DIMENSIONALITY BASED SCALE SELECTION IN 3D LIDAR POINT CLOUDS

Jérôme Demantké, Clément Mallet, Nicolas David and Bruno Vallet

Université Paris Est, Laboratoire MATIS, IGN
73, avenue de Paris 94165 Saint-Mandé, Cedex – firstname.lastname@ign.fr

Commission III - WG III/2

KEY WORDS: point cloud, adaptive neighborhood, scale selection, multi-scale analysis, feature, PCA, eigenvalues, dimensionality

ABSTRACT:

This papers presents a multi-scale method that computes robust geometric features on lidar point clouds in order to retrieve the optimal neighborhood size for each point. Three dimensionality features are calculated on spherical neighborhoods at various radius sizes. Based on combinations of the eigenvalues of the local structure tensor, they describe the shape of the neighborhood, indicating whether the local geometry is more linear (1D), planar (2D) or volumetric (3D). Two *radius-selection* criteria have been tested and compared for finding automatically the optimal neighborhood radius for each point. Besides, such procedure allows a dimensionality labelling, giving significant hints for classification and segmentation purposes. The method is successfully applied to 3D point clouds from airborne, terrestrial, and mobile mapping systems since no a priori knowledge on the distribution of the 3D points is required. Extracted dimensionality features and labellings are then favorably compared to those computed from constant size neighborhoods.

1 INTRODUCTION

Point cloud data from airborne and terrestrial devices provide a direct geometrical description of the 3D space. Such information is reliable, of high accuracy but irregular and not dense. However, the underlying structures and objects may be detected among sets of close 3D points. The local geometry is estimated by the distribution of points in the neighborhood. Finding the best neighborhood for each point is a main issue for a large variety of common processes: data downsampling, template fitting, feature detection and computation, interpolation, registration, segmentation, or modelling purposes. The notion of neighborhood and its fundamental properties are fully described in (Filin and Pfeifer, 2005).

The neighbors of a lidar point are traditionally retrieved by finding the k nearest neighbors or all the points included in a small restricted environment (sphere or cylinder) centered on the point of interest. The main problem stems from the fact that the k and environment radius values are (1) usually heuristically chosen, and (2) assumed to be constant for the whole point cloud, instead of being guided by the data. This does not ensure that all theses neighbors belong to the same object as the current point. Therefore, its local description may be biased when including several distinct structures, and provides erroneous feature descriptors. Moreover, the relative variation in the spatial extent of geometrical structures is ignored. For aerial datasets, problems will occur at the borders between objects and for objects which size is inferior or close to the neighborhood size. For terrestrial datasets, in addition, the point density may significantly fluctuate due to foreground object occlusion, dependence on distance and relative orientation of the objects (Soudarissanane et al., 2009), leading to data sparseness or irregular sampling.

This paper aims at proposing a methodology to find the optimal neighborhood radius for each 3D point on a lidar point cloud. An "optimal" neighborhood is defined as the largest set of spatially close points that belong to the same object as the point of interest. The inclusion of points lying on different surfaces is prohibited. The context of the study is rather general: in order to be applicable both on terrestrial (TLS) and airborne (ALS) datasets, the method is simply based on the point location, without requiring knowledge on intensity, echo number or full waveforms. The

topology resulting from the sequential acquisition of the data is also considered to be lost ("unorganized" point cloud), preventing the adoption of specific scan line grouping methods (Hadjiliadis and Stamos, 2010). Furthermore, the process is designed out of the scope of any application, even if the final goal is indeed to be beneficial to any application requiring a correct local description around each 3D point.

The problem of scale selection has been mainly tackled for surface reconstruction and feature extraction of scanned opaque objects. Several approaches have therefore been developed for noisy point clouds and irregular sampling issues. Most of them are surface-based *i.e.*, they try to fit a curve or a surface of some form to the 3D point cloud (Pauly et al., 2006). Finding the optimal group that well represents the local geometrical properties is performed using indicators such as the normal and/or the curvature (Hoppe et al., 1992; Zwickler et al., 2002; Dey and Sun, 2005; Belton and Lichti, 2006). For instance, it is retrieved by minimizing the upper bound on angular error between the true normal and the estimated one. Starting for the minimal possible subset around the point of interest, the neighborhood is iteratively increased until the angular variance reaches a predefined threshold (Mitra et al., 2004). Such works have been theoretically improved by Lalonde et al. (2005), no more requiring knowledge on the data distribution, and applied to mobile mapping datasets. An alternative work, based of the expression of the positional uncertainty, is introduced in (Bae et al., 2009) for processing TLS datasets. However, these methods are effective for smoothly varying surfaces and may not be adapted to anthropic surfaces acquired with various kinds of lidar systems. Furthermore, the model-based assumption does not hold when dealing with objects without predefined shapes (*e.g.*, vegetated areas) or with noise stemming for relief high frequencies (*e.g.*, chimneys and facades for ALS data or pedestrians and points inside buildings for TLS data). Consequently, in our context, a more suitable solution is to directly compute shape features (Gumhold et al., 2001; Belton and Lichti, 2006), *i.e.*, low-level primitives that may capture the variability of natural environments.

The proposed methodology is developed in Section 2. The shape features are described in Section 3. The computation of the optimal neighborhood radius embedded in a multi-scale framework is proposed in Section 4. Results on various laser scanning datasets

are presented in Section 5, and conclusions are drawn in Section 6.

2 PROPOSED METHOD

2.1 Description

The methodology aims at finding the optimal neighborhood radius for each lidar point, working directly and exclusively in the 3D domain, without relying on surface descriptors (such as normals) or structures (such as triangulations or polygonal meshes). It is composed on two main steps:

1. Computation of three dimensionality features for each point, between predefined minimal and maximal neighborhood scale. These features describe the distribution of the points in the 3D space, and more exactly, the matching between the local point cloud and each of the three dimensionalities (linear, planar or volumetric).

2. Scale selection: retrieval of the neighborhood radius for which one dimensionality is most dominant over the two others.

The three dimensionality features (a_{1D}-a_{2D}-a_{3D}) are computed exhaustively, at each point and for each acceptable neighborhood scale, from the local covariance matrix. An isotropic spherical neighborhood, centered on the point of interest, is adopted for this purpose. These low-level features, as well as the automatic set up of the radius lower and upper bounds are described in Section 3.

Then, the optimal radius is retrieved by comparing the behaviours of these three features between the minimal and maximal acceptable radius. Two *radius-selection* criteria are tested to evaluate each scale and find the most relevant value. This multi-scale analysis is performed in order to capture variation in shape when aggregating points for an object distinct from the object of interest (edge effect or outliers). It is also useful in case of significant density variation and lack of support data for gathering points over a large volume while ensuring the conservation of the favorite dimensionality.

Finally, our method presents three interesting characteristics:

- Definition of a confidence index of the saliency of one dimensionality over the two other ones.

- Multi-scale analysis and automatic set up of the bounding scales.

- Labelling of each point according to its privileged dimensionality, providing an interesting basis for segmentation and classification algorithms.

2.2 Datasets

In order to assess the relevance of the proposed approach for various point densities, point distributions and points of view, three kinds of lidar datasets are tested: airborne, terrestrial static, and acquired with a mobile mapping system (named ALS, TLS, and MMS, respectively).

ALS : Three datasets are used. The first one (ALS_G) has been acquired over Biberach (Germany), covering both residential and industrial areas as well as a city center with small buildings (point density of 5 pts/m^2). The second dataset (ALS_R) concerns a residential area in Russia, with 5 pts/m^2 (Shapovalov et al., 2010). Finally, the third one covers the dense city center of Marseille (France), with high buildings, and thus sparse points on the building facades (ALS_F). Three parallel strips are present: the point density therefore varies between 2 and 4 pts/m^2 (for one strip and for the overlapping areas, respectively).

TLS : The terrestrial scans acquired over the Agia Sanmarina church (Greece) have been processed (Bae et al., 2009). The dataset first offers a large variety of structures of various sizes as well as sparse vegetation on the ground. Furthermore, the point density significantly varies with the orientation of the surfaces with respect to the scanner position.

MMS : Datasets over two urban areas (France and United States, respectively MMS_F and MMS_U) from distinct mobile mapping systems have been processed (Munoz et al., 2009). Such datasets also include man-made objects of various sizes and shapes, with varying point densities. The two specificities of MMS datasets are (1) gaps in the point cloud due to the occlusion of foreground objects, and (2) vertical privileged directions in the point clouds due to the sequential acquisition by lines.

3 SHAPE FEATURES

In this work, we focus on shape features computed on the neighborhood of a lidar point. The neighborhood $\mathcal{V}_P^r$ of a point P at scale r is defined as the set of points P_k verifying :

$$P_k \in \mathcal{V}_P^r \Leftrightarrow \|P - P_k\| \leq r. \tag{1}$$

Thus, the neighborhood environments are 3D spherical volumes. They ensure isotropy and rotation invariance, such that the computed shape descriptors are not biased by the shape of the neighborhood. Furthermore, r is the single parameter to be optimized. A classical approach consists in performing a Principal Component Analysis (PCA) of the 3D coordinates of $\mathcal{V}_P^r$ (Belton and Lichti, 2006; Gross et al., 2007). This statistical analysis uses the first and second moments of $\mathcal{V}_P^r$, and results in three orthogonal vectors centered on the centroid of the neighborhood. The PCA synthesizes the distribution of points along the three dimensions (Tang et al., 2004), and thus models the principal directions and magnitudes of variation of the point distribution around the center of gravity. These magnitudes are combined to provide shape descriptors for each of the three dimensions. More advanced features based on harmonics or spin images (Frome et al., 2004; Golovinskiy et al., 2009) are not necessary since the segmentation task, which indeed requires contextual knowledge, is not tackled in this paper.

3.1 Principal Component Analysis

Let $\mathbf{x}_i = \begin{pmatrix} x_i & y_i & z_i \end{pmatrix}^T$ and $\bar{\mathbf{x}} = \frac{1}{n} \sum_{i=1,n} \mathbf{x}_i$ the center of gravity of the n lidar points of $\mathcal{V}_P^r$.

Given $\mathbf{M} = \begin{pmatrix} \mathbf{x}_1 - \bar{\mathbf{x}} & ... & \mathbf{x}_n - \bar{\mathbf{x}} \end{pmatrix}^T$, the 3D structure tensor is defined by $\mathbf{C} = \frac{1}{n} \mathbf{M}^T \mathbf{M}$. Since $\mathbf{C}$ is a symmetric positive definite matrix, an eigenvalue decomposition exists and can be expressed as $\mathbf{C} = \mathbf{R} \mathbf{\Lambda} \mathbf{R}^T$, where $\mathbf{R}$ is a rotation matrix, and $\mathbf{\Lambda}$ a diagonal, positive definite matrix, known as *eigenvector* and *eigenvalue* matrices, respectively. The eigenvalues are positive and ordered so that $\lambda_1 \geq \lambda_2 \geq \lambda_3 > 0$. $\forall j \in [1,3]$, $\sigma_j = \sqrt{\lambda_j}$, denotes the standard deviation along the corresponding eigenvector $\vec{v_j}$. Thus, the PCA allows to retrieve the three principal directions of $\mathcal{V}_P^r$, and the eigenvalues provide their magnitude. The average distance, all around the center of gravity, can also be modeled by a surface. The shape of $\mathcal{V}_P^r$ is then represented by an oriented ellipsoid. The orientation and the size informations are divided between $\mathbf{R}$ and $\mathbf{\Lambda}$: $\mathbf{R}$ turns the canonical basis into the orthonormal basis $(\vec{v_1}, \vec{v_2}, \vec{v_3})$ and $\sqrt{\mathbf{\Lambda}}$ transforms the unit sphere to an ellipsoid (σ_1, σ_2 and σ_3 being the lengths of the semi-axes). As enhanced in Figure 1 and for instance in (Gross et al., 2007), such an ellipsoid reveals the linear, planar or volumetric behaviour of the neighborhood *i.e.*, whether the point

set is spread in one, two or three dimensions (blue, gray, and green ellipsoids in Figure 1, respectively).
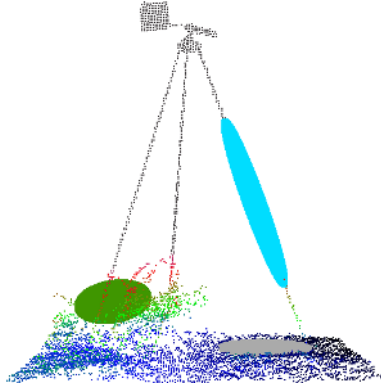

Figure 1: Three examples of ellipsoids computed over three areas of interest of distinct dimensionalities for the *TLS* dataset (a tripod over a low and sparse vegetation – height colored).

3.2 Dimensionality features and labelling

Various geometrical features can be derived from the eigenvalues. Several indicators have already been proposed (West et al., 2004; Toshev et al., 2010), and the following ones have been selected (Figure 2) to describe the linear (a_{1D}), planar (a_{2D}), and scatter (a_{3D}) behaviors within $\mathcal{V}_P^r$:

$$a_{1D} = \frac{\sigma_1 - \sigma_2}{\mu}, \quad a_{2D} = \frac{\sigma_2 - \sigma_3}{\mu}, \quad a_{3D} = \frac{\sigma_3}{\mu},$$

where μ is the normalization coefficient. Both choices $\mu = \sigma_1$ and $\mu = \sum_{d=1,3} \sigma_d$ are conceivable and imply $a_{1D}, a_{2D}, a_{3D} \in [0,1]$. The dimensionality labelling (1D, 2D or 3D) of $\mathcal{V}_P^r$ is defined by:

$$d^*(\mathcal{V}_P^r) = \arg\max_{d \in [1,3]} [a_{dD}]. \qquad (2)$$

If $\sigma_1 \gg \sigma_2, \sigma_3 \simeq 0$, a_{1D} will be greater than the the two others so that the dimensionality labelling $d^*(\mathcal{V}_P^r)$ results to 1. Contrariwise, if $\sigma_1, \sigma_2 \gg \sigma_3 \simeq 0$, a_{2D} *i.e.*, the planar behavior will prevail. At last, $\sigma_1 \simeq \sigma_2 \simeq \sigma_3$ implies $d^*(\mathcal{V}_P^r) = 3$.

We chose $\mu = \sigma_1$ because then : $a_{1D} + a_{2D} + a_{3D} = 1$, so that the three features can be considered as the probabilities of each point to be labelled as 1D, 2D, or 3D. It will help us to select the most appropriate neighborhood size by finding which radius favors the most one dimensionality (see Equation 3).
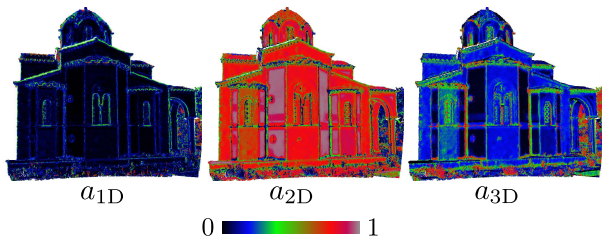

Figure 2: Behaviours of the three shape features (*TLS* dataset).

4 SCALE SELECTION

4.1 Optimal neighborhood radius

As presented in Section 2, the dimensionality features are computed for increasing radius values between a lower bound and an

upper bound (r_{min} and r_{max}, respectively). Their set up is presented in Section 4.2. They represent the minimal and maximal acceptable neighborhood radius according to the area of interest and the sensor. The $[r_{min}, r_{max}]$ space has been sampled in 16 values, which is a suitable trade-off between tuning accuracy and computing time. Since the radius of interest is usually closer to r_{min} than to r_{max}, the r values are not linearly increased but with a square factor. This allows to have more samples near the radius of interest and less when reaching the maximal values.

Two *radius-selection* criteria have been developed, namely the entropy feature (E_f) and the similarity index (Si).
First, a measure of unpredictability is given by the Shannon entropy of the discrete probability distribution $\{a_{1D}, a_{2D}, a_{3D}\}$:

$$E_f(\mathcal{V}_P^r) = -a_{1D} \ln(a_{1D}) - a_{2D} \ln(a_{2D}) - a_{3D} \ln(a_{3D}). \qquad (3)$$

The lower $E_f(\mathcal{V}_P^r)$ is, the more one dimensionality prevails over the two other ones. The relevance of scale selection is demonstrated in Figure 3. This criterion allows to define an optimal radius $r_{E_f}^*$ that minimizes $E_f(\mathcal{V}_P^r)$ in the $[r_{min}, r_{max}]$ space :

$$r_{E_f}^* = \arg\min_{r \in [r_{min}, r_{max}]} E_f(\mathcal{V}_P^r). \qquad (4)$$

A dimensionality labelling is then provided by : $d^*(\mathcal{V}_P^{r_{E_f}^*})$.
Second, $Si(\mathcal{V}_P^r)$ is defined as the ratio of neighbors P_k which dimensionality labelling is the same as P at scale r:

$$Si(\mathcal{V}_P^r) = \frac{1}{n} \sum_{P_k \in \mathcal{V}_P^r} \mathbf{1}_{\{d^*(\mathcal{V}_P^r) = d^*(\mathcal{V}_{P_k}^r)\}}. \qquad (5)$$

where $\mathbf{1}_{\{.\}}$ is the Indicator function and n is the number of points within $\mathcal{V}_P^r$. Si evaluates the homogeneity of the labelling within $\mathcal{V}_P^r$ by measuring the labelling similarity between neighboring points. Finally, we have:

$$r_{Si}^* = \arg\max_{r \in [r_{min}, r_{max}]} Si(\mathcal{V}_P^r). \qquad (6)$$

Si aims at reducing the noise of the results since r_{Si}^* is the scale for which the labelling of the current point is the most similar to the labellings of its neighbors at the same scale. Both criteria have been tested for our datasets. Results and conclusions are presented in Section 5.1.
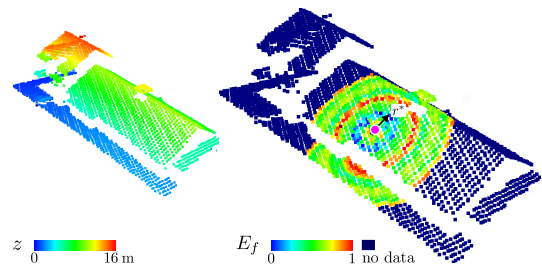

Figure 3: Illustration of the relevance of the entropy feature for scale selection (building roof in ALS_G). **Left**: 3D point cloud (height colored). **Right**: 3D point cloud colored with E_f computed for one point of interest (pink dot). The color of each neighbor P_k corresponds to the E_f value, computed on the smallest neighborhood containing P_k. The entropy first decreases until the neighborhood reaches the edge of the roof (blue circle – optimal size). Then, the entropy increases and becomes maximum when the neighborhood gathers distinct objects (red circle), here the ground and a chimney.

Figure 4 provides an example of the behaviour of the shape fea-

tures and the *radius-selection* criterion for two areas of the *TLS* dataset. Figure 4c corresponds to an area with higher noise than in Figure 4d due to distinct facade orientations. Thus, for Figure 4c (white rectangle of Figure 4a, for low r values, the point set locally looks scattered and a_{3D} prevails. Then, the increase of the size leads to the decrease of the noise effect. The planar feature becomes predominant. For Figure 4d, the facade is orthogonal to the laser beam resulting in lower noise. The neighborhoods remain planar with the increasing scale.
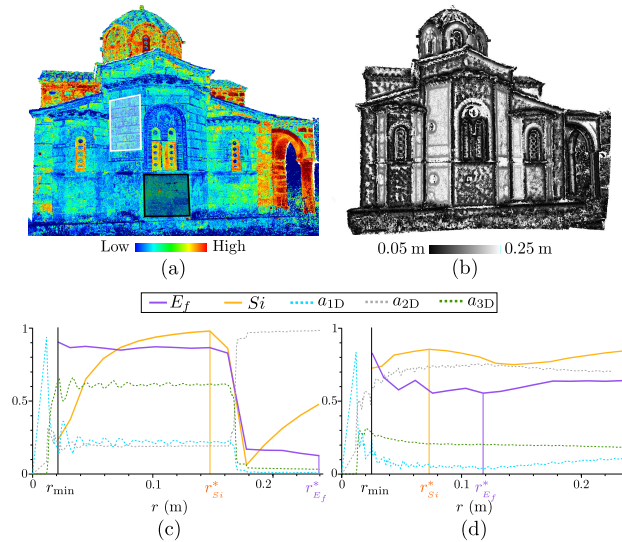

Figure 4: Optimal neighborhood radius retrieval (*TLS* dataset). 3D point cloud colored with **(a)** the intensity and **(b)** r^*. **(c-d)** Variation of the shape features and *radius-selection* criteria over the neighborhood radius over 50 points within the white (c) and black (d) rectangles of (a).

4.2 Bounding scales

The optimal neighborhood radius is retrieved between predefined lower and upper bounds. They depend on various characteristics, and are therefore specific to each dataset.

The choice of the lower bound is driven by (1) the point cloud noise; (2) the sensor specifications; and (3) computational constraints :

1. Local scattering may appear for linear or planar surfaces. It may stem from sensor noise or unfavorable geometrical configurations (*e.g.*, laser beam nearly parallel to a surface). Shape features are then biased and the point cloud may locally be erroneously labelled as "volumetric". The knowledge of the intensity of such noise is a prerequisite for solving this issue. Alternatively, the minimum size can be estimated applying the methods mentioned in Section 1 for this purpose.

2. The laser beam deflection system or the kind of mapping device may create point clouds with very irregular landscape sampling (low values in one direction and larger values in the orthogonal one). This is particularly true for MMS since datasets are acquired line by line with forward motion of the device. To cope with this issue, the r_{min} value is increased until $\mathcal{V}_P^r$ blends points belonging to at least two different scan lines.

3. A minimal number of points is required for the PCA. We consider relevant statistics start with 10 points. If the two first above-mentioned issues are not problematic, the point density directly provides the lower bound.

The selection of the upper bound is not critical and can be tuned with the knowledge of the lower frequencies of the relief *i.e.*, the size of the largest objects in the scene. For TLS and MMS datasets, this corresponds to facades (typically 3 m) whereas for ALS data, ground regions and large buildings are involved. As E_f and Si values may remain constant for these large planar areas, a cut-off value around 5 m works well, even if in practice 3-4 m are sufficient.

5 RESULTS

5.1 Comparison between *radius-selection* criteria

Both criteria E_f and Si have been tested, and radius selection results can be significantly different. The entropy feature criterion E_f has been preferred to the similarity index Si for two main reasons.

First, as illustrated in Figure 4, E_f and Si may provide distinct optimal radius that may correspond to distinct behaviours and privileged dimensionalities. For instance, in Figure 4c, Si erroneously selects a radius for which the volumetric dimensionality is predominant whereas E_f allows to detect the correct planar behaviour of the underlying surface. Si favors the most homogeneous scale, but, in practice, several radii may provide the same result ($Si = 1$). The smallest one is kept which has been revealed to be problematic for the *TLS* dataset since the optimal radius corresponds to the noise level. For *ALS* and *MMS*, such problem of erroneous dimensionality labelling does not occur. However, small radius are favored for large objects since Si may saturate, preventing the selection of larger radius. Consequently, the choice of the lower bound r_{min} for the multi-scale analysis becomes critical.

Second, for each 3D point at each radius of interest, Si depends on the dimensionality labelling of the neighbors at this neighborhood size (*i.e.*, E_f). Two iterations are therefore necessary to retrieve the optimal radius which requires more computational time and memory.

5.2 Scale selection

The six datasets have been processed with the proposed approach. Scale selection results are displayed in Figures 4 and 5, using the E_f criterion. The expected behaviours are reached showing the efficiency of the multi-scale analysis using E_f:

- Planar areas exhibit high values for ALS and TLS datasets (see Figures 4b and 5a). This is particularly true for the TLS dataset since the multi-scale analysis allows to overtake the noise level.

- Lowest scales correspond to border areas, allowing to retrieve building and vegetation edges for ALS datasets, windows in MMS datasets (Figures 5c and d), or small architectural elements in TLS.

- The algorithm allows to select high values when no predominant dimensionality is found for small and medium radius values. A sufficient number of 3D points is collected to retrieve a coherent dimensionality : 1D for wires and poles (Figures 5c and d), 2D for building facades in ALS datasets (*e.g.*, highest values in Figure 5b which correspond to a few number of points lying on the facades of high buildings) or 3D for vegetated areas (all datasets).

- The method deals well with varying point densities. Such variation may come from the number of overlapping strips for ALS datasets (a lower radius size is necessary when two
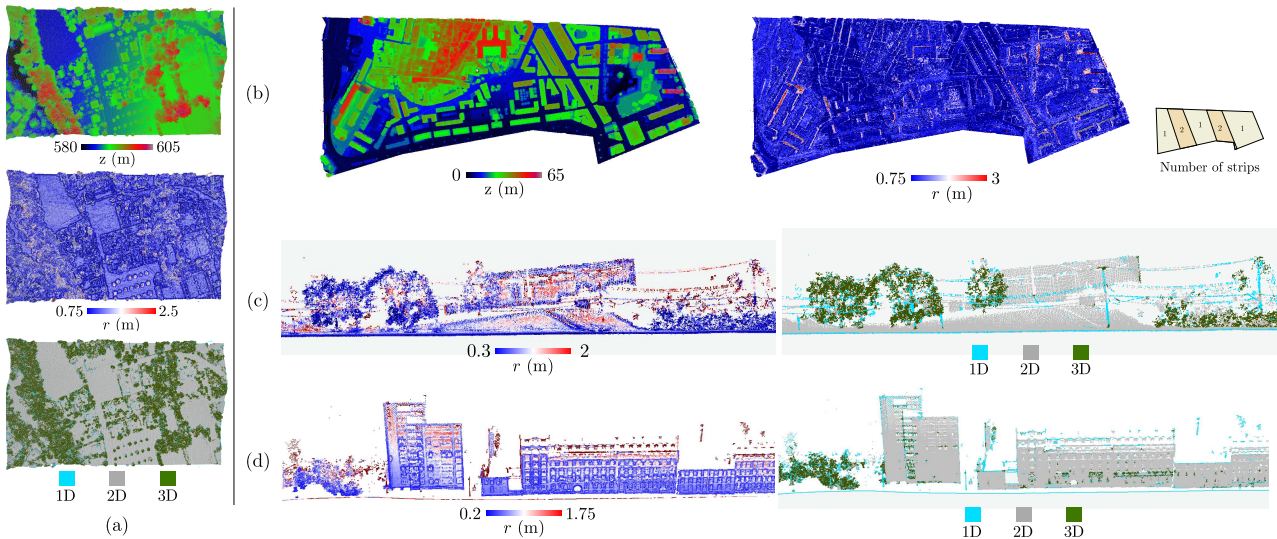
Figure 5: Radius selection (r) and dimensionality labelling for **(a)** ALS_G, **(b)** ALS_F, **(c)** MMS_U, and **(d)** MMS_F datasets.

overlapping strips exist for an area of interest, see Figure 5b), from occlusions or from the fluctuating incidence angle of the laser beam on the surfaces for TLS or MMS datasets (Figure 5d).

Figure 6: dimensionality labelling statistics for two datasets available with ground truth. **(a)** ALS_G dataset: the five main classes have been conserved. **(b)** MMS_U dataset: the 59 classes have been condensed into four classes according to their shape (1D: wires, poles, trunks etc. – 2D: wall, door, facade etc. – 3D: foliage, grass etc. – Clutter: objects without specific shapes). The number of points for each class is indicated inside brackets.

The effectiveness of the scale selection process can be assessed by retrieving the predominant dimensionality that has been retrieved (1D, 2D or 3D ?) for various objects of interest. Several examples are presented in Figure 5, and comparisons with manually labelled ALS and MMS data have been performed (cf. Figure 6). For both datasets, planar objects are correctly retrieved, with few errors corresponding to building edges or small urban items, especially for ALS data. Besides, objects with one privileged dimension such as poles, wires or trunks are mainly labelled as 1D (MMS_U dataset, in Figure 6b). However, since these objects are in fact cylindrical with a small width (*e.g.*, trunks or traffic lights), they may look planar or volumetric for medium-sized neighborhoods. Conversely, volumetric objects such as trees may locally look planar. This happens for large objects low point densities or when no multiple scatterings are found with vegetated areas. Such phenomena is limited for genuine 3D acquisitions (TLS or MMS datasets), whereas it is clearly enhanced for 2.5D

data such as for the ALS_R data (Figure 6a). High vegetation areas are principally labelled as *planar*, which is due to the dense canopy cover. Nevertheless, as displayed in Figure 5a, such effect almost disappears when 3D points are acquired within tree canopies.
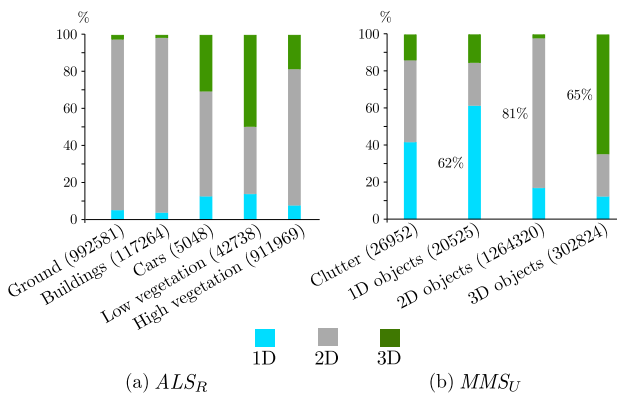
5.3 Comparisons with constant neighborhood size

The adaptive size strategy allows to retrieve a correct number of points to estimate the local dimensionality of the point set. Such strategy may be compared with the results achieved with neighborhoods of fixed size for a whole dataset.

Firstly, the three shape features are computed with both strategies for four classes of interest of the ALS_R dataset (Figure 7). One can see that the planar behaviour of *ground* and *building* points is improved with the adaptive size. Besides, for vegetated areas, the volumetric behaviour is slightly enhanced, however this happens in conjunction with an increase of the planar behaviour (dense canopies).
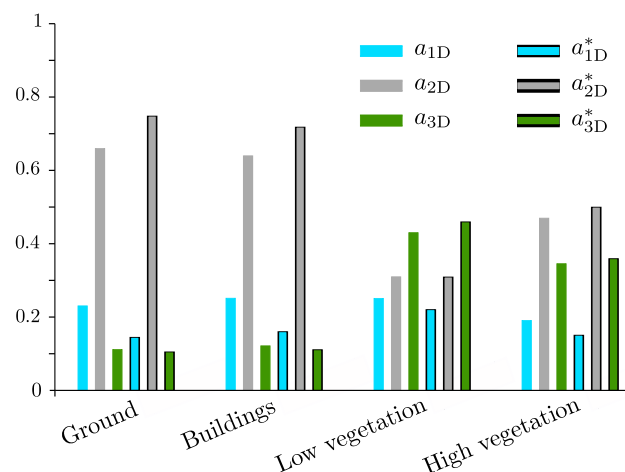

Figure 7: Improvements in shape feature computation using an optimal neighborhood radius per point. a_{dD} and a_{dD}^* are the shape feature for dimensionality d, computed with neighborhoods of constant size and adaptive size, respectively. The constant neighborhood size corresponds to the 30 nearest neighbors of each point.

Secondly, in addition to improved shape features, Figure 8 shows that the privileged dimensionality is also better retrieved with the

adaptive strategy (MMS_F dataset). This is particularly true for trees (1D→3D), small building facades elements (1D→2D), and tree trunks (2D→1D). Furthermore, since the geometrical analysis is less affected by the 3D point distribution, the labelling procedure is far less noisy.
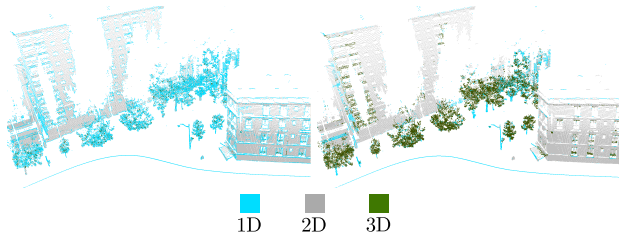

Figure 8: Dimensionality labelling for the MMS_F dataset using a constant neighborhood size for the whole point cloud (**left**), and the optimal value per point selected with E_f (**right**). The constant neighborhood size corresponds to a 1 m radius.

6 CONCLUSIONS AND FUTURE WORK

The objective of this paper was to provide a simple and general method for finding the optimal neighborhood radius for each 3D point of a lidar point cloud, with the simple knowledge of the point locations. The most appropriate scale of analysis has been automatically selected among a range of potential values using three low-level geometrical shape features describing the "dimensionality" of the local point set (1D-2D-3D). Results on distinct lidar point clouds acquired with airborne, terrestrial and mobile mapping systems under various conditions showed the effectiveness of the proposed method. No a priori knowledge or assumption on the point cloud distribution, density, laser scan pattern, object size or underlying structures was required to achieve such results. Several experiments demonstrated how favorably our method performed compared to constant neighborhood sizes, and how relevant such approach is. Finally, the prevailing dimensionality and the scale of interest for the underlying object are found for each point, providing interesting cues for many potential subsequent segmentation or classification algorithms.

Improvements will focus on two main issues. Firstly, as illustrated in Figure 3, in addition to the global minimum of the entropy feature providing the optimal neighborhood radius, the other minima are likely to be relevant for multi-scale characterization. Secondly, each 3D point is analyzed separately resulting in noisy dimensionality labelled point clouds. Therefore, such results would be improved and fastened by taking into account for close points the scale already computed for optimal neighbors.
Further work will consist in detecting and characterizing boundaries between objects by using and quantifying variations in the entropy feature. Finally, segmentation and classification issues will be tackled in order to fully benefit from the computation of non biased shape features and from the retrieval of a correct neighborhood for each point.

ACKNOWLEDGEMENTS

References

Bae, K.-H., Belton, D. and Lichti, D., 2009. A closed-form expression of the positional uncertainty for 3D point clouds. IEEE TPAMI 31(4), pp. 577–590.

Belton, D. and Lichti, D., 2006. Classification and segmentation of terrestrial laser scanner point clouds using local variance information. International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences 36 (Part 5), pp. 44–49.

Dey, T. and Sun, J., 2005. An adaptive MLS surface for reconstruction with guarantees. In: Symposium on Geometry Processing, Vienna, Austria, pp. 43–52.

Filin, S. and Pfeifer, N., 2005. Neighborhood systems for aiborne laser data. Photogrammetric Engineering and Remote Sensing 71(6), pp. 743–755.

Frome, A., Huber, D., Kolluri, R., Bülow, T. and Malik, J., 2004. Recognizing objects in range data using regional point descriptors. In: IEEE ECCV, Prague, Czech Republic.

Golovinskiy, A., Kim, V. and Funkhouser, T., 2009. Shape-based recognition of 3D point clouds in urban environments. In: IEEE ICCV, Kyoto, Japan.

Gross, H., Jutzi, B. and Thoennessen, U., 2007. Segmentation of tree regions using data of a full-waveform laser. International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences 36 (Part 3/W49A), pp. 57–62.

Gumhold, S., Wang, X. and Macleod, R., 2001. Feature extraction from point clouds. In: International Meshing Roundtable, Newport Beach, CA, USA, pp. 293–305.

Hadjiliadis, O. and Stamos, I., 2010. Sequential classification in point clouds of urban scenes. In: International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT), Paris, France.

Hoppe, H., DeRose, T., Duchamp, T., McDonald, J. and Stuetzle, W., 1992. Surface reconstruction for unorganized point clouds. Computer Graphics 26(2), pp. 71–78.

Lalonde, J., Unnikrishnan, R., Vandapel, N. and Herbert, M., 2005. Scale selection for classification of point-sampled 3-D surfaces. Technical Report CMU-RI-TR-05-01, Robotics Institute, Pittsburgh, PA, USA.

Mitra, N., Nguyen, A. and Guibas, L., 2004. Estimating surface normals in noisy point cloud data. International Journal of Computational Geometry and Applications 14(4-5), pp. 261–276.

Munoz, D., Bagnell, D., Vandapel, N. and Hebert, M., 2009. Contextual classification with functional max-margin markov networks. In: IEEE CVPR, Miami, FL, USA.

Pauly, M., Kobbelt, L. and Gross, M., 2006. Point-based multiscale surface representation. ACM Transactions on Graphics 25(2), pp. 177–193.

Shapovalov, R., Velizhev, A. and Barinova, O., 2010. Non-associative markov networks for 3D point cloud classification. International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences 38 (Part 3A), pp. 103–108.

Soudarissanane, S., Lindenbergh, R., Menenti, M. and Teunissen, P., 2009. Incidence angle influence on the quality of terrestrial laser scanning points. International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences 38 (Part 3/W8), pp. 183–188.

Tang, C., Medioni, G., Mordohai, P. and Tong, W., 2004. First order augmentations to tensor voting for boundary inference and multiscale analysis in 3-D. IEEE TPAMI 26(5), pp. 594–611.

Toshev, A., Mordohai, P. and Taskar, B., 2010. Detecting and parsing architecture at city scale. In: IEEE CVPR, San Francisco, CA, USA.

West, K., Webb, B., Lersch, J., Pothier, S., Triscari, J. and Iverson, A., 2004. Context-driven automated target detection in 3D data. In: SPIE 5426, Orlando, FL, USA, pp. 133–143.

Zwickler, M., Pauly, M., Knoll, O. and Gross, M., 2002. Pointshop 3d: An interactive system for point-based surface editing. In: ACM SIGGRAPH, San Antonio, TX, USA, pp. 322–329.