# Peeling Meshed Potatoes 

Boris Aronov • Marc van Kreveld • Maarten Löffler • Rodrigo I. Silveira

Received: 31 March 2009 / Accepted: 21 July 2009 / Published online: 11 September 2009
© The Author(s) 2009. This article is published with open access at Springerlink.com


#### Abstract

We study variants of the potato peeling problem on meshed (triangulated) polygons. Given a polygon with holes, and a triangular mesh that covers its interior (possibly using additional vertices), we want to find a largest-area connected set of triangles of the mesh that is convex, or has some other shape-related property. In particular, we consider (i) convexity, (ii) monotonicity, (iii) bounded backturn, and (iv) bounded total turning angle. The first three problems are solved in polynomial time, whereas the fourth problem is shown to be NP-hard.


Keywords Geometric optimization • Dynamic programming • Potato peeling

[^0]
## 1 Introduction

Suppose you live in a mountainous area, and you are interested in playing golf and drinking wine. Then you would like to find a large subregion of the area that is suitable to turn into a golf course or a vineyard. In the first case, this means your region should be reasonably flat. In the second case, it probably means your region should 'face south' (assuming that you live in the Northern Hemisphere). In addition, it is reasonable to want your region to have a compact shape and, at the same time, to be as large as possible.

In Geographic Information Science, terrain data is often stored as a triangulated irregular network (TIN): a planar triangulation with additional height data on the vertices. In this context, it is easy to evaluate properties such as 'flatness' and 'facing south' on a triangle-by-triangle basis: each triangle is part of a plane and thus has a constant angle with respect to the horizontal or to the optimal south-facing orientation. We can identify for each triangle whether its angle is below a certain prespecified threshold, and discard it if it is not. After this, we are left with a collection of triangles in the plane, of which we want to find a largest-area subset that has a good shapefrom here on, the height information can be ignored. Note that we can assume that the collection of triangles we get as input is connected, since otherwise we can just search for good regions within each connected component separately.

Now consider a seemingly unrelated problem. In computer graphics, before a scene is rendered an occlusion culling algorithm is often used, which serves to quickly remove a large portion of the scene that is hidden behind some object (occluder). Such occluders are often large and contain many triangles. However, culling algorithms work much faster with convex occluders, so it is interesting to find a large convex subset of a given occluder ('convex' should be interpreted in the 2dimensional sense, as seen from the current point of view).

Both of these problems can be abstracted in the same way. We have a triangulated polygon (with extra interior vertices) with holes, of which we want to find a large subregion with a good shape. Finding well-shaped subregions of polygons is a problem that received quite some attention in computational geometry before, and one of its representatives is the well-known potato peeling problem [14], which asks for the largest convex subpolygon of a given polygon (potato). In our case, the problem is the same, except that we have an additional triangulation or mesh on the potato, and we require the solution to respect that mesh. Asking for convexity, in the context of our first example, can be a way to model the vague criterion of good shape for a region. However, because convexity is a rather restrictive condition in combination with respecting the triangular mesh, and because many applications may allow some flexibility on this requirement, we also study some relaxations of convexity that can be used to model the idea of a well-shaped region.

Problem Description We are given a polygon with holes (or potato) $\mathcal{P}$, and a triangular mesh $\mathcal{M}$ that covers the interior of $\mathcal{P}$, and possibly has additional interior vertices. The total complexity of $\mathcal{P}$ and $\mathcal{M}$ is denoted by $n$, and is defined as their total number of vertices. We want to find the largest-area simply-connected region $R$ comprised of triangles of $\mathcal{M}$ such that:
(i) it is convex (see Fig. 1(b));
(ii) it is monotone in some direction (see Fig. 1(c));
(iii) it has a maximum backturn of at most $\gamma$ (see Fig. 1(d)); or
(iv) it has total turning angle of at most $\beta$ (see Fig. 1(e)).

That is, we treat four different shape restrictions: convexity and three possible relaxations, explained below. Since the region boundary is required to follow the edges of $\mathcal{M}$, forcing $R$ to be convex can be quite restrictive, while intuitively this should not be the case. For example, in Fig. 2 a situation is depicted in which a very large area is available, but only rather small convex subregions exist. The problem is that because of the mesh, the boundary needs to make a small turn at every vertex, and in a convex polygon we are only allowed to turn in one direction. To counteract this problem, we study three more variants.

A polygon is said to be monotone if it is possible to rotate the polygon such that any vertical line intersects the polygon boundary at most twice. This allows for flexibility in the shape of $R$, but it comes at a price: There are monotone shapes that do not look nice and compact at all. Therefore we introduce two other measures with a parameter that specifies how 'close to convex' a region is.

We define a backturn in a polygon as a turn of the polygon boundary that is made in the wrong direction. More formally, let $R$ be defined by the edges $\vec{e}_{1}, \vec{e}_{2}, \ldots, \vec{e}_{m}$ in counterclockwise order along its boundary, see Fig. 3. We define the turning angle from edge $\vec{e}_{i}$ to edge $\vec{e}_{j}$ for $i \neq j$ as $\alpha\left(\vec{e}_{i}, \vec{e}_{j}\right)=\sum_{k=i}^{j-1} \measuredangle\left(\vec{e}_{k}, \vec{e}_{k+1}\right)$ (indices wrap around, the sum is taken over at most $n-1$ angles $)$, where $\measuredangle\left(\vec{e}_{k}, \vec{e}_{k+1}\right) \in(-\pi, \pi)$


Fig. 1 (a) A meshed potato with a hole. We will refer to the unbounded face and to the holes as the exterior of the polygon. (b) The largest convex subpotato. (c) The largest monotone subpotato (the direction of monotonicity is shown with an arrow) (d) The largest backturn-constrained subpotato with a maximum backturn of $\approx \pi / 2$. (e) The largest subpotato with a total turning angle of at most $3 \pi$

Fig. 2 The largest convex subregion of a meshed polygon is restricted by the mesh, and not by the polygon boundary


Fig. 3 The maximum backturn (negative turning angle) of the polygon is $-\frac{2}{3} \pi$, and is achieved between $\vec{e}_{i}$ and $\vec{e}_{j}$

is the counterclockwise turn from $\vec{e}_{k}$ to $\vec{e}_{k+1}$ (a clockwise turn yields a negative $\left.\measuredangle\left(\vec{e}_{k}, \vec{e}_{k+1}\right)\right)$. When $i=j$, we set $\alpha\left(\vec{e}_{i}, \vec{e}_{i}\right)=0$. The maximum backturn of a simple polygon $R, \Phi_{R}$, is defined as $\min _{i, j} \alpha\left(\vec{e}_{i}, \vec{e}_{j}\right)$. It represents the maximum amount that part of the polygon boundary turns in the wrong direction. Note that $\Phi_{R} \leq 0$, and that a simple polygon $R$ is convex if and only if $\Phi_{R}=0$. When the polygon under consideration is clear, we will write $\Phi$ instead of $\Phi_{R}$.

Our last relaxation of convexity also deals with the idea of not being too concave. We define the total turning angle of a polygon as the sum of the absolute values of the turning angles at the vertices. We will look for a polygon whose total turning angle does not exceed a predefined constant $\beta \geq 2 \pi$. A convex polygon has a turning angle of exactly $2 \pi$, but we allow the polygon to turn 'the wrong way' a little bit, as long as the total turning angle is less than $\beta$.

Results We show that problem (i) can be solved in $O\left(n^{2}\right)$ time, problem (ii) in $O\left(n^{2}\right)$ time if the direction is given and in $O\left(n^{3}\right)$ time otherwise, problem (iii) in $O\left(n^{6}\right)$ time, and problem (iv) is NP-hard.

Related Work As mentioned earlier, if the largest-area simple polygon we aim to compute were not constrained to be a union of triangles of $\mathcal{M}$, but just to be contained in $\mathcal{P}$, problem (i) would turn into the potato peeling problem [14], also known as the convex skull problem [20]. The problem consists in, given a simple polygon $\mathcal{P}$, determining the largest-area convex polygon inside $\mathcal{P}$. If $\mathcal{P}$ has $n$ vertices, it can be solved in time $O\left(n^{7}\right)$ for a simple polygon, and in time $O\left(n^{8}\right)$ if the polygon has holes [7]. Approximation algorithms were studied by Hall-Holt et al. [15], who give an $O(1)$-approximation algorithm that runs in $O(n \log n)$ time. The orthogonal version of the problem, where both the given polygon and the subpolygon sought are bounded only by horizontal and vertical line segments, can be solved exactly in $O\left(n^{2}\right)$ time with an algorithm by Wood and Yap [21].

Also similar to the problems in this paper are many containment problems, in particular finding largest-area subpolygons. There is a large number of articles on different variants of this problem; we name just a few of them here. Barequet and Rogol [5] find a maximum area axially symmetric polygon contained in a given simple polygon. Daniels et al. [11] present an algorithm to find the largest-area axis-parallel rectangle in a general $n$-vertex polygon (allowing holes), which can be improved to run faster if the polygon is simple [6]. Another related problem is finding the largest similar copy of a polygonal pattern inside a 2 D environment with obstacles [8]. In
particular, several algorithms deal with the case when the pattern is a convex polygon [19]. Arkin et al. [4] study the problem of finding the smallest superpolygon of a given polygon, rather than the largest subpolygon. For convex polygons, this is simply the convex hull; they show that a smallest monotone or star-shaped superpolygon can be computed in $O(n / \log n)$ or $O\left(n^{2}\right)$ time, respectively.

Our problems (iii) and (iv) employ turning angles to capture a weaker notion of convexity. Turning angles have been studied before for optimization of geometric tours [1, 12] and also in contexts more similar to ours, for describing polygonal shapes. In particular, the closely-related turning functions have been used to compare the shape of polygons [2,16]. Closer to our application, Reinbacher et al. [18] consider minimizing the total angular change as a shape criterion to delineate imprecise regions with compact shapes.

Occlusion culling is a well known technique in computer graphics, and occluders that are large and convex give faster results [10]. Cohen-Or et al. [9] use this fact by covering a simple polygon with a few large convex subpolygons. They present a heuristic that first subdivides the polygon into convex regions and then selects a subset of those regions to be combined into a larger subpolygon. Papaioannou et al. [17] construct large convex occluders by grouping partially overlapping smaller occluders together into convex shapes.

Outline of the Paper The next section studies the first two problems: finding a largest-area polygon that is the union of triangles and is convex, and finding one that is monotone in some direction. Section 3 deals with problem (iii), where the restriction is now on the maximum backturn, and gives a polynomial-time algorithm to solve it. Section 4 studies problem (iv), finding a largest-area simple polygon with bounded total turning angle, and shows the problem is NP-hard. Finally, in Sect. 5 we make some concluding remarks and mention some problems left open.

Henceforth, we will refer to the holes in the polygon and to its unbounded face as the exterior of the polygon, or simply as the exterior. We assume that the triangles are closed, and therefore, that the exterior is open. Slightly abusing the notation, we will not distinguish between a subset of triangles and their union.

## 2 The Largest Convex Polygon

In this section we study the problems of computing the largest-area convex polygon and the largest-area monotone polygon formed as a union of triangles of $\mathcal{M}$. Most of the section is devoted to the first problem; adaptations needed for the second problem are discussed in the last subsection. Without loss of generality, we assume that no two vertices of $\mathcal{M}$ have the same $x$-coordinate. Before presenting the algorithm, we will introduce some notation and definitions.

Let $e \in \mathcal{M}$ be an edge of the triangulation. We denote by $X(e)$ the $x$-span of the edge, that is, the projection of the edge onto the $x$-axis. For edges $e, e^{\prime} \in \mathcal{M}$, we say that $e$ is above $e^{\prime}$ if there is a vertical line that meets $e$ at a point $p$ and $e^{\prime}$ at a point $p^{\prime}$ with $p$ above $p^{\prime}$. Assume that $e$ is above $e^{\prime}$. Define $X\left(e, e^{\prime}\right):=X(e) \cap X\left(e^{\prime}\right)$. We now define $T\left(e, e^{\prime}\right)$ as the trapezoid enclosed by $e$ from above, by $e^{\prime}$ from below, and


Fig. 4 (Color online) (a) A promising pair of edges and their trapezoid $T\left(e, e^{\prime}\right)$. (b) An admissible collection $C$ of triangles (gray) for the edge pair ( $e, e^{\prime}$ )
by the vertical lines that bound the interval $X\left(e, e^{\prime}\right)$. If $T\left(e, e^{\prime}\right)$ avoids the exterior of $\mathcal{P}$, we say that the pair ( $e, e^{\prime}$ ) is promising (see Fig. 4(a)).

Let $\left(e, e^{\prime}\right)$ be a promising pair of edges, and let $H\left(e, e^{\prime}\right)$ be the halfplane to the left of the line $x=\max X\left(e, e^{\prime}\right)$ (the same line that bounds $T\left(e, e^{\prime}\right)$ from the right). Let $C$ be a collection of triangles of $\mathcal{M}$. We call $C$ admissible for $\left(e, e^{\prime}\right)$ if the following conditions hold (see Fig. 4(b)):

- The pair $\left(e, e^{\prime}\right)$ is promising.
- $C$ contains the triangle below $e$, but not the triangle above $e$.
- $C$ contains the triangle above $e^{\prime}$, but not the triangle below $e^{\prime}$.
- Every triangle of $C$ intersects $H\left(e, e^{\prime}\right)$.
- The polygon $U=U(C):=\bigcup C \cap H\left(e, e^{\prime}\right)$ is convex, where $\bigcup C$ is the union of the triangles in $C$.

The last condition implies that $U$ cannot contain holes or any other part of the exterior. Also note that a pair of edges $\left(e, e^{\prime}\right)$ can have more than one collection of admissible triangles $C$, thus a pair of edges can have more than one polygon $U$.

We will now define a value $Q\left(e, e^{\prime}\right)$ for any pair of edges. Set $Q\left(e, e^{\prime}\right):=-\infty$ if ( $e, e^{\prime}$ ) has no admissible collection of triangles, otherwise $Q\left(e, e^{\prime}\right)$ is defined to be the largest area $(U(C))$ achieved by any admissible collection $C$. Clearly, any pair that is not promising has a $Q$-value of $-\infty$. The converse, however, is not true: a pair can be promising, but not have any admissible collection of triangles.

We will compute $Q\left(e, e^{\prime}\right)$ by dynamic programming. Notice that if $e$ and $e^{\prime}$ share their right endpoints, $Q\left(e, e^{\prime}\right)$ is precisely the area of the largest convex polygon consisting of triangles from $\mathcal{M}$ and having $e$ and $e^{\prime}$ as its rightmost top and bottom edges, respectively. Hence, examining $Q\left(e, e^{\prime}\right)$ over all pairs of edges with the same right endpoint, one can find the area of the desired largest-area convex polygon. The polygon itself can be extracted by standard dynamic programming methods.

We will first describe an algorithm to identify all the promising pairs in $O\left(n^{2}\right)$ time. After this, we show how to compute the value $Q$ for all those pairs, again in $O\left(n^{2}\right)$ time. The two algorithms rely on a plane sweep and can in principle be combined into a single pass, but we keep them separate for clarity.

### 2.1 Computing the Promising Pairs

We sweep $\mathcal{M}$ by a vertical sweep-line $\ell$ moving from left to right and stopping at every vertex. Using standard sweep techniques, we maintain the sorted order of the edges of $\mathcal{M}$ currently intersected by $\ell$. At a given point of the sweep, when we are at a sweep-line $\ell: x=t$, all the promising pairs with $X\left(e, e^{\prime}\right)$ lying to the left of $t$ have already been discovered. All pairs with $X\left(e, e^{\prime}\right) \ni t$ for which the part of $T\left(e, e^{\prime}\right)$ to the left of $\ell$ does not meet the exterior of $\mathcal{P}$ are called potentially promising and are currently being maintained.

Along the sweep, we maintain the following information. For each edge $e$ meeting the sweep-line, we keep two sorted lists $E_{a}(e)$ and $E_{b}(e) . E_{b}(e)$ contains all the edges $e^{\prime}$ meeting $\ell$ below $e$, which have their left endpoint to the left of that of $e$ and such that the intersection of $T\left(e, e^{\prime}\right)$ with the halfplane $x<t$ lies completely inside $\mathcal{P}$. $E_{b}(e)$ is stored sorted in downward direction along $\ell . E_{a}(e)$ is analogous but for the edges lying above $e$. In addition to these two lists, we also keep an inverted list $E_{I}(e)$ that contains all the edges $e^{\prime}$ such that $E_{a}\left(e^{\prime}\right)$ or $E_{b}\left(e^{\prime}\right)$ contains $e$. The need for these cross-references will become clear later.

We discern three different types of events that all occur when the sweep-line passes a vertex $v$. At this vertex, some edges end, which means potentially promising pairs using those edges are indeed promising and must be output and removed. Secondly, some new edges begin at $v$, which means that new potentially promising pairs must be generated and maintained. Thirdly, we may encounter parts of the exterior of $\mathcal{P}$ at $v$, which means that any pairs we are currently maintaining that have one edge above and one edge below $v$ are no longer potentially promising. We process the events in the order they are described here: we first process the beginning edges, then the ones ending, and finally the new exterior parts.

A New Beginning When a new edge $e$ begins at $v$, we want to find all potentially promising pairs of the form $\left(e, e^{\prime}\right)$ or $\left(e^{\prime}, e\right)$ that cross $\ell$. At this point, this just means that the exterior cannot intersect $\ell$ between them. To process multiple edges starting at $v$ correctly, we must treat them in the order in which these edges would cross $\ell$ if it were slightly to the right of $v$. We walk upwards along $\ell$, collecting all the edges crossing $\ell$ above $e$, and adding them to $E_{a}(e)$ (and the corresponding cross-references to the $E_{I}$ lists). This continues until a part of the boundary of $\mathcal{P}$ is reached: after that, no more promising pairs exist. In the same way, we walk downwards along $\ell$ to fill in $E_{b}(e)$. We spend constant time per added pair.

The End When an edge $e$ ends at $v$, we want to find all potentially promising pairs that use $e$ that we have stored, report them, and then remove them from the data structure. There are two types of such pairs: the ones for which the left endpoint of $e$ lies to the right of the left endpoint of the other edge are stored in the lists $E_{a}(e)$ and $E_{b}(e)$ of $e$. The other pairs, which have the left endpoint of $e$ to the left of the left endpoint of the other edge, are stored in the lists of the other edges. For these, we will use the cross-pointers in $E_{I}(e)$ to find them. In time proportional to the sizes of the lists stored with $e$, we find all such pairs. Since the lists are linked, we can remove the references in the lists of the other edges in constant time per entry, so we spend only constant time per removed pair.

Exterior Parts The last type of event occurs when a part of the exterior begins at the vertex $v$. Intuitively we want to remove all pairs $\left(e, e^{\prime}\right)$ for which $T\left(e, e^{\prime}\right)$ contains $v$. Each such no-longer-promising pair has a top edge $e$ above $v$ and a bottom edge $e^{\prime}$ below $v$. By the invariant we have been maintaining, we know that no other point from the exterior appears in $T\left(e, e^{\prime}\right) \cap\{x<t\}$. We walk up $\ell$, from $v$, looking for candidates $e$, until we hit a part of the exterior. For each $e$ found, we scan its list $E_{b}(e)$, from the tail backwards. If the last element $e^{\prime}$ is above $v$, the entire list is safe. If it is below $v$, the pair $\left(e, e^{\prime}\right)$ is not promising, and we drop $e^{\prime}$ from $E_{b}(e)$ and $e$ from $E_{I}\left(e^{\prime}\right)$, continuing upwards until the first $e^{\prime}$ above $v$. This is repeated until all no longer promising pairs have been discarded. Similarly, we walk down $\ell$ from $v$ and scan $E_{a}(e)$.

For each list that does not need to be truncated we spend only $O(1)$ time, for $O(n)$ lists at each vertex, taking in total $O\left(n^{2}\right)$ time. In addition, constant time is spent on each pair that is discarded, which can occur at most once over the execution of the algorithm for each pair, so we do at most $O\left(n^{2}\right)$ work overall. It follows that all the promising pairs can be computed in a total of $O\left(n^{2}\right)$ time.

### 2.2 Computing the $Q$-Values

Having established how to identify all promising pairs ( $e, e^{\prime}$ ), we explain how to compute the values $Q\left(e, e^{\prime}\right)$ for those pairs. Recall that this value corresponds to the area of a convex polygon bounded from the right by the line $x=\max X\left(e, e^{\prime}\right)$, from the top by $e$ and from the bottom by $e^{\prime}$, and following edges of $\mathcal{M}$ for the rest. We view this polygon as a concatenation of vertical strips, the regions $T(\cdot, \cdot)$ of promising pairs, which guarantees that only triangles, and no parts of the exterior, are used. This allows us to incrementally compute the $Q$-values by dynamic programming.

Fix a promising pair $\left(e, e^{\prime}\right)$. To compute $Q\left(e, e^{\prime}\right)$, consider the left endpoints of both segments. If they are the same point, the value is simply the area of the triangle delimited by $e, e^{\prime}$ and the vertical line $x=\max X\left(e, e^{\prime}\right)$. Otherwise, assume that the $x$-coordinate of the leftmost endpoint of $e$ lies to the left of that of $e^{\prime}$ (the other case is symmetric), and let $v$ be the left endpoint of $e^{\prime}$. Then $Q\left(e, e^{\prime}\right)=\operatorname{area}\left(T\left(e, e^{\prime}\right)\right)+$ $Q\left(e, e^{\prime \prime}\right)$, where $e^{\prime \prime}$ is an edge incident to $v=\left(v_{x}, v_{y}\right)$ with its other endpoint to the left of $v_{x}$ such that $\measuredangle\left(e^{\prime \prime}, e^{\prime}\right) \geq 0$ and $Q\left(e, e^{\prime \prime}\right)$ is maximized over all possible edges $e^{\prime \prime}$.

However, computing the values in this way we may spend up to $O\left(n^{3}\right)$ time, since for each promising pair there may be up to a linear number of candidate edges $e^{\prime \prime}$. We can avoid this by making the following observation.

We consider three sets of edges $E_{A}, E_{R}$ and $E_{L}$, which are defined for a fixed $v$. $E_{A}$ contains those edges that intersect $\ell$ above $v$ (with no intervening exterior points), $E_{R}$ contains all edges incident to and to the right of $v$, and $E_{L}$ contains all edges incident to and to the left of $v$, see Fig. 5(a). For each combination of an edge $e \in E_{A}$ and an edge $e^{\prime} \in E_{R}$ that form a promising pair, we must find the edge $e^{\prime \prime} \in E_{L}$ that optimizes $Q\left(e, e^{\prime \prime}\right)$. We refer to this edge $e^{\prime \prime}$ as best $\left(e, e^{\prime}\right)$.

Given an edge $e_{a} \in E_{R}$, we define the opposite edge of $e_{a}, \operatorname{opp}\left(e_{a}\right)$, as the edge in $E_{L}$ with smallest non-negative turning angle $\measuredangle\left(\operatorname{opp}\left(e_{a}\right), e_{a}\right) \geq 0$ (that is, it is convex), see Fig. 5(b). Note that best $\left(e, e^{\prime}\right)$ is always equal to or is above opp $\left(e^{\prime}\right)$ (that is, it has a higher left endpoint). We make the following observation.


Fig. 5 (a) The three sets of edges of interest when processing vertex $v$. (b) An edge $e_{a} \in E_{R}$ and its opposite edge, $\operatorname{opp}\left(e_{a}\right)$

Observation 1 The function best $\left(e, e^{\prime}\right)$, when defined, is monotone in the sense that if $e_{a}, e_{b} \in E_{R}$ are such that if $e_{a}$ is above $e_{b}$, then $\operatorname{best}\left(e, e_{b}\right)$ is equal to or above $\operatorname{best}\left(e, e_{a}\right)$.

We now describe how to organize the computation to use this observation, and make the whole algorithm run in quadratic time.

Sweep Algorithm We sweep the points from left to right, stopping at each vertex of $\mathcal{M}$. Let the current vertex be $v=\left(v_{x}, v_{y}\right)$. At this point, we will at once compute the values of $Q\left(e, e^{\prime}\right)$ for all promising pairs with $v_{x}=\min X\left(e, e^{\prime}\right)$, in time proportional to the degree of $v$. Recall that for each pair, $e$ lies above $e^{\prime}$. First we only consider pairs for which $e$ intersects $\ell$ above $v$, and $e^{\prime} \in E_{R}$. We separately handle the pairs where $e$ is incident to $v$ and $e^{\prime}$ intersects $\ell$ below $v$, in a symmetric way.

We compute for each $e \in E_{A}$ a list $L(e)$ of the edges in $E_{L}$ sorted by angle (we can assume the edges are stored in sorted order around the vertices). The first entry of the list corresponds to the topmost edge (smallest slope). We go through this list and store at each entry the maximum of the largest value of $Q\left(e, e^{\prime \prime}\right)$ seen so far and the value for the current edge (storing also appropriate pointers to allow reconstructing the solution later). Since the values of $Q\left(e, e^{\prime \prime}\right)$ that are needed have been already computed, we can fill in $L(e)$, for all $e \in E_{A}$, in $O\left(\left|E_{A}\right| \cdot\left|E_{L}\right|\right)$ time.

The second step consists in using this list to compute $Q\left(e, e^{\prime}\right)$ for all promising pairs of edges $e \in E_{A}$ and $e^{\prime} \in E_{R}$. In linear time, we can preprocess the mesh $\mathcal{M}$ to compute and store the opposite edge of each edge. Now, for each promising pair $\left(e, e^{\prime}\right)$, we look up $\operatorname{opp}\left(e^{\prime}\right)$ in the list $L(e)$. The edge $e^{\prime \prime}=\operatorname{best}\left(e, e^{\prime}\right)$ we are looking for is, by Observation 1, the one associated with the $Q$-value stored with $\operatorname{opp}\left(e^{\prime}\right)$. (Note that best $\left(e, e^{\prime}\right)$ may be different from $\operatorname{opp}\left(e^{\prime}\right)$.) Therefore we set $Q\left(e, e^{\prime}\right)=$ $Q\left(e, e^{\prime \prime}\right)+\operatorname{area}\left(T\left(e, e^{\prime}\right)\right)$.

All the $Q$-values can be computed in quadratic time in total: for a vertex $v$, of degree $\operatorname{deg}(v), O(\operatorname{deg}(v) \cdot n)$ time is spent on filling the list of all edges crossing $\ell$ above $v$, which leads to quadratic time for all vertices. In addition, constant time is spent to fill in each entry of $Q\left(e, e^{\prime}\right)$. The total running time is hence quadratic. Together with the algorithm for computing the promising pairs, we conclude with the following theorem:

Theorem 1 Given a polygon with holes $\mathcal{P}$, and a triangular mesh $\mathcal{M}$ with $n$ vertices that covers the interior of $\mathcal{P}$, a maximum-area convex subpolygon that is the union of triangles of $\mathcal{M}$ can be computed in $O\left(n^{2}\right)$ time.

### 2.3 Monotone Polygons

Turning to problem (ii), we note that the algorithm given above can easily be adapted to compute the largest-area polygon that is monotone in a given direction. Assume without loss of generality that the given direction is horizontal. Then computing $Q\left(e, e^{\prime}\right)$ remains the same for edges with shared starting or ending endpoints. When "continuing through" from edge $e^{\prime}$ to $e^{\prime \prime}$, however, we need to remove the condition that the slope of $e^{\prime \prime}$ must be less than that of $e^{\prime}$. When the direction of monotonicity is not given, we can check all $O(n)$ directions perpendicular to the edges of $\mathcal{M}$. To summarize, we obtain:

Theorem 2 Given a polygon with holes $\mathcal{P}$, and a triangular mesh $\mathcal{M}$ with $n$ vertices that covers the interior of $\mathcal{P}$, a maximum-area monotone subpolygon that is the union of triangles of $\mathcal{M}$ can be computed in $O\left(n^{2}\right)$ time if the direction is given, or in $O\left(n^{3}\right)$ time otherwise.

## 3 Maximum Negative Turning Angle

Next we study the problem of finding the collection of triangles of largest total area such that their union $R$ is a simple polygon and the maximum backturn $\Phi$ of the boundary of $R$ is bounded by some constant $0 \leq \gamma \leq \pi$. The backturn is defined as $\Phi=\min _{i, j} \alpha\left(\vec{e}_{i}, \vec{e}_{j}\right)$; see Section 1 for more details. The goal is to find a largest area subregion that does not include any part of the exterior and has $\Phi \geq-\gamma$. Note that it follows from the definition of $\Phi$ that there are $O\left(n^{2}\right)$ relevant values for $\gamma$.

For a given path from vertex $v$ to vertex $u$, we define the extreme edges $\vec{s}_{v u}$ and $\vec{t}_{v u}$ as follows: $\vec{s}_{v u}$ is the edge of the path with the largest turning angle with respect to the first edge of the path, whereas $\vec{t}_{v u}$ is the edge with the smallest one, see Fig. 6. Extreme edges will play an important role in the algorithm. We make the following simple observation, illustrated in Fig. 7.

Observation 2 Let $\left\{\vec{e}_{1}, \ldots, \vec{e}_{m}\right\}$ be a path from $u$ to $v$, and let $\vec{s}_{v u}$ and $\vec{t}_{v u}$ be the extreme edges of the path. Then $\alpha\left(\vec{s}_{v u}, \vec{e}_{m}\right) \leq 0$ and $\alpha\left(\vec{e}_{1}, \vec{t}_{v u}\right) \leq 0$.

We present a dynamic programming algorithm that computes a function $Q$ which assigns to each subproblem the value of an optimal solution of the subproblem. A sub-

Fig. 6 A path from $v$ to $u$ showing the turning angle of each edge, with respect to the first edge of the path. The extreme edges are the ones with the maximum and minimum turning angle



Fig. 7 Example illustrating the two possible cases for Observation 2: (a) $\vec{s}_{v u}$ appearing after $\vec{t}_{v u}$, and (b) $\vec{s}_{v u}$ appearing before $\vec{t}_{v u}$. The total turning angle of each of the highlighted subpaths, $\alpha\left(\vec{e}_{1}, \vec{t}_{v u}\right)$ and $\alpha\left(\vec{s}_{v u}, \vec{e}_{m}\right)$, is always less than or equal to zero


Fig. 8 (Color online) Schematic view of an instance ( $\left.v, u, \vec{e}_{v, \text { out }}, \vec{e}_{u, \text { in }}, \sigma, \tau\right)$ of the algorithm. Vertices $u$ and $v$, together with directions $\tau$ and $\sigma$, separate the plane into two areas. The feasible region $F_{I}$, in white, contains the solution to the instance. The gray arrows indicate the orientation of $\sigma$ and $\tau$
problem is defined on a pair of vertices, a pair of edges that are incident to those vertices, and two angle bounds for the extreme edges. Let $\left(v, u, \vec{e}_{v, \text { out }}, \vec{e}_{u, \text { in }}, \sigma, \tau\right)$ be such an instance; Fig. 8 shows a visual representation.

The goal is to find a path from $v$ to $u$ that uses $\vec{e}_{v, \text { out }}$ as the first edge and $\vec{e}_{u, \text { in }}$ as the last edge, and such that the turning angle of the extreme edges $\vec{s}_{v u}$ and $\vec{t}_{v u}$ of the path are bounded by $\sigma$ and $\tau$, respectively. More precisely, we require $\alpha\left(\vec{e}_{v, \text { out }}, \vec{t}_{v u}\right) \geq-\tau$ and $\alpha\left(\vec{s}_{v u}, \vec{e}_{u \text {,in }}\right) \geq-\sigma$ (for $\sigma, \tau \geq 0$ ). In other words, we are limiting the turning angle from the first edge to the edge with smallest angle, and from the edge with largest angle to the last edge. Since $\vec{s}_{v u}$ and $\vec{v}_{v u}$ are extreme edges, this also limits the angles of all the other edges on the path from $v$ to $u$. With some abuse of notation, we sometimes consider $\sigma$ and $\tau$ as directions, and sometimes as angles. Note that $\sigma$ and $\tau$ restrict the extreme directions of edges of our solution polygon $R$, so the only relevant values they take are the linear number of directions of the edges of $\mathcal{M}$.

Moreover, we require the polygon consisting of this path and completed by $\overrightarrow{u v}$ to be simple, to contain no parts of the exterior, and to have optimal area. Note that $\overrightarrow{u v}$ is the only edge of this polygon that is not necessarily an edge of the triangulation.

### 3.1 Description of the Algorithm

The algorithm is based on the fact that the optimal polygon $R^{*}$ has a triangulation (not to be confused with the given mesh $\mathcal{M}$ of the input polygon), and we can use it to recursively define the solution to an instance in terms of the solutions to smaller instances. In particular, the path from $v$ to $u$ in the instance $I=\left(v, u, \vec{e}_{v, \text { out }}, \vec{e}_{u, \text { in }}, \sigma, \tau\right)$ is composed of a path from $v$ to some vertex $w$ and a path from $w$ to the vertex $u$, and such that $\overrightarrow{v w}$ and $\overrightarrow{w u}$ are diagonals of a triangulation of the solution to the instance $I$

Fig. 9 Example showing the decomposition of the problem into two smaller subproblems

(or, in some cases, edges of the solution itself), see Fig. 9. We can maximize over all vertices $w$, the edges incident to $w$ that are used on the paths, and the way the bounds on the angles of the extreme edges may occur in the subproblems, as long as the solutions do not contain part of the exterior and the angle bounds are respected.

Consider an optimal solution $R^{*}$ and a triangulation of $R^{*}$. Let $\overrightarrow{u v}$ be a diagonal of this triangulation. We define $Q$ as follows:

$$
\begin{aligned}
& Q\left(v, u, \vec{e}_{v, \text { out }}, \vec{e}_{u, \text { in }}, \sigma, \tau\right) \\
& =\max _{w, \vec{e}_{w, \text { out }}, \vec{e}_{w, \text { in }}, \sigma^{\prime}, \tau^{\prime}}\left(Q\left(v, w, \vec{e}_{v, \text { out }}, \vec{e}_{w, \text { in }}, \sigma^{\prime}, \tau\right)\right. \\
& \left.\quad+Q\left(w, u, \vec{e}_{w, \text { out }}, \vec{e}_{u, \text { in }}, \sigma, \tau^{\prime}\right)+\operatorname{area}(\Delta u v w)\right)
\end{aligned}
$$

as long as the following conditions hold:
(1) $\alpha\left(\vec{e}_{v, \text { out }}, \vec{e}_{w, \text { in }}\right) \geq-\tau$ and $\alpha\left(\vec{e}_{w, \text { out }}, \vec{e}_{u, \text { in }}\right) \geq-\sigma$.
(2) $\sigma^{\prime}+\tau^{\prime}-\measuredangle\left(\vec{e}_{w, \text { in }}, \vec{e}_{w, \text { out }}\right) \leq \gamma$.
(3) $\tau^{\prime} \leq \tau+\alpha\left(\vec{e}_{v, \text { out }}, \vec{e}_{w, \text { out }}\right)$ and $\sigma^{\prime} \leq \sigma+\alpha\left(\vec{e}_{w, \text { in }}, \vec{e}_{u, \text { in }}\right)$.
(4) $\Delta u v w$ is contained in $\mathcal{P}$, and $w$ lies to the left of $\overrightarrow{u v}$.

The first condition ensures that $\vec{e}_{w, \text { in }}$ and $\vec{e}_{w, \text { out }}$ are within the bounds imposed by $\tau$ and $\sigma$. The next two conditions guarantee that $\alpha\left(\vec{s}_{v u}, \vec{t}_{v u}\right) \geq-\gamma$ : if both $\vec{s}_{v u}$ and $\vec{t}_{v u}$ lie in the same subproblem, then the recursion ensures their turning angle is within the right bound. Otherwise, if they lie in different subproblems, then the second condition relates the two angle bounds of the subproblems and the angle at $w$ with $\gamma$, whereas the third condition enforces that the angle bounds of the subproblems do not exceed the existing ones. Together they guarantee that $\alpha\left(\vec{s}_{v u}, \vec{t}_{v u}\right) \geq-\gamma$. Last but not least, the fourth condition makes sure that the final solution will not contain parts of the exterior.

If no vertex $w$ exists that satisfies these conditions, then $Q(\cdot)$ is defined to be $-\infty$. The base case is $Q(v, u, \overrightarrow{v u}, \overrightarrow{v u}, \sigma, \tau)$ and has value zero for all combinations of $v$, $u, \sigma$ and $\tau$ (of course, $\overrightarrow{v u}$ must be an edge of $\mathcal{M}$ ).

Now we are ready to give a description of the algorithm. In Sect. 3.2 we argue that the algorithm indeed returns a solution to our problem, and in Sect. 3.3 we analyze the running time. We maintain a 6-dimensional table with one entry for each instance $I=\left(v, u, \vec{e}_{v, \text { out }}, \vec{e}_{u, \text { in }}, \sigma, \tau\right)$. In each cell of the table, we store the value $Q(I)$, as well as pointers to the entries of the two subproblems that give this value (unless $I$ is a base
case). For each entry, we compute the value by trying all possibilities for $w, \vec{e}_{w, \text { in }}$, and $\vec{e}_{w, \text { out }}$, looking up their solutions in the table and checking whether conditions (1) to (4) hold. We start by filling in all base case entries of the table. If we fill the remaining entries in any order that complies with the partial order induced by the inclusion of feasible regions, then we are guaranteed that the subproblems we need to look up have already been solved before (this is proven in the next section).

### 3.2 Correctness of the Algorithm

In this section, we show that the solution returned by the algorithm is an optimal solution to our problem, and that the algorithm always terminates. The following lemma shows that the solutions computed by the dynamic programming algorithm respect the angle constraint.

Lemma 1 Let $I=\left(v, u, \vec{e}_{v, \text { out }}, \vec{e}_{u, \text { in }}, \sigma, \tau\right)$ be an instance of the dynamic programming algorithm above, such that $Q(I)>-\infty$ and $\sigma+\tau-\alpha\left(\vec{e}_{u, \text { in }}, \vec{e}_{v, \text { out }}\right) \leq \gamma$. Let $R$ be the polygon defined by $\overrightarrow{u v}$ and the path from $v$ to $u$ associated with $Q(I)$. Then $\Phi_{R} \geq-\gamma$.

Proof Let $R$ be the polygon whose boundary is the concatenation of $\overrightarrow{u v}$ and $\rho_{v u}$, where $\rho_{v u}$ is the path associated with $Q(I)$. Assume $I$ does not fall into the base case (otherwise the result holds immediately). Since it is the result of applying the dynamic programming algorithm, there must be a vertex $w$ such that $\rho_{v u}$ is the concatenation of $\rho_{v w}$ and $\rho_{w u}$, where $\rho_{v w}$ and $\rho_{w u}$ are the paths which resulted from evaluating $Q\left(v, w, \vec{e}_{v, \text { out }}, \vec{e}_{w, \text { in }}, \sigma^{\prime}, \tau\right)$ and $Q\left(w, u, \vec{e}_{w, \text { out }}, \vec{e}_{u, \text { in }}, \sigma, \tau^{\prime}\right)$, for some combination of $\vec{e}_{w, \text { in }}, \vec{e}_{w, \text { out }}, \sigma^{\prime}$ and $\tau^{\prime}$.

Let $\vec{e}_{k}, \vec{e}_{l}$ be a pair of edges with the maximum backturn in $R$. We will show that $\alpha\left(\vec{e}_{k}, \vec{e}_{l}\right) \geq-\gamma$. We distinguish between several cases depending on where $\vec{e}_{k}$ and $\vec{e}_{l}$ lie:

- Both $\vec{e}_{k}$ and $\vec{e}_{l}$ lie in $\rho_{v w}$ or both lie in $\rho_{w u}$. Then both edges appear in the same subproblem, and it follows by induction that $\alpha\left(\vec{e}_{k}, \vec{e}_{l}\right) \geq-\gamma$.
- $\vec{e}_{k} \in \rho_{v w}$ and $\vec{e}_{l} \in \rho_{w u}$. Then $w$ lies between $\vec{e}_{k}$ and $\vec{e}_{l}$ (see Fig. 10(a)). This implies that $\vec{s}_{v w}=\vec{e}_{k}$ and $\vec{t}_{w u}=\vec{e}_{l}$. To see why, assume that $\vec{t}_{w u} \neq \vec{e}_{l}$. By definition


Fig. 10 Illustrations for proof of Lemma 1. The path associated with $\alpha\left(\vec{e}_{k}, \vec{e}_{l}\right)$ is highlighted. (a) Case $\vec{e}_{k} \in \rho_{v w}$ and $\vec{e}_{l} \in \rho_{w u}$. (b) Case $\vec{e}_{k} \in \rho_{w u}$ and $\vec{e}_{l} \in \rho_{u w}$
of $\vec{t}_{w u}, \vec{t}_{w u}$ achieves the minimum turning angle between $\vec{e}_{w, \text { out }}$ and any edge in $\rho_{w u}$. Therefore $e_{l} \neq \vec{t}_{w u}$ immediately implies $\alpha\left(\vec{e}_{k}, \vec{t}_{w u}\right)<\alpha\left(\vec{e}_{k}, \vec{e}_{l}\right)$, leading to a contradiction. A similar argument shows that $\vec{s}_{v w}=\vec{e}_{k}$.

The turning angle $\alpha\left(\vec{e}_{k}, \vec{e}_{l}\right)$ can be decomposed into $\alpha\left(\vec{e}_{k}, \vec{e}_{l}\right)=\alpha\left(\vec{s}_{v w}, \vec{t}_{w u}\right)=$ $\alpha\left(\vec{s}_{v w}, \vec{e}_{w, \text { in }}\right)+\measuredangle\left(\vec{e}_{w, \text { in }}, \vec{e}_{w, \text { out }}\right)+\alpha\left(\vec{e}_{w, \text { out }}, \vec{t}_{w u}\right)$. From the definition of the angle bounds $\sigma^{\prime}$ and $\tau^{\prime}$ we have $\alpha\left(\vec{s}_{v w}, \vec{e}_{w, \text { in }}\right) \geq-\sigma^{\prime}$ and $\alpha\left(\vec{e}_{w, \text { out }}, \vec{t}_{w u}\right) \geq-\tau^{\prime}$. In combination with condition (2) they yield $\alpha\left(\vec{s}_{v w}, \vec{t}_{w u}\right) \geq-\gamma$.

- $\vec{e}_{k} \in \rho_{w u}$ and $\vec{e}_{l} \in \rho_{v w}$. Similar to the previous case, it is easy to verify that $s_{w u}=\vec{e}_{k}$ and $t_{v w}=\vec{e}_{l}$ (see Fig. 10(b)). The path from $\vec{s}_{v w}$ to $\vec{t}_{v w}$, with total angular change $\alpha\left(\vec{e}_{k}, \vec{e}_{l}\right)$, can be split into three parts: $\vec{s}_{w u}$ to $\vec{e}_{u, \text { in }}, \vec{e}_{u, \text { in }}$ to $\vec{e}_{v, \text { out }}$, and $\vec{e}_{v, \text { out }}$ to $\vec{t}_{v w}$.

From the definition of the angle bounds $\sigma$ and $\tau$ we have that $\alpha\left(\vec{s}_{w u}, \vec{e}_{u, \text { in }}\right) \geq$ $-\sigma$ and $\alpha\left(\vec{e}_{v, \text { out }}, \vec{v}_{v w}\right) \geq-\tau$. Combining these with the lemma hypothesis $\sigma+\tau-$ $\alpha\left(\vec{e}_{u, \text { in }}, \vec{e}_{v, \text { out }}\right) \leq \gamma$, we get $\alpha\left(\vec{s}_{w u}, \vec{t}_{v w}\right)=\alpha\left(\vec{e}_{k}, \vec{e}_{l}\right) \geq-\gamma$.

- $\vec{e}_{k}=\overrightarrow{u v}$ and $\vec{e}_{l} \in \rho_{v u}$ (or vice versa). The lemma hypothesis $\sigma+\tau-$ $\alpha\left(\vec{e}_{u, \text { in }}, \vec{e}_{v, \text { out }}\right) \leq \gamma$ ensures that angles between $\overrightarrow{u v}$ and $\vec{e}_{l}$ that respect $\sigma$ and $\tau$ will result in a total angular change $\geq-\gamma$, therefore $\alpha\left(\vec{e}_{k}, \vec{e}_{l}\right) \geq-\gamma$ follows.

Therefore in all cases $\alpha\left(\vec{e}_{k}, \vec{e}_{l}\right)=\Phi_{R} \geq-\gamma$, and the result follows.
The previous result implies that the polygon induced by the solution of a subproblem must be simple (that is, without self-intersections).

Corollary 1 The polygon R, as defined in Lemma 1, is simple.

Proof Suppose that the solution to $I$ uses vertex $w$ and the solutions to both subproblems of $I$ are simple. Since $\vec{e}_{w, \text { in }}$ comes before $\vec{e}_{w, \text { out }}$, and since the total backturn is not more than $\pi$, the solutions do not intersect except at $w$, and therefore the solution to $I$ is also simple.

We further observe that the definition of the recursive algorithm, together with conditions (1) to (4), imply that the final solution does not contain any part of the exterior (that is, it is valid), and that its area is maximum.

It remains to show that the subproblems needed to solve a given instance $I$ of the recursive algorithm are in some sense smaller than $I$, and therefore can be assumed to have been solved in earlier steps. This can be seen in the following way. For any subproblem $I=\left(v, u, \vec{e}_{v, \text { out }}, \vec{e}_{u, \text { in }}, \sigma, \tau\right)$, we can define a feasible region $F_{I}$ that must contain any possible solution to that subproblem. Consider the oriented Zorro consisting of the halfline to $u$ in the direction of $\sigma$, the line segment $\overrightarrow{u v}$, and the halfline from $v$ in the direction of $\tau$ (see Fig. 8). The feasible region $F_{I}$ is the region to the left of this Zorro.

Lemma 2 In the recursive step for subproblem $I=\left(v, u, \vec{e}_{v, \text { out }}, \vec{e}_{u, \text { in }}, \sigma, \tau\right)$, the feasible regions of the subproblems are contained in the feasible region of $F_{I}$. Furthermore, the number of vertices of the triangulated polygon inside the region strictly decreases.


Fig. 11 (a) The feasible regions of the two subproblems are both contained in the feasible region of the original problem. (b-c) The two situations in which the left subproblem has a feasible region that is not contained in the original problem. Case (b) violates condition (3) ( $\sigma^{\prime}$ is restricted by $\sigma$ ). Case (c) has a turning angle between $\vec{s}_{v w}$ and $\vec{e}_{w, \text { out }}$ that is smaller than $-\pi$

Proof To prove the first part of the lemma, assume for contradiction that there is a subproblem $J$ whose feasible region $F_{J}$ is not contained in $F_{I}$. Further assume without loss of generality that this happens with subproblem $J=\left(v, w, \vec{e}_{v, \text { out }}\right.$, $\left.\vec{e}_{w, \text { in }}, \sigma^{\prime}, \tau\right)$. Clearly, the halflines from $v$ in the Zorros of $F_{I}$ and $F_{J}$ are the same, and $w$ must lie inside the feasible region $F_{I}$ so the middle segment of the Zorro of $F_{J}, v \vec{w}$, is always inside $F_{I}$. If $F_{J}$ is not contained in $F_{I}$, this means that the halfline of $F_{J}$ from $w$ intersects the Zorro of $F_{I}$.

There are essentially two ways in which $F_{J}$ can intersect $F_{I}$, shown in Fig. 11. Since the direction $\sigma^{\prime}$ is restricted by $\sigma$ (by condition (3)), the halfline from $w$ in direction $\sigma^{\prime}$ cannot intersect the halfline from $u$ in direction $\sigma$ without also intersecting the middle segment of the Zorro of $F_{I}, \overrightarrow{v u}$. Thus the case shown in Fig. 11(b) cannot arise.

To see that other possibility, shown in Fig. 11(c), cannot happen as part of a valid problem, assume that there is an edge on the path from $v$ to $w$ that achieves a direction of $\sigma^{\prime}$. If that is not the case, then we can discard this particular subproblem, since there will be another one with the same solution and smaller $\sigma^{\prime}$ (that is, more clockwise). Therefore we can assume that $\sigma^{\prime}$ is attained by some edge, in particular, by $\vec{s}_{v w}$. Then $\alpha\left(\vec{s}_{v w}, \vec{e}_{w, \text { out }}\right) \leq \alpha\left(\vec{s}_{v w}, \overrightarrow{w u}\right)<-\pi$, violating the angle constraint of $\Phi \geq-\gamma \geq-\pi$. It follows that such a solution does not fulfill Conditions (1) to (4) for $I$, thus it will never be considered by the algorithm.

To prove the second part of the lemma, we note that from the same arguments it follows that $v$ cannot be in the interior of the feasible region of $\left(w, u, \vec{e}_{w, \text { out }}, \vec{e}_{u, \text { in }}, \sigma, \tau^{\prime}\right)$, and $u$ cannot be in the feasible region of $\left(v, w, \vec{e}_{v, \text { out }}\right.$, $\left.\vec{e}_{w, \text { in }}, \sigma^{\prime}, \tau\right)$. Therefore the number of vertices in the subproblems decreases.

Because of Lemma 2, the feasible regions define a partial order on the subproblems in which we can evaluate them with the guarantee that all information to solve the current subproblem is available. Since the total number of subproblems is finite (we only need to check finitely many directions for $\tau$ and $\sigma$ since they must be aligned with the point set), the algorithm terminates.

Fig. 12 (Color online) Example showing the decomposition of the problem into two smaller subproblems (darker gray)


### 3.3 Running Time Analysis

In order to analyze the running time of the algorithm, some more detail on the steps to be performed is needed. Recall that a subproblem is specified by a pair of vertices, a pair of edges that are incident to those vertices, and two angle bounds. Let ( $\left.v, u, \vec{e}_{v, \text { out }}, \vec{e}_{u, \text { in }}, \sigma, \tau\right)$ denote the current problem.

Since $\overrightarrow{u v}$ is a diagonal of a triangulation of $R^{*}$, there must be a triangle $\Delta u v w$ to the left of $\overrightarrow{u v}$. We will try all the possible candidates for $w$. Each triangle $\Delta u v w$ subdivides the problem into two new subproblems (see Fig. 12).

Any pair of subproblems that needs to be combined must fulfill conditions (1) to (4). This can be easily checked in constant time for the first three conditions. In order to perform the check for the fourth condition in constant time, we need to do some preprocessing first. We can trivially check and store this for any triple of vertices in $O\left(n^{4}\right)$ time and $O\left(n^{3}\right)$ space. If one of the conditions does not hold, we assign the value $-\infty$ to this instance.

When maximizing over the values of $\sigma^{\prime}$ and $\tau^{\prime}$, we only need to take the combinations into account where $\sigma^{\prime}+\tau^{\prime}-\measuredangle\left(\vec{e}_{w, \text { in }}, \vec{e}_{w, \text { out }}\right)$ is the largest possible value that is still smaller or equal to $\gamma$. As observed earlier, $\sigma^{\prime}$ and $\tau^{\prime}$ can take on a linear number of values. So, for each value of $\sigma^{\prime}$ there is only one value of $\tau^{\prime}$ to consider, so there are only a linear number of combinations possible for choosing $\sigma^{\prime}$ and $\tau^{\prime}$ together. Note that the two subproblems are independent as long as $\gamma \leq \pi$. So, together with the choice of $w$, each instance is solved by taking the maximum over $O\left(n^{2}\right)$ combinations of smaller instances.

As for the overall running time of the recursive algorithm, the base case can be resolved in constant time. For the general case we have a table with $O\left(n^{4}\right)$ entries: there are a linear number of possibilities for the pair ( $v, \vec{e}_{v, \text { out }}$ ), a linear number for the pair ( $u, \vec{e}_{u \text {, in }}$ ), a linear number for $\sigma$, and a linear number for $\tau$. We can fill in an entry in quadratic time, so we need $O\left(n^{6}\right)$ time in total.

Theorem 3 Given a constant $0 \leq \gamma \leq \pi$, a polygon with holes $\mathcal{P}$, and a triangular mesh $\mathcal{M}$ with $n$ vertices that covers the interior of $\mathcal{P}$, a maximum-area polygon that is the union of triangles of $\mathcal{M}$ such that the maximum negative turning angle (or backturn) $\Phi$ is $\Phi \geq-\gamma$ can be computed in $O\left(n^{6}\right)$ time.

## 4 The Largest-Area Polygon with Bounded Turning Angle

We now consider the problem of finding the largest-area simple polygon formed as a union of triangles of $\mathcal{M}$, so that the total angular change of the perimeter is at most some constant $\beta$. We show that this problem is NP-hard for a total angular change of, say, $3 \pi$. The reduction is from the NP-hard problem KNAPSACK [13]: Given a set of $n$ integer pairs $\left(c_{1}, p_{1}\right), \ldots,\left(c_{n}, p_{n}\right)$, where $c_{i}$ is the size of the $i^{\text {th }}$ item and $p_{i}$ is its profit, and a maximum capacity $C$, select the subset of maximum total profit of which the total size is at most $C$. Our reduction is similar to the one used in [3]. Without loss of generality, we scale the sizes $c_{1}, \ldots, c_{n}$ and $C$ so that $C=\pi$.

We construct a triangulated polygon as follows, see Fig. 13. Take a pair $\left(c_{i}, p_{i}\right)$. If $c_{i}>C=\pi$ then we discard it. Otherwise we construct an isosceles triangle $T_{i}$ with a horizontal base of width $2 / \tan \left(c_{i} / 4\right)$, and its top vertex $m_{i}$ at height 1 above the middle of the base. Scale $T_{i}$ so that its area is $2 p_{i}$. The total angular change of a path that arrives horizontally from the left at the lower left corner of $T_{i}$, goes diagonally up to $m_{i}$, and then diagonally down to the lower right corner of $T_{i}$, and then continues horizontally to the right, is $c_{i}$, and the area in $T_{i}$ and below the path is $p_{i}$. This corresponds to choosing the $i$ th item in the knapsack: the profit (area) is $p_{i}$ and the cost (extra angular change) is $c_{i}$. Now we combine all these triangles so that all lower sides lie on a horizontal line $\ell$, spaced sufficiently far apart; their order is irrelevant. Below the line $\ell$ we make large triangles whose union is a large rectangle $R$ with the following properties:

- The top side of $R$ is on $\ell$.
- The upper left corner of $R$ is strictly to the left of all the triangles $T_{i}$.
- The upper right corner of $R$ is strictly to the right of all the triangles $T_{i}$.
- Every triangle in the polygon triangulation of $R$ has area greater than $\sum_{i=1}^{n} p_{i}$ (note that $R$ has two vertices of every triangle $T_{i}$ on its top side).

The first three conditions can be trivially satisfied, whereas the last one can be easily fulfilled by making $R$ very tall. By construction, a maximum area solution will contain at least all triangles inside $R$, because leaving one out cannot be compensated even by using all the $T_{i}$ 's. The total angular change of $R$ is exactly $2 \pi$, and we can choose any remaining triangle $T_{i}$, at the cost of $c_{i}$, which gives an extra area $p_{i}$. Hence, the subset of items with maximum profit in the knapsack with total size at most $\pi$ corresponds precisely to the subset of extra triangles.


Fig. 13 Reduction of knapsack to maximum area subtriangulation with bounded total angular change

Although we assumed that the weights of the knapsack items can be scaled to add up to $\pi$ and we compute coordinates of the form $2 / \tan \left(c_{i} / 4\right)$, it is not essential that these are computed exactly. Knapsack is already NP-hard for items with integer values, so if we use these we can afford to make small errors in the placements of the points, as long as they do not add up to something more than 1. Instead of scaling the weights to add up to $\pi$, we make them add up to a rational number slightly smaller than $\pi$. Still we cannot place the points exactly, but if we make sure the errors are not more than $O(1 / n)$ we will not influence the solutions of the knapsack problem. This can increase the description complexity of the points by at most a factor of $O(n)$.

Theorem 4 Given a polygon with holes $\mathcal{P}$, and a triangular mesh $\mathcal{M}$ that covers the interior of $\mathcal{P}$, computing a maximum area subpolygon, comprised of triangles of $\mathcal{M}$, with total angular change at most $3 \pi$ is NP-hard.

## 5 Conclusions

We studied various area optimization problems on triangulated polygons with holes, where the boundary of the optimal simple polygon is restricted to the triangle boundaries. We focused on convexity as a way to form well-shaped areas, and given that convexity is a strong restriction in combination with forcing the polygon to be composed only of whole triangles, we also studied three different relaxations of convexity. For three versions we gave polynomial algorithms and for one, an NP-hardness proof. It is worth to note that our meshed version of the potato peeling problem, by restricting the convex subpolygon to be composed of whole triangles, allows for a much faster algorithm than the unrestricted version: $O\left(n^{2}\right)$ time against the $O\left(n^{8}\right)$ time needed for the original potato peeling problem.

This paper opens up several directions for further research. Regarding the quadratic-time algorithm of Sect. 2, a natural question is whether a faster algorithm exists or a lower bound can be proved. Improvements should be possible for the algorithm of Sect. 3. The current algorithm is rather involved and has a high running time. It would be interesting to know if there is a different way to approach it that results in a simpler or faster algorithm. The NP-hardness proof of Sect. 4 leads to the question of approximation algorithms, a topic also worth studying.

Studying other relaxations of convexity or different properties that result in wellshaped subregions is also a promising direction of further research. Moreover, several other extensions are possible, such as studying the problem in three dimensions. That is, given a tetrahedralization of a solid, possibly with holes, find a subset of the tetrahedra that forms a convex (or nearly convex) solid. This would have applications, for example, in finding large occluders for hidden surface removal algorithms [9]. The algorithms presented here clearly serve as building blocks for the 3D version, although a deeper study of the problem is needed.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

## References

1. Aggarwal, A., Coppersmith, D., Khanna, S., Motwani, R., Schieber, B.: The angular-metric traveling salesman problem. SIAM J. Comput. 29, 697-711 (1999)
2. Arkin, E.M., Chew, L.P., Huttenlocher, D.P., Kedem, K., Mitchell, J.S.B.: An efficiently computable metric for comparing polygonal shapes. IEEE Trans. Pattern Anal. Mach. Intell. 13(3), 209-216 (1991)
3. Arkin, E.M., Khuller, S., Mitchell, J.S.B.: Geometric knapsack problems. Algorithmica 10, 399-427 (1993)
4. Arkin, E.M., Chiang, Y.-J., Held, M., Mitchell, J.S.B., Sacristán, V., Skiena, S.S., Yang, T.-C.: On minimum-area hulls. Algorithmica 21, 119-136 (1998)
5. Barequet, G., Rogol, V.: Maximizing the area of an axially symmetric polygon inscribed in a simple polygon. Comput. Graph. 31(1), 127-136 (2007)
6. Boland, R., Urrutia, J.: Finding the largest axis aligned rectangle in a polygon in $O(n \log n)$ time. In: Proceedings of the 13th Canadian Conference on Computational Geometry (CCCG'01), pp. 41-44 (2001)
7. Chang, J.S., Yap, C.K.: A polynomial solution for the potato-peeling problem. Discrete Comput. Geom. 1, 155-182 (1986)
8. Chazelle, B.: The polygon containment problem. In: Preparata, F.P. (ed.) Computational Geometry. Adv. Comput. Res., vol. 1, pp. 1-33. JAI Press, Greenwich (1983)
9. Cohen-Or, D., Lev-Yehudi, S., Karol, A., Tal, A.: Inner-cover of non-convex shapes. Int. J. Shape Model. 9, 223-238 (2003)
10. Coorg, S.R., Teller, S.J.: Real-time occlusion culling for models with large occluders. In: Symposium on Interactive 3D Graphics, pp. 83-90, 189, (1997)
11. Daniels, K., Milenkovic, V., Roth, D.: Finding the maximum area axis-parallel rectangle in a polygon. Comput. Geom. Theory Appl. 7, 125-148 (1997)
12. Fekete, S.P., Woeginger, G.J.: Angle-restricted tours in the plane. Comput. Geom. Theory Appl. 8(4), 195-218 (1997)
13. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-Completeness. Freeman, New York (1979)
14. Goodman, J.E.: On the largest convex polygon contained in a non-convex $n$-gon or how to peel a potato. Geom. Dedicata 11, 99-106 (1981)
15. Hall-Holt, O., Katz, M.J., Kumar, P., Mitchell, J.S.B., Sityon, A.: Finding large sticks and potatoes in polygons. In: SODA '06: Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 474-483 (2006)
16. Niblack, W., Yin, J.: A pseudo-distance measure for 2D shapes based on turning angle. In: Proceedings of the International Conference on Image Processing (ICIP), pp. 352-355 (1995)
17. Papaioannou, D.C.G., Gaitatzes, A.: Efficient occlusion culling using solid occluders. In: Proceedings of the 14th International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision (2006)
18. Reinbacher, I., Benkert, M., van Kreveld, M., Mitchell, J.S.B., Snoeyink, J., Wolff, A.: Delineating boundaries for imprecise regions. Algorithmica 50(3), 386-414 (2008)
19. Sharir, M., Toledo, S.: Extremal polygon containment problems. Comput. Geom. Theory Appl. 4, 99-118 (1994)
20. Woo, T.C.: The convex skull problem. Technical report TR 86-31, Department of Industrial and Operations Engineering, University of Michigan (1986)
21. Wood, D., Yap, C.K.: The orthogonal convex skull problem. Discrete Comput. Geom. 3(4), 349-365 (1988)

[^0]:    Work by B.A. has been supported in part by a grant from the U.S.-Israeli Binational Science Foundation and by NSA MSP Grant H98230-06-1-0016. Partially funded by the Netherlands Organization for Scientific Research (NWO) under FOCUS/BRICKS grant number 642.065.503, and under the project GOGO.
    A preliminary version of this paper was presented at the 19th Canadian Conference on Computational Geometry (CCCG 2007), under the title "Largest Subsets of Triangles in a Triangulation".
    B. Aronov

    Department of Computer Science and Engineering, Polytechnic Institute of NYU, Brooklyn, New York, USA
    e-mail: aronov@poly.edu
    M. van Kreveld • M. Löffler • R.I. Silveira ( $\boxtimes$ )

    Department of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands.
    e-mail: rodrigo@cs.uu.nl
    M. van Kreveld
    e-mail: marc@cs.uu.nl
    M. Löffler
    e-mail: loffler@cs.uu.nl

