Computing Large Planar Regions in Terrains, with an Application to Fracture Surfaces[⋆]

Michiel Smid

Carleton University, Ottawa, Ontario, Canada K1S 5B6.

Rahul Ray

Max-Planck-Institute for Computer Science, D-66123 Saarbrücken, Germany.

Ulrich Wendt

Department of Materials Science, University of Magdeburg, D-39106 Magdeburg, Germany.

Katharina Lange

Department of Materials Science, University of Magdeburg, D-39106 Magdeburg, Germany.

Abstract

We consider the problem of computing the largest region in a terrain that is approximately contained in some two-dimensional plane. We reduce this problem to the following one. Given an embedding of a degree-3 graph G on the unit sphere $\mathbb{S}^2$, whose vertices are weighted, compute a connected subgraph of maximum weight that is contained in some spherical disk of a fixed radius. We give an algorithm that solves this problem in $O(n^2 \log n (\log \log n)^3)$ time, where n denotes the number of vertices of G or, alternatively, the number of faces of the terrain. We also give a heuristic that can be used to compute sufficiently large regions in a terrain that are approximately planar. We discuss an implementation of this heuristic, and show some experimental results for terrains representing three-dimensional (topographical) images of fracture surfaces of metals obtained by confocal laser scanning microscopy.

Key words: Computational geometry, terrain, planar region, optimization

1 Introduction

A *terrain* is the surface in $\mathbb{R}^3$ described by a real-valued function in two real variables. If this function is piecewise linear and the surface consists of a collection of triangles, then the terrain is called a *triangulated irregular network (TIN)*. We consider the problem of computing the largest connected region in a TIN that is approximately planar. To be more precise, given a TIN, we want to compute that subset T of the triangles of the TIN such that (i) for each triangle $t \in T$ there is another triangle $t' \in T$ such that t and t' share an edge, (ii) there is a two-dimensional plane that approximately contains all triangles of T, and (iii) the total area of the triangles of T is as large as possible.

In order to define this problem rigorously, we have to specify the notion of being approximately contained in a plane. In this paper, we define this notion as follows. Let ϵ be a small positive real number, and let T be a set of triangles satisfying (i) above. We say that T is ϵ-*planar* if there is a vector $\mathbf{c}$ such that the normal vectors of all triangles in T make an angle of at most ϵ with $\mathbf{c}$. We now show how this notion can be used to reformulate our problem.

The *unit sphere*, i.e., the boundary of the three-dimensional ball centered at the origin and having radius one, is denoted by $\mathbb{S}^2$. The *upper hemisphere* is defined as $\mathbb{S}^2_+ := \mathbb{S}^2 \cap \{(x, y, z) \in \mathbb{R}^3 : z > 0\}$. We can regard the normal vector of any triangle in a terrain as a point on $\mathbb{S}^2_+$. The lower hemisphere $\mathbb{S}^2_-$ is of no interest to us because, due to the imaging principle, each vertical line can intersect the surface described by the image in only one point. Hence, no normal can have a negative z-coordinate.

Consider a TIN, and let $S \subseteq \mathbb{S}^2_+$ be the set of normal vectors of the triangles in this TIN. Let $G = (S, E)$ be the undirected graph having vertex set S and in which any two vertices are connected by an edge if and only if the corresponding triangles in the TIN share an edge. (Actually, S is a multiset because different triangles may have the same normal. Equal normals are treated as different vertices in G.) Observe that each vertex of G has degree at most three. We give each vertex p of G a weight $wt(p)$ which is equal to the area of the triangle that gives rise to p. The weight $wt(C)$ of any subset C of S is defined as $wt(C) := \sum_{p \in C} wt(p)$.

For any point $x \in \mathbb{S}^2_+$, let D_x denote the spherical disk of radius ϵ centered at x. That is, D_x is the set of all points $y \in \mathbb{S}^2$ such that the angle between the vectors $\mathbf{x}$ and $\mathbf{y}$ is less than or equal to ϵ. Furthermore, let G_x denote the subgraph of G having $S \cap D_x$ as its vertex set and whose edge set is the set of all edges $(p, q) \in E$ for which p and q are both contained in D_x. Finally, we

rahul@mpi-sb.mpg.de (Rahul Ray), Ulrich.Wendt@mb.Uni-Magdeburg.DE (Ulrich Wendt), Katharina.Lange@mb.Uni-Magdeburg.DE (Katharina Lange).

define W_x to be the maximum weight of any connected component of the graph G_x. Using this terminology, our problem can be formally stated as follows.

Problem 1 *Given a graph G as above having n vertices, and a real constant $\epsilon > 0$, compute a point $x \in \mathbb{S}_+^2$ such that W_x is maximum.*

Let (p, q) be any edge of the graph G. If the angle between the vectors $\mathbf{p}$ and $\mathbf{q}$ is larger than 2ϵ, then it is clear that (p, q) can be ignored when solving Problem 1. Therefore, we may assume without loss of generality that the angle between the endpoints (when regarded as vectors) of any edge of G is at most 2ϵ.

1.1 Motivation

The problem considered in this paper arose in a collaboration between the Departments of Computer Science and Materials Science at the University of Magdeburg. The goal of this project is to design and implement algorithms that can be used for the quantification of fracture surface topographies, given as three-dimensional images obtained by confocal laser scanning microscopy.

The images we used are taken from fracture surfaces of metals. Knowledge of fracture surface topographies of materials can be used to interpret the mechanical properties and the fracture mechanism of those materials. This knowledge can help to detect the weak points in a material's micro-structure with regard to the mechanical properties. These relationships can be a guideline for the improvement and to the modification of known materials, as well as to new design criteria of materials. This holds for all kind of materials such as metals, polymers, ceramics, and composites.

The images are given as 512×512 arrays of pixels, where the grey value stored at each entry is equal to the height of the corresponding terrain point. One goal in our research is to find large connected regions in this image that are approximately planar. The normal vectors and areas of these planar regions give useful information about the fracture surface generating process. In Section 4, we will see how we converted the array of pixels to a TIN.

1.2 Our results

In Section 2, we will show how computational geometry and dynamic graph algorithms can be used to solve Problem 1 in $O(n^2 \log n (\log \log n)^3)$ time.

It is unlikely that Problem 1 can be solved in subquadratic time. In fact,

it seems that even the problem of computing a point $y \in \mathbb{S}^2_+$ such that W_y approximates the optimal solution cannot be solved in subquadratic time. Therefore, in Section 3, we describe a simple grid-based heuristic. We have implemented this heuristic and discuss some details about it in Section 4. We also present some experimental results on images of fracture surfaces obtained by confocal laser scanning microscopy. These results show that the heuristic is able to find ϵ-planar regions whose area is sufficiently large.

1.3 Related work

The problem considered in this paper is related to the *terrain simplification problem*. In this problem, we want to approximate a polyhedral terrain by a "smaller" terrain, i.e., one having the minimum number of vertices. Although this problem has been studied in, for example, the computer graphics community [12], the main reference we are aware of that considers this problem from a complexity point of view is Agarwal and Suri [2]. They give evidence that the terrain simplification problem is hard by proving that a strongly related problem is NP-hard. They also give a polynomial-time algorithm for approximating the minimum terrain.

The problem 3SUM is defined as follows: Given three sets A, B, and C, each consisting of n real numbers, decide whether there are elements $a \in A$, $b \in B$, and $c \in C$ such that $a + b = c$. This problem can be solved in $O(n^2)$ time, and it is widely believed that it cannot be solved in subquadratic time. A problem P is called 3SUM-*hard* if 3SUM can be reduced to P in subquadratic time; see Gajentaan and Overmars [8].

If we consider Problem 1 for the case when the graph G is complete and all vertices have equal weights, then we get the problem of computing a placement of a spherical disk that contains the largest subset of a given set S of n points on $\mathbb{S}^2$. This problem has been considered by Chazelle and Lee [5] for the case when S is a set of points in the Euclidean plane. They showed that the problem can be solved in $O(n^2)$ time. The related problem of computing the deepest point in an arrangement of halfplanes is 3SUM-hard, see [8]. This indicates that it is unlikely that the disk placement problem can be solved in subquadratic time. Recently, Agarwal *et al.* [1] gave an alternative $O(n^2)$-time algorithm for the optimal disk placement problem, as well as randomized approximation algorithms whose running times are close to linear. Our algorithm in Section 2 for solving Problem 1 has been inspired by the $O(n^2)$-time algorithm in [1].

2 Solving Problem 1

In this section, we give an algorithm that solves Problem 1. Consider the graph $G = (S, E)$, and consider the spherical disks D_p centered at the points p of S. Let A be the arrangement on $\mathbb{S}^2$ defined by the boundaries of the disks D_p, where $p \in S$. That is, A is the subdivision of $\mathbb{S}^2$ into vertices, edges and faces defined by the overlay of the boundaries of the disks D_p, $p \in S$. Since $p \in D_x$ if and only if $x \in D_p$, we have $W_x = W_y$ for any two points x and y that are in the interior of the same face f of A. Also, for each vertex z of f, we have $W_z \geq W_x$. This proves the following lemma.

Lemma 2 *To solve Problem 1, it suffices to consider points $x \in \mathbb{S}^2_+$ that are vertices of the arrangement A.*

Throughout this section, we make the following *general-position* assumption about the set S. We assume that the elements of S are pairwise distinct. Moreover, we assume that for any two distinct points p and q of S, the spherical disks D_p and D_q are either disjoint or have an intersection of positive area (hence, D_p and D_q do not touch each other). Finally, for any three distinct points p, q, and r of S the spherical disks D_p, D_q, and D_r do not intersect in a single point. We make this assumption only to simplify the description of our algorithm. This algorithm can easily be extended to handle arbitrary sets of points.

The discussion above leads to the following preliminary algorithm for solving Problem 1.

Step 1: Compute the arrangement A.

Step 2: Let $W := 0$. For each vertex x of A, do the following.

- Compute the graph G_x.
- Compute the connected components of G_x, together with their weights. Let W_x be the maximum weight of any of these connected components.
- Set $W := \max(W, W_x)$.

Step 3: Return W.

It is clear that this algorithm correctly solves Problem 1. Let us analyze its running time. Recall that n denotes the number of elements of the point set S. For any $p \in S$, let $deg(p)$ denote the degree of p in G. Observe that $\sum_{p \in S} deg(p)$ is equal to twice the number of edges of G, and that $deg(p) \leq 3$ for any $p \in S$. Therefore, G has at most $3n/2$ edges.

Step 1, i.e., computing the arrangement A, takes $O(n^2)$ time using, e.g., the algorithm of Amato *et al.* [3]. Consider any vertex x of A. The graph G_x can clearly be computed in time proportional to the number of vertices and edges of G; hence, G_x can be computed in $O(n)$ time. Given G_x, its connected components and the value W_x can be computed in $O(n)$ time; see, e.g., the book by Cormen *et al.* [6]. Hence, for each vertex x of A, $O(n)$ time is spent in Step 2. Since A has $O(n^2)$ vertices, the overall time for Step 2 is $O(n^3)$. Hence, the entire algorithm takes $O(n^3)$ time.

We now show how to improve the running time considerably. Note that the bottleneck of the previous algorithm is Step 2. The idea of the improved algorithm is to traverse the arrangement A and maintain the connected components of G_x in a data structure. Consider what happens when we walk along an edge of A from a vertex x to a vertex y. Assume that we know the connected components of the graph G_x. Our goal is to compute the connected components of G_y as fast as possible. Observe that walking from x to y means that we move the spherical disk D_x along an edge of A to the position D_y. During this move, at most one point of S can enter or leave the spherical disk. (Here we use our general-position assumption.) Since the graph G has degree three, it follows that the graph G_y can be obtained from G_x by performing at most a constant number of edge insertions and deletions.

Assume that we have a data structure CC that stores the connected components, together with their weights, of a given graph, and that supports edge insertions and deletions, and queries of the form "report the maximum weight of any connected component". (We will specify this data structure later. For the moment, we use it as a black box.) For any point $x \in \mathbb{S}^2_+$, we denote by CC_x the instance of this data structure for the graph G_x.

Our improved algorithm does the following.

Step 1: Compute the arrangement A.

Step 2: Let x be an arbitrary vertex of A.

- Compute the graph G_x.
- Compute the connected components of G_x, together with their weights. Compute W_x as the maximum weight of any connected component of G_x.
- Set $W := W_x$.
- Construct the data structure CC_x.

Step 3: Starting at x, traverse the vertices of the arrangement A, e.g., in depth-first order. In a generic step, we walk from a vertex y, along an edge of A, to a neighboring vertex z. At the moment when we leave y, we have

the data structure CC_y, storing the connected components of the graph G_y, together with their weights. The graph G_z can be obtained by inserting and deleting at most a constant number of edges in the current graph G_y. Hence, we obtain the data structure CC_z by performing these updates in the data structure CC_y. Afterwards, we query CC_z to find the value of W_z, and set $W := \max(W, W_z)$.

Step 4: Return W.

The correctness of this algorithm is clear. Steps 1 and 4 take $O(n^2)$ and $O(1)$ time, respectively. The times for Steps 2 and 3 depend on the data structure CC. Let $P(n)$, $U(n)$, and $Q(n)$ denote the preprocessing time, update time, and query time of this data structure, respectively. Then Step 2 takes $O(n + P(n))$ time. In Step 3, we spend $O(U(n) + Q(n))$ time for each edge of A. Since this arrangement has $O(n^2)$ edges, it follows that the total running time of the algorithm is

$$O\left(P(n) + n^2(U(n) + Q(n))\right).$$

It remains to specify the data structure CC. In [14], Thorup gives a data structure for maintaining a spanning forest of a graph under insertions and deletions of edges, in $O(\log n(\log\log n)^3)$ amortized time per update, where n denotes the number of vertices of the graph. Given any two vertices of this graph, it can be decided in $O(\log n/\log\log\log n)$ time if they are in the same connected component. (A simpler but theoretically slightly less efficient data structure was given by Holm *et al.* [10].) This data structure can easily be extended so that it maintains the weights of all connected components within the same time bound. If we store these weights in a heap, then we can extract the weight of the largest connected component in $O(1)$ time. Moreover, this heap can be updated in $O(\log n)$ time per operation. The data structure can be built by successively inserting all edges into an initially empty graph. Hence, we have $P(n) = O(n \log n(\log\log n)^3)$, $U(n) = O(\log n(\log\log n)^3)$, and $Q(n) = O(1)$. Thus, we have proved the following result.

Theorem 3 *Problem 1 can be solved in $O(n^2 \log n(\log\log n)^3)$ time.*

It is not clear if Problem 1 can be solved in subquadratic time; see Section 1.3. Instead, one can ask about the time complexity for approximating the optimal solution. That is, let δ be a real number such that $0 < \delta < 1$, and let $x \in \mathbb{S}_+^2$ be a point for which W_x is maximum. In the approximation version of Problem 1, we have to compute a point $y \in \mathbb{S}_+^2$ such that $W_y \geq (1 - \delta)W_x$.

Consider the following example. Let D be a spherical disk of radius ϵ, and let p and q be two diametrally opposite points on the boundary of D. Let m be

7

a large integer. For each i, $1 \leq i \leq m$, let $a_i := p$ and $b_i := q$. Consider the edges (a_i, b_i), $1 \leq i \leq m$, and (b_i, a_{i+1}), $1 \leq i < m$. Note that these edges form a path between a_1 and b_m that alternates between the points p and q. Let G be a graph containing the points a_i and b_i, $1 \leq i \leq m$, as vertices, and the above edges. All other vertices of this graph are "far" away from p and q, and have "large" distances among each other. We assume that all vertices of G have unit-weight. (It is easy to construct a TIN for which G is the corresponding graph.) For this graph, the center of D gives the optimal solution to Problem 1. If we move the disk D, then either the point p or the point q leaves the disk and, hence, each connected subgraph of G that is contained in the disk consists of one single vertex. This shows that any approximation algorithm for Problem 1 must return the center of D. Because of this, the approximation version of Problem 1 has the same time complexity as Problem 1 itself.

3   A heuristic for finding large ϵ-planar regions

In this section, we give a simple heuristic approach to compute large connected regions in a TIN that are ϵ-planar. (Different regions need not be approximately contained in the same two-dimensional plane.) We cannot prove any non-trivial bounds on the quality of its output, but experiments have shown that the output is good in practice, and that the heuristic is fast.

Consider a TIN consisting of n triangles, let $\epsilon > 0$, let S be the set of normal vectors of these triangles, and let $G = (S, E)$ be the graph as defined in Section 1. Recall that the weight $wt(p)$ of any element p of S is equal to the area of the triangle whose normal vector is p. Also, recall that S is actually a multiset. Our goal is to find all ϵ-planar regions in the TIN having area at least A, where A is some given positive real number.

We define a grid on the upper hemisphere $\mathbb{S}_+^2$ using lines of longitude and latitude such that every grid cell is contained in some spherical disk of radius ϵ. To be more precise, we choose an appropriate real number δ with $\delta = \Theta(\epsilon)$ and use $X := \lceil \pi/\delta \rceil$ equally spaced lines of longitude lo_i, $0 \leq i < X$, and $Y := \lceil \pi/(2\delta) \rceil$ equally spaced lines of latitude la_j, $0 \leq j < Y$.

For any two indices i and j with $0 \leq i < X$ and $0 \leq j < Y$, we call the pair (i, j) the *index* of the grid cell bounded by lo_i, lo_{i+1}, la_j, and la_{j+1}. For any point $p \in \mathbb{S}_+^2$, we can in $O(1)$ time compute the index of the grid cell containing p. Note that each grid cell is adjacent to at most four other cells, except for those that are incident to the north pole.

Our heuristic takes as input the TIN, the positive real numbers ϵ and A,

and the graph $G = (S, E)$. It starts by computing a subset of all ϵ-planar regions in the TIN having area at least $A/4$. (Below, it will become clear why we choose $A/4$ instead of A. In fact, the factor $1/4$ can be replaced by any constant between zero and one.) Then it makes a boundary correction step, in which each region is enlarged, while remaining ϵ-planar. To be more precise, the heuristic makes the following seven steps.

Step 1: Initialize an array $C[0..X-1, 0..Y-1]$, and store with each entry an empty list of points and an empty list of edges.

Step 2: For each point $p \in S$, compute the index (i_p, j_p) of the grid cell that contains p, and add p to the point list of $C[i_p, j_p]$.

Step 3: For each edge $(p, q) \in E$ with $(i_p, j_p) = (i_q, j_q)$, add (p, q) to the edge list of $C[i_p, j_p]$. (Hence, in this step, all edges that are completely contained within one grid cell are extracted.)

Step 4: Initialize an empty list L.

Step 5: For each i and j with $0 \leq i < X$ and $0 \leq j < Y$, do the following.

- Let G_{ij} be the graph having the points of $C[i, j]$ as its vertices, and the edges of $C[i, j]$ as its edges. Compute the connected components of G_{ij}, together with their weights.
- For each connected component of G_{ij}, add it to the list L if it has weight at least $A/4$.

Consider the list L when Step 5 has been completed. Each element of L corresponds to a connected subgraph of G having weight at least $A/4$ and that is contained in one grid cell, i.e., it corresponds to an ϵ-planar region in the TIN having area at least $A/4$. It may happen that such a region R can be enlarged by adding triangles that are adjacent to R and whose normals are in a grid cell that is adjacent to the grid cell that gave rise to R. Of course, the enlarged region should still be ϵ-planar. Below, we describe a "boundary correction" step that does exactly this.

Step 7 (Boundary correction step): This step is performed for each connected subgraph of G that is stored in the list L. Let G' be any such subgraph.

We first compute the smallest enclosing spherical disk D' containing all vertices of G', using the linear-time algorithm as presented, e.g., in de Berg *et al.* [7]. Let c be the center of D', and let D be the spherical disk of radius ϵ centered at c. Note that G' is contained in D, because D' is contained in D.

We now go back to the TIN and mark the triangles corresponding to the vertices of G'. Let T be the set of all marked triangles. Note that all normals of triangles in T are contained in D. Now we consider each unmarked triangle t of the TIN that shares an edge with at least one marked triangle, and mark t if its normal is contained in D. Let T' be the set of all marked triangles after we have considered all such triangles t. If $T = T'$, then we stop the boundary correction step for this subgraph G'. Otherwise, we repeat this process by considering unmarked triangles that share an edge with at least one marked triangle.

Step 8: After having completed the boundary correction step for each element in L, we have a collection of ϵ-planar regions, each one having area at least $A/4$. Since these regions may overlap, we continue as follows. We sort all these enlarged regions according to their areas. If the largest region has area at least A, then we report it, and mark all its triangles in the TIN. Then we consider the second largest region. If its area is at least A, and if none of its triangles has been marked yet, we also report it and mark all its triangles in the TIN. We continue this until we have reported all regions whose area is at least A and that do not overlap any of the previously reported regions.

What can we say about the quality of the output? Let R be any ϵ-planar region in the TIN, and let G' be the corresponding connected subgraph of G. Our algorithm reports the region R if and only if G' contains a connected subgraph G'' of weight at least $A/4$ that is completely contained in one of our grid cells. If each edge of G' crosses a cell boundary, then R will not be reported. Therefore, in order to improve the chances of finding R, we run the algorithm several times using shifted copies of the grid.

4   Implementation and experimental results

We have implemented the heuristic that was discussed in the previous section. The program, which we call *PlaneFinder*, takes as input a $k \times \ell$ array of pixels in `tif` (tagged image file) format. The value stored at each entry is the height of the corresponding surface point, i.e., the z-value of the voxel position.

The application is open to all kind of imaging methods which deliver topographical images. Besides confocal laser microscopy that we used to obtain our images, examples include scanning force microscopy [13], white light interference microscopy [9], extended focus conventional light microscopy [15], and photogrammetry based on stereo pairs received by scanning electron microscopy [4].

Fig. 1. *Converting an array of pixels to a TIN.*

PlaneFinder starts by computing a TIN as illustrated in Figure 1. The vertices of the TIN are (i) the centers of the pixels, where the z-coordinate is given by the height of the pixel, and (ii) the centers of all 2×2 blocks of pixels, where the z-coordinate is given by the average height of the four pixels comprising the block. These vertices are joined by edges as indicated in the right part of Figure 1. Note that the total number of triangles, which we denote by n, is equal to $n = 4(k - 1)(\ell - 1)$. Given this TIN, *PlaneFinder* proceeds as described in Section 3.

The program is written in C++ and uses the LEDA library [11]. In the current version, the following operations are supported.

- Show the k largest ϵ-planar regions (as found by the heuristic), for some values k and ϵ provided by the user.
- Show all ϵ-planar regions having area at least A (again as found by the heuristic), for some values ϵ and A provided by the user.

The output is the original `tif`-image in which all regions that have been detected are colored. (A pixel in the image is colored if and only if at least one of the triangles overlapping the pixel belongs to a region.) The program takes care that adjacent regions are drawn in visibly distinct colors. The program also lists the regions found in increasing order of their area.

We have run *PlaneFinder* on various images of fracture surfaces obtained by confocal laser scanning microscopy. As mentioned already, these images are in `tif` format, and consist of 512×512 pixels. Hence, the corresponding TINs consist of $n = 1,044,484$ triangles. Most of the surfaces are very rough and contain only small ϵ-planar regions. For a typical image *PlaneFinder* takes about 35 seconds.

An example is given in Figure 2, which shows the original image (on the left) as well as the output of *PlaneFinder* (on the right). The program was run with ϵ equal to five degrees and A equal to 1500 square units. The right part shows, in white, the four ϵ-planar regions found having area at least $A = 1500$ square

Fig. 2. *On the left, a confocal laser scanning micrography (topographical image) of a steel fracture surface is displayed. The colors represent the heights of the pixels. The four ϵ-planar regions computed by PlaneFinder are marked white in the right image. The image size is $100\mu m \times 100\mu m$.*

units. (Note that area refers to the true area in space and not the projected area.)

Figure 3 shows an example of an image in which we have added four large slanted triangles. One of these is ϵ-planar for $\epsilon = 5$ degrees, whereas the others are ϵ'-planar for a value of ϵ' that is slightly larger than five degrees. *PlaneFinder* was run on this image with $A = 3000$ square units. It correctly detected only one of the triangles.

In order to flatten local roughness in the images, we have included the option to apply a mean-filter in a preprocessing step. When we apply this filter to the image of Figure 3, then *PlaneFinder* finds all four triangles.

5 Concluding remarks

We have considered the problem of computing large regions in a terrain that are connected and approximately planar. We showed that the problem of computing the largest such region can be solved in a time that is roughly quadratic in the number of triangles in the terrain, and argued that it is unlikely to solve the problem faster. We leave open the problem of proving this rigorously. We also argued that it may even be hard to approximate this largest region. Proving this claim formally is also left as an open problem.

We defined a connected set of triangles to be "approximately planar" if their

Fig. 3. *On the top, a topographical image and its parallel projection onto the xy-plane are displayed. One "approximately planar" triangle and three "not so approximately planar" triangles have been inserted for testing purposes. The ϵ-planar regions computed by PlaneFinder are marked white in the bottom image.*

normals are contained in a spherical disk of a fixed small radius. It would be interesting to solve the problem for other notions of being approximately planar. For example, we could require the angular diameter of the normals to be at most ϵ, or the set of triangles to be contained between two parallel planes having distance ϵ.

A first application of our software is the detection of fracture facets in brittle fracture surfaces of steel and greycast iron.

Our software is in a preliminary stage and has a lot of scope for improvements. We will add various features to improve the user interaction. In a future version of *PlaneFinder*, we will provide the user with more information about the planar areas. For example, by clicking on a region, information such as the normal vector and the area of the region will be displayed. Also, we will add a function that provides a three-dimensional view of the image with the ability to rotate it.

13

Acknowledgements

References

[1] P. K. Agarwal, T. Hagerup, R. Ray, M. Sharir, M. Smid, and E. Welzl. Translating a planar object to maximize point containment. In *Proceedings of the 10th European Symposium on Algorithms*, volume 2461 of *Lecture Notes in Computer Science*, pages 42–53, Berlin, 2002. Springer-Verlag.

[2] P. K. Agarwal and S. Suri. Surface approximation and geometric partitions. *SIAM J. Comput.*, 27:1016–1035, 1998.

[3] N. M. Amato, M. T. Goodrich, and E. A. Ramos. Computing the arrangement of curve segments: Divide-and-conquer algorithms via sampling. In *Proc. 11th ACM-SIAM Sympos. Discrete Algorithms*, pages 705–706, 2000.

[4] A. Boyde. Quantitative photogrammetric analysis and qualitative stereoscopic analysis of SEM images. *J. of Microsc.*, 98:452–471, 1973.

[5] B. Chazelle and D. T. Lee. On a circle placement problem. *Computing*, 36:1–16, 1986.

[6] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 2nd edition, 2001.

[7] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, 2nd edition, 2000.

[8] A. Gajentaan and M. H. Overmars. On a class of $O(n^2)$ problems in computational geometry. *Comput. Geom. Theory Appl.*, 5:165–185, 1995.

[9] A. Harasaki, J. Schmit, and J. C. Wyant. Improved vertical scanning interferometry. *Appl. Opt.*, 39:2107–2115, 2000.

[10] J. Holm, K. de Lichtenberg, and M. Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. In *Proc. 30th Annu. ACM Sympos. Theory Comput.*, pages 79–89, 1998.

[11] K. Mehlhorn and S. Näher. *LEDA: A Platform for Combinatorial and Geometric Computing*. Cambridge University Press, Cambridge, U.K., 1999.

[12] N. S. Sapidis and P. J. Besl. Direct construction of polynomial surfaces from dense range images through region growing. *ACM Transactions on Graphics.*, 14:171–200, 1995.

[13] U. D. Schwarz. Scanning force microscopy. In S. Amelinckx, D. van Dyck, and J. van Landuyt van G. van Tendeloo, editors, *Handbook of Microscopy, Vol. II*, page 827. VCH Verlagsgesellschaft, Weinheim, 1997.

[14] M. Thorup. Near-optimal fully-dynamic graph connectivity. In *Proc. 32nd Annu. ACM Sympos. Theory Comput.*, 2000.

[15] Y. Yamaguchi, M. K. Weldon, and M. D. Morris. Fractal characterization of SERS-active electrodes using extended focus reflectance microscopy. *Appl. Spectrosc.*, 53:127–132, 1999.