



Activity Report 2016

# Project-Team TITANE

Geometric Modeling of 3D Environments

RESEARCH CENTER  
Sophia Antipolis - Méditerranée

THEME  
Interaction and visualization



## Table of contents

<b>1. Members</b>	<b>1</b>
<b>2. Overall Objectives</b>	<b>2</b>
<b>3. Research Program</b>	<b>2</b>
3.1. Context	2
3.2. Analysis	3
3.3. Approximation	3
3.4. Reconstruction	4
<b>4. Application Domains</b>	<b>4</b>
<b>5. Highlights of the Year</b>	<b>5</b>
<b>6. New Software and Platforms</b>	<b>5</b>
6.1. CGAL Barycentric_coordinates_2	5
6.2. MeshMantics	5
6.3. Module CGAL : Point Set Processing	6
6.4. Module CGAL : Scale space surface reconstruction	6
6.5. Skeleton-Blockers	6
6.6. APP Deposits	6
<b>7. New Results</b>	<b>6</b>
7.1. Analysis	6
7.1.1. Object Classification via Planar Abstraction	6
7.1.2. Fidelity vs. Simplicity: a Global Approach to Line Drawing Vectorization	7
7.1.3. High-Resolution Semantic Labeling with Convolutional Neural Networks	7
7.1.4. Learning Iterative Processes with Recurrent Neural Networks to Correct Satellite Image Classification Maps	7
7.1.5. Convolutional Neural Networks for Large-Scale Remote-Sensing Image Classification	10
7.1.6. Fully Convolutional Neural Networks for Remote Sensing Image Classification	11
7.1.7. Large-scale Remote Sensing Image Segmentation and Classification	11
7.2. Reconstruction	12
7.2.1. Towards Large-scale City Reconstruction from Satellites	13
7.2.2. A Survey of Surface Reconstruction from Point Clouds	13
7.3. Approximation	14
7.3.1. A Line/Trimmed NURBS Surface Intersection Algorithm Using Matrix Representations	14
7.3.2. Optimal Voronoi Tessellations with Hessian-based Anisotropy	15
7.3.3. Symmetry and Orbit Detection via Lie-Algebra Voting	15
<b>8. Bilateral Contracts and Grants with Industry</b>	<b>15</b>
8.1. Bilateral Contracts with Industry	15
8.1.1. Geoimage	15
8.1.2. CSTB 1	15
8.1.3. CSTB 2	17
8.1.4. Luxcarta	17
8.2. Bilateral Grants with Industry	17
<b>9. Partnerships and Cooperations</b>	<b>18</b>
9.1. European Initiatives	18
9.2. International Initiatives	18
9.2.1.1. Declared Inria International Partners	18
9.2.1.2. Informal International Partners	18
9.3. International Research Visitors	18
<b>10. Dissemination</b>	<b>19</b>
10.1. Promoting Scientific Activities	19
10.1.1. Scientific Events Organisation	19

---

10.1.1.1. General Chair, Scientific Chair	19
10.1.1.2. Member of the Organizing Committees	19
10.1.2. Scientific Events Selection	19
10.1.2.1. Member of the Conference Program Committees	19
10.1.2.2. Reviewer	19
10.1.3. Journal	19
10.1.3.1. Member of the Editorial Boards	19
10.1.3.2. Reviewer - Reviewing Activities	19
10.1.4. Invited Talks	19
10.1.5. Scientific Expertise	20
10.1.6. Research Administration	20
10.2. Teaching - Supervision - Juries	20
10.2.1. Teaching	20
10.2.2. Supervision	20
10.2.3. Juries	20
10.3. Popularization	21
<b>11. Bibliography</b> .....	<b>21</b>

# Project-Team TITANE

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

## Keywords:

### Computer Science and Digital Science:

- 5.3. - Image processing and analysis
- 5.4.4. - 3D and spatio-temporal reconstruction
- 5.5.1. - Geometrical modeling
- 7.5. - Geometry, Topology

### Other Research Topics and Application Domains:

- 3.3. - Geosciences
- 5. - Industry of the future
- 8. - Smart Cities and Territories

## 1. Members

### Research Scientists

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

### Engineers

Nicolas Douillet [Inria, until Oct 2016]  
Come Le Breton [Inria, since Apr 2016]  
Sven Oesau [Inria, until Feb 2016, funded by CSTB]  
Cédric Portaneri [Inria, since Dec 2016]  
Maxim Torgonskiy [Inria, until Feb 2016]

### PhD Students

Jean-Philippe Bauchet [Luxcarta, since Sep 2016, funded by CIFRE]  
Liuyun Duan [Geoimage, until Sep 2016, funded by CIFRE, then thesis relay from Oct 2016]  
Oussama Ennaffi [IGN, from Oct 2016, funded by IGN]  
Hao Fang [Inria, since Feb, funded by CIFRE CSTB]  
Jean-Dominique Favreau [UNS]  
Manish Mandad [RWTH Aachen, thesis relay until Dec 2016]  
Emmanuel Maggiori [Inria]

### Post-Doctoral Fellow

Mohammad Rouhani [Inria, until Apr 2016]

### Visiting Scientists

Jingjing Shen [Fitzwilliam College, from May to Jun 2016]  
David Bommès [RWTH Aachen, Research professor, Sep 2016]  
Martin Heistermann [RWTH Aachen, PhD student, Sep 2016]

### Administrative Assistant

Florence Barbara [Inria]

### Others

Mathieu Desbrun [Inria, International Chair from Aug. to Nov. 2016 - Associate Researcher from Caltech]  
Chunlin Xiao [Inria, Master intern from UNS, from Mar to Sep 2016, funded by CNES]

## 2. Overall Objectives

### 2.1. General Presentation

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

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

## 3. Research Program

### 3.1. Context

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

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

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

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

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

## 3.2. Analysis

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

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

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

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

## 3.3. Approximation

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

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

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

### 3.4. Reconstruction

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

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

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

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

## 4. Application Domains

### 4.1. Applications

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



Our first ambition for technology transfer is to consolidate the components of our research experiments in the form of new software components for the CGAL (Computational Geometry Algorithms Library) 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 provide technical support.

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

## 5. Highlights of the Year

### 5.1. Highlights of the Year

#### 5.1.1. Awards

We obtained a Proof of Concept grant from the European Research Council, entitled TITANIUM (Software Components for Robust Geometry Processing). The TITANIUM project aims to develop a software demonstrator for geometry processing and 3D urban modeling, in order to facilitate the pre-commercialization of novel software components for the Computational Geometry Algorithms Library. The demonstrator will include novel approaches resulting from the ERC-funded IRON project.

BEST PAPER AWARD:

[7]

Z. SHI, P. ALLIEZ, M. DESBRUN, H. BAO, J. HUANG. *Symmetry and Orbit Detection via Lie-Algebra Voting*, in "Computer Graphics Forum", June 2016, 12 p. , <https://hal.inria.fr/hal-01344293>

## 6. New Software and Platforms

### 6.1. CGAL Barycentric\_coordinates\_2

This CGAL software component offers an efficient and robust implementation of two-dimensional closed-form generalized barycentric coordinates defined for simple two-dimensional polygons.

- Participants: Pierre Alliez
- Contact: Pierre Alliez
- URL: [http://doc.cgal.org/latest/Barycentric\\_coordinates\\_2/index.html#Chapter\\_2D\\_Generalized\\_Barycentric\\_Coordinates](http://doc.cgal.org/latest/Barycentric_coordinates_2/index.html#Chapter_2D_Generalized_Barycentric_Coordinates)

### 6.2. MeshMantics

This software component implements an approach that reconstructs 3D urban scenes in the form of levels of detail (LODs). Starting from raw data sets such as surface meshes generated by multi-view stereo systems, the algorithm proceeds in three main steps: classification, abstraction and reconstruction. From geometric attributes and a set of semantic rules combined with a Markov random field, we classify the scene into four meaningful classes. The abstraction step detects and regularizes planar structures on buildings, fits icons on trees, roofs and facades, and performs filtering and simplification for LOD generation. The abstracted data are then provided as input to the reconstruction step which generates watertight buildings through a min-cut formulation on a set of 3D arrangements.

- Participants: Florent Lafarge and Pierre Alliez
- Contact: Pierre Alliez
- URL: <https://bil.inria.fr>

### 6.3. Module CGAL : Point Set Processing

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

- Participants: Pierre Alliez and Clément Jamin
- Contact: Pierre Alliez
- URL: [http://doc.cgal.org/latest/Point\\_set\\_processing\\_3/index.html#Chapter\\_Point\\_Set\\_Processing](http://doc.cgal.org/latest/Point_set_processing_3/index.html#Chapter_Point_Set_Processing)

### 6.4. Module CGAL : Scale space surface reconstruction

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

- Participants: Pierre Alliez
- Contact: Pierre Alliez
- URL: [http://doc.cgal.org/latest/Scale\\_space\\_reconstruction\\_3/index.html#Chapter\\_Scale\\_space\\_reconstruction](http://doc.cgal.org/latest/Scale_space_reconstruction_3/index.html#Chapter_Scale_space_reconstruction)

### 6.5. Skeleton-Blockers

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

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

### 6.6. APP Deposits

WALLEMME is a software for classifying large-scale urban areas from dense textured 3D meshes in a supervised manner.

- Participants: Mohammad Rouhani, Florent Lafarge and Pierre Alliez.

DIMUVIC is a software for reconstructing in 3D a polyline-sketch using contextual knowledge contained in multiview stereo images.

- Participants: Jean-Dominique Favreau, Florent Lafarge and Adrien Bousseau.

ROOFEXTRACTOR is a software for reconstructing roofs from dense defect-laden meshes as compact piecewise-planar surface representations.

- Participants: Sven Oesau and Florent Lafarge.

## 7. New Results

### 7.1. Analysis

#### 7.1.1. Object Classification via Planar Abstraction

**Participants:** Sven Oesau, Florent Lafarge, Pierre Alliez.

*In collaboration with EADS ASTRIUM*

We contributed a supervised machine learning approach for classification of objects from sampled point data. The main idea consists in first abstracting the input object into planar parts at several scales, then discriminate between the different classes of objects solely through features derived from these planar shapes. Abstracting into planar shapes provides a means to both reduce the computational complexity and improve robustness to defects inherent to the acquisition process. Measuring statistical properties and relationships between planar shapes offers invariance to scale and orientation. A random forest is then used for solving the multiclass classification problem. We demonstrate the potential of our approach on a set of indoor objects from the Princeton shape benchmark and on objects acquired from indoor scenes and compare the performance of our method with other point-based shape descriptors [10]. This work was published in the proceedings of ISPRS.

### 7.1.2. *Fidelity vs. Simplicity: a Global Approach to Line Drawing Vectorization*

**Participants:** Jean-Dominique Favreau, Florent Lafarge.

*In collaboration with Adrien Bousseau (GraphDeco Inria team)*

Vector drawing is a popular representation in graphic design because of the precision, compactness and editability offered by parametric curves. However, prior work on line drawing vectorization focused solely on faithfully capturing input bitmaps, and largely overlooked the problem of producing a compact and editable curve network. As a result, existing algorithms tend to produce overly-complex drawings composed of many short curves and control points, especially in the presence of thick or sketchy lines that yield spurious curves at junctions. We propose the first vectorization algorithm that explicitly balances fidelity to the input bitmap with simplicity of the output, as measured by the number of curves and their degree. By casting this trade-off as a global optimization, our algorithm generates few yet accurate curves, and also disambiguates curve topology at junctions by favoring the simplest interpretations overall. We demonstrate the robustness of our algorithm on a variety of drawings, sketchy cartoons and rough design sketches (See Figure 1). This work was published at ACM SIGGRAPH 2016 [4].

### 7.1.3. *High-Resolution Semantic Labeling with Convolutional Neural Networks*

**Participants:** Emmanuel Maggiori, Yuliya Tarabalka, Pierre Alliez.


*In collaboration with Guillaume Charpiat (Inria TAO team)*

Convolutional neural networks (CNNs) were initially conceived for image categorization, i.e., the problem of assigning a semantic label to an entire input image. We have address the problem of dense semantic labeling, which consists in assigning a semantic label to *every* pixel in an image. Since this requires a high spatial accuracy to determine *where* labels are assigned, categorization CNNs, intended to be highly robust to local deformations, are not directly applicable. By adapting categorization networks, many semantic labeling CNNs have been recently proposed. Our first contribution is an in-depth analysis of these architectures. We establish the desired properties of an ideal semantic labeling CNN, and assess how those methods stand with regard to these properties. We observe that even though they provide competitive results, these CNNs often do not leverage properties of semantic labeling that could lead to more effective and efficient architectures. Out of these observations, we then derive a CNN framework specifically adapted to the semantic labeling problem [13]. In addition to learning features at different resolutions, it learns how to combine these features. By integrating local and global information in an efficient and flexible manner, it outperforms previous techniques. We evaluate the proposed framework and compare it with state-of-the-art architectures on public benchmarks of high-resolution aerial image labeling.

### 7.1.4. *Learning Iterative Processes with Recurrent Neural Networks to Correct Satellite Image Classification Maps*

**Participants:** Emmanuel Maggiori, Yuliya Tarabalka, Pierre Alliez.

*In collaboration with Guillaume Charpiat (Inria TAO team)*



*Figure 1. Line Drawing Vectorization. Rough sketches often contain overlapping strokes (left). Since existing algorithms analyze junctions locally, they cannot recover the proper topology of these seemingly similar line configurations. By adopting a global formulation that optimizes for both fidelity to the input sketch and simplicity of the output curve network, our algorithm recovers proper topology while significantly reducing the overall number of curves and control points (right). Design sketch after Sori Yanagi Butterfly stool.*




Figure 2. Our MLP network architecture (left) learns features at different resolutions and also learns how to combine those features. The technique was evaluated on the ISPRS 2D Semantic Segmentation Contest (right), providing competitive results.

While initially devised for image categorization, convolutional neural networks (CNNs) are being increasingly used for the pixelwise semantic labeling of images. However, the proper nature of the most common CNN architectures makes them good at recognizing but poor at localizing objects precisely. This problem is magnified in the context of aerial and satellite image labeling, where a spatially fine object outlining is of paramount importance.

Different iterative enhancement algorithms have been presented in the literature to progressively improve the coarse CNN outputs, seeking to sharpen object boundaries around real image edges. However, one must carefully design, choose and tune such algorithms. Instead, our goal is to directly learn the iterative process itself. For this, we formulate a generic iterative enhancement process inspired from partial differential equations, and observe that it can be expressed as a recurrent neural network (RNN). Consequently, we train such a network from manually labeled data for our enhancement task. In a series of experiments we show that our RNN effectively learns an iterative process that significantly improves the quality of satellite image classification maps [12].




Figure 3. A recurrent neural network (RNN) learns an algorithm to iteratively correct the output of a coarse classification map. As a result, the satellite image classification maps become finer and better aligned to the real objects.

### 7.1.5. Convolutional Neural Networks for Large-Scale Remote-Sensing Image Classification

**Participants:** Emmanuel Maggiori, Yuliya Tarabalka, Pierre Alliez.

*In collaboration with Guillaume Charpiat (Inria TAO team)*

We propose an end-to-end framework for the dense, pixelwise classification of satellite imagery with convolutional neural networks (CNNs). In our framework, CNNs are directly trained to produce classification maps out of the input images. We first devise a *fully convolutional* architecture and demonstrate its relevance to the dense classification problem. We then address the issue of imperfect training data through a two-step training

approach: CNNs are first initialized by using a large amount of possibly inaccurate reference data, then refined on a small amount of accurately labeled data. To complete our framework we design a multi-scale neuron module that alleviates the common trade-off between recognition and precise localization. A series of experiments show that our networks take into account a large amount of context to provide fine-grained classification maps. This work was published in IEEE Transactions on Geoscience and Remote Sensing (TGRS) [5].




Figure 4. We train in a two-step scheme: first we train a fully convolutional network on large amounts of imperfect training data, to capture the generalities of the problem, which leads to coarse classification maps (1). In a second stage we fine-tune the network for few iterations on a precise manually labeled dataset, outputting fine classification maps as a results (2). The overall system is efficient and scalable.

### 7.1.6. Fully Convolutional Neural Networks for Remote Sensing Image Classification

**Participants:** Emmanuel Maggiori, Yuliya Tarabalka, Pierre Alliez.

*In collaboration with Guillaume Charpiat (Inria TAO team)*

We propose a convolutional neural network (CNN) model for remote sensing image classification, i.e. the assignment of a class to every pixel in an image. Using CNNs provides us with a means of learning contextual features for large-scale image labeling. Our network consists of four stacked convolutional layers that downsample the image and extract relevant features. On top of these, a deconvolutional layer upsamples the data back to the initial resolution, producing a final dense image labeling. Contrary to previous frameworks, our architecture is a fully convolutional network (FCN), contains only convolution and deconvolution operations and no fully connected layers as in previous work. The fact of being fully convolutional removes the artifacts present in previous work by construction and is considerably more efficient. Experiments on aerial images show that our network produces more accurate classifications in lower computational time. This work was published in the proceedings of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS) [9].

### 7.1.7. Large-scale Remote Sensing Image Segmentation and Classification

**Participants:** Chunlin Xiao, Emmanuel Maggiori, Yuliya Tarabalka.

*In collaboration with Guillaume Charpiat (Inria TAO team)*

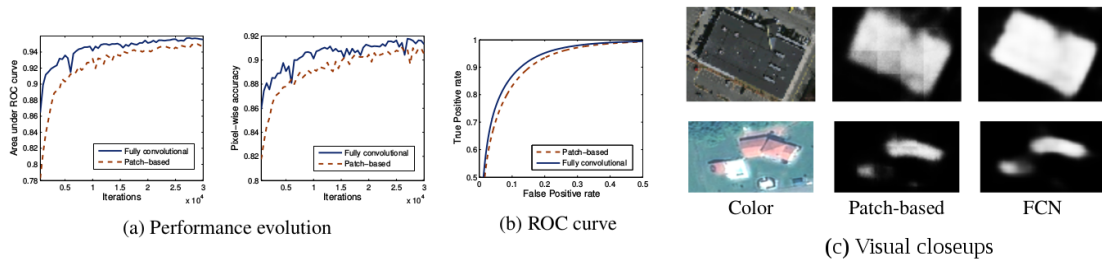


Figure 5. Our fully convolutional network (FCN) provides a better accuracy compared to a previous method (the “patch-based” network), as observed by the evolution of the accuracy through the training iterations (a) and the final precision/recall curve (b). We can also observe, visually, that our FCN network removes the artifacts at the border of patches (c). Besides the improved performance, the architecture drastically reduces the number of trainable parameters, being 10 times faster to run compared to the patch-based counterpart.

The representation of images with binary partition trees (BPTs) has proven to be very efficient for multiscale analysis, object detection and classification of high-resolution images. We propose a new framework for multi-class image segmentation using a binary partition tree. The region model is composed of three components : color component, probability component and shape component, some of which can be used or omitted depending on the information available and the application itself. The problem to extract a segmentation is formulated as the minimization of an energy function which can be solved with dynamical programming efficiently. However, BPT represents a hierarchy of the image regions at different scales. For large-scale images such representation can be demanding in terms of both memory and computation resources. We propose a tile-based scheme to extend the framework for processing arbitrarily large images. Experiments (see Fig. 6) prove that the algorithm can segment large images efficiently while ensuring quite similar results with respect to processing the whole image at once. This work has not been published yet.




Figure 6. Results of (left) unsupervised segmentation and (right) supervised segmentation of the image into building (red) and non-building (green) regions, using  $4 \times 4$  tiling scheme.



## 7.2. Reconstruction

### 7.2.1. Towards Large-scale City Reconstruction from Satellites

**Participants:** Liuyun Duan, Florent Lafarge.

*In collaboration with Geoimage.*

Automatic city modeling from satellite imagery is one of the biggest challenges in urban reconstruction. Existing methods produce at best rough and dense Digital Surface Models. Inspired by recent works on semantic 3D reconstruction and region-based stereovision, we propose a method for producing compact, semantic-aware and geometrically accurate 3D city models from stereo pair of satellite images [8]. Our approach relies on two key ingredients. First, geometry and semantics are retrieved simultaneously bringing robustness to occlusions and to low image quality. Second, we operate at the scale of geometric atomic region which allows the shape of urban objects to be well preserved, and a gain in scalability and efficiency. We demonstrate the potential of our algorithm by reconstructing different cities around the world in a few minutes (See Figure 7). This work has been published in the proceedings of the European Conference on Computer Vision (ECCV).

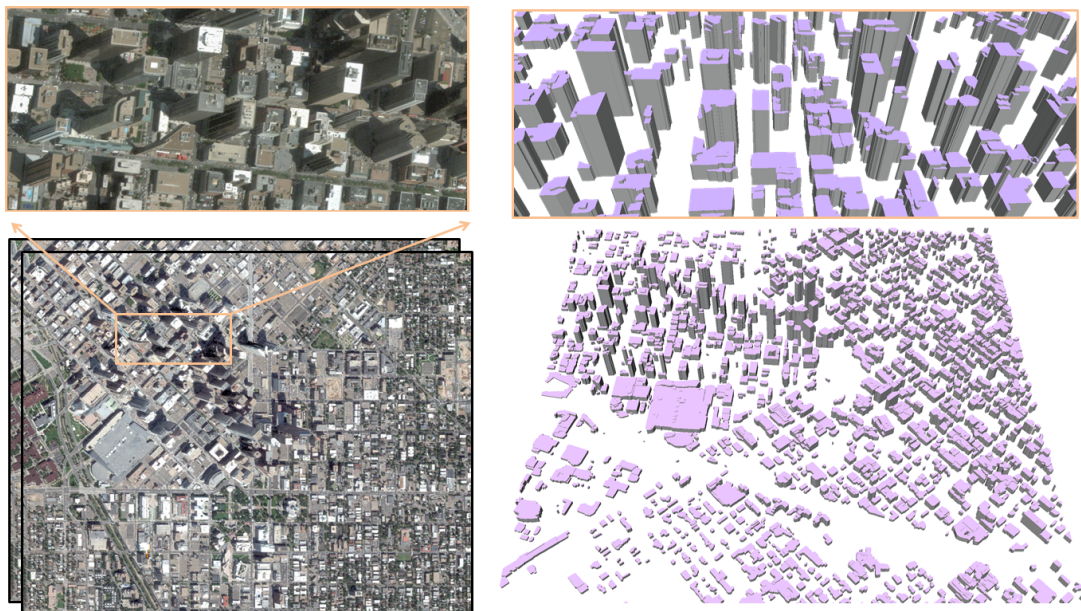


Figure 7. City reconstruction from satellites. Starting from a stereo pair of satellite images (left), our algorithm produces a compact and semantic-aware 3D model (right) in a few minutes.

### 7.2.2. A Survey of Surface Reconstruction from Point Clouds

**Participant:** Pierre Alliez.

*In collaboration with Matthew Berger, Andrea Tagliasacchi, Lee Seversky, Gael Guennebaud (Inria MANAO), Joshua Levine, Andrei Sharf and Claudio Silva.*

The area of surface reconstruction has seen substantial progress in the past two decades. The traditional problem addressed by surface reconstruction is to recover the digital representation of a physical shape that has been scanned, where the scanned data contains a wide variety of defects. While much of the earlier work has been focused on reconstructing a piece-wise smooth representation of the original shape, recent work has taken on more specialized priors to address significantly challenging data imperfections, where the reconstruction can take on different representations – not necessarily the explicit geometry. We survey the field of surface reconstruction, and provide a categorization with respect to priors, data imperfections, and reconstruction output. By considering a holistic view of surface reconstruction, we show a detailed characterization of the field, highlight similarities between diverse reconstruction techniques, and provide directions for future work in surface reconstruction. This survey was published in Computer Graphics Forum [2].


## 7.3. Approximation

### 7.3.1. A Line/Trimmed NURBS Surface Intersection Algorithm Using Matrix Representations

**Participant:** Pierre Alliez.

*In collaboration with Laurent Busé from Inria AROMATH, and Jingjing Shen and Neil Dodgson from Cambridge University (UK).*

We contribute a reliable line/surface intersection method for trimmed NURBS surfaces, based on a novel matrix-based implicit representation and numerical methods in linear algebra such as singular value decomposition and the computation of generalized eigenvalues and eigenvectors. A careful treatment of degenerate cases makes our approach robust to intersection points with multiple pre-images. We then apply our intersection algorithm to seamlessly mesh NURBS surfaces through Delaunay refinement (see Figure 8). We demonstrate the added value of our approach in terms of accuracy and treatment of degenerate cases, by providing comparisons with other intersection approaches as well as a variety of meshing experiments. This work was published in Computer Aided Geometric Design [6].



*Figure 8. Seamless meshing. Top: meshing with two mesh sizing values. The initial point set is generated by sampling along the open boundary. Bottom: meshing across smooth edges (red). The initial point set is generated by sampling along the boundary edge (blue). Meshes generated with two sizing values (side and front view).*

### 7.3.2. *Optimal Voronoi Tessellations with Hessian-based Anisotropy*

**Participants:** Pierre Alliez, Mathieu Desbrun.

*In collaboration with Max Budninskiy and Beibei Liu from Caltech, Fernando de Goes from Pixar and Yiying Tong from Michigan State University.*

We contribute a variational method to generate cell complexes with local anisotropy conforming to the Hessian of any given convex function and for any given local mesh density. Our formulation builds upon approximation theory to offer an anisotropic extension of Centroidal Voronoi Tessellations which can be seen as a dual form of Optimal Delaunay Triangulation. We thus refer to the resulting anisotropic polytopal meshes as Optimal Voronoi Tessellations. Our approach sharply contrasts with previous anisotropic versions of Voronoi diagrams as it employs first-type Bregman diagrams, a generalization of power diagrams where sites are augmented with not only a scalar-valued weight but also a vector-valued shift. As such, our OVT meshes contain only convex cells with straight edges (Figure 9), and admit an embedded dual triangulation that is combinatorially-regular. We show the effectiveness of our technique using off-the-shelf computational geometry libraries. This work was published at ACM SIGGRAPH Asia [3].

### 7.3.3. *Symmetry and Orbit Detection via Lie-Algebra Voting*

**Participants:** Pierre Alliez, Mathieu Desbrun.

*In collaboration with Zeyun Shi, Hujun Bao and Jin Huang from Zhejiang University.*

We formulate an automatic approach to the detection of partial, local, and global symmetries and orbits in arbitrary 3D datasets. We improve upon existing voting-based symmetry detection techniques by leveraging the Lie group structure of geometric transformations. In particular, we introduce a logarithmic mapping that ensures that orbits are mapped to linear subspaces, hence unifying and extending many existing mappings in a single Lie-algebra voting formulation (Figure 10). Compared to previous work, our resulting method offers significantly improved robustness as it guarantees that our symmetry detection of an input model is frame, scale, and reflection invariant. As a consequence, we demonstrate that our approach efficiently and reliably discovers symmetries and orbits of geometric datasets without requiring heavy parameter tuning. This work was published in the proceedings of the EUROGRAPHICS Symposium on Geometry Processing [7].

## 8. Bilateral Contracts and Grants with Industry

### 8.1. Bilateral Contracts with Industry

#### 8.1.1. *Geoimage*

**Participants:** Liuyun Duan, Florent Lafarge.

The aim of this collaboration is to devise a new type of 2.5D representation from satellite multi-view stereo images which is more accurate, compact and meaningful than the conventional digital elevation models (DEMs). A key direction consists in incorporating semantic information directly during the image matching process. This semantic is related to the type of components of the scene, such as vegetation, roofs, building edges, roads and land.


- Starting date: November 2013 - Duration: 4 years

#### 8.1.2. *CSTB 1*

**Participants:** Sven Oesau, Florent Lafarge.

The goal of this collaboration was to consolidate and integrate research codes developed in Titane on urban semantization and reconstruction into the CSTB reconstruction platform.

- Starting date: September 2015 - Duration: 6 months



*Figure 9. Optimal Voronoi Tessellations with Hessian-based Anisotropy. We show that the construction of an optimal piecewise-linear approximation of a function over a cell complex (left) extends the isotropic notion of Centroidal Voronoi Tessellations (CVT, top) to an anisotropic variant (middle and bottom) we call Optimal Voronoi Tessellation (OVT), to stress its duality to Optimal Delaunay Triangulation (ODT). Cell anisotropy (indicated by tightest ellipses) and density are independently controlled, and the dual triangulation based on cell barycenters is embedded and combinatorially-regular.*

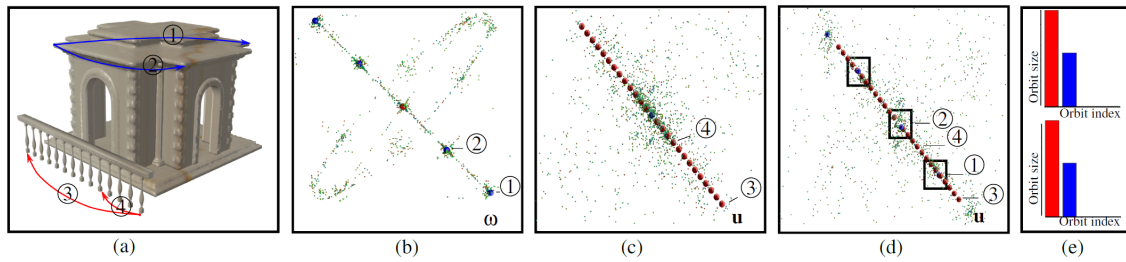


Figure 10. Our Lie algebra voting approach to symmetry and orbit detection maps  $SE(3)$  transformations into points in a logarithmic space composed of a rotation part and a translation part. The rotational orbit of the church and the translational orbit of the side railing (a) are mapped into collinear blue and red spheres respectively (a few transformations within these two orbits are marked with circled numbers to enhance comprehension). When the scene is centered, the two lines are orthogonal to each other and easy to distinguish (b). However, after a rigid translation of the scene, the rotational orbit now has translation-values near the translation orbit points, making it impossible to automatically distinguish these two orbits using a Euclidean distance (d), while our adjoint invariant distance for orbit shows no discernible difference in results as evidenced by a binning of detected orbit sizes for both situations (e).

### 8.1.3. CSTB 2

**Participants:** Hao Fang, Florent Lafarge.

The goal of this recent collaboration is to develop methods for analyzing and exploring scale-spaces into urban 3D data.

- Starting date: March 2016 - Duration: 3 years

### 8.1.4. Luxcarta

**Participants:** Jean-Philippe Bauchet, Florent Lafarge.

The goal of this recent collaboration is to design automatic approaches for producing LOD2 city models from the last generation of satellites.

- Starting date: October 2016 - Duration: 3 years

## 8.2. Bilateral Grants with Industry

### 8.2.1. CNES Toulouse

**Participants:** Emmanuel Maggiori, Yuliya Tarabalka [PI].

Hierarchical approaches for object-oriented classification of multi-source images. Contract 150490/00.

- Starting date: November 2015

- Duration: 2 years

## 9. Partnerships and Cooperations

### 9.1. European Initiatives

#### 9.1.1. FP7 & H2020 Projects

##### 9.1.1.1. TITANIUM - Software Components for Robust Geometry Processing

Type: IDEAS

Instrument: ERC Proof of concept

Duration: 18 months

Coordinator: Pierre Alliez

Inria contact: Pierre Alliez

Abstract: The TITANIUM project aims to develop a software demonstrator for geometry processing and 3D urban modeling, in order to facilitate the pre-commercialization of novel software components for the Computational Geometry Algorithms Library. The demonstrator will include novel approaches resulting from the ERC-funded IRON project (Robust Geometry Processing, StG-2010-257474), which are illustrated by publications presented at premier conferences in our field and a patent submitted in 2015. The expected outcomes of TITANIUM will be versatile methods for 3D reconstruction and simplification of data gathered from geometric measurements, as well as related methods specifically tailored to urban modeling. These methods represent a significant step forward by offering unrivaled levels of robustness, and automated generation of levels of detail that are semantically meaningful. The acronym TITANIUM, a robust and lightweight material, conveys our wish to streamline the geometric modeling pipeline through robust algorithms and lightweight representations. This Proof of Concept project will also implement the steps required for pre-commercialization. In view of this goal, we have included an industrial partner, GeometryFactory, a spinoff from Inria. We have already established preliminary contacts in the fields of metrology and geographic information systems. These contacts will provide real-world industrial case studies.

### 9.2. International Initiatives

#### 9.2.1. Inria International Partners

##### 9.2.1.1. Declared Inria International Partners

We have a long standing collaboration with Prof. Mathieu Desbrun from Caltech.

##### 9.2.1.2. Informal International Partners

We collaborate with researchers from RWTH Aachen.

### 9.3. International Research Visitors

#### 9.3.1. Visits of International Scientists

Prof. Mathieu Desbrun visited us for 3 months between August and November, within the framework of the Inria international chair.

##### 9.3.1.1. Internships

Chunlin Xiao (University of Nice Sophia-Antipolis and University of L'Aquila): large-scale remote sensing image segmentation and classification.

## 10. Dissemination

### 10.1. Promoting Scientific Activities

#### 10.1.1. Scientific Events Organisation

##### 10.1.1.1. General Chair, Scientific Chair

Florent Lafarge was co-chair of the ISPRS working group on point cloud processing.

##### 10.1.1.2. Member of the Organizing Committees

Pierre Alliez is a member of the Steering Board of the EUROGRAPHICS Workshop on Graphics and Cultural Heritage.

Yuliya Tarabalka chaired a session “Classification of Hyperspectral Image” at IEEE IGARSS 2016.

#### 10.1.2. Scientific Events Selection

##### 10.1.2.1. Member of the Conference Program Committees

Pierre Alliez was a program committee member of ACM SIGGRAPH, EUROGRAPHICS annual conference, EUROGRAPHICS Symposium on Geometry Processing, Geometric Modeling and Processing, Shape Modeling International and EUROGRAPHICS Workshop on Graphics and Cultural Heritage.

Florent Lafarge was a program committee member for the ISPRS congress in 2016.

Yuliya Tarabalka was a program committee member for ACIVS 2016, Lecce, Italy.

##### 10.1.2.2. Reviewer

Pierre Alliez was a reviewer for the conferences listed above, and for around 25 additional papers in several journals ranging from geometric modeling to computer vision. Florent Lafarge was a reviewer for JPRS.

Yuliya Tarabalka was a reviewer for the conferences IEEE IGARSS (student competition selection committee) and ACIVS 2016.

#### 10.1.3. Journal

##### 10.1.3.1. Member of the Editorial Boards

Pierre Alliez is an associate editor of ACM Transactions on Graphics, Elsevier Graphical Models, Computer Aided Geometric Design and Visual Informatics (since 2016). He is also a member of the editorial board of the CGAL open source project.

Florent Lafarge is an associate editor of The Visual Computer since 2015.

Yuliya Tarabalka was a co-editor of the special issue “Hyperspectral Imaging and Image Processing” of the Springer journal Sensing and Imaging.

##### 10.1.3.2. Reviewer - Reviewing Activities

Florent Lafarge was a reviewer for CVPR, ECCV, and SIGGRAPH.

Yuliya Tarabalka was a reviewer for the journals IEEE TGRS and IEEE JSTARS.

#### 10.1.4. Invited Talks

Pierre Alliez gave an invited talk at the CVPR Workshop on Large Scale 3D Data: Acquisition, Modelling and Analysis. The talk was entitled: *Shape Reconstruction and Approximation: Robustness and Guarantees*. See [http://www.multimediautis.org/3DWorkshop\\_CVPR2016/Program.html](http://www.multimediautis.org/3DWorkshop_CVPR2016/Program.html).

He also gave a keynote at STAG 2016: Smart Tools and Apps in computer Graphics, organized October 3-4, in Genova, Italy. The talk was entitled: *Low Distortion Inter-surface Mapping via Optimal Mass Transport*.

He also gave a course at ACM SIGGRAPH on the CGAL library, in collaboration with Andreas Fabri from GeometryFactory.

### 10.1.5. Scientific Expertise

Pierre Alliez was an evaluator and reviewer for the H2020, EuroSTARS and ERC programmes from the European commission. He participated to the Scientific Commissions of Belgium FNRS (Fund for Scientific Research). He was a member of the panel GEV 01 (mathematics and computer science) of ANVUR (Italian research evaluation agency). We evaluated the period 2011-2014 through bibliometry and peer-review. He was also a reviewer for PRIN projects (fundamental research projects) for the MIUR (the Italian Ministry for Education, University and Research).

Yuliya Tarabalka was an expert evaluator for an ANR project submission in May 2016. She is since November 2016 a member of the expert panel SBWT (signal processing) of the FWO (Belgian research funding foundation).

### 10.1.6. Research Administration

Pierre Alliez: member of the BCP (bureau du CP) since 2015, comité MASTIC (popularization), and comité espace immersif.

Florent Lafarge was an elected member of the Comité de Centre from 2013 to 2016, and a member of the CSD (Commission de Suivi Doctorale) from 2011 to 2016.

Yuliya Tarabalka is an elected member of the Comité de Centre since 2016.

## 10.2. Teaching - Supervision - Juries

### 10.2.1. Teaching

Master: Pierre Alliez and Florent Lafarge, Ingénierie 3D, 21h, M2, university Nice Sophia Antipolis, France.

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

Master: Pierre Alliez, Mathématiques pour la géométrie, 24h, M2, EFREI, France.

Master: Florent Lafarge, Traitement d'images numériques, 9h, M2, university Nice Sophia Antipolis, France.

Master: Yuliya Tarabalka, Mathematical methods, 25h, MSc in data sciences and business analytics ESSEC-CS, CentraleSupélec, France.

L2: Yuliya Tarabalka, Advanced algorithms, 28.5h, L2 Networks and Telecoms, IUT Nice Cote d'Azur, Sophia-Antipolis, France.

### 10.2.2. Supervision

PhD defended November 29th: Manish Mandad, Shape approximation with guarantees, since October 2012, Pierre Alliez.

PhD in progress: Dorothy Duan, Semantized Elevation Maps, since October 2013, Florent Lafarge.

PhD in progress: Jean-Dominique Favreau, Sketch-based modeling in multi-view context, since October 2014, Florent Lafarge and Adrien Bousseau.

PhD in progress: Hao Fang, Scale-space understanding in urban scenes, since March 2016, Florent Lafarge.

PhD in progress: Jean-Philippe Bauchet, City modelling from high resolution satellite images, since October 2016, Florent Lafarge.

PhD in progress: Emmanuel Maggiori, Representation and Analysis of Multisensor Remote Sensing Images with Partition Trees, University of Nice-Sophia Antipolis, since January 2015, Yuliya Tarabalka and Pierre Alliez.

### 10.2.3. Juries

Pierre Alliez:

- Thesis reviewer: Ludovic Blache (CReSTIC-SIC Reims).
- Thesis reviewer: Ruqi Huang (Inria Saclay).
- Thesis examiner: Mael Rouxel-Labbé (Inria Sophia Antipolis)



Florent Lafarge:

- Thesis examiner: Remi Cura (University Paris Est).

Yuliya Tarabalka:

- Monitoring committee for the thesis of Amine Bohi, Southern University of Toulon-Var in October 2016.

### 10.3. Popularization

Pierre Alliez participated to the "Fête de la science" in Juan-les-Pins, October 22 and 23. He also gave two talks in high schools: April 28 in Celony, and February 25th in Aix Valabre, and organized a workshop for MathC2+ in June.

## 11. Bibliography

### Publications of the year

#### Doctoral Dissertations and Habilitation Theses

- [1] M. MANDAD. *Robust Shape Approximation and Mapping between Surfaces*, Université de Nice-Sophia Antipolis, November 2016, <https://hal.inria.fr/tel-01396443>

#### Articles in International Peer-Reviewed Journals

- [2] M. BERGER, A. TAGLIASACCHI, L. SEVERSKY, P. ALLIEZ, G. GUENNEBAUD, J. LEVINE, A. SHARF, C. SILVA. *A Survey of Surface Reconstruction from Point Clouds*, in "Computer Graphics Forum", 2016, 27 p. [DOI : 10.1111/CGF.12802], <https://hal.inria.fr/hal-01348404>
- [3] M. BUDNINSKIY, B. LIU, F. DE GOES, Y. TONG, P. ALLIEZ, M. DESBRUN. *Optimal Voronoi Tessellations with Hessian-based Anisotropy*, in "ACM Transactions on Graphics", December 2016, 12 p. , <https://hal.inria.fr/hal-01376243>
- [4] J.-D. FAVREAU, F. LAFARGE, A. BOUSSEAU. *Fidelity vs. Simplicity: a Global Approach to Line Drawing Vectorization*, in "ACM Transactions on Graphics", 2016 [DOI : 10.1145/2897824.2925946], <https://hal.inria.fr/hal-01309271>
- [5] E. MAGGIORI, Y. TARABALKA, G. CHARPIAT, P. ALLIEZ. *Convolutional Neural Networks for Large-Scale Remote Sensing Image Classification*, in "IEEE Transactions on Geoscience and Remote Sensing", September 2016, <https://hal.inria.fr/hal-01369906>
- [6] J. SHEN, L. BUSÉ, P. ALLIEZ, N. DODGSON. *A Line/Trimmed NURBS Surface Intersection Algorithm Using Matrix Representations*, in "Computer Aided Geometric Design", 2016, vol. 48, pp. 1-16 [DOI : 10.1016/J.CAGD.2016.07.002], <https://hal.inria.fr/hal-01268109>
- [7] *Best Paper*  
Z. SHI, P. ALLIEZ, M. DESBRUN, H. BAO, J. HUANG. *Symmetry and Orbit Detection via Lie-Algebra Voting*, in "Computer Graphics Forum", June 2016, 12 p. , <https://hal.inria.fr/hal-01344293>.

### International Conferences with Proceedings

- [8] L. DUAN, F. LAFARGE. *Towards large-scale city reconstruction from satellites*, in "European Conference on Computer Vision (ECCV)", Amsterdam, Netherlands, October 2016, <https://hal.inria.fr/hal-01352466>
- [9] E. MAGGIORI, Y. TARABALKA, G. CHARPIAT, P. ALLIEZ. *Fully Convolutional Neural Networks For Remote Sensing Image Classification*, in "IEEE International Geoscience and Remote Sensing Symposium", Beijing, China, Proc. IEEE International Geoscience and Remote Sensing Symposium (IGARSS), IEEE, July 2016, pp. 5071-5074, <https://hal.inria.fr/hal-01350706>
- [10] S. OESAU, F. LAFARGE, P. ALLIEZ. *Object classification via planar abstraction*, in "ISPRS congress", Prague, Czech Republic, July 2016, <https://hal.inria.fr/hal-01318637>

### Other Publications

- [11] S.-G. JEONG, Y. TARABALKA, N. NISSE, J. ZERUBIA. *Progressive Tree-like Curvilinear Structure Reconstruction with Structured Ranking Learning and Graph Algorithm*, December 2016, working paper or preprint, <https://hal.inria.fr/hal-01414864>
- [12] E. MAGGIORI, G. CHARPIAT, Y. TARABALKA, P. ALLIEZ. *Learning Iterative Processes with Recurrent Neural Networks to Correct Satellite Image Classification Maps*, October 2016, working paper or preprint, <https://hal.inria.fr/hal-01388551>
- [13] E. MAGGIORI, Y. TARABALKA, G. CHARPIAT, P. ALLIEZ. *High-Resolution Semantic Labeling with Convolutional Neural Networks*, November 2016, working paper or preprint, <https://hal.inria.fr/hal-01393279>