Mathematical Morphology on Irregularly Sampled Data Applied to Segmentation of 3D Point Clouds of Urban Scenes

Teo Asplund[1], Andrés Serna[2], Beatriz Marcotegui[3], Robin Strand[1], and Cris L. Luengo Hendriks[4]

[1] Centre for Image Analysis, Division of Visual Information and Interaction, Uppsala University, Uppsala, Sweden.
[2] Terra3D 36 Boulevard de la Bastille, 75012 Paris, France
[3] MINES ParisTech, PSL Research University, CMM - Centre for Mathematical Morphology 35 rue Saint Honoré - Fontainebleau, France
[4] Flagship Biosciences Inc., Westminster, Colorado, USA.

Abstract. This paper proposes an extension of mathematical morphology on irregularly sampled signals to 3D point clouds. The proposed method is applied to the segmentation of urban scenes to show its applicability to the analysis of point cloud data. Applying the proposed operators has the desirable side-effect of homogenizing signals that are sampled heterogeneously. In experiments we show that the proposed segmentation algorithm yields good results on the Paris-rue-Madame database and is robust in terms of sampling density, i.e. yielding similar labelings for more sparse samplings of the same scene.

1 Introduction

Asplund *et al.* introduced morphological operators able to deal with irregularly sampled data. These operators, initially defined for 1D-signals, better approximate the underlying continuous nature of the world [1, 2] than regular, discrete morphology.

Large amounts of 3D data, acquired by Mobile Mapping System, are available nowadays (Oakland [6], KITTI [4], Rue Madame [10] datasets). Such data is of interest for a wide range of applications, for example: urban planning, documentation of cultural heritage, city modeling, among others [14]. 3D point clouds represent discrete versions of real world scenes. The sampling is by construction irregular, as it depends on the distance from the surface to the device, the angle of the surface, occlusions, vehicle speeds, etc. Thus, such data are particularly suitable to illustrate the potential of *irregular morphological operators* (i.e. operators on irregularly sampled data).

Previous morphological approaches to the problem of segmenting urban scenes have made use of regular, discrete morphology, by projecting the point cloud onto a regular grid [9, 11]. A disadvantage of this approach is that the segmentation depends on the coarseness of this grid (since nearby points can no longer be differentiated). Moreover, the projection loses information where the scene contains
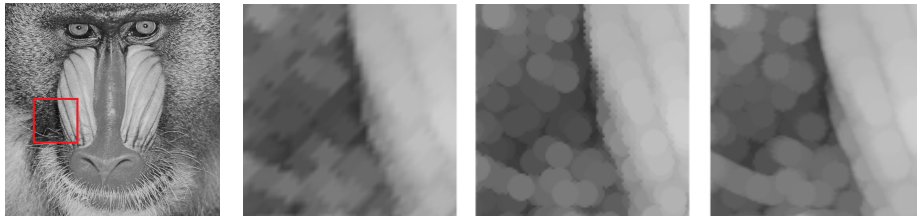
Fig. 1. From left to right: input image, regular dilation on sampled input, irregular dilation on sampled input, and regular dilation on non-sampled input.

one part overshadowing another, for example where ground is occluded by a roof. A 3D morphological approach is proposed by Calderon and Boubekeur [3], that requires the estimation of the underlying surface.

In this paper we introduce a framework for applying 2D, real-valued morphology to irregularly sampled signals, extending previous work by Asplund *et al.* [1]. We show that our proposed approach can be applied to 3D point clouds with promising results for segmenting point clouds of urban scenes.

2   Morphology on Irregularly Sampled Data

Morphology on irregularly sampled data is related to morphology on graphs. Najman and Cousty provide a nice survey of a large number of papers regarding this issue [7]. Here we propose a new approach to morphology on point clouds.

Irregular morphology in the specific case of rectangular structuring elements (SE) for 2D real-valued signals was dealt with in a previous paper [1] focusing mostly on regularly sampled input. Moreover, the previously presented method has some limitations regarding the sampling of the transformed signal. The main issue is that edges end up being "fuzzy" because the number of samples that are generated near edges is too low. A proposed solution is to iterate operators. In this section we present a similar idea for irregular morphology, which however does not require such iterations.

Morphology on irregularly sampled signals has a number of advantages. For one, in the regular case, the SE depends on the sampling grid, i.e. a continuous shape must first be discretized, which creates artifacts. This is illustrated on a regularly sampled image in Fig. 1, where an input image is slightly blurred (Gaussian kernel with $\sigma = 1.0$ pixels) in order to remove high frequencies and then subsampled by taking every third pixel in every third row. The subsampled image is then dilated with a small disk ($r = 2$ pixels in the subsampled grid) using both regular, discrete morphology, as well as the proposed irregular approach detailed below. The figure also shows the result of applying a regular, discrete dilation to the original, blurred image (before subsampling), with a SE of corresponding radius. As can be seen, the small disk, after discretization, ends up looking more like a diamond. It is also worth noting that irregular morphology has a number of advantages over regular, discrete morphology when it comes

to approximating continuous morphology on images in the discrete domain [1]. However, this is outside the scope of this paper.

We investigate the applicability of the proposed irregular morphological operators to point clouds. This seems a natural application, since they are naturally irregularly sampled.

2.1   Dilations of Irregularly Sampled Signals

The basic idea behind the proposed algorithm is to sample two disks at each input point. One with radius r, and one slightly larger with radius $r + \varepsilon$. Additionally, for each of these samples, examine whether a disk of radius $r + \varepsilon$ placed at any other input sample at a larger z-value overlaps this sample, and if so, discard it. The remaining samples from the larger disk are then projected downward, until they hit a disk (of radius r) placed at a lower sample, or falls through to $-\infty$. The remaining samples of the smaller disk are kept, defining a border between transitions in the value of the dilation. Here ε defines the width of this border.

We will use $B_{\mathbf{x}}$ to denote the set B translated by $\mathbf{x}$, i.e. $B_{\mathbf{x}} = \{\mathbf{b} + \mathbf{x} \mid \mathbf{b} \in B\}$. To perform a dilation by a disk of radius r of an irregularly sampled signal, we propose the following algorithm:

Choose some small $\varepsilon > 0$. For a disk shaped structuring element of radius r

$$D(r) = \{(x, y) \mid x^2 + y^2 \leq r\} \tag{1}$$

let $\check{D}$ denote a sampled version of D. In this paper, we will only consider the case where the disks are sampled uniformly along the border with 8 samples, and a ninth sample at the origin. This means that smaller circles are sampled more densely than larger ones. This keeps the number of samples down and larger SEs lead to larger plateaus, which probably do not need as many samples. However, less naive ways of sampling the disk are obviously of interest, but outside the scope of this paper.

To dilate some input $I = \{(x_i, y_i, z_i)\}_{i \in \{1,2,...,N\}}$ containing N samples at positions (x_i, y_i, z_i) do the following:

1. Let $I_\delta^0 = \emptyset$. Here I_δ^0 is empty, but will be filled with samples from the dilation as we process the samples in I.
2. In this step we iteratively select an unprocessed sample $(x_c, y_c, z_c) \in I$ and mark it as processed. This is analogous to sliding the SE over the image. In Fig. 2 the blue X is the sample being processed.
3. For each sample (x, y) of the translated disk $\check{D}_{(x_c, y_c)}(r)$, and each sample $(x_i, y_i, z_i) \in I \setminus \{x_c, y_c, z_c\}$, such that $z_i \geq z_c$, check if the distance from (x, y) to (x_i, y_i) is at most $r + \varepsilon$. If so, discard (x, y) (see the blue star in Fig. 2). Let $\hat{D}_{(x_c, y_c)}(r)$ denote the remaining set of samples (blue circle). I.e. discard samples from the smaller disk that are overlapped.
4. Similarly to the previous step, for each sample (x, y) of the translated disk with radius $r + \varepsilon$, and each sample in $I \setminus \{x_c, y_c, z_c\}$ no lower than z_c, check if the distance from (x, y) to (x_i, y_i) is at most $r + \varepsilon$. If so, discard
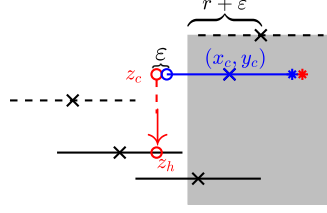
Fig. 2. Illustration of the algorithm from a side view. The disk shaped SE is located on top of each sample. The red circle is a sample (x, y) from $D(r + \epsilon)$. The black dashed disks do not belong to H since they are too far from (x, y) (and one of them is above z_c). Other two disks, placed below z_c, belong to H. z_h is found by projecting down until hitting the highest disk in H. In case no disk is hit, such a sample is discarded. The shaded area illustrates the part that is covered by a disk at a higher level. The samples that are in this shadow are discarded (blue and red stars as well as the blue X). The remaining samples are added to the dilation. This approach is similar to the paper by Asplund *et al.* [2].

(x, y) (see the red star in Fig. 2). For each remaining sample (x, y) and each $(x_i, y_i, z_i) \in I$, such that $z_i < z_c$, let $H(x, y)$ be the set of values of samples for which the distance between (x, y) and (x_i, y_i) is at most r (see the red circle and black solid disks in Fig. 2), i.e.

$$H(x, y) = \{z_i \mid (x_i, y_i, z_i) \in I \wedge z_i < z_c \wedge \sqrt{(x - x_i)^2 + (y - y_i)^2} \leq r\} \quad (2)$$

Let $z_h(x, y) = \bigvee H(x, y)$. Conceptually, z_h is the value of a given sample (x, y) of the bigger disk, after projecting downward until hitting a disk at a lower value (or dropping out to $-\infty$). See red circle in Fig. 2.

5. Finally, let

$$\hat{D}^c_{(x_c, y_c)}(r) = \{(x, y, z_c) \mid (x, y) \in \hat{D}_{(x_c, y_c)}(r)\}, \quad (3)$$

$$\hat{D}^h_{(x_c, y_c)}(r + \varepsilon) = \{(x, y, z_h(x, y)) \mid (x, y) \in \hat{D}_{(x_c, y_c)}(r + \varepsilon)\}, \text{ and} \quad (4)$$

$$I^{i+1}_\delta = I^i_\delta \cup \hat{D}^c_{(x_c, y_c)}(r) \cup \hat{D}^h_{(x_c, y_c)}(r + \varepsilon) \quad (5)$$

I.e. add the samples spawned by the processed sample to I^{i+1}_δ, where i is the number of previously processed samples. These steps are then iterated until there are no unprocessed samples left.

Then I^N_δ is the dilated signal. That is, I^N_δ is the irregular dilation of I by a disk shaped structuring element of radius r. The parameter ε controls the distance between samples where translated disks overlap. This parameter should be small (as a rule, at least smaller than the distance between samples in the input).

Figure 3 illustrates an example of applying this algorithm to an input containing five irregularly sampled points when seen from above. The samples that fall in the shadow of any SE placed at a higher sample are discarded. The remaining samples from the larger SE are projected down until they hit a SE at a lower input sample (or until they fall through to $-\infty$).
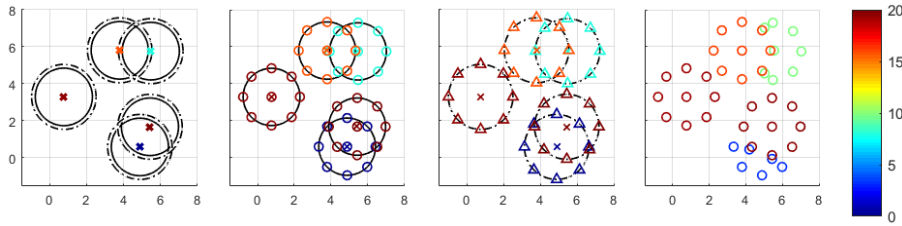
Fig. 3. A dilation of an irregularly sampled signal containing five samples (illustrated by crosses), by a disk shaped structuring element (solid black line). The first figure shows a SE placed on top of each input sample. The dashed circles illustrate the larger disk in step 3 (with radius $r + \varepsilon$). The second figure shows the sampled SEs (i.e. $\check{D}(r)$). The third figure shows the sampled larger disk (i.e. $\check{D}(r + \varepsilon)$). The last figure shows the final result (i.e. I_δ^N). Here the SE samples that fell under a large disk at a higher location (i.e. one of the disks in the third figure) have been suppressed, and the samples of the larger disks have been projected down, helping define the border between the sampled disks (see for example the bottom right of the last figure).

2.2   Erosions, Openings, and Closings

To get the erosion $\varepsilon(I)$ of an input signal I, one can simply compute the dilation $\delta(I^{\complement})$, where $I^{\complement} = \{(x_i, y_i, -z_i) \mid (x_i, y_i, z_i) \in I\}$, the erosion is then $\delta(I^{\complement})^{\complement}$. Openings and closings can be formed by composition [12].

3   Irregular morphology for point clouds

Figure 4 shows some examples of different morphological operators applied to a simple point cloud depicting a signpost. The signpost is first eroded by a disk of radius 0.016m, followed by a dilation with the same SE (yielding an opening). The result is interpolated back onto the original input positions, and then subtracted from the input, resulting in the tophat, which has enhanced the thin pole holding the sign. The value of ε is set to 10^{-6}m in all these cases.

We use nearest neighbor interpolation for a few reasons: Firstly, for the sake of simplicity. Secondly, empirically this seems to give good results. Finally the morphological operators will, in general, create cusps when applied to a signal, even if the input is smooth [1]. Therefore, a smooth interpolation of the opening is problematic. Some other possible candidates would be linear interpolation, or natural interpolation [13], however in some preliminary testing, these seem to give worse results for our data.

An interesting property of the irregular operators is their tendency to homogenize the sampling density. The sampling of the ground in the input is less dense further from the camera (since these samples are further from the sensor). The eroded point cloud, on the other hand, has a similar sampling density close to, and further away from the camera. In essence, this behavior stems from not sampling inside the disks (except for one sample in the center) thereby reducing the
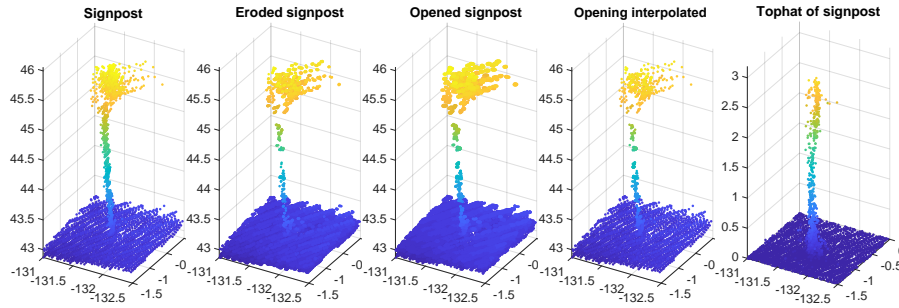
Fig. 4. Examples of irregular morphology applied to a simple point cloud of a signpost. Here, the SE is a disk with radius 0.016m.

sampling density where it is high (since disks are likely to overlap many nearby samples) accelerating further processing, while increasing the density where it is low (since one input sample gives rise to many output samples and there are fewer neighboring samples in the input that are likely to be suppressed). This homogenization may better define surfaces, enabling more accurately detecting object boundaries. Moreover, homogenizing the density introduces an implicit normalization of features and a more robust feature description, avoiding the excessive contribution of high density areas.

4   Application to Urban Point-Cloud Segmentation

As previously mentioned, the problem of segmenting point clouds of urban scenes has already been addressed using morphology on elevation images. However, using elevation images has some drawbacks. The images being projected onto a regular grid means that the segmentation will depend on the coarseness of said grid. Furthermore there is a question about how to deal with holes in the projected image. Thus, this problem seems like a natural candidate for applying irregular morphology.

We will consider a problem where the goal is to segment point clouds of urban scenes, to illustrate the applicability of our proposed method. In order to segment the point cloud, we first focus on separating out the facades and objects, followed by a refinement which yields the final segmentation. Figure 5 shows a diagram outlining the workflow.

We make use of point clouds from the Paris-rue-Madame dataset[5] [10], containing two point clouds from a 3D Mobile Laser Scanning (MLS) of a section of a street in Paris. Apart from the 3D-location, each point is also equipped with three additional pieces of data: reflectance, label, and class. However, we only use the position of the points for segmenting. We use the provided classes to test our segmentation. The reflectance and label information is ignored.

---

[5] © 2014 MINES ParisTech. MINES ParisTech created this special set of 3D MLS data for the purpose of detection-segmentation-classification research activities.
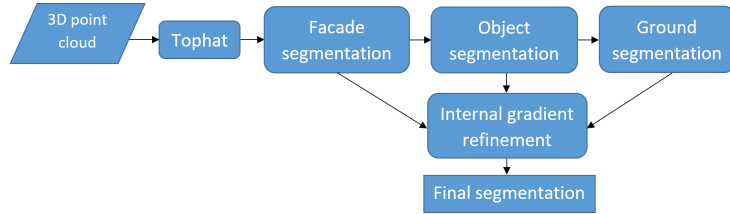
Fig. 5. Workflow diagram of the proposed algorithm. The output is a segmented point cloud with labels: ground, facade and object.

4.1 Tophat

An initial preprocessing step will be applied to the entire point cloud, the purpose being to flatten out the surface, removing large-scale variations in height (e.g. slopes and hills), while preserving local height differences (e.g. cars, pedestrians, facades). This preprocessing consists of a tophat applied to the point cloud. First we apply an opening to the set of points (by applying an erosion followed by a dilation), the result is then subtracted from the original signal (via interpolation as described in Sec. 3).

The size of the structuring element used by the opening controls the size of the support of height variations to be preserved. Based on the size of objects of interest one can make a decision on the radius of the SE. In this case, our objects of interest are on the scale of cars or smaller (we also capture facades, which are generally quite thin). Therefore, we set the radius of the disk to be 1.5m, since a disk of this size cannot fit inside a normal vehicle. Figure 6 shows an example of the result of taking the tophat of a point cloud from the dataset [10].

4.2 Segmenting Facades

Facades are initially segmented by thresholding the tophat at a large value h_{facade}. We are assuming the highest objects in the scene are house facades. If this does not hold, one could use some more sophisticated feature, such as considering the elongation of objects [5,11], or learned features, but for our purposes, this assumption is fine. The value of h_{facade} is not very sensitive. The important part is that the threshold contains the top part of facades. We simply fix this value to be 5 meters. This initial segmentation fails to include a large number of points that are actually parts of facades, however, it contains very few false positives (very tall objects might be included, e.g. lampposts, but they are easily eliminated). The purpose of this initial segmentation is to localize the facades. Following this localization, we can find candidate facade points by thresholding the result of the tophat at a lower value h'_{facade} as described in Section 4.5. These candidates are then filtered based on their distance (in 2D) from the initially selected facade points.

Specifically, let

$$F_{h_{\text{facade}}} = \{(x,y,z) \in I \mid z > h_{\text{facade}}\}, \ F_{h'_{\text{facade}}} = \{(x,y,z) \in I \mid z > h'_{\text{facade}}\} \ (6)$$
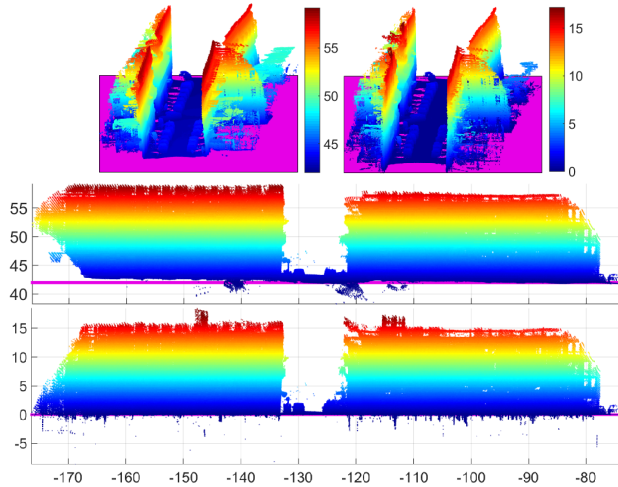
Fig. 6. The first row shows the tophat (right) of the input (left). Also shown is a horizontal plane in purple that is approximately at ground level for the input figure, and at 0 for the figure of the tophat, to illustrate slope. Rows two and three show the same point clouds from the side. Notice that the slope of the input has been removed in the tophat (bottom), and large scale variations in height have been suppressed. This makes thresholding for segmentation feasible. The tophat contains some negative values (points below the plain), which stems from interpolation artifacts. These negative points happen mostly inside buildings, where the sensor has sampled parts of the insides of buildings so that it has some samples from rooms on top of each other.

where I is the set of input samples and $F_{h_{\mathrm{facade}}}$ denotes the initial segmentation of the facades associated with h_{facade}. The refined segmentation F' is given by the following set of points:

$$F' = \{(x, y, z) \in F_{h'_{\mathrm{facade}}} \mid d((x, y, z), F_{h_{\mathrm{facade}}}) < \varepsilon_f\} \tag{7}$$

here $\varepsilon_f > 0$ is a parameter that controls the distance that the initial facade labels can be propagated. For our experiments we use $\varepsilon_f = 0.05$m. The function d denotes the minimum 2D-distance between all points in a set and some sample:

$$d((x, y, z), F) = \min_{(x', y', z') \in F} \sqrt{(x - x')^2 + (y - y')^2} \tag{8}$$

By restricting the potential facade points to those that are near the initially selected set of points, we can rule out a lot of false positives, meaning that we can push the threshold value h'_{facade} close to 0, i.e. close to the height of the ground, even though the number of ground points that are above the threshold increases as h'_{facade} approaches 0. Thresholding at $h'_{facades}$ (for reasonable values) will clearly find points that are parts of objects too (e.g. pedestrians), since the lower parts of facades are at the same height as objects. However, such points will generally not be immediately next to facades (although it is possible) and are therefore discarded, since they are too far ($\geq \varepsilon_f$) from the initial facade points. The remaining points are labeled "facade". Figure 7 illustrates this process.
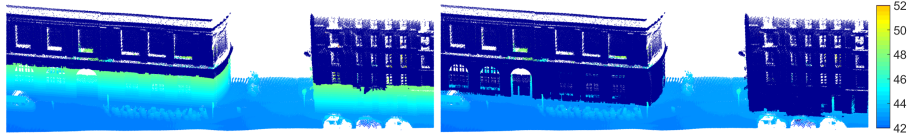
Fig. 7. The left figure shows the initial facade segmentation in dark blue (other colors indicate height). On the right: points labeled facade after the process in Sec. 4.2.

4.3   Segmenting Objects

We now focus on the objects, i.e. everything that is neither facade nor ground. At this point, we only need consider the non-facade points. Again, we follow the procedure outlined in the previous section. Objects are initially segmented by selecting a threshold h_{object}. Again, this initial threshold is set at a height such that there are few non-object points above it, but many object points fall below it. To remedy this, a threshold of h'_{object} is selected (see Section 4.5) and applied to the remaining points. Finally, points that are too far from the initial set of object points (in 2D) are discarded (as in Eqn. 7 with appropriate substitutions). The remaining points are labeled "object".

4.4   Segmenting Ground and Refining Labeling

The points that are not labeled "object" or "facade" are labeled "ground". However, there may be some points that are close to objects and facades that are mislabeled as "ground," since the thresholds h'_{facade} and h'_{object} cannot be set too close to ground level without introducing noisy labels close to the borders between ground and not ground. An additional step to refine the segmentation is performed by finding the internal gradient G_i [8] of the set of points, G, labeled ground. I.e. $G_i = G - \varepsilon_r(G)$, where ε_r denotes the erosion with a disk shaped structuring element of radius $r > 0$. Here r is set to a small value which depends on the average minimum distance from any ground point p to the rest of the ground points $G \setminus \{p\}$, i.e.

$$r = \frac{C}{|G|} \sum_{p \in G} d(p, G \setminus \{p\}) \tag{9}$$

where $C \geq 1$ is a constant increasing the radius to give more context (we fix $C = 10$). In other words, this radius depends on the point cloud density. We want it to be small, in order to get a good approximation of the internal gradient, but it cannot be too small, since this would essentially lead to each point being looked at in isolation, which is a problem, since the gradient obviously requires information about values at nearby points. The internal gradient G_i enhances the edges, which are the ground points that are likely mislabeled. A final threshold h_{edge} is chosen as described in Section 4.5. The set of points above this threshold E, are then used instead of $F_{h'_{\mathrm{facade}}}$ and $F_{h'_{\mathrm{object}}}$ (see Eqns. (6) – (7)). As a result, more points close to the ground near edges can be correctly labeled, though this may lead to a small number of ground points being mislabeled.

4.5   Choosing the Threshold Value

For the upper thresholds (i.e. h_{facade} and h_{object}) we simply choose values by hand, based on the height of objects that we are interested in. The method is not particularly sensitive to the choice for these values. The main concern is to set values that capture the top part of facades and objects respectively. In this paper the values used are $h_{facade} = 5$m, and $h_{object} = 0.5$m. I.e., we are assuming that facades are higher than 5 meters, and objects are at least half a meter tall.

For the lower thresholds (i.e. h'_{facade} h'_{object} and h_{edge}), ideally the value would be 0, but even after leveling the scene using the tophat, the ground is still not perfectly flat. For simplicity, we set the values of h'_{facade}, h'_{object}, and h_{edge} to the same value of 0.4m. One may select a threshold automatically by examining the histogram of heights of the tophat for values between 0 and h_{object}. The histogram contains a large peak around the height of the ground, which can be used to set the other thresholds. It is possible to push the threshold value down, thus improving the labeling of the lower parts of objects and walls, at the cost of a somewhat noisier ground segmentation.

4.6   Results

The top of Figure 8 shows the segmentation result for `GT_Madame1_2` [9] for different sampling densities of the input (achieved by discarding every n:th sample). The result seems reasonably robust to changes in sampling density. To quantify the performance for different densities we consider the percentage of unchanged labels as more and more input samples are discarded (see bottom left of Fig. 8).

Additionally, we make use of the class labels provided in the dataset [9] to create a ground truth for our labeling by combining the class labels for facade, wall light, balcony plant, and wall sign into a single class label for the walls. Similarly, the ground label is comprised of the background class label and the ground class label, the remaining classes are considered objects. We then look at the percentage of correctly labeled samples for each of these three classes as the sampling density of the input is decreased (bottom right of Fig. 8).

5   Discussion and Conclusions

This is a first attempt at applying irregular morphology to 3D point cloud analysis. A previous morphological approach [9, 11] to the problem studied in this paper first projected the point cloud onto a regularly sampled range image, which causes issues where multiple points in the point cloud correspond to the same pixel in the range image. Working directly on the irregular domain means that there is no need to resample onto a regular grid, thus avoiding this problem. It also allows for morphological operators of any scale, unrestricted by any grid.

We argue that another possible benefit of this approach is the homogenizing tendency of the proposed operators on the point cloud density. This can speed up
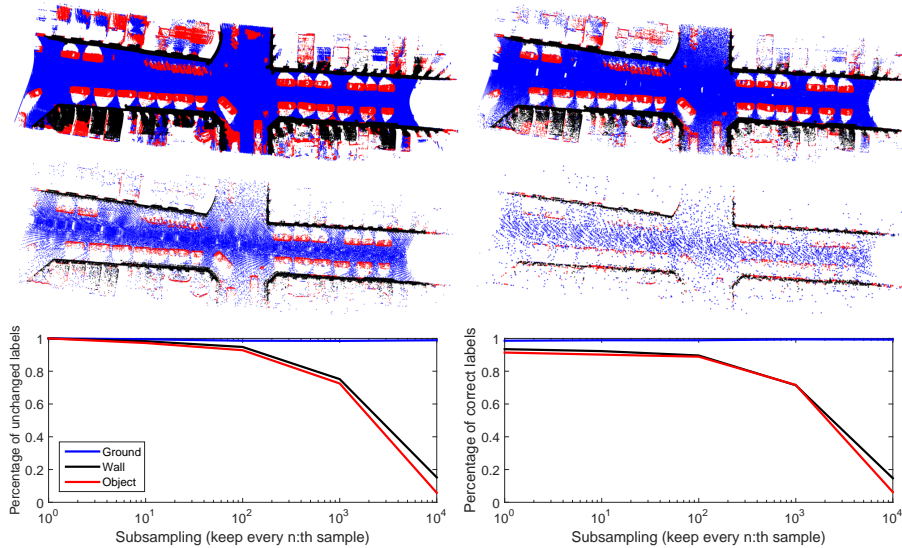
Fig. 8. Top: Labeled scene for different point cloud densities. From left to right and top to bottom, we have: 10^7, 10^6, 10^5, 10^4 samples. Red indicates object, blue indicates ground, and black indicates facade. Bottom left: The percentage of samples in the labeled point cloud, per class, that have changed when comparing less densely sampled inputs to the most densely sampled input. Bottom right: The percentage of correct labels for each class and different point cloud densities. The behavior of the curves for the ground label is explained by its simple nature. In general, the ground is relatively flat and does not need many samples to be well represented.

further processing, since oversampled parts of the signal become sparser, while better defining the surface where the density is low, which may allow for more accurate detection of object boundaries.

In experiments we show that the proposed method is relatively robust, yielding similar labels for the remaining samples as we decrease the input sampling, thus keeping the percentage of correctly labeled samples high even as the density drops to a hundredth. The most robust class is ground, which is to be expected, since large and smooth surfaces do not require a large amount of samples.

Worth noting is that the proposed approach is slower than regular, discrete morphology. Computation time depends on a number of factors, including point density, size of the SE, and size of the input. Processing inputs on the order of 10^6 samples, using a k-d tree to find neighbors, takes seconds for smaller SEs (on an Intel Xeon 3.40GHz), but may take several minutes if the SE is large (e.g. in the case of the large tophat used in our experiments), though further processing is normally faster, since the density is likely decreased, as previously described.

In the future, we would like to improve the sampling strategy. Instead of generating the same number of tentative output samples at each input sample, one could let this number depend on local properties of the signal (e.g. smooth-

ness). Allowing for more general shapes of SEs is also of interest. This may also have an affect on the sampling strategy (e.g. pointy parts of the SE may need a denser sampling, while smoother parts can be sampled more sparsely). With regard to the segmentation approach, comparing our results should be the next step. It would be interesting to take into account some other features. Local thresholding could be useful. Finally, a postprocessing step could likely further improve the segmentation (e.g. via some regularization method).

References

1. Asplund, T., Luengo Hendriks, C.L., Thurley, M.J., Strand, R.: Mathematical morphology on irregularly sampled signals. In: Asian Conference on Computer Vision - DGMM4CV Workshop. pp. 506–520. Springer (2016)
2. Asplund, T., Luengo Hendriks, C.L., Thurley, M.J., Strand, R.: Mathematical morphology on irregularly sampled data in one dimension. Mathematical Morphology-Theory and Applications **2**(1), 1–24 (2017)
3. Calderon, S., Boubekeur, T.: Point morphology. ACM Transactions on Graphics (TOG) **33**(4), 45 (2014)
4. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? The KITTI vision benchmark suite. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2012)
5. Morard, V., Decenciere, E., Dokladal, P.: Geodesic attributes thinnings and thickenings. In: International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing. pp. 200–211. Springer (2011)
6. Munoz, D., Bagnell, J.A., Vandapel, N., Hebert, M.: Contextual classification with functional Max-Margin Markov Networks. In: Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on. pp. 975–982. IEEE (2009)
7. Najman, L., Cousty, J.: A graph-based mathematical morphology reader. Pattern Recognition Letters **47**, 3–17 (2014)
8. Rivest, J.F., Soille, P., Beucher, S.: Morphological gradients. Journal of Electronic Imaging **2**(4), 326–337 (1993)
9. Serna, A., Marcotegui, B.: Detection, segmentation and classification of 3D urban objects using mathematical morphology and supervised learning. ISPRS Journal of Photogrammetry and Remote Sensing **93**, 243–255 (2014)
10. Serna, A., Marcotegui, B., Goulette, F., Deschaud, J.E.: Paris-rue-Madame database: a 3D mobile laser scanner dataset for benchmarking urban detection, segmentation and classification methods. In: 4th International Conference on Pattern Recognition, Applications and Methods ICPRAM 2014 (2014)
11. Serna, A., Marcotegui, B., Hernández, J.: Segmentation of façades from urban 3D point clouds using geometrical and morphological attribute-based operators. ISPRS International Journal of Geo-Information **5**(1), 6 (2016)
12. Serra, J.: Image analysis and mathematical morphology. Academic Press, Inc. (1983)
13. Sibson, R.: A brief description of natural neighbour interpolation. V. Barnett, editor, Interpreting multivariate data (1981)
14. Vosselman, G., Maas, H.G.: Airborne and terrestrial laser scanning. CRC (2010)