

Polylidar - Polygons from Triangular Meshes*

Jeremy Castagno¹ and Ella Atkins²

Abstract—This paper presents Polylidar, an efficient algorithm to extract non-convex polygons from 2D point sets, including interior holes. Plane segmented point clouds can be input into Polylidar to extract their polygonal counterpart, thereby reducing map size and improving visualization. The algorithm begins by triangulating the point set and filtering triangles by user configurable parameters such as triangle edge length. Next, connected triangles are extracted into triangular mesh regions representing the shape of the point set. Finally each region is converted to a polygon through a novel boundary following method which accounts for holes. Real-world and synthetic benchmarks are presented to comparatively evaluate Polylidar speed and accuracy. Results show comparable accuracy and more than four times speedup compared to other concave polygon extraction methods.

I. INTRODUCTION

Video and LiDAR data are widely used in robotics to provide rich information about the environment. LiDAR and RGBD cameras generate point clouds for localization and mapping [1], 3D modelling [2], and scene classification for autonomous navigation [3]. Flat surfaces such as walls and floors are key environmental elements to identify; they are often extracted using planar segmentation techniques [4], [5], [6]. However points clouds are dense incurring a high computational cost when used directly. A common simplifying approach transforms point clouds into lower dimensional representations such as lines and planes [7]. Furthermore, polygonal representations of planes reduces map size and may accelerate matching for localization [8]. Convex polygon representations of planar segments were proposed by [7]. Convex polygons are simple and efficient to generate but ignore boundary concavities and overestimate area of the enclosed point set. Non-convex polygon representations may be generated using techniques such as boundary following outlined in [8] or α -shapes as proposed in [9]. However few methods also capture the interior holes within non-convex polygons. Safe robot navigation demands accurate capture of non-convex polygons with interior holes in real-time, requiring both speed and robustness.

This paper presents Polylidar, an efficient algorithm to transform 2D point sets into simplified non-convex (i.e. concave) polygons with holes. Polylidar begins by triangulating the point set and filtering triangles given user-specified parameters such as maximum triangle edge length. Once filtering is complete, edge-connected triangles are combined into regions creating a set of triangular meshes representing

the shape of the point set. Next, Polylidar converts each mesh region to a polygon through a novel boundary following method which accounts for holes. Figure 1b shows Polylidar applied to a 2D point set while (c) shows Polylidar used on a plane segmented point cloud from an RGBD image.

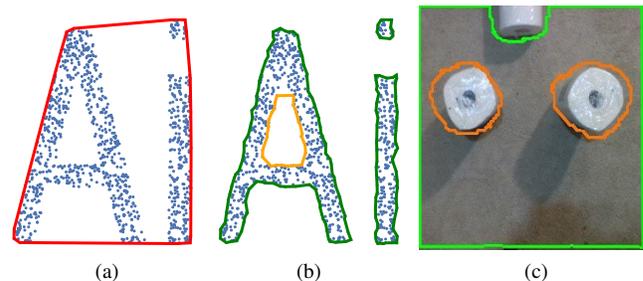


Fig. 1. (a) Convex hull of a point set (red); (b) MultiPolygon extraction using Polylidar (green). (c) Polygon extraction from a plane segmented point cloud from an Intel RealSense RGBD camera capturing paper towel rolls on a basement floor. Note that Polylidar also identifies holes (orange).

We show that the Polylidar algorithm is approximately 4 times faster than leading open source approaches for concave polygon extraction. Polylidar’s speed is attributed to rapidly identifying boundary edges (shell and holes) and then performing boundary following to ensure a valid polygon is returned. Contributions of this paper are:

- A faster open source [10] concave (multi)polygon extraction algorithm from 2D point sets.
- A benchmark comparison of leading concave polygon extraction techniques in terms of accuracy and speed.

Below, Sections II and III provide background on non-convex shape generation and mathematical preliminaries, respectively. Section IV describes Polylidar algorithms, while Section V shows benchmark test results of Polylidar versus other methods. Section VI describes test results. Sections VII and VIII provide discussion and conclusions.

II. BACKGROUND

Characterizing the shape of a set of 2D points \mathcal{P} has been a long-term focus of computational geometry research. A convex hull is defined as the smallest convex polygon that fully encapsulates all points in a set \mathcal{P} . Although widely used to estimate shape, point sets with non-convex distributions are poorly characterized by a convex hull [11]. Convex hull over-estimation can be a serious issue when the points represent physical objects, e.g., obstacle free navigable areas. Several algorithms have been developed to construct shapes that “fit” or “cover” point sets more closely.

*This work was supported in part by NSF I/UCRC Award 1738714.

¹Jeremy Castagno is a Robotics Institute PhD candidate, University of Michigan, jdcasta@umich.edu

²Ella Atkins is a Professor of Aerospace Engineering and Robotics, University of Michigan, ematkins@umich.edu

Figure 1 compares convex and concave hulls. Figure 1b is the multipolygon output of PolyLidar described below. While there is a unique convex hull, there is no true or unique concave hull. Concave hull algorithm implementations can also have different output types. Some return only an unordered set of edges while others return a single polygon. Some algorithms return multiple disconnected polygons (multipolygon), and some can generate holes inside a polygon.

The α -shape algorithm is an early strategy to generate a family of shapes ranging from a convex hull to a point set [12]. The parameter α dictates the radius of a closed disk used to prune/remove area in the convex hull. This disk is allowed to move freely shaving off the excess shape until it finds points. When disk radius is large, ideally infinite, the convex hull is produced; when disk radius is infinitesimally small only the points remain. A common implementation of α -shape organizes points using Delaunay triangulation and filters triangles whose circumcircle radius is less than α . The final shape is represented by the remaining edges and triangles. Note that the α -shape method creates multiple non-intersecting shapes with the possibility of holes.

The geospatial software library Spatialite [13], an extension to SQLite [14], contains a concave hull extraction procedure. The algorithm again starts with Delaunay triangulation then analyzes the distribution of each triangle’s edge length to determine mean μ_l and standard deviation σ_l . Any triangle with edge length greater than $C \cdot \sigma_l + \mu_l$ is removed, where C is a user-defined parameter. The final geometry returned is the union of all triangles computed with GEOS, a high performance open source geometry engine. The output may be a multipolygon (i.e., multiple disjoint polygons) with the possibility of holes inside each.

PostGIS is a geospatial database of computational geometry routines such as the concave hull method in [15]. This algorithm first calculates the convex hull and then shrinks the hull by adjusting vertex connections to closer points which “cave in” the hull. This process recursively shrinks a boundary until a user-specified percent reduction in area from the convex hull is achieved. The resulting shape is a single polygon with the possibility of holes.

TABLE I
CONCAVE HULL EXTRACTION METHODS

Algorithm	Output	Holes?
CGAL α -shape	unordered set of edges	Yes
Spatialite	(multi)polygon	Yes
PostGIS	polygon	Yes
PolyLidar (new)	(multi)polygon	Yes

Table I provides a summary of the concave hull algorithms discussed above. The Computational Geometry Algorithms Library (CGAL) is used as the implementation of the α -shape method [16]. Note that the time complexity of all algorithm implementations, with the exception of PostGIS, is $\mathcal{O}(n \log n)$. Our paper contributes a procedure to more rapidly compute (multi)polygon output with the possibility of holes. Though this is a complex output to generate, we show

through benchmarks that our algorithm and implementation outperforms other available approaches.

III. PRELIMINARIES

A 2D *point set* is an arbitrarily ordered set of two dimensional points in a Cartesian reference frame. Each point is defined by orthogonal bases \hat{e}_x and \hat{e}_y with

$$\vec{p}_i = x \hat{e}_x + y \hat{e}_y = [x, y] \quad (1)$$

where x, y are plane coordinates.

An n -point array $\mathcal{P} = \{\vec{p}_1, \vec{p}_2, \dots, \vec{p}_n\}$ contains points $\vec{p}_i \in \mathbb{R}^2$ indexed by i . A triangular mesh \mathcal{T} is defined by

$$\mathcal{T} = \{t_1, t_2, \dots, t_k\} \quad (2)$$

where each t_i is a triangle with vertices defined by three point indices $\{i_1, i_2, i_3\} \in [1, n]$ referencing points in \mathcal{P} .

We follow the Open Geospatial Consortium (OGC) standard [17] for defining *linear ring* and *polygon*. A linear ring is a consecutive list of points that is both closed and simple. This requires a linear ring to have non-intersecting line segments that join to form a closed path. The key components of a valid polygon are a single exterior linear ring representing the *shell* of the polygon and a set of linear rings (possibly empty) representing *holes* inside the polygon.

IV. METHODS

Sections IV-A, IV-B, and IV-C describe the triangulation data structures, filtering, and mesh extraction respectively. Section IV-D describes polygon extraction.

A. Triangulation with Half-Edge Decomposition

PolyLidar begins with the Delaunator library [18] performing a Delaunay triangulation of point set \mathcal{P} . The original algorithm was written in JavaScript but a C++ port of the library is used in PolyLidar [19]. Note that we have modified Delaunator to use robust geometric predicates to ensure correctness during triangulation [20]. Delaunator was chosen for its ease of integration, speed, and output data structure which returns a *half-edge* triangulation. A half-edge triangulation decomposes a shared edge using two half-edges $A \rightarrow B$ and $B \rightarrow A$. An example of this decomposition and resulting data structures is shown in Figure 2.

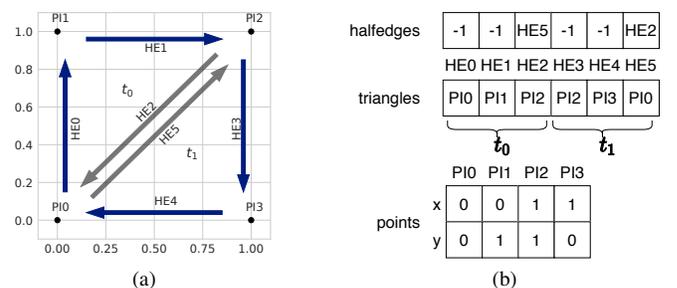


Fig. 2. (a) Triangulation of a square point set using Delaunator [18] with output data structure indexed by half-edge ids in (b). HE=half-edge, PI=point index, t =triangle. Grey edges show shared edges decomposed individually.

Figure 2a triangulates point set $\{PI0, PI1, PI2, PI3\}$. Triangulation produces two triangles, t_0 and t_1 , with half-edges $\{HE0, HE1, HE2\}$ and $\{HE3, HE4, HE5\}$, respectively. Each half-edge supports clockwise travel to the next half-edge in that triangle's edge set. Figure 2b lists the resulting *halfedges*, *triangles*, and *points* data structures. The *halfedges* array is indexed by a half-edge reference id. It provides the opposite half-edge of a shared edge if it exists; otherwise -1 is returned. The *triangles* array is also indexed by half-edge id and gives the starting point index of the associated half edge. The relationship between half-edge and triangle indices is $t = \text{floor}(he/3)$.

B. Triangle Filtering

As with Spatialite and α -shape methods the initial shape starts with k triangles in \mathcal{T} per Eqn. 2 returned from Delaunay triangulation. Also similar to α -shape and Spatialite methods, PolyLidar filters triangles by configurable criteria for each triangle. PolyLidar allows the user to perform triangle filtering using either the α parameter or maximum triangle edge length parameter l_{max} . The filtered triangle set is denoted \mathcal{T}_f .

C. Triangular Mesh Region Extraction

An iterative plane extraction procedure inspired from [21] generates subsets of \mathcal{T}_f that are spatially connected. These subsets are denoted \mathcal{T}_r which represent triangular mesh regions. A spatial connection between triangles exists when they share an edge. A random seed triangle is selected from \mathcal{T}_f where a new region is created and expanded by its adjacent edge neighbors from the *halfedges* data structure. Region growth halts when no more triangles in \mathcal{T}_f connect to the region. The process repeats with another seed triangle until all triangles in \mathcal{T}_f have been examined.

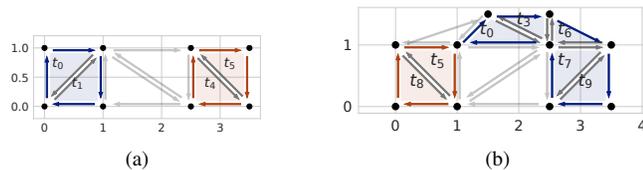


Fig. 3. (a) Example of two regions extracted denoted by orange and blue. Triangles t_0 and t_1 are one region while t_4 and t_5 are another. (b) Two regions are also extracted even with a shared vertex.

Figure 3a shows triangular mesh region examples. Distinct regions are shown in orange and blue; light grey edges denote triangles that have been filtered out. The output of this step is a set of spatially connected triangular mesh regions, \mathcal{T}_R , where each specific region, $\mathcal{T}_{r,i}$, is a set of triangle indices. We denote the set of m triangular mesh regions as:

$$\mathcal{T}_R = \{\mathcal{T}_{r,1}, \mathcal{T}_{r,i}, \dots, \mathcal{T}_{r,m}\} \quad (3)$$

$$\mathcal{T}_{r,i} = \{t_i, \dots, t_j\} \quad (4)$$

D. 2D Polygon Extraction

Polygon extraction has three steps: data structure initialization, concave shell extraction, and hole(s) extraction. Each of these steps is described below. Note that polygon extraction is independent of the specific triangular mesh regions $\mathcal{T}_{r,i}$, thus subsequent notation will drop the i index for brevity when used in algorithms. The following steps are executed for each of the m regions in \mathcal{T}_R to generate m polygons.

1) *Data structure initialization*: Data structure initialization is shown in Algorithm 1 which produces three data structures: a boundary half-edge set, a point index hash map, and the extreme point. A visual example of these data structures is shown in Figure 4. Boundary half-edge set \mathcal{HE} contains the half-edge indices that are on the exterior border of a region, marked in blue in Figure 4a. A half-edge is marked as a boundary if it has no opposite half-edge (meaning it is on the convex hull of the full triangulated set) or if its adjacent triangle is not in $\mathcal{T}_{r,i}$. The last check is important because a half-edge may share an edge with an interior triangle that is not part of $\mathcal{T}_{r,i}$ as seen in the rightmost edge for the blue region in Figure 3a. The *halfedges* data structure is fixed at triangulation and is not aware of filtered triangles or the regions discussed in Section IV-C.

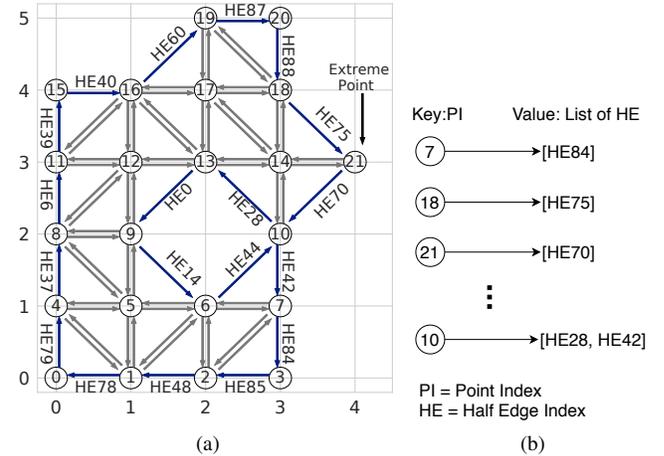


Fig. 4. (a) The boundary half-edge set is marked in blue and point index 21 (PI21), the farthest point on the x-axis, is noted. (b) A sample of the resulting point index hash map, *PtE* is shown. Note that the display order has been arbitrarily chosen.

The second data structure is a point index hash map, *PtE*, whose key is a point index and value is a list of outgoing boundary half-edges from the keyed point index. This *unordered* hash map is represented in Figure 4b; note the keyed point index 7 mapping to the single element list containing half-edge 84. The final data structure represents an extreme point in the triangle mesh, referring to the point farthest to the right on the x -axis. This point will be used as the starting point index when extracting the concave hull to help ensure extraction does not start on a hole edge. Multiple points may exist on the extreme edge; the algorithm will track the first one found in this case.

2) *Concave Shell Extraction*: Outer shell extraction begins by traversing the half-edge graph, starting with the

Algorithm 1: Initialize

Input : Triangular Mesh Region , $\mathcal{T}_r = \{t_i, \dots, t_k\}$
Shared Halfedges, *halfedges*
Triangles Point Index, *triangles*

Output: Half Edge Set , $\mathcal{HE} = \{he_i, \dots, he_n\}$
Point Index Hash Map, *PtE*
Extreme Point, pi_{xp}

```
1  $\mathcal{HE} = \emptyset$ ; // boundary half-edge set
2  $PtE = \emptyset$ ; // Point to half-edge hashmap
3  $pi_{xp} = 0$ ; // will be overwritten
4 for  $t_i \in \mathcal{T}_r$  do
5   for  $he_i \in t_i$  do
6      $he_j = halfedges[he_i]$ ; // opposite edge
7      $t_j = floor(he_j/3)$ ; // adjacent tri
8     if  $t_j \notin \mathcal{T}_r$  then
9        $\mathcal{HE} = \mathcal{HE} + he_i$ ; // boundary edge
10       $pi = triangles[he_i]$ 
11       $pi_{xp} = TrackXp(pi, pi_{xp})$ 
12      if  $pi \notin PtE$  then
13        /* create half-edge list */
14         $PtE[pi] = [he_i]$ 
15      else
16        Append( $PtE[pi], he_i$ )
17    end
18 end
19 return  $\mathcal{HE}, PtE, pi_{xp}$ 
```

half-edge provided by the extreme point. As the edges are traversed the point indices are recorded in a list representing the linear ring of the concave hull. Edges are removed from the half-edge set, \mathcal{HE} , as they are traversed. In Figure 4a the extreme point index is PI21 and the starting half-edge is HE70. This starting half edge and start point index are arguments to the `ExtractLinearRing` procedure in Algorithm 2, with the procedure halting when edge traversal returns back to the starting point index, indicating a closed linear ring has been extracted. The hole in this shape, represented by edges (HE28, HE0, HE14, HE44), with a shared vertex at PI10, must be carefully handled as explained below. This is an example of a non-manifold mesh.

The example in Figure 4 begins with HE70 traversing to PI10. The outgoing boundary half-edges for this point index are determined from *PtE* which provides a list of both HE28 and HE42. However HE28 is an edge for a hole in this polygon while HE42 is the correct half-edge to traverse for the outer shell. The `SelectEdge` procedure determines which of these edges to choose and is visually outlined in Figure 5a. Angles between the proposed edges and previous edge HE70 are calculated and the edge with the largest angle is chosen which guarantees the largest concave hull. This edge cannot be a hole edge because that would imply that the hole is outside the concave shell, which is invalid.

On rare occasions the extreme point may have more than one outgoing half edge, meaning that a hole is connected to it. This can be handled in the same way stated above

by using the `SelectEdge` procedure. The only difference is that the previous hull edge is not known (the procedure has just started), but since we know we are on the far right of the hull we can substitute the previous edge for the unit vector $[0,1]$ per Figure 5b. This unit vector is guaranteed to provide a stable order of the angle differences which would have been provided by the actual previous hull edge.

Algorithm 2: ExtractLinearRing

Input : Half Edge Set , $\mathcal{HE} = \{he_i, \dots, he_n\}$
Point Index Hash Map, *PtE*
Starting half-edge, *he*
Start point index, *startPI*
Triangles Point Index, *triangles*

Output: Linear Ring , $lr = [pi_1, \dots, pi_k]$

```
1  $lr = []$ ; /* empty linear ring */
2 while True do
3    $\mathcal{HE} = \mathcal{HE} \setminus he$ 
4    $he_t = NextTriangleEdge(he)$ 
5    $pi = triangles[he_t]$ 
6   Append( $lr, pi$ )
7   if  $pi$  is startPI then
8     /* closed linear ring */
9     break
10     $nextEdges = PtE[pi]$ 
11     $he = SelectEdge(he, nextEdges)$ 
12 end
13 return  $lr$ 
```

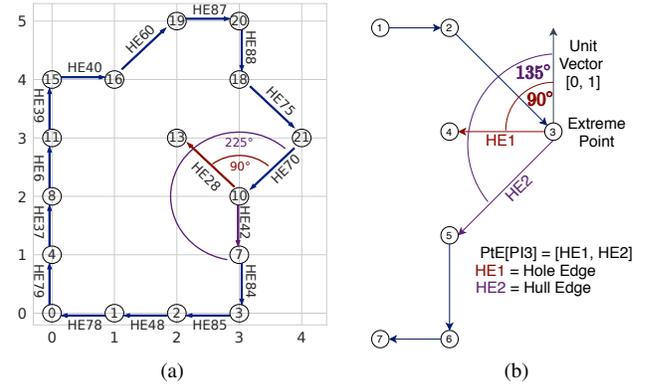


Fig. 5. (a) Edge selection for Fig. 4a. HE70 leads to point index PI10 during shell extraction. Half-edges HE28 and HE42 leave PI10. The correct edge to follow (HE42) has the greatest angle with HE70. (b) If the extreme point has two outgoing edges (HE1, HE2), choose the edge with largest angle difference with the unit vector $[0,1]$. This is edge HE2.

3) *Hole(s) Extraction:* After the outer shell of the concave hull has been determined, only the holes remain to be found (if any holes exist). Any edges that remains inside \mathcal{HE} are hole edges and will be extracted using Algorithm 3. A half-edge is randomly chosen from \mathcal{HE} for which the same `ExtractLinearRing` procedure is run. Figure 6a shows a corner case of a non-manifold mesh that must be handled if two holes share the same vertex. The previously extracted

concave shell is displayed in green while the remaining half-edges to be processed are in blue; note the shared vertex at P116. Figure 6b shows the event when HE19 is randomly chosen for hole extraction leading to P116. HE0 or HE29 is chosen in the manner previously discussed: the edge with largest angle guarantees the smallest hole thus is chosen. If the other edge was chosen this would indicate a hole inside a hole which is invalid.

Algorithm 3: Extract Holes

Input : Half Edge Set , $\mathcal{HE} = \{he_i, \dots, he_n\}$
Point Index Hash Map, PtE
Triangles Point Index, $triangles$

Output: Set of Linear Ring Holes ,
 $\mathcal{HR} = \{lr_1, \dots, lr_k\}$

- 1 $\mathcal{HR} = \emptyset$; /* empty hole set */
- 2 **while** \mathcal{HE} is not empty **do**
- 3 $he = \text{RandomChoice}(\mathcal{HE})$
- 4 $pi = \text{triangles}[he]$
- 5 $lr =$
 $\text{ExtractLinearRing}(\mathcal{HE}, PtE, he, pi, triangles)$
- 6 $\mathcal{HR} = \mathcal{HR} + lr$
- 7 **end**
- 8 **return** \mathcal{HR}

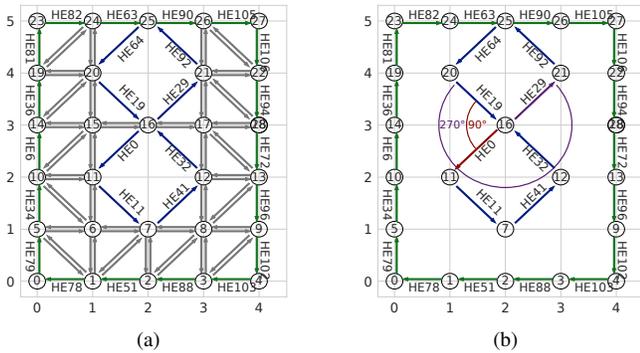


Fig. 6. (a) Edge case of two holes sharing the same vertex at P116. The outer shell (green) is already extracted. (b) When traversing from HE19 to point index P116, two outgoing edges (HE0 and HE29) are found. Edge HE29 with the largest angle difference from HE19 is chosen.

V. BENCHMARKING COMPARISONS

This section benchmarks PolyLidar against other common concave hull extraction methods which also extract holes; all code is open source¹. Three other implementations are tested: CGAL’s Alpha Shape function and the ST_ConcaveHull function from PostGIS and Spatialite. For uniformity, PolyLidar and CGAL are set to use the same α parameter to guarantee exact shape reproduction. Note that CGAL’s Alpha Shape returns an unordered set of boundary edges; it does not convert these edges into a valid (multi)polygon. These edges produce the same shape as PolyLidar when drawn on a

¹<https://github.com/JeremyBYU/concavehull-evaluation>

canvas, but lack the desired polygon semantic data structure. PostGIS’s concave hull implementation only returns single polygons, so MultiPolygon test cases are not evaluated against it. Both PostGIS and Spatialite are databases which require upload of the point set prior to algorithm execution; benchmark timing does not include data upload time.

Section V-A provides a benchmark from plane segmented point clouds produced by an RGBD camera. Section V-B generates synthetic 2D point sets from the state shapes of California (CA) and Hawaii (HI) to explore how the algorithms scale with respect to point size. Section V-C shows a similar benchmark but with the English alphabet. All utilize ground truth (multi)polygon shape GT to evaluate shape accuracy. Each implementation takes as input a point set and produces a concave shape, CS , which is similar to the ground truth polygon. The L^2 error norm, the area of the symmetric difference between GT and CS , is computed to enable evaluation of shape error $\frac{\text{area}((GT-CS) \cup (CS-GT))}{\text{area}(CS)}$.

Each implementation contains its own parameter(s) modified to minimize L^2 error. Shape accuracy is therefore subject to parameter selection. Table II displays the parameters chosen and used for all test cases (RGBD, CA, HI, Alphabet). Rows with two parameters separated by a semicolon indicate parameters for use with non-hole and hole cases. PolyLidar and CGAL use the same α parameter adjusted on a case by case basis. For each case we calculate point density p_d and compute parameter α as $2p_d^{-1}$. This gives reasonable but not necessarily optimal results. Spatialite’s concave hull implementation has parameter C which at its default value ($C = 3$) produces excellent results. C is adjusted as needed (for CA, HI) to further reduce error. PostGIS’ *target percent* is set to provide the optimal accuracy based on percent area reduction required. The most important takeaway when interpreting accuracy is thus trends in accuracy, not small numerical differences.

TABLE II
PARAMETERS FOR TEST CASES

Algorithm	Parameter	RGBD	CA	HI	Alph.
CGAL/PolyLidar	α	$2p_d^{-1}$	$2p_d^{-1}$	$2p_d^{-1}$	$2p_d^{-1}$
Spatialite	C	3.0	2.0	2.0;1.3	3
PostGIS	<i>target %</i>	Varies	0.76;0.72	-	Varies

A. Plane Segmented Point Clouds from RGBD Images

Point clouds were generated with an Intel RealSense D435i camera at 424X240 resolution from eleven different scenes. Ten scenes were taken with the camera 1.5m above ground level pointing directly downward as shown in the top of Figure 8. Floor obstacle positions and orientations were changed in each scene. The camera was placed higher and angled for the eleventh scene shown in the bottom of Figure 8. The floor can be quickly segmented using planar segmentation techniques [4], [5]. However for this experiment the floor was manually segmented, rotated to align with the XY image plane, and subsequently projected. This creates a

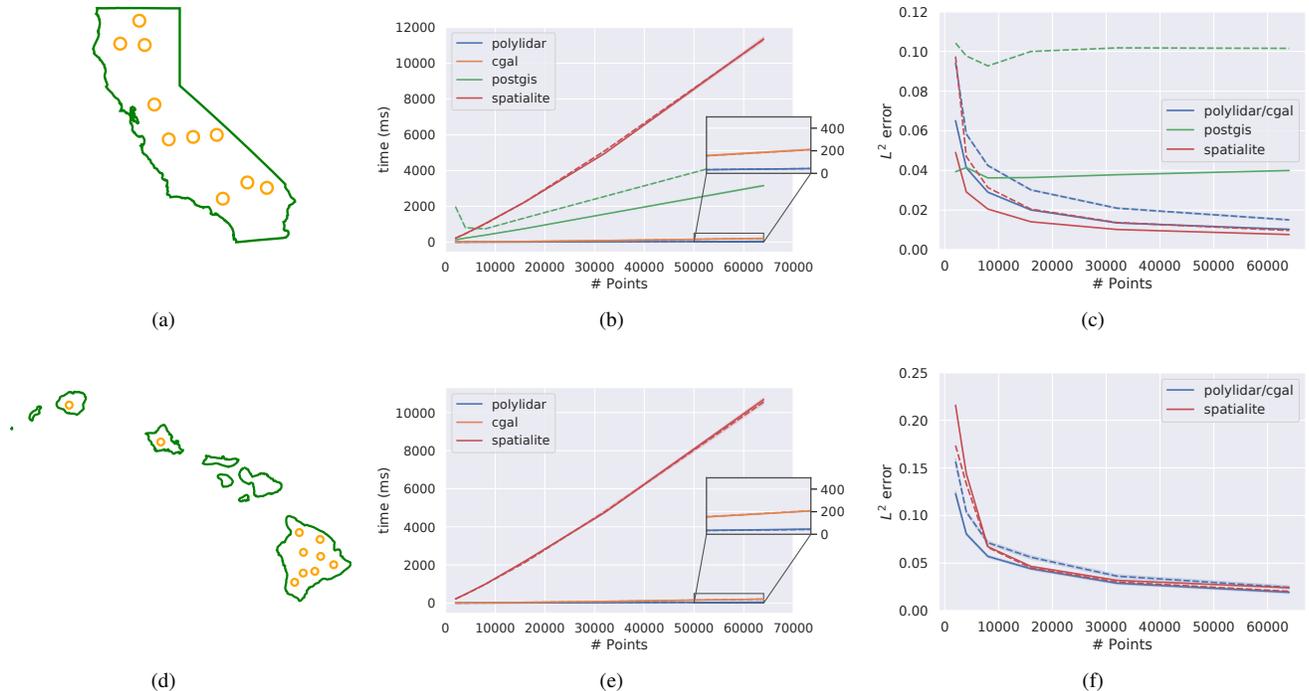


Fig. 7. Concave hull extraction results. Rows from top to bottom correspond to outlines of California (CA) (a, b, c), and Hawaii (HI) (d, e, f) with random holes inserted. The first column shows ground truth polygons with circular holes in orange. The second column shows execution time as a function of number of 2D points provided. The third column shows shape error as a function of number of 2D points provided. Dashed lines show results where holes were placed inside the polygon outline, while solid lines show results with no holes. PostGIS cannot handle MultiPolygons thus was not tested for HI.

2D point set of the floors 3D point cloud. The ground truth polygon of each segmented point cloud was labeled by hand to provide accuracy scores. The average size of the eleven segmented point clouds is 83,184 points. Table III displays the aggregate execution timings and accuracy results of all eleven points clouds for each algorithm. Polylidar is fastest. Polylidar, CGAL, and Spatialite have similar accuracies. Note that Polylidar and CGAL are configured to produce the same shape and therefore have the same L^2 error values.

TABLE III
RGBD PLANE SEGMENTED POINT CLOUDS

Algorithm	L^2 error %			Time (ms)		
	mean	std	max	mean	std	max
Polylidar	2.2	1.5	6.4	47.9	4.3	50.9
CGAL	2.2	1.5	6.4	248.3	25.0	267.7
PostGIS	7.5	1.6	9.9	2734.7	249.3	2939.9
Spatialite	2.2	1.5	6.3	13333.0	2486.6	16386.5

B. State Shapes

Figure 7 shows CA and HI test case geometries (first column), execution times (second column), and error results (third column). Each state shape is processed with and without random holes ; dashed lines indicate results where holes are included in the ground truth polygon. Point sets are randomly sampled from the state shapes. Each test was run 10 times with input point set sizes ranging from (2, 4, 8, 16, 32, 64) thousand points with mean timing and

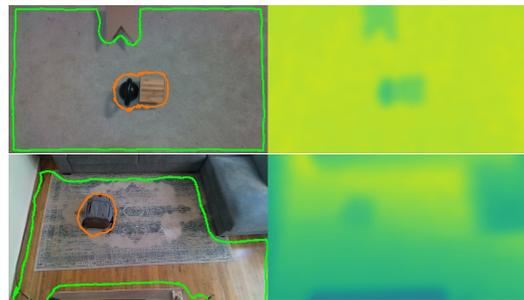


Fig. 8. Two example scenes (top/bottom) from RGBD benchmark. A point cloud is generated from depth image (right) and manually segmented to include only the ground floor. The polygonal output of Polylidar is shown in the RGB image (left). Green is the hull, orange represents holes.

error plotted. Confidence intervals are provided for execution timing, however they are almost imperceptible because the variance is low at this scale. Polylidar and CGAL are significantly faster than the other methods, with Spatialite having the slowest implementation. An inset (zoomed) box that focuses solely on CGAL and Polylidar is shown in the second column, showing that on average Polylidar is ~ 4 times faster than CGAL. The presence of holes affected each method differently: decreased time in Spatialite (fewer triangles to union), increased time for PostGIS (a decrease in *target percent* increases run-time). No significant changes were noted for CGAL and Polylidar.

Spatialite produced shapes with the least error, followed by Polylidar/CGAL and then PostGIS. Spatialite has the

lowest error because it incorporates triangle edge length statistics into its triangle filtering which better handles random sampling. In contrast, Polyliidar/CGAL offer comparable accuracies with RGBD data due to the more uniform point distribution in top-down RGBD imagery. PostGIS error increased markedly with holes since it did not accurately reproduce them. Figure 9 shows a visual comparison of CA concave polygon outputs for each algorithm.

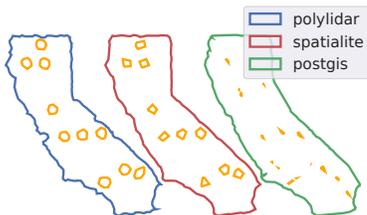


Fig. 9. Concave polygon output from Polyliidar/CGAL (left), Spatialite (center), and PostGIS (right). Input to each algorithm was a 4000 point set sampled from the California (CA) polygon with holes per Figure 7a.

C. Alphabet Shapes

Polygons from 26 capital letters of the English alphabet were generated and 2000 points randomly sampled inside. The “A” in Figure 1b shows an example capital letter with the output of Polyliidar’s concave hull. Table IV provides aggregate statistics of all 26 test cases. Polyliidar continues to lead in speed. Spatialite leads in accuracy by a marginal amount. The alphabet shapes are significantly more concave than previous benchmarks. Documentation of PostGIS indicates that the run time grows quadratically as concavity increases leading to the high execution times observed [15].

TABLE IV
ALPHABET LETTER RESULTS, 26 SHAPES

Algorithm	L^2 error %			Time (ms)		
	mean	std	max	mean	std	max
Polyliidar	12.8	1.8	16.8	1.2	0.3	2.4
CGAL	12.8	1.8	16.8	5.4	0.9	7.2
PostGIS	36.5	9.9	53.7	13091.8	7500.6	28451.0
Spatialite	11.2	4.5	22.1	230.2	6.3	242.9

VI. RANDOM POLYGON TESTS

More than 19,600 polygons were randomly generated to test Polyliidar. Half the test cases had random holes. Polygon complexity is characterized by convexity metric

$$CV = \frac{Area(P)}{Area(CH(P))}$$

where P is the polygon and $CH()$ is the convex hull function. A convexity of 1 indicates the sample polygon is its convex hull. 8,000 points were randomly sampled for each polygon and input to Polyliidar with the α parameter from Table II. Execution time and accuracy are summarized in Table V. The table is partitioned into high, medium, and low ground truth polygon convexity defined by $CV \geq 0.75$,

$0.75 < CV \leq 0.55$, and $CV < 0.55$ respectively. Every polygon produced by Polyliidar was confirmed valid independently by the GEOS geometry library. As polygon convexity (CV) decreases Polyliidar shape estimation accuracy also decreases. Polygons in our “low” convexity class have extremely non-convex shapes, the lowest with $CV = 0.26$ per Figure 10.

TABLE V
RANDOM TESTS; $CV =$ CONVEXITY METRIC

CV	L^2 error %			Time (ms)		
	mean	std	max	mean	std	max
hi	4.4	0.5	6.1	4.6	0.1	5.0
mid	8.0	1.1	13.0	4.6	0.1	8.1
low	15.5	3.0	25.0	4.7	0.2	9.9



Fig. 10. (a) Example of a high convexity polygon; $CV = 86.1\%$ (b) and a low convexity polygon; $CV = 26.2\%$

VII. DISCUSSION

The benchmarks above indicate that Polyliidar is the faster concave (multi)polygon extraction algorithm with the possibility of holes. This section discusses why Polyliidar was faster in comparison to others. We specifically analyze the execution time of the major steps in Polyliidar in comparison to other triangulation-based methods, namely CGAL and Spatialite. The three major steps are:

- 1) Triangulation - The point set is triangulated creating a mesh of faces, edges, and vertices.
- 2) Shape Extraction - Mesh simplices are removed based upon the α parameter or edge length. Remaining triangles, edges, and vertices represent the “shape”.
- 3) Polygon Extraction - The “shape” is converted to a (multi)polygon with the possibility of holes.

Triangulation All perform Delaunay triangulation using robust geometric predicates but use different libraries to do so. Polyliidar uses Delaunator, CGAL uses its own 2D triangulation, and Spatialite uses GEOS.

Shape Extraction Polyliidar and Spatialite are most similar, focusing only on filtering triangles in the mesh. However Polyliidar goes further with region growing (Section IV-C) that isolates disconnected regions in the mesh. For memory efficiency and speed we represent the filtered triangle set \mathcal{T}_f as a bit array with 1/0 indicating in/out of set. This allows rapid triangle filtering and region growing which was previously profiled to be slower when using hashmaps. On the other hand CGAL first creates “interval hashmaps” for its simplices, including triangles, edges, and vertices. These

hashmaps store data detailing at what α -interval a specific simplex would be in the α -complex. These ordered hashmaps give the ability to more quickly compute a *family* of α -shapes from a point set. These data structures are implemented as C++ multimaps with $\mathcal{O}(\log n)$ for insertion/look-up in comparison to unordered maps having $\mathcal{O}(1)$. This design choice leads to shape extraction having an $\mathcal{O}(n \log n)$ complexity for CGAL. By creating hash maps for edges and vertices CGAL can also return the *singular* points and edges which are isolated and not attached to any triangle in the α -complex (e.g., a single point far removed from all others). PolyliDAR need not do this because singular points and edges cannot be polygons thus are not required steps in shape extraction.

Polygon Extraction PolyliDAR independently converts each region into a polygon. Algorithm 2 quickly identifies all border edges and uses efficient unordered contiguous memory hashmaps to store this information in \mathcal{HE} and PtE . The essence of Algorithms 3 and 4 are entirely border-edge based leading to a significant speed up compared to triangle based methods (i.e., perimeter vs. area). Spatialite uses GEOS to take the union of all unfiltered triangles to generate a valid multipolygon. CGAL's Alpha Shape produces an unordered list of the boundary edges of the α -shape. However CGAL does not provide any explicit function to convert this list to a valid (multi)polygon.

TABLE VI
ALGORITHM TIMINGS - MEAN OF 30 RUNS IN MILLISECONDS

Algorithm	triangulation	shape extraction	polygon extraction	total
PolyliDAR	36.0	4.4	1.0	41.4
CGAL	44.5	154.0	–	198.5
Spatialite	234.2	135.3	10788.7	11158.1

Table VI summarizes mean execution timings for each of the main steps for PolyliDAR, CGAL, and Spatialite. The 64,000 point set in the shape of California (with holes) is used, with each algorithm executed 30 times with the mean presented. Relative execution times with other point sets are similar. Delaunator in PolyliDAR triangulated this specific point set fastest with CGAL a close second. PolyliDAR achieves a more significant speed-up in shape extraction for which PolyliDAR is 35 and 32 times faster than CGAL and Spatialite, respectively. Also, PolyliDAR's polygon extraction is about four orders of magnitude faster than Spatialite whereas CGAL does not extract polygons. CGAL instead offers a general purpose α -shape construction routine to compute a family of shapes from different α -values.

VIII. CONCLUSION

This paper has introduced PolyliDAR, an efficient 2D concave hull extraction algorithm which produces (multi)polygon output with holes. Comparison benchmarks of numerous test sets, similarly done in [11], show PolyliDAR is faster than competing approaches with comparable or better accuracy. Additionally we perform random polygon tests that confirm every polygon produced by PolyliDAR is valid. In

future work we will extend PolyliDAR to operate directly on 3D point cloud data by performing both planar segmentation and polygon extraction. We will remove PolyliDAR's reliance on Delaunay triangulation when used with organized point clouds similar to [9] and explore task-based parallelization through the data independence between polygons.

REFERENCES

- [1] Kaustubh Pathak, Andreas Birk, Narunas Vaskevicius, Max Pfingsthorn, Sören Schwertfeger, and Jann Poppinga. Online three-dimensional SLAM by registration of large planar surface segments and closed-form pose-graph relaxation. *Journal of Field Robotics*, 27(1):52–84, 2010.
- [2] S. Malihi, M. J. Valadan Zoej, M. Hahn, M. Mokhtarzade, and H. Arefi. 3D Building Reconstruction Using Dense Photogrammetric Point Cloud. *Intl. Archives of the Photogrammetry, Remote Sensing and Spatial Inf. Sciences*, pages 71–74, Jun 2016.
- [3] M. Himmelsbach, F. v. Hundelshausen, and H.-J. Wuensche. Fast segmentation of 3d point clouds for ground vehicles. In *2010 IEEE Intelligent Vehicles Symposium*. IEEE, 2010.
- [4] Chen Feng, Yuichi Taguchi, and Vineet R. Kamat. Fast plane extraction in organized point clouds using agglomerative hierarchical clustering. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6218–6225, May 2014. ISSN: 1050-4729.
- [5] Trung T. Pham, Markus Eich, Ian Reid, and Gordon Wyeth. Geometrically consistent plane extraction for dense indoor 3D maps segmentation. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4199–4204, Daejeon, South Korea, October 2016. IEEE.
- [6] Alexander Schaefer, Johan Vertens, Daniel Buscher, and Wolfram Burgard. A Maximum Likelihood Approach to Extract Finite Planes from 3-D Laser Scans. In *ICRA*, pages 72–78, Montreal, QC, Canada, May 2019. IEEE.
- [7] Joydeep Biswas and Manuela Veloso. Planar polygon extraction and merging from depth images. In *IEEE/RSJ IROS*, pages 3859–3864, October 2012. ISSN: 2153-0858.
- [8] Tae-kyeong Lee, Seungwook Lim, Seongsoo Lee, Shounan An, and Se-young Oh. Indoor mapping using planes extracted from noisy RGB-D sensors. In *IEEE/RSJ IROS*, pages 1727–1733, October 2012.
- [9] Dirk Holz and Sven Behnke. Fast Range Image Segmentation and Smoothing Using Approximate Surface Reconstruction and Region Growing. In *Intelligent Autonomous Systems 12*, volume 194, pages 61–73. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.
- [10] Github - PolyliDAR. <https://github.com/JeremyBYU/polyliDAR>, 2019. Accessed: 2019-01-05.
- [11] M. Duckham, L. Kulik, M. Worboys, and A. Galton. Efficient generation of simple polygons for characterizing the shape of a set of points in the plane. *Pattern Recognition*, 41(10):3224–3236, 2008.
- [12] H. Edelsbrunner, D. Kirkpatrick, and R. Seidel. On the shape of a set of points in the plane. *IEEE Transactions on Information Theory*, 29(4):551–559, July 1983.
- [13] A. Furieri. Spatialite. <https://www.gaia-gis.it/fossil/libspatialite/index>, 2017. Visited on 2017-01-14.
- [14] D. R. Hipp, D. Kennedy, and J. Mistachkin. SQLite (version 3.28) sqlite development team. <https://www.sqlite.org/index.html>, 2020. Visited on 2020-06-14.
- [15] OSGeo. Postgis. https://postgis.net/docs/ST_ConcaveHull.html, 2019. Visited on 2019-01-14.
- [16] The CGAL Project. *CGAL User and Reference Manual*. CGAL Editorial Board, 4.14 edition, 2019.
- [17] John R Herring. OpenGIS implementation specification for geographic information-Simple feature access- Part 1: Common architecture. *Open Geospatial Consortium*, page 95, 2006.
- [18] Github - Delaunator. <https://github.com/mapbox/delaunator>, 2018. Accessed: 2018-01-05.
- [19] Github - Port of Delaunator to C++. <https://github.com/delfrerr/delaunator-cpp>, 2018. Accessed: 2018-01-05.
- [20] Jonathan Richard Shewchuk. Adaptive precision floating-point arithmetic and fast robust geometric predicates. *Discrete & Computational Geometry*, 18(3):305–363, 1997.
- [21] Rujun Cao, Yongjun Zhang, Xinyi Liu, and Zongze Zhao. Roof plane extraction from airborne lidar point clouds. *International Journal of Remote Sensing*, 38(12):3684–3703, 2017.