

VISIBILITY MAPS OF REALISTIC TERRAINS HAVE LINEAR SMOOTHED COMPLEXITY

Mark de Berg^{*}, Herman Haverkort^{*}, Constantinos P. Tsirigiannis^{*}

ABSTRACT. We study the complexity of the visibility map of terrains whose triangles are fat, not too steep and have roughly the same size. It is known that the complexity of the visibility map of such a terrain with n triangles is $\Theta(n^2)$ in the worst case. We prove that if the elevations of the vertices of the terrain are subject to uniform noise which is proportional to the edge lengths, then the worst-case expected (smoothed) complexity is only $\Theta(n)$. We also prove non-trivial bounds for the smoothed complexity of instances where some triangles do not satisfy the above properties. Our results provide an explanation why visibility maps of superlinear complexity are unlikely to be encountered in practice.

1 Introduction

Motivation. A (triangulated) *terrain* is a polyhedral surface obtained by assigning elevations to the vertices of a triangulation of a planar point set. In geographic information science such terrain models are known as triangulated irregular networks, or TINs for short. Terrains can be used to model mountainous regions, but in fact they can serve as (piecewise-linear approximations of) models of any scalar function defined over a planar region.

Often it is desirable to compute which parts of a terrain \mathcal{T} are visible from a given viewing point p_{view} . More precisely, for each triangle of \mathcal{T} , one wants to know exactly which parts are visible and which parts are invisible from p_{view} . The projections of the visible triangle parts onto a viewing plane form the so-called *visibility map* of \mathcal{T} with respect to p_{view} . Visibility maps are useful for visualisation purposes; for example, they can be used for hidden-surface removal or shadow generation. There are several algorithms for computing visibility maps of terrains, the most efficient of which runs in time $O((n\alpha(n) + k) \log n)$ [7] where $\alpha(\cdot)$ is the inverse Ackermann function. Here n is the number of triangles in \mathcal{T} and k is the output size. In other words, k is the *complexity* of the visibility map, which can be defined as the number of vertices¹ of the map. Each vertex of the map either corresponds to a triangle vertex, or to two edges whose projections onto the viewing plane intersect. In the worst case, $\Theta(n^2)$ pairs of edges have intersecting projections and all of these intersections are visible, so that the visibility map has complexity $\Theta(n^2)$. In most applications a quadratic complexity would make an explicit computation of the visibility

^{*}TU Eindhoven, PO Box 513, 5600 MB, Eindhoven, the Netherlands.
mberg@win.tue.nl, cs.herman@haverkort.nl, ctsirogi@win.tue.nl.

MdB and CPT were supported by the Netherlands' Organisation for Scientific Research (NWO) under project no. 639.023.301.

¹Formally, the complexity would be defined as the total number of vertices, edges, and faces of the map. In our setting this is always linear in the number of vertices, so we restrict ourselves to this quantity.




Figure 1: Two views of the same terrain defined by a regular grid. The second view gives a visibility map of complexity $\Theta(n\sqrt{n})$. Note that the terrain can be flattened further without changing the view combinatorially.

map infeasible. Fortunately such high complexity is seldom encountered. In fact, in practice it seems that the complexity of visibility maps is close to linear. This is the goal of our paper: to understand why visibility maps of terrains in practice often have low complexity.

Realistic input models. One possible approach to explain the low complexity is using a so-called *realistic input model* [3]. Here one assumes that the input has certain properties that are hopefully satisfied by inputs encountered in practice, and that rule out contrived worst-case inputs. This approach works well for many problems, and Moet et al. [10] have applied it to visibility maps of terrains. In particular, Moet et al. make the following three assumptions on the terrain: the triangles of the planar triangulation defining the terrain are fat (as defined below), the edges of these triangles differ in length by not more than a constant factor, and the domain of the triangulation is a rectangle of constant aspect ratio. Unfortunately, the assumptions do not explain why visibility maps of terrains would have near-linear complexity in practice: Moet et al. showed that the worst-case complexity of the visibility map of a terrain that satisfies their assumptions is $\Theta(n\sqrt{n})$. In fact, one can even assign elevations to the vertices of a triangulated grid in such a way that the triangles do not become steep while the visibility map has complexity $\Theta(n\sqrt{n})$ for certain viewing directions—see Fig. 1. Thus, to explain the linear behavior, an alternative approach is needed.

Smoothed analysis. The idea behind the alternative approach is to study how sensitive worst-case inputs are to small perturbations. If a small random perturbation of the input is likely to turn any input (whether worst-case or good-case) into a good-case input, then one may argue that worst-case inputs are unlikely to be found in practice—especially if the input is subject to small measurement or rounding errors. *Smoothed analysis* formalizes this idea.

Let $\mathcal{I}(n)$ be the set of all possible input instances (in our case: terrains) of size n .

For an input $I \in \mathcal{I}(n)$, let $C(I)$ denote the quantity we want to analyse (the complexity of the visibility map). Furthermore, for any input $I \in \mathcal{I}(n)$ we define a neighbourhood $\mathcal{N}(I) \subset \mathcal{I}(n)$ of input instances, and we define a probability distribution over $\mathcal{N}(I)$ that indicates for every $I' \in \mathcal{N}(I)$ the probability that perturbing the input I will result in the input I' . Now the *smoothed complexity* of an instance I is defined as

$$C_{smooth}(I) = \mathbb{E}_{I' \in \mathcal{N}(I)}[C(I')],$$

where the expectation is according to the given probability distribution on $\mathcal{N}(I)$. The worst-case smoothed complexity —this is what we are interested in— is then defined as

$$C_{smooth}(n) = \sup_{I \in \mathcal{I}(n)} C_{smooth}(I).$$

When $\mathcal{N}(I)$ is defined to be the full set of possible inputs and exactly the same probability distribution is used for each I , then the above complexity measure is just the average-case complexity under the given distribution. However, it is often unclear what a reasonable probability distribution is. Moreover, it does not say anything about whether cases that are significantly worse than average may be expected to occur in practice. When $\mathcal{N}(I)$ is narrowly defined as $\{I\}$, then the above complexity measure is just the (possibly unrealistic) worst-case complexity. By making a good choice for $\mathcal{N}(I)$ between these extremes, one may get a more realistic estimate of the output complexity for the problem being studied.

Smoothed analysis was introduced by Spielman and Teng [12]. So far there have only been a few applications in computational geometry (see e.g. [4, 5, 6]), none of which deals with terrains. An up-to-date collection of the published works related to smoothed analysis is maintained by Spielman [11].

Our results. We study the smoothed complexity of visibility maps of terrains under the following model:

- To the elevation of each vertex we add a noise value that follows a uniform distribution in an interval $[-c, c]$, where $c = c' \cdot \eta$ with η being the minimum edge length of the triangulation underlying the terrain and c' a positive parameter which we consider to be constant.²

Our noise model defines for each input terrain \mathcal{T} a neighbourhood $\mathcal{N}(\mathcal{T})$ consisting of those terrain instances that can be obtained by changing the elevation of each vertex by at most c , and a probability distribution on $\mathcal{N}(\mathcal{T})$. It is easy to see that this model alone is not sufficient to explain the linear complexity of the visibility map. Indeed, by applying a small perturbation one does not get rid of peaks that are unrealistically skinny and high, and so the smoothed visibility-map complexity of arbitrary terrains is still quadratic. Hence, we combine the power of smoothed analysis with the ideas of realistic input models. In particular, we define the following parameters of terrains:

²Terrain data is usually acquired using techniques such as LIDAR. Typically, the vertical accuracy of the measurements is in the range 3–15cm while the horizontal resolution is in the range 10–100cm [8].

- *Fatness*: the smallest angle in the triangles of the underlying triangulation (or in other words, the smallest angle of any triangle’s projection onto the horizontal plane);
- *Steepness*: the largest dihedral angle between any triangle and the horizontal plane;
- *Scale factor*: the length of the longest edge divided by the length of the shortest edge of the triangulation.

We assume that the fatness ϕ , steepness θ , and scale factor σ of the unperturbed terrain are constants that are independent of the number of triangles n , with $\phi > 0$, and $\theta < \pi/2$, and $\sigma \geq 1$. These assumptions are also used in other papers [1, 10], although the steepness assumption is not needed for the specific result on visibility maps by Moet et al. In her thesis, Moet [9] experimentally investigates terrain models of various mountainous regions in the US. She concludes that, at least to a large extent, they satisfy our assumptions. In itself, these assumptions do not lead to the desired result: there are terrains satisfying these assumptions with quadratic-complexity visibility maps. For example, we can take a slice of the construction of Fig. 1, with the aspect ratio of the domain being $\Theta(n)$. Our main result is that the smoothed complexity of any visibility map of a terrain satisfying the abovementioned assumptions is only $\Theta(n)$. This result can be generalized to certain non-uniform noise distributions. We also prove $O(nk)$ smoothed complexity for the visibility map of terrains that contain k triangles that do not fulfil our assumptions; the bound becomes $O(k^2 + n)$ when additionally the xy -domain of the terrain is a square. To avoid technical details regarding what happens if one looks at the boundary of the terrain from the side, we focus on the case of perspective views with the view point being located above the terrain.

2 Visibility maps resulting from perspective projection

Let \mathcal{T} be a terrain with n triangles, and let E be the set of edges of \mathcal{T} . Let the vertices be specified by three coordinates x , y and z , where the z -axis is the vertical axis on which the elevation is specified; the x - and y -axis are orthogonal to the z -axis and to each other. Let $\overline{\mathcal{T}}$ denote the triangulation in the xy -plane defining \mathcal{T} . Without loss of generality we assume that the minimum edge length in $\overline{\mathcal{T}}$ is 1. Hence, the maximum edge length equals σ , the scale factor of the terrain. We assume that $\overline{\mathcal{T}}$ is a ϕ -fat triangulation—that is, a triangulation in which all angles are at least ϕ , for some fixed constant $\phi > 0$ —and that the steepness of \mathcal{T} is bounded by θ . We study the smoothed complexity of the visibility map of \mathcal{T} for perspective views, that is, the map as it appears in the projection on a viewing plane h_{view} for a given viewing point p_{view} . We assume that p_{view} is located above the terrain.

Notation, terminology and basic properties. We denote the projection of an object o onto h_{view} by $\text{pr}(o)$. For an edge $e \in E$ we use $h_{\text{align}}(e)$ to denote the plane containing e and p_{view} ; thus $h_{\text{align}}(e) \cap h_{\text{view}}$ contains $\text{pr}(e)$. If e is collinear with p_{view} , there are many such planes: in that case we define $h_{\text{align}}(e)$ as the vertical plane containing p_{view} and e . The *steepness* $\theta(t)$ of a triangle t is defined as the dihedral angle of the plane containing t with the xy -plane, and the steepness $\theta(s)$ of a segment s is defined as the smallest acute

angle of the line containing s with the xy -plane. Observe that the steepness of a triangle equals the maximum steepness of any segment contained in it. Recall that θ denotes the maximum steepness of any triangle in \mathcal{T} . The following lemma shows that the steepness of terrains that satisfy our assumptions does not change much if the vertex elevations are subject to a small perturbation.

Lemma 2.1 *Let \mathcal{T} be a terrain constructed from a ϕ -fat triangulation $\bar{\mathcal{T}}$, and let θ be the maximum steepness of any triangle in \mathcal{T} , where $\theta < \pi/2$ is a constant. Then, after raising or lowering each vertex independently by a distance of at most c , no triangle is steeper than $\theta_{\max} = \arctan(\tan(\theta) + \frac{2c}{\sin \phi})$.*

Proof. Consider any triangle $\triangle(u, v, w)$ of \mathcal{T} . Then there must be a vertex of this triangle, say v , and a point p on the edge opposite to v , such that the segment pv is parallel to the direction of steepest descent on the triangle after the perturbation. Let \bar{u} , \bar{v} and \bar{p} denote the projections of u , v , and p onto the xy -plane, respectively. Since $|\bar{uv}| \geq 1$ and the angle at \bar{u} is at least ϕ , we have $|\bar{pv}| \geq \sin \phi$ —see Fig. 2(a).




Figure 2: Illustrations for the proof of Lemma 2.1.

Denote the elevations of v and p before the perturbation by $z(v)$ and $z(p)$. Then, since before the perturbation the steepness of the triangle is at most θ , we have

$$\frac{|z(v) - z(p)|}{|\bar{pv}|} \leq \tan \theta$$

In the worst case, the difference in elevation between v and p can increase by at most $2c$ due to the perturbation. Hence, as illustrated in Fig. 2(b), the steepness after the perturbation is at most

$$\arctan \left(\frac{|z(v) - z(p)| + 2c}{|\bar{pv}|} \right) \leq \arctan \left(\tan \theta + \frac{2c}{\sin \phi} \right).$$

□

Note that because ϕ , θ and c are constants with $\phi > 0$ and $\theta < \pi/2$, we know that θ_{\max} is a constant strictly smaller than $\pi/2$. We may assume that $\theta_{\max} \geq \pi/4$; otherwise we simply replace θ_{\max} by $\pi/4$ and the bounds proven in this paper will still hold.

The *perceived steepness* $\theta_{\text{view}}(e)$ of an edge e is the steepness of $\text{pr}(e)$ in the plane h_{view} . In other words, $\theta_{\text{view}}(e)$ is the smallest angle between the line containing $\text{pr}(e)$ and a

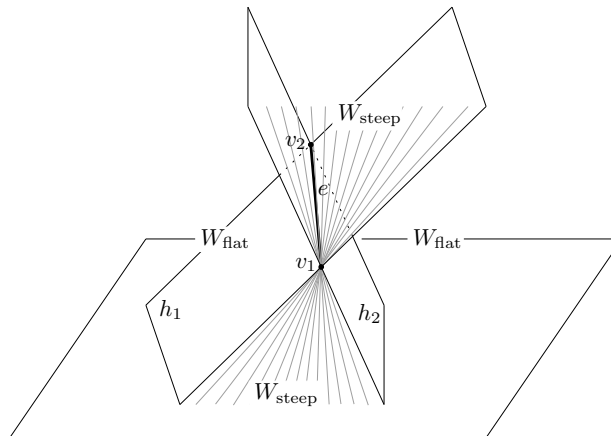


Figure 3: Partitioning the space around an edge in two double wedges: one containing all steep planes that contain the edge, one containing all other planes that contain the edge.

horizontal line on h_{view} . Even though $\theta(e) \leq \theta_{\text{max}}$ by Lemma 2.1, $\theta_{\text{view}}(e)$ may be greater than θ_{max} and in fact be equal to $\pi/2$. Indeed, even an edge that is horizontal can appear vertical when projected onto h_{view} . We say that an edge e *appears steep* when $\theta_{\text{view}}(e)$ is well-defined— $\text{pr}(e)$ is not a single point—and $\theta_{\text{view}}(e) > \theta_{\text{max}}$, otherwise e *appears flat*.

We say that an edge e lies *in front of* an edge e' , if there is a ray from p_{view} that hits e before hitting e' . A *silhouette edge* is an edge e such that the two triangles of \mathcal{T} that share e are on the same side of $h_{\text{align}}(e)$. Note that this is equivalent to saying that one of the incident triangles of a silhouette edge is front-facing while the other is back-facing. We say that two edges e and e' *create a visible intersection* if $\text{pr}(e) \cap \text{pr}(e')$ constitutes a vertex of the visibility map. For this to be possible, the edge hit first by a ray from p_{view} —say this edge is e —must be a silhouette edge. Otherwise it would have to be an edge on the boundary of the terrain or a non-boundary, non-silhouette edge, but both cases would lead to a contradiction: in the first case e could only be the *last* edge hit by any directed line through p_{view} because there is nothing beyond e —this is true because the domain of \mathcal{T} is a triangulation and, hence, convex—while in the second case the two triangles incident to e would hide e' from view locally around the point of e' projecting onto $\text{pr}(e) \cap \text{pr}(e')$.

Above we observed that even though the edges of the terrain are not steeper than θ_{max} , they can still appear steep on h_{view} . The next lemma, which is the key to bounding the number of visible intersections, states that this cannot happen for silhouette edges.

Lemma 2.2 *Let \mathcal{T} be a terrain whose triangles have steepness at most θ_{max} and let p_{view} be a fixed viewing point. Then the perceived steepness (on any vertical viewing plane) of any silhouette edge of \mathcal{T} is at most θ_{max} .*

Proof. Consider an edge $e = v_1v_2$. Let h_1 and h_2 denote the two planes containing e that have steepness exactly θ_{max} . These two planes partition the space into two double wedges: $W_{\text{steep}} := (h_1^+ \cap h_2^+) \cup (h_1^- \cap h_2^-)$ and $W_{\text{flat}} := (h_1^+ \cap h_2^-) \cup (h_1^- \cap h_2^+)$, where h_i^+ and h_i^- denote the half-spaces above and below h_i , respectively. Note that W_{steep} is the union of all planes

containing e that are steeper than θ_{\max} , while W_{flat} is the union of all planes containing e that are less steep than θ_{\max} —see Fig. 3.

Let $\Delta(u, v_1, v_2)$ and $\Delta(v_1, v_2, w)$ be the triangles sharing the edge e . Since no terrain triangle is steeper than θ_{\max} , the vertices u and w must lie in W_{flat} . Moreover, they must lie in different parts (that is, wedges) of W_{flat} , since otherwise there would be a vertical line intersecting the interiors of both $\Delta(u, v_1, v_2)$ and $\Delta(v_1, v_2, w)$, which cannot happen because \mathcal{T} is a terrain.

Now consider $h_{\text{align}}(e)$, the plane that contains e and p_{view} and whose intersection with h_{view} contains $\text{pr}(e)$. Suppose e appears steep, that is, $\theta_{\text{view}}(e) > \theta_{\max}$. Because $h_{\text{align}}(e)$ contains $\text{pr}(e)$ it is at least as steep as $\text{pr}(e)$, so we have $\theta(h_{\text{align}}(e)) > \theta_{\max}$. Since $h_{\text{align}}(e)$ contains e , this means $h_{\text{align}}(e)$ is contained in W_{steep} . Since u and w lie in different parts of W_{flat} , this implies that $\Delta(u, v_1, v_2)$ and $\Delta(v_1, v_2, w)$ lie on different sides of $h_{\text{align}}(e)$. Thus e is not a silhouette edge if it appears steep, which proves the lemma. \square

Counting intersections. Since the number of terrain vertices is $O(n)$, we only need to worry about bounding the number of visible intersections created by pairs of terrain edges. We will charge each visible intersection to the edge furthest from the viewer (of the two edges creating it). Thus we need to count, for each edge e in the perturbed terrain, the number of visible intersections it creates with edges in front of it. We denote this number by $K(e)$.

Consider the situation in which we have already perturbed the edges in front of e , and we wish to analyse the effect of perturbing e . Consider the triangle whose vertices are p_{view} and the endpoints of e , and define $\overline{\Delta}(p_{\text{view}}, e)$ to be the projection of this triangle onto the xy -plane. Let $E_{\text{fr}}(e)$ be the set of silhouette edges whose projection onto the xy -plane intersects $\overline{\Delta}(p_{\text{view}}, e)$, clipped to the part whose projection is contained in $\overline{\Delta}(p_{\text{view}}, e)$. We exclude the edges sharing a vertex with e ; such edges cannot create a visible intersection with e . Then the visible intersections charged to e are intersections of $\text{pr}(e)$ with the upper envelope of $\{\text{pr}(s) : s \in E_{\text{fr}}(e)\}$ —see Fig. 4(a). We denote this upper envelope by $H(e)$. Now consider a fragment $f \in E_{\text{fr}}(e)$ that appears on this upper envelope. We wish to bound the probability that after perturbation of e , the edge e creates a visible intersection with f .

Observe that the combinatorial structure of the visibility map on a viewing plane h_{view} does not depend on the location and orientation of the viewing plane, provided h_{view} does not contain p_{view} (in which case one would not see anything). We can therefore assume without loss of generality that h_{view} is a vertical plane that contains e . For an object o , define $\text{span}_e(o)$ to be the projection of $\text{pr}(o)$ onto a horizontal line on h_{view} , and define $\text{width}_e(o)$ to be the length of $\text{span}_e(o)$. We have the following lemma.

Lemma 2.3 *Let e be an edge of \mathcal{T} and let f be a fragment of a silhouette edge e' not incident to and in front of e such that $\text{span}_e(f) \subset \text{span}_e(e)$ and $\text{width}_e(f) \leq \text{width}_e(e)/3$. Now suppose we independently perturb the elevations of the vertices of e , where the perturbations are chosen uniformly at random from the range $[-c, c]$. Then*

$$\Pr[e \text{ creates a visible intersection with } f] \leq \frac{3 \text{width}_e(f) \cdot \tan \theta_{\max}}{c}.$$

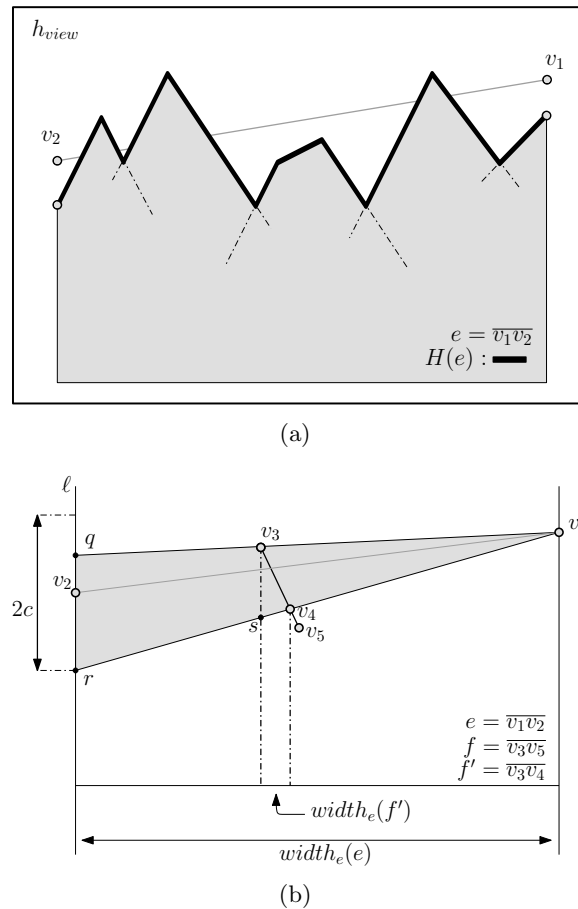


Figure 4: (a) The horizon in front of e as it appears on the visibility map. (b) The shaded cone indicates all the possible positions of e that may induce a visible intersection between e and f for a fixed position of v_1 . The point s is the point on the image of v_1r that has the same x -coordinate as v_3 .

Proof. Assume without loss of generality that the projection of e on the xy -plane is parallel to the x -axis. Let v_1 and v_2 be the vertices of e . Without loss of generality assume v_2 is the vertex closest to f in the projection onto the x -axis, with ties broken arbitrarily. Since $\text{width}_e(f) \leq \text{width}_e(e)/3$, the distance from v_1 's projection to $\text{span}_e(f)$ is at least $\text{width}_e(e)/3$. We will now show that for any elevation of v_1 after the perturbation, the probability that e intersects f when v_2 is perturbed, is at most $3 \text{width}_e(f) \cdot \tan(\theta_{\max})/c$.

Let ℓ be the vertical line on h_{view} through $\text{pr}(v_2)$. Assume v_1 has already been perturbed, so that v_1 's position is now fixed. Consider the set of all possible positions of the projection of the perturbed vertex v_2 that induce an intersection between f and e . This set is a segment on ℓ ; we denote its upper endpoint by q and its lower endpoint by r —see Fig. 4(b). The probability that, after perturbing v_2 , the edge e creates a visible intersection with f is bounded by $|qr|/(2c)$.

We will now compute an upper bound on $|qr|$. Notice that the triangle $\Delta(v_1, q, r)$

does not necessarily contain f completely: there may be parts of f where e cannot create an intersection (for the given, fixed position of v_1), because v_2 could not be raised or lowered far enough without going beyond the given bounds on the perturbation. Let f' be the part of f inside $\Delta(v_1, q, r)$ and let v_3 and v_4 be the endpoints of f' , with v_3 being the endpoint closest to ℓ . Let s be the point such that $\Delta(v_1, v_3, s)$ and $\Delta(v_1, q, r)$ are similar triangles. The line through v_1 and r has steepness at most θ_{\max} by Lemma 2.1 and the fact that h_{view} contains e ; fragment f' has steepness at most θ_{\max} by Lemma 2.2. This implies that

$$|v_3s| \leq 2 \text{width}_e(f') \cdot \tan \theta_{\max}.$$

Moreover, we have $|v_1v_3| \geq |v_1q|/3$. Now

$$|qr| = |v_3s| \cdot \frac{|v_1q|}{|v_1v_3|} \leq 6 \text{width}_e(f') \cdot \tan \theta_{\max} \leq 6 \text{width}_e(f) \cdot \tan \theta_{\max}.$$

Hence, the probability of intersection between e and f for any fixed position of v_1 is at most

$$\frac{|qr|}{2c} \leq \frac{3 \text{width}_e(f) \cdot \tan \theta_{\max}}{c}.$$

This bound is independent of the position of v_1 . Therefore the probability that, when e is perturbed, e creates a visible intersection with f is at most $3 \text{width}_e(f) \cdot \tan(\theta_{\max})/c$. \square

Lemma 2.4 *Let $K(e)$ be the number of visible intersections of an edge e of \mathcal{T} with edges in front of it. Then*

$$\mathbb{E}[K(e)] \leq \frac{3 \text{width}_e(e) \cdot \tan \theta_{\max}}{c} + 2.$$

Proof. We distinguish two cases: either $h_{\text{align}}(e)$ is vertical, or not. If $h_{\text{align}}(e)$ is vertical, then $\text{pr}(e)$ is vertical too. This implies $K(e) \leq 1$: the part of e above the intersection point would be visible and the part below it would be hidden from view.

It remains to discuss the case in which $h_{\text{align}}(e)$ is not vertical. We now need to bound the expected number of visible intersections created by e with any edge e' in front of it. Consider the situation where we already perturbed the edges in front of e , but not yet the edge e itself. Recall that e' must be a silhouette edge. Hence, we can restrict our attention to $E_{\text{fr}}(e)$, the set of segments which are parts of silhouette edges lying in front of e in the projection onto the xy -plane, excluding the edges sharing a vertex with e . We defined $H(e)$ to be the upper envelope of the projections of those segments onto h_{view} —see Fig. 4(a). Then we can bound $\mathbb{E}[K(e)]$ by analysing the number of intersections of $\text{pr}(e)$ with $H(e)$.

For a fragment f on the upper envelope $H(e)$, we define an indicator random variable X_f :

$$X_f = \begin{cases} 1 & e \text{ and } f \text{ intersect} \\ 0 & \text{otherwise} \end{cases}$$

We have $K(e) = \sum_f X_f$. Observe that there can be at most two such fragments with $\text{width}_e(f) > \text{width}_e(e)/3$. Now consider the other fragments. By Lemma 2.3, we have

$$\Pr[e \text{ creates a visible intersection with } f] \leq \frac{3 \text{width}_e(f) \tan \theta_{\max}}{c}.$$

Adding the at most two fragments with $\text{width}_e(f) > \text{width}_e(e)/3$, and summing over all fragments with $\text{width}_e(f) \leq \text{width}_e(e)/3$, we get

$$\mathbb{E}[K(e)] = \mathbb{E}\left[\sum_f X_f\right] = \sum_f \mathbb{E}[X_f] \leq 2 + \sum_f \frac{3 \text{width}_e(f) \tan \theta_{\max}}{c}.$$

Clearly we have $\sum_f \text{width}_e(f) \leq \text{width}_e(e)$, which finishes the proof. \square

Using that $\text{width}_e(e) \leq \sigma$ and $\tan \theta_{\max} \leq \tan(\theta) + \frac{2c}{\sin \phi}$ (by Lemma 2.1) we obtain our final result:

Theorem 2.5 *Let \mathcal{T} be a terrain of n triangles with fatness ϕ , steepness θ , and scale factor σ . Suppose we add noise to each vertex's elevation independently, where the noise value is taken uniformly at random from the interval $[-c, c]$, and c is a fixed constant fraction of the minimum edge length of the triangulation underlying the terrain model. Then a visibility map of \mathcal{T} under perspective projection has smoothed complexity*

$$O\left(\left(\frac{\tan \theta}{c} + \frac{1}{\sin \phi}\right) \sigma n\right).$$

Remark 2.6 The result easily generalises to other noise distributions over the interval $[-c, c]$, provided that the density function of the distribution is upper-bounded. More precisely, if the density is upper-bounded by τ , then for any interval $[a, b] \subset [-c, c]$ the probability that the noise is in $[a, b]$ is at most $\tau \cdot |b - a|$ instead of $\frac{|b-a|}{2c}$. The bound on the probability given in Lemma 2.3 and, hence, also the final bound obtained in Theorem 2.5 then simply have to be multiplied by $2c\tau$.

Remark 2.7 One can also consider the setting where not only the elevations of the sample points are subject to noise, but also their xy -positions. In a practical setting, however, the 2D locations of the sample points determine the triangulation (because one uses the Delaunay triangulation, for instance). This means that a very different analysis will be necessary. If one ignores this aspect and simply works with the given triangulation, then our results still hold provided that (i) the noise is small enough that the fatness, steepness, and scale factor of the terrain are still bounded, and (ii) the perturbation in the xy -plane is independent of the perturbation of the elevation.

3 Terrains that almost satisfy the assumptions

In previous works on realistic input models [3, 9] it has been observed that even in well-formed scenes, where the majority of the objects fulfills the properties of some model, there

may exist a few objects that do not conform with the model. This leads us to consider the case where there exists a set of $k \leq n$ terrain triangles that do not follow the model assumptions. We call the triangles that have fatness ϕ , slope at most θ and edge length $\in [1, \sigma]$ *good* and the rest *bad*. Furthermore, we call a terrain edge *bad* if it is incident to a bad triangle; otherwise the edge is called *good*. We now analyse the smoothed complexity of the visibility map of a terrain with k bad triangles.

Theorem 3.1 *Let \mathcal{T} be a terrain containing $n - k$ good triangles with fatness ϕ , steepness θ , and edge lengths $[1, \sigma]$ and k bad triangles. Suppose we add noise to each vertex's elevation independently, where the noise value is taken uniformly at random from the interval $[-c, c]$, and c is a constant. Then a visibility map of \mathcal{T} under perspective projection has smoothed complexity*

$$O\left(nk + \left(\frac{\tan \theta}{c} + \frac{1}{\sin \phi}\right) \sigma n\right).$$

Proof. For every terrain edge e we will count how many visible intersections $\text{pr}(e)$ can create with edges in front of it. Recall that $\text{pr}(e)$ can only create a visible intersection with the images of the segments in $s \in E_{\text{fr}}(e)$ that is, the silhouette edge parts for which it holds that $\text{span}_e(s) \subset \text{span}_e(e)$. More precisely $\text{pr}(e)$ can only create a visible intersection with $H(e)$, the upper envelope of these segments. Because \mathcal{T} contains bad triangles it is possible in this setting that $H(e)$ contains steep segments, which belong to the images of bad edges.

We can partition the intersections that appear in the visibility map into two categories, which we analyze separately.

- *Category (i): at least one of the involved edges is bad.*
Since there are at most $3k$ bad edges, and any edge can generate only $O(n)$ intersections, the number of visible intersections of the first category is $O(nk)$.

- *Category (ii): both involved edges are good.*
In Lemma 2.4, in a setting where no bad triangles exist in \mathcal{T} , we proved that a good edge e can be charged with a smoothed number of at most $\frac{3 \cdot \text{width}_e(e) \cdot \tan \theta_{\max}}{c} + 2$ intersections with the segments of $H(e)$. In fact, this can be rephrased as an upper bound for the smoothed number of intersections between e and the flat segments of $H(e)$; the proof itself is independent of the existence of steep segments in $H(e)$ and thus this bound directly applies for the cases that \mathcal{T} contains bad triangles. Since e is a good edge, $\text{width}_e(e) \leq \sigma$ and e can be charged with a smoothed number of at most $\frac{3\sigma \cdot \tan \theta_{\max}}{c} + 2$ intersections with other good edges. Hence the total number of intersections between good edges is $O\left(\left(\frac{\tan \theta}{c} + \frac{1}{\sin \phi}\right) \sigma n\right)$.

Summing up the intersections of categories (i) and (ii) the theorem follows. □

Fig. 5 shows a view of a construction where we place $\Theta(n)$ realistic triangles behind k skinny peaks in a way that the noise cannot cause their incident edges to hide behind each other. The smoothed complexity of this view is $\Omega(nk)$, thus showing that the bound from Theorem 3.1 is tight in the worst case.




Figure 5: A terrain with a visibility map of $\Theta(nk)$ smoothed complexity. In this view, each of the $\Theta(k)$ thin and long spikes in the front appears to intersect with each of the $\Theta(n)$ good triangles at the back.

A more restricted model. The result of Theorem 3.1 is a bit disappointing: even with, say, $\Theta(\log n)$ bad edges, the smoothed complexity can be superlinear. Next we show that we can get much stronger bounds if we make one more (very reasonable) assumption: we assume that the xy -domain of the terrain is a square of size $\Theta(\sqrt{n}) \times \Theta(\sqrt{n})$. Consequently, our model becomes a special case of the one presented by Moet et al. [10]. Below, we mention an interesting property that they prove for a terrain \mathcal{T} that conforms with their model.

Property 3.2 *Let s be a straight line segment that intersects the projection of \mathcal{T} . Then s intersects $O(\sqrt{n})$ triangles of \mathcal{T} .*

This property is based on a packing argument that provides an upper bound for the number of good triangles that can fit in a rectangle of a certain size. The existence of bad triangles therefore does not affect this argument when applied to the good triangles. Hence, the number of good triangles intersecting any straight line segment is still $O(\sqrt{n})$.

We prove a tight upper bound for the smoothed complexity of the visibility map of such a terrain that also contains at most k bad triangles.

Theorem 3.3 *Let \mathcal{T} be a terrain with an xy -domain that is a square of size $\Theta(\sqrt{n}) \times \Theta(\sqrt{n})$. Let \mathcal{T} contain $n - k$ good triangles with fatness ϕ , steepness θ , edge lengths $\in [1, \sigma]$ and k bad triangles. Suppose we add noise to each vertex's elevation independently, where the noise value is taken uniformly at random from the interval $[-c, c]$, and c is a constant. Then*

a visibility map of \mathcal{T} under perspective projection has smoothed complexity

$$O\left(k^2 + \left(\frac{\tan \theta}{c} + \frac{1}{\sin \phi}\right) (\sqrt{n}k + \sigma n)\right).$$

Proof. As in the proof of Theorem 3.1, we distinguish several categories for the intersections of the visibility map. Next we analyse each category.

- *Category (i): intersections between a bad edge e and the steep segments of $H(e)$.*
For an edge e , the steep segments of $H(e)$ are parts of the images of bad edges. In the worst case every bad edge may appear to intersect with $\Theta(k)$ other such edges, summing up to $O(k^2)$ intersections of this kind.
- *Category (ii): intersections between a bad edge e and the flat segments of $H(e)$.*
Let e be a bad edge. We consider first the case that $\theta(e) \leq \theta_{\max}$. With an analysis similar to the one of ii) in the proof of Theorem 3.1, e can be charged with a smoothed number of at most $\frac{3 \text{width}_e(e) \cdot \tan \theta_{\max}}{c}$ intersections. Since $\text{width}_e(e) = O(\sqrt{n})$, this number is $O\left(\frac{\sqrt{n} \tan \theta_{\max}}{c}\right)$. We examine now the case that $\theta(e) > \theta_{\max}$. There can be at most $O(k)$ steep edges that can participate in $H(e)$, that means that e can disappear behind $H(e)$ and reappear at most $O(k)$ times. Notice that e can disappear after a steep segment and then reappear after a flat segment or the other way round. Thus e can in fact participate in $O(k)$ intersections with flat segments. Hence we can have in total a smoothed number of $O\left(k^2 + \left(\frac{\tan \theta}{c} + \frac{1}{\sin \phi}\right) \sqrt{n}k\right)$ such intersections.
- *Category (iii): intersections between a good edge e and the flat segments of $H(e)$.*
This category is identical to ii) in the proof of Theorem 3.1. Accordingly, we can have at most $O\left(\left(\frac{\tan \theta}{c} + \frac{1}{\sin \phi}\right) \sigma n\right)$ intersections of this kind.
- *Category (iv): intersections between a good edge e and the steep segments of $H(e)$.*
We now consider the intersections that good edges induce with bad edges that appear in front of them. Let E_{bad} be the set of bad edges in \mathcal{T} and V_{bad} be the set of the vertices incident to bad edges. For a good terrain edge e and a viewing plane h_{view} let $H_{\text{bad}}(e)$ be the upper envelope of the images of the bad edges that appear in front of e . We can divide $H_{\text{bad}}(e)$ into maximal polygonal chains whose endpoints are images of actual vertices in V_{bad} . Each such chain is convex and obviously $\text{pr}(e)$ can create at most two visible intersections with this chain.

For some vertex $v \in V_{\text{bad}}$, consider the ray that starts at p_{view} and passes through v . Consider on this ray the segment that starts at v and ends at the intersection of this ray with the xy -projection of the boundary of \mathcal{T} . We call this segment the *lifespan* of v —see Fig 6(b). If a good edge e intersects the lifespans of m vertices in V_{bad} then in a possible view of the terrain $\text{pr}(e)$ may appear to cross all the $m - 1$ convex chains between the images of those vertices plus at most two other chains. As remarked before, $\text{pr}(e)$ can create at most two visible intersections with each chain. So, we have at most $2(m + 1)$ visible intersections. We charge two of them to e and two to each vertex in V_{bad} whose lifespan is crossed. Hence, each vertex in $v \in V_{\text{bad}}$ can be

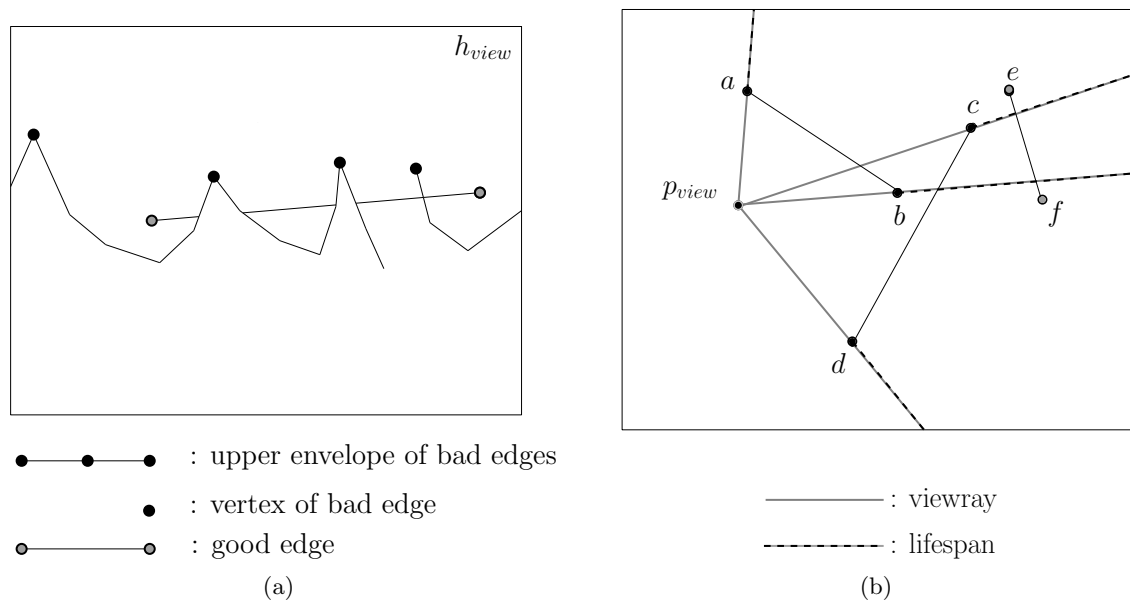


Figure 6: (a) A possible view of the upper envelope of the bad edges and a good edge at the back. (b) The xy -projection of a terrain that contains a good edge \overline{ef} and two bad edges $\overline{ab}, \overline{cd}$. Illustrated are also the lifespans for the vertices of the bad edges.

charged with a constant number of intersections from each good edge that intersects its lifespan.

According to Property 3.2 there can be $O(\sqrt{n})$ edges that cross the lifespan of v and thus there can be $O(\sqrt{n})$ intersections charged to v . That means that there are $O(k\sqrt{n})$ intersections in total for all the vertices in V_{bad} .

Hence, the total number of visible intersections that can appear on the visibility map is $O\left(k^2 + \left(\frac{\tan \theta}{c} + \frac{1}{\sin \phi}\right)(\sqrt{nk} + \sigma n)\right)$.

□

4 Concluding remarks

We proved that the smoothed complexity of the visibility map of not-too-steep terrains with fat triangles of similar size is $O(n)$. This is a possible explanation why in practice terrains with visibility maps of super-linear complexity are unlikely to occur. We also examined the smoothed complexity for terrains that contain a few triangles that do not satisfy the assumptions. This is the first time that realistic input models have been combined with smoothed analysis. We believe this is a promising approach, which could also shed light on the complexity of certain other structures on real-world terrains. For example, the

complexity of the river network on real-world terrains seems to be linear, while the worst-case complexity of the river network on a terrain with the above-mentioned properties is still $\Theta(n^2)$ [2]. Combining these properties with a smoothed analysis may lead to better bounds.

References

- [1] B. Aronov, M. de Berg, and S. Thite. The complexity of bisectors and Voronoi diagrams on realistic terrains. In *Proc. 16th Annual European Symposium on Algorithms*, LNCS 5193, pages 100–111, 2008.
- [2] M. de Berg, O. Cheong, H. Haverkort, J.G. Lim, L. Toma. The complexity of flow on fat terrains and its I/O-efficient computation. *Computational Geometry: Theory and Applications* 43:331–356 (2010).
- [3] M. de Berg, A.F. van der Stappen, J. Vleugels, M.J. Katz. Realistic input models for geometric algorithms. *Algorithmica* 34:81–97 (2002).
- [4] S. Chauduri and V. Koltun. Smoothed analysis of probabilistic roadmaps. *Computational Geometry: Theory and Applications* 42:731–747 (2009).
- [5] V. Damerow, F. Meyer auf der Heide, H. Räcke, C. Scheideler, and C. Sohler. Smoothed motion complexity. In *Proc. 11th Annual European Symposium on Algorithms*, LNCS 2832, pages 161–171, 2003.
- [6] V. Damerow, C. Sohler. Extreme points under random noise. In *Proc. 12th Annual European Symposium of Algorithms*, LNCS 3221, pages 264–274, 2004.
- [7] M.J. Katz, M.H. Overmars, M. Sharir. Efficient hidden surface removal for objects with small union size. In *Proc. 7th ACM Symposium on Computational Geometry*, pages 31–40, 1991.
- [8] G. Mandlbürger, C. Briese. Using airborne laser scanning for improved hydraulic models. In *Proc. International Congress on Modelling and Simulation*, pages 731–738, 2006.
- [9] E. Moet. *Computation and Complexity of Visibility in Geometric Environments*. PhD Thesis, Utrecht University, 2008.
- [10] E. Moet, M. van Kreveld, A.F. van der Stappen. On realistic terrains. *Computational Geometry: Theory and Applications* 41:48–67 (2008).
- [11] D.A. Spielman. <http://www.cs.yale.edu/homes/spielman/SmoothedAnalysis/index.html> Retrieved July 21, 2009.
- [12] D.A. Spielman, S.H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time? *Journal of the ACM* 51:385–463 (2004).