From Points to Prints

Generating Building Roofprints and Footprints from Airborne Lidar Data and Inaccurate Outlines

A4 Presentation

Alexandre Bry, supervised by Hugo Ledoux (TUD), Ravi Peters (TUD) and Bruno Vallet (IGN)

June 26, 2026

Title

—

Context

Context

–

Roofprint vs. footprint

	Roofprint	Footprint
Defined by	Roof	Façades
Comes from	Aerial data (images, ALS)	Terrain measurements (TLS, MLS)
Used for	3D modelling, solar potential	Tax estimation, energy consumption

Table 1: Comparison between roofprints and footprints

Context
– Roofprint vs. footprint

- Roofprints and footprints can serve different purposes, and knowing the difference between them is important when the roof overhangs are significant.

	Roofprint	Footprint
Defined by	Roof	Façades
Comes from	Aerial data (images, ALS)	Terrain measurements (TLS, MLS)
Used for	3D modelling, solar potential	Tax estimation, energy consumption

Table 1: Comparison between roofprints and footprints

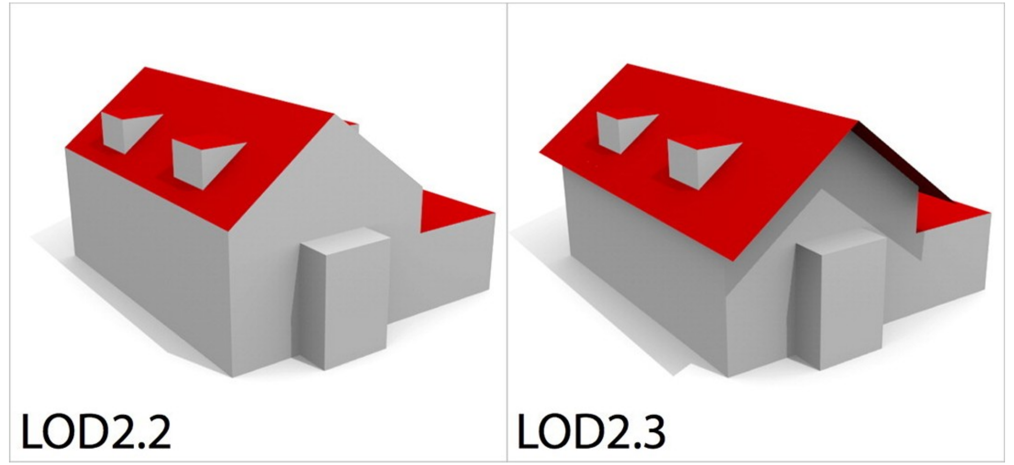

Figure 1: Part of the visual definition of buildings Level of Detail (LoD) containg LoD 2.2 and 2.3 (Biljecki et al., 2016).

- Roofprints and footprints can serve different purposes, and knowing the difference between them is important when the roof overhangs are significant.
- Since roof overhangs are usually not available, it is unclear what the use cases could be, but we can think about:
  - More accurate visualisation and texturing of 3D models
  - More accurate simulations (wind, luminosity, visibility)

Strengths and weaknesses of BD TOPO

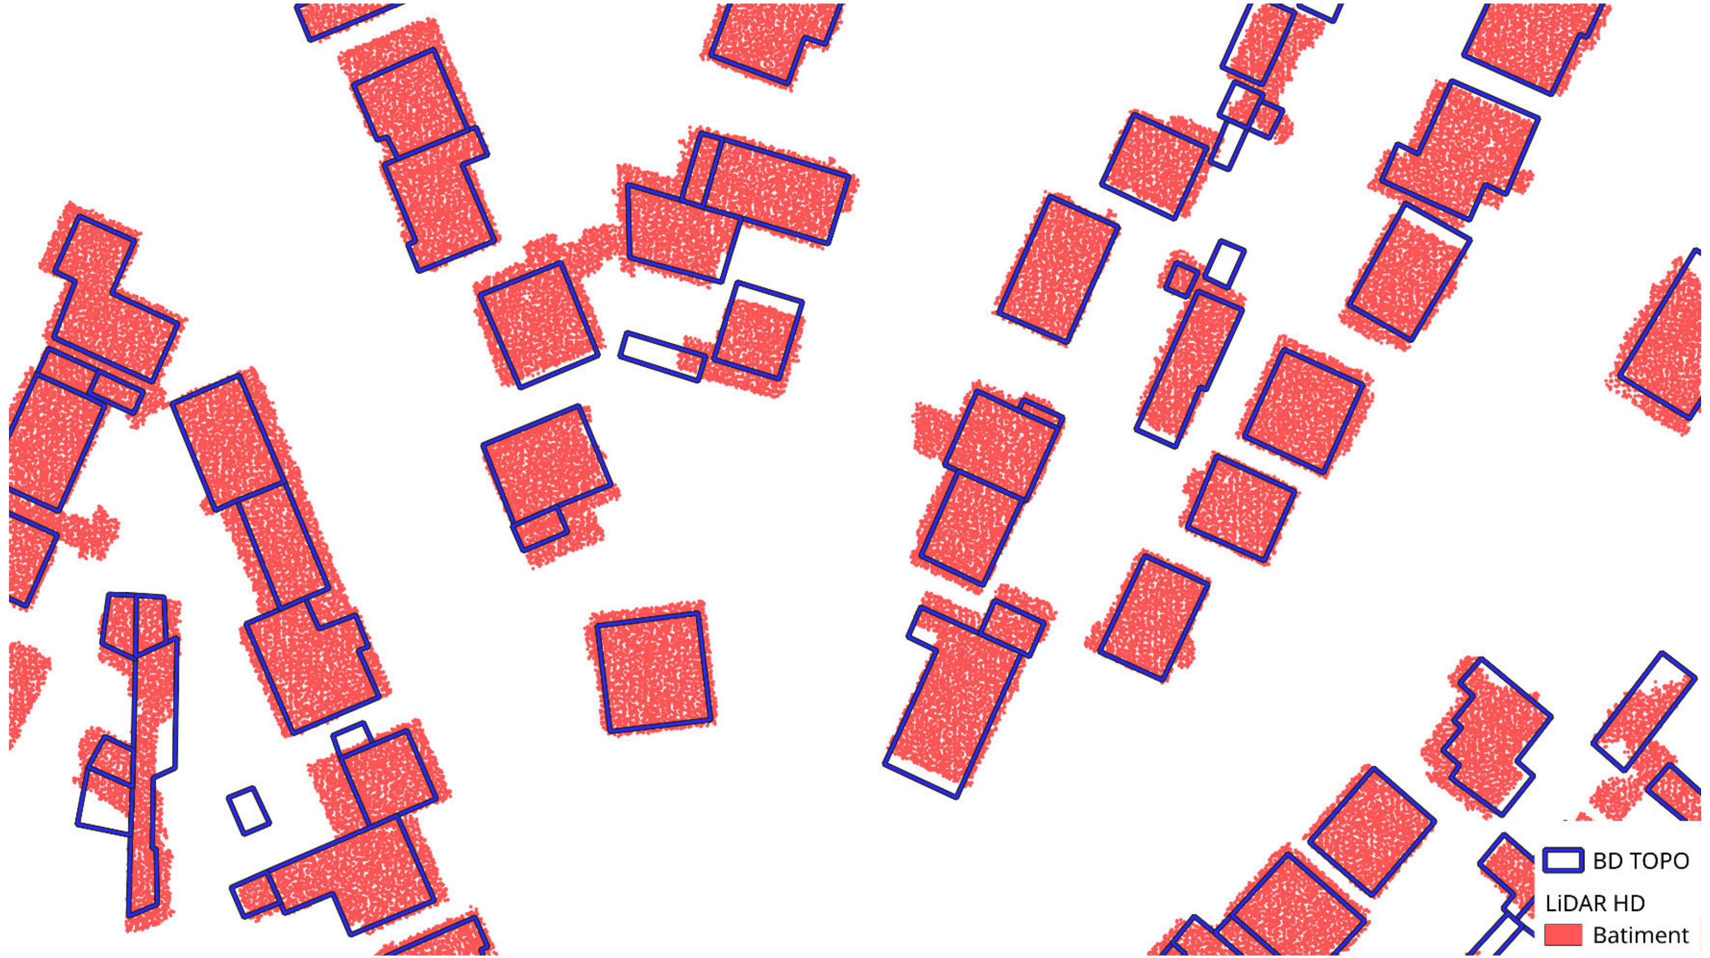

Figure 2: BD TOPO buildings are most of the time footprints, and often shifted.

- Strengths:
  - The BD TOPO contains most of the buildings in France
  - Buildings were manually delineated by humans, so their shape is generally of good quality
- Weaknesses:
  - It is a mix of footprints and roofprints, coming from different sources
  - Georeferencing is not perfect because it was not crucial for the cadastre initially

Research questions

> *How to generate coherent building **roofprints** and **footprints** from high-density **ALS** point clouds and existing **imprecise outlines**?*

Overview of the methodology

1. BD TOPO and LiDAR HD

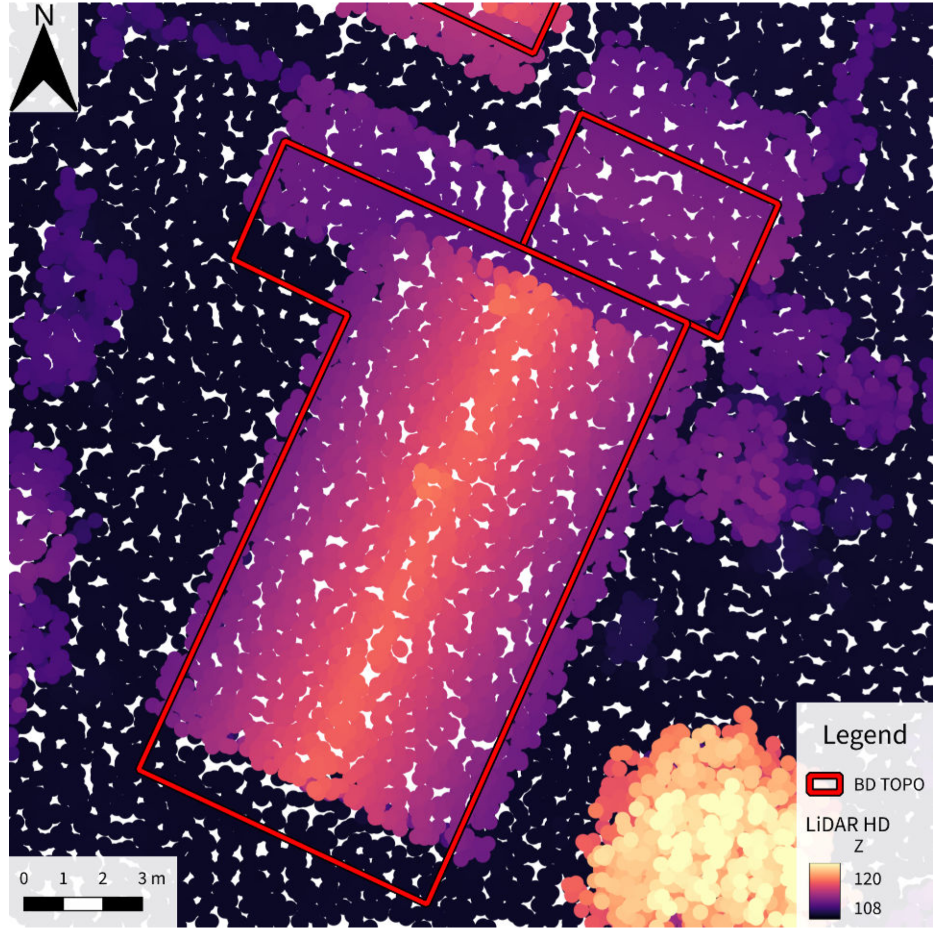

2. Find roof edge points

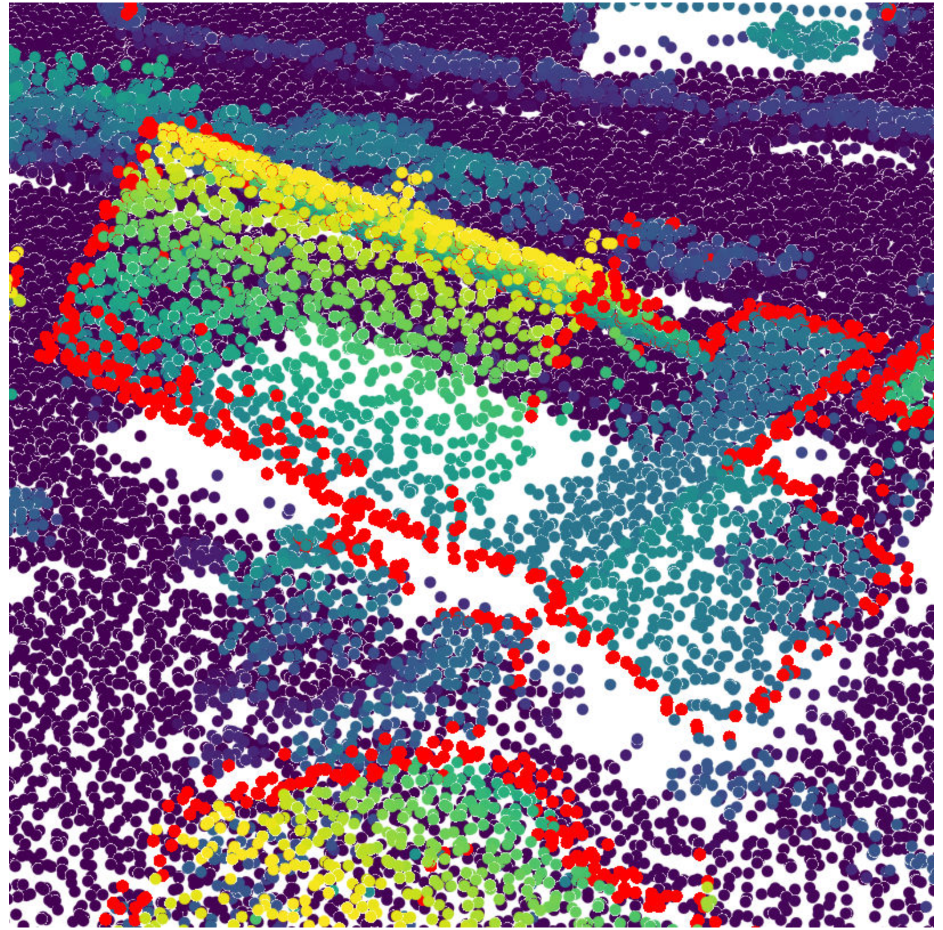

Context
– Overview of the methodology

The BD TOPO and the LiDAR HD datasets are initially not aligned. The first step is to identify the points which are at the edge of the roof.

Overview of the methodology

2. Find roof edge points

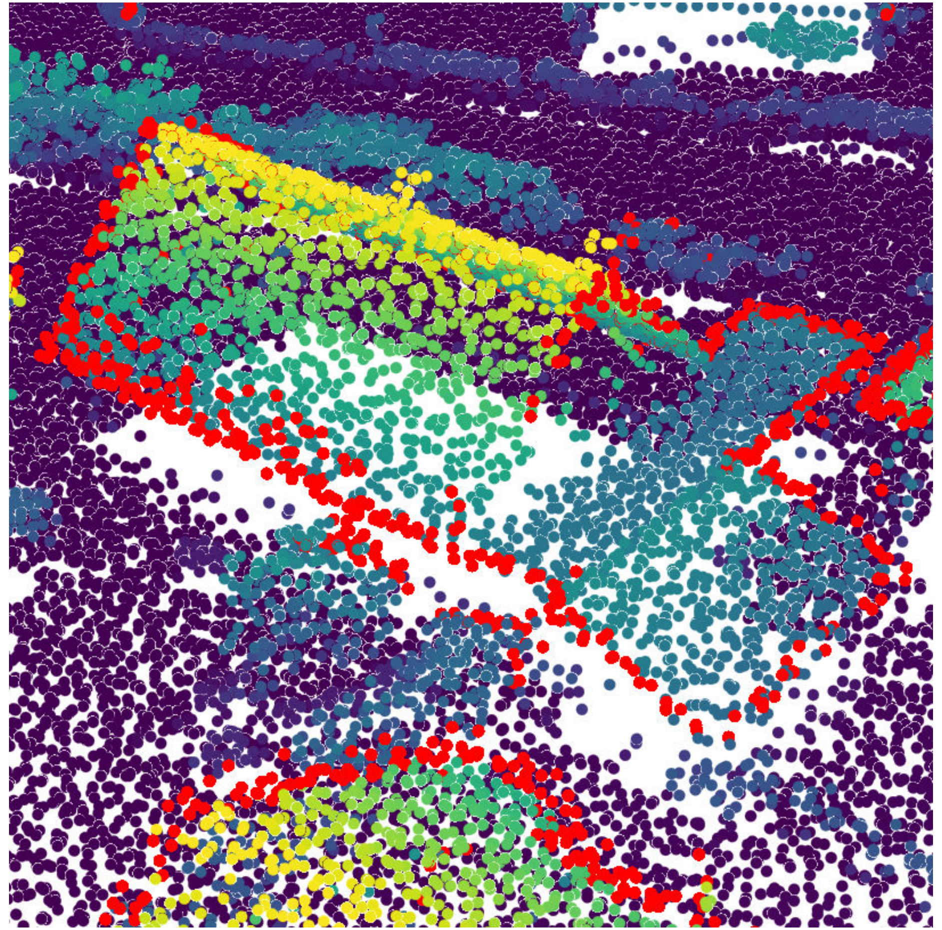

3. Align into a roofprint

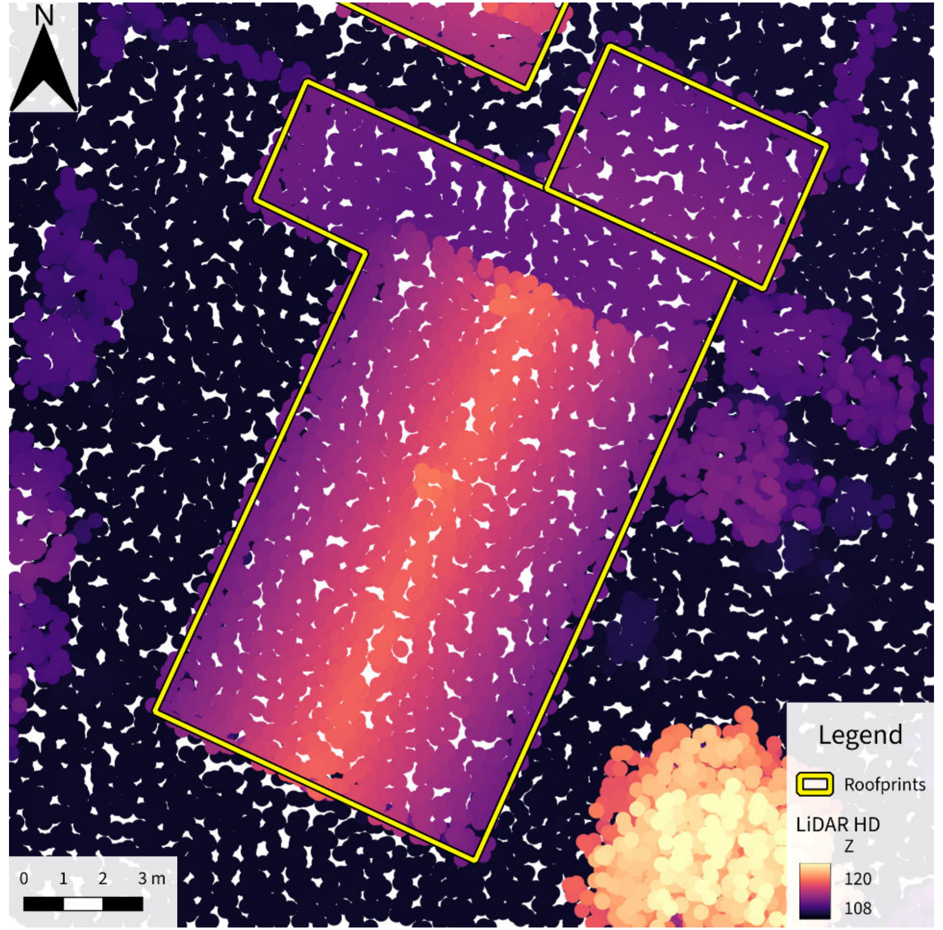

Context
– Overview of the methodology

Then, these points are used to deform the initial outline and align it on the roof edge points to create a roofprint.

Overview of the methodology

3. Align into a roofprint

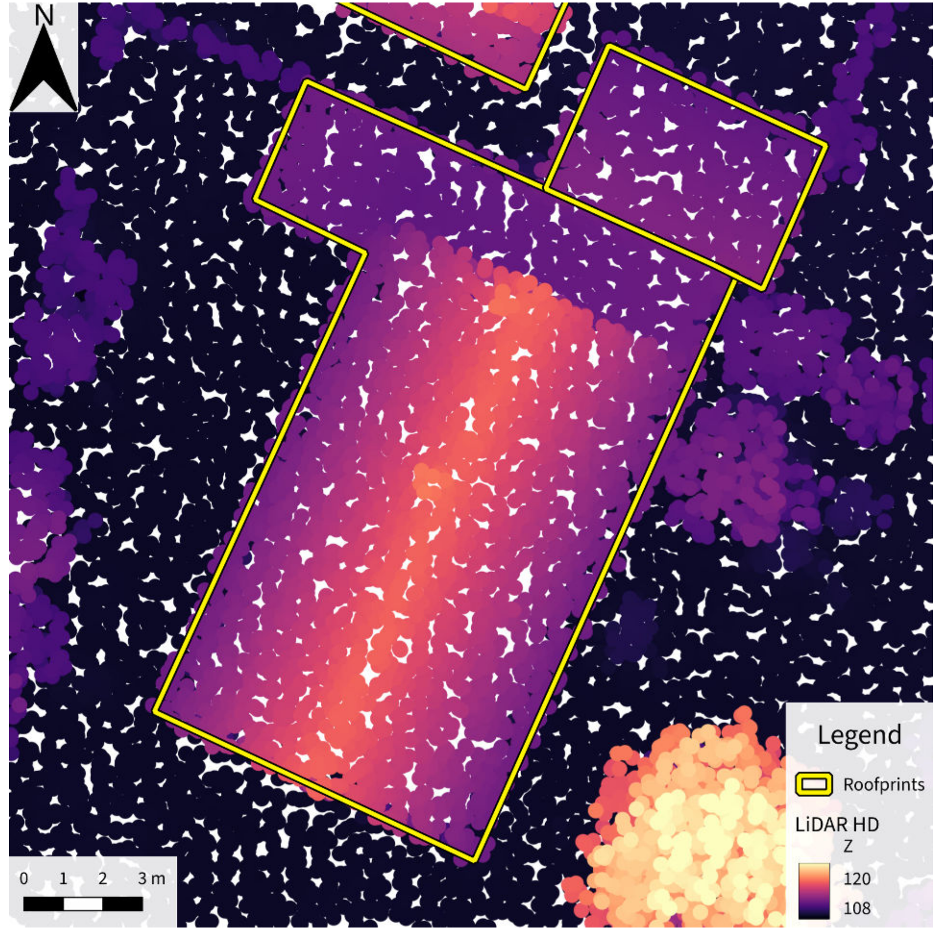

4. Generate 3D roof model

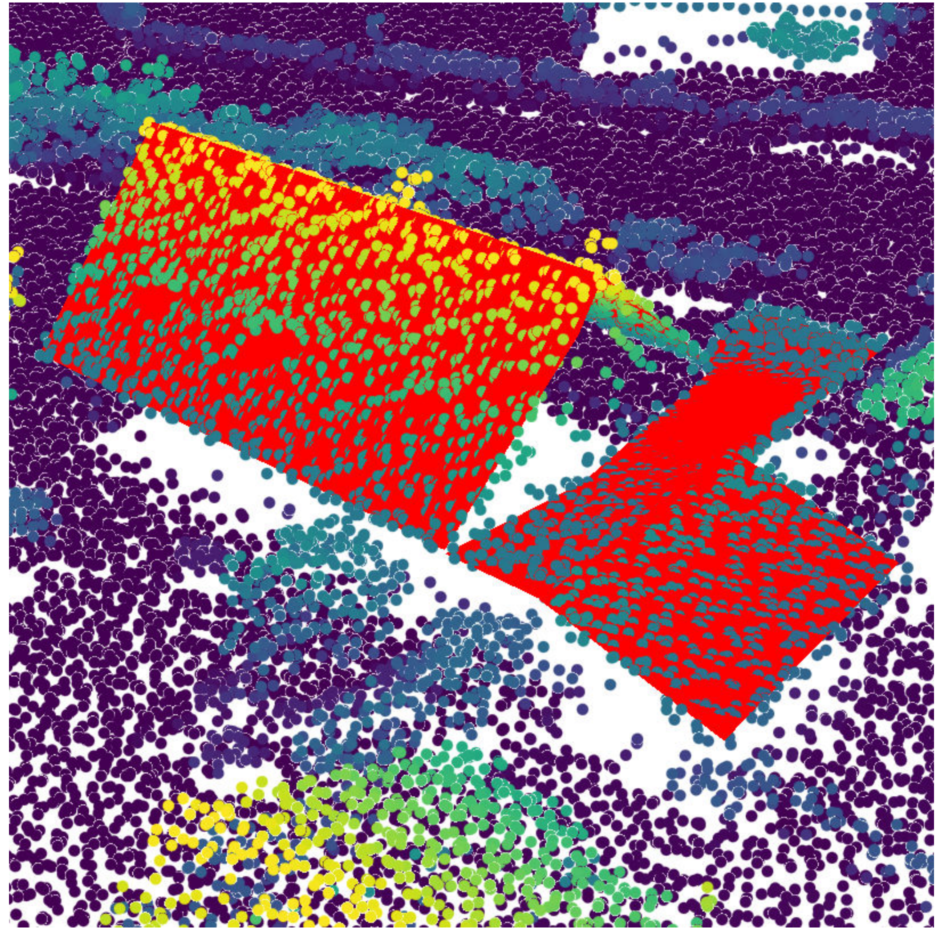

Context
– Overview of the methodology

After that, this clean roofprint can be used to produce and clean 3D roof model.

Overview of the methodology

4. Generate 3D roof model

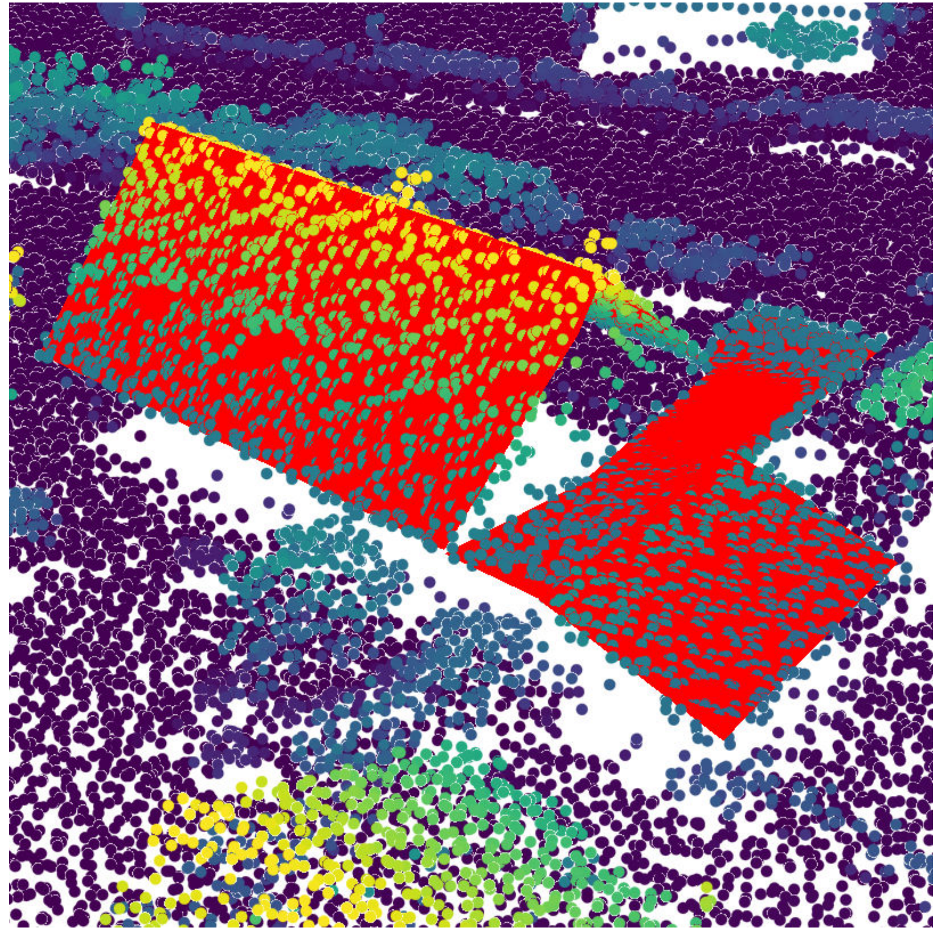

5. Find façade and ground points

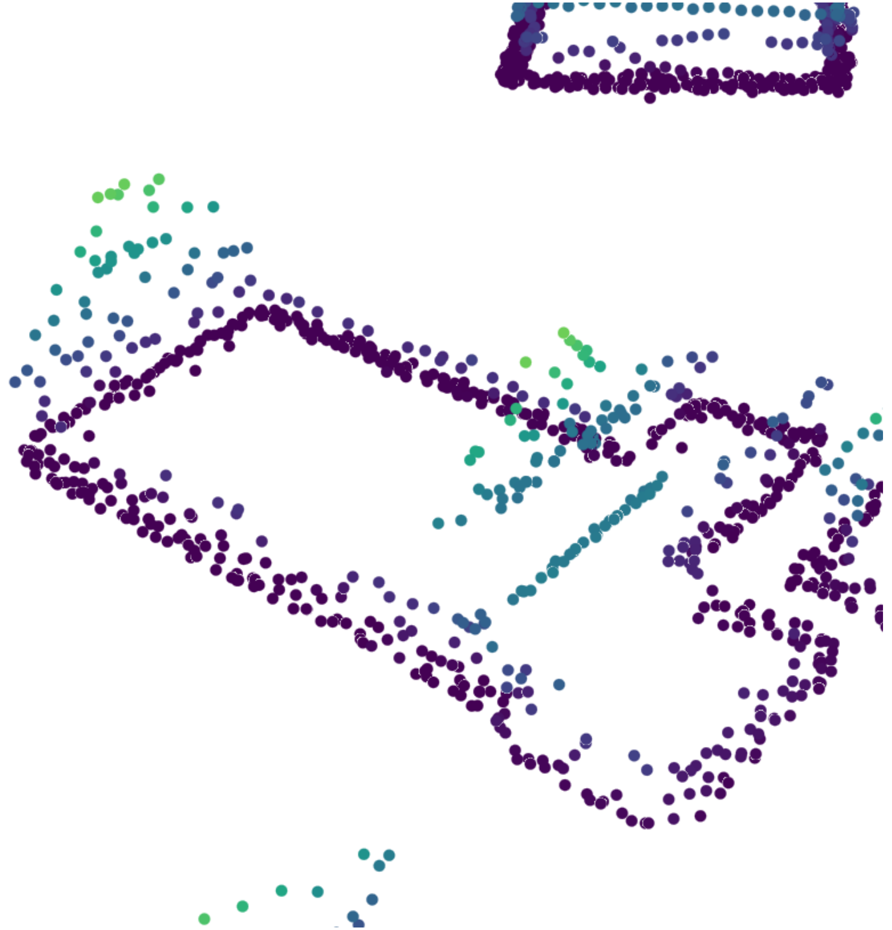

This 3D roof model can then be used to identify useful points on the façades and on the ground, by simply filtering the points geometrically under the roof.

Overview of the methodology

5. Find façade and ground points

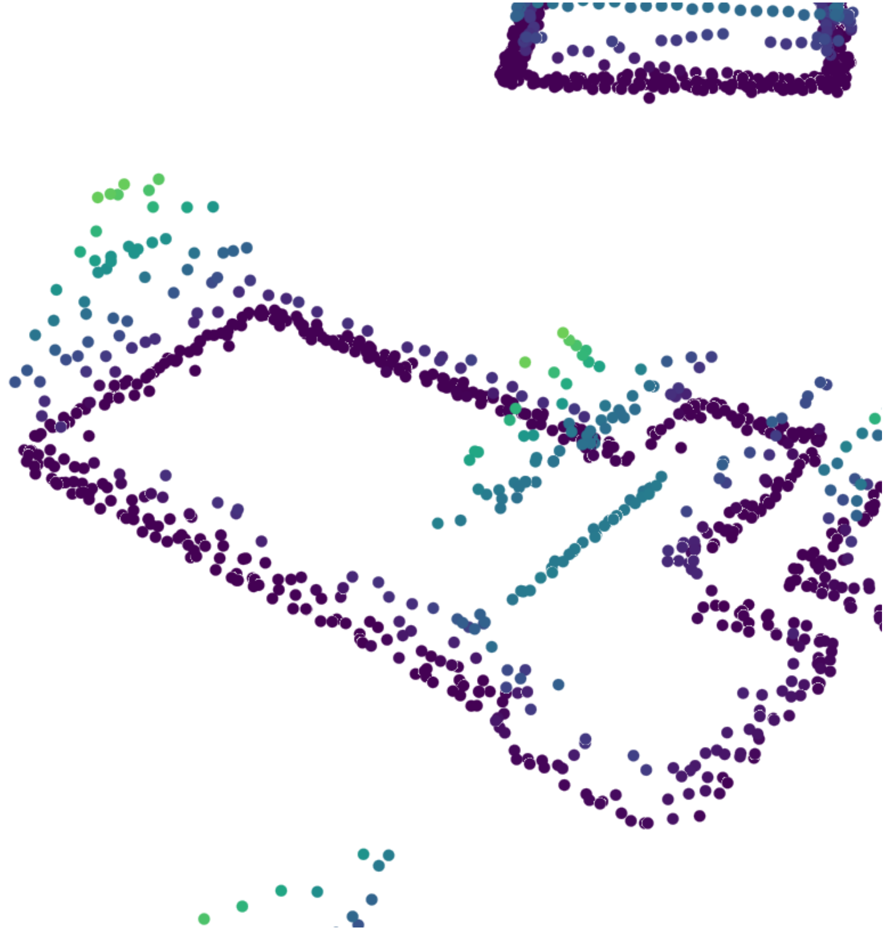

6. Align into a footprint

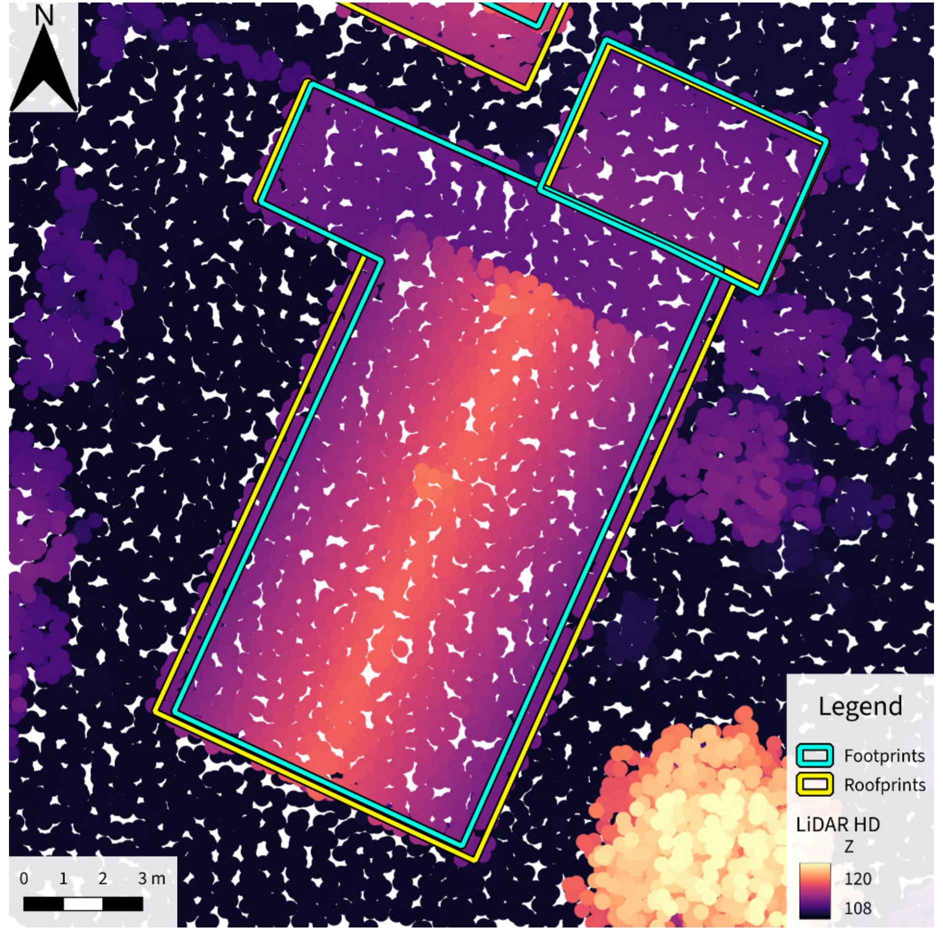

Context
– Overview of the methodology

The filtered points are finally used to deform the previous roofprint into a footprint.

Polygon deformation

–

Our simple approach

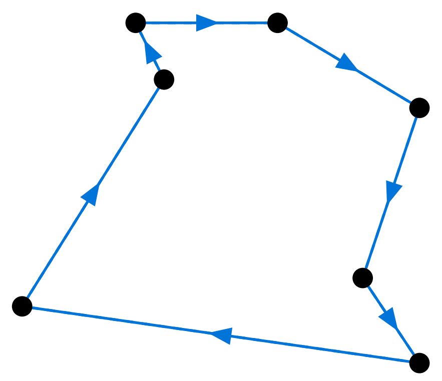

(a) Initial state.

Figure 3: Illustration of the polygon deformation algorithm on a single polygon.

Flipped edges are indicated with red arrows, and shifted edges are green.

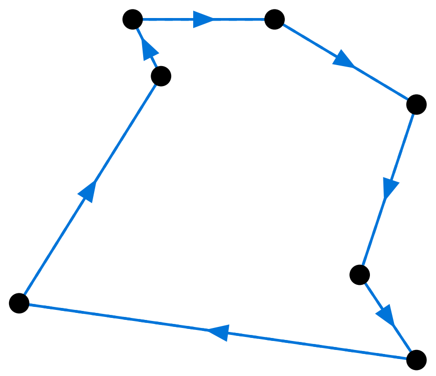

(a) Initial state.

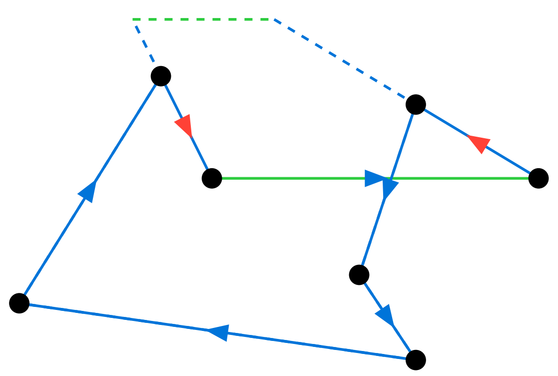

(b) After shifting the focus edge.

Figure 3: Illustration of the polygon deformation algorithm on a single polygon.

Polygon deformation
– Our simple approach

Flipped edges are indicated with red arrows, and shifted edges are green.

Our simple approach

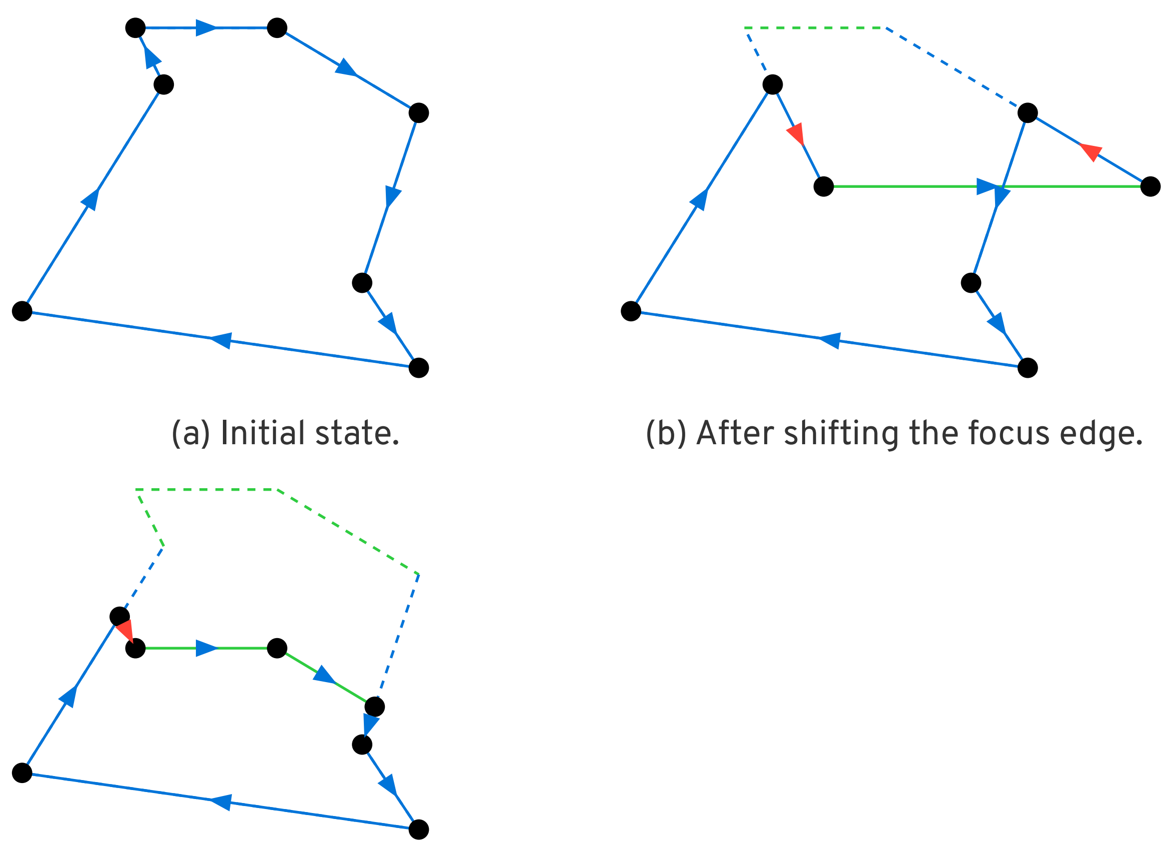

(a) Initial state.

(b) After shifting the focus edge.

(c) First step of the resolution.

Figure 3: Illustration of the polygon deformation algorithm on a single polygon.

Flipped edges are indicated with red arrows, and shifted edges are green.

Our simple approach

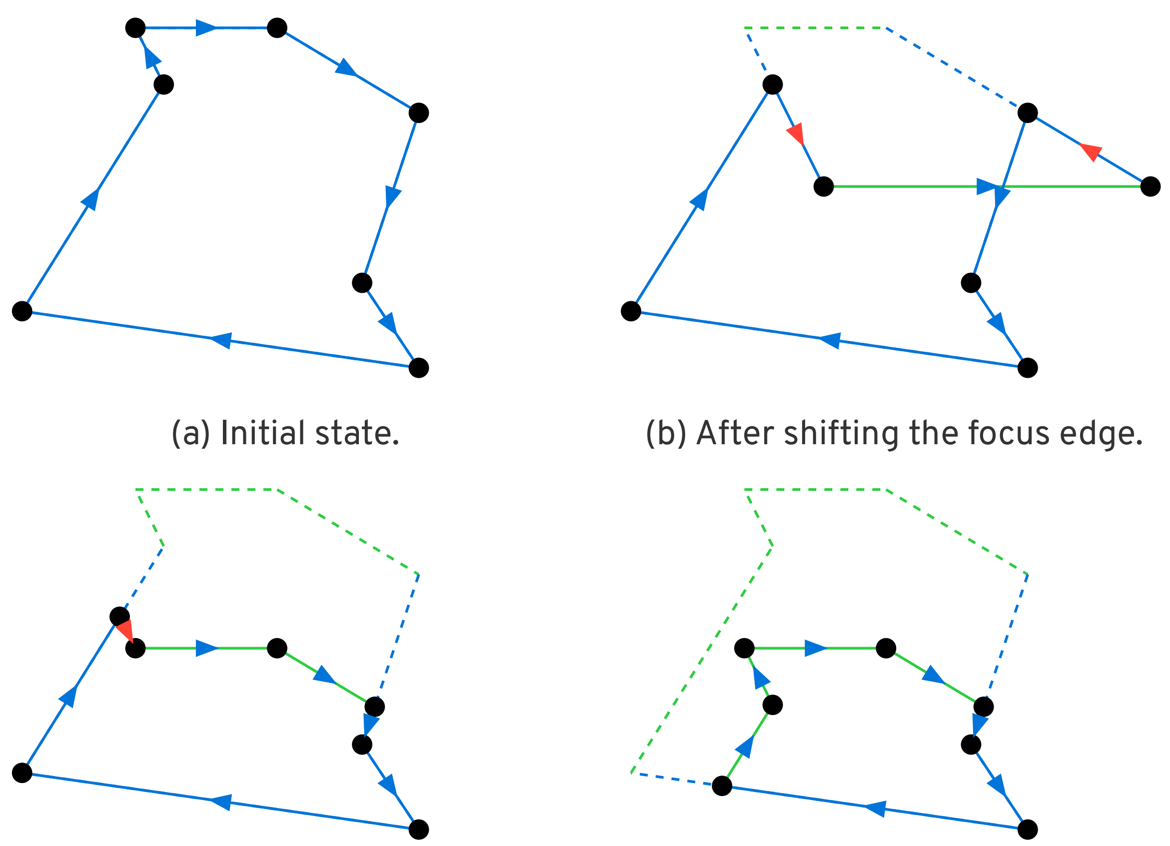

(a) Initial state. (b) After shifting the focus edge.

(c) First step of the resolution. (d) Second and final step of the resolution.

Figure 3: Illustration of the polygon deformation algorithm on a single polygon.

Flipped edges are indicated with red arrows, and shifted edges are green.

Roofprints

Roofprints

—

Point cloud topology

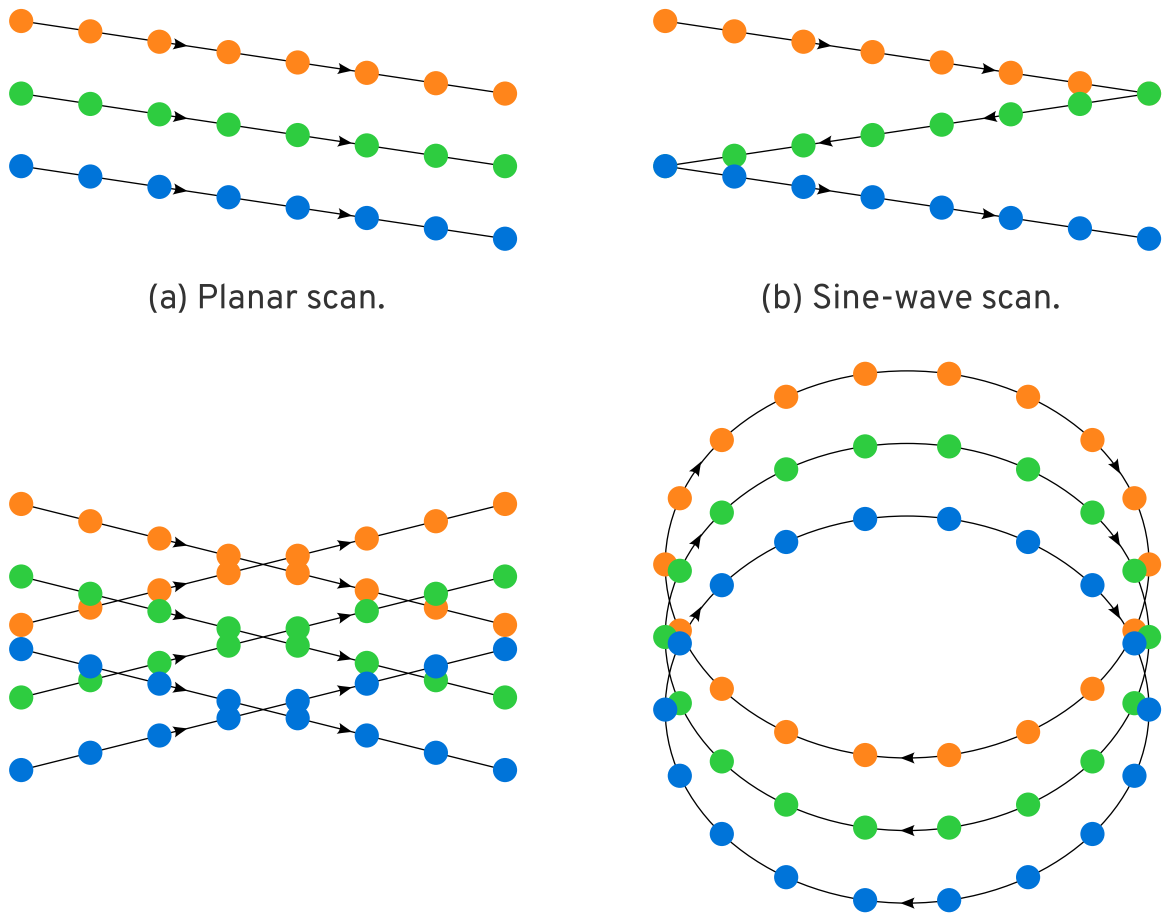

(a) Planar scan.

(b) Sine-wave scan.

(c) Cross scan.

(d) Circular scan.

Figure 4: Illustration in 2D (view from the top) of different types of ALS sensors. inspired from (Wu et al., 2026) .

These are the four types of Airborne Lidar Scanning (ALS) sensors used in the LiDAR HD project.

Roofprints evidence

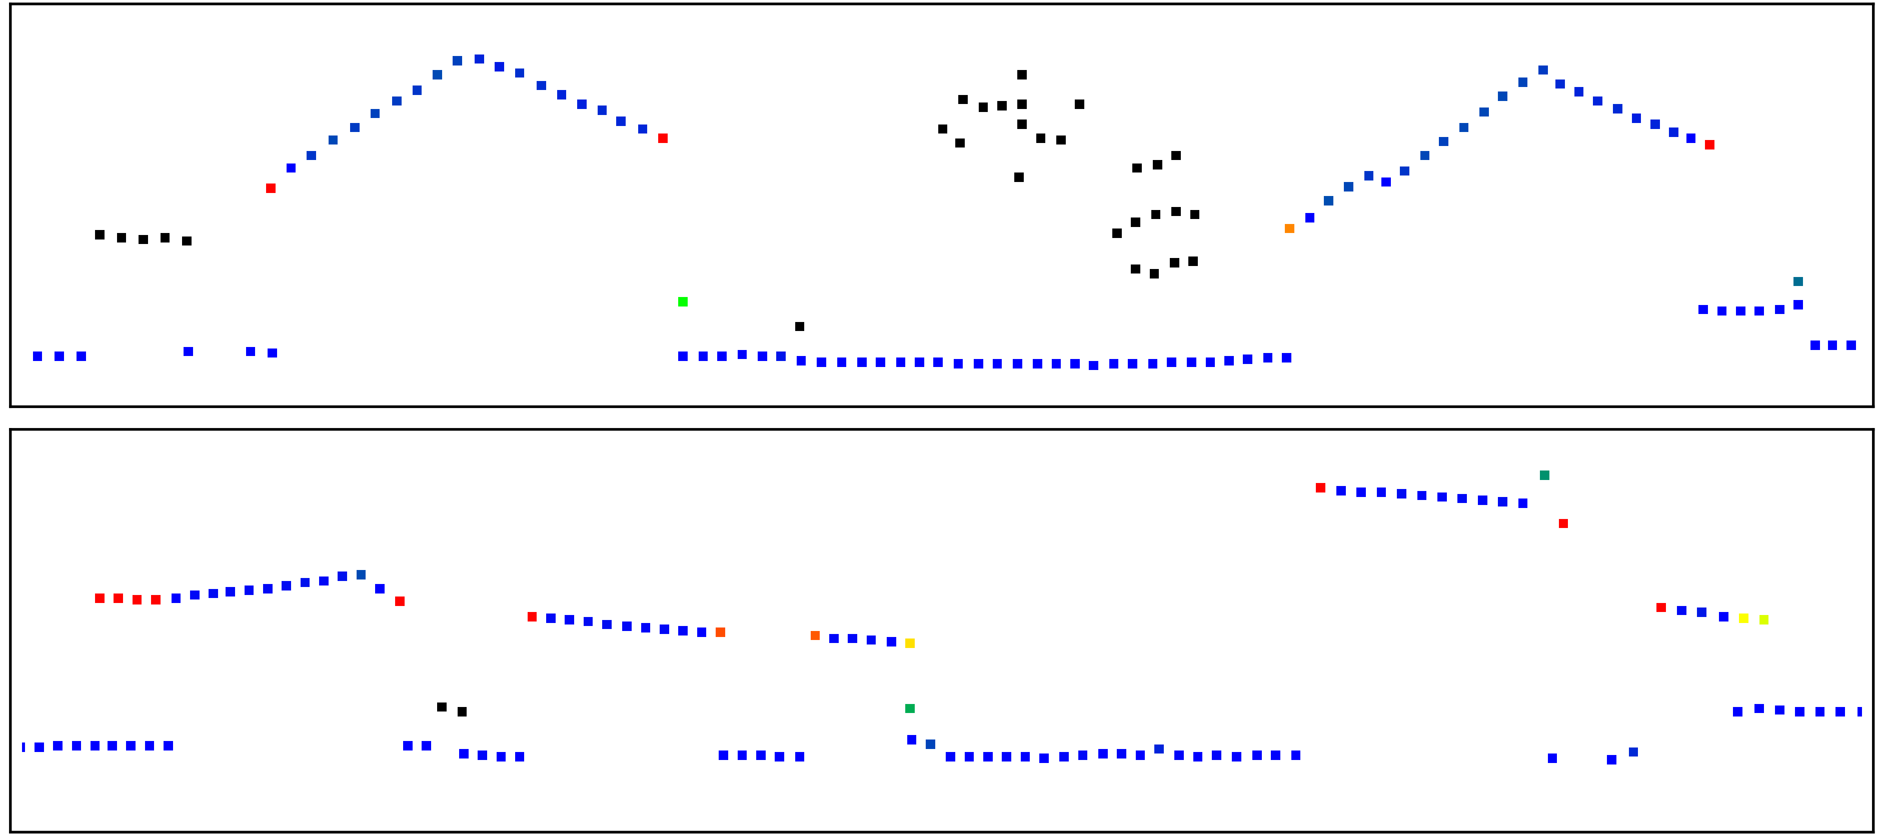

Figure 5: Two examples of the height differences obtained on a scan line. Black points are vegetation points, and for the other points the colour represents the value of Δh, with blue-green-yellow-red from low to high values.

Roofprints
– Roofprints evidence

We use the **height differences Δh between neighbor points** to identify points that are likely to be on the edges of roofs:

- We look at the neighbors in the previous, current and next scan lines
- Points at the edges of building roofs have **high values of Δh**
- Points from **high vegetation** also get high Δh and therefore need to be excluded
- With circular ALS sensors, the scan lines are **curved** when looked at from above, but this is not a problem as the **angle is very small** between consecutive pulses

Assign a score to roofprints

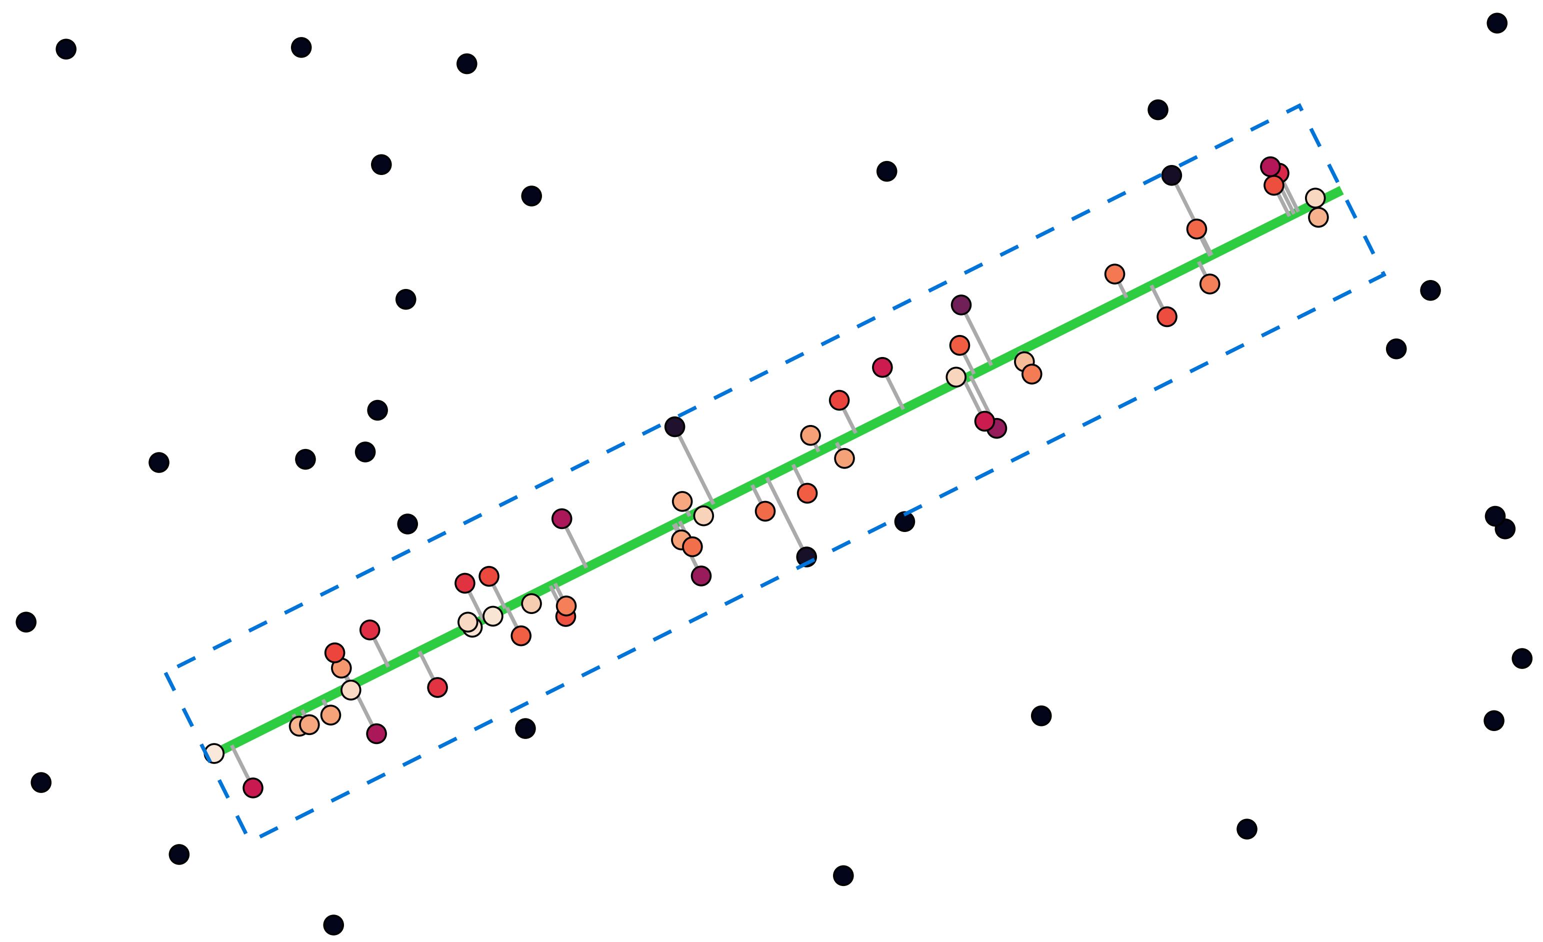

Figure 6: Illustration in 2D (view from the top) of the proximity score for an edge.

The score assigned to each edge is based on the projection of each point on the edge and the distance to it. A comparison between the normal of the edge and the normal of the points is also used to enforce correct orientations.

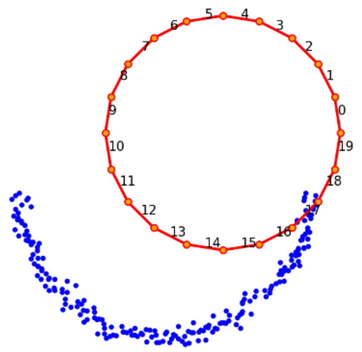

(a) Initial state

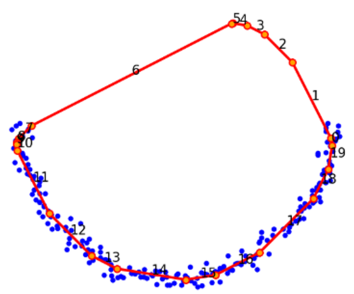

(b) $\alpha = 0.00$ (no regularisation)

Figure 7:

There is also a need for regularization, because without it, the shape can be arbitrarily deformed into a shape very different from the initial one.

Need for regularization

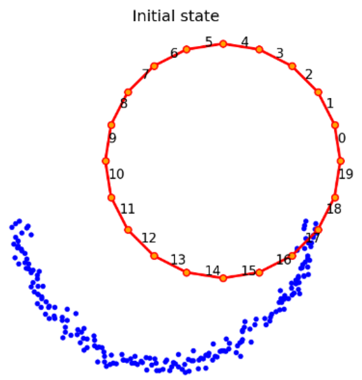

(a) Initial state

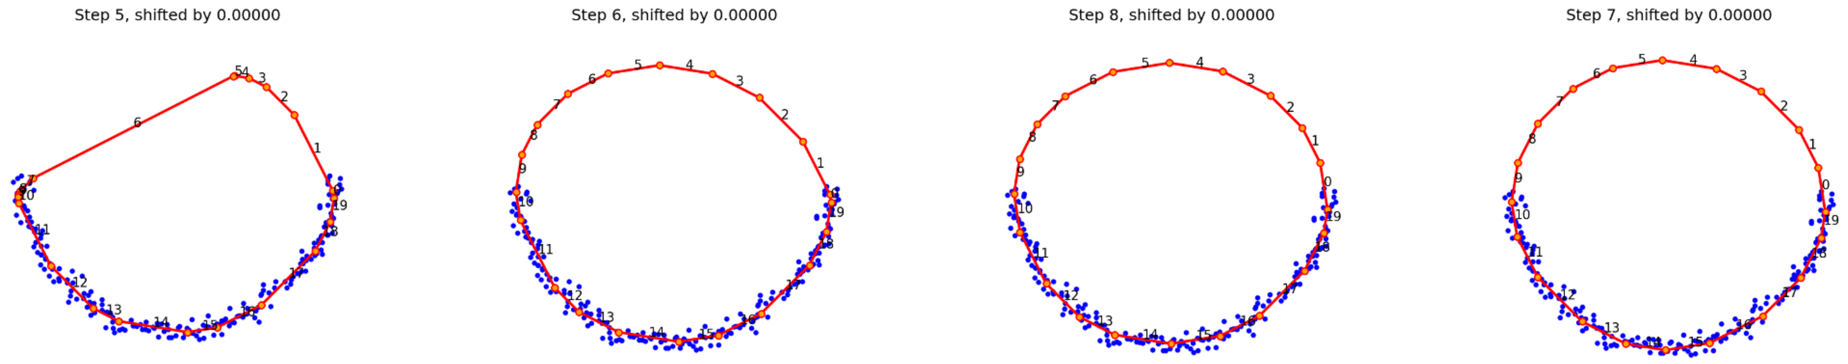

(b) $\alpha = 0.00$ (no regularisation) (c) $\alpha = 0.05$ (d) $\alpha = 0.20$ (e) $\alpha = 1.00$

Figure 25: Matching a circle to a half-circle of points with different values of the regularization parameter α.

Roofprints
– Need for regularization

We can see how regularization helps to get a shape more similar to the target even with half of the points missing.

Need for regularization

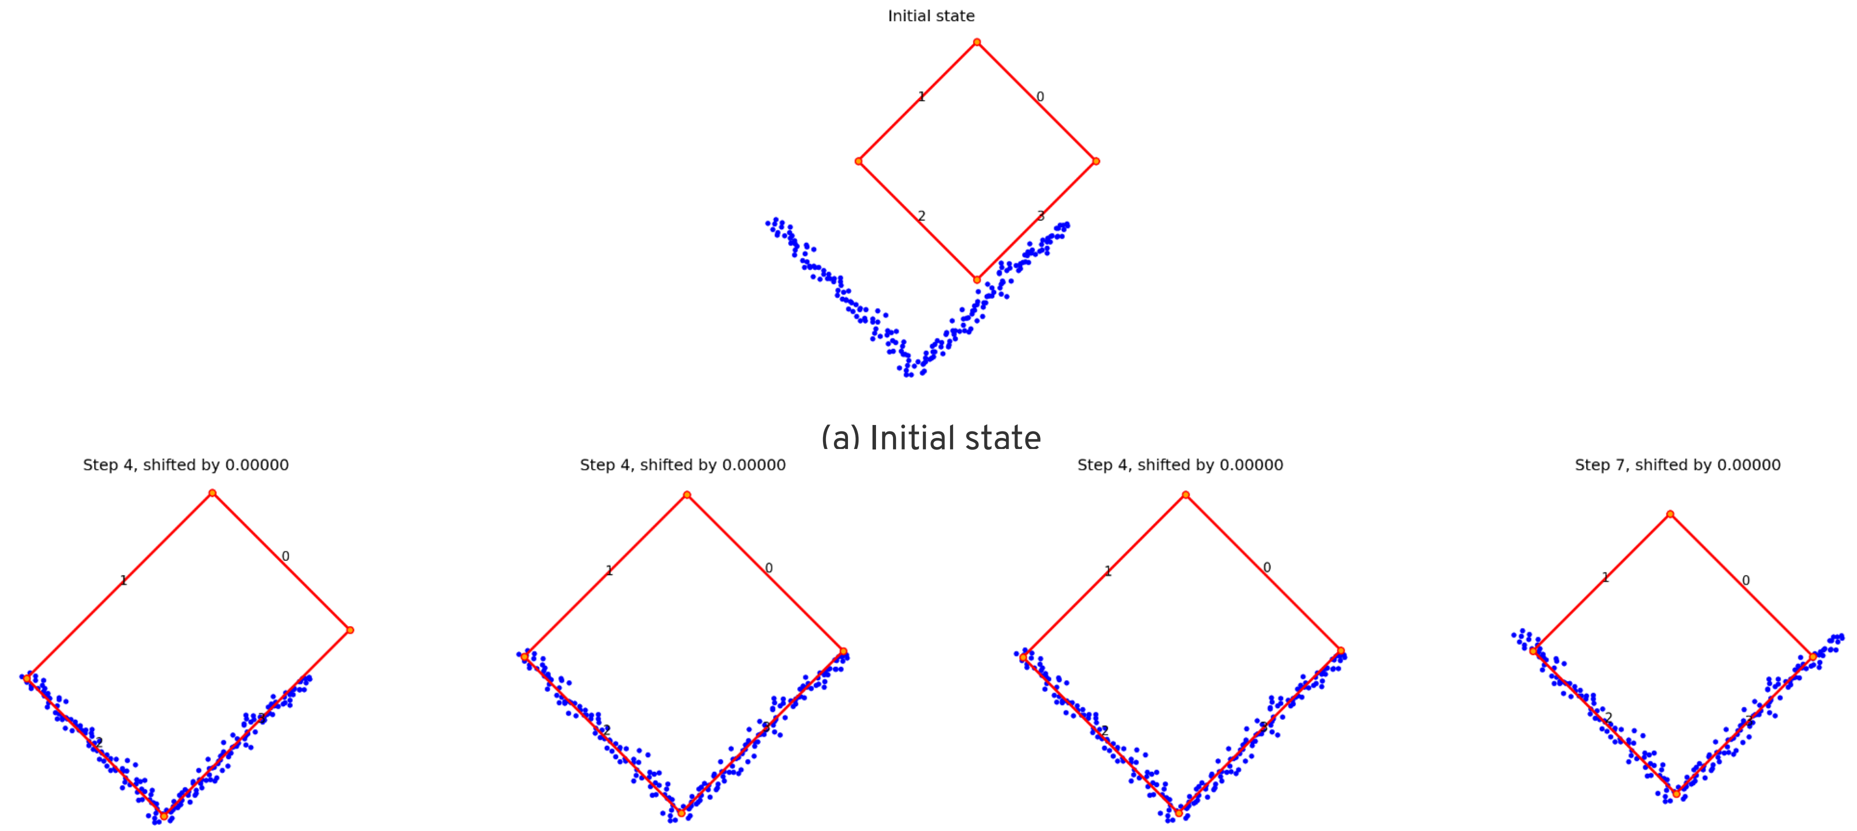

(a) Initial state

(b) $\alpha = 0.00$ (no regularization) (c) $\alpha = 0.05$ (d) $\alpha = 0.20$ (e) $\alpha = 1.00$

Figure 25: Matching a square to points along two sides of a larger square with different values of the regularization parameter α.

$\alpha > 0$ encourages the shape to be closer to the initial one, which is better than $\alpha = 0$, until a certain point where the regularization is too strong and prevents the shape from extending to capture all the points.

Need for regularization

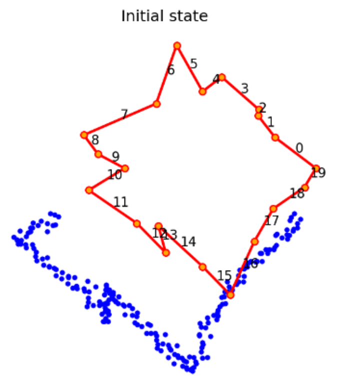

(a) Initial state

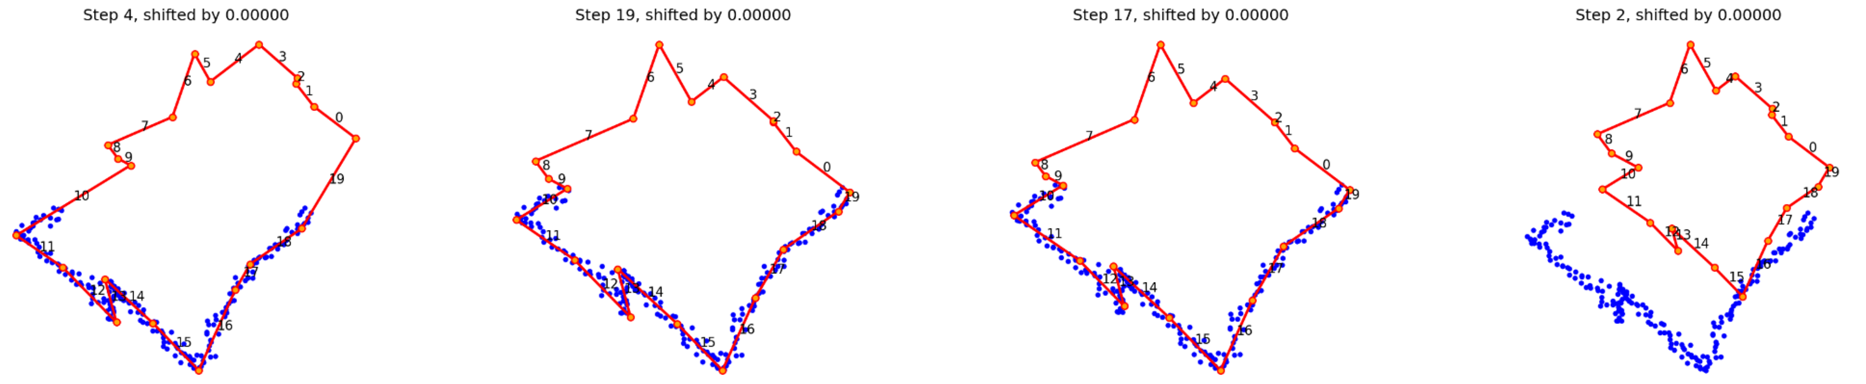

(b) $\alpha = 0.00$ (no regularization) (c) $\alpha = 0.05$ (d) $\alpha = 0.20$ (e) $\alpha = 1.00$

Figure 25: Matching a polygon to half of its shape with points with different values of the regularization parameter α.

- $\alpha > 0$ forces the shape to be closer to the initial one, but it takes many more steps to converge, because the more edges that don't have points take longer to balance the regularization term between each other.
- If α is too high, it prevents the shape from matching the points, because the proximity score is not good enough to balance the regularization term, which is the case here for $\alpha \geq 0.50$.

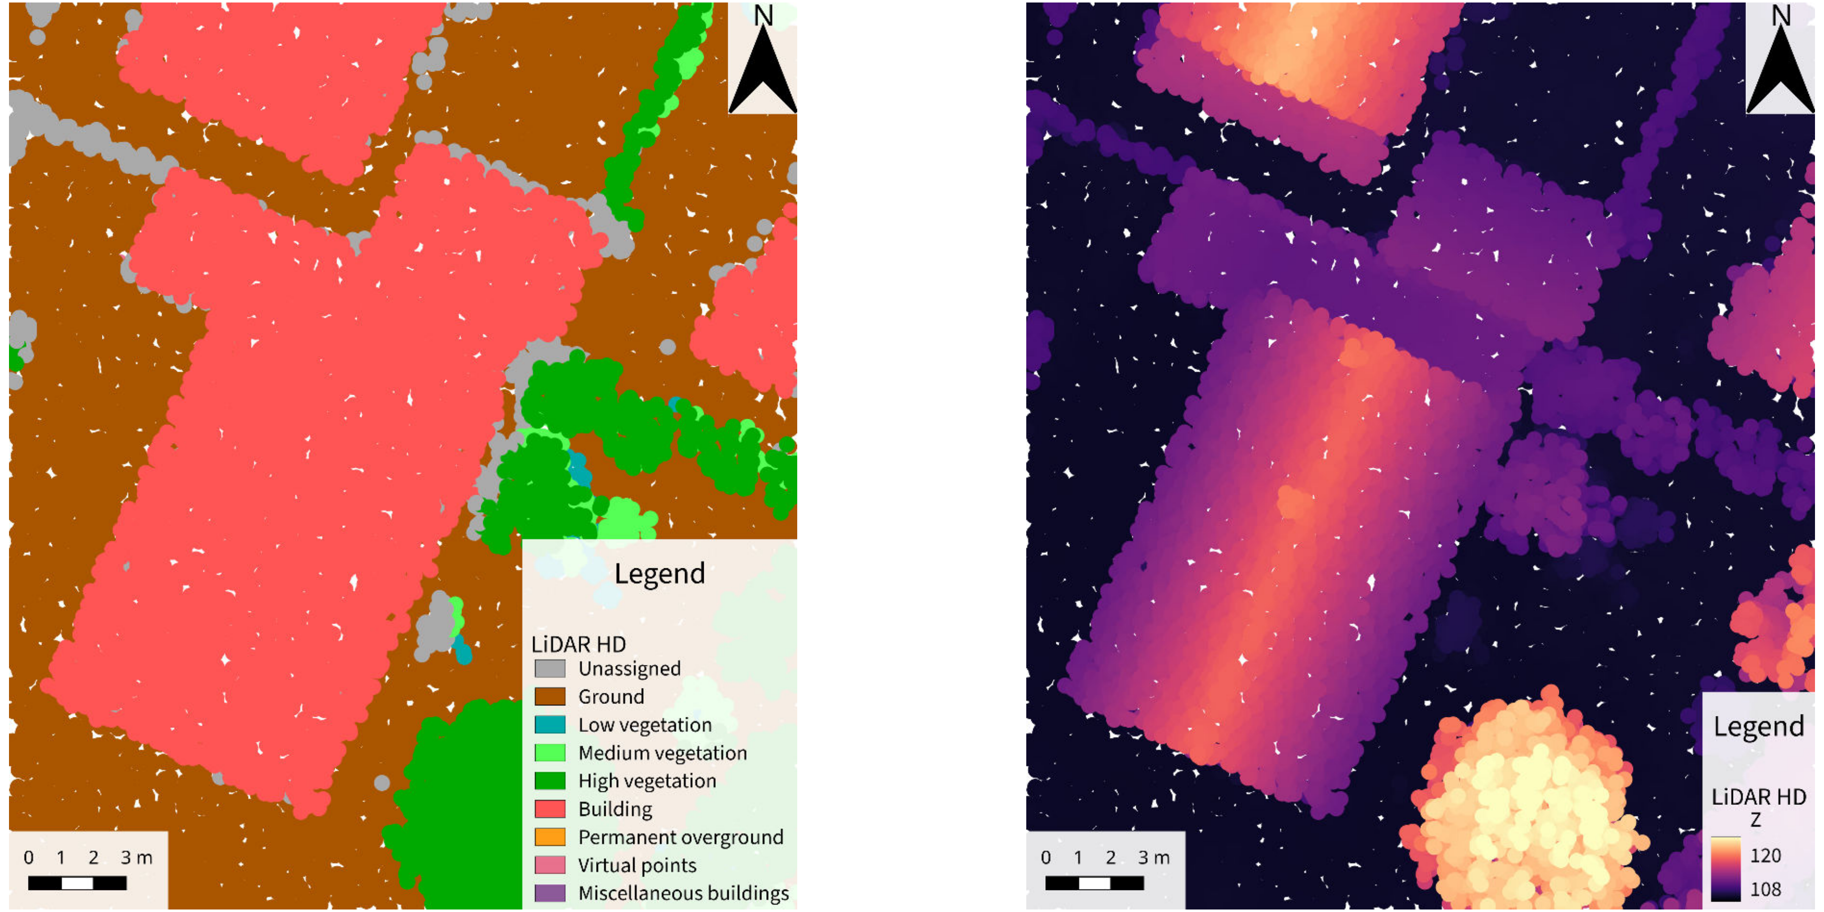

(a) Coloured by classification. (b) Coloured by height.

Figure 25: The LiDAR HD data.

Roofprints
– Illustration with a real building

- We can see **three different building parts** that are connected, one with a pitched roof and the two others with flat roofs.
- The points are well classified overall, but there are a number of points **classified as Unassigned (grey)**, especially at the boundaries between the building and the vegetation.

Illustration with a real building

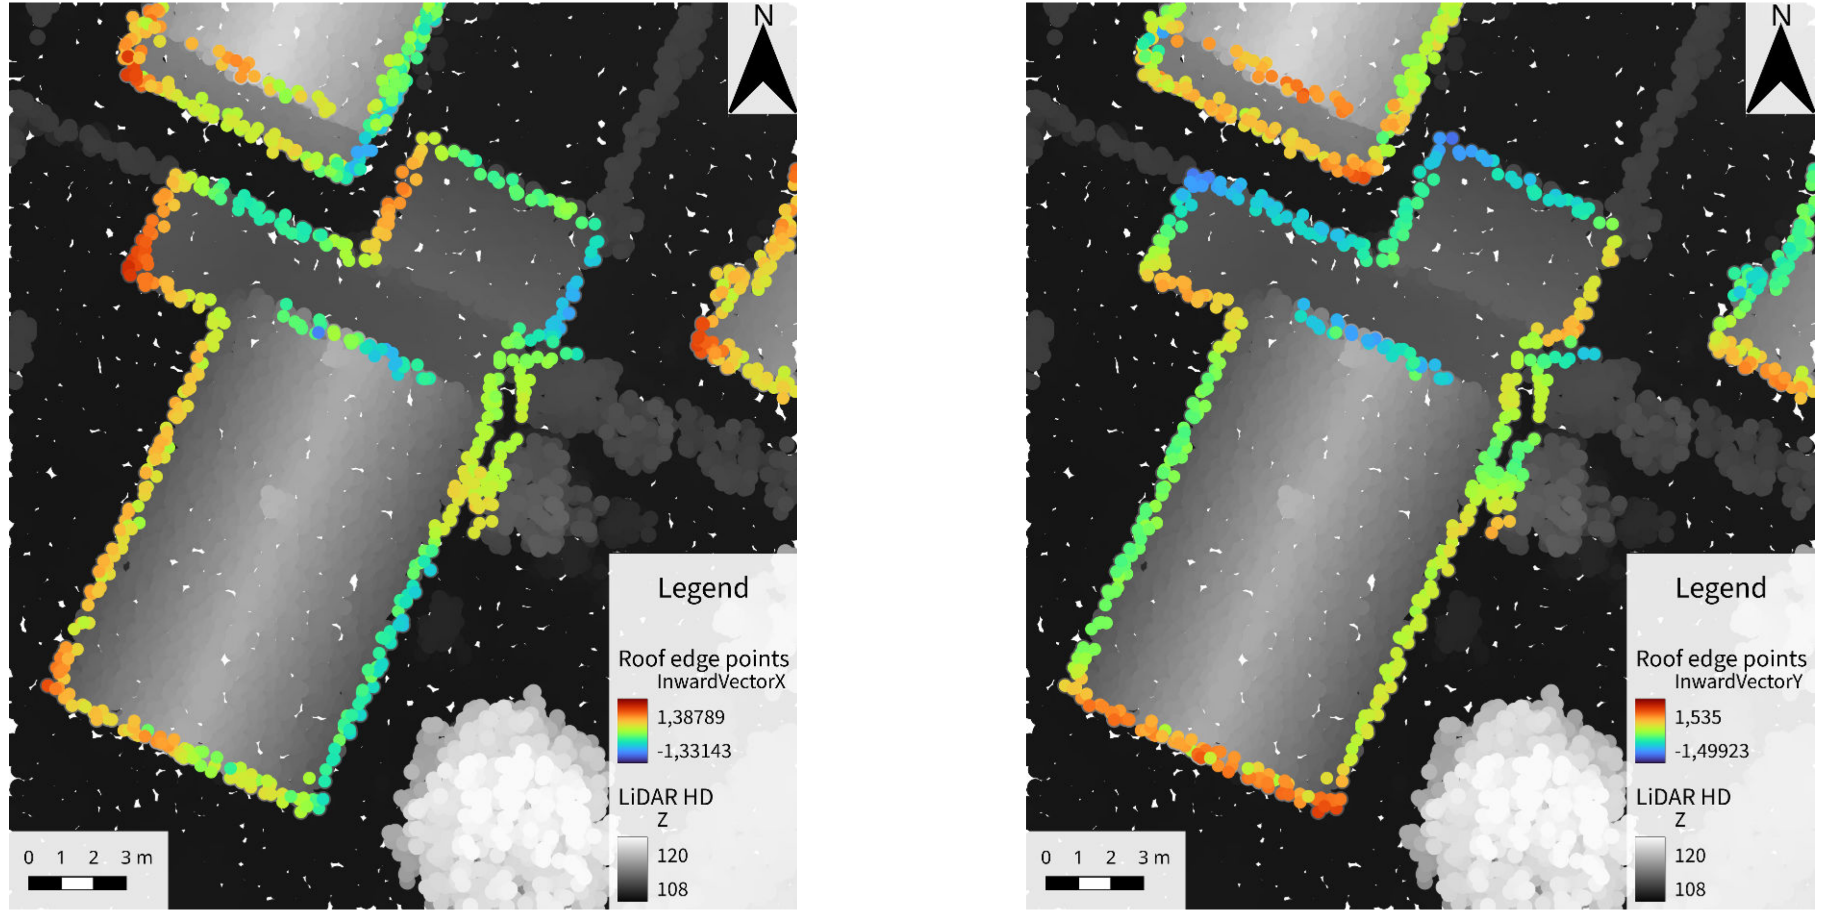

(a) Coloured by the X component. (b) Coloured by the Y component.

Figure 25: The detected roof edge points coloured by their inward vector.

- The detected roof edge points are shown here, coloured by their **inward vector** (the direction in which the building is located from the edge).
- We can see that the inward vectors are **generally well oriented** towards the building, except in areas where the building is **very close to vegetation** at the same height.

Illustration with a real building

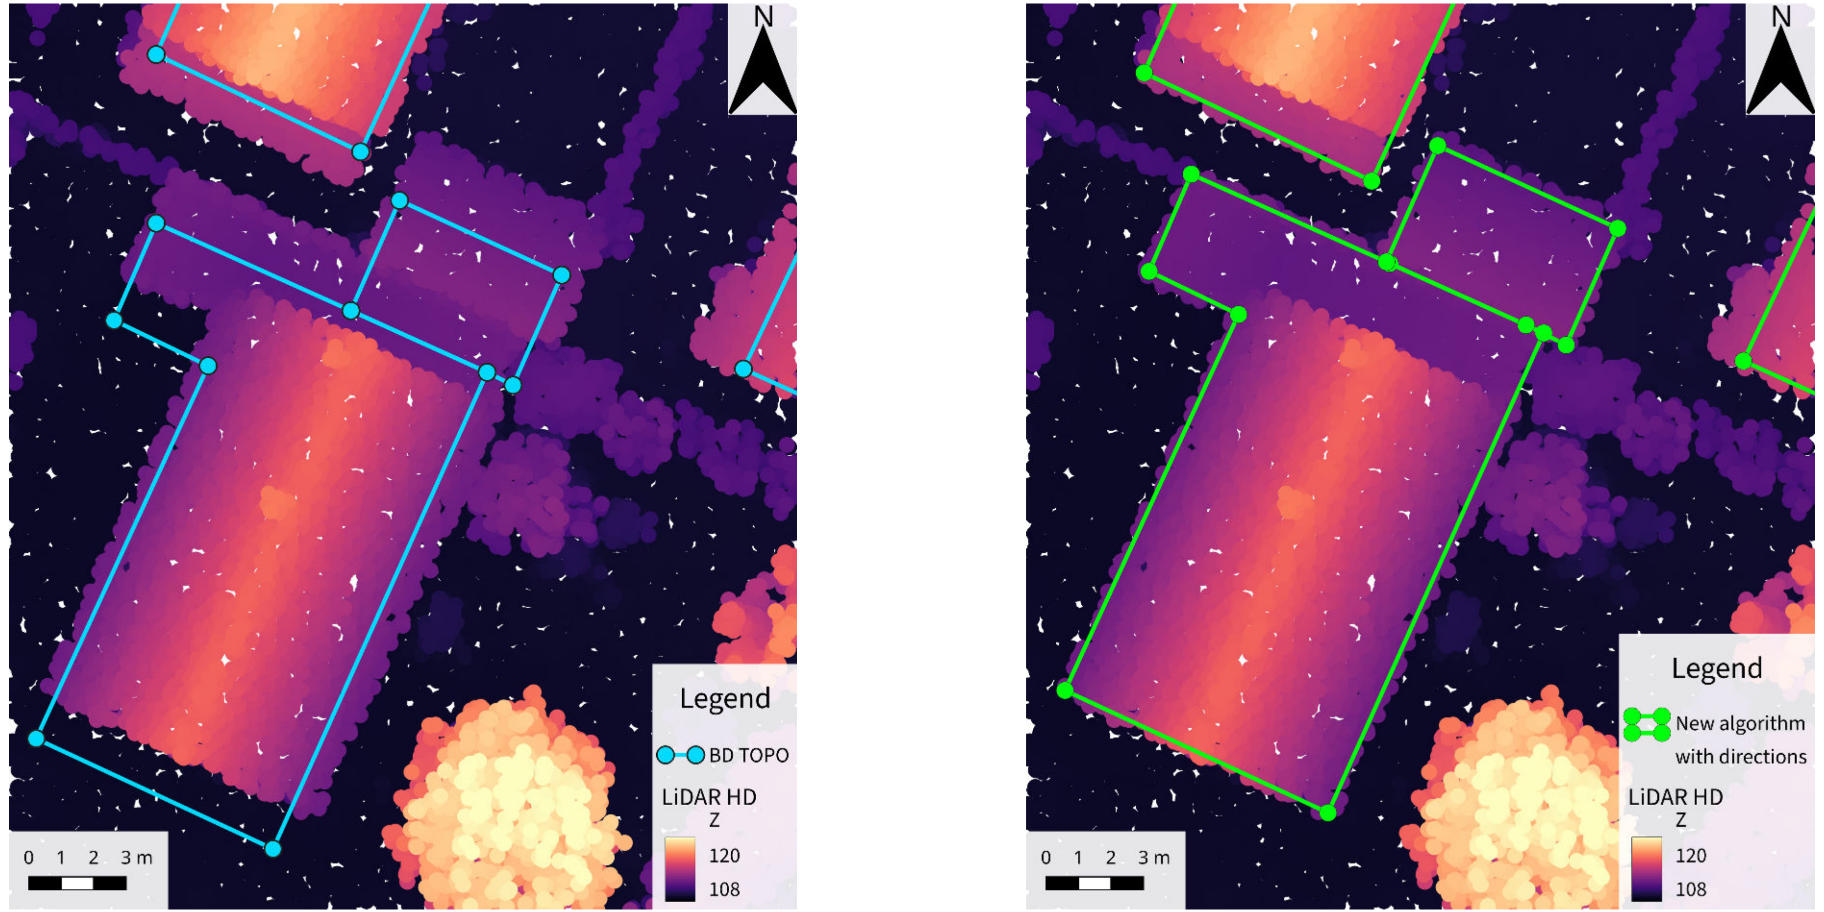

(a) Initial outlines in the BD TOPO.

(b) Final roofprint produced by the algorithm.

Figure 25: Comparison of BD TOPO outlines and roofprints computed with our algorithm.

- Interestingly, this building unit is actually divided into **two parts in the BD TOPO**, even though we could expect 1 or 3 as well.
- The results give a much better alignment, despite:
  - the neighbouring building that is very close
  - the lack of identified points on roof edges between the two polygons
  - the wrong identification of some vegetation points as roof edges

Footprints

Footprints evidence

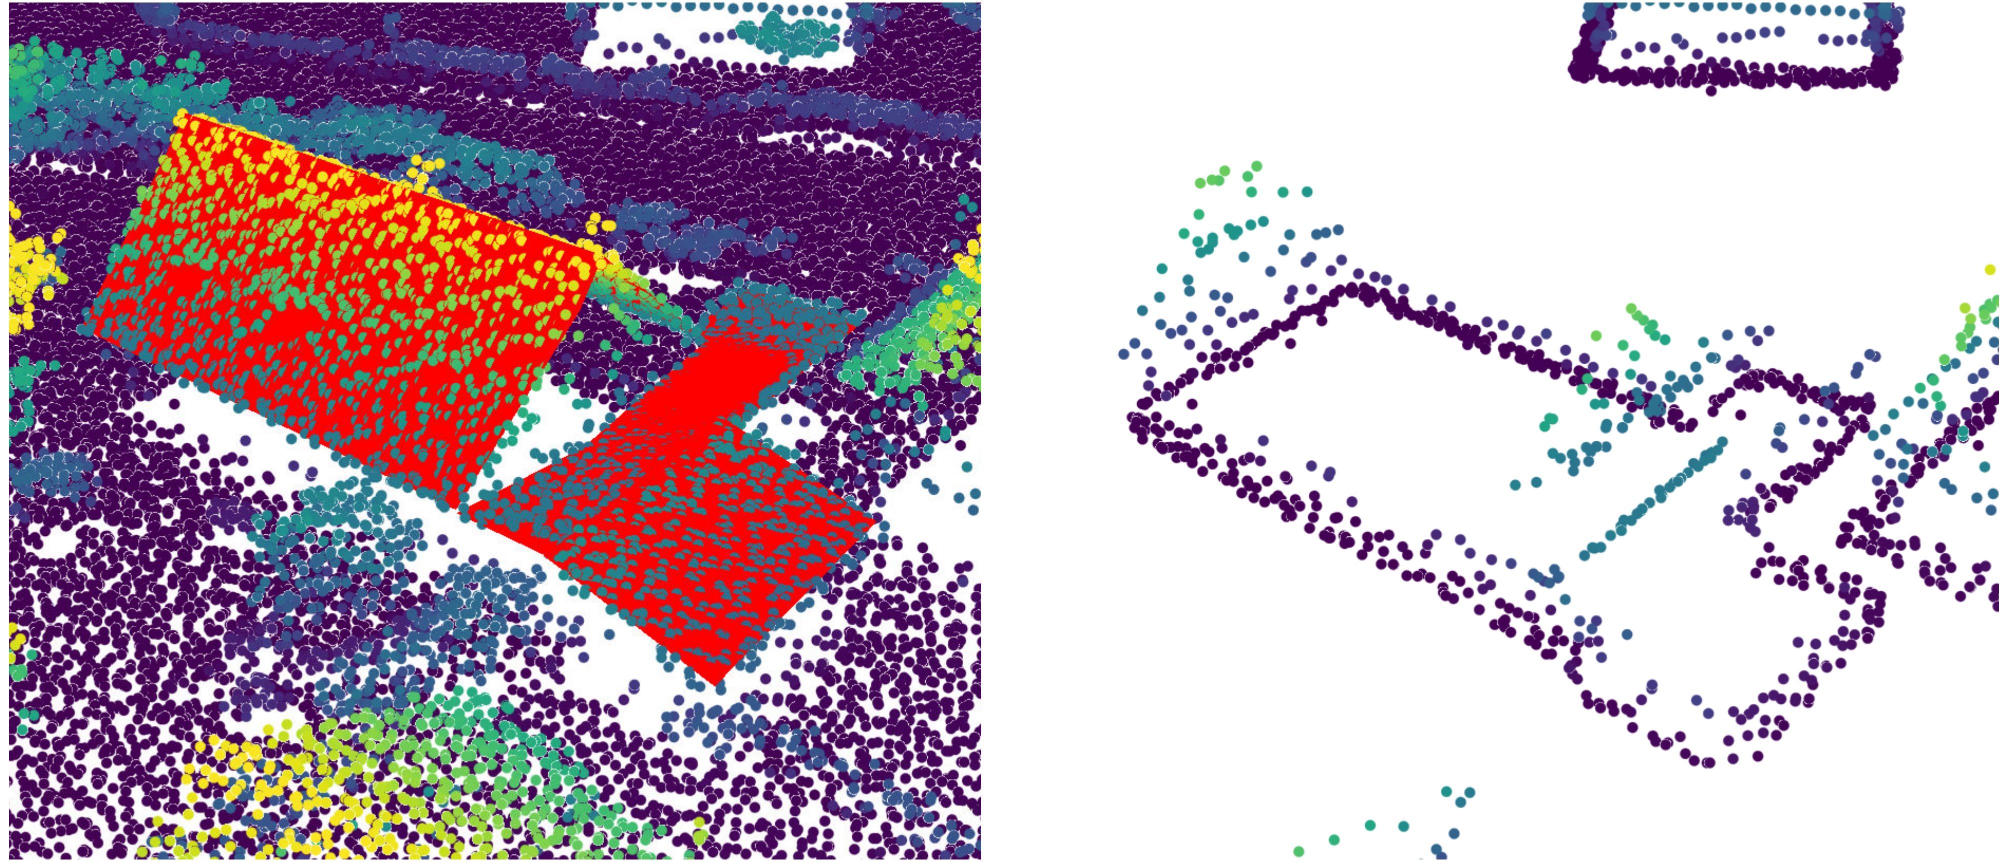

Figure 25: Illustration of the 3D roof structure and the selected points for the footprint computation.

- Process:
  1. Build the roof in 3D using a 3D roof constructor (such as roofer) with the roofprints as input
  2. Select all the points **under** the 3D roof with a small horizontal buffer (for the scoring function) and a small vertical buffer (for roof points slightly below the roof)
- It seems very simple in practice but actually requires to use the complex roof reconstruction algorithms developed in the past few years.

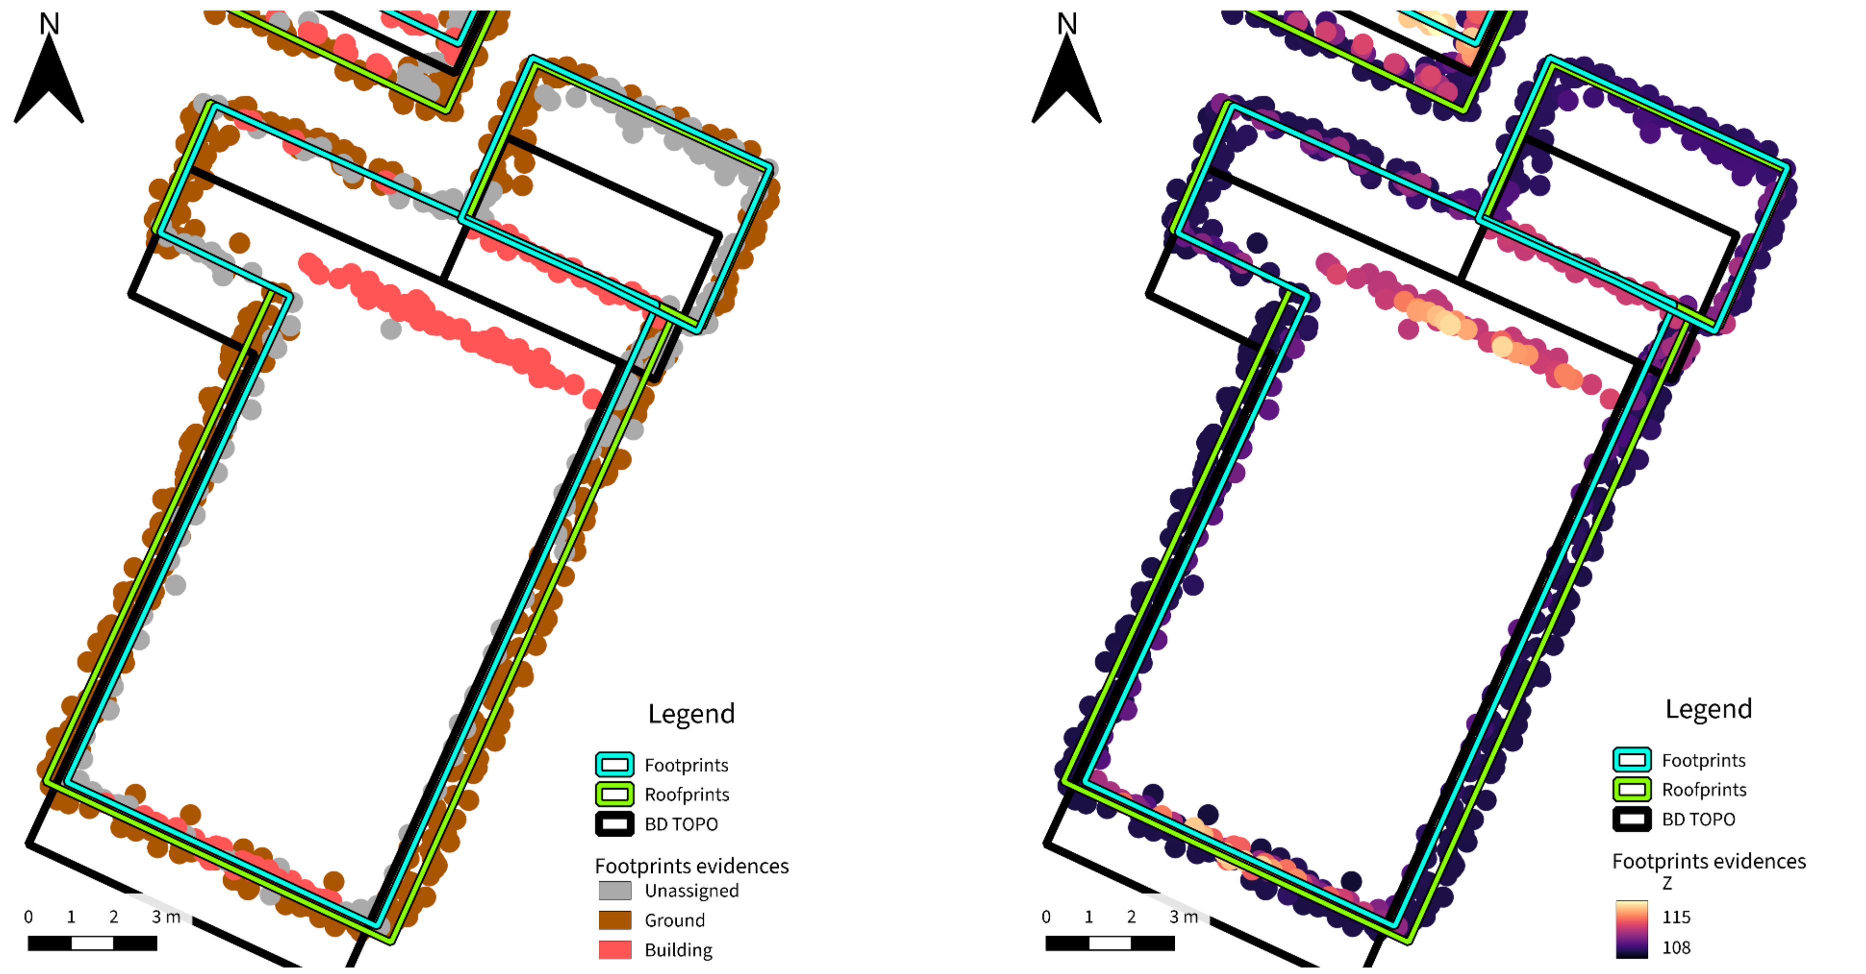

Figure 25: Illustration of the computed footprints.

Example where it works well:

- the footprint is **properly aligned** with the façade points
- we can see **significant roof overhangs for façades oriented west/ east** and no significant roof overhangs for north/south, as expected from the orientation of the slanted roof
- the final footprint is **very close to the initial one from BD TOPO** in terms of shape, with a much better alignment on the point cloud

Illustration with a real building

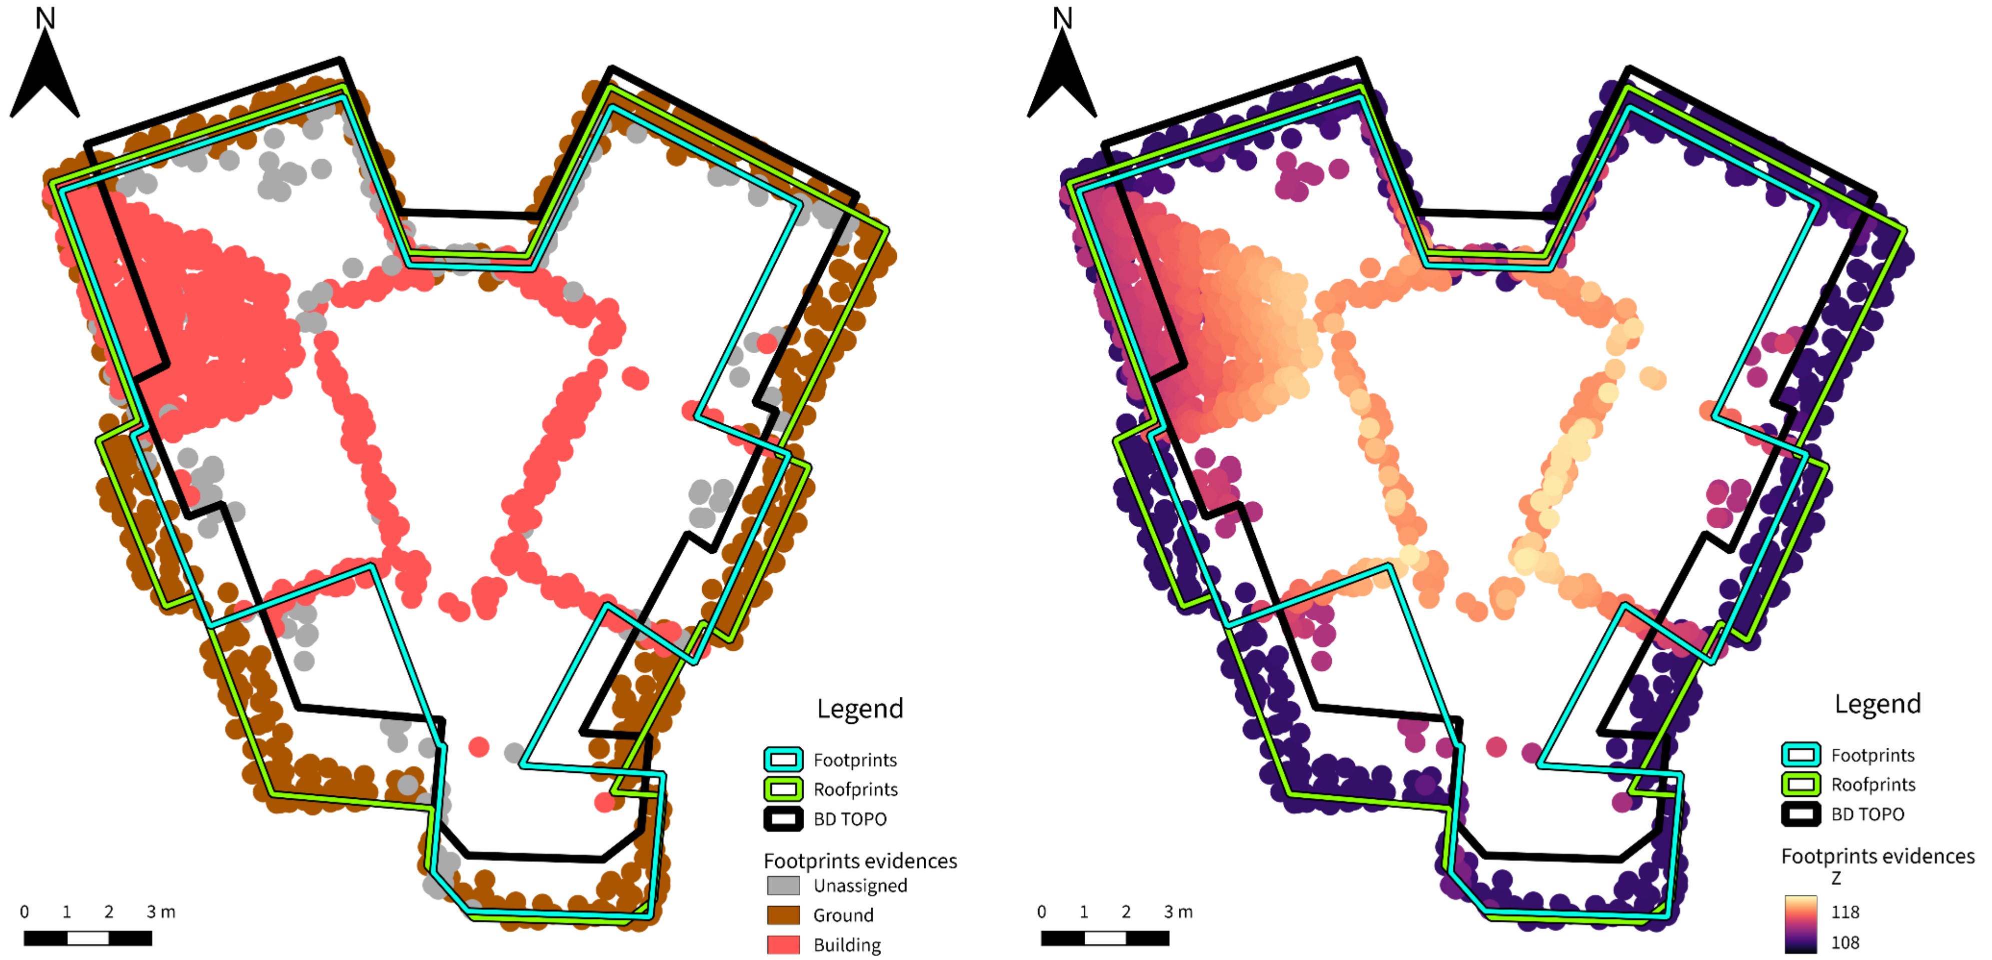

Figure 25: Illustration of the computed footprints.

Footprints
– Illustration with a real building

Example where it doesn't work well:

- many points were selected even though they are actually on the roof
- this resulted in some edges aligning on structures in the roof instead of on the façades

Conclusion

Conclusion

—

Result of the thesis	Benefit
Roofprint and footprint	More precise computations with 2D data
	LoD2.2 → LoD2.3

Conclusion

Result of the thesis	Benefit
Roofprint and footprint	More precise computations with 2D data
	LoD2.2 → LoD2.3
Better georeferencing	Better visualisation in context
	Better 3D roof models

Thank you for your attention!

Link to the slides:

https://alexandre-bry.github.io/MSc_Thesis-Report/slides_pages/2026_06_26-A4.html

alexandre.bry@ign.fr, abry@tudelft.nl

The end

References

Biljecki, F., Ledoux, H., & Stoter, J. (2016). An improved LOD specification for 3D building models. *Computers, Environment and Urban Systems*, *59*, 25–37. https://doi.org/10.1016/j.compenvurbsys.2016.04.005

Wu, T., Vallet, B., Brédif, M., Caraffa, L., & Vosselman, G. (2026). *Reverse engineering LiDAR pulse emission points*.

The end

—

Appendix

Appendix

Validation dataset

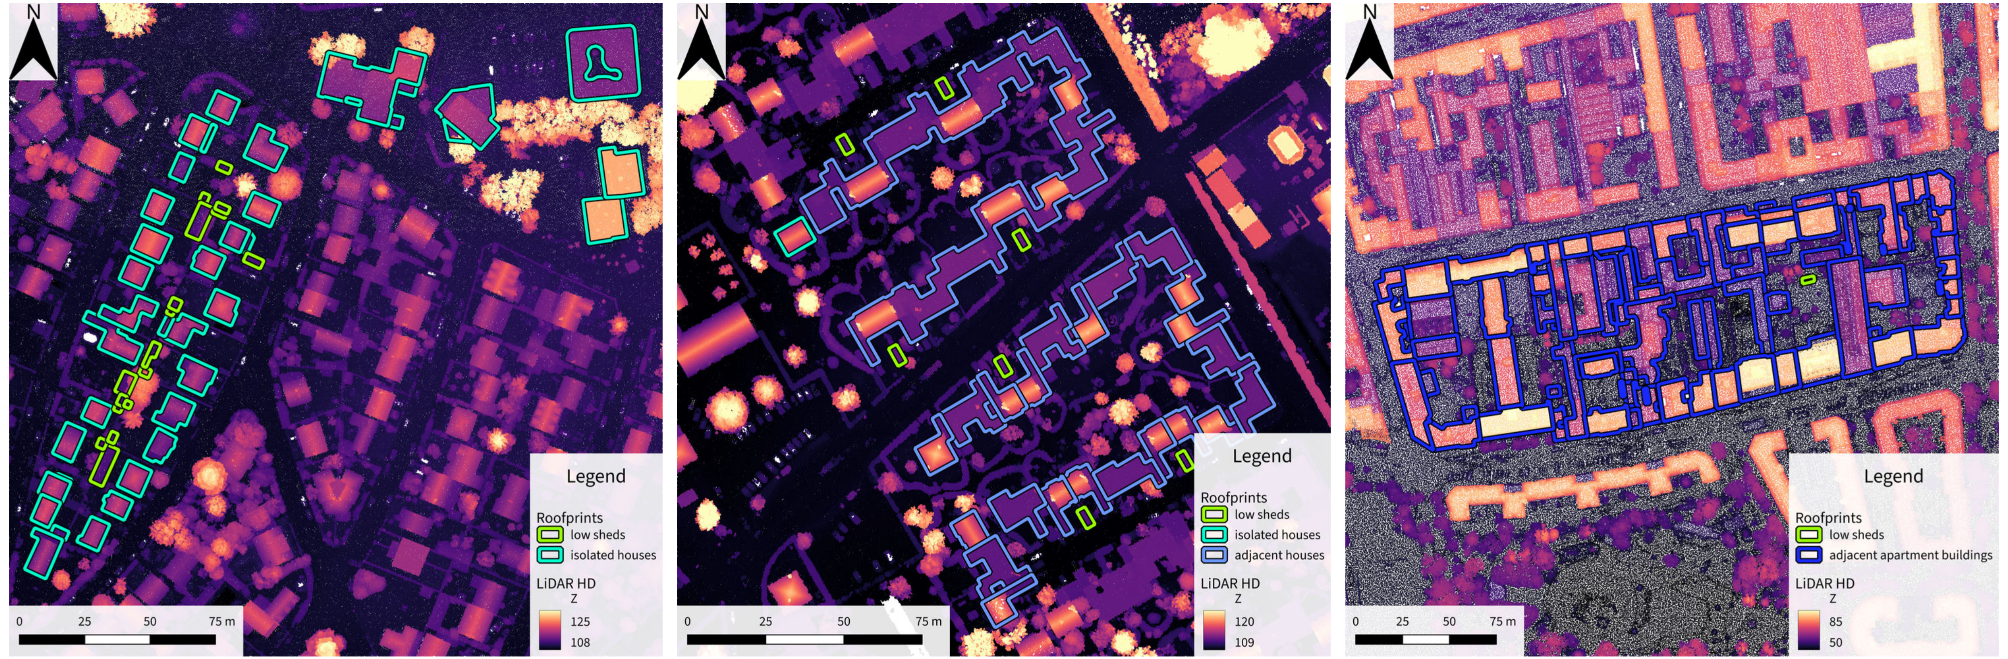

(a) Area in Ozoir-la-Ferrière with isolated houses and small sheds.

(b) Area in Ozoir-la-Ferrière with mainly adjacent houses.

(c) Area in Paris with mainly adjacent blocks of flats.

Figure 25: Validation dataset.

Appendix
– Validation dataset

We made a validation dataset for roofprints with buildings in four different categories (see Figure 17):

- 40 “isolated houses”: medium houses and buildings having up to a few storeys, isolated from the buildings around,
- 25 “adjacent houses”: blocks of adjacent and medium houses and buildings having up to a few storeys,
- 25 “low sheds”: low buildings in height, usually in gardens,
- 80 “adjacent blocks of flats”: blocks of adjacent and high buildings.

Validation results

Dataset	IoU	Chamfer (m)	Centroid (m)
BD TOPO	0.735	0.604	0.939
Rfpt 1 iter	0.842	**0.331**	**0.541**
Rfpt 2 iter	0.851	0.335	0.548
Rfpt 3 iter	**0.853**	0.342	0.566

(a) Whole dataset.

Dataset	IoU	Chamfer (m)	Centroid (m)
BD TOPO	0.776	0.557	0.846
Rfpt 1 iter	0.89	0.243	0.388
Rfpt 2 iter	0.902	0.226	0.362
Rfpt 3 iter	**0.905**	**0.22**	**0.354**

(b) Whole dataset except the “low sheds” category.

Dataset	IoU	Chamfer (m)	Centroid (m)
BD TOPO	0.497	0.876	1.477
Rfpt 1 iter	**0.562**	**0.842**	**1.426**
Rfpt 2 iter	0.557	0.97	1.631
Rfpt 3 iter	0.549	1.045	1.793

(c) Category “low sheds”.

Dataset	IoU	Chamfer (m)	Centroid (m)
BD TOPO	0.744	0.55	0.681
Rfpt 1 iter	0.888	0.165	0.215
Rfpt 2 iter	**0.935**	**0.114**	0.12
Rfpt 3 iter	**0.935**	**0.114**	**0.117**

(d) Category “isolated houses”.

Dataset	IoU	Chamfer (m)	Centroid (m)
BD TOPO	0.797	0.448	0.638
Rfpt 1 iter	0.901	0.196	0.309
Rfpt 2 iter	0.91	0.175	0.269
Rfpt 3 iter	**0.914**	**0.168**	**0.259**

(e) Category “adjacent houses”.

Dataset	IoU	Chamfer (m)	Centroid (m)
BD TOPO	0.785	0.595	0.994
Rfpt 1 iter	**0.888**	0.297	**0.499**
Rfpt 2 iter	0.882	0.298	0.512
Rfpt 3 iter	0.887	**0.289**	0.503

(f) Category “adjacent blocks of flats”.

Figure 25: Average metrics over different subsets of the validation dataset.

Appendix
– Validation results

The different metrics on the validation dataset show that the do improve the outlines by a lot in all categories except “low sheds” where the results are bad for a few buildings and from correct to very good for the rest.

Results

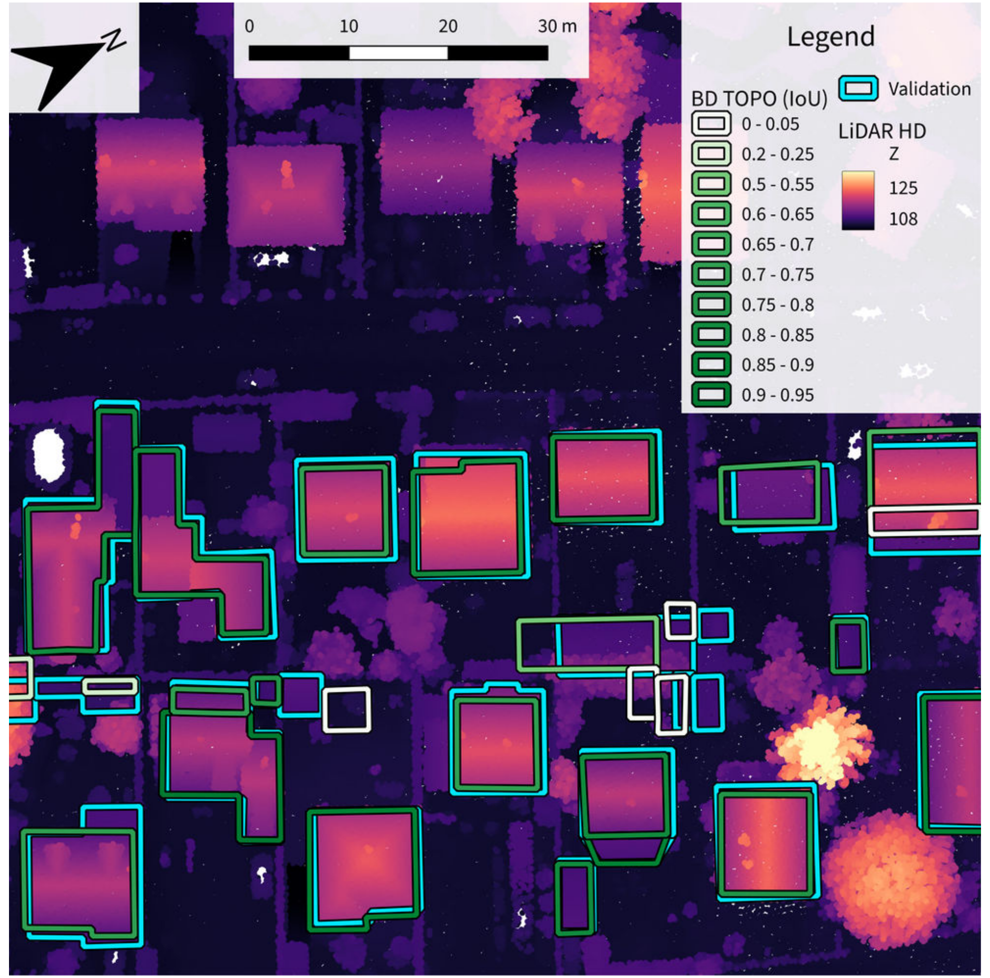

Figure 19: BD TOPO.

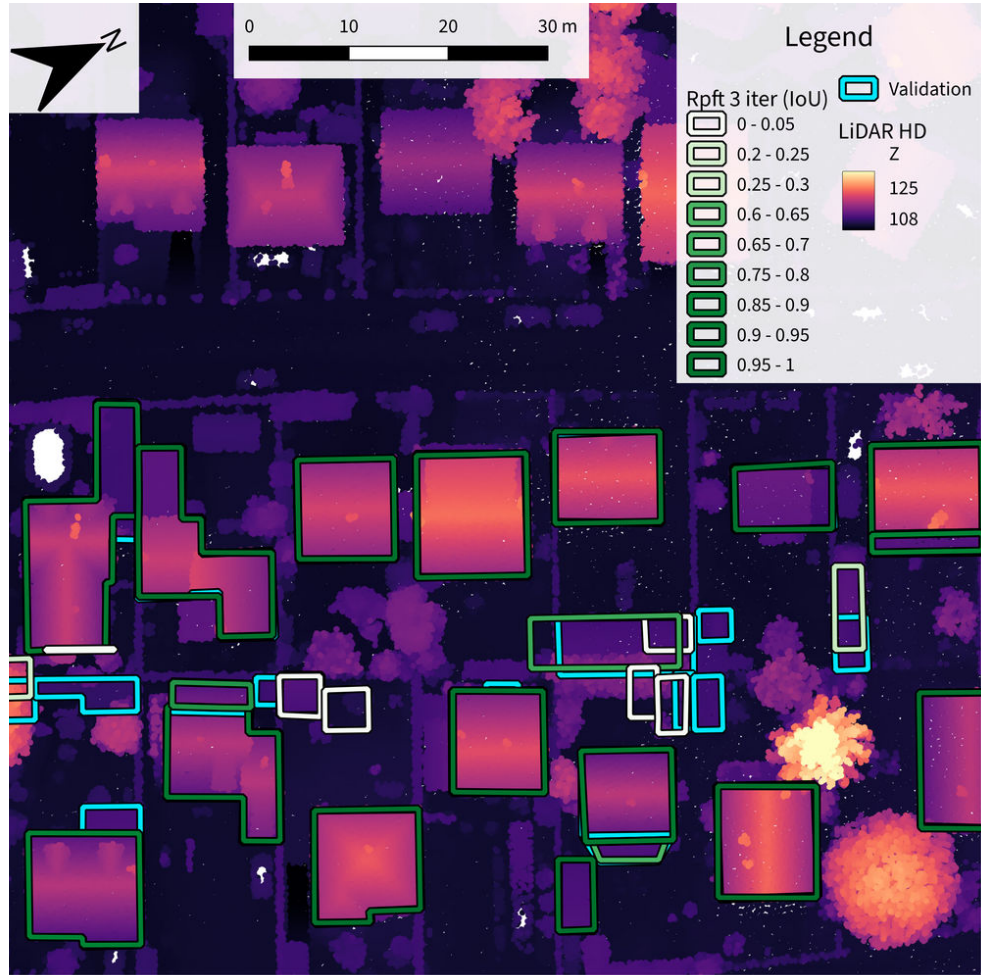

Figure 20: Roofprints after 3 iterations.

Results

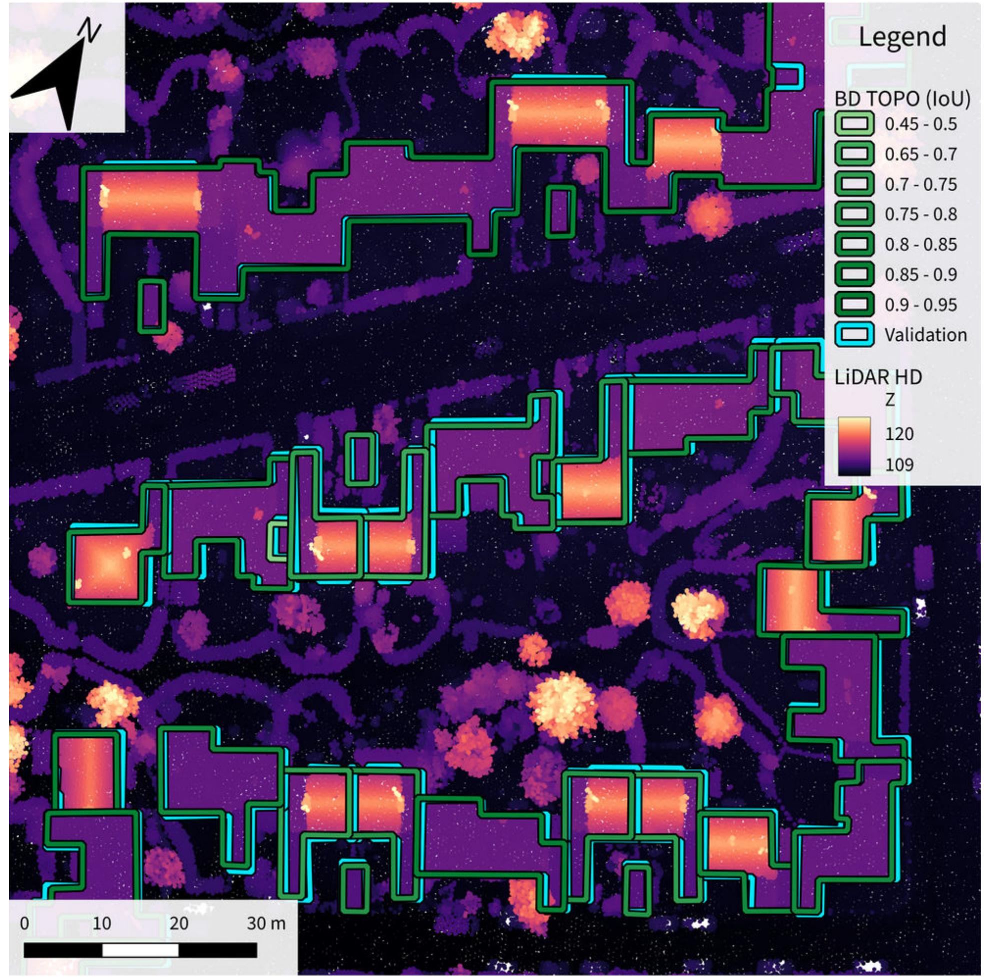

Figure 21: BD TOPO.

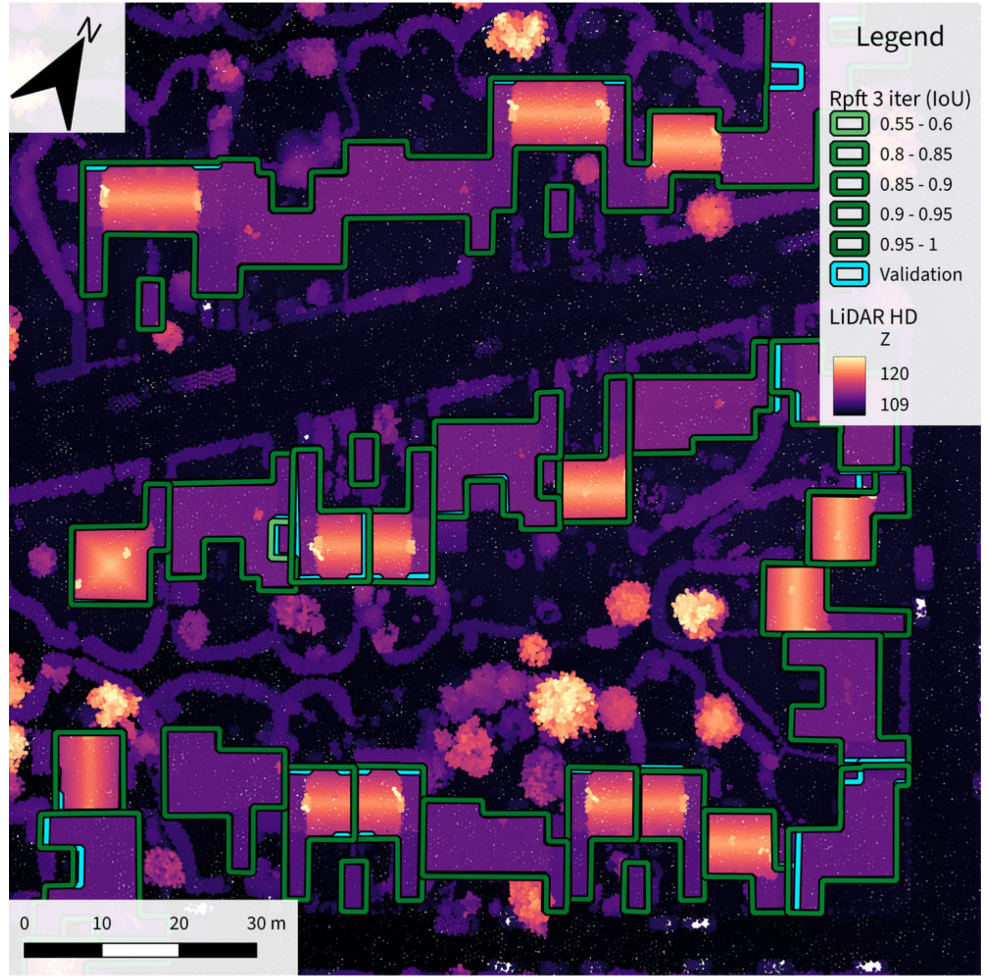

Figure 22: Roofprints after 3 iterations.

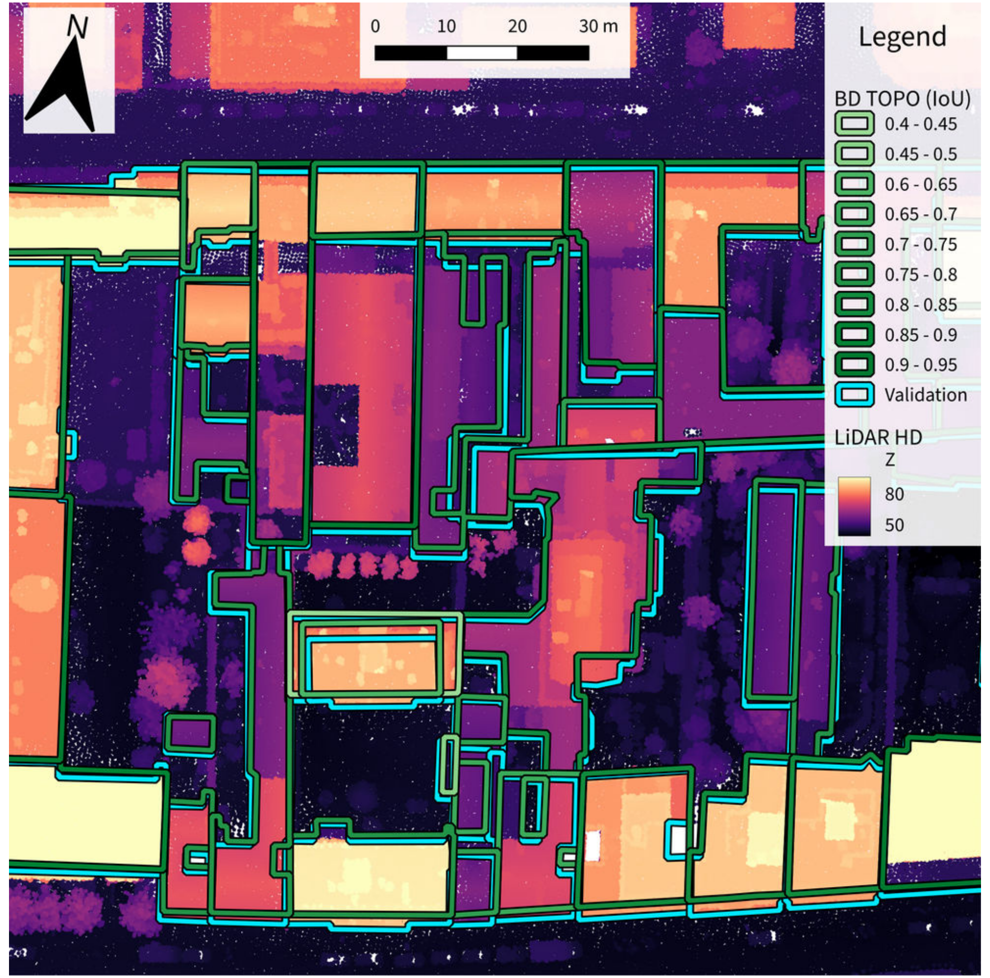

Figure 23: BD TOPO.

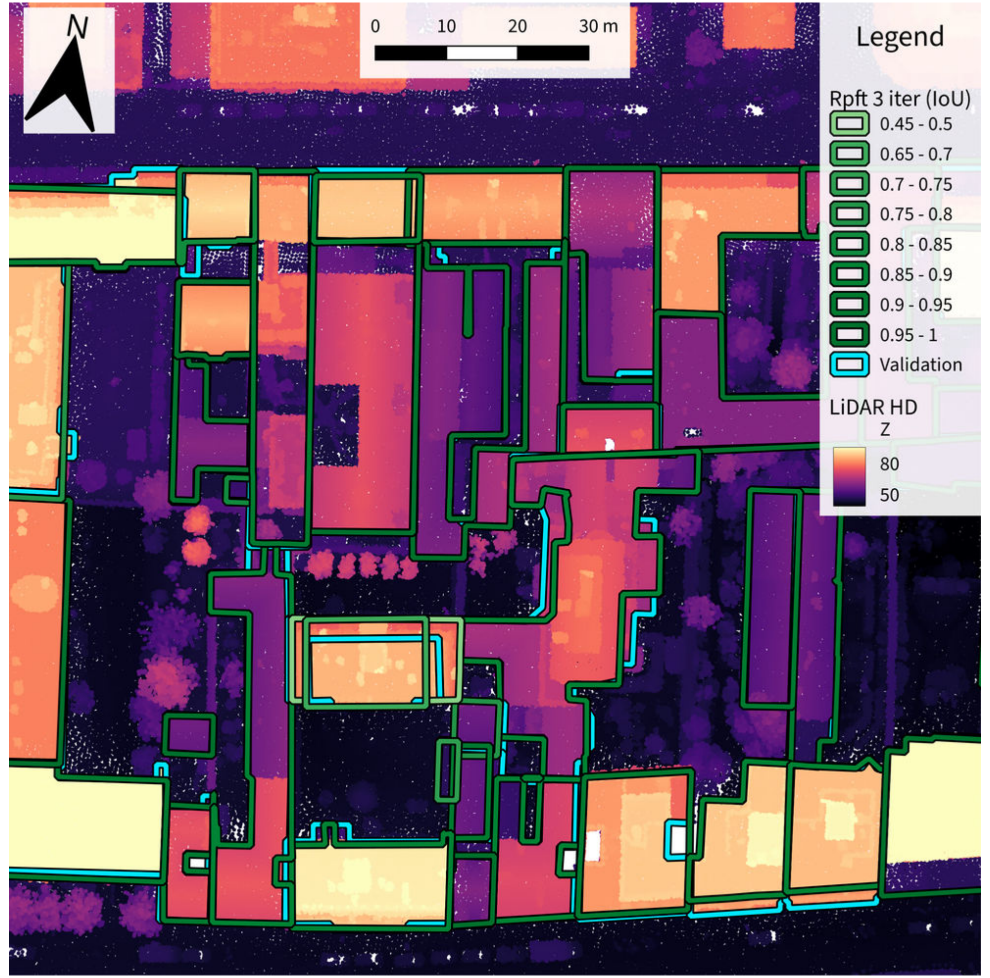

Figure 24: Roofprints after 3 iterations.

Overview of the methodology

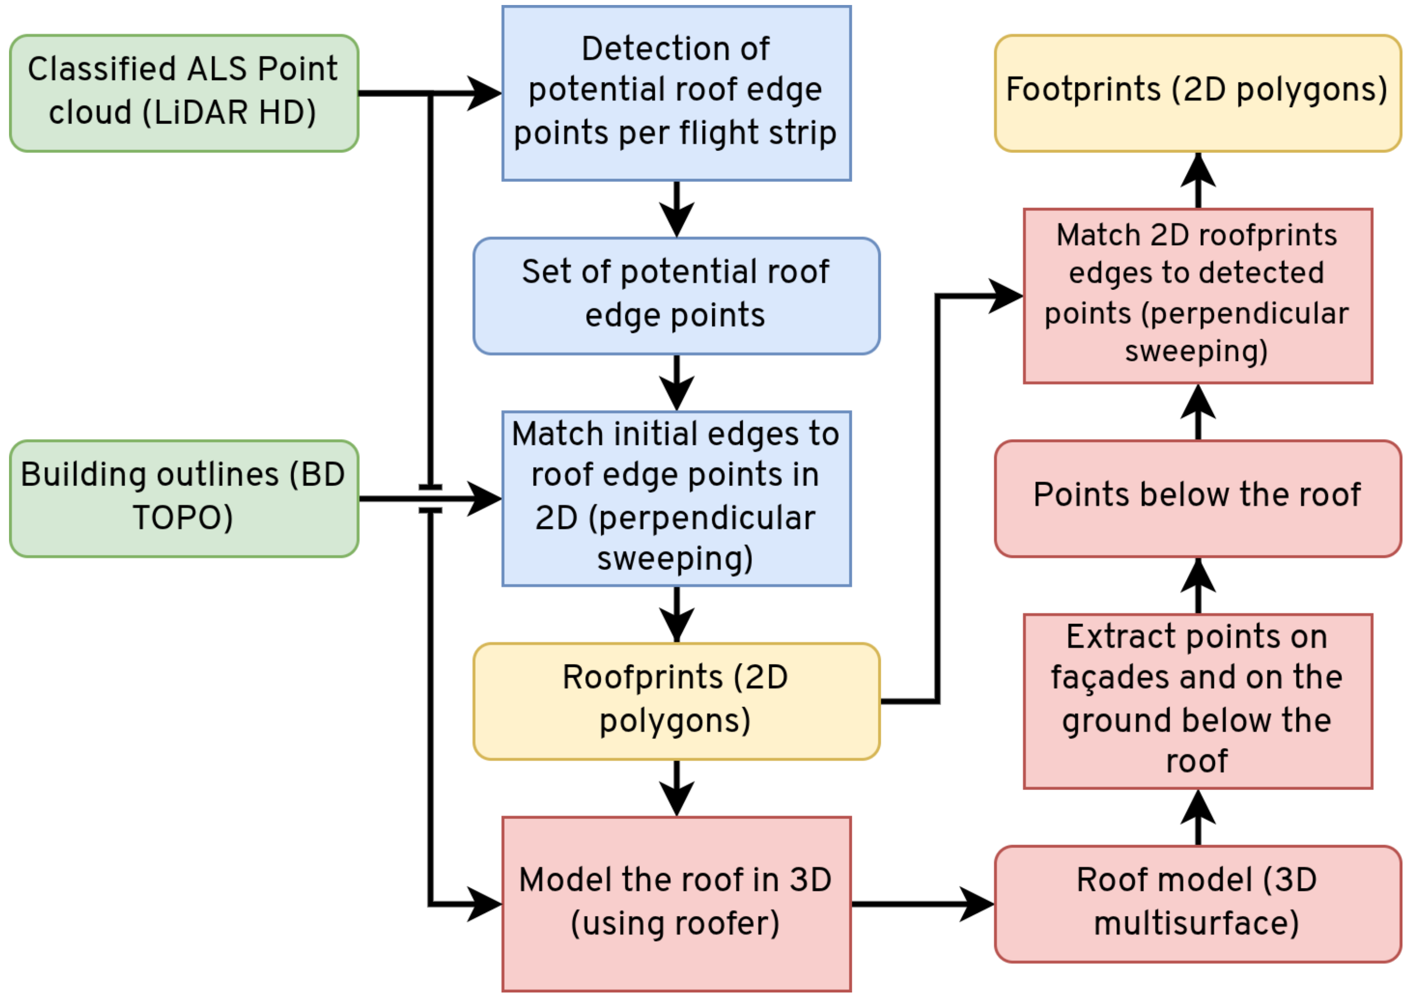

Figure 25: Overview of the pipeline.

- We start with the roofprint because aerial LiDAR data gives a much higher density of points on the roof than on the façades, making the roof much easier to identify.
- Then, once the roofprint is accurately registered on the point cloud, we can use it to generate an accurate 3D roof model that allows to select all the points below the roof (hopefully façade and ground points), which are all the points that provide information about the façades and the roof overhangs.