

Multi-scale Planar Segments Extraction from Point Clouds

T. Lejembre¹, C. Mura², L. Barthe¹ and N. Mellado¹

¹IRIT, Université de Toulouse, CNRS, INPT, UPS, UT1C, UT2J

²Department of Informatics, University of Zurich

Abstract

We propose a flexible method to extract a set of segments from a 3D point cloud that are relevant across several scales and optimal in a planar sense. Since planar geometric primitives are ubiquitous, especially in man-made scene, their accurate detection is crucial for an abstract representation of point-based 3D data.

In this paper, we introduce a new hierarchical graph representation in which each node represents a region at a given scale. The proposed graph is initialized with multiple segmentations performed at different scales and then reduced by collapsing groups of nodes. Each resulting group of nodes defines a meaningful segment and is obtained through an optimization that balances number of extracted segments and accuracy with respect to the input data in a planar sense. The output graph is a compact abstraction of the input cloud into multiple, possibly overlapping, segments, each relevant at a certain scale. The edges of the graph connect nodes whose segments overlap across different scales, thus allowing to represent both detailed and approximating parts of the scene.

1. Introduction

With the increase in accuracy and resolution of scanning devices, captured point-based geometries become more and more detailed and complex. This complexity includes largely varying sampling and features size in an object or a scene. This variation in feature size raises the need of *multi-scale* analysis methods that are able to characterize features in a discrete range from thin details to coarse shapes. In this context, we study the very challenging problem of extracting shape abstractions with a focus on finding the most representative segments that characterize a point set in a planar sense.

Point-Based Shape Abstraction : A popular method abstracting a point cloud by a set of simple geometric primitives is the well-known RANSAC [SWK07]. While being fast and robust, this approach explores the space of solutions in a randomized way and therefore does not guarantee consistent results across different runs. An alternative is the use of voting approaches as the 3D Hough Transform [RDvdHV07], which accumulate votes in a discretized parameter space. However, such methods generally have high memory requirements, which makes them unsuitable for processing large input point sets. A number of approaches significantly improve primitive extraction with the detection and the reinforcement of structural information [LWC*11, MMBM15]. This is unfortunately effective only in the presence of strong regularities in the data. Overall, all these techniques have a common drawback : the notion of *scale* is not explicit and they only deal with primitives with size of the same order of magnitude.

Multi-scale Analysis : Inspired by the scale-space theory introduced in computer vision [Wit87], the multi-scale analysis has been applied to 3D data [PKG03, MGB*12]. The point set is convoluted by a smoothing operator of progressively increasing width. Strong variations in the result correspond to steps in the relevant scales of analysis. While these methods are efficient for local geometry processing, they are intrinsically local and therefore lack a global or regional regularization for a direct use in shape abstraction.

Hierarchical segmentation : A multi-scale abstraction of a point cloud can be obtained by computing an over-segmentation and incrementally merging groups of segments until obtaining a coarse, over-simplified representation [AP10, FLD18]. This process produces redundant sequence of intermediate segmentations and a learning process can be used to detect the meaningful scales of abstraction [FLD18]. However, the generation of candidate segmentations is based on greedy merge operations, which is likely to miss intermediate representations that are meaningful in the context of a global, multi-scale analysis.

There is no effective approach able to extract an optimal set of segments that abstract parts of the input model potentially defined at different scale. Our contribution is a new multi-scale analysis framework optimizing the merging of regions segmented at different scales and stored in a graph. We end up with a set of segments, optimal in a planar sense and possibly describing features of different scale, following the procedure described in Section 2.




Figure 1: Example of segmentations at multiple scales and the associated graph. Left : hierarchical graph representation of segmentation at 20 scales, with arbitrary colors. Right : 4 segmentations from low (a) to high (d) scale.

2. Multi-scale Planar Region Extraction

The input of our algorithm is a set of unstructured points $\{\mathbf{p}_i\} \in \mathbb{R}^3$ with their normals \mathbf{n}_i that sample an unknown surface embedded in 3D. Our first goal is to perform a surface segmentation at multiple scales to obtain a hierarchical representation of the point cloud. In order to extract the initial regions, we perform a region growing algorithm based on a scale-space normal field (Section 2.1). The resulting segmentations – one for each scale level considered – are then expressed as a hierarchical graph encoding similarities between segments at different scales (Section 2.2). Finally, an optimization procedure extracts from the graph a reduced set of meaningful segments that are optimal in a planar sense (Section 2.3).

2.1. Segmentation at Multiple Scales

The first step aims at generating N segmentations of the point cloud, each at a different scale, where segmentations details increase as the scale decreases, from very thin to coarse features. We develop a segmentation process that is parametrized by a *scale parameter* $t \in \mathbb{R}^+$. The segmentation algorithm uses the surface normal vectors as local descriptors, hence the need of a scale-space normal field.

Scale-space Normal Field We compute a scale-space normal field η_i using the Moving Least Square (MLS) paradigm [GG07]. For each point \mathbf{p}_i and its neighbors, an algebraic sphere is fitted and the point is projected onto the 0-isosurface of the sphere. This is iteratively repeated until convergence and the scale-space normal η_i is equal to the normalized gradient of the converged scalar field. Similarly to the Growing Least Square method [MGB*12], the neighborhood of a point is determined by an euclidean ball centered at that point of radius t which defines the scale parameter of our pipeline.

By increasing the scale parameter t , the estimated normal η_i become smoother as required by the scale-space framework. However, since our primitive of interest is the plane,

we combine the previous MLS iterations with several re-weighting steps [ÖGG09]. By doing so, sharp edges are preserved even at high scales.

Region Growing The segmentation is done by a non-seeded deterministic region growing algorithm [RVDHV06] where regions are propagated from one point \mathbf{p}_i to a neighbor \mathbf{p}_j if the similarity between them is high enough. We relate this similarity measure to the angle between the two multi-scale normal vectors η_i and η_j previously computed. A threshold θ is set so that the region is expanded if the following criterion is respected :

$$1 - \eta_i(t) \cdot \eta_j(t) < \theta \quad (1)$$

Scale-space Sampling The theoretical continuous scale-space is discretized in N scale values $\{t_j\}$ between two bounds (t_{\min}, t_{\max}) . The sampling is chosen logarithmic to increase the precision at low scales because this is where variation of normal η_i will be more pronounced. One important property of this scale-space approach is that the scale parameter t (i.e the neighborhood size) is a distance so we can relate these bounds to a reference length like the bounding box diagonal length d_{ref} . Typical values are $t_{\min} = 10^{-4}d_{\text{ref}}$ and $t_{\max} = 10^{-1}d_{\text{ref}}$.

2.2. Hierarchical Graph Representation

The resulting segmentations generated individually at the N scales $\{t_j\}$ are turned into a hierarchical graph representation. Each of the N segmentation gives rise to one level in the graph and one region \mathcal{R}_k^j is represented by one node. We also store in each node the number of points $|\mathcal{R}_k^j|$ presented in the region. This attribute determines the importance of the node and is proportional to the node size in Figure 1 and 2.

The connection between the nodes is done between two consecutive scales level t_j and t_{j+1} . The similarity between the two linked nodes can be expressed as a weight attribute of the edge connecting them. In our case, the similarity is




Figure 2: Output of our approach. (a) Initial hierarchical graph representation recolored according to the grouping obtained with the optimization. Each color denotes one planar region. (b) Reduced graph after optimization, where each node represents the set of nodes of corresponding color in the initial graph. (c) One subset of the extracted regions from two views. (d) The other subset of regions from the same views.

simply equal to the number of points that are commonly shared by the two regions \mathcal{R}_k^j and \mathcal{R}_l^{j+1} :

$$w_{k,l}^j = \sum_{\mathbf{p}_i} \delta_{\{\mathcal{R}_k^j \cap \mathcal{R}_l^{j+1}\}}(\mathbf{p}_i) \quad (2)$$

where δ is the indicator function equal to 1 if \mathbf{p}_i is in the intersection of the regions or 0 otherwise.

2.3. Optimization

The full hierarchy of segmentations contains pertinent planar regions, possibly overlapping, at different scales, as well as irrelevant regions. Therefore the graph must be filtered so as to extract only a reduced set of these numerous regions. This final phase can be expressed as a *relabeling* problem. Given the set of all regions $R = \{\mathcal{R}_k^j\}$, the goal is to assign to each region $r \in R$ a label $f_r \in \mathcal{L}$ from a finite set of labels \mathcal{L} . Each label f_r corresponds to a candidate planar model, obtained by fitting a plane to region r using Principal Component Analysis (PCA). The regions that are assigned the same label define a meaningful segment that abstracts a part of the input model at a certain scale. For an appropriate abstraction we expect that : (i) the resulting segments explain the data in an accurate way, (ii) similar regions across scale levels be assigned to the same label and (iii) few labels be employed. These three conditions correspond to three different energy terms : data cost, multi-scale smoothness cost and label cost, respectively. They can be combined into one single energy that can be optimized efficiently [DOIB12] :

$$E(f) = \sum_{r \in R} D_r(f_r) + \lambda_s \sum_{r_1, r_2 \in R} V_{r_1, r_2}(f_{r_1}, f_{r_2}) + \lambda_l \sum_{L \in \mathcal{L}} \delta_{\{L\}}(f) \quad (3)$$

Data cost In order to obtain regions as planar as possible, the data cost for assigning one region $r \in R$ the label f_r is equal to the sum of squared distance from all of its data point to the PCA plane that approximates region r' :

$$D_r(f_r) = \sum_{\mathbf{p}_i \in r} d(\mathbf{p}_i, r')^2 \quad (4)$$

where d is the orthogonal distance from \mathbf{p}_i to the plane fitted to the region r' using principal component analysis.

Multi-scale smoothness cost The smoothness cost penalizes pairs of similar regions that are assigned to different labels and connected in the graph. This term usually refers to spatial smoothness by penalizing spatially close elements of different labels. In our work, we ensure this constraint across scale-space with the following multi-scale smoothness cost :

$$V_{r_1, r_2}(f_{r_1}, f_{r_2}) = \begin{cases} 0 & \text{if } f_{r_1} = f_{r_2} \\ 1 & \text{otherwise} \end{cases} \quad (5)$$

Label cost The third term of Equation 3 imposes a restriction on the number of used labels with a fixed penalty for each label f_r that exists in the current relabelling. This penalty is inversely proportional to the number of points in the region r . Therefore, the optimization process will significantly reduce the number of final labels, i.e. the number of segments explaining the data.

3. Results

We show the results on the example of an L-shaped house in order to demonstrate the functioning of our planar regions extraction introduced in Section 2. In this experiment, we used as propagation threshold on scale-space normal vectors $\theta = 0.005$ (see Equation 1) and as weights to balance the different energy terms $\lambda_s = 0.01$ and $\lambda_l = 0.001$ (see Equation 3).

Figure 1-right shows the results of the region growing process where we can see 4 of the 20 segmentations performed at different scales. Results go from over segmented to over simplified segmentations, which forms an extensive set of candidate regions suitable for shape abstraction. The associated hierarchical graph is shown on the left of Figure 1 with the scale increasing from bottom to top. The node size is proportional to the number of points contained in the region, and the edge weight is proportional to the number of common points in the two linked regions (see Equation 2).

The output of the optimization process is a relabelling of all the initial regions of the hierarchy, with each label corresponding to a group of initial regions that define an output segment. Each of these segments describes one part of the input model, meaningful at a certain scale. In our example, the 569 regions of the segmentations of Figure 2(a) are compacted into a new set of 15 segments, depicted in Figure 2(b). These segments can be intuitively visualized with the planes of the labels they were assigned by the optimization, as shown in Figure 2(c) and Figure 2(d). Note that the layout of the graph in Figure 2(b) is arbitrary since each node corresponds to a group of initial nodes of possibly different scales; thus, a node in the output graph can not be associated to a single scale value. We can observe in Figure 2(c) and Figure 2(d) that some segments exist that include several smaller ones; this is the case e.g. for the large green segment in Figure 2(d)-top, which overlaps with the two segments shown respectively in cyan and magenta in Figure 2(c)-top. All of these three segments are indeed to be considered meaningful, though at different scales, as they describe the overall shape of the side of the house with two different interpretations: a detailed one at low scale with the cyan and the magenta segments; a global one approximating both the roof and the wall as a same entity with the dark green region.

The overall time execution from the beginning of the scale-space normal field estimation to the end of the optimization step is about 30 seconds for 5035 sample points and 20 scales.

4. Conclusion and Future Work

We introduced a method to extract a set of segments that abstract the parts of an input point cloud at different scales. This results is obtained by using a graph representation of multiple segmentations performed at different scales based on a scale-space normal field. By grouping the nodes in the graph through an optimal relabeling, we obtain potentially overlapping segments corresponding to parts meaningful at different scale and optimal in a planar sense. This flexible shape abstraction framework can provide highly diversified set of surface parts to improve high level analysis such as semantic segmentation or structure detection.

Future Work Although our method does not require any scale parameter, the scale-space normal comparison threshold θ of Equation 1 remains tedious to set manually. We will investigate an automatic way of region propagation, linking for instance the angle between two such normal vectors and the scale at which they have been estimated. Other differential properties, such as curvatures information, can also be added to this propagation criterion since they can be obtained directly from the MLS approach we used. Furthermore, these additional data are a possible way to optimally extract other kinds of geometric primitives. More complex data representation like feature histograms can be associated to each region in order to improve the optimization step. Finally, the execution time remains prohibitive for point clouds composed of several hundreds of thousands of points. In this case, the graph becomes too large to be processed using the current optimization model, which is why an alternative optimization strategy should be investigated as well.

Acknowledgments

This research has been partially supported by the FOLD-Dyn project (ANR-16-CE33-0015-01) and the CIMI Labex (ANR-11-LABX-0040) and by a grant Forschungskredit "Postdoc" of the University of Zurich (grant number FK-18-022)

References

- [AP10] ATTENE M., PATANÈ G. : Hierarchical structure recovery of point-sampled surfaces. In *Computer Graphics Forum* (2010), vol. 29, Wiley Online Library, pp. 1905–1920.
- [DOIB12] DELONG A., OSOKIN A., ISACK H. N., BOYKOV Y. : Fast approximate energy minimization with label costs. *International journal of computer vision*. Vol. 96, Num. 1 (2012), 1–27.
- [FLD18] FANG H., LAFARGE F., DESBRUN M. : Planar shape detection at structural scales. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2018).
- [GG07] GUENNEBAUD G., GROSS M. : Algebraic point set surfaces. In *ACM Transactions on Graphics (TOG)* (2007), vol. 26, ACM, p. 23.
- [LWC*11] LI Y., WU X., CHRYSATHOU Y., SHARF A., COHEN-OR D., MITRA N. J. : Globfit : Consistently fitting primitives by discovering global relations. In *ACM Transactions on Graphics (TOG)* (2011), vol. 30, ACM, p. 52.
- [MGB*12] MELLADO N., GUENNEBAUD G., BARLA P., REUTER P., SCHLICK C. : Growing least squares for the analysis of manifolds in scale-space. In *Computer Graphics Forum* (2012), vol. 31, Wiley Online Library, pp. 1691–1701.
- [MMBM15] MONSZPART A., MELLADO N., BROSTOW G. J., MITRA N. J. : Rapter : rebuilding man-made scenes with regular arrangements of planes. *ACM Trans. Graph.*. Vol. 34, Num. 4 (2015), 103–1.
- [ÖGG09] ÖZTIRELI A. C., GUENNEBAUD G., GROSS M. : Feature preserving point set surfaces based on non-linear kernel regression. In *Computer Graphics Forum* (2009), vol. 28, Wiley Online Library, pp. 493–501.
- [PKG03] PAULY M., KEISER R., GROSS M. : Multi-scale feature extraction on point-sampled surfaces. In *Computer graphics forum* (2003), vol. 22, Wiley Online Library, pp. 281–289.
- [RDvdHV07] RABBANI T., DIJKMAN S., VAN DEN HEUVEL F., VOSSELMAN G. : An integrated approach for modelling and global registration of point clouds. *ISPRS journal of Photogrammetry and Remote Sensing*. Vol. 61, Num. 6 (2007), 355–370.
- [RVDHV06] RABBANI T., VAN DEN HEUVEL F., VOSSELMANN G. : Segmentation of point clouds using smoothness constraint. *International archives of photogrammetry, remote sensing and spatial information sciences*. Vol. 36, Num. 5 (2006), 248–253.
- [SWK07] SCHNABEL R., WAHL R., KLEIN R. : Efficient ransac for point-cloud shape detection. In *Computer graphics forum* (2007), vol. 26, Wiley Online Library, pp. 214–226.
- [Wit87] WITKIN A. P. : Scale-space filtering. In *Readings in Computer Vision*. Elsevier, 1987, pp. 329–332.