(α, δ)-Sleeves for Reconstruction of Rectilinear Building Facets

Marc van Kreveld, Thijs van Lankveld, and Maarten de Rie

Abstract We introduce the concept of (α, δ)-sleeves as a variation on the well-known α-shapes. The concept is used to develop a simple algorithm for constructing a rectilinear polygon inside a plane; such an algorithm can be used to delineate a building facet inside a single plane in 3D from a set of points obtained from LiDAR scanning. We explain the algorithm, analyse different parameter settings on artificial data, and show some results on LiDAR data.

1 Introduction

In recent years, public interest in the use of virtual cityscapes has drastically increased. Applications in a variety of fields like navigation, urban planning, and serious games, increasingly use building models for visualization and simulation purposes. Simultaneously, the quality, complexity, and availability of urban datasets are increasing. Smart-phones are becoming ubiquitous, making photo and video images very easy to obtain, while modern LiDAR devices can capture hundreds of data points per square meter (John Chance Land Surveys and Fugro, 2009).

This wide interest requires efficient automation of urban scene reconstruction to process the vast datasets. The general goal of urban reconstruction is recreating the geometry and visual likeness of the buildings in the scene. Whether applying structure from motion and dense stereo reconstruction to image data (Furukawa et al, 2009; Seitz et al, 2006), or directly using LiDAR data, urban geometry reconstruction is usually aimed at identifying the shapes of buildings from a point cloud.

Earlier methods in photogrammetry would use a predefined collection of parametric models of complete buildings and either try to determine the model that best fits the data (Brenner, 2005; Schwalbe et al, 2005), or only model roofs supported by vertical walls (Rottensteiner, 2003; You et al, 2003; Zhou and Neumann, 2008).

Marc van Kreveld · Thijs van Lankveld · Maarten de Rie
Department of Information and Computing Sciences, Utrecht University, The Netherlands

While these methods are able to construct scenes from very sparse data sets, they are inherently limited by the versatility of the building models in their collection.

Other methods reconstruct free-form triangular meshes that interpolate the data points (Carlberg et al, 2009; Marton et al, 2009; Tseng et al, 2007). While these methods can reconstruct buildings of any shape, most have difficulty dealing with the artifacts inherent in the point data like measurement error and outliers. Additionally, most mesh-based reconstruction methods reconstruct smooth surfaces, removing the sharp edges and simple shapes widely present in urban scenes.

We present a method that is partially parametric and partially free-form to deal specifically with the shapes of urban scenes. While most parametric methods reconstruct the scene per building, we assume that these buildings are composed of planar surfaces and reconstruct these individual surfaces.

Most point clouds measured in an urban scene have inliers and outliers, points measured from a planar surface and the remaining points respectively. The data also contains noise, a small displacement in the point locations. We use Efficient RANSAC (Schnabel et al, 2007) to cluster the points per individual surface, although methods based on region growing (Tseng et al, 2007) could also be used.

Both dense stereo and LiDAR produce point sets densely covering the viewed surfaces. Our method estimates the shape that the points were measured from per individual surface. A surface can be reconstructed by computing a polygon that contains all its points while not containing large empty regions. A popular method for computing such a polygon is the α-shape (Edelsbrunner et al, 1983).

Another prevalent feature of the surfaces in urban data sets is rectilinearity, caused by the predominant use of right angles. In recent work (van Lankveld et al, 2011), we showed that roughly a third of the surfaces in many city scenes are rectangles and we presented a method for reconstructing these surfaces. Here, we broaden this to general rectilinear shapes that tightly bound a point set. To determine these shapes, we present a simple variation of the α-shape, called the (α, δ)-sleeve. This structure creates a buffer around the shape and we search for a rectilinear shape within this buffer. Figure 1 shows an overview of our method.

Finding a rectilinear shape that is close to a given shape is a problem that has been studied in different contexts. For example, restricted-orientation line simplification has been studied for the purpose of schematized map computation (Buchin et al, 2011; Swan et al, 2007; Wolff, 2007). Another example is squarifying, an operation that occurs in ground plan generalization (Mayer, 2005; Regnauld et al, 1999; Ruas, 1999). In these cases, the starting point is a polygonal line or planar subdivision, whereas (initially) we start with a set of points. Furthermore, since sampled points are assumed to be inside the rectilinear polygon to be found, we wish to find an outer approximation of the boundary of the point set. Hence, an area-preserving method like in (Buchin et al, 2011) does not appear suitable. Our method guarantees that the rectilinear polygon is within a specified distance of the α-shape but still outside it; other methods do not have this feature.

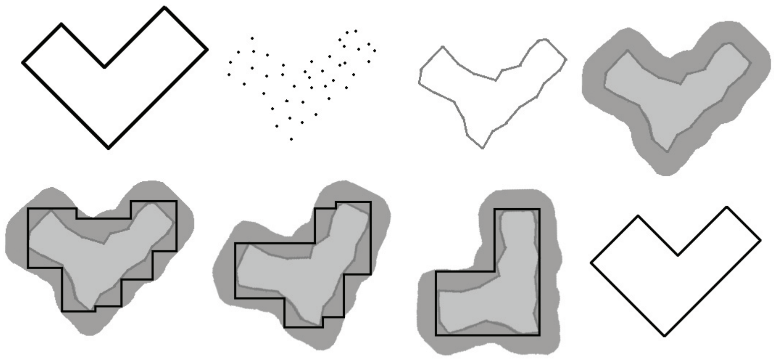

Fig. 1 An overview of our method. From top left to bottom right: the surface, a point sample of the surface, the α-shape, the (α, δ)-sleeve, three rectilinear minimum-link paths within the (α, δ)-sleeve for different rotations, and the path with the fewest links over all rotations

2 (α, δ)-Sleeves and Minimum-Link Paths

In this section we describe the approach for computing a rectilinear polygon that corresponds to the shape of a set of points well. We first define the α-shape (Edelsbrunner et al, 1983), then we introduce a new structure called the (α, δ)-sleeve. We show properties of this new structure and give an efficient algorithm for its construction. Finally, we show how we can use the (α, δ)-sleeve to determine a suitable rectilinear polygon, and give an algorithm to compute it.

Our objective is to bound the point set by a rectilinear shape with few edges. This shape must have all points to the inside or on it, but we must allow the shape to cover some area outside of the α-shape to accomodate a rectilinear shape with few edges.

Like was done for finding rectangles to fit a set S of points in (van Lankveld et al, 2011), we will use the δ-coverage concept. There we defined the δ-coverage region to be the union of the radius-δ disks centered on the points of S. Any point in the plane not in the δ-coverage region is at least at distance δ from all sample points. We required the approximating rectangle to contain all points of the sample, but not be outside the δ-coverage region. The value of δ should be chosen small enough so that the rectangle cannot be too far away from the sampled points and therefore from the likely shape. On the other hand, δ should be chosen large enough so that the irregularities in the sampling do not exclude the existence of a rectangle in the δ-coverage region that encloses the points as well.

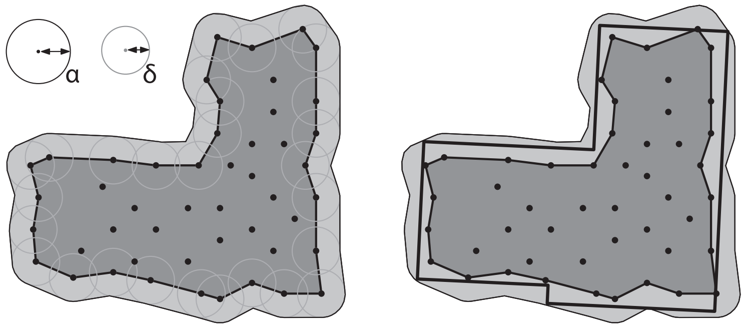

Fig. 2 Left, values of α and δ shown by disks, and the (α,δ)-sleeve of the points shown. Right, a minimum-link rectilinear path in the sleeve

2.1 Definition and Properties of (α,δ)-Sleeves

We adopt these ideas from rectangles to rectilinear shapes. Let S be a set of sampled points in a plane, and let P denote a desired rectilinear polygon. Let $\alpha > 0$ be a real parameter related to the sampling density, typically between 20 and 50 cm. Let $\delta > 0$ be another such parameter, also related to the sampling density.

Definition 1. (Edelsbrunner et al, 1983) Given a point set S, a point $p \in S$ is α-extreme if there exists an empty open disk (i.e., not containing any point from S) of radius α with p on its boundary. Two points $p,q \in S$ are called α-neighbors if they share such an empty disk. The α-*shape of S* is the straight-line graph whose vertices are the α-extreme points and whose edges connect the respective α-neighbors.

We will compute the α-shape A of S and require that P contains A completely. If P was a rectangle, there is no difference between requiring P inside or A inside, but for other rectilinear shapes it can make a difference. Our main reason for using the α-shape is the guarantee that P is not self-intersecting if the α-shape consists of one connected component. Under regular sampling conditions of an individual surface, we can make sure that the α-shape has only one component by carefully choosing α based on the sampling density.

From the α-shape A of S we will compute the (α,δ)-sleeve, defined as follows.

Definition 2. For set S of points in the plane, the (α,δ)-*sleeve* is the Minkowski sum of the α-shape of S with a disk of radius δ, where only the part outside the α-shape is taken.

The Minkowski sum with a disk creates a buffer region around the shape. The (α,δ)-sleeve is an outer proximity region of the α-shape.

Not all values of α and δ, or combinations of values, give nice properties to the (α,δ)-sleeve. Let us assume that α is such that the α-shape is in principle a good approximation of the underlying shape. In particular, let us assume that it is connected and has no holes. This implies that the inner boundary of the (α,δ)-sleeve

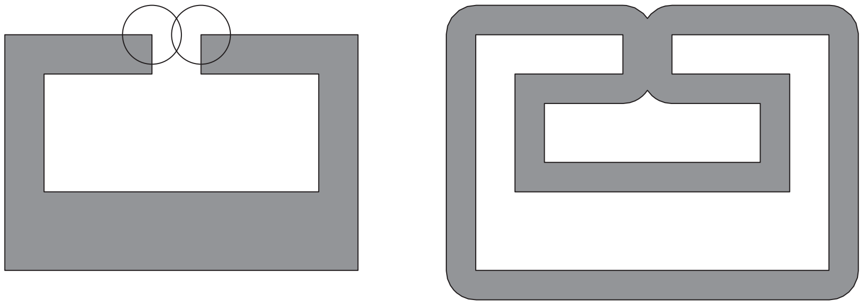

Fig. 3 If the α-shape is the grey interior shown left, and δ corresponds to the radius of the disks shown left, then the (α, δ)-sleeve (shown in grey on the right) has more than one inner boundary

is the boundary of a simple polygon. If δ is sufficiently small, the outer boundary of the (α, δ)-sleeve will be the boundary of a simple polygon as well, but with circular arcs. For some shapes and larger values of δ, the (α, δ)-sleeve can have several inner boundaries: not just those created by subtracting the interior of the α-shape, but also ones where opposite sides on the outside of the α-shape are close. Figure 3 shows an example.

It turns out that if we set $\delta < \frac{4}{5}\alpha$, then the (α, δ)-sleeve will have the desired topology with one outer and one inner boundary. We prove this after proving some more properties of the α-shape.

Lemma 1. *For a given $\alpha > 0$, let A be an α-shape of a set S of points in the plane. Any two points $p, q \in S$ that share the boundary of an empty open disk d of radius $\leq \alpha$ are connected by a line segment inside A.*

Proof. According to Edelsbrunner et al (1983), the interior of the α-shape is the union of all Delaunay triangles with a circumcircle of radius $\leq \alpha$.

The Delaunay triangulation on S must contain the edge p, q as witnessed by d. Now consider the Delaunay triangle t incident to p, q on the side of the center of d; we distinguish two cases. Either (i) this triangle has a circumcircle of radius $\leq \alpha$ and is part of the interior of A, or (ii) there is an empty radius-α disk with p, q on its boundary and p, q are α-neighbors. In both cases, p, q must be either in the interior or on the boundary of A. □

Lemma 2. *For a given $\alpha > 0$, let A be an α-shape of a set S of points in the plane. Any two points $p, q \in S$ at distance at most 2α are connected by a path that lies both inside A and inside the disk with p, q as its diameter*

Proof. In this proof, we use a property of α-shapes that is easily inferred from Lem. 1: for two points $p, q \in S$ at distance at most 2α, either (1) they are α-neighbors, or (2) the connecting line segment is interior to the α-shape, or (3) the disk with p, q as its diameter contains another point of S.

In cases (1) and (2), it is clear there is a connecting path within A and within the disk with p, q as its parameter. In case (3), we consider the collection of empty disks

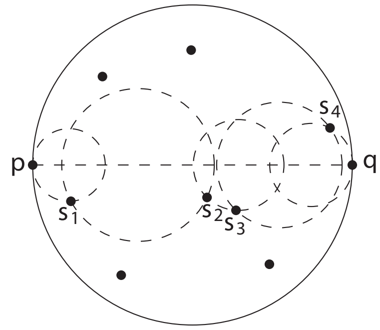

Fig. 4 Illustration of the proof of Lem. 2

with their center on p,q and touching two points of S, as shown in Fig. 4. These disks impose an ordering $(p,s_1,s_2\ldots q)$ on a subset of the points inside the disk. For any two subsequent points x,y in this ordering, the line segment connecting x,y is contained in A, according to Lem. 1. $\square$

Lemma 3. *For a given $\alpha > 0$, let A be an α-shape of a set S of points in the plane, and assume that A is the boundary of a simple polygon. Then for $0 < \delta \leq \frac{4}{5}\alpha$, the (α,δ)-sleeve has the topology of an annulus.*

Proof. Since for very small δ, the topology of the (α,δ)-sleeve is an annulus, we can imagine growing δ to the first (smallest) value δ' where the topology is no longer an annulus. This happens only when there are two points p,q, not necessarily in S, on the α-shape at distance $2\delta'$. Let m be the midpoint of p,q. Then m is at distance δ' from p and q, and no point of the α-shape is closer to m. Hence, the disk Δ centered at m with radius δ' does not intersect the α-shape in points other than p and q.

Assume first that the two points are two vertices $p,q \in S$. According to Lem. 2, Δ contains an edge inside the α-shape. This contradicts the assumption that no point of the α-shape is closer to m than p or q.

Assume next that $p,q \notin S$. Then they lie in the interior of two different α-shape edges. If they are not parallel, it is impossible that the first topological change occurs at m. If they are parallel, then the topological change will occur simultaneously along a stretch of the two edges, and we can choose p and q such that at least one of them is an endpoint and therefore in S. So this case is treated together with the final case.

Assume finally that only one point is in S, say, $p \in S$ and $q \notin S$. Let q_1 and q_2 be the endpoints of the α-shape edge that q lies on, so $q_1,q_2 \in S$. The diametral disk Δ of p and q is tangent to the edge q_1,q_2 at q. Assume without loss of generality that $\overline{q_1q_2}$ is vertical, that q_1 is the lower endpoint of $\overline{q_1q_2}$, that q_1 is closer to p than q_2, and that p is to the right of $\overline{q_1q_2}$, see Fig. 5.

Since the topology of the (α,δ)-sleeve changes for the first time (smallest δ') due to the contact at m, the α-neighbors of p cannot lie to the left of the vertical

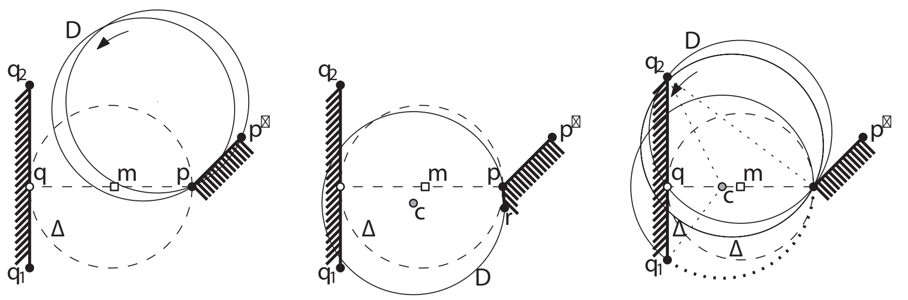

Fig. 5 Illustration of the proof of Lem. 3

line through p. Let p' be the α-neighbor of p that makes the smallest angle with the vertical upward direction. Since p and p' are α-neighbors, an empty α-disk D exists that has p and p' on its boundary and its center to the left of the directed edge from p to p'.

We now rotate D in contact with p in counterclockwise direction, see Fig. 5. Initially D does not contain any point of S inside. Let r be the first point of S that is reached by the boundary of D, such that r would be inside if we were to rotate D further. We distinguish several cases.

If $r = q_2$, then $\overline{pq_2}$ are α-neighbors (as witnessed by the emptiness and current position of D). The implied α-shape edge will intersect Δ because q_2 lies left of p, a contradiction. The same contradiction is obtained when $r = q_1$, or when r is any point that lies left of the vertical line through p. Hence, r lies to the right of this vertical line or on it. This implies that the disk D has rotated beyond the situation where it has a vertical tangent. This is equivalent to stating that the center c of D lies below the horizontal line through p. Also, this center must lie right of the line through q_1 and q_2, because q_1, q_2 is an α-shape edge with the interior to its left.

Because q_1, q_2 are α-neighbors, $\|q_1q_2\| \leq 2\alpha$. We argue that $\|pq_2\| > 2\alpha$. According to Lem. 2, if $\|pq_2\| \leq 2\alpha$, there is a path connecting p, q_2 inside their diametral disk. Because $\|pq_1\| \leq \|pq_2\|$, the same holds for p, q_1. This means that if $\|pq_2\| \leq 2\alpha$, then either there is a point in A closer to m than p, or A is not a simple polygon; both cases contradict an assumtion.

We are interested in the threshold case, where $\|pq\| = 2\delta'$ is as small as possible, and the (α,δ)-sleeve is an annulus for $\delta \leq \delta'$, but not if δ is infinitesimally larger than δ'. Because $\|q_1q_2\| \leq 2\alpha$, $\|pq_2\| > 2\alpha$, and $\|pq_1\| \leq \|pq_2\|$ and $\angle pqq_2 = \frac{\pi}{2}$, the smallest $\|pq\|$ occurs when the angle $\angle q_1q_2p$ is as small as possible. At the same time, the radius-α disk D touching p, r cannot contain q_1 and cannot have its center above $\overline{pq}$. Finally, because r was the first point encounterd by D during its rotation, either $\|pq_1\| \geq 2\alpha$ or c is above $\overline{pq_1}$. If $\|pq_1\| \geq 2\alpha$, then $\|pq\| \geq \sqrt{3}\alpha$ so $\delta' \geq \frac{1}{2}\sqrt{3}\alpha > \frac{4}{5}\alpha$. We continue to show that the other case can give a smaller lower bound for δ. Here $\|pq_1\| < 2\alpha$ and c is above $\overline{pq_1}$. By Lem. 2, A has a sequence of edges connecting p with q_1 via r inside the disk that has $\overline{pq_1}$ as its diameter.

The threshold case is shown in Fig. 5(right). If D stays empty while rotating beyond the point where c lies on p,q, it is possible for the boundary of A to connect q_1 to r and p through a series of edges strictly outside Δ. This would result in a hole in the (α,δ)-sleeve just below m, meaning that the topology of the (α,δ)-sleeve is not an annulus.

In the threshold case, c lies on p,q, and then $\triangle q_2cq_1$ and $\triangle q_2cp$ are mirrored triangles, so the angles $\angle q_2q_1c = \angle q_2pc$, and $\triangle q_2qp$ and $\triangle cqq_1$ are similar triangles. This implies that $\frac{2\alpha}{\alpha} = \frac{\|q_2q\|}{\|cq\|}$ or $\|q_2q\| = 2\|cq\|$. If we combine this with the Pythagorean theorem on $\triangle q_2qp$ we can derive that $2\delta' = \|pq\| = \frac{8}{5}\alpha$:

$$
\begin{aligned}
(2\alpha)^2 &= \|q_2q\|^2 + (\|cq\| + \alpha)^2 \\
4\alpha^2 &= 4\|cq\|^2 + \|cq\|^2 + 2\|cq\|\alpha + \alpha^2 \\
3\alpha^2 - 2\|cq\|\alpha - 5\|cq\|^2 &= 0 \\
(\alpha + \|cq\|)(3\alpha - 5\|cq\|) &= 0 \\
\|cq\| = -\alpha \text{ or } \|cq\| &= \tfrac{3}{5}\alpha
\end{aligned}
$$

The first option leads to a degenerate triangle $\triangle q_1q_2p$. The second option leads to $\|pq\| = \|cq\| + \alpha = \frac{8}{5}\alpha$.

The threshold case presented results in an (α,δ')-sleeve with the topology of an annulus: the value of δ' does not allow D to rotate beyond a vertical tangency at p when it reaches r but before it reaches q_1. Hence, r is not right of p. If r is vertically below p (a degenerate situation), then we can repeat the whole construction with r instead of p, which means we can ignore this case. Hence, in order to get a different topology, r must be strictly right of p, and c must be strictly below p,q. Because D cannot contain q_1, $\|q_1q_2\| \leq 2\alpha$, and $\|pq_2\| > 2\alpha$, this implies that $\|pq\| = 2\delta' > \frac{8}{5}\alpha$. □

We can use known algorithms to compute the (α,δ)-sleeve for given values of α and δ. For a set S of n points in the plane, the α-shape can be computed directly from the Delaunay triangulation of S (Edelsbrunner et al, 1983). Then we use a buffer computation algorithm on the α-shape. Such an algorithm can be based on computing the Voronoi diagram of the line segments of the α-shape first (de Berg et al, 2008; Yap, 1987). Then the buffer boundary can be found in each Voronoi cell, and these can be merged into the boundaries of the δ-buffer of the α-shape. Converting this to the (α,δ)-sleeve is then straightforward. This procedure takes $O(n \log n)$ time in total.

2.2 *Minimum-Link Paths in (α,δ)-Sleeves*

The (α,δ)-sleeve gives a region in which we want to determine a rectilinear shape. The rectilinear shape should separate the inner boundary from the outer boundary.

We will show how to use a minimum-link path algorithm to find the shape. We will assume that the (α, δ)-sleeve has the topology of an annulus.

For any simple polygon and start- and endpoints s and t inside, we can find a minimum-link path that uses only horizontal and vertical edges. This problem has been well-studied in computational geometry, and a linear-time algorithm exists that finds such a path (Hershberger and Snoeyink, 1994). Our problem is different in three aspects: (i) We do not have a simple polygon but a shape with the topology of an annulus. (ii) We do not have a start- and endpoint but we want a rectilinear cycle. (iii) We do not know the orientation of the edges beforehand, we only know that the angles on the path are 90 degrees.

Lemma 4. *In an (α, δ)-sleeve, there exists a minimum-link axis-parallel cycle that passes through the lowest point of the inner boundary of the sleeve (the α-shape).*

Proof. Consider any minimum-link cycle in the sleeve, and let e be its lowest horizontal edge. If e contains a vertex of the α-shape the lemma is true, otherwise we can move e upwards while shortening the two adjacent edges of the cycle. During this move two things can happen: (i) An adjacent edge reduces to length 0, but then a cycle with two fewer links is found. (ii) Edge e is stopped by a vertex of the α-shape, which must be its lowest vertex because all vertices of the α-shape have a y-coordinate at least as high as the lowest edge of the cycle. This proves the lemma. □

Another property of the minimum-link rectilinear path is that it is non-selfintersecting. Otherwise, we could remove the extra loop and obtain a cycle with fewer links.

Lemma 5. *Any minimum-link rectilinear path in an (α, δ)-sleeve is non-selfintersecting.*

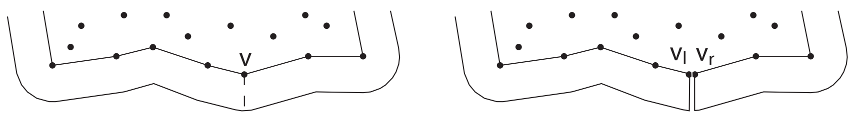

Fig. 6 The bottom part of an (α, δ)-sleeve, the lowest vertex v of the α-shape (left), and the conversion of the sleeve to a simple polygon (right)

Suppose that we know the orientation of the minimum-link path. Then we can rotate the (α, δ)-sleeve so that this orientation becomes the axis-parallel orientation. By the lemma above, we now know a point that we can assume to lie on the minimum-link cycle. To convert the (α, δ)-sleeve to a polygonal region we do the following. We find the lowest vertex v of the α-shape and insert a new edge vertically down from v, until it reaches the outer boundary, see Fig. 6. This edge will split the annulus into a shape with the same topology as a simple polygon. We duplicate the new edge including its endpoints, splitting v into a left copy v_l and a right copy v_r. Now our shape does not have doubly used edges, and it can be treated as a normal simple polygon. We will find a minimum-link axis-parallel path from v_r to v_l using the algorithm in (Hershberger and Snoeyink, 1994). The output path can easily be converted back into a minimum-link cycle.

3 Experiments

In the previous section we presented a method to construct a rectilinear minimum-link path that approximates the α-shape. Here, we evaluate this method on synthetic and real data with the main goal of checking whether the method is suitable for piecewise planar urban scene reconstruction. In particular, we wish to discover whether for a given point density, values of α and δ exist that lead to a rectilinear minimum-link path that is close to the true building facet shape.

This section describes both the setup of these experiments and the results. The first two subsections describe our experiments on synthetic data to determine a value for δ to get the correct number of edges or the best overlap with the ground truth, respectively. The last subsection describes an experiment on urban LiDAR data.

3.1 Universal δ

For the evaluation on synthetic data, we have constructed fifteen test cases of varying shape and complexity. Each test case comprises a ground truth polygon T that should represent a realistic urban surface shape with an area of between 30 and $60\,\mathrm{m}^2$, as shown in Fig. 7. To counter bias towards a certain initial orientation, T is rotated by a random real-valued angle for each test. Each rotated T then yields a point set of predetermined density by uniform sampling from its interior.

For each case, we compute the (α, δ)-sleeve using a fixed α based on the sampling density ρ and a varying δ. We have chosen a fixed α such that the α-shape is connected and without holes irrespective of test case. We vary δ to determine whether there is a single value of δ at each ρ such that our method produces polygons similar to the ground truth in all test cases. To measure the impact of ρ on the results of our method, we have repeated the experiments for three different densities. The chosen α for each density ρ is shown in Table 1. We vary δ using increments of 1 *cm*.

ρ (points/m^2)	α (cm)
50	60
100	20
200	17

Table 1 The different densities used for the point sampling and the associated values of α

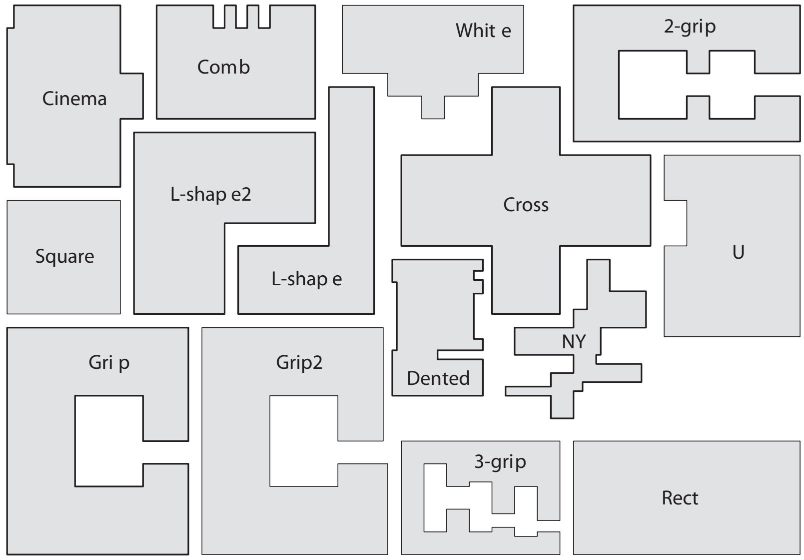

Fig. 7 The different test cases used for the experiments on synthetic data

A rectilinear polygon $\hat{P}$ is constructed inside the (α, δ)-sleeve, according to the algorithm given in Sect. 2, such that $\hat{P}$ has the minimum number of edges over all rotation angles. We use an angular step size of 1 degree, meaning that we run the minimum-link algorithm 90 times for each (α, δ)-sleeve. To get a canonical result for each angle, we 'shrink' $\hat{P}$ to a smaller polygon P by moving each edge inward until it touches the α-shape. We move edges in descending order of edge length. If there are multiple polygons with the minimum number of edges, created for different rotation angles, P is chosen as the one with the smallest area.

In urban reconstruction, the shape of a surface is not known a priori. For this reason, we have analyzed whether our method reconstructs the correct shapes. We measure the correctness of a shape by its number of edges and its angle of rotation, by comparing them to the ground truth. In the optimal case, there is a value for δ at which our method always produces the correct shape. However, it is very likely that the optimal value of δ depends on ρ. Because we want our method to be shape-invariant, we are looking for a δ for which our method performs well, irrespective of ground truth case.

Our experiments showed that the chosen polygons P always have a rotation within 1 degree of the rotation of the ground truth. Figure 8 shows the ranges of δ values for which the constructed polygon has the same number of edges as the ground truth. At the higher densities (100 and 200 points/m^2), we can choose the value of δ at 24 cm and 20 cm respectively to construct a shape with the same number of edges for most cases.

At the lowest density (50 points/m^2), there is no δ value that consistently results in a correct polygon. Additionally, Figure 8 shows that in half of the cases the range

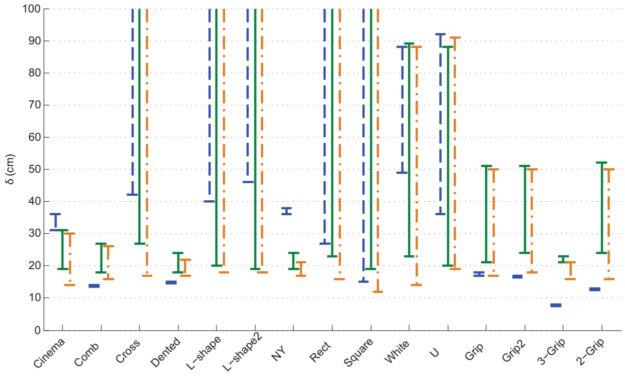

Fig. 8 The range of δ for which the constructed polygon has the same number of edges as the ground truth at point densities 50 (dashed), 100 (solid), and 200 (dash-dotted) points/m^2

of δ resulting in the correct number of edges is very small. These cases all have some small features that add edges to the ground truth while being difficult to make out in the rough point samples. Visual inspection showed that the openings of the "grip" cases were closed off in their α-shape, explaining the problems with constructing the correct shape. For the cases without a small feature, a value of δ between 50 and 85 results in the correct polygons.

Considering the values of α and δ together, we observe that we should choose δ slightly larger than α. However, by Lem. 3, we are not guaranteed to get an (α,δ)-sleeve with the topology of an annulus in this case. Especially in "grip" cases, finding ways to correctly deal with (α,δ)-sleeves that do not have an annulus topology may yield better results.

3.2 Data-Dependent δ

The shapes encountered in urban scenes vary greatly. Small features of the shape combined with insufficient sampling density may make it very difficult to correctly estimate the number of edges of the shape, even for a human modeler. Additionally, the sampling density may vary greatly between surfaces. For this reason, it may be interesting to estimate the best value for δ from the point data itself instead of choosing some fixed value.

One way to determine which δ is best is to analyze how P changes as δ increases. At the smallest δ, P will approximate the α-shape and from some large δ onwards, P is a rectangle. Recall that P is the shrunken version of $\hat{P}$, so we can expect the largest changes in P when its number of edges change. Additionally, given a fixed number of edges, it is likely that the polygon that approximates T best is found immediately after a jump to that number of edges. While growing δ, we use the polygon just after each jump as representative for the polygons with the same number of edges. This leads to a succession of representative polygons $\{R_0 \dots R_k\}$, where each consecutive polygon has fewer edges, culminating in four edges at R_k.

The best R should strike a balance between complexity and approximation of the shape of T. The complexity can again be measured in the number of edges, with a preference for low complexity. How well R approximates T could be measured from their area of symmetric difference. Unfortunately, for real-world data the ground truth is not known, so the symmetric difference cannot be used to determine the best δ. However, because the area of T is fixed, a simple approximation for the symmetric difference is to use the area of R.

Because R always covers the α-shape, we can assume that the reason for large jumps in the area of R is the removal of a large feature of T from R. The goal then becomes to determine when the process of increasing δ stops reducing the complexity of the shape and starts removing large features. If we denote by χ_i the change in area between consecutive representative polygons R_i and R_{i+1}, we search for a threshold value $\bar{\chi}$ for χ_x. The idea is that if the area of R makes a jump of $\chi_i > \bar{\chi}$ due to a decrease in the number of edges, then the polygon loses a key feature of T so R_i is the simplest polygon that contains all important features.

We determine $\bar{\chi}$ by selecting which reference polygon R_i best approximates T by visual inspection and computing χ_i. Choosing $\bar{\chi}$ such that it is smaller than χ_i but larger than χ_j for all $j < i$ would result in terminating the automatic search for δ at the preferred polygon. Because we want the method to be shape-invariant, we are looking for a value of $\bar{\chi}$ for which our method performs well, irrespective of ground truth case.

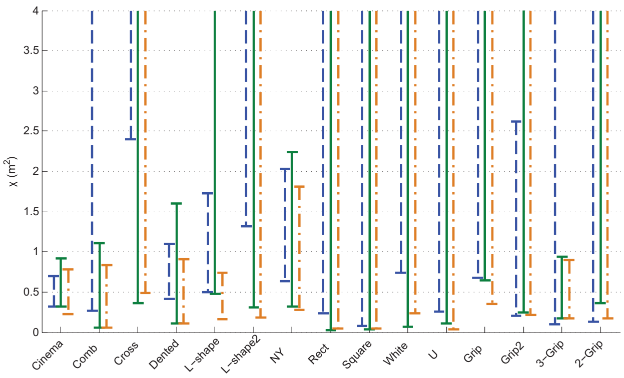

Fig. 9 The range of $\bar{\chi}$ for which the polygon selected by visual inspection is constructed at point densities 50 (dashed), 100 (solid), and 200 (dash-dotted) points/m^2

The ranges of $\bar{\chi}$ resulting in the polygon selected by visual inspection are shown in Fig. 9. At the higher densities (100 and 200 points/m^2), there are values for $\bar{\chi}$ that result in the correct polygon in all the test cases. At a density of 50 points/m^2, there is no $\bar{\chi}$ that produces the correct polygon in all cases, although 0.8 m^2 is a good value for most cases.

3.3 LiDAR Data

Apart from synthetic data, we have also applied our method to an airborne LiDAR data set of a building, as shown in Fig. 10. The point density varies greatly between surfaces, because of the scanning method. However, based on our earlier work on the same data (van Lankveld et al, 2011), we set α to 125 cm for surfaces close to vertical and 60 cm for the other surfaces.

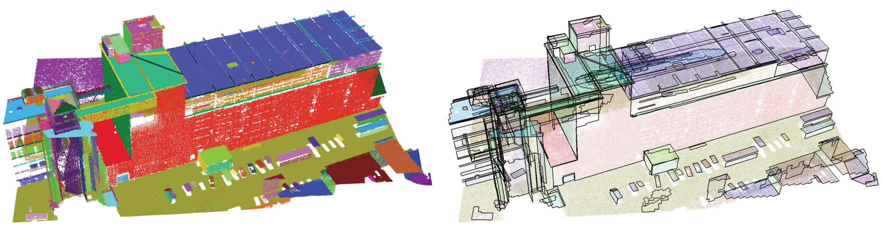

Fig. 10 A LiDAR data set with points colored per surface and the rectilinear boundaries of those surfaces

The building contains many rectilinear surfaces, which we have reconstructed using $\delta = 60$ cm. The surfaces near the edge of the scene were given jagged edges because the buildings are not aligned with the scene. However, our results do favour long straight edges for the remainder of the shape.

Figure 11 shows two interesting surfaces and their rectilinear boundaries. Note how the rotation of the reconstructed polygons matches the neighboring surfaces. A human modeler may construct a similar shape with fewer edges for these surfaces. It seems that these faults in our results are caused by missing data in the input sets.

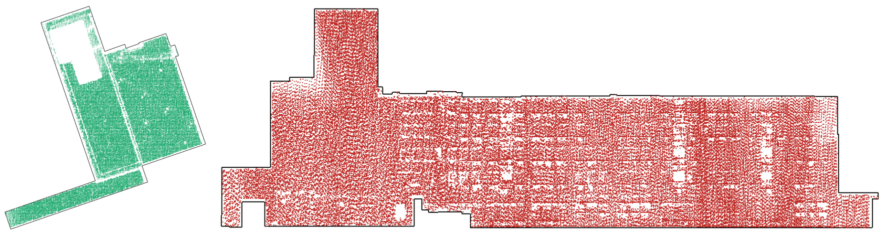

Fig. 11 Two interesting rectilinear surfaces

4 Conclusions

We have presented a novel concept, the (α, δ)-sleeve, which contains a proximity zone around a point set. When combined with a minimum-link path algorithm, this structure can be used to reconstruct simple shapes that contain the point set. We have presented a method aimed at reconstructing surfaces in urban scenes from a point set by combining the (α, δ)-sleeve with a rectilinear minimum-link path. Our experiments showed that when the surfaces are sampled sufficiently dense, there are parameter settings for α and δ that lead to correct reconstructions for artificial data. Finally, we have shown on an urban LiDAR data set that our method produces plausible results.

A number of interesting possibilities for extensions and improvements remain. In most urban scenes there are some surfaces that are not rectilinear. By changing the rectilinear minimum-link path algorithm to allow a few edges that do not follow one of the principal directions, the (α, δ)-sleeve can be used to reconstruct such surfaces as well. Alternatively, we could use a post-processing method that replaces long stair-like parts in the rectilinear surface boundary by one line.

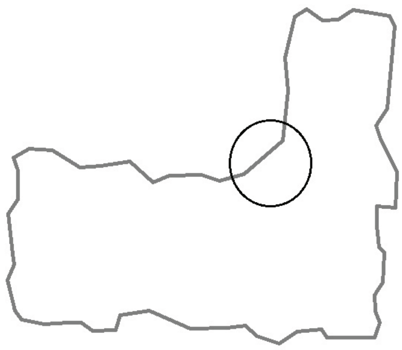

Fig. 12 An example α-shape where a concave corner is rounded off

Another recurring problem we encountered is the rounding of concave corners. As Figure 12 shows, the α-shape can "round off" a concave corner, often going outside the original polygon. In such cases, constructing a polygon within the (α, δ)-sleeve requires either more edges or a larger δ. This problem may be overcome by using pieces of the α-hull (Edelsbrunner et al, 1983) as inner boundary of the (α, δ)-sleeve, instead of the α-shape. The α-hull uses concave circular arcs between its boundary vertices, allowing the polygon to go deeper into concave corners. Unfortunately, using α-hull arcs everywhere may cause the inner boundary to self-intersect, which in turn can result in a self-intersecting rectilinear path. Hence, this solution would require additional steps to ensure that the constructed polygon does not have self-intersections.

Finally, for this work we have ignored holes in the α-shape. Reconstructing those holes can be done in a fashion similar to the method described in this paper. The (α,δ)-sleeve would consist of several components, each one with the topology of an annulus. We can run minimum-link path algorithms in each annulus with the same orientations in order to find an orientation that is best overall. This way we can easily ensure that the rectilinear directions of the outside boundary and of the hole boundaries are the same.

Acknowledgements

This research has been supported by the GATE project, funded by the Netherlands Organization for Scientific Research (NWO) and the Netherlands ICT Research and Innovation Authority (ICT Regie).

References

de Berg M, Cheong O, van Kreveld M, Overmars M (2008) Computational geometry – algorithms and aplications, 3rd edn. Springer-Verlag, Berlin

Brenner C (2005) Building reconstruction from images and laser scanning. Int J Appl Earth Obs 6(3–4):187–198

Buchin K, Meulemans W, Speckmann B (2011) A new method for subdivision simplification with applications to urban-area generalization. In: Proc. SIGSPATIAL, pp 261–270

Carlberg M, Andrews J, Gao P, Zakhor A (2009) Fast surface reconstruction and segmentation with ground-based and airborne lidar range data. Tech. Rep. ADA538860, University of California at Berkeley

Edelsbrunner H, Kirkpatrick DG, Seidel R (1983) On the shape of a set of points in the plane. IEEE Trans on Inf Theory 29(4):551–559

Furukawa Y, Curless B, Seitz SM, Szeliski R (2009) Manhattan-world stereo. IEEE Computer Society, pp 1422–1429

Hershberger J, Snoeyink J (1994) Computing minimum length paths of a given homotopy class. Comp Geom 4:63–97

John Chance Land Surveys, Fugro (2009) Fli-map specifications. http://www.flimap.com/site47.php

van Lankveld T, van Kreveld M, Veltkamp RC (2011) Identifying rectangles in laser range data for urban scene reconstruction. Comput Graph 35(3):719–725

Marton ZC, Rusu RB, Beetz M (2009) On fast surface reconstruction methods for large and noisy point clouds. In: Proc. ICRA, pp 2829–2834

Mayer H (2005) Scale-spaces for generalization of 3D buildings. Int J Geogr Inf Sci 19:975–997

Regnauld N, Edwardes A, Barrault M (1999) Strategies in building generalisation: modelling the sequence, constraining the choice. In: Proc. ICA

Rottensteiner F (2003) Automatic generation of high-quality building models from lidar data. CG&A 23(6):42–50

Ruas A (1999) Modèle de généralisation de données géographiques à base de contraintes et d'autonomie. PhD thesis, Université de Marne la Vallée

Schnabel R, Wahl R, Klein R (2007) Efficient RANSAC for point-cloud shape detection. Computer Graphics Forum 26(2):214–226

Schwalbe E, Maas HG, Seidel F (2005) 3D building model generation from airborne laser scanner data using 2D GIS data and orthogonal point cloud projections. In: Proc. ISPRS, pp 12–14

Seitz S, Curless B, Diebel J, Scharstein D, Szeliski R (2006) A comparison and evaluation of multi-view stereo reconstruction algorithms. In: Proc. CVPR, vol 1, pp 519–528

Swan J, Anand S, Ware M, Jackson M (2007) Automated schematization for web service applications. In: Web and Wireless GISystems, LNCS 4857, pp 216–226

Tseng YH, Tang KP, Chou FC (2007) Surface reconstruction from LiDAR data with extended snake theory. In: Proc. EMMCVPR, pp 479–492

Wolff A (2007) Drawing subway maps: A survey. Inform Forsch Entwickl 22(1):23–44

Yap CK (1987) An $O(n \log n)$ algorithm for the Voronoi diagram of a set of simple curve segments. Discrete Comput Geom 2:365–393

You S, Hu J, Neumann U, Fox P (2003) Urban site modeling from LiDAR. In: Proc. ICCSA, LNCS, vol 2669, pp 579–588

Zhou QY, Neumann U (2008) Fast and extensible building modeling from airborne LiDAR data. In: Proc. SIGSPATIAL, pp 1–8