

RESEARCH CENTRE

Inria Centre
at **Université Côte d'Azur**

2023

ACTIVITY REPORT

Project-Team

TITANE

Geometric Modeling of 3D Environments

DOMAIN

Perception, Cognition and Interaction

THEME

Interaction and visualization

Inria

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 Highlights of the year	6
5.1 Awards	6
6 New software, platforms, open data	6
6.1 New software	6
6.1.1 CGAL - NURBS meshing	6
6.1.2 3D_Roof_Modeler	7
6.1.3 ASIP	7
6.1.4 BSPslicer	7
6.1.5 GoCoPP	8
6.1.6 Feature-Preserving Offset Mesh Generation from Topology-Adapted Octrees	8
6.1.7 PlaneRegul	8
7 New results	9
7.1 Analysis	9
7.1.1 Sharp Feature Consolidation from Raw 3D Point Clouds via Displacement Learning	9
7.1.2 BpNet: Bézier Primitive Segmentation on 3D Point Clouds	9
7.1.3 KIBS: 3D Detection of Planar Roof Sections from a Single Satellite Image	10
7.1.4 Line-segment detection in images with fitting tolerance control	11
7.1.5 Registration for Urban Modeling Based on Linear and Planar Features	11
7.1.6 Robust Multiplatform Image/LiDAR Registration based on Geometric Primitives	12
7.2 Reconstruction	15
7.2.1 VoroMesh: Learning Watertight Surface Meshes with Voronoi Diagrams	15
7.2.2 Variational Shape Reconstruction via Quadric Error Metrics	15
7.2.3 Progressive Shape Reconstruction from Raw 3D Point Clouds	15
7.2.4 A Survey and Benchmark of Automatic Surface Reconstruction from Point Clouds	17
7.2.5 ConSLAM: Construction Data Set for SLAM	18
7.3 Approximation	20
7.3.1 Feature-Preserving Offset Mesh Generation from Topology-Adapted Octrees	20
7.3.2 Constructing Good Convex Polyhedral Decompositions from Polygons	20
7.3.3 Progressive Compression of Triangle Meshes	21
8 Bilateral contracts and grants with industry	22
8.1 Bilateral contracts with industry	22

9 Partnerships and cooperations	26
9.1 International research visitors	26
9.1.1 Visits of international scientists	26
9.2 European initiatives	27
9.2.1 H2020 projects	27
9.3 National initiatives	29
9.3.1 3IA Côte d'Azur	29
9.3.2 Inria challenge ROAD-AI - with CEREMA	30
9.3.3 PEPR NumPex	30
10 Dissemination	31
10.1 Promoting scientific activities	31
10.1.1 Scientific events: organisation	31
10.1.2 Scientific events: selection	31
10.1.3 Journal	31
10.1.4 Leadership within the scientific community	32
10.1.5 Scientific expertise	32
10.1.6 Research administration	32
10.2 Teaching - Supervision - Juries	32
10.2.1 Teaching	32
10.2.2 Supervision	32
10.2.3 Juries	33
11 Scientific production	33
11.1 Major publications	33
11.2 Publications of the year	34
11.3 Other	35

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

- Roberto Dyke [INRIA, Post-Doctoral Fellow]
- Raphael Sulzer [LUXCARTA and INRIA, Post-Doctoral Fellow, from Nov 2023]
- Raphael Sulzer [INRIA, Post-Doctoral Fellow, until Oct 2023]

PhD Students

- Moussa Bendjilali [ALTEIA, CIFRE, from Sep 2023]
- Marion Boyer [CNES and INRIA]
- Rao Fu [GEOMETRY FACTORY (EU ITN GRAPES)]
- Jacopo Iollo [INRIA (Défi ROAD-AI with CEREMA)]
- Nissim Maruani [INRIA (3IA Cote d'Azur)]
- Armand Zampieri [SAMP AI, CIFRE]

Technical Staff

- Abir Affane [INRIA and AI Verse (Inria Spinoff), Engineer, from May 2023]
- Jackson Campolattaro [Delf University, Engineer, from Feb 2023 until May 2023]
- Rahima Djahel [INRIA (EU BIM2TWIN), Engineer, from Mar 2023]
- Kacper Pluta [INRIA (EU BIM2TWIN), Engineer, until Aug 2023]

Interns and Apprentices

- Ashley Chen [INRIA, Intern, from Aug 2023 until Nov 2023, MIT]
- Gauthier Drumel [INRIA, Intern, from Apr 2023 until Jun 2023, Polytech'Nice]
- Kushagra Gupta [INRIA, Intern, from May 2023 until Jul 2023, IIT Delhi]
- Hanna Malet [INRIA, Intern, from Jun 2023 until Jul 2023, Université Dauphine-PSL]
- Vaibhav Seth [INRIA, Intern, from May 2023 until Jul 2023, IIT Delhi]

Administrative Assistant

- Florence Barbara [INRIA]

Visiting Scientists

- Gianmarco Cherchi [UNIV CAGLIARI, from Oct 2023 until Nov 2023]
- Thibault Lejemble [EPITA, from Jul 2023 until Sep 2023]

External Collaborators

- Andreas Fabri [GEOMETRY FACTORY]
- Johann Lussange [ENS PARIS, until May 2023]
- Sven Oesau [GEOMETRY FACTORY]
- Kacper Pluta [ESIEE, from Sep 2023]
- Laurent Rineau [GEOMETRY FACTORY]
- Mael Rouxel-Labbé [GEOMETRY FACTORY]

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, robotics, digital mapping and urban planning. The main outcome 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). 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 provides 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 Highlights of the year

- We are involved in strategic Inria partnerships: Inria challenge with CEREMA, collaboration with CNES on remote sensing, and project with Naval Group on the factory of the future (in collaboration with the Acentauri Inria project-team).
- We increased our collaborations with GeometryFactory with several research and development projects relating to the CGAL library (reliable computation of offsets, feature-preserving alpha wrapping, kinetic space partition and reconstruction, semantic classification of point sets, surface reconstruction via clustering quadric error metrics), and the organization of the CGAL developer meeting with 20 participants.
- We organized the third Inria-DFKI European Summer School on AI (IDESSAI 2023) in coordination with 3IA Côte d’Azur, with around 100 participants.
- Pierre Alliez is chairing the Inria evaluation committee since September 2023.

5.1 Awards

The open source project CGAL (Computational Geometry Algorithms) received the Test of Time award at the International Symposium on Computational Geometry (SoCG) 2023, in Dallas. CGAL is the fruit of close to 30 years of development, with significant contributions from Inria researchers. It features the latest innovations in computational geometry and geometry processing research.

6 New software, platforms, open data

6.1 New software

6.1.1 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](#)

Contact: Pierre Alliez

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

Partner: GeometryFactory

6.1.2 3D_Roof_Modeler

Name: 3D_Roof_Modeler

Keyword: Geometry Processing

Functional Description: This software takes as input (i) a set of 2D polygons representing the footprints of a building, and (ii) additional knowledge on each edge of the polygons to specify how the roof should connect to the facade wall, and provides as output a 3D model of the building in the form of a watertight polygonal surface mesh in which facets of roof, facade and ground are identified.

Author: Florent Lafarge

Contact: Florent Lafarge

6.1.3 ASIP

Name: Approximation of Shape in Images by Polygons

Keywords: Shape approximation, Computer vision, Vectorization, Polygons

Functional Description: This software allows the extraction and vectorization of objects in images with polygons. Departing from a polygonal partition that oversegments an image into convex cells, the algorithm refines the geometry of the partition while labeling its cells by a semantic class. The result is a set of polygons, each capturing an object in the image. The quality of a configuration is measured by an energy that accounts for both the fidelity to input data and the complexity of the output polygons. To efficiently explore the configuration space, splitting and merging operations are performed in tandem on the cells of the polygonal partition. The exploration mechanism is controlled by a priority queue that sorts the operations most likely to decrease the energy.

Publication: [hal-02526028](#)

Authors: Muxingzi Li, Florent Lafarge

Contact: Florent Lafarge

6.1.4 BSPslicer

Keywords: Computer vision, Geometric modeling

Functional Description: This software takes as input a set of planes detected from 3D data measurements and computes a compact plane arrangement that decomposes the space into convex polyhedra. The planes are processed in a specific order that reduces the complexity of the decomposition and of the underlying Binary Space Partitioning (BSP) tree.

Authors: Raphael Sulzer, Florent Lafarge

Contact: Florent Lafarge

6.1.5 GoCoPP

Name: Detection of planar primitives in point clouds by energy-based formulation

Keywords: Computer vision, Point cloud, Geometry Processing

Functional Description: This software detects planar primitives from unorganized 3D point clouds. Departing from an initial configuration, the algorithm refines both the continuous plane parameters and the discrete assignment of input points to them by seeking high fidelity, high simplicity and high completeness. Our key contribution relies upon the design of an exploration mechanism guided by a multi-objective energy function. The transitions within the large solution space are handled by five geometric operators that create, remove and modify primitives.

Publication: [hal-03621896](#)

Authors: Mulin Yu, Florent Lafarge

Contact: Florent Lafarge

Partner: CSTB

6.1.6 Feature-Preserving Offset Mesh Generation from Topology-Adapted Octrees

Keyword: Offset surface

Functional Description: A reliable method to generate offset meshes from input triangle meshes or triangle soups. Our method proceeds in two steps. The first step performs a Dual Contouring method on the offset surface, operating on an adaptive octree that is refined in areas where the offset topology is complex. Our approach substantially reduces memory consumption and runtime compared to isosurfacing methods operating on uniform grids. The second step improves the output Dual Contouring mesh with an offset-aware remeshing algorithm to reduce the normal deviation between the mesh facets and the exact offset. This remeshing process reconstructs concave sharp features and approximates smooth shapes in convex areas up to a user-defined precision. We show the effectiveness and versatility of our method by applying it to a wide range of input meshes. We benchmarked our method on the entire Thingi10k dataset: watertight, 2-manifold offset meshes are obtained for 100% of the cases.

Authors: Pierre Alliez, Daniel Zint, Mael Rouxel-Labbé

Contact: Pierre Alliez

Partner: GeometryFactory

6.1.7 PlaneRegul

Keywords: Geometry Processing, Computer vision

Functional Description: This software takes as input a 3D point cloud that samples the surface of physical objects and scenes, and fits a set of polygons on it while preserving geometric regularities existing in between polygons such as parallelism, orthogonality or co-linearity.

Author: Florent Lafarge

Contact: Florent Lafarge

7 New results

7.1 Analysis

7.1.1 Sharp Feature Consolidation from Raw 3D Point Clouds via Displacement Learning

Participants: Pierre Alliez, Florent Lafarge, Tong Zhao (*former PhD student, now at Dassault Systemes*), Florent Lafarge (*former PhD student, now at Shanghai Artificial Intelligence Laboratory*).

Detecting sharp features in raw point clouds is an essential step in designing efficient priors in several 3D Vision applications. This work presents a deep learning-based approach that learns to detect and consolidate sharp feature points on raw 3D point clouds (See Figure 1). We devise a multi-task neural network architecture that identifies points near sharp features and predicts displacement vectors toward the local sharp features. The so-detected points are thus consolidated via relocation. Our approach is robust against noise by utilizing a dynamic labeling oracle during the training phase. The approach is also flexible and can be combined with several popular point-based network architectures. Our experiments demonstrate that our approach outperforms the previous work in terms of detection accuracy measured on the popular ABC dataset. We show the efficacy of the proposed approach by applying it to several 3D Vision tasks. This work was published in the Computer Aided Geometric Design journal [14].

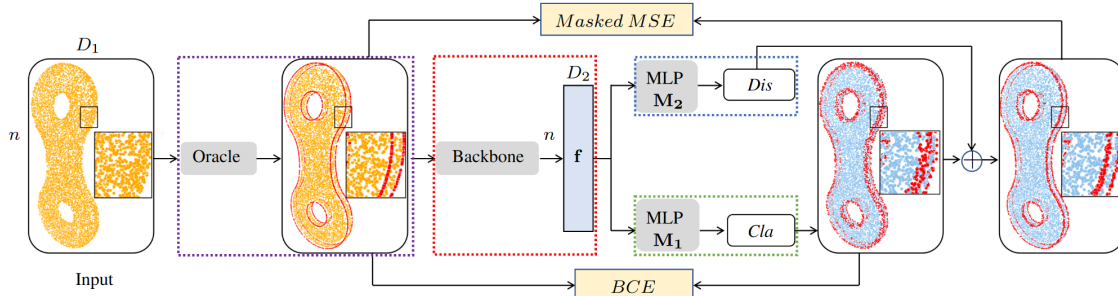


Figure 1: Sharp Feature Consolidation. The proposed architecture comprises four modules: *point-to-feature oracle* (purple dash), *point descriptor extractor* (red dash), *sharp feature detector* (green dash) and *displacement predictor* (blue dash). *Point-to-feature oracle* module is only used for training. It introduces noise to the input, computes the ground truth classification labels and the ground truth displacement vectors. A backbone is used to extract high dimensional representations \mathbf{f} (size: $n \times D_2$) in the *point descriptor extractor* module. **MLP** stands for multi-layer perceptron. *Cl*a and *Dis* represent binary sharp feature labels and displacement vectors, respectively. \oplus adds the corresponding predicted displacement vectors to detected sharp feature points to make them lie closer to the ground truth sharp features. *BCE* and *Masked MSE* are two losses calculated based on *Cl*a, *Dis* and the ground truth generated by the oracle.

7.1.2 BPNet: Bézier Primitive Segmentation on 3D Point Clouds

Participants: Rao Fu, Pierre Alliez, Cheng wen (*University of Sydney*), Qian Li (*Mimetic Inria project-team*), Xiao Xiao (*former postdoctoral fellow, now at Shanghai Jiao Tong University*).

This work proposes BPNet, a novel end-to-end deep learning framework to learn Bézier primitive segmentation on 3D point clouds. The existing works treat different primitive types separately, thus limiting them to finite shape categories. To address this issue, we seek a generalized primitive segmentation on

point clouds. Taking inspiration from Bézier decomposition on NURBS models, we transfer it to guide point cloud segmentation casting off primitive types. A joint optimization framework is proposed to learn Bézier primitive segmentation and geometric fitting simultaneously on a cascaded architecture (see Figure 2). Specifically, we introduce a soft voting regularizer to improve primitive segmentation and propose an auto-weight embedding module to cluster point features, making the network more robust and generic. We also introduce a reconstruction module where we successfully process multiple CAD models with different primitives simultaneously. We conducted extensive experiments on the synthetic ABC dataset and real-scan datasets to validate and compare our approach with different baseline methods. Experiments show superior performance over previous work in terms of segmentation, with a substantially faster inference speed. This work was presented at the International Joint Conference on Artificial Intelligence (IJCAI) [16].

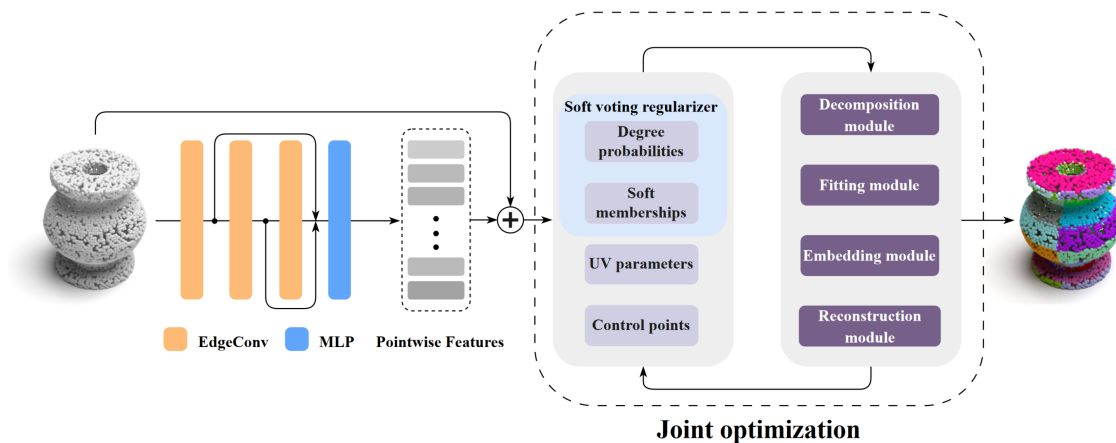


Figure 2: BpNet architecture. The network takes a point cloud as input. It outputs pointwise features followed by four modules to predict Bézier geometry and topology: (1) The decomposition module decomposes point clouds into multiple patches. (2) The fitting module regresses uv parameters for each point and control points for each Bézier patch. (3) The embedding module clusters pointwise features assigned to the same patch. (4) The reconstruction module re-evaluates the input point clouds from the above predictions.

7.1.3 KIBS: 3D Detection of Planar Roof Sections from a Single Satellite Image

Participants: Johann Lussange, Florent Lafarge, Yuliya Tarabalka (*Luxcarta*), Mulin Yu (*former PhD student, now at Shanghai Artificial Intelligence Laboratory*).

Reconstructing urban areas in 3D from satellite raster images has been a long-standing problem for both academical and industrial research. While automatic methods achieving this objective at a Level Of Details (LOD) 1 are efficient today, producing LOD2 models is still a scientific challenge. In particular, the quality and resolution of satellite data is too low to infer accurately the planar roof sections in 3D by traditional plane detection algorithms. Only inverse procedural modeling methods and skeletonization algorithms can approximate LOD2 models that often strongly deviate from the real scenes. In this work, we address this issue with KIBS (*Keypoints Inference By Segmentation*), a method that detects planar roof sections in 3D from a single satellite image (see Figure 3). By exploiting recent deep learning architectures with existing large-scale LOD2 databases produced by human operators, we achieve to both segment roof sections in images and lift keypoints enclosing these sections in 3D to form 3D-polygons. The output set of 3D-polygons can be used to reconstruct LOD2 models of buildings when combined to a plane assembly method. We demonstrate the potential of KIBS by reconstructing different urban areas in a few minutes, with a Jaccard index for the 2D segmentation of individual roof sections of approximately 80%,

and an altimetric error of the reconstructed LOD2 model inferior to 2 meters. A preprint version of this work is available on arXiv [22], currently under submission.



Figure 3: 3D detection of planar roof sections. KIBS detects the planar roof sections of buildings in large-scale satellite image under the form of 3D polygons. In these visual results, polygons are backprojected onto the input image (blue shapes). Red points refer to the vertices of the polygons.

7.1.4 Line-segment detection in images with fitting tolerance control

Participants: Marion Boyer, Florent Lafarge, David Youssefi (CNES).

This work proposes a geometric reformulation of the popular line segment detection problem from images. In this reformulation, a line segment detected in the image does not necessarily refer to line-like structures of objects, but can also approximate the local geometry of organic shapes within a fitting tolerance. Our approach simultaneously seeks high fidelity, completeness and regularity in line segment configurations with an energy minimization framework. This approach brings precision and global consistency in the landscape of existing algorithms (see Figure 4). We show its competitiveness against different families of methods on various datasets, especially for producing regularized configurations of line segments. Operating from an image gradient map, our method can be combined with recent deep learning approaches for the fitting step or simply used as post-processing for refining a result. This work is currently under submission.

7.1.5 Registration for Urban Modeling Based on Linear and Planar Features

Participants: Rahima Djahel, Pascal Monasse (Ecole des Ponts ParisTech), Bruno Vallet (IGN).

The production of a Building Information Model (BIM) from an existing asset is currently expensive and needs automation of the registration of the different acquisition data, including the registration

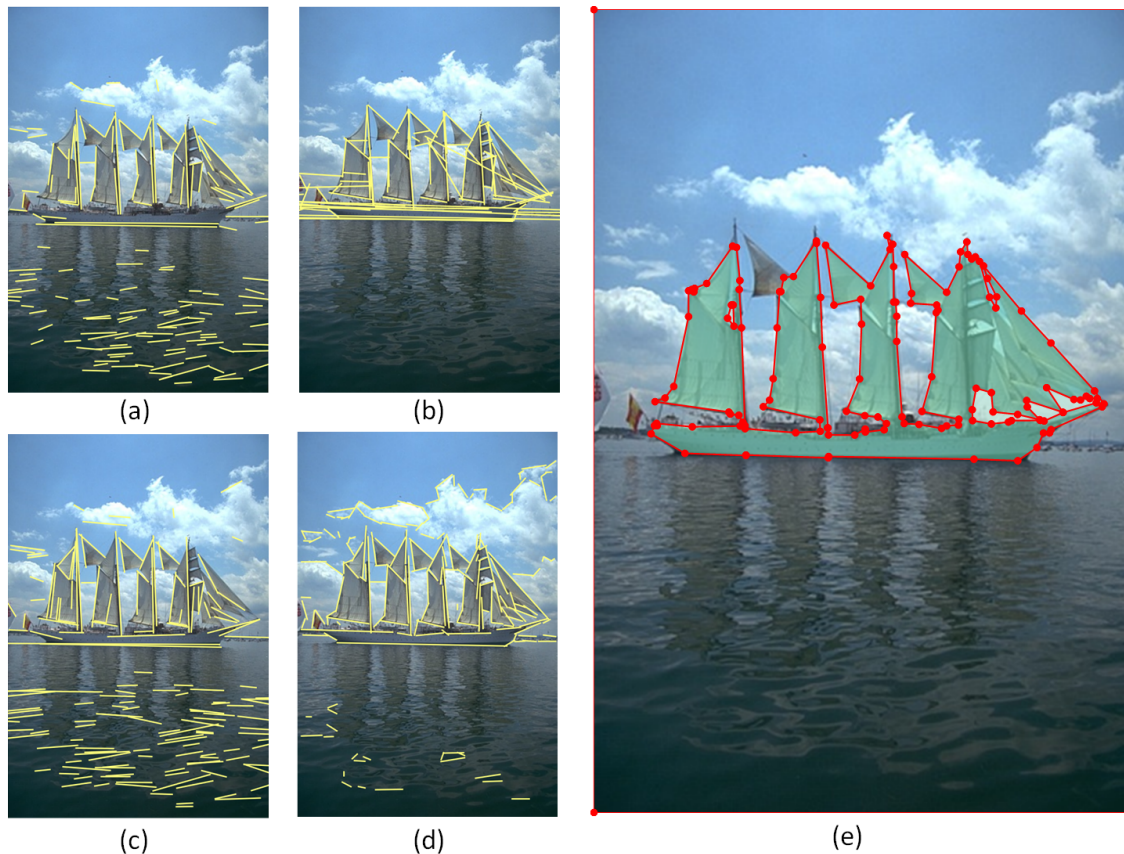


Figure 4: Line-segment detection in images. Fitting tolerance control allows us to accurately locate line segments on line-like structures while approximating curved shapes by chains of line segments (d). Existing methods, ie incremental mechanism (a), end-to-end neural approach (b) and hybrid model (c), sparsely describes freeform parts of the boat and overly detects on regular structures. Combined with a vectorization method that assembles line-segments into intersection-free polygons, our algorithm can be used to polygonize the silhouette of boat with high accuracy (e).

of indoor and outdoor data. This kind of registration is considered a challenging problem, especially when both data sets are acquired separately and use different types of sensors (see Figure 5). Besides, comparing a BIM to as-built data is an important factor to perform building progress monitoring and quality control. To carry out this comparison, the data sets must be registered. In order to solve both registration problems, we introduce two efficient algorithms. The first offers a potential solution for indoor/outdoor registration based on heterogeneous features (openings and planes). The second is based on linear features and proposes a potential solution for LiDAR data/BIM model registration. The common point between the approaches consists in the definition of a global robust distance between two segment sets and the minimization of this distance based on the RANSAC paradigm, finding the rigid geometric transformation that is the most consistent with all the information in the data sets. This work was presented at the European Workshop on Visual Information Processing [19].

7.1.6 Robust Multiplatform Image/LiDAR Registration based on Geometric Primitives

Participants: Rahima Djahel, Pascal Monasse (*Ecole des Ponts ParisTech*), Bruno Vallet (*IGN*).

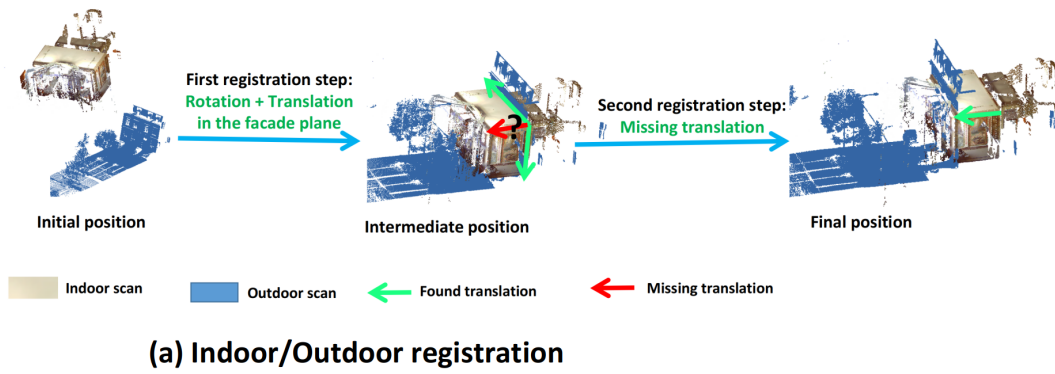


Figure 5: Heterogeneous data registration for urban modeling: (a) indoor / outdoor data registration to simplify the production of a Building Information Model (BIM) from an existing asset. (b) BIM model / LiDAR data registration to simplify the performance of building progress monitoring and Quality control.

Many applications benefit from data acquisition from various platforms and with different modalities. A global approach can take advantage of such heterogeneous data sources by combining them to take advantage of their semantic and geometric complementary. The challenge tackled in this work is to register such data that can be heterogeneous in both platform and modality. We propose a generic primitive based registration algorithm that can be adapted to tackle specific heterogeneous registration problems. The type of primitives and the extraction method is chosen according to the properties of the data, the environment and the acquisition platforms. The proposed generic registration relies on the definition of a global energy between the primitives extracted from the data-sets to register, and on a robust method to optimize this energy based on the RANSAC paradigm. The obtained results show the ability of our algorithms to deal with the registration of heterogeneous data and to solve challenging problems including indoor/outdoor registration, image/LiDAR registration and aerial/terrestrial registration. See Figure 6. This work is being prepared for submission.

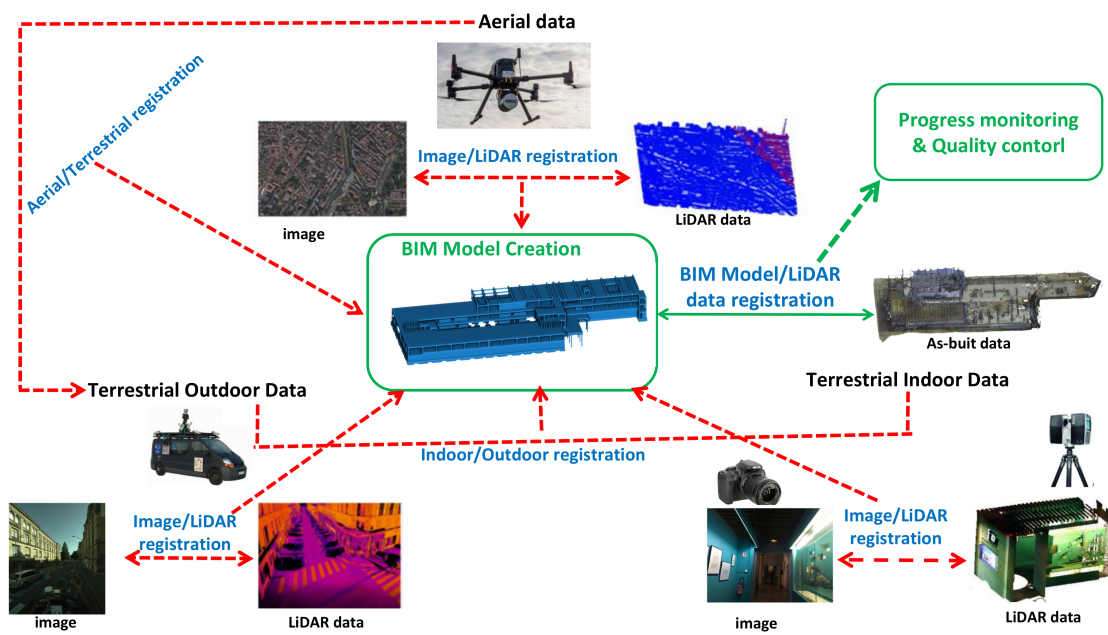


Figure 6: Overview of the registration methods studied in this work and their possible applications.

7.2 Reconstruction

7.2.1 VoroMesh: Learning Watertight Surface Meshes with Voronoi Diagrams

Participants: Nissim Maruani, Pierre Alliez, Mathieu Desbrun (*Geomerix Inria project team*), Roman Klokov (*Ecole Polytechnique*), Maks Ovsjanikov (*Ecole Polytechnique*).

In stark contrast to the case of images, finding a concise, learnable discrete representation of 3D surfaces remains a challenge. In particular, while polygon meshes are arguably the most common surface representation used in geometry processing, their irregular and combinatorial structure often make them unsuitable for learning-based applications. In this work, we present VoroMesh, a novel and differentiable Voronoi-based representation of watertight 3D shape surfaces. From a set of 3D points (called generators) and their associated occupancy, we define our boundary representation through the Voronoi diagram of the generators as the subset of Voronoi faces whose two associated (equidistant) generators are of opposite occupancy: the resulting polygon mesh forms a watertight approximation of the target shape's boundary (see Figure 7). To learn the position of the generators, we propose a novel loss function, dubbed VoroLoss, that minimizes the distance from ground truth surface samples to the closest faces of the Voronoi diagram which does not require an explicit construction of the entire Voronoi diagram. A direct optimization of the VoroLoss to obtain generators on the Thingi32 dataset demonstrates the geometric efficiency of our representation compared to axiomatic meshing algorithms and recent learning-based mesh representations. We further use VoroMesh in a learning-based mesh prediction task from input SDF grids on the ABC dataset, and show comparable performance to state-of-the-art methods while guaranteeing closed output surfaces free of self-intersections. This work was presented at the International Conference on Computer Vision [17].

7.2.2 Variational Shape Reconstruction via Quadric Error Metrics

Participants: Tong Zhao, Pierre Alliez, Laurent Busé (*AROMATH Inria project-team*), David Cohen-Steiner (*DATASHAPE Inria project-team*), Tamy Boubekeur (*Adobe Research*), Jean-Marc Thiery (*Adobe Research*).

Inspired by the strengths of quadric error metrics initially designed for mesh decimation, we propose a concise mesh reconstruction approach for 3D point clouds. Our approach proceeds by clustering the input points enriched with quadric error metrics, where the generator of each cluster is the optimal 3D point for the sum of its quadric error metrics. This approach favors the placement of generators on sharp features, and tends to equidistribute the error among clusters (see Figure 8). We reconstruct the output surface mesh from the adjacency between clusters and a constrained binary solver. We combine our clustering process with an adaptive refinement driven by the error. Compared to prior art, our method avoids dense reconstruction prior to simplification and produces immediately an optimized mesh. This work was presented at the ACM SIGGRAPH conference [18]. We are currently working on a new component for the CGAL library, with support from the Google summer of code program and cooperation with GeometryFactory.

7.2.3 Progressive Shape Reconstruction from Raw 3D Point Clouds

Participants: Tong Zhao.

This Ph.D. thesis funded by 3IA Côte d'Azur was co-advised by Pierre Alliez and Laurent Busé (Aromath Inria project-team).




Figure 7: VoroMesh at a glance. To fit our representation to (a) a target shape (top: curve; bottom: surface), (b) the positions of generators around the shape boundary are optimized based on our VoroLoss, before a binary occupancy is assigned to each generator; (c) the VoroMesh is then extracted as a subset of the Voronoi diagram of the generators containing only the Voronoi faces between pairs of generators of opposite occupancy.

With the enthusiasm for digital twins in the fourth industrial revolution, surface reconstruction from raw 3D point clouds is increasingly needed while facing multifaceted challenges. Advanced data acquisition devices makes it possible to obtain 3D point clouds with multi-scale features. Users expect controllable surface reconstruction approaches with meaningful criteria such as preservation of sharp features or satisfactory accuracy-complexity tradeoffs. This thesis addresses these issues by contributing several approaches to the problem of surface reconstruction from defect-laden point clouds. We first propose the notion of progressive discrete domain for global implicit reconstruction approaches that refines and optimizes a discrete 3D domain in accordance to both input and output, and to user-defined criteria. Based on such a domain discretization, we devise a progressive primitive-aware surface reconstruction approach with capacity to refine the implicit function and its representation, in which the most ill-posed parts of the reconstruction problem are postponed to later stages of the reconstruction, and where the fine geometric details are resolved after discovering the topology. Secondly, we contribute a deep learning-based approach that learns to detect and consolidate sharp feature points on raw 3D point clouds, whose results can be taken as an additional input to consolidate sharp features for the previous reconstruction approach. Finally, we contribute a coarse-to-fine piecewise smooth surface reconstruction approach that proceeds by clustering quadric error metrics. This approach outputs a simplified reconstructed surface mesh, whose vertices are located on sharp features and whose connectivity is




Figure 8: Variational Shape Reconstruction. The clustering of the points is randomly initialized, then alternates partitioning and generator updating. Some generators relocate to sharp features after one iteration. New generators are then added. The clustering converges after five iterations. A set of candidate edges is derived from the adjacency between clusters, and candidate facets (red) are generated. The output mesh is reconstructed via a constrained binary solver that selects a subset of the red facets.

solved by a binary problem solver. In summary, this thesis seeks for effective surface reconstruction from a global and progressive perspective. By combining multiple priors and designing meaningful criteria, the contributed approaches can deal with various defects and multi-scale features [20].

7.2.4 A Survey and Benchmark of Automatic Surface Reconstruction from Point Clouds

Participants: Raphael Sulzer, Loic Landrieu (*Ecole des Ponts Paristech*), Renaud Marlet (*Valeo.ai*), Bruno Vallet (*IGN*).

We survey and benchmark traditional and novel learning-based algorithms that address the problem of surface reconstruction from point clouds. Surface reconstruction from point clouds is particularly challenging when applied to real-world acquisitions, due to noise, outliers, non-uniform sampling and missing data. Traditionally, different handcrafted priors of the input points or the output surface have been proposed to make the problem more tractable. However, hyperparameter tuning for adjusting priors to different acquisition defects can be a tedious task. To this end, the deep learning community has recently addressed the surface reconstruction problem. In contrast to traditional approaches, deep surface reconstruction methods can learn priors directly from a training set of point clouds and corresponding true surfaces. In our survey, we detail how different handcrafted and learned priors affect the robustness of methods to defect-laden input and their capability to generate geometric and topologically

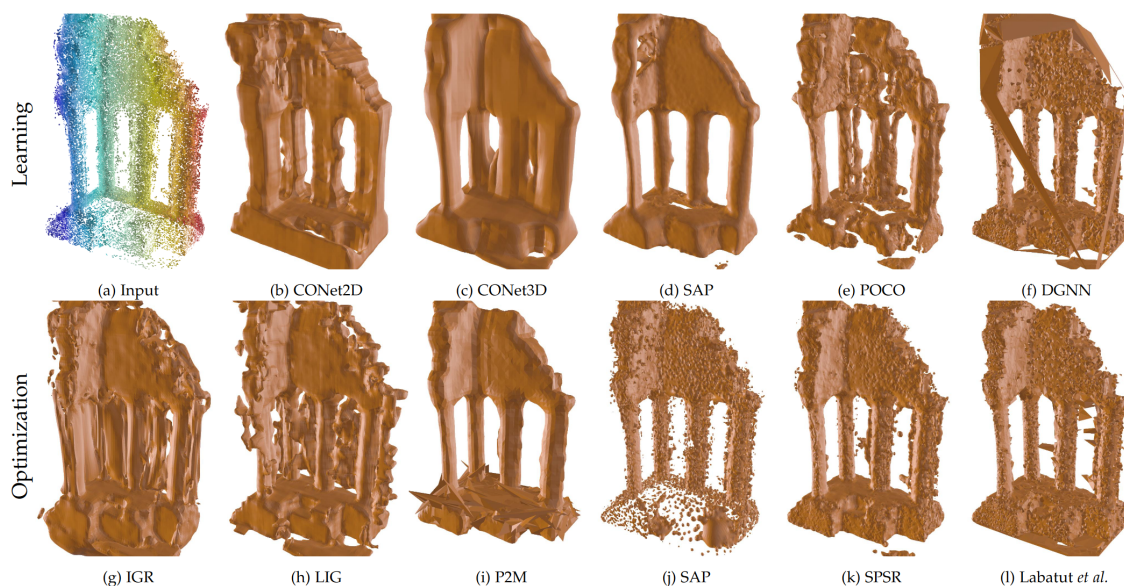


Figure 9: Benchmark of Automatic Surface Reconstruction from Point Clouds. Learning- and optimization-based reconstructions of a real-world MVS point cloud. The learning methods (b-j) were trained on synthetic MVS scans from ShapeNet. Optimization-based methods are optimized per scan using standard settings. The two traditional methods SPSR (k) and Labatut et al. (l) reconstruct the visually best surfaces.

accurate reconstructions. In our benchmark, we evaluate the reconstructions of several traditional and learning-based methods on the same grounds (see Figure 9). We show that learning-based methods can generalize to unseen shape categories, but their training and test sets must share the same point cloud characteristics. We also provide the code and data to compete in our benchmark and to further stimulate the development of learning-based surface reconstruction ([link](#)). A preprint version of this work is available on arXiv [21], currently under revision at TPAMI.

7.2.5 ConSLAM: Construction Data Set for SLAM

Participants: Kacper Pluta, Pierre Alliez, Maciej Trzeciak (*University of Cambridge*), Yasmin Fathy (*University of Cambridge*), Ioannis Brilakis (*University of Cambridge*), Lucio Alcalde (*Laing O'Rourke*), Stanley Chee (*Laing O'Rourke*), Antony Bromley (*Laing O'Rourke*).


This work presents a dataset collected periodically on a construction site. The dataset aims to evaluate the performance of SLAM algorithms used by mobile scanners or autonomous robots. It includes ground-truth scans of a construction site collected using a terrestrial laser scanner along with five sequences of spatially registered and time-synchronized images, LiDAR scans and inertial data coming from our prototypical hand-held scanner (see Figure 10). We also recover the ground-truth trajectory of the mobile scanner by registering the sequential LiDAR scans to the ground-truth scans and show how to use a popular software package to measure the accuracy of SLAM algorithms against our trajectory automatically. To the best of our knowledge, this is the first publicly accessible dataset consisting of periodically collected sequential data on a construction site. This work was published in the *Journal of Computing in Civil Engineering* [12].




(a) Close-up view of the construction site from GT_1



(b) Corridor view of the construction site from GT_1



(c) Close-up view of the construction site from GT_2



(d) Corridor view of the construction site from GT_2



(e) Close-up view of the construction site from GT_3



(f) Corridor view of the construction site from GT_3




(g) Close-up view of the construction site from GT_4



(h) Corridor view of the construction site from GT_4



(i) Close-up view of the construction site from GT_5



(j) Corridor view of the construction site from GT_5

Figure 10: Close-up visualization of ground-truth datasets. The images include also the LiDAR positions depicted by red spheres. The positions have been connected to provide approximated paths/trajectories.

7.3 Approximation

7.3.1 Feature-Preserving Offset Mesh Generation from Topology-Adapted Octrees

Participants: Nissim Maruani, Pierre Alliez, Daniel Zint (*New York University*), Mael Rouxel-Labbé (*GeometryFactory*).

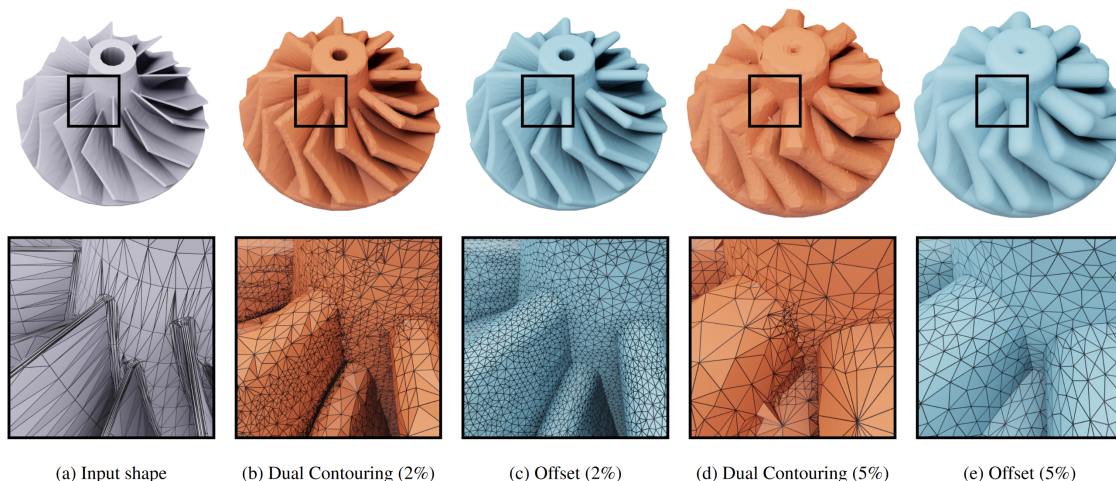


Figure 11: Feature-Preserving Offset Mesh Generation. (a) The input consists of a mesh and two main user-defined parameters: offset radius and maximum normal deviation; (b,d) The topology of the initial offset mesh is computed with Dual Contouring on a topology-adapted octree; (c,e) A high-quality offset mesh is obtained through remeshing.

We introduce a reliable method to generate offset meshes from input triangle meshes or triangle soups. Our method proceeds in two steps. The first step performs a Dual Contouring method on the offset surface, operating on an adaptive octree that is refined in areas where the offset topology is complex. Our approach substantially reduces memory consumption and runtime compared to isosurfacing methods operating on uniform grids. The second step improves the output Dual Contouring mesh with an offset-aware remeshing algorithm to reduce the normal deviation between the mesh facets and the exact offset. This remeshing process reconstructs concave sharp features and approximates smooth shapes in convex areas up to a user-defined precision (see Figure 11). We show the effectiveness and versatility of our method by applying it to a wide range of input meshes. We also benchmark our method on the Thingi10k dataset: watertight and topologically 2-manifold offset meshes are obtained for 100% of the cases. This work was presented at the Eurographics Symposium on Geometry Processing [15].

7.3.2 Constructing Good Convex Polyhedral Decompositions from Polygons

Participants: Raphael Sulzer, Florent Lafarge.

Decomposing a 3D domain into convex polyhedra from a set of polygons is a long standing problem in piecewise planar geometry. Existing works target either quality of the output decomposition or performance of its construction. We present an algorithm that addresses both. Departing from the algorithm of Murali and Funkhouser that iteratively slices polyhedra that contain input polygons only, we propose a series of technical improvements that allow us to produce more compact and meaningful partitions, in less time, than existing methods. In particular, one of our key ideas is to give priority to slicing operations that minimize the number of intersection computations necessary to construct the

decomposition. We also propose a fast analysis of relative spatial positioning of polygons via point sampling and a cell simplification process based on a Binary Space Partitioning tree analysis. We evaluate and compare our approach with the existing polyhedral decomposition methods on a collection of input polygons captured from real-world data measurements. We also show its efficiency on two applicative tasks: compact mesh reconstruction and convex decomposition of 3D meshes (see Figure 12). This work is currently under submission.

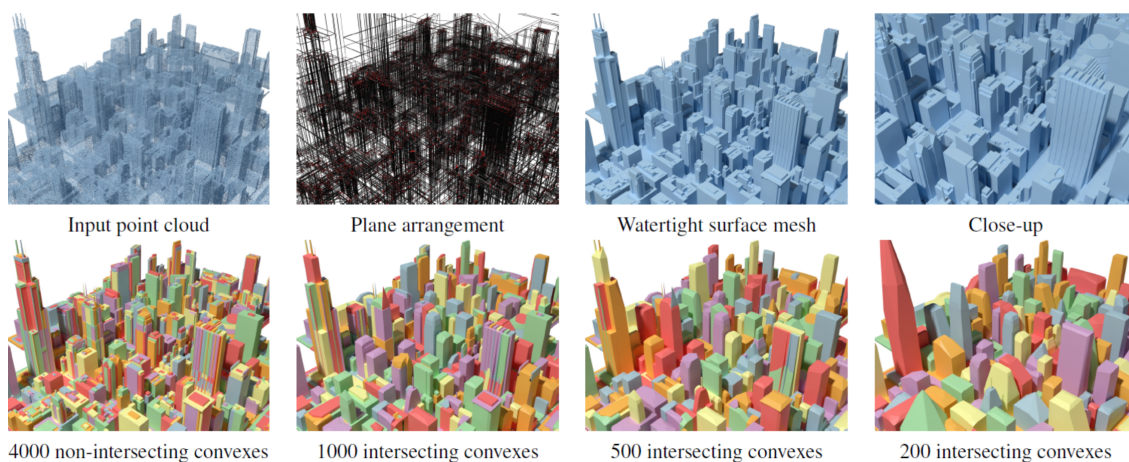


Figure 12: Convex Polyhedral Decompositions from Polygons. Our plane arrangement built from polygons detected from the input point cloud allows the reconstruction of buildings with many geometric details at the scale of a city. The 4,000 non-intersecting convexes that finely describe this urban landscape can be simplified by aggregating the convexes and authorizing their overlap. A simplification with 1,000 convexes typically removes facade and roof details. At 500 convexes, only complex buildings are still described by more than one convex cell. Reducing to 200 convexes allows us to represent each building block or skyscraper by a single convex cell.

7.3.3 Progressive Compression of Triangle Meshes

Participants: Pierre Alliez, Lucas Dubouchet (*Former Inria research engineer*), Vincent Vidal (*LIRIS Laboratory, Lyon*), Guillaume Lavoué (*LIRIS Laboratory, Lyon*).

This paper details the first publicly available implementation of the progressive mesh compression algorithm described in the paper entitled "Compressed Progressive Meshes" (R. Pajarola and J. Rossignac, *IEEE Transactions on Visualization and Computer Graphics*, 6 (2000)). Our implementation is generic, modular, and includes several improvements in the stopping criteria and final encoding. Given an input 2-manifold triangle mesh, an iterative simplification is performed, involving batches of edge collapse operations guided by an error metric. During this compression step, all the information necessary for the reconstruction (at the decompression step) is recorded and compressed using several key features: geometric quantization, prediction, and spanning tree encoding. Our implementation enabled to carry out an experimental comparison of several settings for the key parameters of the algorithm: the local error metric, the position type of the resulting vertex (after collapse), and the geometric predictor. The proposed implementation is publicly available through the MEPP2 platform. The algorithm can be used either as a command-line executable or integrated into the MEPP2 GUI. The source code is written in C++ and is accessible on the IPOL web page of this article, as well as on the GitHub page of MEPP2 (MEPP-team/MEPP2 project 2). This work was published in the *Image Processing On Line (IPOL) journal* [13].

8 Bilateral contracts and grants with industry

8.1 Bilateral contracts with industry

AlteIA - Cifre thesis

Participants: Pierre Alliez, Moussa Bendjilali.

This Cifre Ph.D. thesis project, entitled “From 3D Point Clouds to Cognitive 3D Models”, started in September 2023, for a total duration of 3 years.

Context. Digital twinning of large-scale scenes and industrial sites requires acquiring the geometry and appearance of terrains, large-scale structures and objects. In this thesis, we assume that the acquisition is performed by an aerial laser scanner (LiDAR) attached to a helicopter. The generated 3D point clouds record the point coordinates, intensities of reflectivity, and possibly other attributes such as colors. These raw measurement data can be massive, in the order of 70 points per meter square, summing up to several hundred million points for scenes extending over kilometers. Analyzing large-scale scenes requires (1) enriching these data with an abstraction layer, (2) understanding the structure and semantic of the scenes, (3) producing so-called cognitive 3D models that are searchable, and (4) detecting the changes in the scenes over time.

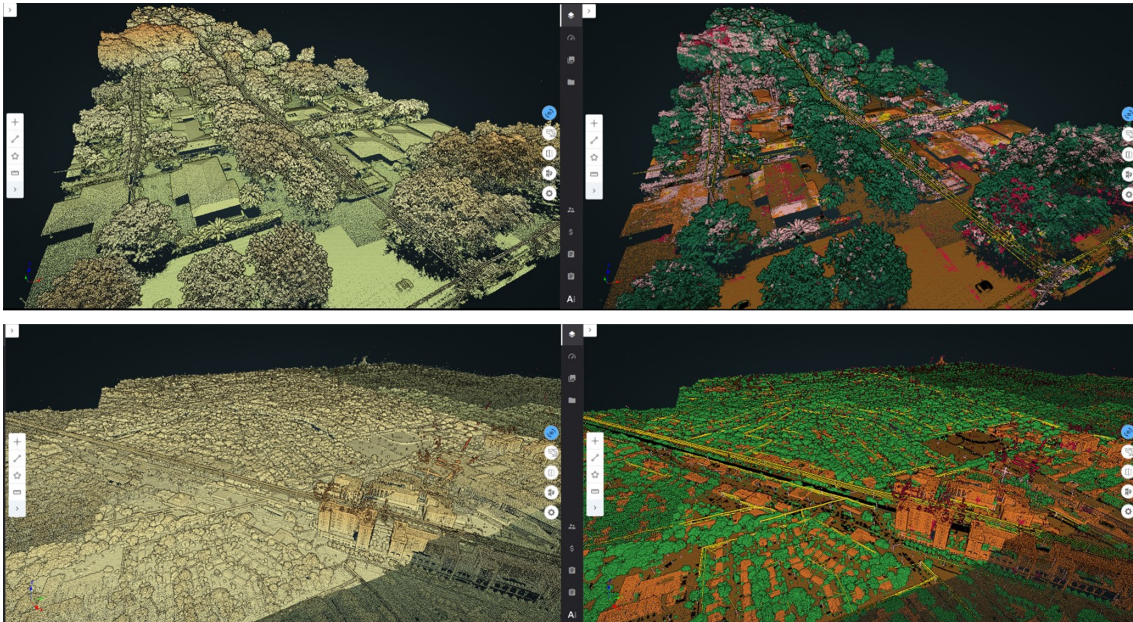


Figure 13: Large-scale scene of an electric infrastructure. The 3D point clouds are obtained with a helicopter equipped with a LiDAR sensor. Left: point cloud colored by the elevation; Right: dense semantic segmentation prediction (of a deep learning model) with seven classes of interest: Man made object, ground, vegetation, building, noise, powerline, pole.

We want to explore novel algorithms for addressing the aforementioned requirements. The input is assumed to be a LiDAR 3D point cloud recording the point coordinates and surface reflectivity intensity of a point set sampled over a physical scene. The objectives are multifaceted:

- Semantic segmentation. We wish to retrieve the semantic of the scene as well as the location, pose and semantic classes of objects of interest. Ideally, we seek for dense semantic segmentation, where the class of each sample point is inferred in order to yield a complete description of the physical scene.

- Instance segmentation. We wish to distinguish each instance of a semantic class so that all objects of a scene can be enumerated.
- Abstraction. We wish to detect canonical primitives (planes, cylinders, spheres, etc.) that approximate (abstract) the objects of the physical scenes, and understand the relationships between them (e.g., adjacency, containment, parallelism, orthogonality, coplanarity, concentricity). Abstraction goes beyond the notion of geometric simplification, with the capability to forget the exquisite details present in dense 3D point clouds.
- Object detection. Knowing a library of CAD models, we wish to detect and enumerate all occurrences of an object of interest used as a query. We assume some known information either on the primitive shape and its parametric description, or on the shape of the query.
- Change detection. Series of LiDAR data acquired over time are required for monitoring changes of industrial sites. The objective is to detect and characterize changes in the physical scene, with or without any explicit taxonomy.
- Levels of details. The above tasks should be tackled with levels of details (LODs), both for the geometry and semantic. LODs will relate to geometric tolerance errors and to a hierarchy of semantic ontologies.

CNES - Airbus

Participants: Marion Boyer, Florent Lafarge.

The goal of this collaboration is to design an automatic pipeline for modeling in 3D urban scenes under a CityGML LOD2 formalism using the new generation of stereoscopic satellite images, in particular the 30cm resolution images offered by Pléiades Néo satellites (see Figure 14). Traditional methods usually start with a semantic classification step followed by a 3D reconstruction of objects composing the urban scene. Too often in traditional approaches, inaccuracies and errors from the classification phase are propagated and amplified throughout the reconstruction process without any possible subsequent correction. In contrast, our main idea consists in extracting semantic information of the urban scene and in reconstructing the geometry of objects simultaneously. This project started in October 2022, for a total duration of 3 years.




Figure 14: Samples of Pléiades Néo satellite images with Ground Truth building wireframes (yellow lines). The goal of this collaboration is to automatically extract building wireframes and lift them in 3D to form 3D models under a CityGML LOD2 formalism.

Luxcarta

Participants: Raphael Sulzer, Florent Lafarge.

The main goal of this collaboration with Luxcarta is to develop new algorithms that improve the geometry and topology of 3D models of buildings produced by the company. The algorithms should detect and enforce (i) geometric regularities in 3D models, such as parallelism of roof edges or connection of polygonal facets at exactly one vertex, and (ii) simplify models, e.g. by removing undesired facets or incoherent roof typologies. These two objectives will have to be fulfilled under the constraint that the geometric accuracy of the solution must remain close to the one of the input models. Assuming the input 3D models are valid watertight polygon surface meshes, we will investigate metrics to quantify the geometric regularities and the simplicity of the input models and will propose a formulation that allows both continuous and discrete modifications. Continuous modifications will aim to better align vertices and edges, in particular to reinforce geometric regularities. Discrete modifications will handle changes in roof topology with, for instance, removal or creation of roof sections or new adjacencies between them. The output model will have to be valid 2-manifold watertight polygon surface meshes. We will also investigate optimization mechanisms to explore the large solution space of the problem efficiently. This project started in November 2023, for a total duration of 1 year.

Naval group

Participants: Pierre Alliez.

This project is in collaboration with the Acentauri project-team (a research engineer is co-advised since November 2023). The context is that of the factory of the future for Naval Group, for submarines and surface ships. As input, we have a digital model (e.g., a frigate), the equipment assembly schedule and measurement data (images or Lidar). We wish to monitor the assembly site to compare the "as-designed" with the "as-built" model. The main problem is a need for coordination on the construction sites for decision making and planning optimization. We wish to enable following the progress of a real project and to check its conformity with the help of a digital twin. Currently, since Naval Group needs to verify on site, rounds of visits are required to validate the progress as well as the equipment. These rounds are time consuming and not to mention the constraints of the site, such as the temporary absence of electricity or the numerous temporary assembly and safety equipment. The objective of the project is to automate the monitoring, with sensor intelligence to validate the work. Fixed sensor systems (e.g. cameras or Lidar) or mobile (e.g. drones) sensors will be used, with the addition of smartphones/tablets carried by operators on site.

Samp AI - Cifre thesis

Participants: Armand Zampieri, Pierre Alliez.

This project (Cifre Ph.D. thesis) investigates algorithms for progressive, random access compression of 3D point clouds for digital twinning of industrial sites (see figure 15). The context is as follows. Laser scans of industrial facilities contain massive amounts of 3D points which are very reliable visual representations of these facilities. These points represent the positions, color attributes and other metadata about object surfaces. Samp turns these 3D points into intelligent digital twins that can be searched and interactively visualized on heterogeneous clients, via streamed services on the cloud. The 3D point clouds are first structured and enriched with additional attributes such as semantic labels or hierarchical group identifiers. Efficient streaming motivates progressive compression of these 3D point clouds, while local search or visualization queries call for random accessible capabilities. Given an enriched 3D point cloud

as input, the objective is to generate as output a compressed version with the following properties: (1) Random accessible compressed (RAC) file format that provides random access (not just sequential access) to the decompressed content, so that each compressed chunk can be decompressed independently, depending on the area in focus on the client side. (2) Progressive compression: the enriched point cloud is converted into a stream of refinements, with an optimized rate-distortion tradeoff allowing for early access to a faithful version of the content during streaming. (3) Structure-preserving: all structural and semantic information are preserved during compression, allowing for structure-aware queries on the client. (4) Lossless at the end: allows the original data to be perfectly reconstructed from the compressed data, when streaming is completed.




Figure 15: Digital twin of an industrial site. Image courtesy Samp AI.

Urbassist

Participants: Florent Lafarge.

This project aims to develop and mature a roof modeler software for the company Urbassist. Given a polygonal footprint of a building and some information related to (i) the height of the facade, (ii) the slope of the roof, and (iii) the type of connexion between the facade walls and the planar roof sections, we implemented a skeletonization algorithm that reconstructs the building in 3D under the form of a polygonal surface mesh. The software has been tested and matured by using random footprints from the French cadastral map database. This project started in March 2021, for a total duration of 2 years.

AI Verse

Participants: Pierre Alliez, Abir Affane.

In collaboration with Guillaume Charpiat (Tau Inria project-team, Saclay) and AI Verse (Inria spin-off).

We are currently pursuing our collaboration with AI Verse on statistics for improving sampling quality, applied to the parameters of a generative model for 3D scenes. A postdoctoral fellow (Abir Affane) funded for three years has started this year. The research topic is described next. The AI Verse technology is devised to create infinitely random and semantically consistent 3D scenes. This creation is fast, consuming less than 4 seconds per labeled image. From these 3D scenes, the system is able to build quality synthetic images that come with rich labels that are unbiased unlike manually annotated labels (Figure 16). As for real data, no metric exists to evaluate the performance of the synthetic datasets to train a neural network. We thus tend to favor the photorealism of the images but such a criterion is far from being the best. The current technology provides a means to control a rich list of additional parameters (quality of lighting, trajectory and intrinsic parameters of the virtual camera, selection and placement of assets, degrees of occlusion of the objects, choice of materials, etc). Since the generation engine can modify all of these parameters at will to generate many samples, we will explore optimization methods for improving the sampling quality. Most likely, a set of samples generated randomly by the generative model does not cover well the whole space of interesting situations, because of unsuited sampling laws or of sampling realization issues in high dimensions. The main question is how to improve the quality of this generated dataset, that one would like to be close somehow to the given target dataset (consisting of examples of images that one would like to generate). For this, statistical analyses of these two datasets and of their differences are required, in order to spot possible issues such as strongly under-represented areas of the target domain. Then, sampling laws can be adjusted accordingly, possibly by optimizing their hyper-parameters, if any.



Figure 16: Examples of beds generated by AI Verse technology.

9 Partnerships and cooperations

Participants: Pierre Alliez, Florent Lafarge.

9.1 International research visitors

9.1.1 Visits of international scientists

Other international visits to the team

Gianmarco Cherchi

Status: Researcher

Institution of origin: University of Cagliari

Country: Italy

Dates: From October 2023 until November 2023

Context of the visit: Collaboration with the team on low-complexity mesh reconstruction from 3D point clouds

Mobility program/type of mobility: Research stay

Thibault Lejemble

Status: Researcher

Institution of origin: Epita

Country: France

Dates: From July 2023 until September 2023

Context of the visit: Collaboration with the team on semantic analysis of 3D point clouds

Mobility program/type of mobility: Research stay

9.2 European initiatives

9.2.1 H2020 projects

BIM2TWIN

Participants: Pierre Alliez, Kacper Pluta, Rahima Djahel.

[BIM2TWIN project on cordis.europa.eu](https://cordis.europa.eu/project/BIM2TWIN)

Title: BIM2TWIN: Optimal Construction Management & Production Control

Duration: From November 1, 2020 to April 30, 2024

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- SPADA CONSTRUCTION (SPADA CONSTRUCTION), France
- AARHUS UNIVERSITET (AU), Denmark
- FIRA GROUP OY, Finland
- INTSITE LTD (INTSITE), Israel
- THE CHANCELLOR MASTERS AND SCHOLARS OF THE UNIVERSITY OF CAMBRIDGE, United Kingdom
- ORANGE SA (Orange), France
- UNISMART - FONDAZIONE UNIVERSITA DEGLI STUDI DI PADOVA (UNISMART), Italy
- FUNDACION TECNALIA RESEARCH & INNOVATION (TECNALIA), Spain
- TECHNISCHE UNIVERSITAET MUENCHEN (TUM), Germany

- IDP INGENIERIA Y ARQUITECTURA IBERIA SL (IDP), Spain
- SITEDRIVE OY, Finland
- UNIVERSITA DEGLI STUDI DI PADOVA (UNIPD), Italy
- FIRA OY, Finland
- RUHR-UNIVERSITAET BOCHUM (RUHR-UNIVERSITAET BOCHUM), Germany
- CENTRE SCIENTIFIQUE ET TECHNIQUE DU BATIMENT (CSTB), France
- DANMARKS TEKNISKE UNIVERSITET (TECHNICAL UNIVERSITY OF DENMARK DTU), Denmark
- TECHNION - ISRAEL INSTITUTE OF TECHNOLOGY, Israel
- UNIVERSITA POLITECNICA DELLE MARCHE (UNIVPM), Italy
- ACCIONA CONSTRUCCION SA, Spain
- SIEMENS AKTIENGESELLSCHAFT, Germany

Inria contact: Pierre Alliez

Coordinator: Bruno Fies, CSTB

Summary: BIM2TWIN aims to build a Digital Building Twin (DBT) platform for construction management that implements lean principles to reduce operational waste of all kinds, shortening schedules, reducing costs, enhancing quality and safety and reducing carbon footprint. BIM2TWIN proposes a comprehensive, holistic approach. It consists of a (DBT) platform that provides full situational awareness and an extensible set of construction management applications. It supports a closed loop Plan-Do-Check-Act mode of construction. Its key features are:

1> Grounded conceptual analysis of data, information and knowledge in the context of DBTs, which underpins a robust system architecture

2> A common platform for data acquisition and complex event processing to interpret multiple monitored data streams from construction site and supply chain to establish real-time project status in a Project Status Model (PSM)

3> Exposure of the PSM to a suite of construction management applications through an easily accessible application programming interface (API) and directly to users through a visual information dashboard

4> Applications include monitoring of schedule, quantities & budget, quality, safety, and environmental impact.

5> PSM representation based on property graph semantically linked to the Building Information Model (BIM) and all project management data. The property graph enables flexible, scalable storage of raw monitoring data in different formats, as well as storage of interpreted information. It enables smooth transition from construction to operation.

BIM2TWIN is a broad, multidisciplinary consortium with hand-picked partners who together provide an optimal combination of knowledge, expertise and experience in a variety of monitoring technologies, artificial intelligence, computer vision, information schema and graph databases, construction management, equipment automation and occupational safety. The DBT platform will be experimented on 3 demo sites (Spain, France, Finland).

GRAPES

Participants: Pierre Alliez, Rao Fu.

Title: GRAPES: Learning, Processing, and Optimising Shapes

Duration: From December 1, 2019 to May 31, 2024

Partners:

- NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- UNIVERSITA DEGLI STUDI DI ROMA TOR VERGATA (UNITOV), Italy
- RHEINISCH-WESTFAELISCHE TECHNISCHE HOCHSCHULE AACHEN (RWTH AACHEN), Germany
- ATHINA-EREVNTIKO KENTRO KAINOTOMIAS STIS TECHNOLOGIES TIS PLIROFORIAS, TON EPIKOINONION KAI TIS GNOSIS (ATHENA - RESEARCH AND INNOVATION CENTER), Greece
- UNIVERSITAT LINZ (JOHANNES KEPLER UNIVERSITAT LINZ UNIVERSITY OF LINZ JOHANNES KEPLER UNIVERSITY OF LINZ JKU), Austria
- SINTEF AS (SINTEF), Norway
- VILNIAUS UNIVERSITETAS (Vilniaus universitetas), Lithuania
- UNIVERSITA DELLA SVIZZERA ITALIANA (USI), Switzerland
- UNIVERSITAT DE BARCELONA (UB), Spain
- GEOMETRY FACTORY SARL, France
- UNIVERSITY OF STRATHCLYDE, United Kingdom

Inria contact: Laurent Busé

Coordinator: Ioannis Emiris (ATHENA, Greece)

Summary: GRAPES aims at considerably advancing the state of the art in Mathematics, Computer-Aided Design, 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. Recent advances in the above domains have solved numerous tasks concerning multimedia and 2D data. However, automation of 3D geometry processing and analysis lags severely behind, despite their importance in science, technology and everyday life, and the well-understood underlying mathematical principles. The CAD industry, although well established for more than 20 years, urgently requires advanced methods and tools for addressing new challenges.

The scientific goal of GRAPES is to bridge this gap based on a multidisciplinary consortium composed of leaders in their respective fields. Top-notch research is also instrumental in forming the new generation of European scientists and engineers. Their disciplines span the spectrum from Computational Mathematics, Numerical Analysis, and Algorithm Design, up to Geometric Modeling, Shape Optimization, and Deep Learning. This allows the 15 PhD candidates to follow either a theoretical or an applied track and to gain knowledge from both research and innovation through a nexus of intersectoral secondments and Network-wide workshops.

Horizontally, our results lead to open-source, prototype implementations, software integrated into commercial libraries as well as open benchmark datasets. These are indispensable for dissemination and training but also to promote innovation and technology transfer. Innovation relies on the active participation of SMEs, either as a beneficiary hosting an ESR or as associate partners hosting secondments. Concrete applications include simulation and fabrication, hydrodynamics and marine design, manufacturing and 3D printing, retrieval and mining, reconstruction and visualization, urban planning and autonomous driving.

9.3 National initiatives

9.3.1 3IA Côte d’Azur

Pierre Alliez holds a senior chair from 3IA Côte d’Azur (Interdisciplinary Institute for Artificial Intelligence). The topic of his chair is “3D modeling of large-scale environments for the smart territory”. In addition, he is the scientific head of the fourth research axis entitled “AI for smart and secure territories”. Two PhD thesis students have been funded by this program: Tong Zhao (Ecole des Ponts ParisTech, master MVA, now at Dassault Systemes Paris) and Nissim Maruani (Ecole Polytechnique, Master MVA, first year PhD student).

9.3.2 Inria challenge ROAD-AI - with CEREMA

The road network is one of the most important elements of public heritage. Road operators are responsible for maintaining, operating, upgrading, replacing and preserving this heritage, while ensuring careful management of budgetary and human resources. Today, the infrastructure is being severely tested by climate change (increased frequency of extreme events, particularly flooding and landslides). Users are also extremely attentive to issues of safety and comfort linked to the use of infrastructure, as well as to environmental issues relating to infrastructure construction and maintenance. Roads and structures are complex systems: numerous sub-systems, non-linear evolution of intrinsic characteristics, assets subject to extreme events, strong changes in use or constraints (e.g. climate change) due to their very long lifespan. This Inria challenge takes into account the following difficulties : volume and nature of data (heterogeneity and incompleteness, multi-sources...), multiple scales of analysis (from the sub-systems of an asset to the complete national network), prediction of very local behaviors on the scale of a global network, complexity of road objects due to their operation and weather conditions, 3D modeling with semantization, and expert knowledge modeling. This Inria challenge aims at providing the scientific building blocks for upstream data acquisition and downstream decision support. Cerema's use cases for this challenge separate into four parts:

- Build a dynamic “digital twin” of the road and its environment on the scale of a complete network;
- the behavioral "laws" of pavements and engineering structures using data from surface monitoring or structure visits, sensors and environmental data;
- Invent the concept of connected bridges and tunnels on a system scale;
- Define strategic investment and maintenance planning methods (predictive, prescriptive and autonomous).

9.3.3 PEPR NumPex

The Digital PEPR for the Exascale (NumPEx) aims to design and develop the software components and tools that will equip future exascale machines and to prepare the major application domains to fully exploit the capabilities of these machines. These major application domains include both scientific research and the industrial sector. This project therefore contributes to France's response to the next EuroHPC call for expressions of interest (AMI) (Exascale France Project), with a view to hosting one of the two European exascale machines planned in Europe by 2024. The French consortium has chosen GENCI as its "Hosting Entity" and the CEA TGCC as its "Hosting Center". This PEPR contributes to the constitution of a set of tools, software, applications but also training that will allow France to remain one of the leaders in the field by developing a national ecosystem for Exascale coordinated with the European strategy.

Our team is mainly involved in the first workpackage entitled “Model and discretize at Exascale”. The main motivation is that geometric representations and their discrete counterparts (such as meshes) are usually the starting point for simulation. These include adaptive, possibly multiresolution, robust to defects, and efficient parallel representations of large-scale models. The physics-based models need to include multiple phenomena, or process couplings at multiple scales in space and time. Space and time adaptivity are then mandatory. Time integration requires special care to become more parallel, more asynchronous, and more accurate for long-time simulations. The requirement of adaptivity in geometry and physics creates an imbalance that needs to be overcome. AI-driven, data-driven, reduced-order, and more generally surrogate models are now mandatory and take various forms. Data must be understood in a broad sense: from observations of the real physical system to synthetic data generated by the physical models. Surrogates enable orders of magnitude faster model evaluations through the extraction and compression of salient features in a very intensive training/learning (supervised, unsupervised or reinforced) phase. Research topics include: handling multiphysics and multiscale coupling, learning parametric dependencies, differential operators or underlying physical laws, mitigating intensive communications during the compression process, and exploiting the underlying computer and data architectures. Multi-fidelity models include hierarchies of models to provide a multi-fidelity problem-solving approach to address very intensive problems beyond the current computing capabilities. The challenge is to switch

between representations to efficiently compute and handle bias in so-called “many-query” problems, which include design, optimization, UQ and other high dimensional PDE problems.

For our team, this project is funding a postdoctoral fellow (to be recruited in 2024), in collaboration with Strasbourg University.

10 Dissemination

Participants: Pierre Alliez, Florent Lafarge, Roberto Dyke, Armand Zampieri.

10.1 Promoting scientific activities

10.1.1 Scientific events: organisation

Member of the organizing committees

- Pierre Alliez was an advisory board member for the EUROGRAPHICS 2023 conference.
- Pierre Alliez was a member of the scientific committee for the SophIA Summit conference 2023.
- Pierre Alliez organized the third Inria-DFKI summer school on AI, at Inria Sophia Antipolis (around 100 participants), in collaboration with Roberto Dyke and Armand Zampieri.

10.1.2 Scientific events: selection

Member of the conference program committees Pierre Alliez was a committee member of the Eurographics Symposium on Geometry Processing 2023.

Reviewer

- Pierre Alliez was a reviewer for ACM SIGGRAPH, ACM SIGGRAPH Asia and the Eurographics Symposium on Geometry Processing.
- Florent Lafarge was a reviewer for CVPR, ICCV, ACM SIGGRAPH and ACM SIGGRAPH ASIA.

10.1.3 Journal

Member of the editorial boards

- Pierre Alliez is the editor in chief of the Computer Graphics Forum (CGF), in tandem with Helwig Hauser, since January 2022 and for a four-year term. CGF is one of the leading journals in Computer Graphics and the official journal of the Eurographics association.
- Pierre Alliez 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 and for the Revue Française de Photogrammétrie et de Télédétection.

Reviewer - reviewing activities

- Florent Lafarge was a reviewer for IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI).
- Pierre Alliez declined most invitations for reviewing journal papers, since he took the role of editor in chief of the Computer Graphics Forum.

10.1.4 Leadership within the scientific community

Pierre Alliez is a member of the following steering committees: Eurographics (via his role of editor-in-chief of the Computer Graphics Forum journal), Eurographics Symposium on Geometry Processing and Eurographics Workshop on Graphics and Cultural Heritage (until September 2023). He was chair of the Ph.D. award committee of the Eurographics Association from 2020 to 2023.

10.1.5 Scientific expertise

- Pierre Alliez was a reviewer for the ERC and Eureka programs. He 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 (CRI - crédit impôt recherche and JEI - jeunes entreprises innovantes).

10.1.6 Research administration

- Pierre Alliez is president of the Inria Evaluation Committee since September 2023, in collaboration with Christine Morin and Luce Brotcorne. This role is a half-time position.
- Pierre Alliez is scientific head for the Inria-DFKI partnership since September 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.
- Florent Lafarge is a member of the technical committee of the UCA Academy of Excellence "Space, Environment, Risk and Resilience".

10.2 Teaching - Supervision - Juries

10.2.1 Teaching

- Master: Pierre Alliez (with Xavier Descombes, Marc Antonini and Laure Blanc-Féraud), advanced machine learning, 12h, M2, Université Côte d'Azur, France.
- Master: Pierre Alliez (with Gaétan Bahl and Armand Zampieri), deep learning, 32h, M2, Université Côte d'Azur, France.
- Master: Florent Lafarge, Applied AI, 8h, M2, Université Côte d'Azur, France.
- Master: Florent Lafarge (with Angelos Mantzaflaris), Numerical Interpolation, 55h, M1, Université Côte d'Azur, France.

10.2.2 Supervision

- PhD defended in March 2023: Tong Zhao, Progressive shape reconstruction from raw 3D point clouds, co-advised by Pierre Alliez and Laurent Busé (Aromath Inria project-team), funded by 3IA Côte d'Azur, in collaboration with Jean-Marc Thiery and Tamy Boubekeur (formerly at Telecom ParisTech, now at Adobe research), and David Cohen-Steiner (Datashape Inria project-team).
- PhD in progress: Rao Fu, Surface reconstruction from raw 3D point clouds, thesis funded by EU project GRAPES, in collaboration with Geometry Factory, since January 2021, advised by Pierre Alliez. Rao should defend in February 2024.
- PhD in progress: Jacopo Iollo, Sequential Bayesian Optimal Design, funded by Inria challenge ROAD-AI, in collaboration with CEREMA, since January 2022, co-advised by Florence Forbes (Inria Grenoble), Christophe Heinkele (CEREMA) and Pierre Alliez.

- PhD in progress: Marion Boyer, Geometric modeling of urban scenes with LOD2 formalism from satellite images, funded by CNES and Airbus, since October 2022, advised by Florent Lafarge.
- PhD in progress: Nissim Maruani, Learnable representations for 3D shapes, funded by 3IA, since November 2022, co-advised by Pierre Alliez and Mathieu Desbrun (Geomerix Inria project-team, Inria Saclay), and in collaboraton with Maks Ovsjanikov (Ecole Polytechnique).
- PhD in progress: Armand Zampieri, Compression and visibility of 3D point clouds, Cifre thesis with Samp AI, since December 2022, co-advised by Pierre Alliez and Guillaume Delarue (Samp AI).
- PhD in progress: Moussa Bendjilali, From 3D point clouds to cognitive 3D models, Cifre thesis with AlteIA, since September 2023, co-advised by Pierre Alliez and Nicola Luminari (AlteIA).

10.2.3 Juries

- Pierre Alliez was a reviewer for the HDR committee of Loïc Landrieu (IGN and Ecole des Ponts ParisTech).
- Pierre Alliez was a HDR committee member for Frédéric Payan (I3S, Université Côte d'Azur).
- Pierre Alliez chaired the PhD thesis committee of Morten Pedersen (University of Copenhagen and Epione project-team, Inria center at Université Côte d'Azur).
- Pierre Alliez was a member of the "Comité de Suivi Doctoral" for three Phd theses: Di Yang (Stars Inria project-team), Arnaud Gueze (Ecole Polytechnique) and Stefan Larsen (Acentauri Inria project-team).
- Florent Lafarge was a reviewer for the PhD of Mariem Mezghanni (LIX).
- Florent Lafarge was a member of the "Comité de Suivi Doctoral" for the PhD thesis of Amine Ouasfi (Mimetic Inria project-team) and Grégoire Grzeczkwicz (IGN).

11 Scientific production

11.1 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] C. Portaneri, M. Rouxel-Labbé, M. Hemmer, D. Cohen-Steiner and P. Alliez. ‘Alpha Wrapping with an Offset’. In: *ACM Transactions on Graphics*. Proceedings of the ACM SIGGRAPH conference 41.4 (1st June 2022), pp. 1–22. DOI: [10.1145/3528223.3530152](https://doi.org/10.1145/3528223.3530152). URL: <https://hal.inria.fr/hal-03688637>.
- [9] 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>.
- [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>.
- [11] T. Zhao, L. Busé, D. Cohen-Steiner, T. Boubekeur, J.-M. Thiery and P. Alliez. ‘Variational Shape Reconstruction via Quadric Error Metrics’. In: SIGGRAPH 2023 - The 50th International Conference & Exhibition On Computer Graphics & Interactive Techniques. Los Angeles, United States, 6th Aug. 2023. DOI: [10.1145/3588432.3591529](https://doi.org/10.1145/3588432.3591529). URL: <https://inria.hal.science/hal-04131765>.

11.2 Publications of the year

International journals

- [12] M. Trzeciak, K. Pluta, Y. Fathy, L. Alcalde, S. Chee, A. Bromley, I. Brilakis and P. Alliez. ‘ConSLAM: Construction Data Set for SLAM’. In: *Journal of Computing in Civil Engineering* 37.3 (May 2023). DOI: [10.1061/JCCEE5.CPENG-5212](https://doi.org/10.1061/JCCEE5.CPENG-5212). URL: <https://inria.hal.science/hal-04039392>.
- [13] V. Vidal, L. Dubouchet, G. Lavoué and P. Alliez. ‘Progressive Compression of Triangle Meshes’. In: *Image Processing On Line* 13 (2023), pp. 1–21. DOI: [10.5201/ipol.2023.418](https://doi.org/10.5201/ipol.2023.418). URL: <https://inria.hal.science/hal-03924042>.
- [14] T. Zhao, M. Yu, P. Alliez and F. Lafarge. ‘Sharp Feature Consolidation from Raw 3D Point Clouds via Displacement Learning’. In: *Computer Aided Geometric Design* 13 (June 2023), p. 102204. DOI: [10.1016/j.cagd.2023.102204](https://doi.org/10.1016/j.cagd.2023.102204). URL: <https://inria.hal.science/hal-03747150>.
- [15] D. Zint, N. Maruani, M. Rouxel-Labbé and P. Alliez. ‘Feature-Preserving Offset Mesh Generation from Topology-Adapted Octrees’. In: *Computer Graphics Forum*. Proceedings of Eurographics Symposium on Geometry Processing 2023 42.5 (2023), p. 12. DOI: [10.1111/cgf.14906](https://doi.org/10.1111/cgf.14906). URL: <https://inria.hal.science/hal-04135266>.

International peer-reviewed conferences

- [16] R. Fu, C. Wen, Q. Li, X. Xiao and P. Alliez. ‘BPNet: Bézier Primitive Segmentation on 3D Point Clouds’. In: IJCAI 2023 - 32nd International Joint Conference on Artificial Intelligence. MACAO, Macau SAR China: International Joint Conferences on Artificial Intelligence Organization, 2023, pp. 1–9. DOI: [10.24963/ijcai.2023/84](https://doi.org/10.24963/ijcai.2023/84). URL: <https://hal.science/ineris-04151504>.
- [17] N. Maruani, R. Klokov, M. Ovsjanikov, P. Alliez and M. Desbrun. ‘VoroMesh: Learning Watertight Surface Meshes with Voronoi Diagrams’. In: ICCV 2023 - International Conference on Computer Vision. Paris, France, 28th Aug. 2023. URL: <https://inria.hal.science/hal-04199011>.
- [18] T. Zhao, L. Busé, D. Cohen-Steiner, T. Boubekeur, J.-M. Thiery and P. Alliez. ‘Variational Shape Reconstruction via Quadric Error Metrics’. In: SIGGRAPH 2023 - The 50th International Conference & Exhibition On Computer Graphics & Interactive Techniques. Los Angeles, United States, 6th Aug. 2023. DOI: [10.1145/3588432.3591529](https://doi.org/10.1145/3588432.3591529). URL: <https://inria.hal.science/hal-04131765>.

Conferences without proceedings

- [19] P. Monasse, R. Djahel and B. Vallet. ‘Registration for Urban Modeling Based on Linear and Planar Features’. In: EUVIP 2023 - 11th European Workshop on Visual Information Processing. Gjøvik, Norway, 2023. URL: <https://enpc.hal.science/hal-04217550>.

Doctoral dissertations and habilitation theses

- [20] T. Zhao. ‘Progressive shape reconstruction from raw 3D point clouds’. Université Côte d’Azur, 27th Mar. 2023. URL: <https://inria.hal.science/tel-04131762>.

Reports & preprints

- [21] R. Sulzer, L. Landrieu, R. Marlet and B. Vallet. *A Survey and Benchmark of Automatic Surface Reconstruction from Point Clouds*. 2023. URL: <https://hal.science/hal-03968453>.

Other scientific publications

- [22] J. Lussange, M. Yu, Y. Tarabalka and F. Lafarge. *3D detection of roof sections from a single satellite image and application to LOD2-building reconstruction*. July 2023. URL: <https://inria.hal.science/hal-04343905>.

11.3 Other

Softwares

- [23] [SW] D. Anisimov, D. Bommès, K. Hormann and P. Alliez, *2D Generalized Barycentric Coordinates* version C++, 13th Jan. 2023. LIC: GNU General Public License v3.0 or later. HAL: [hal-03938534](https://hal.science/hal-03938534), URL: <https://inria.hal.science/hal-03938534>, VCS: https://github.com/CGAL/cgal/tree/master/Barycentric_coordinates_2, SWHID: [sw:1:dir:695cebe5d4ab078642bb746399b7372b62d9dc77](https://sw.hal.science/sw/1/dir/695cebe5d4ab078642bb746399b7372b62d9dc77);origin=<https://github.com/CGAL/cgal>;visit=sw:1:snp:08830c7520f20568b2ce9e0b4b05ce76c9012817;anchor=sw:1:rev:1d6830501d42d6f464779ea632f0944a424560f5;path=/Barycentric_coordinates_2/).