
Anisotropic Triangulation Methods in Adaptive Image Approximation

L. Demaret¹ and A. Iske²

¹ HelmholtzZentrum münchen, Neuherberg, Germany,
laurent.demaret@helmholtz-zentrum.de

² Department of Mathematics, University of Hamburg, D-20146 Hamburg,
Germany, iske@math.uni-hamburg.de

Summary. Anisotropic triangulations are utilized in recent methods for sparse representation and adaptive approximation of image data. This article first addresses selected computational aspects concerning image approximation on triangular meshes, before four recent image approximation algorithms, each relying on anisotropic triangulations, are discussed. The discussion includes generic triangulations obtained by simulated annealing, adaptive thinning on Delaunay triangulations, anisotropic geodesic triangulations, and greedy triangle bisections. Numerical examples are presented for illustration.

1 Introduction

This article surveys recent triangulation methods for adaptive approximations and sparse representations of images. The main purpose is to give an elementary insight into recently developed methods that were particularly designed for the construction of suitable triangulations adapted to the specific features of images, especially to their geometrical contents. In this context, the use of anisotropic triangulations appears to be a very productive paradigm. Their construction, however, leads to many interesting and open questions. To better understand the problems being addressed in current research, selected aspects concerning adaptive approximations on triangulations are discussed. In particular, relevant algorithmic and computational issues are addressed, where special emphasis is placed on the representation of geometrical information contained in images.

Finite element methods (FEM) are among the most popular classical techniques for the numerical solution of partial differential equations, enjoying a large variety of relevant applications in computational science and engineering. FEM are relying on a partitioning of the computational domain, where triangulations are commonly used. In fact, FEM on triangulations provide very flexible and efficient computational methods, which are easy to imple-

ment. Moreover, from a theoretical viewpoint, the convergence and stability properties of FEM on regular meshes are well understood.

Current research concerning different classes of mesh-based approximation methods, including FEM, is focused on the design of suitable *adaptive* meshes to improve their convergence and stability properties over previous methods. Adaptivity is, for instance, particularly relevant for the numerical simulation of multiscale phenomena in time-dependent nonlinear evolution processes. More specific examples are convection-diffusion processes with space-dependent diffusivity, or the simulation of shock front behaviour in hyperbolic problems. In this case, the construction of *well-shaped* adaptive triangulations with *good* anisotropic properties (by the shape and alignment of their locally adapted triangles) is a key issue for the methods' numerical stability (see e.g. [14, 23] for further details).

As regards the central topic of this survey, *adaptive image approximation by anisotropic triangulations*, the basic problem is

- the construction of anisotropic triangular meshes that are *locally adapted* to the geometrical contents of the input image data
in combination with
- the selection of a reconstruction method leading to an *optimal* or *nearby optimal* image approximation scheme.

Note that the above problem formulation is rather general and informal. In fact, to further discuss this, we essentially require a more formal definition of the terms “*locally adaptive*” and “*nearby optimal*”. Although a comprehensive discussion on the approximation theoretical background of this problem is far beyond the aims and scope of this survey, we shall be more specific on the relevant basics later in Section 2.

The image approximation viewpoint taken here is very similar to that of terrain modelling, as investigated in our previous work [7, 12]. In that particular application, greyscale images may be considered as elevation fields, where the terrain's surface is replaced by locally adapted surface patches (e.g. polynomials) over triangulations. But this is only one related application example. In a more general context, we are concerned with the approximation of (certain classes of) bivariate functions by using suitable models for locally adapted approximations that are piecewise defined over anisotropic triangulations.

The rich variety of contemporary applications for mesh-based image approximation methods lead to different requirements for the construction of the utilized (triangular) meshes. This has provided a diversity of approximation methods, which, however, are not yet gathered in a unified theory. Despite of this apparent variety of different methods, we observe that their construction relies on merely a few common principles, two of which are as follows.

One construction principle is based on a very simple idea: sharp edges between objects visible in an image correspond to crucial information. The abstract mathematical modelling leads to measures of regularity which take into

account singularities along curves. But triangulations are only one possible method for representing geometrical singularities of images. In fact, different methods were proposed for the rendering of contours in images. For a recent account of these methods, we refer to [13].

Another common principle requires the image representation to be sparse. In our particular situation, this means that the triangulation should be as small as possible. Note that the requirement of sparsity is not necessarily reflected by the number of vertices (or edges or triangles) in the triangulation, but it can also be characterized by some suitable entropy measure related to a compression scheme.

The construction of sparse triangulations for edge-adapted image approximation, according to the above construction principles, leads to an abstract approximation problem, whose general framework is briefly introduced in Section 2. In Section 3, we present four conceptually different approximation methods to solve the abstract approximation problem of Section 2.

This leads to four different image approximation algorithms, each of which achieves to combine the following desirable properties:

- good approximation behaviour;
- edge-preservation;
- efficient (sparse) image representation;
- small computational complexity.

Selected computational aspects concerning the implementation of the presented image approximation algorithms are discussed in Section 3. Supporting graphical illustrations and numerical simulations are presented in Section 3 and in Section 4. A short conclusion and directions for future work are finally provided in Section 5.

2 Image Approximation on Triangulations

2.1 Triangulations and Function Spaces

We first fix some notation and introduce basic definitions concerning triangulations and their associated function spaces.

Let $\Omega \subset \mathbb{R}^2$ denote a compact planar domain with polygonal boundary. Although we consider keeping this introduction more general, most of the following discussion assumes $\Omega = [0, 1]^2$ for the (continuous) computational image domain.

Definition 1. *A continuous image is a bounded and measurable function $f : \Omega \rightarrow [0, \infty)$, so that f lies in the L^∞ -class of measurable functions, i.e., $f \in L^\infty(\Omega)$.*

Although images are bounded (i.e., are lying in $L^\infty(\Omega) \subset L^p(\Omega)$), we distinguish between the different function spaces $L^p(\Omega)$, corresponding to different norms $\|\cdot\|_{L^p(\Omega)}$ for measuring the reconstruction error on Ω . We further remark that a natural image can always be represented by a bounded function. However, the converse is (trivially) not true: for any fixed $p \in [1, \infty]$, functions in $L^p(\Omega)$ do often *not* correspond to natural images. One of the main tasks of functional analysis methods in image processing is to define function classes, being given by some suitable regularity conditions, which are as small as possible but contain relevant images. In the context of triangulation methods, this immediately leads us to one central question: which image classes may be well-recovered by approximation methods relying on triangular meshes? Later in this section, we shall briefly discuss this question and give pointers to the relevant literature.

But let us first define triangulations and their associated function spaces.

Definition 2. A triangulation \mathcal{T} of the domain Ω is a finite set $\{T\}_{T \in \mathcal{T}}$ of closed triangles $T \in \mathbb{R}^2$ satisfying the following conditions.

(a) The union of the triangles in \mathcal{T} covers the domain Ω , i.e.,

$$\Omega = \bigcup_{T \in \mathcal{T}} T$$

(b) for any pair $T, T' \in \mathcal{T}$ of two distinct triangles, $T \neq T'$, the intersection of their interior is empty, i.e.,

$$\overset{\circ}{T} \cap \overset{\circ}{T'} = \emptyset \quad \text{for } T \neq T'.$$

We denote the set of triangulations of Ω by $\mathcal{T}(\Omega)$.

Note that condition (b) in the above definition disallows overlaps between different triangles in \mathcal{T} . Nevertheless, according to Definition 2, a triangulation \mathcal{T} may contain *hanging vertices*, where a hanging vertex is a vertex of a triangle in \mathcal{T} which is lying on the interior of an adjacent triangle's edge. Triangulations without hanging nodes are *conform*, which leads us to the following definition.

Definition 3. A triangulation \mathcal{T} of Ω is a **conforming triangulation** of Ω , if any pair of two distinct triangles in \mathcal{T} intersect at most at one common vertex or along one common edge. We denote the set of conforming triangulations of Ω by $\mathcal{T}_c(\Omega)$.

For the purpose of illustration, Figure 1 shows one non-conforming triangulation and one conforming triangulation.

On given planar domain Ω , there are many different (conforming) triangulations of Ω . We remark that many relevant applications rely on triangulations, where long and thin triangles need to be avoided, e.g. in finite element methods for the sake of numerical stability. In this case, *Delaunay triangulations* are a popular choice.




Fig. 1. (a) a non-conforming triangulation with three triangles and five vertices, where one vertex is a hanging vertex; (b) a conforming triangulation with six triangles and six vertices.

Definition 4. A Delaunay triangulation \mathcal{D} of Ω is a conforming triangulation of Ω , such that for any triangle in \mathcal{D} its circumcircle does not contain any vertex from \mathcal{D} in its interior.

We remark that the Delaunay triangulation \mathcal{D} of Ω with vertices $X \subset \Omega$ maximizes the minimal angle among all possible triangulations of Ω with vertices X . In this sense, Delaunay triangulations are *optimal* triangulations. Moreover, the Delaunay triangulation $\mathcal{D} \equiv \mathcal{D}(X)$ is unique among all triangulations with vertices X , provided that no four points in X are co-circular [20]. Figure 2 shows one Delaunay triangulation for the square domain $\Omega = [0, 1]^2$ with $|X| = 30$ vertices.




Fig. 2. Delaunay triangulation of a square domain with 30 vertices.

Now let us turn to functions on triangulations, i.e., bivariate functions $f : \Omega \rightarrow \mathbb{R}$ to be defined over a fixed triangulation \mathcal{T} . One possibility for doing so is by using piecewise polynomial functions. In this case, the restriction $f|_T$ to any triangle $T \in \mathcal{T}$ is a bivariate polynomial of a certain degree. To define piecewise polynomial functions on triangulations, their maximum degree is usually fixed beforehand and additional boundary or smoothness conditions are utilized. This then gives a finite dimensional linear function space.

In the situation of image approximation, we prefer to work with piecewise *linear* polynomial functions f . This is due to the simple representation of f . Moreover, we do not require any global smoothness conditions for f apart from global continuity. Since natural images are typically discontinuous, it also makes sense to refrain from assuming global continuity for f . This leads us to the following definition of two suitable function spaces of piecewise linear polynomials, one with requiring global continuity, the other without requiring global continuity. In this definition, \mathcal{P}_1 denotes the linear space of bivariate linear polynomials.

Definition 5. Let $\mathcal{T} \in \mathcal{T}(\Omega)$ be a triangulation of Ω . The set of **piecewise linear functions on \mathcal{T}** ,

$$\mathcal{S}_{\mathcal{T}} := \left\{ f : \Omega \rightarrow \mathbb{R} : f|_{\overset{\circ}{T}} \in \mathcal{P}_1 \right\}$$

is given by all functions whose restriction to the interior $\overset{\circ}{T}$ of any triangle $T \in \mathcal{T}$ is a linear polynomial.

Note that in the above definition, the restriction $f|_{\overset{\circ}{T}}$ of f may, for any individual triangle $T \in \mathcal{T}$, be extended from $\overset{\circ}{T}$ to T , in which case f may not be well-defined. For a conforming triangulation $\mathcal{T}(\Omega)$, however, f will be well-defined on Ω , if we require global continuity. In the following definition, $\mathcal{C}(\Omega)$ denotes the linear space of continuous functions on Ω .

Definition 6. Let $\mathcal{T} \in \mathcal{T}_c(\Omega)$ be a conforming triangulation of Ω . The set of **continuous piecewise linear functions on \mathcal{T}** ,

$$\mathcal{S}_{\mathcal{T}}^0 := \{ f \in \mathcal{C}(\Omega) : f|_T \in \mathcal{P}_1 \},$$

is given by all continuous functions on Ω whose restriction to any triangle $T \in \mathcal{T}$ is a linear polynomial.

Note that (for any conforming triangulation \mathcal{T}) $\mathcal{S}_{\mathcal{T}}^0$ is a linear subspace of $\mathcal{S}_{\mathcal{T}}$, i.e., $\mathcal{S}_{\mathcal{T}}^0 \subset \mathcal{S}_{\mathcal{T}}$. Moreover, note that $\mathcal{S}_{\mathcal{T}}^0$ has finite dimension. Indeed, the (finite) set of *Courant elements* $\varphi_v \in \mathcal{S}_{\mathcal{T}}^0$, for v in \mathcal{T} , each of which being uniquely defined by

$$\varphi_v(x) = \begin{cases} 1 & \text{for } x = v; \\ 0 & \text{for } x \neq v; \end{cases} \quad \text{for any vertex } x \text{ in } \mathcal{T},$$

is a basis of $\mathcal{S}_{\mathcal{T}}^0$. Therefore, the dimension of $\mathcal{S}_{\mathcal{T}}^0$ is equal to the number of vertices in \mathcal{T} . Similarly, the linear function space $\mathcal{S}_{\mathcal{T}}$ has dimension $3|\mathcal{T}|$, where $|\mathcal{T}|$ denotes the number of triangles in \mathcal{T} .

We remark that one important differences between the two approximation spaces, $\mathcal{S}_{\mathcal{T}}^0$ and $\mathcal{S}_{\mathcal{T}}$, is that $\mathcal{S}_{\mathcal{T}}^0$ requires conforming triangulations, whereas for $\mathcal{S}_{\mathcal{T}}$ the triangulation \mathcal{T} may contain hanging vertices. Note that this leads to different approximation schemes, as illustrated in Figure 3.

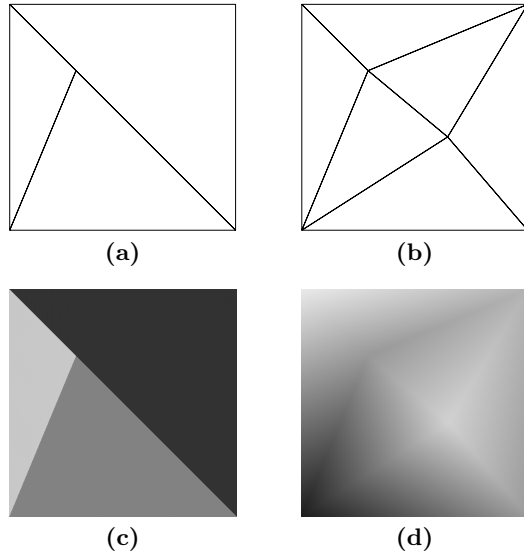


Fig. 3. (a) a non-conforming triangulation \mathcal{T} ; (b) a conforming triangulation \mathcal{T}' (cf. Figure 1); (c) a piecewise constant function $f \in \mathcal{S}_{\mathcal{T}}$; (d) a continuous piecewise linear function $g \in \mathcal{S}_{\mathcal{T}}^0$.

2.2 Isotropic and Anisotropic Approximation Methods

When approximating an image f by (best approximating) functions from $\mathcal{S}_{\mathcal{T}}^0$ or $\mathcal{S}_{\mathcal{T}}$, the resulting approximation quality heavily depends on the *quality* of the triangulation \mathcal{T} , in particular on the shape of the triangles in \mathcal{T} . Several alternative quality measures for triangulations are proposed in [23].

In this subsection, we discuss relevant principles concerning the construction of well-adapted triangulations, leading to adaptive image approximation methods. In this construction, it is essential to take (possible) anisotropic features of the image data into account. Second derivatives of the image (if they exist) are not necessarily of comparable amplitude (i.e., magnitude) in all directions. Furthermore, the ratio between the maximal and the minimal

amplitudes may vary quite significantly. Due to typical such heterogeneity (and other related multiscale phenomena) in image data, this gives rise to prefer *anisotropic* triangulations.

To be more specific, image approximation methods by triangulations are split into non-adaptive methods (relying on uniform triangulations) and adaptive methods (relying on non-uniform triangulations). Another distinction between image approximation methods (and their underlying triangulations) is by *isotropic* and *anisotropic* methods. For illustration, Figure 4 shows examples for uniform vs non-uniform and isotropic vs anisotropic triangulations.




Fig. 4. (a) a uniform triangulation; (b) a non-uniform isotropic (non-conforming) triangulation; (c) an anisotropic (conforming) triangulation.

- As regards characteristic features of different triangulation types, Figure 4,
- (a) triangles in a *uniform triangulation* have comparable sizes and shapes, so that the shape of each triangle is similar to that of an equilateral triangle.
 - (b) *non-uniform isotropic* triangulations comprise triangles of varying sizes, but the triangles have similar shapes.
 - (c) *anisotropic* triangulations comprise triangles of varying sizes and shapes.

We remark that non-adaptive image approximations (relying on uniform triangulations) are, in terms of their approximation quality, inferior to adaptive methods. In adaptive image approximation methods, however, the construction of their underlying non-uniform triangulations requires particular care. While non-uniform isotropic triangulations are suitable in situations, where the target image has point singularities at isolated points, anisotropic triangulations are particularly useful to locally adapt singularities of the image along curves, or other features of the image that may be reflected by a heterogeneous distribution of directional derivatives of first and second order. Therefore, anisotropic triangulations are usually preferred, especially when it comes to design suitable triangulations for adaptive image approximation methods.

2.3 Basic Techniques for Proving Error Estimates

Let us finally discuss available techniques for proving error estimates for image approximation. The following discussion includes classical results from finite elements as well as more recent results concerning Besov spaces and the related smoothness spaces, relying on the Mumford-Shah model. We keep the discussion of this section rather short, but we give pointers to the relevant literature. Further details are deferred to a follow-up paper.

The Bramble-Hilbert Lemma. Let us first recall classical error estimates from finite element methods. To obtain estimates for the global approximation error of a function by a piecewise function on a triangulation \mathcal{T} , standard analysis on finite element methods provides estimates for a single triangle $T \in \mathcal{T}$ (see [2]), where the key estimate is derived from the Bramble-Hilbert lemma [3]. For any image f from Sobolev space $W^{2,2}(T)$, the basic error estimate of the Bramble-Hilbert lemma reads as

$$\|f - \Pi_{\mathcal{S}_T} f\|_{L^2(T)} \leq |f|_{W^{2,2}(T)}, \quad \text{for } f \in W^{2,2}(T),$$

where $\Pi_{\mathcal{S}_T} f$ is the orthogonal L^2 -projection of f onto \mathcal{S}_T .

Slim and Skinny Besov Spaces. Although classical isotropic Besov spaces offer a more suitable framework for adaptive approximation schemes, they usually fail to represent approximation classes relying on anisotropic triangulations. Just recently, an more flexible concept was proposed to remedy this problem.

In [15], Karavainov and Petrushev introduced two different classes of anisotropic Besov spaces, *slim* and *skinny* Besov spaces. The construction of such Besov spaces relies on subdivision schemes leading to a family of nested triangulations. The approach taken in [15] is rather technical, but their main results can loosely be explained as follows.

The set of functions belonging to a slim Besov space (relying on a specific subdivision scheme) are the functions which can be approximated at a given convergence rate. Skinny Besov spaces are obtained by using similar construction principles for the special case of piecewise linear approximations. The quality of the resulting image approximation, however, heavily relies on the properties of the utilized subdivision scheme. The construction of suitable subdivision schemes remains a rather critical and challenging task. For further details, we refer to [15].

Bivariate Smoothness Spaces. Cartoon models lead to a large family of approximation methods which are based on the celebrated Mumford-Shah model [19]. These methods essentially take sharp edges of images into account, and their basic idea is to regard images as piecewise regular functions being separated by piecewise smooth curves. A generalization of the Mumford-Shah model has been proposed by Dekel, Leviatan & Sharir in [6]. In their work [6], smoothness spaces are defined by using a combination of two distinct notions

of smoothness, one for the inner pieces (away from the edge singularities), the other for the singularity supporting curves.

The corresponding smoothness spaces, \mathcal{B} -spaces, are in [6] defined in a similar way as by the standard interpolation techniques used for classical isotropic Besov spaces. In the present setting, we are interested in the use of smoothness spaces for the characterization of the approximation spaces

$$\mathcal{A}^\alpha := \left\{ f \in L^2(\Omega) : \inf_{|T|=n, T \in \mathcal{T}(\Omega)} \|f - \Pi_{S_T} f\|_{L^2(\Omega)} \leq \frac{C}{n^\alpha} \text{ for some } C > 0 \right\}.$$

Further in this context, we consider the smoothness space $\mathcal{B}_q^{\alpha, r_1, r_2}(L^p(\Omega))$ in [6, Definition 1.3] with $p = 2$, $r_1 = r_2 = 2$ and $q = \infty$. This combination of parameters is in [6] used for the special case of piecewise affine functions ($r_1 = 1$) on triangles ($r_2 = 2$) by measuring the error in $L^2(\Omega)$ ($p = 2$), and the approximation rates in the L^∞ -norm ($q = \infty$).

Useful error estimates for \mathcal{B} -spaces are proven in [6], where also a suitable characterization for the relevant approximation classes $\mathcal{A}^\alpha \equiv \mathcal{A}^\alpha(\Omega)$ is given. In fact, the main result in [6, Theorem 1.9] states that the inclusion $\mathcal{A}^\alpha \subset \mathcal{B}^\alpha$ holds for any planar domain $\Omega \subset \mathbb{R}^2$, and, conversely, we have the inclusion

$$\mathcal{B}^\alpha(\Omega) \cap L^\infty(\Omega) \subset \mathcal{A}^\alpha(\Omega),$$

which provides an almost sharp characterization of the approximation class \mathcal{A}^α . For further details, we refer to [6].

3 Four Algorithms for Adaptive Image Approximation

In this section we discuss four different algorithmic concepts for adaptive image approximation, which were proposed during the last five years. Each of the resulting approximation algorithms, to be discussed in Subsections 3.1-3.4, aims at the construction of suitable anisotropic triangulations to obtain sufficiently accurate image approximations. Moreover, their utilized adaptation strategies achieve a well-balanced trade-off between essential requirements concerning computational costs, approximation properties, and information redundancy. Computational details and key features of the five different concepts are explained in the following Subsections 3.1-3.4.

3.1 Generic Triangulations and Simulated Annealing

A naive approach to solve the image approximation problem of the previous Section 2 would consist in finding an *optimal* triangulation (among all possible triangulations of equal size) to obtain a *best* image approximation. However, the problem of finding such an optimal triangulation is clearly intractable. In fact, the set of all possible triangulations is huge, and so it would be far too costly to traverse all triangulations to locate an optimal one.

An alternative way for (approximately) solving this basic approximation problem is to traverse a smaller set of generic triangulations, according to a suitable set of traversing rules, to compute only a *suboptimal* triangulation but at much smaller computational complexity.

Recently, Lehner, Umlauf & Hamann [18] have introduced such a method for traversing a set of generic triangulations, where their basic algorithm relies on simulated annealing. The triangulations output by the method in [18] are very sparse, i.e., for a target approximation error, the output triangulation (whose resulting approximation error is below the given tolerance), requires only a small number of vertices (cf. the numerical comparison in Figure 5).




Fig. 5. (a) Triangulation and (b) image reconstruction obtained by simulated annealing [18]; (c) triangulation and (d) image reconstruction obtained by adaptive thinning [8] (Section 3.2). In either test example, the triangulation has 979 vertices.

The generic algorithm of [18] is iterative and can briefly be explained as follows. On given initial triangulation \mathcal{T}_0 , local modifications are performed, which yields a sequence $\{\mathcal{T}_n\}_n$ of triangulations. Any local modification on a

current triangulation \mathcal{T}_k is accomplished by using one of the following three basic operations (edge flip, local and global vertex move), yielding the subsequent triangulation \mathcal{T}_{k+1} .

- For two triangles sharing a common edge in a convex quadrilateral, an *edge flip* replaces the diagonal of the quadrilateral by the opposite diagonal;
- *local vertex move*: a vertex is moved to a new position in its neighbourhood;
- *global vertex move*: a vertex is moved and its cell is retriangulated.

To each of the three elementary operations, corresponding probabilities p_e, p_ℓ, p_g , satisfying

$$p_e + p_\ell + p_g = 1,$$

are assigned, according to which the next operation (edge flip, local or global vertex move) is performed. But the selection of edges to be flipped or vertices to be moved is by random.

To avoid local optima, the next triangulation is either taken or rejected, according another probability,

$$p_n = p(\Delta E_n, n),$$

where ΔE_n is the difference between approximation errors induced by the triangulations \mathcal{T}_n and \mathcal{T}_{n+1} . For further details concerning the probability measures we refer to [18].

The iteration of [18] was shown to be very flexible, but the computational efficiency is rather critical. This is due to the construction of the greedy algorithm, which navigates through a very large set of triangulations. Further improvements may be used to speed-up the convergence of the simulated annealing procedure, where the reduction of computational complexity is done by working with *local* probabilities [18].

Finally, the numerical example in Figure 5 **(a)**-**(b)** shows the efficacy of the simulated annealing method. Note that, especially from the viewpoint of anisotropy, the triangles in Figure 5 **(a)** are well-aligned with the sharp edges in the test image.

3.2 Adaptive Thinning Algorithms

Adaptive thinning algorithms [12] are a class of greedy point removal schemes for bivariate scattered data. In our recent work [8, 9, 10], adaptive thinning algorithms were applied to image data, and more recently, also to video data [11] to obtain an adaptive approximation scheme for images and videos. The resulting compression methods were shown to be competitive with JPEG2000 (for images) and MPEG4-H264 (for videos).

To explain the basic ideas of adaptive thinning for image data, let X denote the set of image pixels, whose corresponding pixel positions are lying on a rectangular two-dimensional grid. This, in combination with the pixels' luminance values defines a bivariate discrete function $f : X \rightarrow \mathbb{R}$. Now the aim

of the adaptive thinning algorithm is to select a small subset $Y \subset X$ of *significant pixels*, whose corresponding Delaunay triangulation $\mathcal{D} \equiv \mathcal{D}(Y)$ gives a suitable finite-dimensional ansatz space $\mathcal{S}_{\mathcal{D}}^0$ of globally continuous piecewise linear functions.

The image approximation $s : [X] \rightarrow \mathbb{R}$ to f is then given by the *best approximation* $s^* \in \mathcal{S}_{\mathcal{D}}^0$ in the least squares sense, i.e. s^* minimizes the ℓ^2 -error

$$\|s - f\|_2^2 := \sum_{x \in X} |s(x) - f(x)|^2$$

among all functions $s \in \mathcal{S}_{\mathcal{D}}^0$. Note that s^* is unique and can be computed efficiently by standard least squares approximation.

The challenge of this particular approximation method is to determine a *good* adaptive spline space $s \in \mathcal{S}_{\mathcal{D}}^0$, by the selection of a suitable subset $Y \subset X$, such that the resulting least squares error

$$\eta \equiv \eta(Y) = \|s^* - f\|_2^2$$

is as small as possible. Ideally, one wishes to compute an optimal $Y^* \subset X$ which minimizes $\eta(Y)$ among all subsets $Y \subset X$ of equal size. But the problem of computing Y^* is NP-complete. This requires greedy approximation algorithms for the selection of suitable (sub-optimal) solutions $Y \subset X$.

In greedy adaptive thinning algorithms, a subset $Y \subset X$ of significant pixels is computed by the recursive removal of pixel points, one after the other. The generic formulation of the adaptive thinning algorithm is as follows.

Algorithm 1

INPUT: *set of pixel positions X and luminance values $\{f(x)\}_{x \in X}$.*

- (1) Let $X_N = X$;
- (2) **FOR** $k = 1, \dots, N - n$
 - (2a) Find a least significant pixel $x \in X_{N-k+1}$;
 - (2b) Let $X_{N-k} = X_{N-k+1} \setminus x$;

OUTPUT: *subset $Y = X_n \subset X$ of significant pixels.*

To implement an adaptive thinning algorithm, it remains to give, for any $Y \subset X$, a definition for *least significant* pixels in Y . To this end, several different significance measures were proposed in [7, 8, 9, 10, 12]. Each of the utilized significance measures are relying on (an estimate) of the *anticipated error* that is incurred by the removal of the pixel point. The anticipated error for a pixel y is a local measure $\sigma(y)$ for the incurred ℓ^2 -error due to its removal. In the greedy implementation of adaptive thinning, a pixel y^* is least significant (in any step of the algorithm), whenever its anticipated error $\sigma(y^*)$ is smallest among all points in the current subset $Y \subset X$. Since $\sigma(y)$ can be computed and updated in constant time, this allows for an efficient implementation of adaptive thinning in only $\mathcal{O}(N \log(N))$ operations, where $N = |X|$.

3.3 Anisotropic Geodesic Triangulations

The anisotropic meshing problem, as pointed out in [1], can be interpreted as the search for a criterion based on a locally modified metric, according to which triangulations are then constructed. To connect this viewpoint with the concept of anisotropic triangulations, the Euclidean metric corresponds to a uniform triangulation, whereas metrics whose unit balls are disks of varying sizes lead to isotropic adaptive triangulations. Finally, anisotropic triangulations can be generated by using metrics whose unit balls are ellipses of varying sizes, shapes and directions. A suitable triangulation algorithm is then required to produce triangulations which are *aligned* with this modified metric, i.e., each triangle should be included in such an ellipse. Moreover, direction and ratio between the major and minor radii are directed by the local structure of data, while the size of the ellipse (i.e., the radius of the ball in the modified metric) is a parameter depending on the global target reconstruction quality.

The local metric is commonly defined by a positive definite tensor field, i.e., a mapping which associates to each point $x \in \Omega$ a symmetric positive definite tensor matrix $H(x) \in \mathbb{R}^{2 \times 2}$. For any point $x_0 \in \Omega$, the local metric $H(x_0)$ is then defined by the distance between a point $x \in \Omega$ and x_0 ,

$$\|x - x_0\|_{H(x_0)} = \sqrt{(x - x_0)^t H(x_0) (x - x_0)}.$$

Note that this concept enables us to define the length of a piecewise smooth curve $\gamma : [0, 1] \rightarrow \Omega$ w.r.t. metric H by

$$L_H(\gamma) = \int_0^1 \|\gamma'(t)\|_{H(\gamma(t))} dt$$

and so the *geodesic distance* between two points $x, y \in \Omega$ by

$$d_H(x, y) = \min_{\gamma} L_H(\gamma),$$

where the minimum is taken over all piecewise smooth curves joining x and y .

A quite natural choice for H seems to be the Hessian. One construction of anisotropic triangulations for image approximation, based on an anisotropic geodesic metric, has recently been proposed in [1]. Instead of taking the Hessian matrix as tensor structure matrix, a regularized version of the gradient tensor is used in [1]. Regularization is then performed by the convolution with a Gaussian kernel; the role of this regularization is to ensure a robust estimation in the presence of noise. Let us denote this regularized gradient tensor by $T(x)$ and assume that $T(x)$ can be diagonalized in an orthonormal basis, i.e., for a suitable basis we have (with λ_1, λ_2 depending on x):

$$T(x) = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}.$$

The geodesic metric is then defined as follows: the matrix T is slightly perturbed and then set to a power α , which is an ad hoc parameter controlling the anisotropy of the triangulation. This leads to

$$H(x) = \begin{bmatrix} (\lambda_1 + \varepsilon)^\alpha & 0 \\ 0 & (\lambda_2 + \varepsilon)^\alpha \end{bmatrix}.$$

Using this particular definition for a locally modified metric, the following difficult problem needs to be solved: determine a small as possible finite set of vertices $V \subset \Omega$ satisfying

$$\inf_{y \in V} d_H(x, y) \leq \delta \text{ for some } \delta > 0.$$

Here, δ is a global parameter which controls the reconstruction quality. The meshing of this set of points is then based on the use of the anisotropic Voronoi diagram, where the Voronoi cells are defined by the modified metric d_H rather than by the Euclidean one. Since the dual of a so obtained Voronoi diagram is not necessarily a Delaunay triangulation, some effort is necessary to construct a valid triangulation. This can be achieved by greedy insertion algorithms, where at each step the farthest point (w.r.t. the modified metric) to the current set of vertices is added to this set. This strategy is coupled with a suitable triangulation technique. In Figure 6, it is shown how this method works for a smooth image with steep gradients in an area around a regular curve. The result is an anisotropic triangulation. For more details, see [1, 16, 17].




Fig. 6. Geodesic triangulation with (a) 50, (b) 200, and (c) 800 vertices.

3.4 Greedy Triangle Bisections

In [4], Cohen, Dyn, Hecht & Mirebeau propose a greedy algorithm which is based on a very simple but effective rule for recursive subdivisions of triangles. The main operations in their method are *bisections* of triangles, where a bisection of a triangle T is given by a subdivision of T in two smaller triangles, obtained by the insertion of an edge which connects one vertex in T with the midpoint of the opposite edge. Therefore, for any triangle there are three possible bisections, with the (inserted) edges e_1, e_2 and e_3 , say. An example of the two first steps of such a recursive subdivision is shown in Figure 7. Note that this method produces non-conforming triangulations. Therefore, the approximation is performed w.r.t. the reconstruction space $\mathcal{S}_{\mathcal{T}}$.




Fig. 7. (a) Initial triangulation; (b),(c) triangulations by greedy bisection [4].

To derive a refinement algorithm from this bisection rule, a suitable criterion is required for selecting the triangle to be subdivided along with the bisection rule. The criterion proposed in [4] is straight forward: it takes one triangle with maximal approximation error, along with a bisection whose resulting approximation error is minimal. Therefore, an edge e^* corresponding to an optimal bisection of a triangle T is given by

$$e^* = \operatorname{argmin}_{e \in \{e_1, e_2, e_3\}} \left(\|f - \Pi_{\mathcal{S}_{T_1(e)}}\|_{L^2(T_1(e))}^2 + \|f - \Pi_{\mathcal{S}_{T_2(e)}}\|_{L^2(T_2(e))}^2 \right),$$

where $T_1(e)$ and $T_2(e)$ are the triangles resulting from the bisection of T by e .

This procedure outputs a sequence of refined triangulations, which we denote by $\mathcal{T}_{b,n}$ (b for bisection). In [5, Theorem 5.1], optimality of the bisection algorithm is proven for strictly convex functions. We can formulate the result in the relevant L^2 -setting as follows. For a strictly convex function $f \in \mathcal{C}^2(\Omega)$, there exists a constant $C > 0$ satisfying

$$\|f - \Pi_{\mathcal{S}_{\mathcal{T}_{b,n}}} f\|_{L^2(\Omega)} \leq \frac{C}{n} \|\sqrt{\det(D^2 f)}\|_{L^\tau(\Omega)}, \text{ with } \tau = \frac{2}{3},$$

where $\mathcal{T}_{b,n}$ is the sequence of triangulations produced by greedy bisection in combination with an L^1 -based selection criterion [5].

4 Numerical Simulations

In this final section, we wish to demonstrate the utility of anisotropic triangulation methods for image approximation. To this end, we apply adaptive thinning of Subsection 3.2 to four different images of size 512×512 , as shown in Figure 8: (a) one artificial image generated by a piecewise quadratic function, PQuad, and three natural images, (b) Aerial, (c) Game, and (d) Boats.

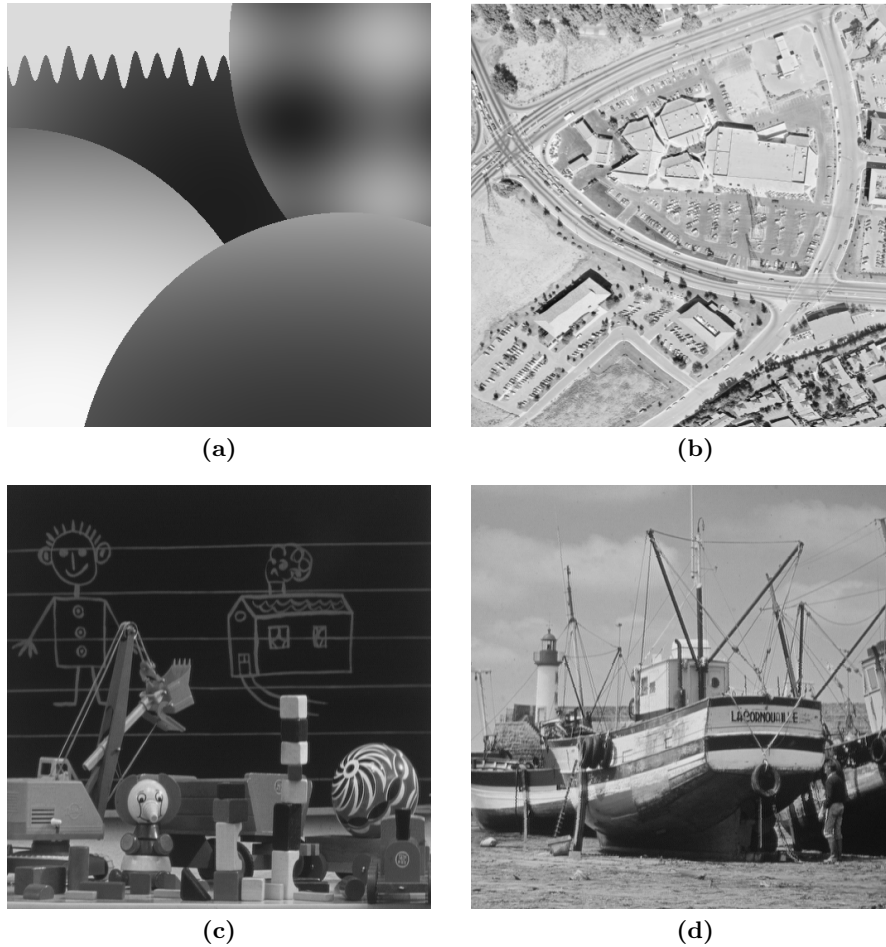


Fig. 8. Four images of size 512×512 : (a) PQuad; (b) Aerial; (c) Game; (d) Boats.

The Delaunay triangulations of the significant pixels, output by adaptive thinning, are shown in Figure 9. The quality of the image approximations, shown in Figure 10, is measured in dB (decibel) by the *peak signal to noise*

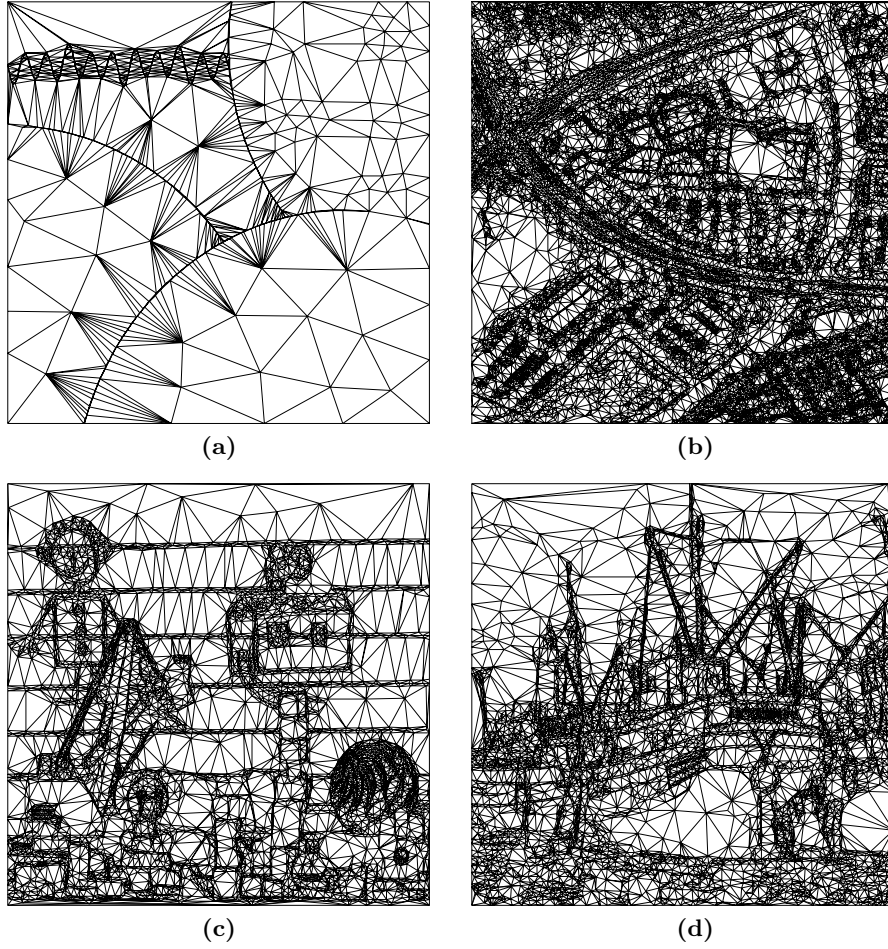


Fig. 9. Anisotropic Delaunay triangulations. (a) PQuad with 800 vertices (b) Aerial: 16000 vertices (c) Game: 6000 vertices (d) Boats: 7000 vertices.

ratio,

$$\text{PSNR} = 10 * \log_{10} \left(\frac{2^r \times 2^r}{\bar{\eta}^2(Y, X)} \right),$$

where the *mean square error* (MSE) is given by

$$\bar{\eta}^2(Y, X) = \frac{1}{|X|} \sum_{x \in X} |s(x) - f(x)|^2.$$

Note that for each test case the anisotropic triangulations achieve to capture the image geometry fairly well. This results in image approximations

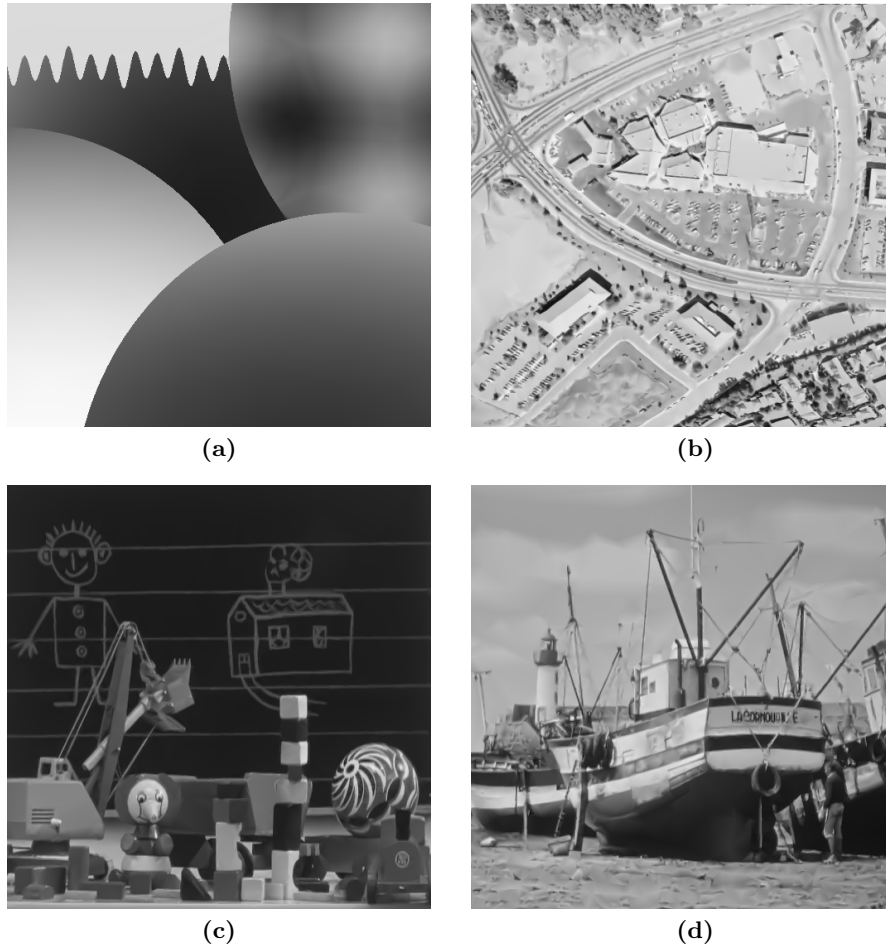


Fig. 10. Image approximation by adaptive thinning. (a) PQuad at PSNR 42.85 dB (b) Aerial: PSNR 30.33 dB (c) Game: PSNR 36.54 dB (d) Boats: PSNR 31.83 dB.

where the key features of the images (e.g. sharp edges) and finer details are recovered very well, at reasonable coding costs, as reflected the small number of significant pixels. In fact, the efficient distribution of sparse significant pixels is one key feature of adaptive thinning, which yields a very competitive compression method [8, 9, 11].

Note that the test case **PQuad** reflects the behaviour of adaptive thinning for an artificial cartoon image: inside on the smooth parts, the triangles are uniform and close to equilateral, whereas long and thin triangles are obtained along the discontinuities. As regards the two test cases **Game** and **Boats**, it is

shown how anisotropic triangulations help concentrate the representing triangles in the content-rich areas of these natural images.

Finally, the performance of adaptive thinning for the test case **Aerial** shows the potential of anisotropic triangulations yet once more, but now in a context where much more information are available. The anisotropy of the triangles vary according to the kind of feature they represent, and so the corresponding triangulation helps reproduce the geometrical properties of the underlying features very well, in particular the roads and buildings.

5 Final Remarks and Future Work

One of the practical issues of anisotropic meshing is related to the high computational costs induced by the search for nearby optimal approximating triangulations. In this article, we have considered methods which are too slow for applied fields where real-time computational costs are indispensable. Recently, very fast methods, coming from a slightly different world, closer to the preoccupations of engineering applications, have been developed [21, 22, 24]. These methods rely on some heuristic intended to find an adapted sampling set of pixels together with corresponding meshing and reconstruction techniques. In [22], a quite detailed comparison of these methods in terms of number of triangles versus quality and in terms of computational costs is provided, including comparisons with adaptive thinning. In comparison with the methods discussed in the present survey (in Subsections 3.1-3.4), the number of vertices required in the methods of [21, 22, 24] is much higher for a given target quality, but they allow for very fast implementations. One of the most challenging tasks for future research remains to bridge the gap between these rather pragmatic but highly efficient methods and the mathematically well-motivated but much slower methods of Subsections 3.1-3.4.

Acknowledgement

We gratefully thank Jean-Marie Mirebeau and Gabriel Peyré for several fruitful discussions concerning image approximation. The graphical illustrations in Subsection 3.1 were provided by Burkhard Lehner, those in Subsection 3.3 were provided by Gabriel Peyré.

References

1. S. Bougleux, G. Peyré, and L. Cohen: Image compression with anisotropic geodesic triangulations. Proceedings of ICCV'09, Oct. 2009.
2. D. Braess: *Finite Elements. Theory, Fast Solvers and Applications in Solid Mechanics*. 3rd edition, Cambridge University Press, Cambridge, UK, 2007.
3. J. Bramble and S. Hilbert: Estimation of linear functionals on Sobolev spaces with application to Fourier transforms and spline interpolation. *SIAM J. Numer. Anal.* **7**(1), 1970.
4. A. Cohen, N. Dyn, F. Hecht, and J.-M. Mirebeau: Adaptive multiresolution analysis based on anisotropic triangulations. Preprint.
5. A. Cohen and J.-M. Mirebeau: Greedy bisection generates optimally adapted triangulations. Preprint.
6. S. Dekel, D. Leviatan, and M. Sharir: On bivariate smoothness spaces associated with nonlinear approximation. *Constructive Approximation* **20**, 2004, 625–646.
7. L. Demaret, N. Dyn, M.S. Floater, and A. Iske: Adaptive thinning for terrain modelling and image compression. In: *Advances in Multiresolution for Geometric Modelling*, N.A. Dodgson, M.S. Floater, and M.A. Sabin (eds.), Springer-Verlag, Berlin, 2004, 321–340.
8. L. Demaret, N. Dyn, and A. Iske: Image compression by linear splines over adaptive triangulations. *Signal Processing* **86**(7), July 2006, 1604–1616.
9. L. Demaret and A. Iske: Adaptive image approximation by linear splines over locally optimal Delaunay triangulations. *IEEE Signal Processing Letters* **13**(5), May 2006, 281–284.
10. L. Demaret and A. Iske: Scattered data coding in digital image compression. In: *Curve and Surface Fitting: Saint-Malo 2002*, A. Cohen, J.-L. Merrien, and L.L. Schumaker (eds.), Nashboro Press, Brentwood, 2003, 107–117.
11. L. Demaret, A. Iske, and W. Khachabi: Sparse representation of video data by adaptive tetrahedralisations. In: *Locally Adaptive Filters in Signal and Image Processing*, L. Florack, R. Duits, G. Jongbloed, M.-C. van Lieshout, and L. Davies (eds.) Springer, Dordrecht, 2010, 197–220.
12. N. Dyn, M.S. Floater, and A. Iske: Adaptive thinning for bivariate scattered data. *Journal of Computational and Applied Mathematics* **145**(2), 2002, 505–517.
13. H. Führ, L. Demaret, and F. Friedrich: Beyond wavelets: new image representation paradigms. In: *Document and Image Compression*, M. Barni and F. Bartolini (eds.), 2006, 179–206.
14. P.-L. George and H. Borouchaki: *Delaunay Triangulations and Meshing: Application to Finite Elements*. Hermes, Paris, 1998.
15. B. Karavainov and P. Petrushev: Nonlinear piecewise polynomial approximation beyond Besov spaces. *Appl. Comput. Harmon. Anal.* **15**(3), 2003, 177–223.
16. F. Labelle and J. Shewchuk: Anisotropic Voronoi diagrams and guaranteed-quality anisotropic mesh generation. Proc. 19th Annual Symp. on Computational Geometry, ACM, 2003, 191–200.
17. G. Leibon and D. Letscher: Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. Proc. 16th Annual Symp. on Computational Geometry, ACM, 2009, 341–349.
18. B. Lehner, G. Umlauf, and B. Hamann: Image compression using data-dependent triangulations. In: *Advances in Visual Computing*, G. Bebis et al. (eds.), Springer, LNCS **4841**, 2007, 351–362.

19. D. Mumford and J. Shah: Optimal approximation of piecewise smooth functions and associated variational problems. *Comm. in Pure and Appl. Math.* **43**, 1989, 577–685.
20. F.P. Preparata and M.I. Shamos: *Computational Geometry*. Springer, New York, 1988.
21. G. Ramponi and S. Carrato: An adaptive sampling algorithm and its application to image coding. *Image and Vision Computing* **19**(7), 2001, 451–460.
22. M. Sarkis and K. Diepold: Content adaptive mesh representation of images using binary space partitions. *IEEE Transactions on Image Processing* **18**(5), 2009, 1069–1079.
23. J. Shewchuk: What is a good linear finite element? Interpolation, conditioning, anisotropy and quality measures. Preprint, December 2002.
24. Y. Yang, M. Wernick, and J. Brankov: A fast approach for accurate content-adaptive mesh generation. *IEEE Transaction on Image Processing* **12**(8), 2003, 866–880.