An improved method for generating triangle-mesh models of images

Xiao (Brian) Ma* Michael D. Adams*

Abstract

In earlier work, Yang et al. proposed a highly-effective technique for generating triangle-mesh models of digital images, known as the error diffusion (ED) method. In this paper, we propose a modified version of the ED method that better exploits triangulation connectivity. Through experimental results, our proposed method is shown to generate image approximations of substantially higher quality than those obtained with the ED scheme, by an average margin of about 3 decibels in terms of the peak-signal-to-noise ratio. Moreover, this improvement in quality comes at a relatively modest computational cost, with the proposed method typically requiring only a few seconds of computation time.

1 Introduction

In recent years, there has been a growing interest in representations of digital images that are based on nonuniform (i.e., content-adaptive) sampling. Because most real-world images are non-stationary, uniform sampling of images (e.g., using a rectangular grid) is almost guaranteed to be highly suboptimal, as it places too few sample points in regions of rapid change and too many sample points in regions of slow variation. For this reason, it is desirable to select the sample points in a nonuniform manner that is dependent on the image content.

Although many classes of methods for nonuniform sampling have been considered over the years, a particularly effective one is the class based on triangle meshes. In this approach, the (nonuniformly) chosen sample points are triangulated, partitioning the image domain into triangular faces, and then an approximating function is constructed over each face of the triangulation. One key difference between the various triangle-mesh-based approaches is in how they select the triangulation connectivity (i.e., how vertices are connected by edges). In this regard, the most common approach is to choose the connectivity by using a Delaunay triangulation [5, 6]. In such a case, the connectivity is determined solely by the set of points being triangulated. Another approach is to use a data-dependent triangulation (DDT), which chooses the connectivity using information in the dataset from which the points being

---
*Department of Electrical and Computer Engineering, University of Victoria, Victoria, BC, Canada, {xiaom,mdadams}@ece.uvic.ca

triangulated were chosen. Since the connectivity of a DDT may be chosen in an arbitrary manner, DDTs offer much greater flexibility than Delaunay triangulations.

In [12], Yang et al. proposed a simple technique for generating triangle-mesh models of images, known as the error-diffusion (ED) method, which chooses the connectivity of the sample points using a Delaunay triangulation. Although quite effective, the ED method has the weakness that it often yields triangulations in which a significant number of triangulation edges crosscut image edges (i.e., discontinuities in the image), leading to a degradation in approximation quality. In this paper, we propose a modified version of the ED method that utilizes DDTs instead of Delaunay triangulations. Through experimental results, we show that our proposed method yields image approximations of much higher quality (i.e., lower squared error) than the ED method, with a relatively modest computational cost.

The remainder of this paper is organized as follows. To begin, Section 2 provides some background information on triangle meshes for image representation and introduces some key methods related to our work. Our proposed mesh-generation method is presented in Section 3. This is done by starting with a general computational framework having a free parameter, and then explaining how this free parameter was chosen in order to arrive at our proposed method. Through experimental results, Section 4 shows our proposed method to yield image approximations of much higher quality than the ED method, with a relatively modest computational cost. Finally, Section 5 concludes with a brief summary of our key results and some closing remarks.

2 Background

In what follows, the cardinality of a set S is denoted $|S|$. Consider an image function ϕ defined on the domain $I = [0, W-1] \times [0, H-1]$ and sampled on the truncated two-dimensional integer lattice $\Lambda = \{0, 1, \ldots, W-1\} \times \{0, 1, \ldots, H-1\}$ (i.e., a rectangular grid of width W and height H). A (triangle) mesh model of ϕ consists of: 1) a set $P = \{p_i\} \subset \Lambda$ of sample points and their corresponding function values $\{z_i = \phi(p_i)\}$; and 2) a triangulation T of P. As a matter of terminology, the size and sampling density of the model are defined as $|P|$ and $|P| / |\Lambda|$, respectively. The mesh model is associated with a piecewise linear approximating function $\hat{\phi}$ that interpolates ϕ at each point in P. More specifi-

cally, over each face in T, $\hat{\phi}$ corresponds to the unique linear function that interpolates ϕ at the three vertices of the face. In our work, for a given model size (i.e., number of sample points), we want to choose the model to minimize the mean squared error (MSE) ϵ between $\hat{\phi}$ and ϕ, where

$$\epsilon = |\Lambda|^{-1} \sum_{p \in \Lambda} \left(\hat{\phi}(p) - \phi(p) \right)^2. \tag{1}$$

For convenience, we will express the MSE in terms of the peak-signal-to-noise ratio (PSNR), which is defined as $\mathsf{PSNR} = 20 \log_{10}[(2^\rho - 1)/\sqrt{\epsilon}]$, where ρ is the number of bits per sample used by the image ϕ. The PSNR expresses the MSE relative to a signal's dynamic range using a decibel (dB) scale, with higher PSNR corresponding to lower MSE.

ED Method. As mentioned earlier, one highly effective method for generating mesh models of images is the ED method [12]. Since our work builds on the ED method, it is worthwhile to make a few brief comments about this method here. In general terms, the ED method consists of two steps:

1. Sample-point selection. Use Floyd-Steinberg error diffusion [8] in order to select the set P of sample points (for the model) such that they are distributed with a density approximately proportional to the maximum-magnitude second-order directional derivative of the image.
2. Triangulation. Triangulate the sample points using a Delaunay triangulation.

In step 1, the set P is always chosen to include all of the extreme convex hull points of the image domain. This ensures that the triangulation generated in step 2 covers the entire image domain. Since several variants of the ED scheme are presented in [12], it is worth noting, for the sake of completeness, that we consider the variant with the following characteristics herein: 1) a third-order binomial filter is used for smoothing; 2) non-leaky error diffusion is used with a serpentine scan order; 3) the sensitivity parameter γ is chosen as 1; and 4) the error diffusion algorithm is performed iteratively in order to achieve exactly the desired number of sample points.

LOP. Next, we present some details about triangulations with the goal of introducing a key method of interest in our work which relates to DDTs. An edge e of a triangulation is said to be flippable if e has two incident faces (i.e., is not on the triangulation boundary) and the union of these two faces is a strictly convex quadrilateral q. For a flippable edge e, an edge flip is an operation that replaces the edge e in the triangulation by the other diagonal of q. The fact that every triangulation of a set of points is reachable from every other triangulation (of the same set of points) via a finite sequence of edge flips motivated Lawson to propose

the so called local optimization procedure (LOP) [10]. The LOP is a technique for selecting the connectivity of a triangulation so as to be optimal with respect to some prescribed criterion. As a matter of terminology, a triangulation is said to be optimal if every flippable edge in the triangulation is optimal with respect to the prescribed optimality criterion. The LOP simply applies edge flips to flippable edges that are not optimal until the triangulation is optimal. Many different optimality criteria have been proposed in the literature, numerous of which can be found in [7, 11]. One particularly important criterion in the context of our work is the squared error (SE) optimality criterion. The SE criterion deems an edge e to be optimal if applying an edge flip to e would not lead to a strict decrease in the MSE as given by (1). For most choices of optimality criterion (including SE), the optimal solution produced by the LOP is not uniquely determined. The nonuniqueness of the solution is important because it implies that some optimal solutions may be (and, in practice, are) much better than others. The LOP is frequently employed to choose triangulation connectivity in the case of DDTs.

3 Proposed Method and Its Development

Having introduced the necessary background, we now turn our attention to introducing the mesh-generation method proposed in this paper. As mentioned previously, our method is essentially a modified version of the ED scheme. The ED method, as explained earlier, chooses triangulation connectivity using a Delaunay triangulation. Experimentally, however, we have observed that selecting the connectivity in this way results in a mesh in which triangulation edges often crosscut image edges (i.e., discontinuities in the image), leading to a degradation in approximation quality. This motivated us to consider choosing triangulation connectivity in a more flexible manner, using a DDT instead of a Delaunay triangulation.

In what follows, we will first introduce the general computational framework associated with our method, where this framework has one free parameter. In order to arrive at the specific method proposed in this paper, we advocate a particular choice for this parameter. Since it is likely helpful for the reader to see how we arrived at this choice, we provide a few details in this regard, including some experimental results.

The general computational framework associated with our proposed method consists of the following steps:

1. Sample-point selection. Select the set P of sample points for a mesh model of the desired size, using the same sample-point selection strategy in step 1 of the ED method (as introduced earlier in Section 2).

2. Initial mesh construction. Construct the triangulation T by inserting the points in P in a triangulation. More specifically, for each point $p \in P$, starting with the extreme convex hull points of P (i.e., the four corners of the image bounding box) and followed by the remaining points in randomized order:

   (a) Insert p in the triangulation T. This is accomplished by deleting any faces containing p and retriangulating the resulting hole.

   (b) Adjust the connectivity of T by applying the LOP (described in Section 2) with the triangulation optimality criterion chosen as `oc`, where `oc` is a free parameter of our framework.

3. Final connectivity adjustment. Adjust the connectivity of T by applying the LOP with the triangulation optimality criterion chosen as SE (i.e., squared error).

In step 2b of the above framework, the choice of the optimality criterion `oc` is critical, as different choices of `oc` will typically lead to vastly differing meshes. One of the optimality criteria considered in our work is the SE criterion introduced in Section 2. We also considered numerous other criteria, which we will introduce shortly. Before proceeding further, however, there is a very important comment that we must make regarding our above computational framework. Since our objective is to produce a mesh that minimizes the MSE (as given by (1)), this suggests the "obvious" solution of choosing the optimality criterion `oc` as SE and simply skipping final connectivity adjustment (i.e., step 3) altogether (since final connectivity adjustment would not change anything if `oc` were chosen as SE). In other words, the obvious solution would be to simply optimize for squared error using the LOP after the insertion of each point in step 2. As it turns out, this obvious solution performs extremely poorly. This poor performance is due to an interplay between point insertion and the SE criterion in step 2b, which leads to triangulations with many poorly-chosen sliver (i.e., long thin) triangles, severely degrading approximation quality. In effect, this interplay causes the mesh-generation optimization process to become trapped at a very poor local optimum. To combat the above problem, our framework allows the parameter `oc` to be chosen differently from SE, and then adds a final connectivity-adjustment step employing the SE criterion in order to reduce the squared error for the final mesh.

Test data. Shortly, we will have the need to present some experimental results obtained with various test images. So, before proceeding further, a brief digression is in order to introduce the test images used herein. In our work, we have employed 40 images, taken mostly from standard test sets such as [1], [3], and [2]. For the most part, the results that we present herein focus

Table 1: Test images

Image	Size, Bits/Sample	Description
animal	1238×1195, 8	cartoon character (computer generated)
cr	1744×2048, 10	x-ray [1]
lena	512×512, 8	woman [3]
peppers	512×512, 8	collection of peppers [3]

on the representative subset of these images listed in Table 1. This particular subset was chosen to contain a variety of image types (i.e., photographic, medical, and computer-generated imagery).

Choice of `oc` parameter. As we saw above, our computational framework has the free parameter `oc`, which corresponds to the choice of optimality criterion used for the LOP (in step 2b). We will now briefly introduce the possibilities that we considered for `oc` and describe the experiments that led to our recommended choice for `oc`, which will be presented shortly.

In our work, we considered ten possibilities for `oc`:

1. SE (as introduced earlier);
2. Delaunay [6, 10];
3. angle between normals (ABN) [7];
4. jump in normal derivatives (JND) [7];
5. deviations from linear polynomials (DLP) [7];
6. distances from planes (DP) [7];
7. edge-length-weighted ABN (ELABN) [4, Section 2.2] (called absolute mean curvature in [4]);
8. Garland-Heckbert hybrid (GHH) [9, 11];
9. shape-quality-weighted SE (SQSE) [11]; and
10. JND-weighted SE (JNDSE) [11].

In the above list, we have provided references in which the formal mathematical definition of each of these optimality criteria can be found. In the interest of brevity, however, we have not included these definitions herein. Note that, if `oc` is chosen as SE, final connectivity adjustment (i.e., step 3) effectively does nothing, as the triangulation that is input to final connectivity adjustment is already optimal with respect to the SE criterion.

To determine which optimality criterion is best, we performed the following experiment. For each of the 40 images in our test set and five sampling densities per image, we generated a mesh with each of the ten possible choices for the optimality criterion `oc` and measured the resulting approximation error in terms of PSNR. A representative subset of the results (namely, the results for the images listed in Table 1) is given in Table 2, with the best result in each case shown in italic font. From Table 2, it is clear that the JNDSE criterion performs best, followed by the SQSE and GHH criteria. In particular, the JNDSE criterion yields the best result in 17/20 of the test cases. From Table 2, it is also evident that the SE criterion performs very poorly (as we claimed

to be the case earlier). In terms of the overall results taken over all $40 \cdot 5 = 200$ test cases (i.e., 40 images with five sampling densities per image), the JNDSE criterion performs best and second best in 164/200 (82%) and 26/200 (13%) of the test cases, respectively. Based on these results, the best choice for the oc parameter is clearly JNDSE.

As the above experimental results show, choosing oc as SE leads to very poor performance. Earlier, we explained that this behavior was due to an interplay between point insertion and the SE criterion, which leads to triangulations with many poorly-chosen sliver triangles. At this point, we provide an example to illustrate this phenomenon. For one of our test cases (from above), part of the triangulation and corresponding image approximation is shown in Figure 1 for the cases of SE (which performs poorly) and JNDSE (which performs well). From Figures 1(a) and (b), it is clear that the result obtained with the SE criterion has a large number of poorly-chosen sliver triangles, which leads to large error in the image approximation. In contrast, the result obtained from the JNDSE criterion does not exhibit such a problem, as can be seen from Figures 1(c) and (d).

Although the interplay between point insertion and the LOP optimality criterion is complex, we can suggest one possible reason why the SE criterion has a propensity to produce triangulations with many poorly-chosen sliver triangles. When a point is inserted in the triangulation, one or more sliver triangles can sometimes result. In some circumstances, when the SE optimality criterion is used, the LOP is less likely to eliminate such sliver triangles. This could, for example, be due to the squared error cost associated with an edge and its flipped version both being zero (e.g., when no points from the sampling grid Λ fall inside the associated quadrilateral), in which case no edge-flip is performed. In contrast, the JNDSE criterion does not suffer from this type of problem as it depends on a geometric criterion (namely, JND) in addition to squared error.

Proposed method. The preceding experimental results show the best choice for the optimality-criterion parameter oc to be JNDSE. Therefore, we recommend that oc be chosen as JNDSE in our framework, and the specific mesh-generation method that we propose in this paper is the one that uses our framework with this particular choice for oc.

4 Results

Having introduced our proposed method, we now evaluate its performance by comparing it in terms of mesh quality to the ED method. For test data, we again employ the same 40 images described earlier in Section 3 (under the "Test data" heading).
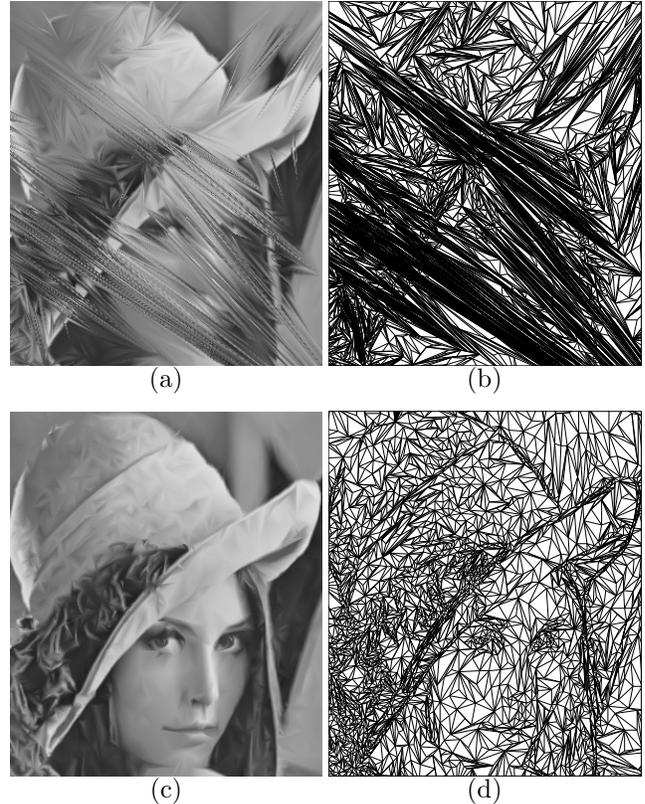

Figure 1: Part of the image approximation obtained for the lena image at a sampling density of 2% with the (a) SE (20.72 dB) and (c) JNDSE (29.99 dB) optimality criteria and (b) and (d) their corresponding triangulations.

To evaluate mesh-quality performance, we proceeded as follows. For all 40 images in our test set and five sampling densities per image, we used each of the proposed and ED methods to generate a mesh, and then measured the resulting approximation error in terms of PSNR. A representative subset of the results obtained (namely, the results for the four images listed in Table 1) is shown in Table 3. From Table 3, we can see that our proposed method outperforms the ED method in all 20 cases by a significant margin, namely, a margin of at least 1.88 dB. In terms of the overall results taken over all $40 \cdot 5 = 200$ test cases (i.e., 40 images with five sampling densities per image), our proposed method beats the ED scheme in all 200 test cases by a margin of 1.5 to 6.7 dB, with the average margin being approximately 3 dB. From the preceding results, it is clear that our proposed method yields meshes of much higher quality, in terms of PSNR, than the ED scheme.

In the above evaluation, PSNR was found to correlate reasonably well with subjective quality as perceived by the human visual system. For the benefit of the reader, however, we provide an example illustrating visual quality. For one of the test cases from Table 3, Figure 2

Table 2: Comparison of the mesh quality obtained with various choices of the `oc` parameter

Image	Sampling Density (%)	PSNR (dB)									
		SE	Delaunay	ABN	JND	DLP	DP	ELABN	GHH	SQSE	JNDSE
animal	0.50	28.45	37.32	30.44	37.09	33.34	28.92	36.40	*37.73*	37.71	37.70
	1.00	28.95	40.58	35.74	40.60	38.02	30.24	40.45	40.36	*40.67*	40.61
	2.00	34.60	42.93	36.66	42.76	40.27	29.23	42.74	42.86	43.26	*43.26*
	3.00	32.99	44.24	39.35	43.85	41.70	34.65	43.22	44.23	44.46	*44.49*
	4.00	36.52	45.23	39.91	44.79	41.58	35.89	44.32	45.27	45.46	*45.50*
cr	0.50	31.19	34.40	30.38	34.45	32.42	30.23	34.22	34.30	34.81	*34.84*
	1.00	32.41	36.33	33.01	36.35	34.42	31.16	36.37	36.48	37.13	*37.16*
	2.00	33.33	38.68	34.34	38.36	36.33	32.52	38.24	38.75	38.95	*39.01*
	3.00	34.12	39.57	34.95	39.32	36.96	33.78	39.17	39.62	39.76	*39.82*
	4.00	35.63	40.10	36.29	39.89	37.56	33.54	39.70	40.19	40.31	*40.36*
lena	0.50	17.61	21.17	19.22	20.55	19.61	18.07	20.51	21.20	21.75	*21.82*
	1.00	21.50	25.21	20.69	24.91	21.86	19.91	24.58	25.30	25.89	*25.92*
	2.00	20.72	29.48	24.36	29.09	26.25	21.04	27.67	29.26	29.91	*29.99*
	3.00	23.43	31.26	24.62	30.99	27.22	22.34	30.15	31.21	31.58	*31.62*
	4.00	23.67	32.39	26.30	32.17	29.13	24.06	31.45	32.47	32.78	*32.84*
peppers	0.50	17.78	19.38	17.81	19.18	18.26	16.38	19.25	19.61	20.52	*20.54*
	1.00	19.84	24.97	21.53	24.57	21.99	18.95	24.43	25.07	25.39	*25.64*
	2.00	22.66	29.39	22.55	29.06	26.40	20.64	28.74	29.22	29.67	*29.74*
	3.00	23.15	31.27	24.18	30.87	26.67	21.41	30.73	30.75	*31.52*	31.50
	4.00	22.68	32.08	26.63	31.88	28.65	23.01	31.55	32.00	32.32	*32.37*

Table 3: Comparison of the mesh quality obtained with the proposed and ED methods

Image	Sampling Density (%)	PSNR (dB)	
		Proposed	ED
animal	0.50	*37.70*	33.86
	1.00	*40.61*	37.66
	2.00	*43.26*	40.46
	3.00	*44.49*	41.91
	4.00	*45.50*	42.23
cr	0.50	*34.84*	31.96
	1.00	*37.16*	33.84
	2.00	*39.01*	35.72
	3.00	*39.82*	37.63
	4.00	*40.36*	38.48
lena	0.50	*21.82*	17.76
	1.00	*25.92*	21.50
	2.00	*29.99*	26.38
	3.00	*31.62*	28.50
	4.00	*32.84*	29.83
peppers	0.50	*20.54*	17.04
	1.00	*25.64*	21.74
	2.00	*29.74*	26.79
	3.00	*31.50*	28.90
	4.00	*32.37*	29.85

shows part of the image approximation and correspond-

ing image-domain triangulation obtained with each of the proposed and ED methods. From this figure, we can see that the image approximation produced by the proposed method in Figure 2(a) is clearly of higher visual quality than the image obtained with the ED scheme in Figure 2(c). In particular, the proposed method better preserves detail, such as image edges, compared to the ED scheme. The reason for the better performance in the case of the proposed method can be seen by examining the image-domain triangulations in Figures 2(b) and (d). A careful examination of these triangulations shows that, in the case of the proposed method, the triangulation edges are better aligned with image edges (i.e., discontinuities in the image). This leads to improved mesh quality.

In terms of complexity, it is worthwhile to note that our proposed method has a relatively low computational cost. On the first author's very-modest eight-year-old Lenovo notebook computer, the proposed method typically requires only a few seconds of computation time for mesh generation. For example, to generate a mesh for the lena image with a sampling density of 2%, the proposed method requires 1.5 seconds on this old Lenovo machine. So, the use of DDTs in our proposed method does not lead to a mesh-generation scheme requiring computation times on the order of minutes (or tens of minutes), as is the case with some other DDT-based mesh-generation approaches. The amount of time required by our proposed method is, in fact, quite modest.
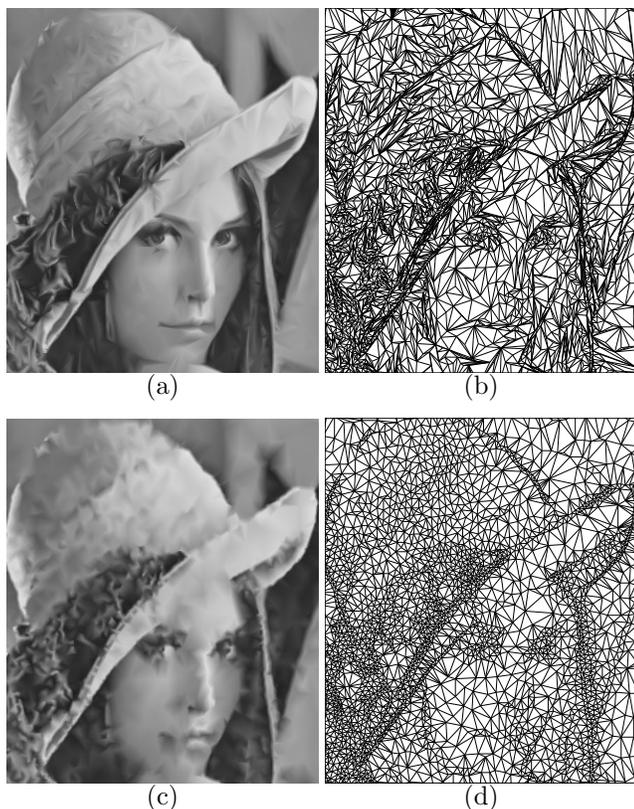
Figure 2: Part of the image approximation obtained for the lena image at a sampling density of 2% with the (a) proposed (29.99 dB) and (c) ED (26.38 dB) methods and (b) and (d) their corresponding triangulations.

5   Conclusions

In this paper, we proposed an improved method for generating mesh models of images, based on the ED scheme. Our proposed method makes use of DDTs in order to better exploit triangulation connectivity for improved approximation quality. Through experimental results, our proposed method was shown to produce image approximations of significantly higher quality than those obtained with the ED scheme, both in terms of PSNR (typically, by more than 1.5 dB) and visual quality. The improved approximation quality yielded by our proposed method comes at a relatively modest cost in terms of computation time. Thus, our proposed method is of great value to the many applications that require mesh models of images.

References

[1] JPEG-2000 test images. ISO/IEC JTC 1/SC 29/WG 1N 545, July 1997.

[2] Kodak lossless true color image suite. http://r0k.us/graphics/kodak, 2014.

[3] USC-SIPI image database. http://sipi.usc.edu/database, 2014.

[4] L. Alboul, G. Kloosterman, C. Traas, and R. van Damme. Best data-dependent triangulations. *Journal of Computational and Applied Mathematics*, 119:1–12, 2000.

[5] B. Delaunay. Sur la sphere vide. *Bulletin of the Academy of Sciences of the USSR, Classe des Sciences Mathematiques et Naturelle*, 7(6):793–800, 1934.

[6] C. Dyken and M. S. Floater. Preferred directions for resolving the non-uniqueness of Delaunay triangulations. *Computational Geometry—Theory and Applications*, 34:96–101, 2006.

[7] N. Dyn, D. Levin, and S. Rippa. Data dependent triangulations for piecewise liner interpolations. *IMA Journal of Numerical Analysis*, 10(1):137–154, 1990.

[8] R. W. Floyd and L. Steinberg. An adaptive algorithm for spatial greyscale. *Proc. of the Society for Information Display*, 17(2):75–77, 1976.

[9] M. Garland and P. S. Heckbert. Fast polygonal approximation of terrains and height fields. Technical Report CMU-CS-95-181, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA, Sept. 1995.

[10] C. L. Lawson. Software for C^1 surface interpolation. In J. R. Rice, editor, *Mathematical Software III*, pages 161–194. Academic Press, New York, NY, USA, 1977.

[11] P. Li and M. D. Adams. A tuned mesh-generation strategy for image representation based on data-dependent triangulation. *IEEE Trans. on Image Processing*, 22(5):2004–2018, May 2013.

[12] Y. Yang, M. N. Wernick, and J. G. Brankov. A fast approach for accurate content-adaptive mesh generation. *IEEE Trans. on Image Processing*, 12(8):866–881, Aug. 2003.