

RESEARCH CENTRE

Sophia Antipolis - Méditerranée

2021

ACTIVITY REPORT

Project-Team

TITANE

Geometric Modeling of 3D Environments

DOMAIN

Perception, Cognition and Interaction

THEME

Interaction and visualization

Contents

Project-Team TITANE	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	3
2.1 General Presentation	3
3 Research program	3
3.1 Context	3
3.2 Analysis	4
3.3 Approximation	4
3.4 Reconstruction	5
4 Application domains	5
5 Social and environmental responsibility	6
5.1 Impact of research results	6
6 Highlights of the year	6
6.1 Awards	6
7 New software and platforms	6
7.1 New software	6
7.1.1 Cifre PhD Thesis Titane - DS	6
7.1.2 CGAL - Orthtree	7
7.1.3 CGAL - 3D Alpha Wrapping	7
7.1.4 CGAL - NURBS meshing	8
7.2 New platforms	8
8 New results	8
8.1 Analysis	8
8.1.1 Binary Graph Neural Networks	8
8.1.2 Planar Shape Based Registration for Multi-modal Geometry	9
8.1.3 Polygonal Building Segmentation by Frame Field Learning	10
8.1.4 An efficient representation of 3D buildings: application to the evaluation of city models	11
8.1.5 Automatic fault mapping in remote optical images and topographic data with deep learning	12
8.1.6 Geometric approximation of structured scenes from images	14
8.1.7 Wearable Cooperative SLAM System for Real-time Indoor Localization Under Challenging Conditions	14
8.2 Reconstruction	15
8.2.1 Floorplan generation from 3D point clouds: A space partitioning approach	15
8.2.2 Progressive discrete domains for implicit surface reconstruction	16
8.2.3 Efficient open surface reconstruction from lexicographic optimal chains and critical bases	17
8.2.4 Planimetric simplification and lexicographic optimal chains for 3D urban scene reconstruction	18
8.3 Approximation	19
8.3.1 Delaunay Meshing and Repairing of NURBS Models	19
8.3.2 Alpha wrapping with an offset	20
8.3.3 Geometric Model Reduction Driven by Numerical Simulation Accuracy	20
9 Bilateral contracts and grants with industry	22
9.1 Bilateral contracts with industry	22

10 Partnerships and cooperations	24
10.1 European initiatives	24
10.1.1 FP7 & H2020 projects	24
10.2 National initiatives	24
10.2.1 ANR	24
11 Dissemination	25
11.1 Promoting scientific activities	26
11.1.1 Scientific events: organisation	26
11.1.2 Scientific events: selection	26
11.1.3 Journal	26
11.1.4 Invited talks	26
11.1.5 Leadership within the scientific community	26
11.1.6 Scientific expertise	27
11.1.7 Research administration	27
11.2 Teaching - Supervision - Juries	27
11.2.1 Teaching	27
11.2.2 Supervision	27
11.2.3 Juries	28
12 Scientific production	28
12.1 Major publications	28
12.2 Publications of the year	29

Project-Team TITANE

Creation of the Project-Team: 2014 January 01

Keywords

Computer sciences and digital sciences

- A5.3. – Image processing and analysis
- A5.3.2. – Sparse modeling and image representation
- A5.3.3. – Pattern recognition
- A5.5.1. – Geometrical modeling
- 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 members, visitors, external collaborators

Research Scientists

- Pierre Alliez [Team leader, Inria, Senior Researcher, HDR]
- Florent Lafarge [Inria, Senior Researcher, HDR]

Post-Doctoral Fellows

- Shenlu Jiang [Inria, until Mar 2021]
- Johann Lussange [Inria]
- Xiao Xiao [Inria, until Oct 2021]

PhD Students

- Gaetan Bahl [IRT Saint Exupéry]
- Rao Fu [Geometry Factory]
- Muxingzi Li [Inria, until Aug 2021]
- Vincent Vadez [Dorea Technology]
- Julien Vuillamy [Dassault Systèmes, until Aug 2021]
- Mulin Yu [Inria]
- Tong Zhao [Inria]
- Alexandre Zoppis [Univ Côte d'Azur, from Oct 2021]

Technical Staff

- Lucas Dubouchet [Inria, Engineer, until Apr 2021]
- Fernando Ireta Munoz [Inria, Engineer, until Mar 2021]
- Kacper Pluta [Inria, Engineer, from Jul 2021]
- Cédric Portaneri [Inria, Engineer]

Interns and Apprentices

- Jackson Campolattaro [Inria, from June 6 to August 31]
- Paul Vinh Le [Inria, until Feb 2021]
- Nissim Maruani [Inria, from Mar 2021 until Jul 2021]

Administrative Assistant

- Florence Barbara [Inria]

Visiting Scientist

- Benoit Morisset [Ai Verse, until Sep 2021]

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 subtending 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

In addition to tackling enduring scientific challenges, our research on geometric modeling and processing is motivated by applications to computational engineering, reverse engineering, digital mapping and urban planning. The main deliverable of our research will be algorithms with theoretical foundations. Ultimately, we 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). 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-off company 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. theses and postdoctoral fellows on topics that match our research program.

5 Social and environmental responsibility

5.1 Impact of research results

We are collaborating with Ekinnox, an Inria spin-off which develops human movement analysis solutions for healthcare institutions. We provide scientific advice on automated analysis of human walking, in collaboration with Laurent Busé from the Aromath project-team.

6 Highlights of the year

Titane has hosted the AI Verse startup project led by Benoit Morisset, focused on the generation of 3D training data for deep learning. AI Verse has now secured enough funding to mature the technology into an industrial product.

Pierre Alliez has been promoted editor in chief of the Computer Graphics Forum, one of the leading journals for technical articles on computer graphics. The journal features a mix of original research, computer graphics applications, conference reports, state-of-the-art surveys and workshops.

6.1 Awards

Muxingzi Li received the best PhD thesis award from doctoral school EDSTIC (accessit prize).

7 New software and platforms

We intensified our contributions to the CGAL library, with three novel components for meshing and repairing NURBS surfaces, computing orthotrees (quadrees in 2D, octrees in 3D, etc) and for wrapping sets of finite 3D primitives with a valid surface triangle mesh. Valid herein translates into watertight, intersection-free and combinatorially 2-manifold.

7.1 New software

7.1.1 Cifre PhD Thesis Titane - DS

Name: Software for Cifre PhD thesis Titane - Dassault Systemes

Keyword: Geometric computing

Functional Description: This software addresses the problem of simplifying two-dimensional polygonal partitions that exhibit strong regularities. Such partitions are relevant for reconstructing urban scenes in a concise way. Preserving long linear structures spanning several partition cells motivates a point-line projective duality approach in which points represent line intersections and lines possibly carry multiple points. We implement a simplification algorithm that seeks a balance between the fidelity to the input partition, the enforcement of canonical relationships between lines (orthogonality or parallelism) and a low complexity output. Our methodology alternates continuous optimization by Riemannian gradient descent with combinatorial reduction, resulting in a progressive simplification scheme. Our experiments show that preserving canonical relationships helps gracefully degrade partitions of urban scenes, and yields more concise and regularity-preserving meshes than previous mesh-based simplification approaches.

Contact: Pierre Alliez

Participants: Pierre Alliez, Florent Lafarge

Partner: Dassault Systèmes

7.1.2 CGAL - Orthtree

Name: Quadtrees, Octrees, and Orthtrees

Keyword: Octree/Quadtree

Functional Description: Quadtrees are tree data structures in which each node encloses a square section of space, and each internal node has exactly 4 children. Octrees are a similar data structure in 3D in which each node encloses a cubic section of space, and each internal node has exactly 8 children. We call the generalization of such data structure "orthtrees", as orthants are generalizations of quadrants and octants. The term "hyperoctree" can also be found in literature to name such data structures in dimensions 4 and higher.

This package provides a general data structure Orthtree along with aliases for Quadtree and Octree. These trees can be constructed with custom point ranges and split predicates, and iterated on with various traversal methods.

Release Contributions: Initial version

Contact: Pierre Alliez

Participants: Pierre Alliez, Jackson Campolattaro, Simon Giraudot

Partner: GeometryFactory

7.1.3 CGAL - 3D Alpha Wrapping

Keyword: Shrink-wrapping

Scientific Description: Various tasks in geometric modeling and processing require 3D objects represented as valid surface meshes, where "valid" refers to meshes that are watertight, intersection-free, orientable, and combinatorially 2-manifold. Such representations offer well-defined notions of interior/ exterior and geodesic neighborhoods.

3D data are usually acquired through measurements followed by reconstruction, designed by humans, or generated through imperfect automated processes. As a result, they can exhibit a wide variety of defects including gaps, missing data, self-intersections, degeneracies such as zero-volume structures, and non-manifold features.

Given the large repertoire of possible defects, many methods and data structures have been proposed to repair specific defects, usually with the goal of guaranteeing specific properties in the repaired 3D model. Reliably repairing all types of defects is notoriously difficult and is often an ill-posed problem as many valid solutions exist for a given 3D model with defects. In addition, the input model can be overly complex with unnecessary geometric details, spurious topological structures, nonessential inner components, or excessively fine discretizations. For applications such as collision avoidance, path planning, or simulation, getting an approximation of the input can be more relevant than repairing it. Approximation herein refers to an approach capable of filtering out inner structures, fine details and cavities, as well as wrapping the input within a user-defined offset margin.

Given an input 3D geometry, we address the problem of computing a conservative approximation, where conservative means guaranteeing a strictly enclosed input. We seek unconditional robustness in the sense that the output mesh should be valid (oriented, combinatorially 2-manifold and without self-intersections), even for raw input with many defects and degeneracies. The default input is a soup of 3D triangles, but the generic interface leaves the door open to other types of finite 3D primitives.

Functional Description: The algorithm proceeds by shrink-wrapping and refining a 3D Delaunay triangulation loosely bounding the input. Two user-defined parameters, alpha and offset, offer control over the maximum size of cavities where the shrink-wrapping process can enter, and the tightness of the final surface mesh to the input, respectively. Once combined, these parameters provide a means to trade fidelity to the input for complexity of the output

Contact: Pierre Alliez

Participants: Pierre Alliez, Cédric Portaneri, Mael Rouxel-Labbé

Partner: GeometryFactory

7.1.4 CGAL - NURBS meshing

Name: CGAL - Meshing NURBS surfaces via Delaunay refinement

Keywords: Meshing, NURBS

Scientific Description: NURBS is the dominant boundary representation (B-Rep) in CAD systems. The meshing algorithms of NURBS models available for the industrial applications are based on the meshing of individual NURBS surfaces. This process can be hampered by the inevitable trimming defects in NURBS models, leading to non-watertight surface meshes and further failing to generate volumetric meshes. In order to guarantee the generation of valid surface and volumetric meshes of NURBS models even with the presence of trimming defects, NURBS models are meshed via Delaunay refinement, based on the Delaunay oracle implemented in CGAL. In order to achieve the Delaunay refinement, the trimmed regions of a NURBS model are covered with balls. Protection balls are used to cover sharp features. The ball centres are taken as weighted points in the Delaunay refinement, so that sharp features are preserved in the mesh. The ball sizes are determined with local geometric features. Blending balls are used to cover other trimmed regions which do not need to be preserved. Inside blending balls implicit surfaces are generated with Duchon's interpolating spline with a handful of sampling points. The intersection computation in the Delaunay refinement depends on the region where the intersection is computed. The general line/NURBS surface intersection is computed for intersections away from trimmed regions, and the line/implicit surface intersection is computed for intersections inside blending balls. The resulting mesh is a watertight volumetric mesh, satisfying user-defined size and shape criteria for facets and cells. Sharp features are preserved in the mesh adaptive to local features, and mesh elements cross smooth surface boundaries.

Functional Description: Input: NURBS surface Output: isotropic tetrahedron mesh

Publication: [hal-03276860](https://hal.archives-ouvertes.fr/hal-03276860)

Contact: Pierre Alliez

Participants: Pierre Alliez, Laurent Busé, Xiao Xiao, Laurent Rineau

Partner: GeometryFactory

7.2 New platforms

Participants: Pierre Alliez, Florent Lafarge.

No new platforms in the period.

8 New results

8.1 Analysis

8.1.1 Binary Graph Neural Networks

Participants: Gaetan Bahl.

In collaboration with Mehdi Bahri and Stefanos Zafeiriou (Imperial College London).

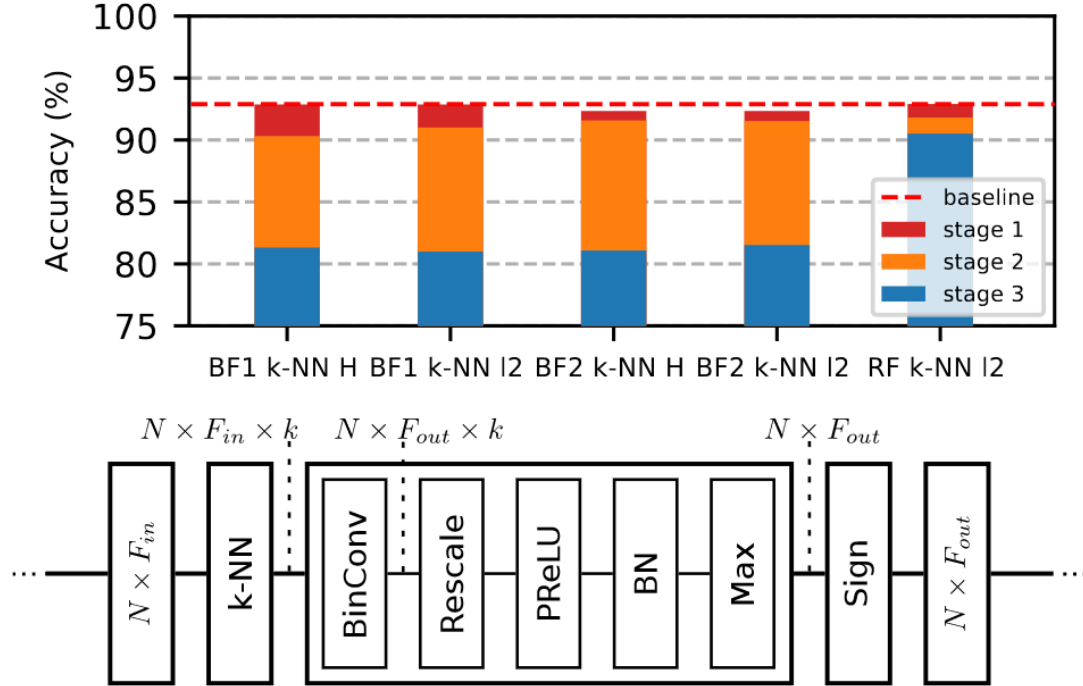


Figure 1: Binary Graph Neural Networks. Top: Test accuracy of different binarization schemes at all stages of our cascaded distillation protocol (baseline:92.89%). Bottom:The "BF2" variant of our XorEdgeConv operator.

Graph Neural Networks (GNNs) have emerged as a powerful and flexible framework for representation learning on irregular data. As they generalize the operations of classical CNNs on grids to arbitrary topologies, GNNs also bring much of the implementation challenges of their Euclidean counterparts. Model size, memory footprint, and energy consumption are common concerns for many real world applications. Network binarization allocates a single bit to parameters and activations, thus dramatically reducing the memory requirements (up to 32x compared to single-precision floating-point numbers) and maximizing the benefits of fast SIMD instructions on modern hardware for measurable speedups. However, in spite of the large body of work on binarization for classical CNNs, this area remains largely unexplored in geometric deep learning. In this paper, we present and evaluate different strategies for the binarization of graph neural networks. We show that through careful design of the models, and control of the training process, binary graph neural networks can be trained at only a moderate cost in accuracy on challenging benchmarks (See Figure 1). In particular, we present the first dynamic graph neural network in Hamming space, able to leverage efficient k -NN search on binary vectors to speed-up the construction of the dynamic graph. We further verify that the binary models offer significant savings on embedded devices. This work was presented at the Conference on Computer Vision and Pattern Recognition [19].

8.1.2 Planar Shape Based Registration for Multi-modal Geometry

Participants: Muzingzi Li, Florent Lafarge.

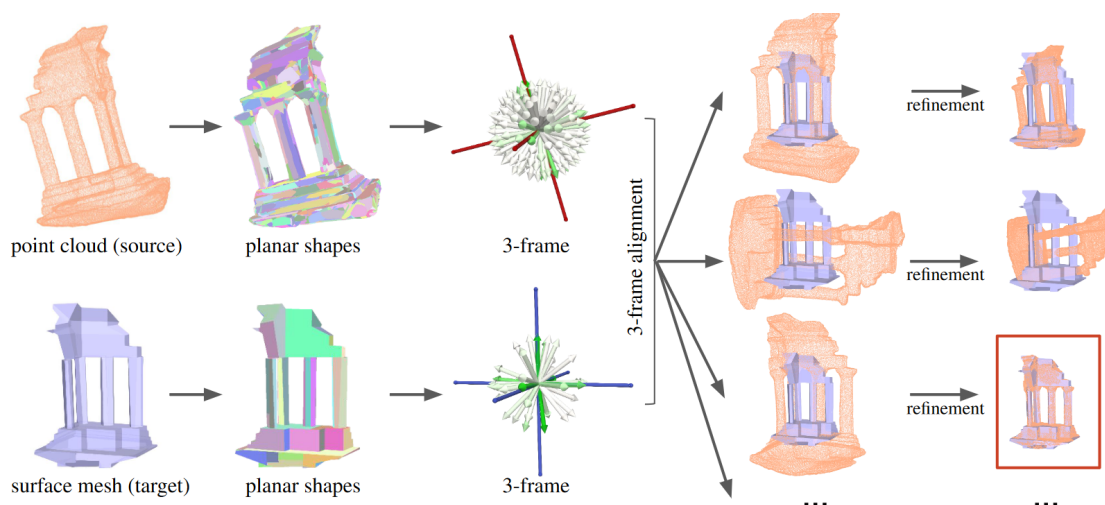


Figure 2: Planar Shape Based Registration. Planar shapes are first extracted from the input point cloud and surface mesh. From the surface-normal distribution of planar shapes (green corresponds to a high portion of planar area with normal pointing towards the arrow direction), three dominant directions are estimated, called a 3-frame. The 3-frames are aligned between the source and the target, leading to 24 possible rotations (only three are represented here). The refinement step takes each candidate rotation and estimate a final similarity transform. The alignment with the minimal loss is kept as the final result (see red frame).

We present a global registration algorithm for multi-modal geometric data, typically 3D point clouds and meshes. Existing feature-based methods and recent deep learning based approaches typically rely upon point-to-point matching strategies that often fail to deliver accurate results from defect-laden data. In contrast, we reason at the scale of planar shapes whose detection from input data offers robustness on a range of defects, from noise to outliers through heterogeneous sampling. The detected planar shapes are projected into an accumulation space from which a rotational alignment is operated. A second step then refines the result with a local continuous optimization which also estimates the scale (See Figure 2). We demonstrate the robustness and efficacy of our algorithm on challenging real-world data. In particular, we show that our algorithm competes well against state-of-the-art methods, especially on piece-wise planar objects and scenes. This work was presented at the British Machine Vision Conference [17].

8.1.3 Polygonal Building Segmentation by Frame Field Learning

Participants: Nicolas Girard.

In collaboration with Dmitriy Smirnov and Justin Solomon (MIT) and Yuliya Tarabalka (Luxcarta).

While state of the art image segmentation models typically output segmentations in raster format, applications in geographic information systems often require vector polygons. To help bridge the gap between deep network output and the format used in downstream tasks, we add a frame field output to a deep segmentation model for extracting buildings from remote sensing images. We train a deep

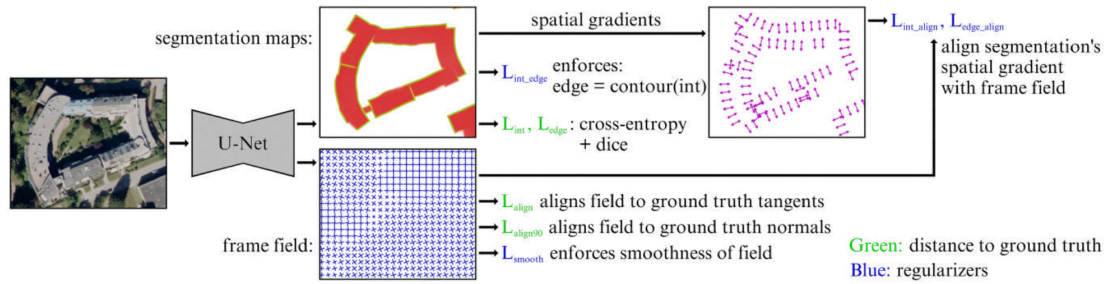


Figure 3: Polygonal Building Segmentation. Given an overhead image, our model outputs an edge mask, interior mask, and frame field. The loss aligns the masks and field to ground truth data, enforces smoothness of the frame field, and ensures consistency between the outputs.

neural network that aligns a predicted frame field to ground truth contours (See Figure 3). This additional objective improves segmentation quality by leveraging multi-task learning and provides structural information that later facilitates polygonization; we also introduce a polygonization algorithm that utilizes the frame field along with the raster segmentation. Our code is available [here](#). This work was presented at the Conference on Computer Vision and Pattern Recognition [16].

8.1.4 An efficient representation of 3D buildings: application to the evaluation of city models

Participants: Florent Lafarge.

In collaboration with Oussama Ennaffi (Gambi-M), Arnaud Le Bris and Clément Mallet (IGN).

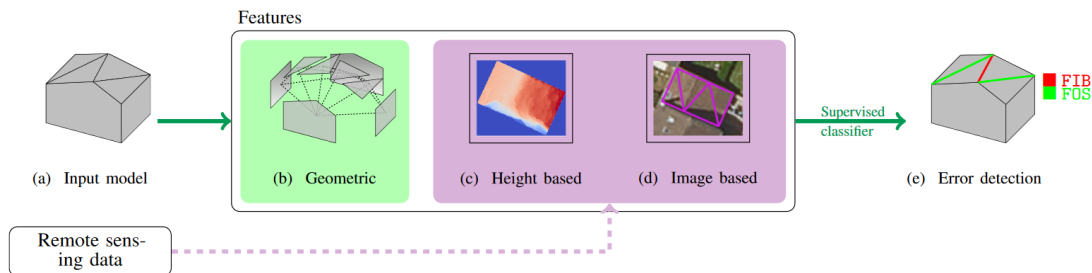


Figure 4: Evaluation of city models. Both intrinsic attributes (b) and external modalities (c and d) are used to describe the input model (a). Based on these features, a trained classifier predicts the presence of each error label (e).

City modeling consists in building a semantic generalized model of the surface of urban objects. These could be seen as a special case of Boundary representation surfaces. Most modeling methods focus on 3D buildings with Very High Resolution overhead data (images and/or 3D point clouds). The literature abundantly addresses 3D mesh processing but frequently ignores the analysis of such models. This requires an efficient representation of 3D buildings. In particular, for them to be used in supervised learning tasks, such a representation should be scalable and transferable to various environments as only a few reference training instances would be available. In this paper, we propose two solutions that take into account the specificity of 3D urban models (See Figure 4). They are based on graph kernels and scattering networks. They are here evaluated in the challenging framework of quality evaluation of

building models. The latter is formulated as a supervised multilabel classification problem, where error labels are predicted at building level. The experiments show for both feature extraction strategy strong and complementary results (F-score > 74 % for most labels). Transferability of the classification is also examined in order to assess the scalability of the evaluation process yielding very encouraging scores (F-score > 86 % for most labels). This work was presented at the ISPRS congress [20].

8.1.5 Automatic fault mapping in remote optical images and topographic data with deep learning

Participants: Lionel Mattéo, Nicolas Girard, Onur Tasar.

In collaboration with Isabelle Manighetti, Martijn van den Ende, Antoine Mercier, Frederique Leclerc and Tiziano Giampetro (Géoazur), Yuliya Tarabalka (Luxcarta), Jean-michel Gaucel (Thalès Alenia Space), Stéphane Dominguez and Jacques Malavieille (Université de Montpellier).

Faults form dense, complex multi-scale networks generally featuring a master fault and myriads of smaller-scale faults and fractures off its trace, often referred to as damage. Quantification of the architecture of these complex networks is critical to understanding fault and earthquake mechanics. Commonly, faults are mapped manually in the field or from optical images and topographic data through the recognition of the specific curvilinear traces they form at the ground surface. However, manual mapping is time-consuming, which limits our capacity to produce complete representations and measurements of the fault networks. To overcome this problem, we have adopted a machine learning approach, namely a U-Net Convolutional Neural Network, to automate the identification and mapping of fractures and faults in optical images and topographic data. Intentionally, we trained the CNN with a moderate amount of manually created fracture and fault maps of low resolution and basic quality, extracted from one type of optical images (standard camera photographs of the ground surface). Based on a number of performance tests, we select the best performing model, MRef, and demonstrate its capacity to predict fractures and faults accurately in image data of various types and resolutions (ground photographs, drone and satellite images and topographic data). MRef exhibits good generalization capacities, making it a viable tool for fast and accurate mapping of fracture and fault networks in image and topographic data. The MRef model can thus be used to analyze fault organization, geometry, and statistics at various scales, key information to understand fault and earthquake mechanics (see Figure 5). This work was published in the Journal of Geophysical Research [13].

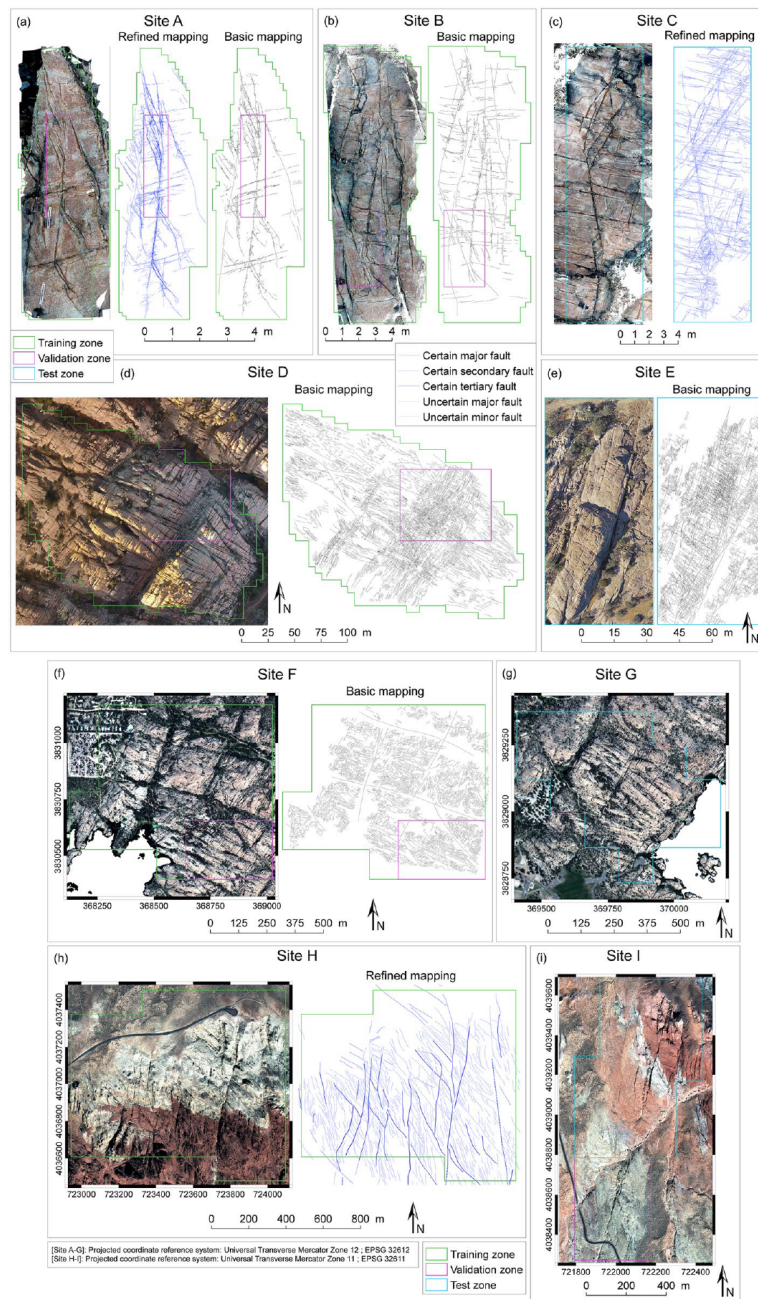


Figure 5: Sites, images, and ground truth fault maps used in present study. (a–g) from Granite Dells site, (h–i) from Valley of Fire site. (a–c): hand-held camera ground images, (d and e): drone images, (f–i): Pléiades satellite images (ID in data statement). Training, validation and test zones are indicated in green, purple and blue, respectively. Sites C and G are only used as test zones. Basic (black) and refined (blue) mapping were done manually (not shown on Site I for figure clarity). Fault hierarchy is indicated only in refined maps in (a, c, and h) through different line thicknesses, but is most visible when zooming the figures out. Dotted lines are for uncertain fractures and faults. Note, in most sites, the very high density of fractures and faults, the existence of different fault families with various orientations, and the presence of non-tectonic features.

8.1.6 Geometric approximation of structured scenes from images

Participants: Muzingzi Li.

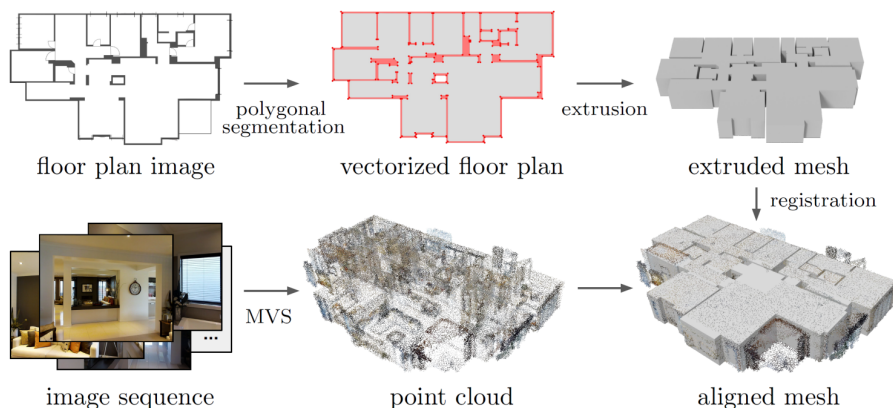


Figure 6: Overview of our pipeline. The floor plan image is first converted into a vectorized form, where the interior region is depicted in gray. The 2.5D model is generated by extruding polygons corresponding to the interior region given an estimated height. Finally, the model is registered to the point cloud, computed from the image sequence, to output an aligned model.

Geometric approximation of urban objects with compact and accurate representation is a challenging problem that concerns both computer vision and computer graphics communities. Existing literature mainly focuses on reconstruction from high-quality point clouds obtained by laser scanning which are too costly for many practical scenarios. This motivates the investigation into the problem of geometric approximation from low-budget image data. Dense reconstruction from a collection of images is made possible by recent advances in multi-view stereo techniques, yet the resulting point cloud is often far from perfect for generating a compact model. In particular, our goal is to describe the captured scene with a compact and accurate representation. In this thesis, we propose two generic algorithms which address different aspects of image-based geometric approximation. First, we present an algorithm for extracting and vectorizing objects in images with polygons. Second, we present a global registration algorithm from multi-modal geometric data, typically 3D point clouds and surface meshes (Figure 6). Both approaches exploit the detection of linear geometric primitives to approximate either 2D silhouettes with polygons consisting of line segments, or 3D point sets with a collection of planar shapes. The proposed algorithms could be used sequentially to form a pipeline for geometric approximation of an urban object from a set of image data, consisting of an overhead shot for coarse model extraction and multi-view stereo data for generation of point clouds. We demonstrate the robustness and scalability of our methods for structured scenes and objects, as well as applicative potential for free-form objects [22].

8.1.7 Wearable Cooperative SLAM System for Real-time Indoor Localization Under Challenging Conditions

Participants: Pierre Alliez, Fernando Ireta Munoz.

This research has been funded by the ANR/DGA MALIN challenge. In collaboration with the IBSIC lab (Evry) and Innodura TB (SEM from Lyon).



Figure 7: WeCo-SLAM system. The group of LVI sensors is mounted on the shoulder of a tactical waistcoat. This allows its wearability by a trained agent that performs recon missions under challenging conditions. Example results are shown right, where the followed trajectory is shown in green and the outdoor/indoor reconstruction of visited buildings is shown in gray and red, respectively. Predicted GPS coordinates are sent via radio to a ground station that monitors the agent's position in real-time.

Real-time globally consistent GPS tracking is critical for an accurate localization and is crucial for applications such as autonomous navigation or multi-robot mapping. However, under challenging environment conditions such as indoor/outdoor transitions, GPS signals are partially available or not consistent over time. We contribute a real-time tracking system for continuously locating emergency response agents in challenging conditions (Figure 7). A cooperative localization method based on Laser-Visual-Inertial (LVI) and GPS sensors is achieved by communicating optimization events between a LiDAR-Inertial-SLAM (LI-SLAM) and Visual-Inertial-SLAM (VI-SLAM) that operate simultaneously. The estimation of the pose assisted by multiple SLAM approaches provides the GPS localization of the agent when a stand-alone GPS fails. The system has been tested under the terms of the MALIN Challenge, which aims to globally localize agents across outdoor and indoor environments under challenging conditions (such as smoked rooms, stairs, indoor/outdoor transitions, repetitive patterns, extreme lighting changes) where it is well known that a stand-alone SLAM will not be enough to maintain the localization. The system achieved an absolute trajectory error of 0.48%, with a pose update rate between 15 and 20 Hz. Furthermore, the system is able to build a global consistent 3D LiDAR Map that is post-processed to create a 3D reconstruction at different levels of details [12].

8.2 Reconstruction

8.2.1 Floorplan generation from 3D point clouds: A space partitioning approach

Participants: Florent Lafarge.

In collaboration with Hao Fang and Hui Huang (Shenzhen University) and Cihui Pan (BeiKe).

We propose a novel approach to automatically reconstruct the floorplan of indoor environments from raw sensor data. In contrast to existing methods that generate floorplans under the form of a



Figure 8: Floorplan generation. Left: the algorithm departs from raw point clouds as input data. Right: the floorplan of the indoor scene is reconstructed as a planar graph where each simple cycle represents the polygonal boundary of a room. Note that the 2D floorplan is converted into a 3D CAD model for visualization purposes.

planar graph by detecting corner points and connecting them, our framework employs a strategy that decomposes the space into a polygonal partition and selects edges that belong to wall structures by energy minimization. By relying on an efficient space-partitioning data structure instead of a traditional and delicate corner detection task, our framework offers a high robustness to imperfect data. We demonstrate the potential of our algorithm on both RGBD and LIDAR points scanned from simple to complex scenes. Experimental results indicate that our method is competitive with respect to existing methods in terms of geometric accuracy and output simplicity (See Figure 8). This work was published in the ISPRS Journal of Photogrammetry and Remote Sensing [11].

8.2.2 Progressive discrete domains for implicit surface reconstruction

Participants: Tong Zhao, Pierre Alliez.

In collaboration with Laurent Busé (Aromath Inria project-team), Tamy Boubekour and Jean-Marc Thiery (Télécom ParisTech and Adobe Research, Paris). Tong Zhao is funded by 3IA Côte d'Azur.

Many global implicit surface reconstruction algorithms formulate the problem as a volumetric energy minimization, trading data fitting for geometric regularization. As a result, the output surfaces may be located arbitrarily far away from the input samples. This is amplified when considering i) strong regularization terms, ii) sparsely distributed samples or iii) missing data. This breaks the strong assumption commonly used by popular octree-based and triangulation-based approaches that the output surface should be located near the input samples. As these approaches refine during a pre-process, their cells near the input samples, the implicit solver deals with a domain discretization not fully adapted to the final isosurface. We relax this assumption and propose a progressive coarse-to-fine approach that jointly refines the implicit function and its representation domain, through iterating solver, optimization and refinement steps applied to a 3D Delaunay triangulation. There are several advantages to this approach: the discretized domain is adapted near the isosurface and optimized to improve both the solver conditioning and the quality of the output surface mesh contoured via marching tetrahedra (see Figure 9). This work was presented at the EUROGRAPHICS Symposium on Geometry Processing [15].

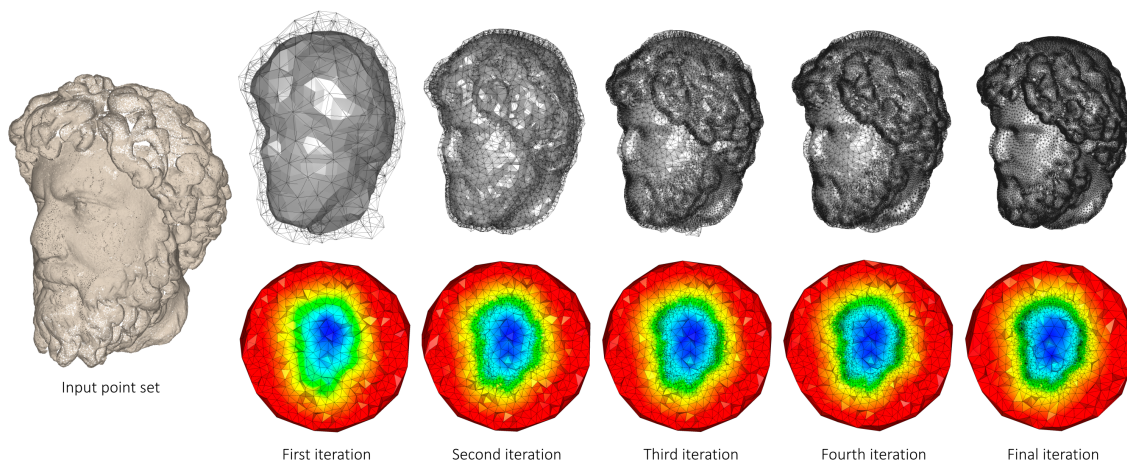


Figure 9: Progressive discrete domain for surface reconstruction. Left: input point set. Right: our approach jointly refines and optimizes the implicit function (bottom) and its discretized domain (a 3D Delaunay triangulation) around the refined isosurface (top). In such a progressive approach, the implicit solver is used iteratively, as a means of consolidating hypotheses emitted in previous iterations. Top: the isosurface and only the set of tetrahedra intersected by the isosurface are shown. Bottom: the implicit function (piecewise-linear over the 3D triangulation) is depicted on the facets intersected by a clipping plane

8.2.3 Efficient open surface reconstruction from lexicographic optimal chains and critical bases

Participants: Julien Vuillamy.

In collaboration with André Lieutier (Dassault Systemes) and David Cohen-Steiner (Inria Datashape).

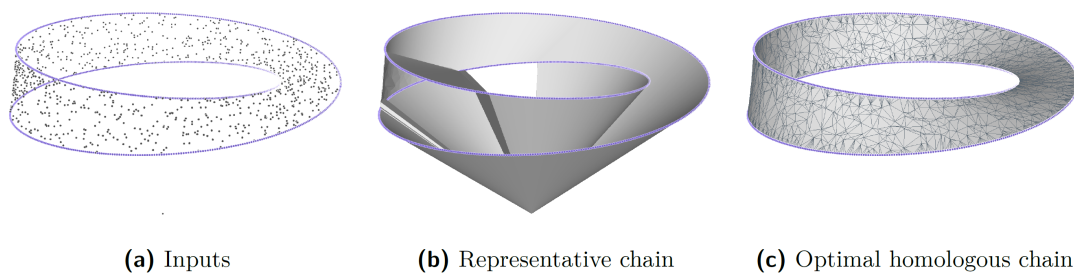


Figure 10: Illustration of the open surface reconstruction steps. From a set of points and a 1-cycle in the Delaunay 3-complex (a), a representative 2-chain bounding this 1-cycle is computed (b). The lexicographic optimal 2-chain homologous to this representative chain is then computed (c).

Previous works on lexicographic optimal chains have shown that they provide meaningful geometric homology representatives while being easier to compute than their l_1 -norm optimal counterparts. This work presents a novel algorithm to efficiently compute lexicographic optimal chains with a given boundary in a triangulation of 3-space, by leveraging Lefschetz duality and an augmented version of the disjoint-set data structure. Furthermore, by observing that lexicographic minimization is a linear operation, we define a canonical basis of lexicographic optimal chains, called critical basis, and show




Figure 11: Open surface meshing on two models from the senseFly dataset. Computing the optimal chain takes 1.5 seconds (resp. 8.0 seconds) to mesh 1 million (resp. 5 millions) points.

how to compute it (Figure 10). In applications, the presented algorithms offer new promising ways of efficiently reconstructing open surfaces in difficult acquisition scenarios (Figure 11).

8.2.4 Planimetric simplification and lexicographic optimal chains for 3D urban scene reconstruction

Participants: Julien Vuillamy.

This Ph.D. thesis was co-advised by Florent Lafarge, Pierre Alliez and David Cohen-Steiner at Inria, and by André Lieutier from Dassault Systemes.

Creating mesh representations for urban scenes is a requirement for numerous modern applications of urban planning ranging from visualization, inspection, to simulation. Adding to the diversity of possible input data – photography, laser-based acquisitions and existing geographical information system (GIS) data, the variety of urban scenes as well as the large-scale nature of the problem makes for a challenging line of research. Working towards an automatic approach to this problem suggests that a one-fits-all method is hardly realistic. Two independent approaches of reconstruction from point clouds have thus been investigated in this work, with radically different points of view intended to cover a large number of use cases. In the spirit of the GIS community, the first approach makes strong assumptions on the reconstructed scenes and creates a 2.5D piecewise-planar representation of buildings using an intermediate 2D cell decomposition. Constructing these decompositions from noisy or incomplete data often leads to overly complex representations, which lack the simplicity or regularity expected in this context of reconstruction. Loosely inspired by clustering problems such as mean-shift, the focus is put on simplifying such partitions by formulating an optimization process based on a tradeoff between attachment to the original partition and objectives striving to simplify and regularize the arrangement. This method involves working with point-line duality, defining local metrics for line movements and optimizing using Riemannian gradient descent.

The second approach is intended to be used in contexts where the strong assumptions on the representation of the first approach do not hold. We strive here to be as general as possible and investigate the problem of point cloud meshing in the context of noisy or incomplete data. By considering a specific minimization, corresponding to lexicographic orderings on simplicial chains, polynomial-time algorithms

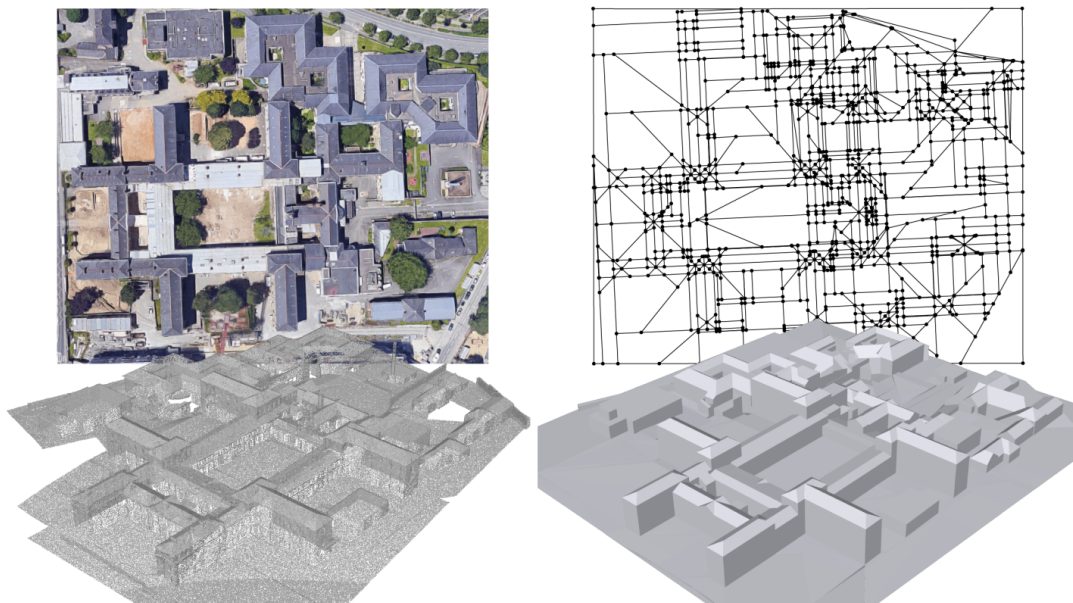


Figure 12: Results on a large regular building complex. Our 3D output model (bottom right) preserves the orthogonality existing between facade and rooftop components. Note how the planimetric partition (top right) aligns well with the aerial image (top left).

finding lexicographic optimal chains, homologous to a given chain or bounded by a given chain, are derived from algorithms for the computation of simplicial persistent homology. For pseudomanifold complexes in codimension 1, leveraging duality and an augmented version of the disjoint-set data structure improves the complexity of these problem instances to quasi-linear time algorithms. By combining its uses with a sharp feature detector in the point cloud, we illustrate different use cases in the context of urban reconstruction [23].

8.3 Approximation

8.3.1 Delaunay Meshing and Repairing of NURBS Models

Participants: Xiao Xiao, Pierre Alliez.

In collaboration with Laurent Busé (Aromath) and Laurent Rineau (Geometry Factory). The postdoctoral stay of Xiao Xiao was funded both by Inria and by Geometry Factory.

CAD models represented by NURBS surface patches are often hampered with defects due to inaccurate representations of trimming curves. Such defects make these models unsuitable to the direct generation of valid volume meshes, and often require trial-and-error processes to fix them. We propose a fully automated Delaunay-based meshing approach which can mesh and repair simultaneously, while being independent of the input NURBS patch layout. Our approach proceeds by Delaunay filtering and refinement, in which trimmed areas are repaired through implicit surfaces. Beyond repair, we demonstrate its capability to smooth out sharp features, defeature small details, and mesh multiple domains in contact (see Figure 13). This work was presented at the EUROGRAPHICS Symposium on Geometry Processing [14]. A novel software component for the CGAL library has been implemented.

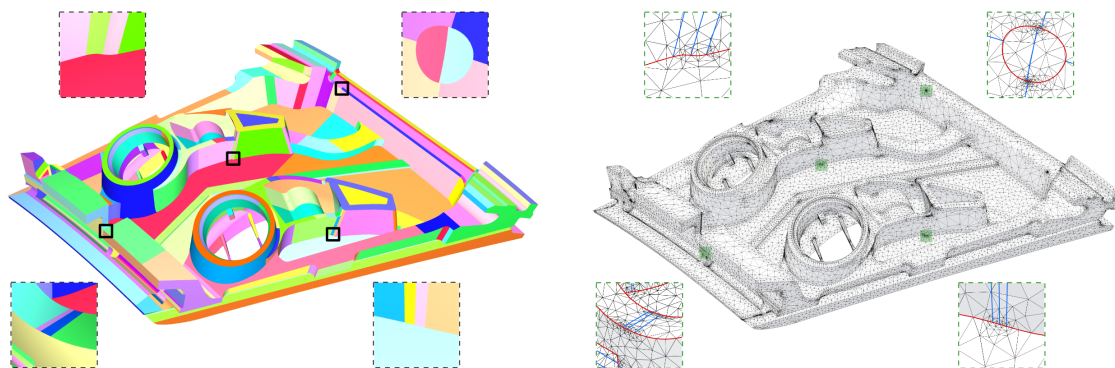


Figure 13: Delaunay meshing. Left: a complex mechanical part (609 NURBS surface patches). Right: its Delaunay surface mesh. Apart from trimming defects, the input model also suffers from topological defects of trimming curves. The locally dense mesh is due to short sharp feature curves, as shown in close-up views.

8.3.2 Alpha wrapping with an offset

Participants: Cedric Portaneri, Pierre Alliez.

In collaboration with Mael Rouxel-Labbe (Geometry Factory), David Cohen-Steiner (Inria DataShape), and Michael Hemmer (Google X).

Various tasks in geometric modeling and processing require 3D objects represented as valid surface meshes, where "valid" refers to meshes that are watertight, intersection-free, orientable, and combinatorially 2-manifold. Such representations offer well-defined notions of interior/interior and geodesic neighborhoods. 3D data are usually acquired through measurements followed by reconstruction, designed by humans, or generated through imperfect automated processes. As a result, they can exhibit a wide variety of defects including gaps, missing data, self-intersections, degeneracies such as zero-volume structures, and non-manifold features. Given the large repertoire of possible defects, many methods and data structures have been proposed to repair specific defects, usually with the goal of guaranteeing specific properties in the repaired 3D model. Reliably repairing all types of defects is notoriously difficult and is often an ill-posed problem as many valid solutions exist for a given 3D model with defects. In addition, the input model can be overly complex with unnecessary geometric details, spurious topological structures, nonessential inner components, or excessively fine discretizations. For applications such as collision avoidance, path planning, or simulation, getting an approximation of the input can be more relevant than repairing it. Approximation herein refers to an approach capable of filtering out inner structures, fine details and cavities, as well as wrapping the input within a user-defined offset margin. Given an input 3D geometry, we address the problem of computing a conservative approximation, where conservative means guaranteeing a strictly enclosed input (Figure 14). We seek unconditional robustness in the sense that the output mesh should be valid (oriented, combinatorially 2-manifold and without self-intersections), even for raw input with many defects and degeneracies. The default input is a soup of 3D triangles, but the generic design of our method leaves the door open to other types of finite 3D primitives such as segments or points, or even a mixture of those. This work is not published yet.

8.3.3 Geometric Model Reduction Driven by Numerical Simulation Accuracy

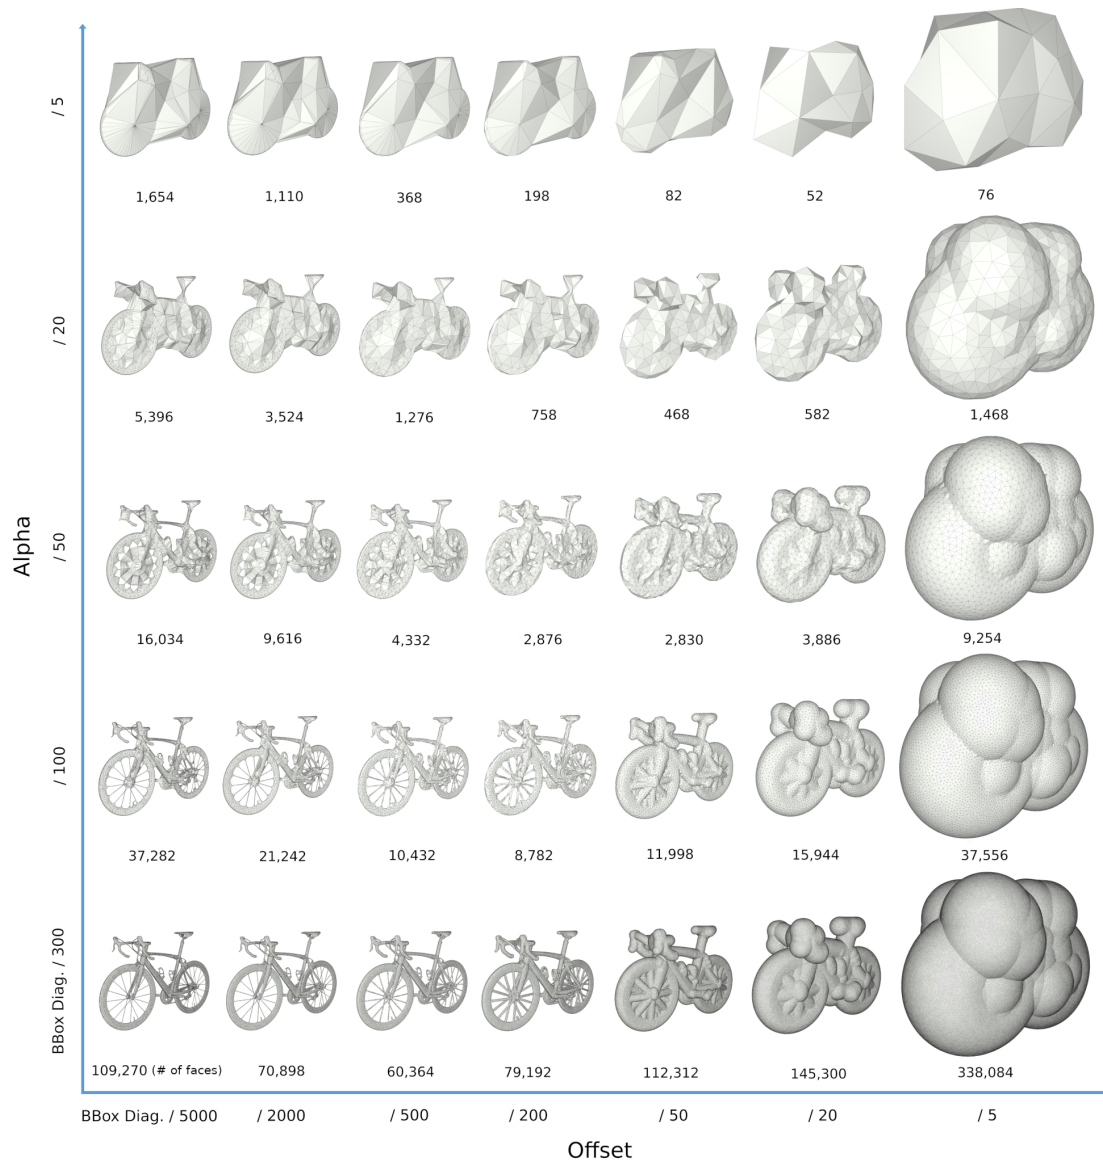


Figure 14: The algorithm takes a 3D triangle soup (or point set) as input, and generates a valid surface triangle mesh that contains the input (watertight, intersection-free and combinatorially 2-manifold). The algorithm proceeds by shrink-wrapping and refining a 3D Delaunay triangulation loosely bounding the input. Two user-defined parameters, alpha and offset (increasing on the figure), offer control over the maximum size of cavities where the shrink-wrapping process can enter, and the tightness of the final surface mesh to the input, respectively. Once combined, these parameters provide a means to trade fidelity to the input for complexity of the output (expressed in number of triangle facets).

Participants: Vincent Vadez, Pierre Alliez.

In collaboration with Francois Brunetti (Dorea technology). This work has been funded by a Cifre collaboration between the Dorea SME and Inria.

When dealing with real-time simulation, radiative thermal computations have always been, and are

still a challenge. Notably, computing view factors is very compute-intensive when the input 3D model is complex and exhibits many holes and occlusions. The task is even more difficult on complex geometries generated through topological optimization and on dense meshes required for finite element simulation. This paper focuses on geometric model reduction through mesh decimation. The decimation algorithm is made accurate to the radiative thermal simulation, in order to trade accuracy for computing times. More specifically, the input model is first decomposed into thermal nodes, then we estimate through radiative thermal simulation the sensitivity of decimating each thermal node against a maximum temperature tolerance. Such estimations are then utilized to render the entire mesh decimation process informed by the physical simulation. These estimations are relevant for predicting the amount of decimation applicable to each thermal node, given a user-defined maximum temperature tolerance (see Figure 15). This work was presented at the International Conference on Environmental Systems [18].

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

X Development LLC (formerly Google X)

Participants: Cédric Portaneri, Pierre Alliez.

We explore an adaptive framework to repair and generate levels of detail from defect-laden 3D measurement data. Repair herein implies to convert raw 3D point sets or triangle soups into watertight solids, which are suitable for computing collision detection and situational awareness of robots. In this context, adaptive stands for the capability to select not only different levels of detail, but also the type of representation best suited to the environment of the scene and the task of the robot. This project focuses on geometry processing for collision detection of robots with mechanical parts and other robots in motion.

Dorea technology

Participants: Vincent Vadez, Pierre Alliez.

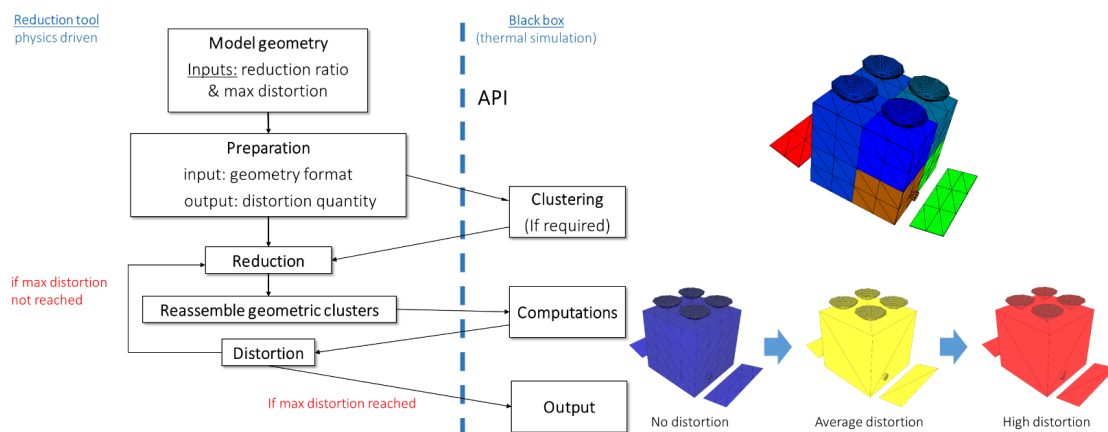


Figure 15: Geometric Model Reduction. Overview of the reduction process. Left: mesh reduction. Right: numerical simulation.

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.

Luxcarta

Participants: Johann Lussange, Florent Lafarge.

The goal of this collaboration is to develop efficient, robust and scalable methods for extracting geometric shapes from the last generation of satellite images (at 0.3m resolution). Besides satellite images, input data will also include classification maps that identify the location of buildings and Digital Surface Models that bring rough pixel-based estimation of the urban object elevation. The geometric shapes will be first restricted to planes that typically describe well the piecewise planar geometry of buildings. The goal will be then to detect and identify each roof section of facade component of a building by a plane in the 3D space. This DGA Rapid project started in September 2020, for a total duration of 2 years.

CSTB

Participants: Mulin Yu, Florent Lafarge.

This collaboration takes the form of an independent contract. The 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 in November 2019, for a total duration of 3 years.

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.

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.

10 Partnerships and cooperations

10.1 European initiatives

10.1.1 FP7 & H2020 projects

We are participating in the BIM2TWIN RIA project (Optimal Construction Management and Production Control), coordinated by CSTB and sponsored by the H2020 “Digital Twins” program. Our overall objective is to build a virtual replica platform, referred to as a digital building twin, to improve construction management, i.e. enhance quality and safety, reduce operational waste, carbon footprints and costs, and shorten schedules. The platform should provide full situational awareness via various sensing modalities in order to monitor human activities during construction and assess conformity to the digital plan. My team will contribute to the conformity assessment part of the platform by providing novel geometry processing pipelines and reliable software components for 3D urban reconstruction.

We are participating in the GRAPES RTN project (learninG, Representing, And oPtimizing shapES), coordinated by ATHENA (Greece). GRAPES aims at advancing the state of the art in Mathematics, CAD, and Machine Learning in order to promote game changing approaches for generating, optimizing, and learning 3D shapes, along with a multisectoral training for young researchers. Concrete applications include simulation and fabrication, hydrodynamics and marine design, manufacturing and 3D printing, retrieval and mining, reconstruction and urban planning. In this context, Pierre Alliez advises a PhD student in collaboration with the beneficiary partner Geometry Factory, an Inria spin-off company commercializing components from the CGAL library. Our objective is to advance the state of the art in applied mathematics, computer-aided design (CAD) and machine learning in order to promote novel approaches for generating, optimizing, and learning 3D shapes. We will focus on dense semantic segmentation of 3D point clouds and semantic-aware reconstruction of 3D scenes. Our goal is to enrich the final reconstructed 3D models with labels that reflect the main semantic class of outdoor scenes such as ground, buildings, roads and vegetation. The Titane project-team will also contribute to the training of young researchers via tutorials on the CGAL library and practical exercises on real-world use cases.

10.2 National initiatives

10.2.1 ANR

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

Participants: Pierre Alliez (*contact*), Lucas Dubouchet, 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 Appliquées (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.

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 localization 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.

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 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.

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, IGP 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.

BIOM: Building Indoor and Outdoor Modeling

Participants: Muxingzi Li, Pierre Alliez, Florent Lafarge (*contact*).

The BIOM project aims at automatic, simultaneous indoor and outdoor modeling 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.

11 Dissemination

Participants: Pierre Alliez, Florent Lafarge.

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

Member of the organizing committees TITANE organized the JFIG (Journées Françaises d'Informatique Graphique), in collaboration with the GRAPHDECO Inria team and Mediacoding team from I3S (CNRS).

11.1.2 Scientific events: selection

Chair of conference program committees

- Florent Lafarge was co-program chair of the ISPRS congress 2021 [21].
- Pierre Alliez was a member of the scientific committee for the SophIA Summit conference.

Member of the conference program committees

- Pierre Alliez was a member of the international program committee for the 2021 Eurographics Symposium on Geometry Processing.

Reviewer

- Florent Lafarge was a reviewer for CVPR, ICCV, BMVC and ACM SIGGRAPH Asia.
- Pierre Alliez was a reviewer for ACM SIGGRAPH.

11.1.3 Journal

- Pierre Alliez was a reviewer for Computer Graphics Forum.

Member of the editorial boards

- Pierre Alliez is an associate editor (and now editor in chief) of the Computer Graphics Forum, Computer-Aided Geometric Design and Computer-Aided Design. He is a member of the editorial board of the CGAL open source project.
- Florent Lafarge is an associate editor for the ISPRS Journal of Photogrammetry and Remote Sensing.

Reviewer - reviewing activities

- Florent Lafarge was a reviewer for TPAMI, CGF and the ISPRS Journal of Photogrammetry and Remote Sensing.
- Pierre Alliez was a reviewer for ACM Transactions on Graphics.

11.1.4 Invited talks

- Pierre Alliez gave an invited talk at CORESA conference (Compression et représentation des signaux audiovisuels).
- Florent Lafarge gave a talk at the Airbus AI day "Journée d'information Intelligence Artificielle" ("IA et Télédétection : apprendre à extraire des informations géométriques dans les images satellitaires") and at the Inria intro (The challenge of automated geometric modeling of structured objects).

11.1.5 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.

11.1.6 Scientific expertise

- Pierre Alliez was a reviewer for the ERC, for the European commission, and is a scientific advisory board member for the Bézout Labex in Paris (Models and algorithms: from the discrete to the continuous).
- Florent Lafarge was a reviewer for the MESR (CIR and JEI expertises).

11.1.7 Research administration

- Pierre Alliez is Head of Science (VP for Science) of the Inria Sophia Antipolis center, in tandem with Fabien Gandon. In this role, he works with the executive team to advise and support the Institute's scientific strategy, to stimulate new research directions and the creation of project-teams, and to manage the scientific evaluation and renewal of the project-teams. This part-time position runs until mid-2022.
- Pierre Alliez is a member of the scientific committee of the 3IA Côte d'Azur.
- Florent Lafarge is a member of the NICE committee. The main actions of the NICE committee are to verify the scientific aspects of the files of postdoctoral students, to give scientific opinions on candidates for national campaigns for postdoctoral stays, delegations, secondments as well as requests for long duration invitations.

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

- Master: Pierre Alliez, advanced machine learning, 10h, M2, Univ. Côte d'Azur, France.
- Master: Pierre Alliez and Gaétan Bahl, deep learning, 21h, M2, Univ. Côte d'Azur, France.
- Master: Pierre Alliez and Florent Lafarge, 3D Meshes and Applications, 32h, M2, Ecole des Ponts ParisTech, France.
- Master: Florent Lafarge, Applied AI, 8h, M2, Univ. Côte d'Azur, France.
- Master: Pierre Alliez and Florent Lafarge, Interpolation numérique, 60h, M1, Univ. Côte d'Azur, France.
- Master: Florent Lafarge and Gaétan Bahl, Mathématiques pour la géométrie, 35h, M1, EFREI, France.

11.2.2 Supervision

- PhD defended in September: Julien Vuillamy, city reconstruction from multi-sourced data, Pierre Alliez and Florent Lafarge.
- PhD defended in September: Muxingzi Li, indoor reconstruction from a smartphone, since February 2018, Florent Lafarge.
- 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: Vincent Vadez, geometric simplification of satellites for thermal simulation, since August 2018, Pierre Alliez.
- PhD in progress: Tong Zhao, shape reconstruction, since November 2019, Pierre Alliez and Laurent Busé (from the Aromath Inria project-team), in collaboration with Jean-Marc Thiery and Tamy Boubekeur (Telecom ParisTech and Adobe research).

- PhD in progress: Rao Fu, since January 2021, Pierre Alliez.
- PhD in progress: Alexandre Zoppis, 3D modeling from satellite images, since October 2021, Florent Lafarge.

11.2.3 Juries

- Pierre Alliez was a PhD thesis reviewer for Pierre-Alain Langlois (Ecole des Ponts ParisTech) and Raphaël Groscolt (Paris Dauphine), and a PhD thesis committee member for Yana Nehmé (INSA Lyon).

12 Scientific production

12.1 Major publications

- [1] J.-P. Bauchet and F. Lafarge. ‘Kinetic Shape Reconstruction’. In: *ACM Transactions on Graphics* (2020). This project was partially funded by Luxcarta Technology. We thank our anonymous reviewers for their input, Qian-Yi Zhou and Arno Knapitsch for providing us datasets from the Tanks and Temples benchmark (Meeting Room, Horse, M60, Barn, Ignatius, Courthouse and Church), and Pierre Alliez, Mathieu Desbrun and George Drettakis for their help and advice. We are also grateful to Liangliang Nan and Hao Fang for sharing comparison materials. Datasets Full thing, Castle, Tower of Piand Hilbert cube originate from Thingi 10K, Hand, Rocker Arm, Fertility and LansfromAim@Shape, and Stanford Bunny and Asian Dragon from the Stanford 3D Scanning Repository. DOI: [10.1145/3376918](https://doi.org/10.1145/3376918). URL: <https://hal.inria.fr/hal-02924409>.
- [2] M. Budninskiy, B. Liu, F. De Goes, Y. Tong, P. Alliez and M. Desbrun. ‘Optimal Voronoi Tessellations with Hessian-based Anisotropy’. In: *ACM Transactions on Graphics*. Proceedings of SIGGRAPH Asia (Dec. 2016), p. 12. URL: <https://hal.inria.fr/hal-01376243>.
- [3] L. Duan and F. Lafarge. ‘Towards large-scale city reconstruction from satellites’. In: *European Conference on Computer Vision (ECCV)*. Amsterdam, Netherlands, Oct. 2016. URL: <https://hal.inria.fr/hal-01352466>.
- [4] L. Feng, P. Alliez, L. Busé, H. Delingette and M. Desbrun. ‘Curved Optimal Delaunay Triangulation’. In: *ACM Transactions on Graphics*. Proceedings of SIGGRAPH 2018 37.4 (Aug. 2018), p. 16. DOI: [10.1145/3197517.3201358](https://doi.org/10.1145/3197517.3201358). URL: <https://hal.inria.fr/hal-01826055>.
- [5] M. Li, F. Lafarge and R. Marlet. ‘Approximating shapes in images with low-complexity polygons’. In: *CVPR 2020 - IEEE Conference on Computer Vision and Pattern Recognition*. Seattle / Virtual, United States, June 2020. URL: <https://hal.inria.fr/hal-02526028>.
- [6] M. Mandad, D. Cohen-Steiner and P. Alliez. ‘Isotopic Approximation within a Tolerance Volume’. In: *ACM Transactions on Graphics*. Proceedings of SIGGRAPH 34.4 (2015), p. 12. DOI: [10.1145/2766950](https://doi.org/10.1145/2766950). URL: <https://hal.inria.fr/hal-01186074>.
- [7] M. Mandad, D. Cohen-Steiner, L. Kobbelt, P. Alliez and M. Desbrun. ‘Variance-Minimizing Transport Plans for Inter-surface Mapping’. In: *ACM Transactions on Graphics* 36 (2017), p. 14. DOI: [10.1145/3072959.3073671](https://doi.org/10.1145/3072959.3073671). URL: <https://hal.inria.fr/hal-01519006>.
- [8] M. Rouhani, F. Lafarge and P. Alliez. ‘Semantic Segmentation of 3D Textured Meshes for Urban Scene Analysis’. In: *ISPRS Journal of Photogrammetry and Remote Sensing* 123 (2017), pp. 124–139. DOI: [10.1016/j.isprsjprs.2016.12.001](https://doi.org/10.1016/j.isprsjprs.2016.12.001). URL: <https://hal.inria.fr/hal-01469502>.
- [9] Z. Shi, P. Alliez, M. Desbrun, H. Bao and J. Huang. ‘Symmetry and Orbit Detection via Lie-Algebra Voting’. In: *Computer Graphics Forum*. Proceedings of EUROGRAPHICS Symposium on Geometry Processing (June 2016), p. 12. URL: <https://hal.inria.fr/hal-01344293>.
- [10] Y. Verdie, F. Lafarge and P. Alliez. ‘LOD Generation for Urban Scenes’. In: *ACM Transactions on Graphics* 34.3 (2015), p. 15. URL: <https://hal.inria.fr/hal-01113078>.

12.2 Publications of the year

International journals

- [11] H. Fang, F. Lafarge, C. Pan and H. Huang. ‘Floorplan generation from 3D point clouds: A space partitioning approach’. In: *ISPRS Journal of Photogrammetry and Remote Sensing* 175 (May 2021), pp. 44–55. DOI: [10.1016/j.isprsjprs.2021.02.012](https://doi.org/10.1016/j.isprsjprs.2021.02.012). URL: <https://hal.inria.fr/hal-03387916>.
- [12] V. Kachurka, B. Rault, F. I. Munoz, D. Roussel, F. Bonardi, J.-Y. Didier, H. HADJ-ABDELKADER, S. BOUCHAFA, P. Alliez and M. Robin. ‘WeCo-SLAM: Wearable Cooperative SLAM System for Real-time Indoor Localization Under Challenging Conditions’. In: *IEEE Sensors Journal* (2021), pp. 1–1. DOI: [10.1109/JSEN.2021.3101121](https://doi.org/10.1109/JSEN.2021.3101121). URL: <https://hal.archives-ouvertes.fr/hal-03341622>.
- [13] L. Mattéo, I. Manighetti, Y. Tarabalka, J.-M. Gaucel, M. van den Ende, A. Mercier, O. Tasar, N. Girard, F. Leclerc, T. Giampetro, S. Dominguez and J. Malavieille. ‘Automatic fault mapping in remote optical images and topographic data with deep learning’. In: *Journal of Geophysical Research : Solid Earth* (6th Apr. 2021). DOI: [10.1029/2020JB021269](https://doi.org/10.1029/2020JB021269). URL: <https://hal.inria.fr/hal-03191101>.
- [14] X. Xiao, P. Alliez, L. Busé and L. Rineau. ‘Delaunay Meshing and Repairing of NURBS Models’. In: *Computer Graphics Forum*. Proceedings of the EUROGRAPHICS Symposium on Geometry Processing 2021 40.5 (2021), pp. 125–142. DOI: [10.1111/cgf.14362](https://doi.org/10.1111/cgf.14362). URL: <https://hal.inria.fr/hal-03276860>.
- [15] T. ZHAO, P. Alliez, T. Boubekeur, L. Busé and J.-M. Thiery. ‘Progressive Discrete Domains for Implicit Surface Reconstruction’. In: *Computer Graphics Forum*. Proceedings of the EUROGRAPHICS Symposium on Geometry Processing 2021 40.5 (2021), pp. 143–156. DOI: [10.1111/cgf.14363](https://doi.org/10.1111/cgf.14363). URL: <https://hal.inria.fr/hal-03276748>.

International peer-reviewed conferences

- [16] N. Girard, D. Smirnov, J. Solomon and Y. Tarabalka. ‘Polygonal Building Segmentation by Frame Field Learning’. In: *CVPR 2021 - IEEE Conference on Computer Vision and Pattern Recognition*. Pittsburgh / Virtual, United States, 19th June 2021. URL: <https://hal.inria.fr/hal-02548545>.
- [17] M. Li and F. Lafarge. ‘Planar Shape Based Registration for Multi-modal Geometry’. In: *British Machine Vision Conference*. virtual, United Kingdom, 2021. URL: <https://hal.inria.fr/hal-03388316>.
- [18] V. Vadez, F. Brunetti and P. Alliez. ‘Geometric Model Reduction Driven by Numerical Simulation Accuracy’. In: *ICES 2021 - International Conference on Environmental Systems*. Lisbon, Portugal, 12th July 2021. URL: <https://hal.inria.fr/hal-03278653>.

Conferences without proceedings

- [19] M. Bahri, G. Bahl and S. Zafeiriou. ‘Binary Graph Neural Networks’. In: *CVPR 2021 - IEEE Conference on Computer Vision and Pattern Recognition*. Nashville / Virtual, United States, 19th June 2021. URL: <https://hal.archives-ouvertes.fr/hal-03184720>.
- [20] O. Ennafii, A. Le Bris, F. Lafarge and C. Mallet. ‘An efficient representation of 3D buildings: application to the evaluation of city models’. In: *ISPRS 2021 - XXIVth Congress of the International Society for Photogrammetry and Remote Sensing*. Vol. XLIII-B2-2021. Nice / Virtual, France, 28th June 2021, pp. 329–336. DOI: [10.5194/isprs-archives-xxliii-b2-2021-329-2021](https://doi.org/10.5194/isprs-archives-xxliii-b2-2021-329-2021). URL: <https://hal.archives-ouvertes.fr/hal-03276198>.
- [21] C. Mallet, F. Lafarge, M. Poreba, T. Wu, G. Bahl, M. Yu, A. Garioud, Y. Chen, S. Jiang, M. Y. Yang and N. Paparoditis. ‘Preface: the 2021 edition of the XXIVth ISPRS congress’. In: *XXIVth ISPRS Congress, 2021 edition*. Nice (virtuel), France, 17th June 2021, pp. 1–5. DOI: [10.5194/isprs-annals-v-1-2021-1-2021](https://doi.org/10.5194/isprs-annals-v-1-2021-1-2021). URL: <https://hal.archives-ouvertes.fr/hal-03276194>.

Doctoral dissertations and habilitation theses

- [22] M. Li. ‘Geometric approximation of structured scenes from images’. Université Côte d’Azur, 8th Oct. 2021. URL: <https://tel.archives-ouvertes.fr/tel-03388295>.
- [23] J. Vuillamy. ‘Planimetric simplification and lexicographic optimal chains for 3D urban scene reconstruction’. Université Côte d’Azur, 17th Sept. 2021. URL: <https://hal.archives-ouvertes.fr/tel-03339931>.

Reports & preprints

- [24] D. Cohen-Steiner, A. Lieutier and J. Vuillamy. *Efficient open surface reconstruction from lexicographic optimal chains and critical bases*. 30th Nov. 2021. URL: <https://hal.archives-ouvertes.fr/hal-03456390>.