man-made objects at different level of details, since March 2016, Florent Lafarge.

PhD defended December 6th: Jean-Philippe Bauchet, Kinetic data structures for the geometric modeling of urban environments, since October 2016, Florent Lafarge.

PhD in progress: Julien Vuillamy, city reconstruction from multi-sourced data, since April 2018, Pierre Alliez and Florent Lafarge.

PhD in progress: Muxingzi Li, indoor reconstruction from a smartphone, since February 2018, Florent Lafarge.

PhD in progress: Lionel Matteo: From Pleiades images to very high resolution topography in complex zones, since September 2017, Yuliya Tarabalka and Isabelle Manighetti.

PhD in progress: Onur Tasar, Using deep learning approaches to devise an efficient representation for large-scale satellite images, since October 2017, Yuliya Tarabalka and Pierre Alliez.

PhD in progress: Nicolas Girard, How to structure satellite data, since November 2017, Yuliya Tarabalka.

PhD in progress: Vincent Vadez, Geometric simplification of satellites for thermal simulation, since August 2018, Pierre Alliez.

PhD in progress: Gaétan Bahl, low-power neural networks, since March 2019, Florent Lafarge.

PhD in progress: Mulin Yu, remeshing urban-specific CAD formats, since November 2019, Florent Lafarge.

PhD in progress: Tong Zhao, shape reconstruction, since November 2019, Pierre Alliez and Laurent Busé (from the Aromath Inria project-team).

PhD stopped after one year: Flora Quilichini, Geometry Compression, between January 2018 and June 2019, Pierre Alliez and Guillaume Lavoué (INSA Lyon).

10.2.3. Juries

- Pierre Alliez was a PhD committee member for Siargey Kachanovich and Jean-Philippe Bauchet (Inria Sophia Antipolis). He was a reviewer for the habilitation committee of Sébastien Valette (INSA Lyon).
- Florent Lafarge was a PhD thesis reviewer for Romain Rombourg (University of Grenoble Alpes) and Thomas Holzmann (TU Graz, Austria).

10.3. Popularization

10.3.1. Interventions

- Pierre Alliez gave a talk on geometric modeling in a high school in Cannes, April 26th.

11. Bibliography

Publications of the year

Doctoral Dissertations and Habilitation Theses

[1] J.-P. BAUCHET. *Kinetic data structures for the geometric modeling of urban environments*, Université Côte d'Azur, Inria, France, December 2019, https://hal.inria.fr/tel-02432386

[2] H. FANG. *Geometric modeling of man-made objects at different level of details*, Université Côte d'Azur, January 2019, https://tel.archives-ouvertes.fr/tel-02406834

Articles in International Peer-Reviewed Journals

[3] P. ALLIEZ, R. DI COSMO, B. GUEDJ, A. GIRAULT, M.-S. HACID, A. LEGRAND, N. P. ROUGIER. *Attributing and Referencing (Research) Software: Best Practices and Outlook from Inria*, in "Computing in Science & Engineering", 2019, pp. 1-14, https://arxiv.org/abs/1905.11123 [*DOI :* 10.1109/MCSE.2019.2949413], https://hal.archives-ouvertes.fr/hal-02135891

[4] G. CHERCHI, P. ALLIEZ, R. SCATENI, M. LYON, D. BOMMES. *Selective Padding for Polycube-Based Hexahedral Meshing*, in "Computer Graphics Forum", January 2019 [*DOI :* 10.1111/CGF.13593], https://hal.inria.fr/hal-01970790

[5] O. ENNAFII, A. LE BRIS, F. LAFARGE, C. MALLET. *A learning approach to evaluate the quality of 3D city models*, in "Photogrammetric engineering and remote sensing", 2019, forthcoming, https://hal.archives-ouvertes.fr/hal-02193116

[6] H. FANG, F. LAFARGE. *Pyramid Scene Parsing Network in 3D: improving semantic segmentation of point clouds with multi-scale contextual information*, in "ISPRS Journal of Photogrammetry and Remote Sensing", 2019, forthcoming, https://hal.inria.fr/hal-02159279

[7] J.-D. FAVREAU, F. LAFARGE, A. BOUSSEAU, A. AUVOLAT. *Extracting Geometric Structures in Images with Delaunay Point Processes*, in "IEEE Transactions on Pattern Analysis and Machine Intelligence", 2019 [*DOI :* 10.1109/TPAMI.2018.2890586], https://hal.inria.fr/hal-01950791

[8] S. LEFÈVRE, D. SHEEREN, O. TASAR. *A Generic Framework for Combining Multiple Segmentations in Geographic Object-Based Image Analysis*, in "ISPRS International Journal of Geo-Information", 2019, vol. 8, n^{o} 2 [*DOI :* 10.3390/IJGI8020070], https://hal.archives-ouvertes.fr/hal-02083265

[9] O. TASAR, Y. TARABALKA, P. ALLIEZ. *Incremental Learning for Semantic Segmentation of Large-Scale Remote Sensing Data*, in "IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing", September 2019, vol. 12, n^{o} 9, pp. 3524-3537 [*DOI :* 10.1109/JSTARS.2019.2925416], https://hal.inria.fr/hal-01909830

Invited Conferences

[10] N. GIRARD, G. CHARPIAT, Y. TARABALKA. *Noisy Supervision for Correcting Misaligned Cadaster Maps Without Perfect Ground Truth Data*, in "IGARSS 2019 - IEEE International Geoscience and Remote Sensing Symposium", Yokohama, Japan, July 2019, https://arxiv.org/abs/1903.06529 , https://hal.inria.fr/hal-02065211

[11] A. KHALEL, O. TASAR, G. CHARPIAT, Y. TARABALKA. *Multi-Task Deep Learning for Satellite Image Pansharpening and Segmentation*, in "IGARSS 2019 - IEEE International Geoscience and Remote Sensing Symposium", Yokohama, Japan, July 2019, https://hal.inria.fr/hal-02276549

[12] O. TASAR, Y. TARABALKA, P. ALLIEZ. *Continual Learning for Dense Labeling of Satellite Images*, in "IGARSS 2019 - IEEE International Geoscience and Remote Sensing Symposium", Yokohama, Japan, July 2019, https://hal.inria.fr/hal-02276543

International Conferences with Proceedings

[13] G. BAHL, L. DANIEL, M. MORETTI, F. LAFARGE. *Low-power neural networks for semantic segmentation of satellite images*, in "ICCV Workshop on Low-Power Computer Vision", Seoul, South Korea, October 2019, https://hal.inria.fr/hal-02277061

[14] J.-P. BAUCHET, F. LAFARGE. *City Reconstruction from Airborne Lidar: A Computational Geometry Approach*, in "3D GeoInfo 2019 - 14thConference 3D GeoInfo", Singapore, Singapore, September 2019, https://hal.inria.fr/hal-02295768

[15] G. CHARPIAT, N. GIRARD, L. FELARDOS, Y. TARABALKA. *Input Similarity from the Neural Network Perspective*, in "Neural Information Processing Systems (NeurIPS)", Vancouver, Canada, December 2019, https://hal.inria.fr/hal-02394647

[16] O. ENNAFII, A. LE BRIS, F. LAFARGE, C. MALLET. *Scalable Evaluation of 3D City Models*, in "IGARSS 2019 - IEEE International Geoscience and Remote Sensing Symposium", Yokohama, Japan, July 2019, https://hal.inria.fr/hal-02165557

[17] O. ENNAFII, C. MALLET, A. LE BRIS, F. LAFARGE. *The necessary yet complex evaluation of 3D city models: a semantic approach*, in "JURSE 2019 - Joint Urban Remote Sensing Event", Vannes, France, May 2019, https://hal.inria.fr/hal-02165562

[18] M. LI, P. TU, W. HEIDRICH. *Robust Joint Image Reconstruction from Color and Monochrome Cameras*, in "British Machine Vision Conference", Cardiff, United Kingdom, Proceedings of the British Machine Vision Conference, September 2019, To appear in the proceedings, https://hal.inria.fr/hal-02370233

[19] C. PORTANERI, P. ALLIEZ, M. HEMMER, L. BIRKLEIN, E. SCHOEMER. *Cost-driven framework for progressive compression of textured meshes*, in "ACM 2019 - 10th Multimedia Systems Conference", Amherst, United States, ACM Press, June 2019, pp. 175-188 [*DOI :* 10.1145/3304109.3306225], https://hal.inria.fr/hal-02273929

Conferences without Proceedings

[20] V. VADEZ, F. BRUNETTI, P. ALLIEZ. *First steps of numerical simulation using Artificial Intelligence*, in "European Space Thermal Engineering Workshop 2019", Amsterdam, Netherlands, October 2019, https://hal.inria.fr/hal-02367540

Research Reports

[21] D. COHEN-STEINER, A. LIEUTIER, J. VUILLAMY. *Lexicographic optimal chains and manifold triangulations*, Inria Sophia-Antipolis Méditerranée, December 2019, https://hal.archives-ouvertes.fr/hal-02391190

Other Publications

[22] D. COHEN-STEINER, A. LIEUTIER, J. VUILLAMY. *Lexicographic optimal homologous chains and applications to point cloud triangulations*, December 2019, working paper or preprint, https://hal.archives-ouvertes.fr/hal-02391240

[23] D. COHEN-STEINER, A. LIEUTIER, J. VUILLAMY. *Regular triangulations as lexicographic optimal chains*, December 2019, working paper or preprint, https://hal.archives-ouvertes.fr/hal-02391285

[24] O. TASAR, S. L. HAPPY, Y. TARABALKA, P. ALLIEZ. *ColorMapGAN: Unsupervised Domain Adaptation for Semantic Segmentation Using Color Mapping Generative Adversarial Networks*, November 2019, working paper or preprint, https://hal.inria.fr/hal-02382155

[25] O. TASAR, Y. TARABALKA, P. ALLIEZ, A. GIROS, S. CLERC. *From Satellite Images to Vector Representations*, November 2019, JC2 - 14èmes Journées CNES Jeunes Chercheurs, Poster, https://hal.inria.fr/hal-02367545