ℓ_0-plane pursuit

Piecewise-planar approximation of large 3D data as graph-structured optimisation

Stéphane Guinard

Univ. Paris-Est, LASTIG ACTE, IGN, ENSG, F-94160 Saint-Mande, France
DGA

Table of Contents

Context

Context

LiDAR point clouds are more and more dense (i.e. millions of points):

- Capture more details
- Harder to process.

Thus, it is interesting to generalize and simplify 3D data.

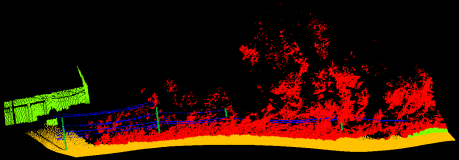

(a) Oakland dataset
1.4×10^5 points
[Munoz et al., 2009]

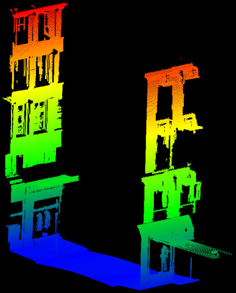

(b) Urban dataset
2.2×10^5 points
[Paparoditis et al., 2012]

Context

- LiDAR point cloud in urban areas, from TLS or MLS
- Urban areas are composed of man-made objects with simple shapes
- Each object can be represented as a set of piecewise-planar polygones.

$\Rightarrow$ **We want to approximate a LiDAR point cloud with a set of planar regions.**

State of the Art - Plane extraction

RANSAC approaches

- RANSAC algorithm [Schnabel et al., 2007]
- Manhattan-like structures [Holzmann et al., 2017]
- Weighted RANSAC [Xu et al., 2015].

Geometric features

- Features based on local point cloud organisation [Chauve et al., 2010, Demantke et al., 2011]
- Local curvature [Ma et al., 2013].

State of the Art - Segmentation

Region growing

- Mesh approximation [Cohen-Steiner et al., 2004]
- Curvature-based [Whelan et al., 2015]
- Alternance of region growing and regularisation [Oesau et al., 2016].

Mean-Shift

- Tree detection [Dai et al., 2018, Ferraz et al., 2012].

Graph-cuts

- ℓ_0-cut pursuit [Landrieu and Obozinski, 2017]
- Normalized Cut [Dutta et al., 2018].

Contribution

Formulation of the piecewise-planar approximation problem as a:

- global problem: one optimisation for the whole cloud
- non-convex optimisation problem.

Although being **global**, the resulting segmentation is **adaptive** to **local** geometry.

Objective

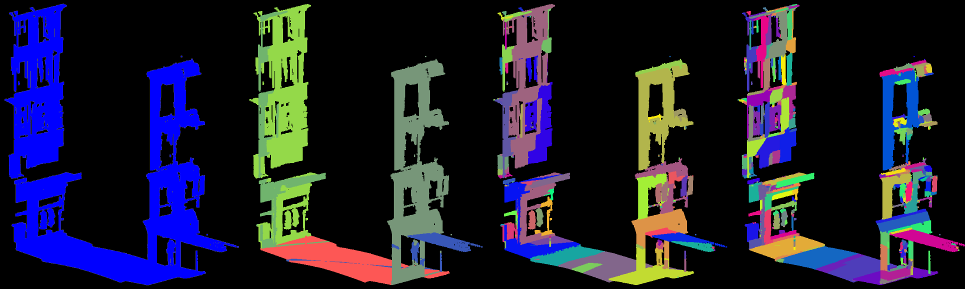

a: Point cloud **b:** Initialization **c:** First iteration **d:** Final result

Illustration of our method. The initialization is made with RANSAC (b). Our method is iterative (c) and produces the output shown in (d).

Method

Problem

Let $G = (V, E, w)$ be a graph with,

- V: the set of 3D points / mesh
- $E \subset V \times V$: the adjacency between vertices
- $w \in \mathbb{R}_+^E$: weight of each edge, encoding their *closeness*.

Let Π be a partition:

- $\Pi = (U_1, \ldots, U_k)$
- $\begin{cases} \cup_{i=1\ldots k} U_i = V \\ (U_i \cap U_j)_{i \neq j} = \varnothing \end{cases}$

Problem

We want to obtain a piecewise-planar partition Π of V, such that:

- we minimize the euclidean distance between each point and the plane supporting its region $\rightarrow$ **data term**.
- we preserve simple and geometrically adaptive regions $\rightarrow$ **regularisation term**.

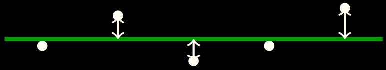

Data term

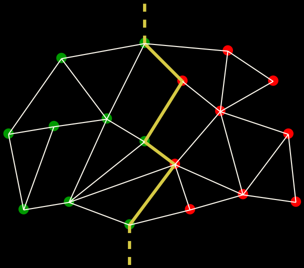

Regularisation term

Adjacency graph

We use the weighted adjacency graph introduced by Guinard and Vallet [2018] as a prefiltering step to filter planar regions and remove linear or scattered ones.

We weight an edge (u, v) of the graph according to its length $d(u, v)$, compared to the mean length of an edge in the graph d_0:

$$w_{u,v} = \frac{1}{\alpha + \frac{d(u,v)}{d_0}} \ , \quad (1)$$

with α taken as 2 here.

Piecewise-planar approximation

The piecewise-planar approximation is computed by solving a graph-structured optimization problem.

We use a modified version of the ℓ_0-cut pursuit algorithm of Landrieu and Obozinski [2017] which we dub ℓ_0-*plane pursuit*.

Energy formulation

For $\Pi \in \mathcal{P}^V$, we propose the following energy:

$$E(\Pi) = \underbrace{\sum_{v \in V} d\,(\,v, \Pi_v\,)^2}_{Data\ \ term} + \mu \underbrace{\sum_{(u,v) \in E} w_{u,v}\,[\,\Pi_u \neq \Pi_v\,]}_{Regularizer}, \quad (2)$$

where $[\pi \neq \pi'] : \mathcal{P}^2 \rightarrow \begin{cases} 0 \ \ if \ \ \pi = \pi' \\ 1 \ \ else \end{cases}$

Energy minimization

$$E(\Pi) = \sum_{v \in V} d(v, \Pi_v)^2 + \mu \sum_{(u,v) \in E} w_{u,v}[\Pi_u \neq \Pi_v].$$

We define the set of approximating planes as the following optimization problem:

$$\Pi^* = \arg\min_{\Pi \in \mathcal{P}^V} E(\Pi). \tag{3}$$

E is non-convex, non-continuous and non-differentiable.

E is very similar to the energy in the *generalized minimal parition problem* [Landrieu and Obozinski, 2017].

ℓ_0-cut pursuit

Introduced by Landrieu and Obozinski [2017].

The algorithm is divided in 2 main steps:

- Split step: cut a region in two new regions
- Merge step: merge two regions into one.

Non-convex problem:

$\Rightarrow$ Local minima

$\Rightarrow$ algorithm can be stuck in meaningless solutions

$\Rightarrow$ A (heuristic) initialization step is fundamental to guide the process.

Initialization

To give a first approximate solution, we use the RANSAC algorithm [Schnabel et al., 2007].

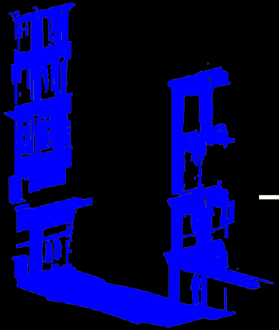

(a) Point cloud

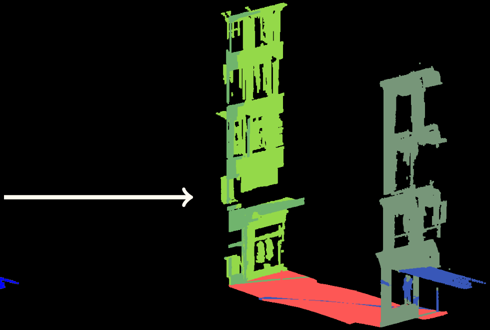

(b) RANSAC initialization

Split step: 2D illustration

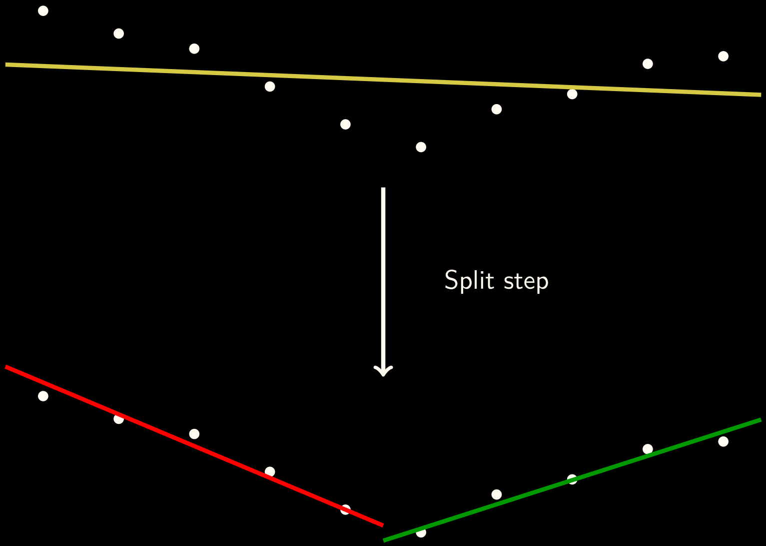

Split step

In ℓ_0-cut pursuit algorithm, components are split according to optimal binary partition criterion [Landrieu and Obozinski, 2017, 3.1.1].

In our case, this amounts to finding the binary partition (B, B^c) of a component U defined as follow:

$$\underset{B \subset U, (\pi, \pi') \in \mathcal{P}^2}{\arg\min} \sum_{v \in B} d(v, \pi) + \sum_{v \in U \setminus B} d(v, \pi') + \mu \sum_{(i,j) \in E^B} w_{i,j} \ , \qquad (4)$$

with $E^B = E \cap (B \times B^c \cup B^c \times B)$ the set of edges linking B and $B^c = U \setminus B$.

Split step

$$\underset{B\subset U,(\pi,\pi')\in\mathcal{P}^2}{\arg\min} \sum_{v\in B} d(v,\pi) + \sum_{v\in U\setminus B} d(v,\pi') + \mu \sum_{(i,j)\in E^B} w_{i,j} \ ,$$

- Approximately solved by an alternated minimization scheme adapted from the original ℓ_0-cut pursuit algorithm
  - Initialization of $[\pi, \pi']$ with RANSAC
  - Optimization wrt B is performed through a graph-cut formulation
  - Optimization wrt π and π' done through least square minimization.
- Converges in 3 iterations

Split step

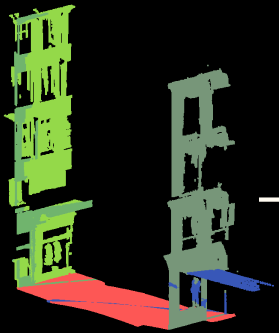

(a) RANSAC initialization

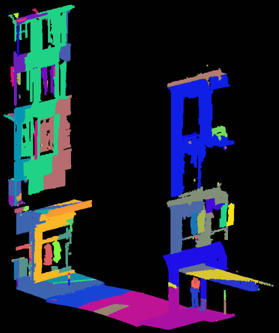

(b) Split step

Merge step: 2D illustration

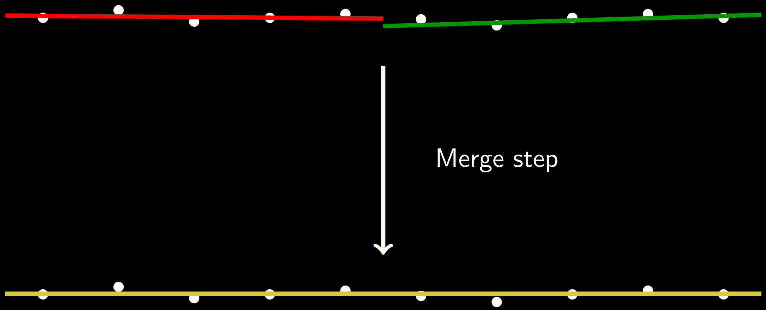

Merge step

ℓ_0-cut pursuit algorithm benefits from allowing the mergeing of adjacent components as long as it decreases the global energy.

In our case, we estimate if two adjacent components should be merged, by computing a common plane and see if it decreases the energy.

Merge step

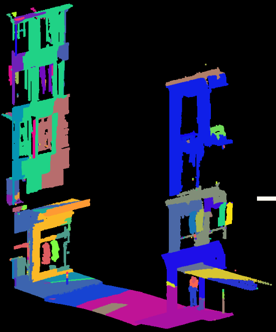

(a) Split step

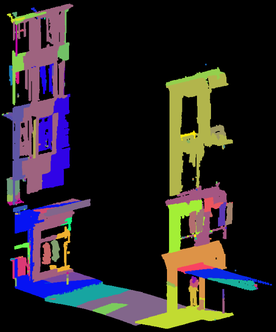

(b) Merge step

Results

Data

Urban dataset

- Acquired with Stereopolis vehicle [Paparoditis et al., 2012]
- $218,546$ points and $418,254$ triangles
- Streets of Paris with large Hausmannian buildings

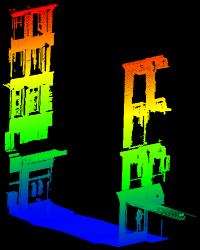

Chapel dataset

- Acquired with TLS by Centre de formation ENSG-Forcalquier
- $1,263,321$ points and $7,313,760$ triangles
- Inside of a chapel with vault portions

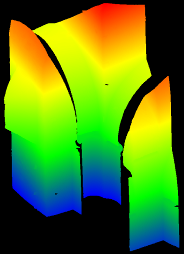

Baseline

We compare our results to our own implementation of the region-growing-based method of Cohen-Steiner et al. [2004], which will serve as baseline.

- Input: mesh and the desired number of regions $\mathcal{N}$
- $\mathcal{N}$ seed triangles are chosen at random
- Growing along each region
- Beginning of iteration:
  - Supporting planes recomputed
  - New triangle seeds

Metric

We want to minimize the distance between each point $p \in V$ and its supporting plane $\pi_v \in \mathcal{P}$:

$$E = \sum_{v \in \mathcal{V}} d(v, \pi_v)^2. \tag{5}$$

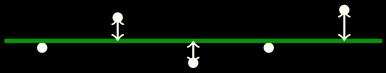

Urban dataset

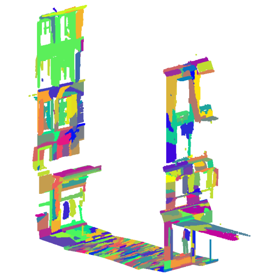

(a) RG, 492 comps, 5 iter, err: $18.3 \cdot 10^3$

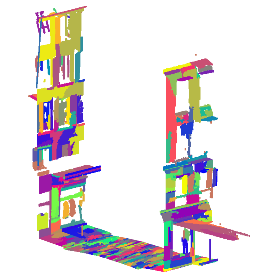

(b) RG, 492 comps, 15 iter, err: $17.5 \cdot 10^3$

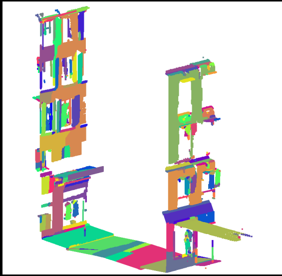

(c) Ours, no merge, 514 comps, err: $1.6 \cdot 10^3$

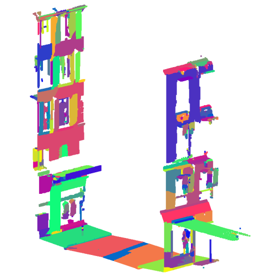

(d) Ours, with 492 comps, err: $1.6 \cdot 10^3$

Urban dataset - Details

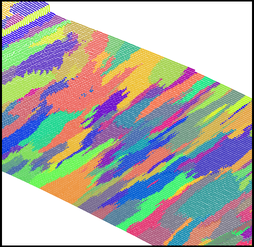

(a) RG, 492 comps, 5 iter, err: $18.3 \cdot 10^3$

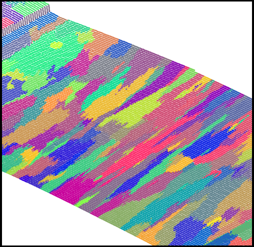

(b) RG, 492 comps, 15 iter, err: $17.5 \cdot 10^3$

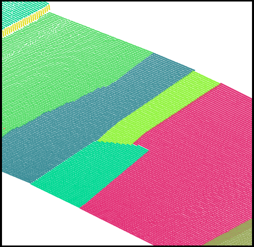

(c) Ours, no merge, 514 comps, err: $1.6 \cdot 10^3$

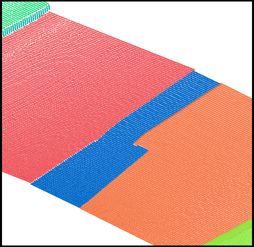

(d) Ours, with 492 comps, err: $1.6 \cdot 10^3$

Urban dataset - Details

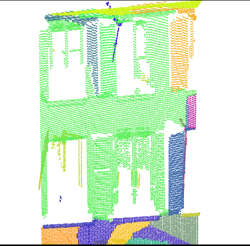

(a) RG, 492 comps, 5 iter, err: $18.3 \cdot 10^3$

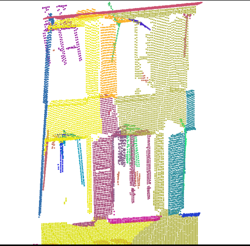

(b) RG, 492 comps, 15 iter, err: $17.5 \cdot 10^3$

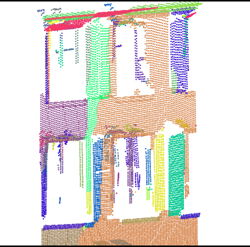

(c) Ours, no merge, 514 comps, err: $1.6 \cdot 10^3$

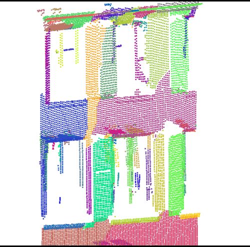

(d) Ours, with 492 comps, err: $1.6 \cdot 10^3$

Details

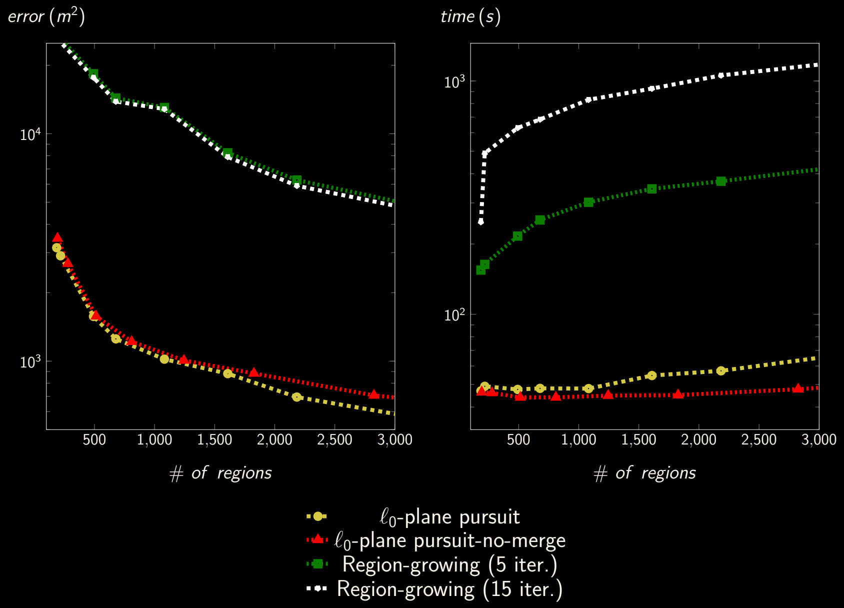

error (m^2)
10^4
10^3
500
1,000
1,500
2,000
2,500
3,000
# of regions
time (s)
10^3
10^2
500
1,000
1,500
2,000
2,500
3,000
# of regions
ℓ_0-plane pursuit
ℓ_0-plane pursuit-no-merge
Region-growing (5 iter.)
Region-growing (15 iter.)

Chapel dataset

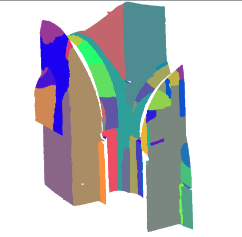

(a) without merge
Reg: 0.1, err: $27 \cdot 10^3$

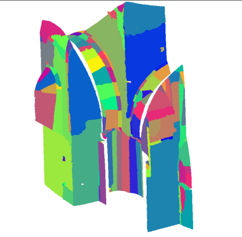

(b) without merge
Reg: 0.01, err: $8 \cdot 10^3$

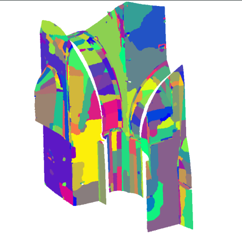

(c) without merge
Reg: 0.001, err: $5 \cdot 10^3$

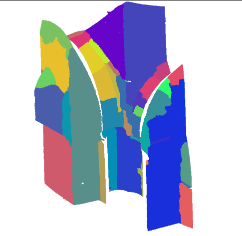

(d) with merge
Reg: 0.1, err: $20 \cdot 10^3$

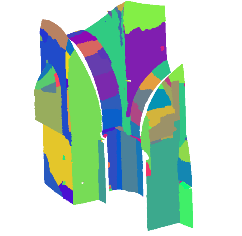

(e) with merge
Reg: 0.01, err: $9.8 \cdot 10^3$

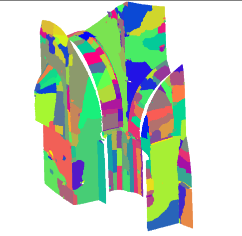

(f) with merge
Reg: 0.001, err: $5 \cdot 10^3$

Details

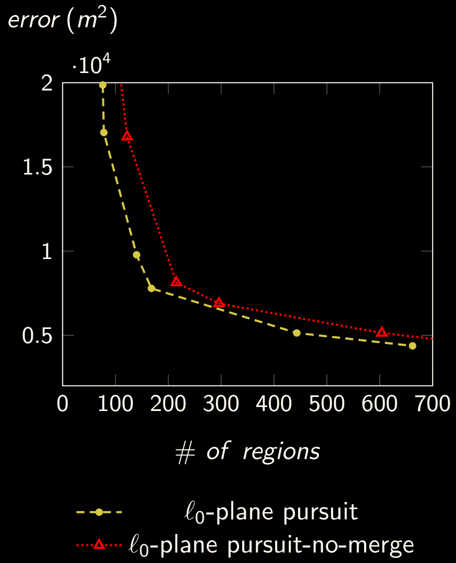

Merge step:

- better fits the data
- decreases number of regions.

Conclusions

Conclusions

ℓ_0-plane pursuit:

- \+ piecewise-planar segmentation of large 3D point clouds
- \+ faster than region growing-based approaches
- \+ better approximations / locally adaptive
- – lose topological connection between adjacent regions.

Improvments include the use of multiple primitives for non-planar parts of the cloud [Vidal et al., 2014]. This method can also serve as a baseline for a polyhedral reconstruction algorithm.

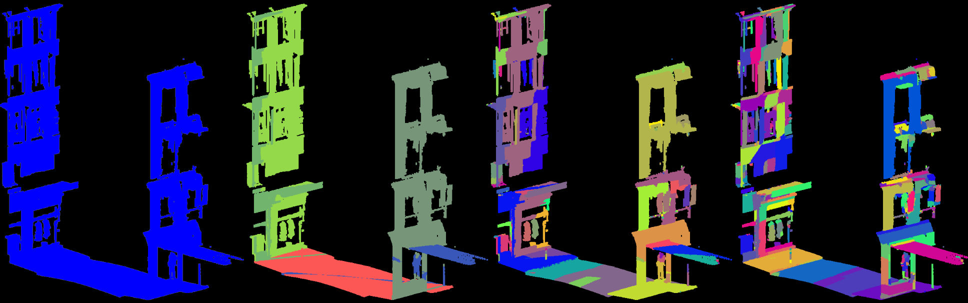

a: Point cloud **b:** Initialization **c:** First iteration **d:** Final result

Bibliography I

A.-L. Chauve, P. Labatut, and J.-P. Pons. Robust piecewise-planar 3d reconstruction and completion from large-scale unstructured point data. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 1261–1268. IEEE, 2010.

D. Cohen-Steiner, P. Alliez, and M. Desbrun. Variational shape approximation. In *ACM Transactions on Graphics (ToG)*, volume 23, pages 905–914. ACM, 2004.

W. Dai, B. Yang, Z. Dong, and A. Shaker. A new method for 3d individual tree extraction using multispectral airborne lidar point clouds. *ISPRS journal of photogrammetry and remote sensing*, 144: 400–411, 2018.

Bibliography II

J. Demantke, C. Mallet, N. David, and B. Vallet. Dimensionality based scale selection in 3d lidar point clouds. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 38(Part 5):W12, 2011.

A. Dutta, J. Engels, and M. Hahn. Segmentation of laser point clouds in urban areas by a modified normalized cut method. *IEEE transactions on pattern analysis and machine intelligence*, 2018.

A. Ferraz, F. Bretar, S. Jacquemoud, G. Gonçalves, L. Pereira, M. Tomé, and P. Soares. 3-d mapping of a multi-layered mediterranean forest using als data. *Remote Sensing of Environment*, 121:210–223, 2012.

Bibliography III

S. Guinard and B. Vallet. Sensor-topology based simplicial complex reconstruction from mobile laser scanning. *ISPRS Annals of Photogrammetry, Remote Sensing & Spatial Information Sciences*, 4 (2), 2018.

T. Holzmann, M. R. Oswald, M. Pollefeys, F. Fraundorfer, and H. Bischof. Plane-based surface regularization for urban 3d reconstruction. In *28th British Machine Vision Conference*, 2017.

L. Landrieu and G. Obozinski. Cut pursuit: Fast algorithms to learn piecewise constant functions on general weighted graphs. *SIAM Journal on Imaging Sciences*, 10(4):1724–1766, 2017.

Bibliography IV

L. Ma, R. Favier, L. Do, E. Bondarev, and P. H. de With. Plane segmentation and decimation of point clouds for 3d environment reconstruction. In *2013 IEEE 10th Consumer Communications and Networking Conference (CCNC)*, pages 43–49. IEEE, 2013.

D. Munoz, J. A. Bagnell, N. Vandapel, and M. Hebert. Contextual classification with functional max-margin markov networks. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 975–982. IEEE, 2009.

S. Oesau, F. Lafarge, and P. Alliez. Planar shape detection and regularization in tandem. In *Computer Graphics Forum*, volume 35, pages 203–215. Wiley Online Library, 2016.

Bibliography V

N. Paparoditis, J.-P. Papelard, B. Cannelle, A. Devaux, B. Soheilian, N. David, and E. Houzay. Stereopolis ii: A multi-purpose and multi-sensor 3d mobile mapping system for street visualisation and 3d metrology. *Revue française de photogrammétrie et de télédétection*, 200(1):69–79, 2012.

R. Schnabel, R. Wahl, and R. Klein. Efficient ransac for point-cloud shape detection. In *Computer graphics forum*, volume 26, pages 214–226. Wiley Online Library, 2007.

V. Vidal, C. Wolf, and F. Dupont. *Mechanical Mesh Segmentation and Global 3D Shape Extraction*. PhD thesis, Université Lyon 1-Claude Bernard; INSA Lyon, 2014.

Bibliography VI

T. Whelan, L. Ma, E. Bondarev, P. de With, and J. McDonald. Incremental and batch planar simplification of dense point cloud maps. *Robotics and Autonomous Systems*, 69:3–14, 2015.

B. Xu, W. Jiang, J. Shan, J. Zhang, and L. Li. Investigation on the weighted ransac approaches for building roof plane segmentation from lidar point clouds. *Remote Sensing*, 8(1):5, 2015.