TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Near-Delaunay Metrics

Citation for published version (APA):
van Beusekom, N., Buchin, K. A., Koerts, H. O., Meulemans, W., Rodatz, B., & Speckmann, B. (2021). *Near-Delaunay Metrics*. 1-11. Abstract from 33rd Canadian Conference on Computational Geometry (CCCG 2021).

Document license:
CC BY

Document status and date:
Published: 10/08/2021

Document Version:
Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

• A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
• The final author version and the galley proof are versions of the publication after peer review.
• The final published version features the final layout of the paper including the volume, issue and page numbers.

Link to publication

General rights
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

• Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
• You may not further distribute the material or use it for any profit-making activity or commercial gain
• You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy
If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Download date: 07. Nov. 2026

Near-Delaunay Metrics

Nathan van Beusekom* Kevin Buchin* Hidde Koerts*
Wouter Meulemans* Benjamin Rodatz† Bettina Speckmann*

Abstract

We study metrics that assess how close a triangulation is to being a Delaunay triangulation, for use in contexts where a good triangulation is desired but constraints (e.g., maximum degree) prevent the use of the Delaunay triangulation itself. Our near-Delaunay metrics derive from common Delaunay properties and satisfy a basic set of design criteria, such as being invariant under similarity transformations. We compare the metrics, showing that each can make different judgments as to which triangulation is closer to Delaunay. We also present a preliminary experiment, showing how optimizing for these metrics under different constraints gives similar, but not necessarily identical results, on random and constructed small point sets.

1 Introduction

Delaunay triangulations are a common construct in computational geometry: practically any class on computational geometry teaches about Delaunay triangulations of point sets and their duality to the Voronoi diagram. They are efficiently computable, and have many desirable properties, for instance, lexicographically maximizing the minimum angle of the corners of the resulting triangles [20], having bounded dilation [4, 11], minimizing the maximum circumcircle [5], maximizing the minimum enclosed circle [5, 18] of its triangles. As such, they form the basis of many algorithms that require some triangulation of a point set.

However, there are various scenarios imaginable where the Delaunay triangulation is not immediately applicable due to constraints on the desired triangulation, such as a limited vertex degree or a set of edges that needs to be included. In such cases, we may want to find a triangulation that is as close as possible to the Delaunay triangulation while adhering to the constraints. This, however, requires a way to measure how *near-Delaunay* a triangulation is. Another scenario in which such a measure would be useful is when a triangulation is already given – for example, the triangulation of a terrain. In such a setting, we could use the measure to assess the quality of the triangulation.

An example of a triangulation that aims to be as close as possible to the Delaunay triangulation given a set of edges that needs to be included is the Constrained Delaunay Triangulation (CDT) [3, 13]. It provides an alternative definition of when an edge or triangle may be part of the triangulation, such that it is "as close as possible" to being Delaunay for the given constraints. However, it does not help in assessing how close a triangulation is to being the actual Delaunay triangulation, nor does it generalize to other forms of constraints.

In this work, we consider various ways to measure how close a triangulation is to being Delaunay. The problem of studying the properties of triangulations that are close to the Delaunay triangulation was proposed at CCCG 2017 by O'Rourke [16]. In this context, two measures were already proposed [16, 14], which we discuss further in Section 2. The aim of our work, is to explore a broader range of measures together with algorithms to compute them and with an analytical and experimental comparison. We identify several criteria for near-Delaunay metrics:

C1 The Delaunay triangulation should obtain the perfect score. We do not want to distinguish how nice Delaunay triangulations are of different point sets – because this would be a factor of the point set itself. A non-Delaunay triangulation should always score less than perfect, such that any non-Delaunay triangulation is considered less Delaunay.

C2 The measure should behave continuously for slight perturbations of the point set. Triangles that "severely" violate properties of a Delaunay triangulation should score worse than those with "slight" violations.

C3 The measure should be invariant under similarity transformations (translations, rotations, scaling and reflections). Though this criteria follows from the first for the Delaunay triangulation (which is invariant under similarity transformations), it is desirable for this to also hold for non-Delaunay triangulations, such that triangulations of different point sets can still be reasonably compared.

*Department of Mathematics and Computer Science, TU Eindhoven, the Netherlands; n.a.c.v.beusekom@tue.nl, k.a.buchin@tue.nl, h.o.koerts@student.tue.nl, w.meulemans@tue.nl, b.speckmann@tue.nl.

†Department of Computer Science, University of Oxford; benjamin.rodatz@cs.ox.ac.uk.

C4 The measure should be decomposable, that is, evaluated separately on different elements of the triangulation. Though not strictly necessary, this allows various forms of aggregation (worst situation, average situation, etc.).

We observe, that the combination of criterion *C3* and *C4* suggests that the metric should also be "locally invariant". That is, even within the same triangulation, a decomposition element that is subject to the same constraints to another element up to similarity transformations should score the same. That is, the metric should not naturally award higher (or lower) scores to elements that are larger; instead small and large elements should contribute equally to the overall metric. In other words, a triangulation should not be considered near-Delaunay, simply because the deviations from the Delaunay property are only at very small elements.

We propose and compare several *near-Delaunay metrics*. These metrics (as shown in Table 1) satisfy our four criteria. They differ in the properties of the Delaunay triangulation they aim to capture, as well as how they decompose the given triangulation. For this, we use decompositions into: (1) quadrilaterals – a quadrilateral here is an edge and its two incident triangles, ignoring all other points; (2) edges – an edge is considered in context of all other points; (3) triangles in context of all other points. We show that these measures behave differently, capturing different aspects of how near-Delaunay a triangulation is. Finally, we briefly compare how our measures relate to the CDT.

Related work. One of the well-known results in computational geometry is that any triangulation can be transformed into the Delaunay triangulation using Lawson flips [12]. A natural consideration would thus be to measure the number of flips necessary to transform a triangulation into a Delaunay triangulation. Though it satisfies *C1* and *C3*, it satisfies neither *C2* – it is a discrete measure and it is not immediately based on the violations of a Delaunay property – nor *C4*. An additional complication is that computing such flip distances is generally hard [17].

Another natural consideration is to measure the number of points in any circumcircle of a triangle in the triangulation. This leads to the notion of higher-order Delaunay triangulations [7]. We do not consider them in this paper, since a small perturbation can greatly change the number of points in a circumcircle. That is, the resulting measure would not adhere to *C2*. But when this criterion is not needed, higher-order Delaunay triangulations are well suited to obtain triangulations close to the Delaunay triangulation that at the same time are optimized for another criterion. Van Kreveld et al. [22] discuss optimizing over first-order Delaunay triangulations for various criteria like minimizing the maximum degree.

Computing a triangulation that is as close as possible to the Delaunay triangulation (given certain constraints) can be seen as optimization problem. There are many papers studying optimal triangulations under various criteria. Bern et al. [2] show how to efficiently compute triangulations under criteria like maximizing the minimum height or minimizing the maximum eccentricity. Unfortunately for other criteria, computing optimal triangulations is more difficult. For instance computing the minimum-weight triangulation is NP-hard [15]. While the complexity of finding the minimum-dilation triangulation is open [8], many related problems on minimizing dilation are NP-hard [6]. Similarly, the complexity of finding the minimum-degree triangulation seems open, while it is NP-hard if as a constraint certain edges have to be included [9, 10].

The fact that for many optimization criteria efficient algorithms are not known, also limits the size of the point sets that we can include in our experiments, in which we for instance want to compute near-Delaunay triangulations with additional constraints like a degree bound or a bound on the weight of the triangulation. We here resort to enumerating triangulations, however, the total number of triangulation for a given point set is exponential [1, 19].

2 Near-Delaunay metrics

We define seven near-Delaunay metrics as shown in Table 1, which all satisfy our four criteria. Throughout, we assume that we are to measure a triangulation T on a point set P. We further assume, for sake of simplicity, that P is in general position: no three points are on a line and no four points are cocircular. We organize them below into three categories, depending on their form of decomposition (quadrilateral, edge and triangle). We always describe the measure just for a single decomposition element, silently assuming some form of aggregation such as taking their sum or extremal value.

2.1 Quadrilateral-based metrics

We start by introducing intuitive metrics that evaluate quadrilaterals. A quadrilateral consists of a non-convex-hull edge and the two triangles that are incident to it, defined by two vertices. Particularly, note that any point in P that is not one of the four defining vertices does not influence the metric on this particular quadrilateral – this contrasts the edge-based and triangle-based metrics. Thus, to compute the metrics it is sufficient to be given only the quadrilateral. Consequently, the metrics we introduce for quadrilaterals can be computed in constant time per quadrilateral and in linear time overall. Note that for each quadrilateral-based metric, lower scores mean closer to Delaunay (contrasting our other two forms of decomposition).

Table 1: Overview of metrics. Listed with each is the form of decomposition, the Delaunay property it is based on, and the running time for computing the measure for a given triangulation.

Opposing Angles	Dual Edge Ratio	Dual Area overlap	Lens	Shrunk Circle	Triangular Lens	Shrunk Circumcircle
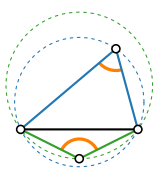	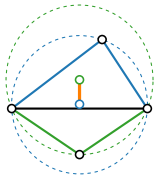	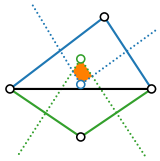	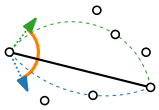	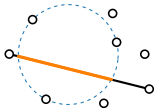	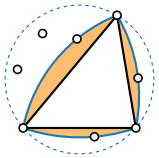	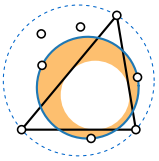
Quadrilateral	Quadrilateral	Quadrilateral	Edge	Edge	Triangle	Triangle
Max-min angle	Voronoi Dual (edges)	Voronoi Dual (faces)	Empty Circle	Empty Circle	Empty Circumcircle	Empty Circumcircle
$O(n)$	$O(n)$	$O(n)$	$O(n^2)$	$O(n^2)$	$O(n^2)$	$O(n^2)$

Throughout, we denote by (u, v) the defining edge of the quadrilateral, and with p and q the *opposing* vertices of the two incident triangles $\triangle uvp$ and $\triangle uvq$. We use c_p and c_q to denote the center of the circumcircles of these triangles. A quadrilateral is *locally Delaunay* if the edge (u, v) is part of the Delaunay triangulation of $\{u, v, p, q\}$ – in other words, if the circumcircle of the one triangle does not contain the other vertex.

Opposing Angles. At CCCG 2017, O'Rourke suggested that a triangulation T is a near-Delaunay triangulation if the opposite angles α and β of a quadrilateral sum to at most $\pi + \varepsilon$ for $\varepsilon \geq 0$ [16] (see Fig. 1). If $\varepsilon = 0$, then T is Delaunay. Hence, it is natural to consider the smallest ε for a triangulation as a metric for how close a triangulation is to a Delaunay triangulation. We can readily interpret ε on a per-quadrilateral basis for a metric in our context; O'Rourke's suggestion is simply to take their maximum as the overall metric.

The intuition behind the metric is, when the sum of two opposing angles is larger than π, the empty circumcircle property is violated. When the sum is close to π, a slight movement in the points can restore the empty circumcircle property. However as the sum grows larger, the points generally need to move further to restore the property, unless p or q is very close to u or v, in which case a small movement is sufficient. Thus, we can use the sum of two opposing angles to evaluate how far from Delaunay a quadrilateral is.

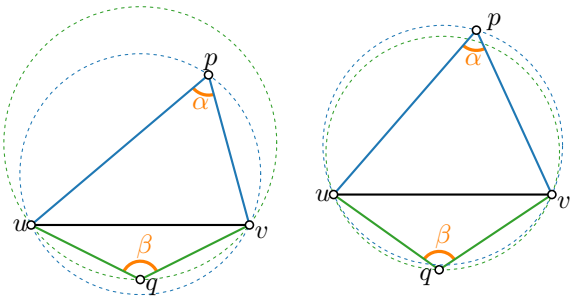

Figure 1: Opposing Angles metric. (left) Sum of opposing angles is less than π for quadrilateral that is not locally Delaunay. (right) It is at least π for quadrilaterals that are locally Delaunay.

Dual Edge Ratio. Additionally, Mitchell [16] suggested "measuring the signed distance between circumcenters of triangles sharing an edge; for Delaunay triangulations this is simply the dual edge length and non-negative, but for non-Delaunay triangulations the circumcenters can be in the wrong order and hence have a negative distance between them. So one could look at the ratio of the dual edge signed-length to the primal edge length (for 2D triangulations) as a continuous measure of how close it is to non-Delaunay." [16] – which are also known as Hodge-optimized triangulations [14].

We adapt this metric to evaluate quadrilaterals: we measure only the negative distance part of the suggestion, to satisfy criterion *C1* and not distinguish between Delaunay triangles. The "wrong order" referred to above matched to the quadrilateral being not locally Delaunay. We thus define the Dual Edge Ratio as

$$\begin{cases} 0, & \text{if the quadrilateral is locally Delaunay.} \\ \frac{d(c_p, c_q)}{d(u,v)}, & \text{otherwise.} \end{cases}$$

Intuitively, if the quadrilateral is not locally Delaunay, ratio of the distance between c_p and c_q and the length of (u, v) roughly corresponds to how skinny the triangles are and thus how far from Delaunay they are as well (see Fig. 2). Clearly, the Delaunay triangulation scores zero on all its quadrilaterals. Any quadrilateral that is not locally Delaunay, and hence does not locally describe the dual of the Voronoi diagram, will score greater than zero. Any non-Delaunay triangulation must have at least one such quadrilateral.

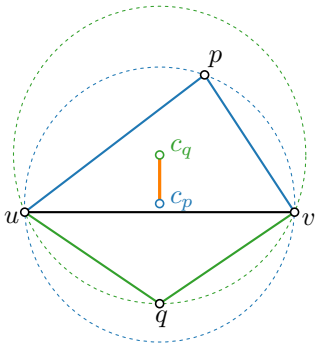

Figure 2: Dual Edge Ratio metric. The dual edge (orange) for a quadrilateral that is not locally Delaunay. We measure its length relative to the length of (u, v).

Dual Area Overlap. Similar to the previous metric, we may use the duality to Voronoi diagrams but in a different manner. Rather than looking at the distance, we may also consider an area-based metric. Intuitively, we consider the incorrect area of overlap between the "local Voronoi cells" that we may construct from the triangles of the quadrilateral.

The local Voronoi cell of p (and q analogously) is defined by the bisectors of u and v with p, which cross in c_p. If the quadrilateral is locally Delaunay, then the cells of p and q are disjoint. However, if the quadrilateral is not locally Delaunay, then they must overlap. The area of this overlap divided by the squared length of (u, v) is our Dual Area Overlap metric (see Fig. 3). We normalize using the squared edge length here, to ensure scale invariance, criterion *C3*.

The intuition behind this metric is similar to the intuition from the Dual Edge Ratio metric. The area of the overlap of the two Voronoi regions implies how far the quadrilateral is from having a proper non-overlapping Voronoi dual. A larger area means that points will have to move further to reach a non-overlapping dual, while a smaller area implies that a small movement in the points can already achieve this.

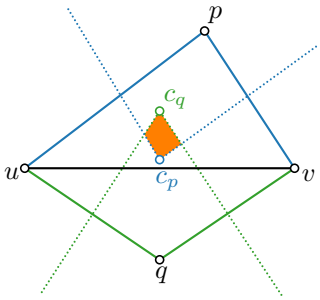

Figure 3: Dual Area Overlap metric. The local Voronoi cells of p and q (dotted) for a quadrilateral that is not locally Delaunay. We measure the overlap (orange).

2.2 Edge-based metrics

We now turn to edge-based metrics, that evaluate each edge (u, v) of T in context of all other points in P. That is, it penalizes edges, even if the defined quadrilateral is locally Delaunay. We present two new metrics below, both of which are based on the same principle: as an edge of the Delaunay triangulation must be the chord of a circle that does not strictly contain any other vertices of P, we consider how much we much deform a circle to find such an empty deformed circle instead. The difference between our two metrics is how they perform this deformation. Note that in both cases, higher scores mean closer to Delaunay, contrasting quadrilateral-based metrics.

Lens. When an empty circle exists, it can be seen as two circulars arcs, one on each side of the edge. The arc in the one halfplane with respect to the line spanned by (u, v) excludes from its interior all vertices of P in that same halfplane. With our Lens metric, we reverse this idea to deform our circle into a lens; the "sharpness" of this lens is then our metric.

Specifically, consider all points $P' \subseteq P$ that lie on one side of the line spanned by edge (u, v). The circular arc from u to v through some point of P' that is minimal in terms of segment area (or equivalently, arc length or central angle) is the largest arc possible on this side of the edge that does not contain any point of P' in its segment area. Let a and a' denote the two circular arcs obtained this way. We consider the "interior" angle α between the tangent directions of a and a' at u as our metric (see Fig. 4). If $\alpha \geq \pi$, the edge is a Delaunay edge and we cap the metric to π to satisfy criterion *C1*; for any non-Delaunay edge, $\alpha < \pi$.

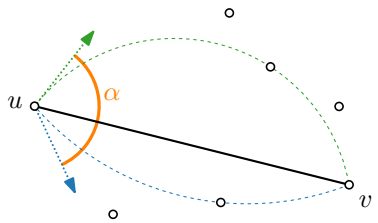

Figure 4: Lens metric. We find the largest empty arc on both sides of (u, v), and measure the angle α between their tangent directions at u.

We can easily compute the metric for a triangulation in quadratic time: for each edge, find the point that gives the smallest arc on both sides; then compute the angle at which the arcs touch.

Shrunk Circle. With our second metric, we consider a different type of deformation: scaling. That is, we aim to find a smaller empty circle that relates to the edge we wish to measure. For a non-Delaunay edge, such a circle cannot have (u, v) as a chord, but it may still over-

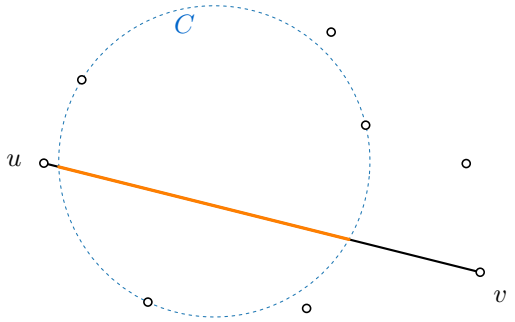

Figure 5: Shrunk Circle metric. We find the empty circle C that covers edge (u, v) most (orange).

lap the edge. As a Delaunay edge would be overlapped fully by an empty circle, we define our Shrunk Circle metric as the maximal fraction of the edge that can be overlapped by an empty circle C (see Fig. 5)). Note that, whereas the Lens metric considers both halfplanes independently, this is not the case here.

To compute this measure for an edge (u, v) of triangulation T on point set P, we observe the following: if an empty circle does not touch at least two points of P, we can readily grow it to a circle C' that does touch two points and strictly encompasses the previous circle and thus the overlap with (u, v) does not decrease. In other words, we need to consider only maximal circles with centers on the Voronoi diagram P. The lemma below argues that the overlap along one edge of the Voronoi diagram is convex and thus we need to test only its endpoints. In fact, we need to test only its vertices since unbounded edges occur only for pairs of points on the convex hull and thus the overlap of such circles with (u, v) are determined only by the circle's part inside the convex hull – which is maximized at bounded side of the unbounded edge. We can thus first compute the Delaunay triangulation for P, and then for every edge (u, v) of T, test the circumcircles of the Delaunay triangulation explicitly. This readily gives an quadratic-time algorithm for the overall metric.

Lemma 1 *Let e be an edge of the Voronoi diagram of P. The overlap of maximal circles along e with edge (u, v) is a convex function.*

Proof. Let p and q denote the points defining e. As the problem is invariant under translation, rotation and scaling, assume without loss of generality that $p = (0, 1)$ and $q = (0, -1)$; this implies that e is along the horizontal axis. Maximal circles C' are thus fully defined by their center $(m, 0)$. Let $\ell\colon y = ax + b$ denote the line spanned by (u, v). We first consider the overlap of C' with ℓ as a function of m.

The two intersection points of $\delta C'$ with ℓ are obtained by solving $|(x, ax + b) - (m, 0)|_2 = |(c, 0) - p|_2$, which simplifies to $(a^2+1)x^2+2(ab-m)x+(b^2-1) = 0$. The difference in x-coordinates between the two solutions to this quadratic equation are given by $\sqrt{D}/(a^2+1)$, where $D = 4(ab-m)^2-4(a^2+1)(b^2-1)$ is the discriminant; the amount over overlap with ℓ is thus $a\sqrt{D}/(a^2+1)$. As D is a convex quadratic function in m, so is $a\sqrt{D}/(a^2+1)$. To note, D is negative if the circle does not overlap ℓ, in which case the overlap is trivially zero: technically, the overlap is thus a function $a\sqrt{\max\{0, D\}}/(a^2+1)$. This proves that the overlap with ℓ is a convex function.

Since we are only interested in the values of m along e, that is, for which C' is empty, the overlap with ℓ is either fully within (u, v) or fully outside. Since the intersections behave continuously, the overlap with (u, v) thus behaves convex as well. □

2.3 Triangle-based metrics

For our triangle-based metrics, we measure a triangle $\triangle uvw$ of T in context of all other points in P. Analogous to our edge-based metrics, we deform the circumcircle C of $\triangle uvw$ (which is empty for a Delaunay triangle) to find a suitable empty deformed circle. In contrast to the edge-based metrics, we now have a single fixed circle C which guides (and constrains) our metric. Specifically, we restrict our deformations to be contained in the circumcircle. As with edge-based metrics, higher scores mean closer to Delaunay.

Triangular Lens. Similarly to the Lens metric, when a empty circumcircle exists, it can be considered as three circular arcs on the outside of the triangle. We again replace each arc by the largest arc that is contained in the arc of the circumcircle and that contains no other points of P.

For our metric, we measure the fraction of the area in C but outside $\triangle uvw$ that is covered by the constructed lens. Let a_{uv}, a_{vw} and a_{wu} denote the segment areas of the three arcs constructed (see Fig. 6). Interpreting C

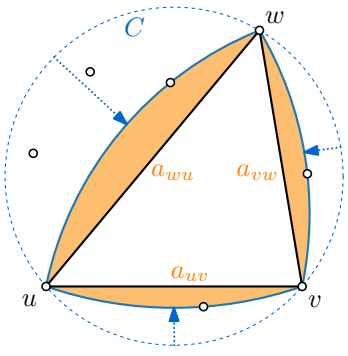

Figure 6: Triangular Lens metric. We compute the largest arc inside C that does not contain any points in its segment area for each edge of the triangle. Note that the three arcs are determined independently.

and $\triangle uvw$ as their enclosed areas, the score is then

$$\frac{a_{uv} + a_{vw} + a_{wu}}{C - \triangle uvw}.$$

We subtract the triangle to be able to assess and to meaningfully aggregate both skinny and fat triangles; this also means that each triangle's score lies in $(0, 1]$, where a score of 0 is only achievable in the limit (a point converging on each edge).

We easily compute the metric for a triangulation in quadratic time: for each triangle, find for each edge the smallest arc by testing all other vertices; then compute the segment areas and compute the resulting fraction.

Shrunk Circumcircle. As with our edge-based measures, the Triangular Lens metric deforms independently in the three segment areas defined by the edges. With our Shrunk Circumcircle metric, we consider a variant that considers all points simultaneously instead.

Specifically, we aim to find an empty circle C' that is contained in the circumcircle C and intersects all three sides of $\triangle uvw$. The score we associate with this circle C' is $\frac{C'-I}{C-I}$, where I is the area of the inscribed circle and we identify C' and C with the areas of these circles as well (see Fig. 7 (left)).

Note that we subtract the areas (numbers), we do not use set subtraction (difference of the shapes) and their resulting area; in our figures, we use examples where the inscribed circle is contained in the largest empty circle such that these two variants are the same, to easily visualize the score, but this is not necessarily the case.

The Shrunk Circumcircle metric is the circle C' with the highest associated score. For any triangle, I and C are constant, and thus it is simply the largest circle satisfying the two constraints. We subtract I as a lower bound on the largest empty circle (since the triangle itself must be empty), and divide by $C - I$ to normalize the the score to the range $(0, 1]$ independent of triangle shape and make the metric scale invariant.

We constrain C' to lie within C, such that the area we count is always an actual circle: if C' was to be allowed to grow outside C, it would either count area outside of the area that the Delaunay triangulation "considers", or the region we use the area of is not a circle in itself.

We constrain C' to intersect all three sides to make the score meaningful, even for very skinny triangles: otherwise, the largest empty circle may simply be fully in one segment, and even yield a relatively high score, though points are very close to the triangle edges (see Fig. 7 (right)).

To compute the metric for a triangle $\triangle uvw$ in context of P, we first argue about the properties of largest C'. Circle C' must either touch two points strictly inside C, or touch one such point and the boundary of C – otherwise, we can grow the circle into one that strictly encloses C' while still adhering to the constraints. As such, its center lies on the edges of the *local Voronoi diagram* of circle C and the points of P inside C (see Fig. 8). Curved segments for circle centers equidistant to the boundary of C and one of the points inside C, and straight segments defined by two points inside C. These curved segments are elliptical: the distance from a point on this curve to the center of C and to the contained point of P sums up to the radius of C. By construction, segments incident to the outerface of the diagram are curved segments; interior segments are straight.

In case the entire ellipse is part of the local Voronoi diagram (and hence in fact the only curve), we may consider it an elliptical arc, starting and ending at the furthest point from its defining point. Then, both along a straight segment as well as along a curved segment of the local Voronoi diagram, the radius of C' behaves unimodally: the function has a single (local and global) minimum at the closest point of the segment to the defining point(s). Along a segment, the maximal circle is hence either one of a small set of *critical placements* (if they exist): the endpoints of the segment, the fur-

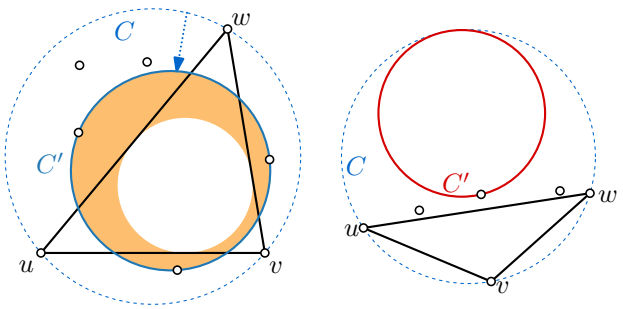

Figure 7: Shrunk Circumcircle metric. (left) We look for the empty circle C' in C of maximal radius. The score (orange) is its area minus the inscribed circle. (right) C' must intersect all three sides, for the measure to be meaningful when points are close to the triangle.

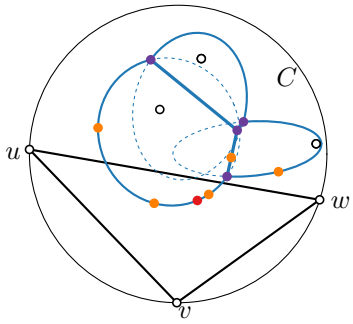

Figure 8: The local Voronoi diagram (blue) consists of elliptical and straight segments. The critical placements: endpoints of straight segments (purple), the furthest point on an ellipse (red), and circle touching a triangle edge (orange).

thest point of the ellipse, or one of the $O(1)$ placements where the defined circle touches a triangle edge. We further observe that the straight segments are a subset of the (full) Voronoi diagram of P.

With the insights above, we can define an algorithm to compute the Shrunk Circumcircle metric for all triangles in T in quadratic time. First, we compute the (full) Voronoi diagram of P. Then, for each triangle, we compute the metric as follows, in linear time. If there are no points inside the circumcircle, we have a Delaunay triangle and the metric is 1. If the circumcircle contains one point, then the local Voronoi diagram is a single ellipse and we test the critical placements. If the circumcircle contains more than one point, we find all straight segments local Voronoi diagram by traversing the full Voronoi diagram, testing the critical placements. For any straight segment found as such, we first shorten it to ensure that it does not define circles extending outside C. Then we test its critical placements. If the segment was shortened, we know that its defining points also define a segment of the local Voronoi diagram. For these segments, we also test the critical placements. By "testing" in the above, we mean testing whether the defined circle intersects all three sides (it is contained in C by construction), and if so, see if its radius is larger than any circle found so far. As we test $O(1)$ critical cases per segment, computing the metric takes linear time per triangle (after computing the full Voronoi diagram in $O(P \log P)$ time once) and thus quadratic time for the entire triangulation.

3 Comparing metrics

In the previous section, we defined seven near-Delaunay metrics. These capture in different ways how close to Delaunay a triangulation is. Though future work may endeavor to establish a standard here, it is not a-priori clear which measure is "the best": this likely depends on context, that is, the purpose of evaluating a triangulation. The question we ask here, is whether these metrics actually capture different facets of being "near-Delaunay". Specifically, given two metrics, do they always evaluate the same triangulation to be closer to Delaunay, for any given pair of triangulations?

Considering a single decomposition element to answer this question, there is clear distinction between the forms of decomposition: quadrilateral-based metrics do not take other points into account, contrasting the other two forms; edge-based metrics use angles and lengths, whereas triangle-based metrics use area ratios instead. We thus focus here on comparisons between metrics using the same form of decomposition. We study how the metrics differ from each other, and what properties they value. Specifically, for each comparison of metrics μ and μ', we show that there are triangulations T and T', such that $\mu(T) = \mu(T')$ and $\mu'(T) > \mu'(T')$. Such an example answers the above question of making the same judgments negatively.

Quadrilateral-based metrics. Our three quadrilateral-based metrics use only four points in measuring one element. Yet, we show that each metric evaluates a different facet of being near-Delaunay. Note that a (convex) quadrilateral is immediately a triangulation on this point set.

The Opposing Angles metric does not consider how the two summed angles are distributed over the two triangles. Even a very flat triangle (angle approaching π) can be offset a very tall triangle (angle approaching 0). Yet, these angles behave very differently from the distances between the circumcenters and thus the Dual Edge Ratio and Dual Area Overlap. Thus, we can construct two quadrilaterals (see Fig. 9): one has very similar triangles, while the other has very different triangles, but both have the same edge (u, v). Whereas the Opposing Angles metric scores the same on both, we see readily that the Dual Edge Ratio and Dual Area Overlap score differently.

Comparing the Dual Edge Ratio to the Dual Area Overlap, we see that the former is based purely on the distance between the circumcenters, whereas the latter depends also on the shape of the triangles. This allows

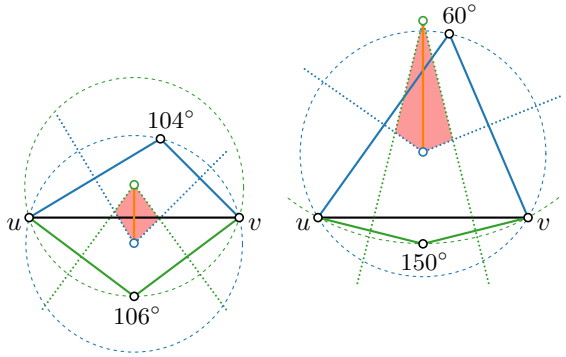

Figure 9: Two quadrilaterals with equal (sum of) Opposing Angles, but different Dual Edge Ratio (orange) and Dual Area Overlap (red).

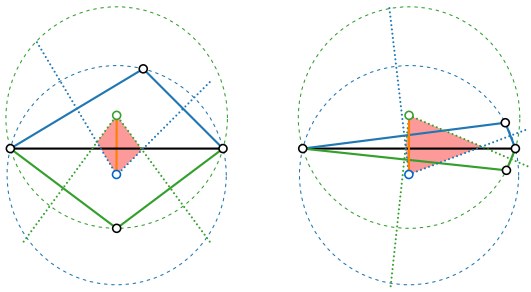

Figure 10: Two quadrilaterals with equal Dual Edge Ratio (orange) and different Dual Area Overlap (Red).

us to construct another two quadrilaterals (see Fig. 10): we move the opposing points along their defined circumcircles so as to not move the circumcenters, while changing the area of overlap.

Edge-based metrics. We have two metrics here: Lens and Shrunk Circle. As already mentioned in the previous sections, these differ by how they treat points on the different sides of the edge. Whereas the Lens measure treats these independently, the Shrunk Circle measure requires an integrated consideration of all points. This allows us to construct two triangulations again, using only four points in convex position (see Fig. 11): we keep the Lens measure constant, by moving one of the vertices over the defining arc, thus keeping the tangent at u constant as well; in contrast, the Shrunk Circle metric in the first example uses a circle that encompasses a large part of the defining arc of the Lens – by placing the point there, we can force the circle to shrink further and cover less of the edge (u, v).

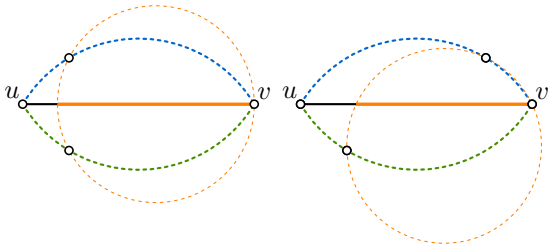

Figure 11: Two edges with equal Lens (blue, green) different Shrunk Circle (orange).

Triangle-based metrics. Finally, we consider the two triangle-based metrics. They behave somewhat similarly with respect to each other as the edge-based metrics do: whereas the Triangular Lens works with independent arcs per edge, the Shrunk Circumcircle uses a single circle that must overlap each of the three edges. We can thus follow the same principle to show two triangles in a point set that score equally on the Triangular Lens, but differently on the Shrunk Circumcircle, by moving the points along the arcs of the former to force the defining circle of the latter to shrink (see Fig. 12).

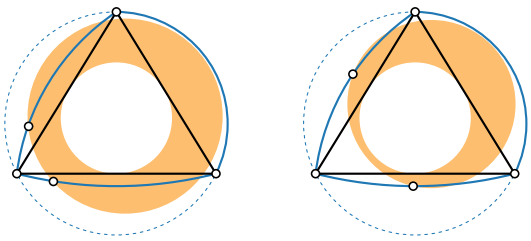

Figure 12: Two triangles with equal Triangular Lens (blue) and different Shrunk Circumcircle (orange).

4 Experiments

Here we explore what the most Delaunay-like triangulation of a point set looks like, for each of our metrics, given different constraints that can force a triangulation to be non-Delaunay. For a point set, we try all triangulations that adhere to the given constraints to find the *optimized* triangulation that scores best according to each metric. We use small sets of only 10 points, to ensure that this is feasible.

Constraints. We consider four types of constraints: constrained edges, a lower bound or upper bound on the total edge length, and a maximum degree.

For constrained edges, we are given a set of edges that must be included in the triangulation. The constrained edges are handpicked edges, which we consider to be "interesting". Most importantly the constrained edges are not in the Delaunay triangulation, are not chord of the convex hull, and lie somewhat close to another point. This is the same constraint as for the CDT and hence we may also compare how this structure compares to our result.

The length of the triangulation is the sum over the lengths of its edges. As the Delaunay triangulation neither minimizes nor maximizes the length, we can use an upper bound (maximum length) or a lower bound (minimum length), to constrain the triangulation to be non-Delaunay. Specifically, we use either 1.2 times the length of the Delaunay triangulation as a lower bound, or 0.8 times the length of the Delaunay triangulation as an upper bound. However, since the Delaunay triangulation is inherently short for many point sets, such triangulations often do not exist. We hence also create a special point set where the Delaunay triangulation is longer than most other triangulations, to see the differences between metrics.

Bounding the maximum degree means we consider only triangulations for which all vertices have degree at most a given constant. We use maximum degree 5 in our experiment. This number is sometimes exceeded by the Delaunay triangulation, but still allows for different triangulations of the same point set.

Aggregation. For our optimized triangulations, we have to evaluate each metric on the entire triangulation. So far, we have left the method of aggregation out of our considerations. For the purpose of our experiment, we consider two methods: using the sum and using bottleneck values.

The sum is a natural way to aggregate the values of a triangulation T, such that all decomposition elements (quadrilaterals, edges or triangles) have an impact on the triangulation. For quadrilateral-based metrics, we minimize the sum; otherwise, we maximize it.

Using bottleneck values means we focus on the worst-case element (maximum for quadrilateral-based metrics,

minimum for our other metrics). The goal of this approach is to let the worst value be the deciding factor. A problem with this method is that there can be two different triangulations that score the same, as they both include the same bottleneck. Hence, we compare triangulations lexicographically: this means that the elements are sorted (increasingly for quadrilateral-based, decreasingly for other metrics) and the first element in which the triangulations differ, determines which triangulation is closer to Delaunay.

Results. Tables 3 through 10 in the full version [21] show our experimental results; Table 2 contains the most relevant excerpts. For the constrained edges, we show the CDT for comparison, using red edges to indicate the constraints. For the other constraints, we show the Delaunay triangulation for comparison. In both cases, the optimized triangulations use green markings to show edges that are different from the comparison.

Table 3 in [21] shows optimized triangulations with constrained edges using sum aggregation. We observe that most metrics are always similar to the CDT for such random point sets. The only differences are for Dual Edge Ratio and Dual Area Overlap: each time Dual Edge Ratio is different from the CDT, Dual Area Overlap is also different, though not necessarily vice versa. One exception is for Shrunk Circle. This suggests to us that, at least in small random cases, our metrics capture near-Delaunay quite well, as the CDT is an established way of getting a Delaunay-like triangulation, for these constraints. Using bottleneck aggregation (Table 4 in [21]), we observe that the quadrilateral metrics often behave similarly and are regularly different from the CDT. With one exception, the other optimized triangulations are again all identical to the CDT.

Tables 5 and 6 in [21] show the optimized triangulations with a minimum-length constraint. We can compare the different optimizations to empirically evaluate the behavior of the metrics. We observe that in the sum aggregation quadrilateral-based metrics generally behave differently than the other metrics. However, this distinction is less clear using the bottleneck aggregation. Furthermore, with bottleneck aggregation the Shrunk Circle metric often produces a unique triangulation. Note that this is the same case that was excepting from the general trend for constrained edges.

Table 7 and 8 in [21] show the optimized triangulation with a maximum-length constraint. If no such triangulation exists, we simply show the Delaunay triangulation. In every point set except the one specifically created for this constraint, there was no such triangulation. For this case, we observe that the Dual Edge Ratio and Dual Area overlap produce different triangulations from the other metrics. In the bottleneck variant, we also see that the Lens measure procedures another distinct triangulation.

Table 9 and 10 in [21] shows the optimization with a maximum-degree constraint. In only few of the random point sets, the maximum degree exceeds the constraint. Most metrics flip the same edge, the notable exception being the Shrunk Circumcircle metric which flips two different edges for one case (shown). The two other point sets shown were specifically created for this case. We see considerable differences here between metrics, as well as between sum and bottleneck aggregation.

With the wheel example here, we may perhaps see the one case where there is a somewhat clear case of a "visually nice" optimized triangulation. Specifically, a somewhat regular pattern emerges for Lens and Shrunk Circle using sum aggregation as well as all metrics except for the Dual Edge Ratio and Dual Area Overlap using bottleneck aggregation – in the other cases we see very skinny triangles occurring. Whether this form of being near-Delaunay, however, is the most useful remains possibly context dependent. This hints at bottleneck aggregation perhaps being more useful – indeed, it matches the traditional lexicographic optimization of the minimum angle that the Delaunay triangulation achieves.

5 Discussion

With a suite of metrics, we now have a common framework to think about situations where we need a good (Delaunay-like) triangulation to compute with, but constraints such as bounded degree prevent us from actually using the Delaunay triangulation itself. We have shown how these metrics differ among themselves as well as how they result in different triangulations when considering optimizing under the constraint of including given edges (like the Constrained Delaunay Triangulation).

We leave to future work to establish efficient algorithms to compute the best triangulation (given one of the metrics) given a set of constraints. It may further be interesting to investigate how humans (or computational geometers) assess the quality of a triangulation, how close it is to being Delaunay, and how this relates to the metrics provided here. This may uncover that at least to match intuition, we may need combinations of metrics, or possibly a different metric altogether.

A possible avenue for further metrics is to consider the empty ellipse as the generalization of the empty circle, trying to optimize for the ellipse's aspect ratio. This seems mostly relevant for edge-based or quadrilateral-based metrics, since triangles do not necessarily allow for empty ellipses passing through their corners. But for a single edge or quadrilateral, a general ellipse seems to provide too much freedom, allowing for somewhat arbitrary-seeming results. Restricting one of the axes of the ellipse to be parallel to the defining edge may offer a solution; yet, this seems to cause a counterintuitive relation between the two sides of the edge.

Table 2: Excerpt with the most relevant results, summarizing the appendix in the full version [21].

Optimizing with constrained edges (Sum aggregation)

CDT	Opposing Angles	Dual Edge Ratio	Dual Area overlap	Lens	Shrunk Circle	Triangular Lens	Shrunk Circumcircle
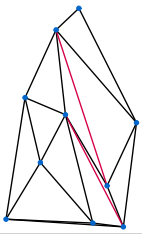	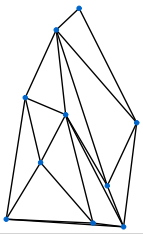	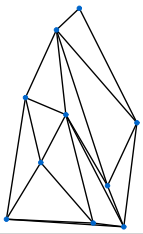	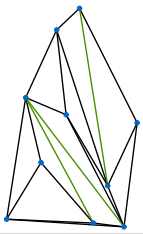	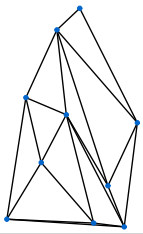	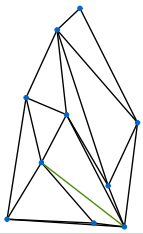	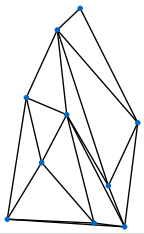	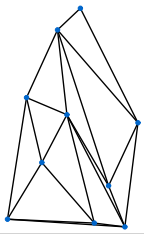

Optimizing with constrained edges (Bottleneck aggregation)

CDT	Opposing Angles	Dual Edge Ratio	Dual Area overlap	Lens	Shrunk Circle	Triangular Lens	Shrunk Circumcircle
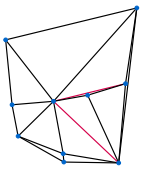	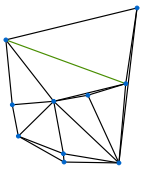	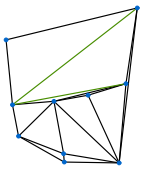	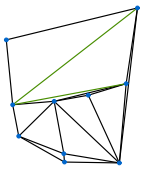	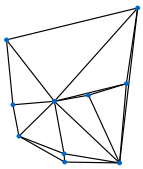	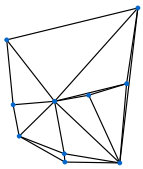	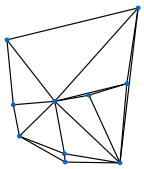	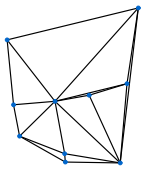

Optimizing with minimum length 1.2 Delaunay length (Sum aggregation)

DT	Opposing Angles	Dual Edge Ratio	Dual Area overlap	Lens	Shrunk Circle	Triangular Lens	Shrunk Circumcircle
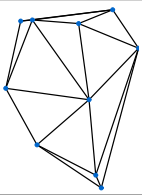	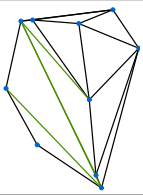	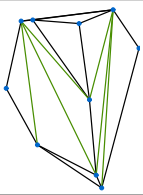	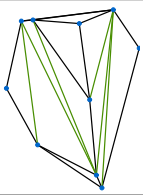	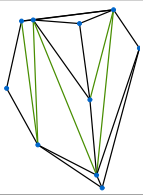	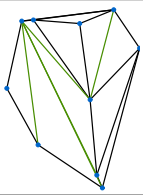	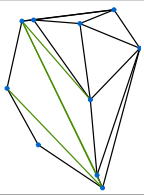	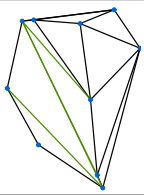

Optimizing with maximum length 0.8 Delaunay length (Bottleneck aggregation)

DT	Opposing Angles	Dual Edge Ratio	Dual Area overlap	Lens	Shrunk Circle	Triangular Lens	Shrunk Circumcircle
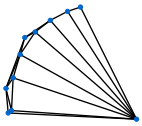	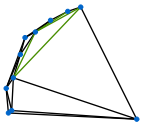	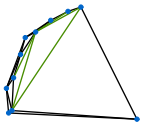	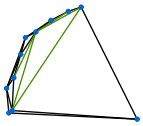	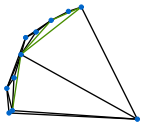	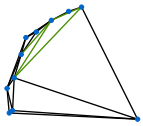	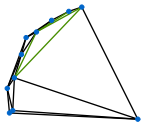	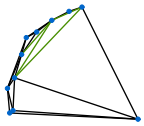

Optimizing with maximum degree 5 (Sum aggregation)

DT	Opposing Angles	Dual Edge Ratio	Dual Area overlap	Lens	Shrunk Circle	Triangular Lens	Shrunk Circumcircle
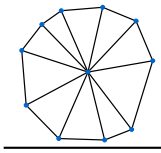	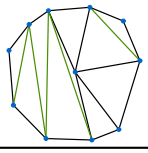	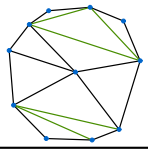	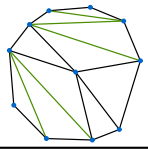	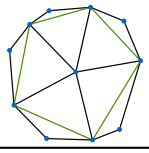	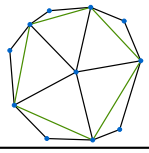	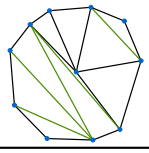	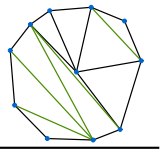

Optimizing with maximum degree 5 (Bottleneck aggregation)

DT	Opposing Angles	Dual Edge Ratio	Dual Area overlap	Lens	Shrunk Circle	Triangular Lens	Shrunk Circumcircle
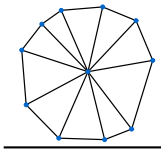	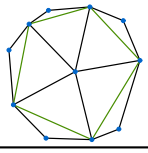	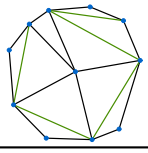	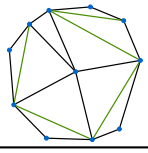	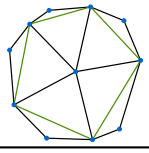	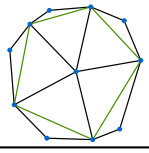	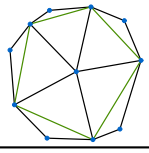	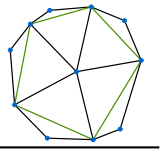

References

[1] O. Aichholzer, F. Hurtado, and M. Noy. A lower bound on the number of triangulations of planar point sets. *Computational Geometry*, 29(2):135–145, 2004.

[2] M. Bern, H. Edelsbrunner, D. Eppstein, S. Mitchell, and T. S. Tan. Edge insertion for optimal triangulations. *Discrete & Computational Geometry*, 10(1):47–65, 1993.

[3] L. P. Chew. Constrained Delaunay triangulations. *Algorithmica*, 4(1-4):97–108, 1989.

[4] D. P. Dobkin, S. J. Friedman, and K. J. Supowit. Delaunay graphs are almost as good as complete graphs. *Discrete & Computational Geometry*, 5(4):399–407, 1990.

[5] E. F. D'Azevedo and R. B. Simpson. On optimal interpolation triangle incidences. *SIAM Journal on Scientific and Statistical Computing*, 10(6):1063–1075, 1989.

[6] P. Giannopoulos, R. Klein, C. Knauer, M. Kutz, and D. Marx. Computing geometric minimum-dilation graphs is NP-hard. *International Journal of Computational Geometry & Applications*, 20(02):147–173, 2010.

[7] J. Gudmundsson, M. Hammar, and M. van Kreveld. Higher order Delaunay triangulations. *Computational Geometry*, 23(1):85–98, 2002.

[8] J. Gudmundsson and C. Knauer. Dilation and detours in geometric networks. In T. F. Gonzalez, editor, *Handbook of Approximation Algorithms and Metaheuristics*. Chapman and Hall/CRC, 2007.

[9] K. Jansen. One strike against the min-max degree triangulation problem. *Computational Geometry*, 3(2):107–120, 1993.

[10] G. Kant and H. L. Bodlaender. Triangulating planar graphs while minimizing the maximum degree. *Information and Computation*, 135(1):1–14, 1997.

[11] J. M. Keil and C. A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete & Computational Geometry*, 7(1):13–28, 1992.

[12] C. L. Lawson. Software for C1 surface interpolation. In *Mathematical software*, pages 161–194. Elsevier, 1977.

[13] D.-T. Lee and A. K. Lin. Generalized delaunay triangulation for planar graphs. *Discrete & Computational Geometry*, 1(3):201–217, 1986.

[14] P. Mullen, P. Memari, F. de Goes, and M. Desbrun. HOT: Hodge-optimized triangulations. In *Proc. ACM SIGGRAPH 2011*, Article 103, pages 1–12, 2011.

[15] W. Mulzer and G. Rote. Minimum-weight triangulation is NP-hard. *Journal of the ACM (JACM)*, 55(2):1–29, 2008.

[16] J. O'Rourke. Open problems from CCCG 2017, 2018.

[17] A. Pilz. Flip distance between triangulations of a planar point set is APX-hard. *Computational Geometry*, 47(5):589–604, 2014.

[18] V. T. Rajan. Optimality of the Delaunay triangulation in $\mathbb{R}^d$. *Discrete & Computational Geometry*, 12(2):189–202, 1994.

[19] M. Sharir and A. Sheffer. Counting triangulations of planar point sets. *The Electronic Journal of Combinatorics*, 18(P70):1, 2011.

[20] R. Sibson. Locally equiangular triangulations. *The Computer Journal*, 21(3):243–245, 1978.

[21] N. van Beusekom, K. Buchin, H. Koerts, W. Meulemans, B. Rodatz, and B. Speckmann. Near-Delaunay Metrics, 2021. `https://arxiv.org/abs/2106.11621`.

[22] M. van Kreveld, M. Löffler, and R. I. Silveira. Optimization for first order Delaunay triangulations. *Computational Geometry*, 43(4):377–394, 2010.