

Multiscale Approximation of Piecewise Smooth Two-Dimensional Functions using Normal Triangulated Meshes

Maarten Jansen ^{a,*} Richard Baraniuk, ^b Sridhar Lavu ^b

^a*Department of Mathematics and Computer Science, Technical University of Eindhoven, PO Box 513 NL-5600 MB, Eindhoven, The Netherlands, and Department of Computer Science K.U.Leuven, Belgium.*

^b*Department of Electrical and Computer Engineering, Rice University, 6100 Main Street, Houston, TX 77005, USA. Supported by NSF, ONR, AFOSR, DARPA, and the Texas Instruments Leadership University Program.*

Abstract

Multiresolution triangulation meshes are widely used in computer graphics for representing three-dimensional (3-d) shapes. We propose to use these tools to represent 2-d piecewise smooth functions such as grayscale images, because triangles have potential to more efficiently approximate the discontinuities between the smooth pieces than other standard tools like wavelets. We show that *normal mesh subdivision* is an efficient triangulation, thanks to its local *adaptivity* to the discontinuities. Indeed, we prove that, within a certain function class, the normal mesh representation has an optimal asymptotic error decay rate as the number of terms in the representation grows. This function class is the so-called *horizon class* comprising constant regions separated by smooth discontinuities, where the line of discontinuity is C^2 continuous. This optimal decay rate is possible because normal meshes automatically generate a polyline (piecewise linear) approximation of each discontinuity, unlike the blocky piecewise constant approximation of tensor product wavelets. In this way, the proposed nonlinear multiscale normal mesh decomposition is an anisotropic representation of the 2-d function. The same idea of anisotropic representations lies at the basis of decompositions such as wedgelet and curvelet transforms, but the proposed normal mesh approach has a unique construction.

Key words: normal offsets, mesh, image, multiresolution, wavelet, approximation
AMS 2000 MSC: 41A46, 41A63, 42C40

* Corresponding author

Email addresses: `mjansen@win.tue.nl` (Maarten Jansen), `richb@rice.edu` (Richard Baraniuk), `lavu@rice.edu` (Sridhar Lavu).

1 Introduction: images with long smooth edges

This paper concerns the representation and approximation of piecewise smooth, two-dimensional (2-d) functions, which consist of smooth regions delineated by step discontinuities along smooth one-dimensional (1-d) contours, which we call *edges*. Many different types of real-world data can be modeled as piecewise smooth. As an important example, a piecewise smooth function is a quite accurate model for a grayscale *image*, which represents the light intensity of a black-and-white visual scene. While we will use images as our central, running example in this paper, other examples abound in statistics and differential equations for a broad spectrum of applications.

By *approximation*, we mean approximating a piecewise smooth function with a finite dimensional representation. Immediate applications of approximation results include compression and noise removal (denoising).

For images and many other kinds of data, an approximation is typically defined on a discrete set of points on some grid. For example, digital images are typically acquired by sampling the light intensity at discrete points on a square grid of pixels (currently using a CCD array), and so image representations and processing algorithms typically operate on this square grid. The square pixel grid is nearly always assumed to be fixed, with the dependent variable of the image the pixel intensity. While the acquisition and processing of image data on a square grid of pixels is simple, it turns out to be very inefficient for representing many important image features, including the *edges*.

Edges are the dominating features in piecewise smooth 2-d functions. Edges contain two types of information: *where* the edge is located, i.e., its location and geometry, and *what* is the step value, i.e., the height of the discontinuity. In 2-d, geometry information plays a crucial role, much more than in 1-d. In 1-d piecewise smooth functions, discontinuities occur at isolated points, and these can be easily captured in a wavelet transform. In 2-d, edge singularities lie along 1-d contours, which are much harder to capture.

The time-scale analysis of the wavelet representation provides a powerful tool for approximating a 1-d function f . Thanks to the local support of the basis functions, under mild conditions, a nonlinear wavelet approximation f_n containing the n largest terms of the wavelet expansion of f performs as well on a piecewise smooth f as on a smooth f [19,20,8,7,13]. Indeed, the L_2 approximation error decays rapidly with increasing n :

$$\|f - f_n^{1-d \text{ wavelet}}\| = \mathcal{O}(n^{-\nu}). \quad (1)$$

In this equation, ν stands for

$$\nu = \min(\tilde{p}, \alpha),$$

with \tilde{p} the number of (dual) vanishing moments of the wavelet analysis and α the Lipschitz regularity of the signal at its non-singular points. Wavelets provide a very efficient representation of 1-d




Fig. 1. Haar wavelet approximation of a 2-d piecewise constant image featuring a smooth 1-d edge singularity. Each shaded square corresponds to the support of a wavelet basis function. At each finer scale, an increasing number of wavelets is necessary to cover (and hence represent) the singularity. This effect does not exist in 1-d and explains the suboptimal performance of wavelets for representing 2-d piecewise smooth functions. Another drawback of tensor product wavelets is that they approximate the edge curve as a piecewise constant. This explains the “blockiness” of wavelet image approximations.

piecewise smooth signals primarily because in 1-d the geometry information consists of merely a few isolated points.

Wavelets are thus well-suited for estimating a piecewise smooth 1-d function in the presence of noise. In the minimax sense, the performance of a simple n -term approximation algorithm comes within a neglectable logarithmic factor of the best possible method involving a piecewise polynomial with knots at the (assumed known) positions of the singularities [12].

Unfortunately, this approximation power does not carry over into two and higher dimensions. Indeed, standard tensor-product wavelet transforms based on a square grid of 2-d sampling points are ill-prepared to represent edges, since many wavelets overlap with the 1-d edge, leading to a preponderance of geometry information (see Fig. 1).

Given a 2-d function f that is smooth except for an edge singularity along a smooth (say C^2) curve, the nonlinear wavelet approximation f_n using the n largest wavelet terms has an L_2 error rate

$$\|f - f_n^{2\text{-d wavelet}}\| = \mathcal{O}(n^{-1/2}).$$

This outperforms a Fourier procedure, where about the best we can do is a linear approximation taking the *first* n Fourier coefficients

$$\|f - f_n^{2\text{-d Fourier}}\| = \mathcal{O}(n^{-1/4}).$$

Nevertheless, neither of these procedures comes close to the 1-d rate of (1). This is partly due to an inherent dimensionality effect: approximation of 2-d data is inevitably more difficult than 1-d data. Yet, wavelets do not obtain the optimal 2-d rate either. They approximate a curved singularity as a piecewise constant. This observation explains the blocky output of wavelet image approximations.

In order to achieve better approximation rates on 2-d edge contours, new, edge-adaptive, multiscale decompositions have been developed in recent years. Due to the numerous possible orientations, lengths and curvatures of edges, it is impossible to catch all possible edges by a basis decomposition. The new multiscale decompositions therefore may proceed in overcomplete representations (frames), such as contourlets [9] or curvelets [2]. Another type of edge-adaptive decompositions is by a best basis selection within a overcomplete collection (called library or dictionary) of candidate bases. Examples of such constructions are wedgelets [10,23], platelets [27], bandelets [18], beamlets [11] and others [6,24,16].

The idea followed in this paper is to treat images as special cases of 3-d surfaces and represent them using *triangular* image patches [22,15]. A triangulation consists of *triangles*, that is, triplets of *vertices* connected by *edges* (not to be confused with image edges). Because the triangles edges can be placed in arbitrary locations and orientations, triangles have the potential to represent arbitrary edge contours (the geometry information) more accurately with a fewer number of patches than a fixed square grid representation. The key is to use an *adaptive triangulation* that places vertices more densely in edge regions for accurate and efficient edge representation, yielding a parsimonious image representation (see Fig. 2).

Indeed, an adaptive triangle-based decomposition can provide a piecewise linear edge approximation, provided that the triangulation adapts itself to the precise locations of the edges. Ideally, this could lead to an error rate of

$$\|f - f_n^{2-d \text{ adapt. tri.}}\| = \mathcal{O}(n^{-1}), \quad (2)$$

For efficient processing of 3-d mesh data, *multiscale* triangulation based on nonlinear *subdivision* has been proposed in computer graphics. Multiscale mesh construction starts from a small number of coarse-scale points on the surface. Finer triangular meshes are formed by subdividing, that is, by gradually adding more data points (vertices, pixels). Unlike the standard subdivision scheme that places new vertices at the midpoints of the triangle edges, we can adapt the location of the new vertices based on local geometry information. The *normal mesh* scheme selects the new points based on the local normal direction computed from the previous coarser scale mesh [17].

Originally developed for efficient 3-d surface representation, we will show that the normal mesh representation shows remarkable adaptivity to the edge structure of 2-d piecewise smooth functions (see Figure 2(c)). Starting with an arbitrary set of initial vertices, we will demonstrate that the normal mesh subdivision algorithm rapidly places more and more vertices directly on the edge contour, enabling a direct representation of the location and geometry information of edges as well as the information on (the height of) the singularity itself. Adaptivity and better approximation with triangles are the keys to the success of normal meshes as opposed to wavelets.

In this paper, we propose a multiscale normal mesh representation for piecewise smooth 2-d functions such as images. We will show that for the idealized *horizon class* of images that are piecewise constant save for a C^2 edge discontinuity we obtain the optimal error decay rate of (2).

This asymptotic error decay is the same as the one obtained with wedgelets [10] and a log-factor

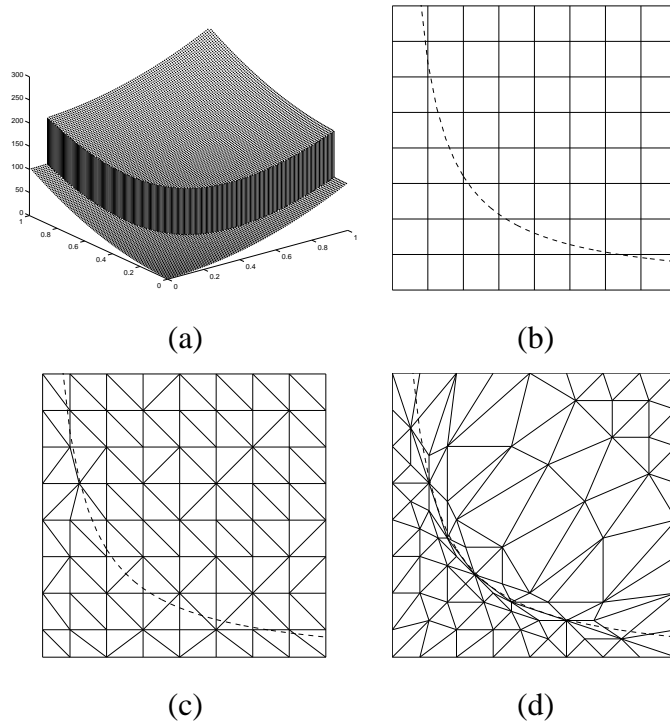


Fig. 2. Representing the 2-d function $f(x, y) = x^2 + y^2 + I_{\{y > 0.1/x\}}$. (a) 3-d mesh plot of the function. (b) A square grid representation such as that used by a tensor-product wavelet transform. The dashed line is the hyperbolic edge in $f(x, y)$. (c) A non-adaptive triangular refinement has the potential of a better approximation of the edges, but it does not exploit this potential: the triangle sides do not constitute an interpolating polyline (i.e., piecewise linear) approximation of the edge. The vertices do not lie on the edge. (d) The combination of triangulation and adaptivity does the job.

faster than curvelet approximation [3]. The normal mesh decomposition of an image of n pixels has linear, i.e., $\mathcal{O}(n)$, computational complexity, compared to $\mathcal{O}(n \log n)$ for a wedgelet transform. Curvelets and wedgelets are inherently overcomplete representations and had first been developed for image *enhancement*, rather than compression. Later, they were adopted in compression algorithms as well [23]. On the other hand, the normal mesh approach proposed in this paper, is specifically oriented towards image compression: the geometry of an edge and its singularity are both stored into the *normal offsets*, i.e., the coefficients of the decomposition. The decomposition requires no additional constructions — such as edgelets in a wedgelet transform — to take care of the location and orientation of edges. A second unique feature of the proposed decomposition is its embedding into the *lifting scheme*. Although beyond the scope of this paper, the proposed scheme can be incorporated in a natural way into an adaptive lifting scheme which applies wavelet steps in textured areas and uses normal meshes in areas with long, smooth edges. Such an adaptive scheme is straightforward and requires no iterative implementation. It is subject of current research. The scheme can also be extended to faster approximate edges whose curves are more than twice differentiable.

The idea of data-adaptive triangulation has been elaborated under different assumptions in [14]. In that paper, the triangulation is constructed based on an *existing* set of vertices. These vertices coincide with samples of the function and remain fixed. The normal offset method on the other




Fig. 3. *The triangular subdivision principle: four triangles at a finer scale are generated by subdividing each triangle at the next coarser scale.*

hand first finds the best locations for the vertices before triangulating them. These locations are chosen for optimal approximation.

Data-adaptive meshes (content-based triangulations) are also popular in video compression, see for instance [26,1]. Apart from the fact that data-adaptive meshes in video are not generated with normal offsets, there are at least two conceptual differences with the normal mesh method proposed in this paper. First, the content-based triangulations in video processing are obviously driven by motion analysis. Second, the triangulation in this paper cannot stand on its own, it is not an independent mesh generation but only a by-product in the multiscale (wavelet-like) decomposition.

This paper is organized as follows. In Section 2 we overview background on triangulations for general 3-d surfaces, present the normal mesh concept, and specialize the construction to 2-d piecewise smooth functions. The rest of the paper conducts a detailed performance analysis for horizon class images in two steps. In Step I (Sections 3 and 4), we analyze the normal “mesh” (polyline) approximation in 1-d. In Step II (Section 6), we leverage this analysis into the 2-d horizon class case. Section 7 presents some practical results on synthetic and real images. Section 8 offers a discussion and conclusions.

2 Multiscale Image Triangulations

2.1 Quadtree triangulations

Consider the construction of a multiscale triangulated function representation using the principle of subdivision (vertex refinement). In standard subdivision for 2-d functions, we introduce vertices for the next finer scale at the midpoint of the existing triangle legs. This results in four child triangles that, in the function domain, cover the same area as their parent (see Fig. 3).

Once we begin and fix the rules for refinement, we need only specify the initial vertex points; finer scale vertices are uniquely defined without extra information. This is valuable in applications such as compression, where a parsimonious image representation is required.

To build a wavelet transform for a multiscale triangulated image representation, we separate the vertices from the finest scale mesh into two groups: those from the previous coarser scale (group

A) and those obtained by the subdivision of these points (group B). With this “decimation” of points, we can apply the lifting scheme [25] to implement a wavelet transform on the triangular grid.

For each point in group B , we define a proper neighborhood window around it, and we predict the value at that point using the points from group A in the neighborhood. The simplest algorithm merely applies simple linear prediction [25]. By subtracting the predicted value from the actual value, we obtain the wavelet coefficient. The scaling coefficients are obtained by updating the points in group A by adding update values computed from the wavelet coefficients.

2.2 Normal subdivision in 1-d

The salient concepts of normal subdivision are easily described in a simple 1-d example. In a standard 1-d wavelet transform, a wavelet coefficient is computed as the vertical offset between a sample value and its prediction based on its neighboring samples. The length of the dashed, vertical line in Fig. 4(a) is such a wavelet coefficient; it tells how far the function value in the middle point deviates from a linear interpolation of its two neighbors at the next coarser scale. Instead of linear interpolation, more sophisticated predictions could be used and more coarse scale neighbors could be involved.

The normal subdivision in Fig. 4(b) is very similar, except that it computes its offset in a direction *normal* to the current, coarse scale prediction. Thus, the detail coefficient tells us how far to go, not just vertically, but in a specific *direction*, supplying geometry information. This normal direction obviously depends on the coarse scale prediction, which makes the procedure nonlinear.

The normal subdivision scheme in 1-d proceeds as follows:

- (1) Choose the initial, coarsest scale data points on the curve and connect them into a piecewise linear approximation, i.e., an interpolating polyline approximation of the curve.
- (2) Repeat the following steps for each pair of successive points:
 - (a) Compute the midpoint of the polyline segment connecting the two successive points. This is the prediction for the new point. Compute the direction of the vector normal to the polyline segment in this midpoint.
 - (b) Find the point on the function curve that is pierced by the normal vector. The new point is defined by the coordinates of this piercing point.
 - (c) Record the displacement between the midpoint prediction and the piercing point as the wavelet coefficient for the new point.

2.3 Normal subdivision in 2-d

We can treat a 2-d piecewise smooth function as a surface in 3-d space and build a normal mesh for it. Step edges in the 2-d function now become vertical 3-d surfaces. The basic steps to building




Fig. 4. 1-d wavelet coefficient computation in (a) the classical lifting scheme and (vertical offsets) and (b) the normal mesh scheme (normal offsets).

a normal mesh of a 2-d function run as follows:

- (1) Choose the initial, coarsest scale vertices and connect them into a triangular mesh. In the simplest case, the initial vertices can be the set of four corner points of the function (assuming it has a finite domain).
- (2) Repeat the following steps for each edge of the triangulation until the approximation converges:
 - (a) Compute the midpoint of the edge. This is the prediction for the new vertex. Compute the direction of the vector normal to the surface that fitting the immediate neighbors of the edge.
 - (b) Find the point on the function surface that is pierced by the normal vector. (In practice, with a sampled 2-d function, the surface is defined locally by fitting the given sample points.) The new vertex is defined by the coordinates of this piercing point.
 - (c) Record the displacement between the midpoint prediction and the piercing point as the wavelet coefficient for the new vertex.
 - (d) Using the new vertices, retriangulate each triangle to obtain the finer mesh through subdivision.

A stopping rule in this refinement process depends upon the application. In digital image processing, for instance, refinement stops as soon as all pixels have been inserted. Note that the refinement generates irregular grids, so it may be necessary to continue in some areas after convergence in other areas.

In some exceptional cases, we cannot use the normal piercing point and we need to go back to the usual midpoint subdivision scheme. This is further elaborated in Section 6. A similar exception handling occurs in a normal mesh decomposition of surfaces [17].

As illustrated in Fig. 2(d), this procedure builds a normal mesh with remarkable adaptivity to the locations of the singularities in the function. Even with an arbitrary choice of initial vertices, the normal mesh algorithm almost immediately starts placing new vertices close to the singularities, making the edges of the triangles align with the function contours. As the triangles refine, the normal mesh provides a successive piecewise linear approximation of each contour, yielding an




Fig. 5. (a) Normal mesh approach at a discontinuity, compared to (b) the classical subdivision scheme. In the normal mesh case, the approximation f_{j+1} at scale $j + 1$ has a point on the exact location of the singularity, because the normal offset points to this singularity.

efficient representation of 2-d singularities with a small number of normal mesh detail coefficients.

The mechanism behind this interesting behavior is explained by considering a simple 1-d step function example, as illustrated in Fig. 5(a). *The normal direction tends to point to the singularities*; if two vertices at a coarse scale j lie on either side of a singularity, then the normal piercing point is always closer to the singularity than the standard midpoint subdivision point. And once the inter-vertex distance r_j becomes smaller than the singularity height Δ , then the newly inserted point is always on the singularity itself, see Section 3.1.

We will now analyze the performance of 2-d nonlinear approximation based on normal mesh triangulations for horizon class images. We proceed in two steps. Step I (Sections 3 and 4) analyzes the performance of normal “meshes” (polylines) at representing 1-d step discontinuities. Step II (Section 6) then extends this analysis to take into account the continuous curve of 1-d step discontinuities present in a horizon class image.

3 Analysis I: Normal Polylines for 1-d Piecewise Constant Functions

This section analyses the behavior of a normal polyline approximation of a piecewise constant function. It turns out that the convergence rate depends on the initial geometry, i.e., the exact location of the singularity. We derive both an average rate, in Theorem 1, and a minimum convergence rate, in Theorem 3. Since the convergence rate depends in a chaotic way on the initial geometry, the analysis of the minimum convergence rate requires some additional lemmas.

3.1 Rapid localization of singularity position

Suppose we have a step function $f(x) = \Delta \cdot 1_{\{x > x_0\}}$ with a discontinuity at $x_0 \in (0, 1)$ and let $\Delta = f(x_{0+}) - f(x_{0-})$ be the *height* of the discontinuity. We wish to construct a multiscale




Fig. 6. Evolution of the approximation near singularities through successive scales.

approximation f_j of this function, where j denotes the scale (or resolution level). The normal mesh subdivision scheme generates an irregular grid \mathbf{x}_j near this singularity. In this analysis, we consider normal offsets with respect to a prediction by polylines, i.e., piecewise linear functions which are continuous in the knots. We call r_j the uncertainty about the location (i.e., geometry) of the jump, i.e., the width of the sample interval $[x_{j,i}, x_{j,i+1}]$, containing the singularity x_0 : $r_j = x_{j,i+1} - x_{j,i}$ and $x_{j,i} \leq x_0 \leq x_{j,i+1}$. By construction in Figure 5 (top), it holds that $r_j \leq r_{j-1}/2$, so

$$r_j = \mathcal{O}(2^{-j}).$$

If r_j becomes smaller than Δ , then for certain $r_{j+1} = 0$, i.e., the next subdivision point lies on the singularity. This can be seen by sketching the situation where $\Delta = r_j$ and $x_{j,i+1} = x_0 + \delta$, with δ arbitrarily small. This situation is the least favorable for a given width r_j , and yet, the next piercing point lies on the edge. This singularity locating property becomes crucial in 2-d, as we discuss later.

3.2 Expected behaviour near singularities

After having located the singularity position, the algorithm breaks down into two independent approximations: the behaviour left of the singularity is independent of, though similar to the behaviour on the right, so we study (the left) one side only, as in Figure 6. For further reference, we let the first subdivision point on the singularity correspond to level $j = 0$.

The procedure gradually reduces the error near the singularity, which essentially happens in two possible ways, depending on the slope of the present approximation, i.e., the angle α_j in the picture. If $\alpha_j < \pi/4$, as in the figure, the normal direction pierces the step function in its flat (horizontal) region, thereby reducing the interval on which the approximation differs from the true function. We call S_j the width of this interval: it is the support of the (left side of) the approximation error. On the other hand, if $\alpha_j > \pi/4$, the normal direction finds a new point on the singularity, and so it reduces the *height* H_j of the error function. H_j is also the L_∞ norm of the (left side) error function.

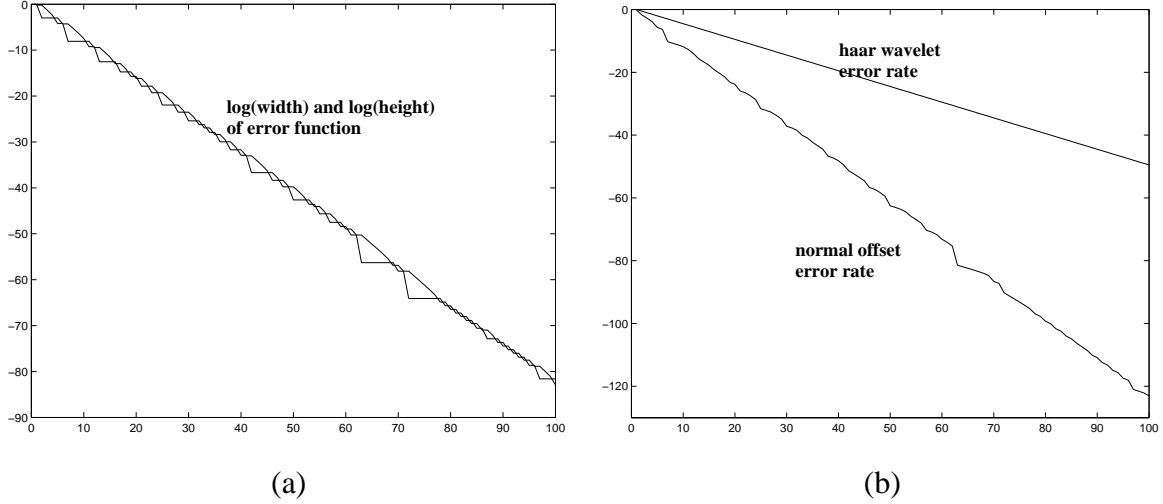


Fig. 7. (a) Evolution of the logarithm of the error support and the error L_∞ -norm for a normal mesh approximation of a step function. (b) Log of the L_2 -norm of error function as function of n , the number of detail coefficients.

It holds that:

- (1) If $\alpha_j < \pi/4$ then
 - $S_{j+1} = \frac{S_j}{2} (1 - \tan^2 \alpha_j) = S_j \frac{\tan \alpha_j}{\tan 2\alpha_j}$
 - $H_{j+1} = H_j$
 - $\alpha_{j+1} = 2\alpha_j$
- (2) If $\alpha_j > \pi/4$ then
 - $S_{j+1} = S_j$
 - $H_{j+1} = \frac{H_j}{2} (1 - \cot^2 \alpha_j)$
 - $\alpha_{j+1} = 2\alpha_j - \frac{\pi}{2}$

In other words, if $\alpha_j > \pi/4$, the support width of the error function remains unchanged, but if it changes, i.e., if $\alpha_j < \pi/4$, it is reduced by a factor of more than a half. This means that the convergence is not monotone, and the precise process depends on the initial angle α_0 . Figure 7(a) plots $\log_2 S_n$ and $\log_2 H_n$ and Figure 7(b) compares the logarithmic error $\log_2 \|\epsilon_n\|_2 = \log_2 \sqrt{S_n H_n^2/3}$ with that of Haar wavelets ($\log_2 \|\epsilon_n\| = \mathcal{O}(-n/2)$). Note that the number of coefficients n equals the scale j . We could state all results in terms of n , but we chose to use j as subscript wherever a result or argument is based on thinking in scales.

It is interesting to analyse the behaviour of a wavelet approximation in this framework. Figure 8 shows that the support of the error function is divided by two in every step, but the height (i.e., the difference between maximum and minimum of the error curve) remains a constant. This leads to a L_2 -error rate of $\mathcal{O}(2^{-n/2})$.

Theorem 1 *If the initial angle α_0 is uniformly distributed on $[0, \pi/2]$, then all subsequent α_j are*

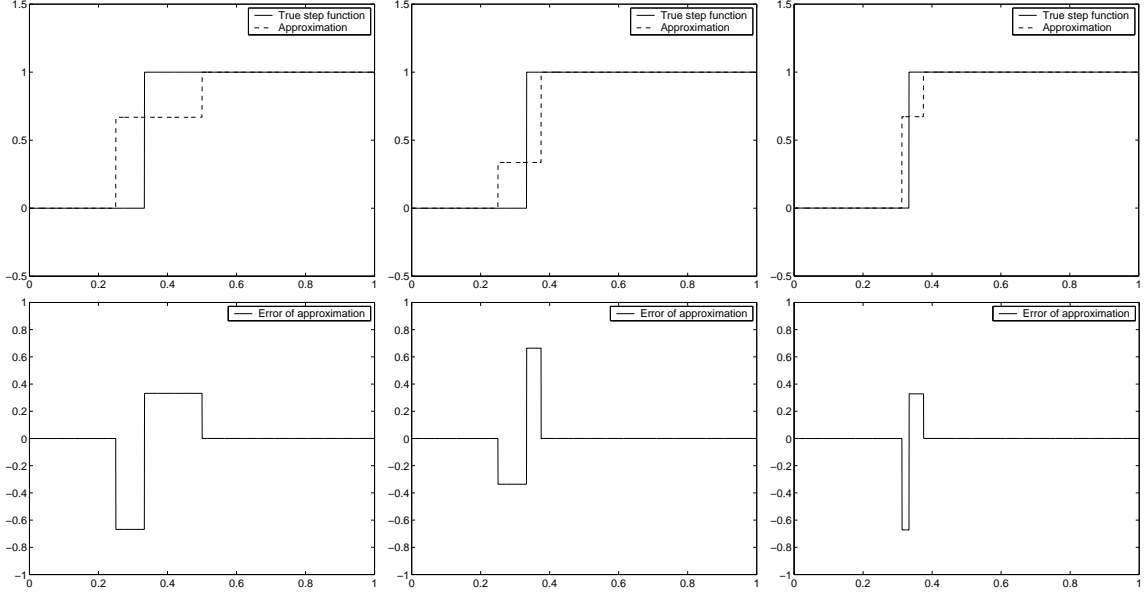


Fig. 8. The evolution of the error of a Haar approximation for a step function on $[0, 1]$ with step at $1/3$. The 3 plots on top compare the approximation with the true function at successive scales. The plots below show the corresponding error function. The error support width is reduced by a factor 2 in every step, but the error height remains a constant.

uniformly distributed, and the expected logarithmic error reduction is:

$$\gamma := \mathbb{E} \log_2 \frac{\|\epsilon_{n+1}\|}{\|\epsilon_n\|} = -\frac{3}{2} - \frac{6}{\pi \log 2} \int_0^{\pi/4} \log \cos \alpha \, d\alpha$$

Proof:

$$\begin{aligned} \gamma &:= \mathbb{E} \log_2 \frac{\|\epsilon_{n+1}\|}{\|\epsilon_n\|} = \mathbb{E} \log_2 \sqrt{\frac{S_{n+1} H_{n+1}^2}{S_n H_n^2}} \\ &= \frac{1}{2} \frac{1}{\pi/2} \left(\int_0^{\pi/4} \log_2 \left(\frac{1 - \tan^2 \alpha}{2} \right) d\alpha + \int_{\pi/4}^{\pi/2} \log_2 \left(\frac{1 - \cot^2 \alpha}{2} \right)^2 d\alpha \right) \\ &= \frac{1}{\pi} \left(3 \int_0^{\pi/4} \log_2 \left(\frac{\cos^2 \alpha - \sin^2 \alpha}{\cos^2 \alpha} \right) d\alpha + 3 \int_0^{\pi/4} \log_2 \frac{1}{2} d\alpha \right) \\ &= \frac{3}{\pi} \left(-\frac{\pi}{4} \log_2 2 + \int_0^{\pi/4} \log_2 \frac{\cos 2\alpha}{\cos^2 \alpha} d\alpha \right) \\ &= \frac{3}{\pi} \left(-\frac{\pi}{4} + \frac{1}{2 \log 2} \int_0^{\pi/2} \log \cos \alpha \, d\alpha - \frac{2}{\log 2} \int_0^{\pi/4} \log \cos \alpha \, d\alpha \right) \end{aligned}$$

$$\begin{aligned}
&= \frac{3}{\pi} \left(-\frac{\pi}{4} - \frac{\pi}{4} - \frac{2}{\log 2} \int_0^{\pi/4} \log \cos \alpha \, d\alpha \right) \\
&\approx -\frac{3}{2} + 0.24 = -1.26
\end{aligned}$$

We used the fact that $\cot \alpha = \tan(\pi/2 - \alpha)$, to reduce all integrals to the interval $[0, \pi/4]$ and we filled in the known integral

$$\int_0^{\pi/2} \log \cos \alpha \, d\alpha = \pi \log 2/2.$$

□

Note that we could make this analysis, because the uniform density is invariant under the refinement scheme for α_j . Starting from any other density with bounded derivative, the successive density functions rapidly converge to a constant:

Lemma 2 *If f_n is the density function of the angle α_n after n refinement steps, and if the initial density f_0 has a bounded derivative on $[0, \pi/2]$, then*

$$\lim_{n \rightarrow \infty} f_n(\alpha) = 2/\pi.$$

Proof: Let

$$A = \max_{\alpha \in [0, \pi/2]} |f_0'(\alpha)|,$$

and note that

$$f_n(\alpha) = \frac{1}{2} f_{n-1} \left(\frac{\alpha}{2} \right) + \frac{1}{2} f_{n-1} \left(\frac{\alpha}{2} + \frac{\pi}{4} \right).$$

It then follows immediately that all f_n are continuous and differentiable and that

$$|f_n'(\alpha)| \leq 2^{-n} A.$$

From the middle value theorem, we then have that for any $x, y \in [0, \pi/2]$, $|f_n(x) - f_n(y)| \leq 2^{-n} A |x - y| \leq 2^{-n-1} A \pi$. From this, it follows immediately that the limiting function $f(x) = \lim_{n \rightarrow \infty} f_n(x)$ satisfies $f(x) - f(y) = 0$ for any x and y . Using the fact that for all n ,

$$\int_0^{\pi/2} f_n(x) dx = 1,$$

we find that the constant is indeed $2/\pi$. □

3.3 Ergodicity

The question arises whether the *average* error reduction per step, i.e., the mean reduction taken over a sequence of scales, converges to the expected error reduction in Theorem 1. Figure 7(a) seems to confirm this for an arbitrary initial α_0 .

It is interesting though to observe that

$$\alpha_j = \text{frac}\left(\frac{\alpha_0}{\pi/2}2^j\right)\frac{\pi}{2},$$

where $\text{frac}(x)$ is the fractional part of a real number x . If we call

$$\beta_j = \frac{\alpha_j}{\pi/2}, \tag{3}$$

then we see that in binary notation β_j is the fractional part of the number obtained by shifting the point in β_0 over j positions to the right. Then it is clear that

$$\beta_0 = .010011000111000011110000011111000000111111\dots$$

induces a sequence with two accumulation points (0 and 1) and that the α_j do not fill up the interval $[0, \pi/2]$ uniformly. This holds even more for $\beta_0 \in \mathbb{Q}$ in which case only a limited subset of possible angles show up ever again. $\alpha_0 = \pi/4$ leads to immediate convergence. Another example is $\beta_0 = .010101010\dots$ corresponding to $\alpha_0 = \pi/6$. In this case α_j alternates between $\alpha_{2j} = \pi/6$ and $\alpha_{2j+1} = \pi/3$, and it is easy to check that the convergence rate for the L_2 -norm is then $\log_2(3\sqrt{3})/2 \approx 1.19$, which is a bit slower than the average. The next section shows that this is the slowest possible convergence.

There is an immediate and interesting connection with the notion of “normal numbers” in number theory [21, Chapter8]. A normal number in base b is a number where any configuration of n digits (for any choice of n) appears with frequency $1/b^n$. If β_0 is normal in the 2-base, this means that any sequence of n bits in its binary notation, appears with frequency $1/2^n$. The β_j following from the refinement scheme on a normal β_0 are uniformly distributed. Indeed, as mentioned before, the uniform density is invariant under the refinement map. In measure theory, it is said that the refinement $T(\beta) = \text{frac}(2\beta) := 2\beta - [2\beta]$ on the unit interval with Lebesgue measure (i.e., uniform distribution) is an *endomorphism*. In this definition $\text{frac}(\cdot)$ denotes the fractional part of a real number and $[\cdot]$ stands for the integer part. It can be proven that this endomorphism is *ergodic*, i.e., $T^{-1}(A) = A$ is only possible if $P(A) = 0$ or $P(A) = 1$. Birkhoff’s ergodic theorem then guarantees that the average of any (measurable) function $f(\beta)$ over the observed

values $\beta_j = T(\beta_{j-1})$ converges and that the limit equals a.s. the expectation of this function with respect to (in our case) the Lebesgue measure, i.e., the uniform density. If we take for $f(\beta)$ the characteristic function on an interval $[0, B]$, for arbitrary $0 \leq B \leq 1$, we see that the frequency of $\beta_j \leq B$ equals B , which is the cumulative of the uniform density. This proves that almost all initialisations β_0 induce a uniformly distributed sequence. The set of such initial β_0 equals the set of normal numbers in $[0, 1]$.

3.4 Minimum convergence rate

Theorem 3 *The error ϵ_j of a normal polyline approximation of a singular point in a piecewise constant function as in Figure 6 satisfies:*

$$\bar{\gamma} := \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=1}^n \log_2 \frac{\|\epsilon_j\|_2}{\|\epsilon_{j-1}\|_2} \leq -\log_2(3\sqrt{3})/2 \approx -1.19.$$

From the analysis in Section 3.2, it is clear that the function $\bar{\gamma}(\alpha_0)$ shows a sort of fractal behaviour and simple optimisation techniques do not apply. In order to prove the theorem, we specifically show that $\alpha_0 = \pi/6$ leads to the slowest convergence rate. To do so, we start with some quick observations:

Lemma 4 *If the sequence α_j is generated as defined in Section 3.2, then $\forall j \in \mathbb{Z}$*

$$\bar{\gamma}(\alpha_j) = \bar{\gamma}(\alpha_0).$$

In particular: the asymptotic behaviour for $\alpha_0 = \beta_0\pi/2$ is the same as that for $\alpha_1 = \beta_1\pi/2$, where β_1 is the fraction obtained by shifting the binary representation of β_0 up to the first 0-1-switch.

So, in $.0000001 \dots \pi/2$, the first zeros can be omitted for asymptotic analysis. This is a trivial asymptotic argument. It limits the search for the maximum value of $\bar{\gamma}(\alpha_0)$ to the interval $(\pi/8, 3\pi/8)$.

Lemma 5 *Measured in L_1 -norm, the error decay starting from initial angle α_0 is exactly the same as the error behaviour for $\pi/2 - \alpha_0$.*

This is clear from Figure 9. Unlike wavelets, the normal polyline approximation does not treat a signal as a function, and therefore it does not see jumps as discontinuities. Horizontal and vertical line segments are treated equally.

In L_1 , investigating β_0 is equivalent to investigating $1 - \beta_0$. In other words, switching all zeros and ones in β_0 's binary representation does not change the L_1 behaviour.

Lemma 6 *Minimising the asymptotic L_p decay is equivalent to minimising the asymptotic L_1 behaviour.*




Fig. 9. Normal polyline approximation near a singularity. The approximation is determined by the angle α_j , but α_j and $\pi/2 - \alpha_j$ correspond to the same L_1 error.

This is because $\|\epsilon_n\|_p^p = S_n H_n^p$, and, once more, the normal polyline approximation does not make a difference between vertical and horizontal sections of a piecewise constant signal. As a consequence, $H_n = \mathcal{O}(S_n)$, from which the Lemma follows. Lemmata 5 and 6 further restrict the maximisation to the interval $(\pi/4, 3\pi/8)$.

Figure 10 plots the two and three steps log-average decay of the L_1 -error as a function of the initial angle α_0 . In general, we call

$$\bar{\eta}_n(\alpha_0) := \frac{1}{n} \sum_{j=1}^n \log_2 \frac{\|\epsilon_j\|_1}{\|\epsilon_{j-1}\|_1}$$

the n steps average decay of the L_1 -error. Note that $\eta_n(\pi/6) = -\log_2 3 = \eta_n(\pi/3)$, independent of n . According to Lemma 6, we prove the theorem for L_2 if we show that this value is the maximum of

$$\bar{\eta}(\alpha_0) := \lim_{n \rightarrow \infty} \bar{\eta}_n(\alpha_0).$$

This $\bar{\eta}(\alpha_0)$ is the L_1 equivalent of the objective function $\bar{\gamma}(\alpha_0)$.

The singular points of $\bar{\eta}_n(\alpha_0)$ are $k \cdot \frac{(\pi/2)}{2^n}$, with $k = 0, \dots, 2^n - 1$. The lobes in between are characterised by the first n digits in the binary representation of $\beta_0 = \alpha_0/(\pi/2)$, as indicated in Figure 10. As n increases, the plot of $\bar{\eta}_n(\alpha_0)$ soon becomes chaotic.

$\alpha_0 = \pi/3$ clearly is not a maximum of $\bar{\eta}_2$ or $\bar{\eta}_3$. Especially starting angles near 0 or $\pi/2$ seem to show a much slower convergence than that of $\alpha_0 = \pi/3$. This is however just an initial behaviour. For the asymptotic analysis, it suffices to analyse within the interval $(\pi/4, 3\pi/8)$, but even there, $\pi/3$ is not a maximum. Even within its own lobe, there are angles that converge more slowly in the first three steps.

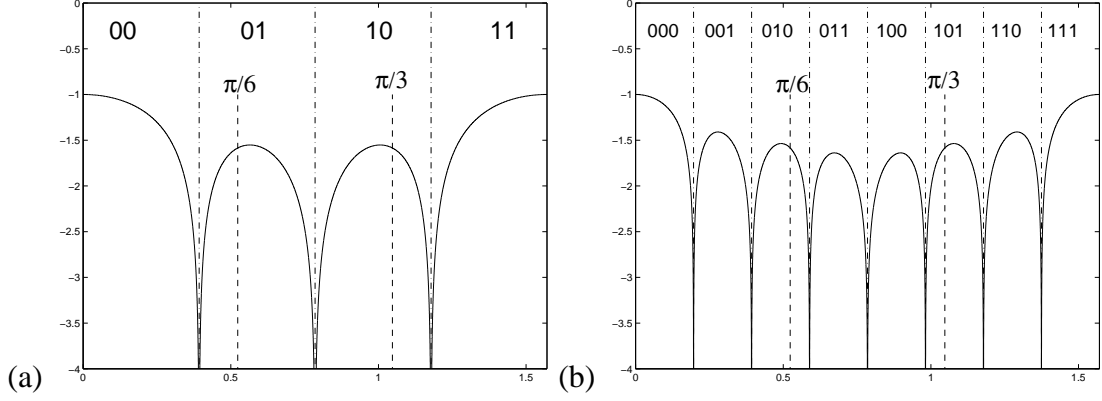


Fig. 10. (a) Mean two steps log-average decay of the L_1 -error as a function of the initial angle α_0 . (b) Mean three steps log-average decay of the L_1 -error as a function of the initial angle α_0 . The binary figures correspond to the first two, resp. three digits in the binary representation of $\beta_0 = \alpha_0/(\pi/2)$.

Nevertheless, we can prove the following lemma:

Lemma 7 Let a_n and b_n be the end points of the lobe containing $\pi/3$ in the n -th subdivision step, i.e., $\lim_{\alpha \rightarrow a_n} \overline{\eta}_n(\alpha) = -\infty = \lim_{\alpha \rightarrow b_n} \overline{\eta}_n(\alpha)$ and $\overline{\eta}_n(\alpha) \in \mathbb{R}, \forall a_n < \alpha < b_n$. Let $m_n = (a_n + b_n)/2$ be the middle point of this lobe and $\widehat{\eta}_n(\alpha) = \overline{\eta}_n(2m_n - \alpha)$ the mirror of $\overline{\eta}_n$ around this center point. Then:

$$\forall \alpha \in (a_n, m_n) : \overline{\eta}_n(\alpha) > \widehat{\eta}_n(\alpha) \text{ iff } a_n < \pi/3 < m_n.$$

In other words, every lobe that contains $\pi/3$ is asymmetric around its center, and the side that contains $\pi/3$ is always higher than its mirror. Moreover, two mirroring angles α_n and $\widehat{\alpha}_n$ within the same lobe in step n satisfy:

$$\widehat{\alpha}_n = \frac{\pi}{2} - \alpha_n,$$

and according to Lemma 5, the subsequent subdivision steps lead to exactly the same L_1 error decay for both α_n and its mirror.

As a consequence, if for instance

$$\alpha_0 = .1010\mathbf{100}10010101 \dots \frac{\pi}{2},$$

then one can always find an initial angle with slower convergence, namely its mirror in the lobe 101010:

$$\alpha_0^{[1]} = .1010\mathbf{101}01101010 \dots \frac{\pi}{2}.$$

As soon as two consecutive digits in the binary notation of β_0 are equal, flip all subsequent digits,




Fig. 11. A normal polyline does not make any distinction between piecewise constant functions (left) and continuous piecewise linear functions with rectangles in the singular points (right). As a matter of fact, the notion of function is not that important at all.

and the resulting angle will converge slower. Repeating this procedure on any arbitrary α_0 leads to $\alpha_0 = \pi/3$ being the slowest possible initial configuration. So, if we can prove Lemma 7, this reasoning concludes the proof of Theorem 3.

A similar situation appears in [4, Sections 7.1.1, 7.1.2] in a totally different context. Our problem, however, does not satisfy the assumptions in Lemmata 7.1.6 and 7.1.7 of [4].

The proof of Lemma 7 can be found in an appendix.

4 Normal polylines for 1-d piecewise linear and piecewise smooth functions

Unlike the wavelet approach, a normal polyline approximation does not use basis functions. It even does not rely on any concept of function whatsoever. A jump in a piecewise constant function leads to the same (L_1) error decay as a rotated version of this function, which is a continuous, piecewise linear with a rectangle in the singular points, as in Figure 11.

The question arises how well normal polylines do the job for general piecewise linear functions. Figure 12 analyses the performance in a singular point with angle θ (sometimes referred to as a *cusp*). Suppose that the normal subdivision point reduces S_j and not H_j , as in the Figure. It holds that $\alpha_{j+1} = \pi - 2(\pi/2 - \alpha_j) = 2\alpha_j$, and since α_{j+1} and θ are the angles of a triangle, this imposes the condition that $\alpha_j < (\pi - \theta)/2$. Otherwise, the new point reduces H_j and not S_j . From Figure 12, we see that

$$\tan \alpha_{j+1}(S_{j+1} - \cos \theta H_j) = \tan \alpha_j(S_j - \cos \theta H_j),$$

and

$$\frac{H_j}{\sin \alpha_j} = \frac{S_j}{\sin(\pi - \theta - \alpha_j)}.$$




Fig. 12. Multiscale normal polyline approximation of piecewise linear function with angle θ . This function (upper left) can be “laid down” to analyse the L_1 error reduction in the j -th step.

This leads to:

$$S_{j+1} = \left(\frac{\cos \theta \sin \alpha}{\cos \theta \sin \alpha + \cos \alpha \sin \theta} + \frac{\cos \alpha \sin \theta}{\cos \theta \sin \alpha + \cos \alpha \sin \theta} \frac{\tan \alpha}{\tan 2\alpha} \right) S_j.$$

If $\alpha_j > (\pi - \theta)/2$, we have that $\alpha_{j+1} = 2\alpha_j - (\pi - \theta)$.

Just as in the rectangular case, we can write α_j as a fraction of the maximal angle, $\alpha_j = \beta_j(\pi - \theta)$. If $\beta_0 = 1/3$ or $\beta_0 = 2/3$, all subsequent β_j alternate between these two values. If $\beta_j = 1/3$, i.e., $\alpha_j = (\pi - \theta)/3$, we have $2\alpha_j = \pi - \theta - \alpha_j$, and so, from the expressions above, we find

$$\frac{H_j}{\sin \alpha_j} = \frac{S_j}{\sin(2\alpha_j)} \Rightarrow \frac{H_j}{S_j} = \frac{1}{2 \cos(\alpha_j)} = \frac{1}{2 \cos\left(\frac{\pi-\theta}{3}\right)}$$

Moreover, $S_{j+1}/S_j = (H_j/S_j)^2$, so the corresponding L_1 -norm reduction in this j -th and all other steps equals:

$$\frac{S_{j+1}}{S_j} = \frac{1}{4 \cos^2\left(\frac{\pi-\theta}{3}\right)}.$$

For small θ , this error rate can come arbitrarily close to 1. At $\theta = \pi/4$, normal polylines catch a singularity at the same rate as wavelets. But even for smaller θ , the approximation of a singularity requires only one coefficient at each scale. This is due to the locality property, which this approach shares with the classical wavelets. As a consequence, the effort for the approximation of isolated singularities vanished when compared to approximating a superposed smooth function.

5 A few words on normal polyline approximation of smooth functions

The convergence rate of a normal polyline approximation for smooth functions (i.e., functions with a certain number of derivatives) has been extensively investigated [5]. For further reference, we list a few results which are special cases of that general analysis.

Lemma 8 *If a function f has a bounded derivative on the interval $[a, b]$, then the subdivision points $x_{j,k}$ at scale j and location k found by normal offsets on linear prediction, starting from a (strictly) increasing sequence $\{x_{0,k}\}_k$, satisfy:*

$$\begin{aligned} |x_{j+1,2k+1} - x_{j+1,2k}| &\leq \delta \cdot |x_{j,k+1} - x_{j,k}| \\ |x_{j+1,2k+2} - x_{j+1,2k+1}| &\leq \delta \cdot |x_{j,k+1} - x_{j,k}|, \end{aligned}$$

for some $0 < \delta < 1$, only dependent on f , not on j .

This means that

$$\lim_{j \rightarrow \infty} \sup |x_{j,k+1} - x_{j,k}| \rightarrow 0,$$

i.e., the normal offsets, though not so regular as classical subdivision, leaves no ‘gaps’, or unrefined subintervals.

The proof relies on a geometrical construction in Figure 13. Referring to this construction, Let $\mathbf{m} = (x_m, y_m) = ((x_{j,k} + x_{j,k+1})/2, (f_{j,k} + f_{j,k+1})/2)$ be the prediction point and $M = \sup |f'(x)|$. Then construct two lines with negative and positive slope M , meeting in point \mathbf{b} . By construction, this polyline is a majorant for f . As a consequence, the intersection $\mathbf{c} = (x_c, y_c)$ of this polyline with the normal direction satisfies: $|x_c - x_m| > |x_{j+1,2k+1} - x_m|$. A further upperbound can be found by replacing the true normal direction with the lowest possible slope, i.e., the normal on the polyline in \mathbf{b} . The intersection $\mathbf{n} = (x_n, y_n)$ satisfies $|x_n - x_m| > |x_{j+1,2k+1} - x_m|$. For finite M , x_n cannot be arbitrarily close to either $x_{j,k}$ or $x_{j,k+1}$, so $|x_{j+1,2k+1} - x_m| \leq \delta' |x_{j,k} - x_m|$ with $\delta' < 1$. Now it follows that $|x_{j+1,2k+1} - x_{j,k}| \leq \delta |x_{j,k+1} - x_{j,k}|$ with $\delta = (\delta' + 1)/2$.

Lemma 9 *Suppose a function f has a bounded second derivative on the interval $[0, h]$, and fix $f(0) = 0$, and let $x(h)$ be the x -coordinate of the normal subdivision point, then*

$$\lim_{h \rightarrow 0} \frac{x(h)}{h} = \frac{1}{2} = \lim_{h \rightarrow 0} \frac{h - x(h)}{h}$$

This means that normal refinement generates asymptotically regular subdivision grids.

Proof: Let $E(x) = f(x) - f(h) \cdot x/h$, be the approximation error function. Given that $E(0) = 0 = E(h)$ and that $|E''(x)| \leq M$, with $M = \sup |f''(x)|$, it follows that $|E(x)|$ is bounded by the quadratic function with zeros in 0 and h and second derivative equal to M . As a consequence,




Fig. 13. Proof of Lemma 8, see the text of this proof for explanation of the symbols.




Fig. 14. At fine scales, normal subdivision creates nearly regular grids.

$|E(x)| \leq Mh^2/8$. This holds for all $x \in [0, h]$, in particular for x the subdivision point, as illustrated in Figure 14. By Pythagoras' theorem, we have for $\epsilon := x - h/2$ that $|\epsilon| < |E| \leq Mh^2/8$. So,

$$\lim_{h \rightarrow 0} \frac{x}{h} = \lim_{h \rightarrow 0} \frac{h/2 + \epsilon}{h} = 1/2.$$

□

Corollary 10 *If f has a bounded second derivative on the interval $[a, b]$, then the approximation error of a normal polyline f_j after j refinement steps satisfies:*

$$|f_j(x) - f(x)| \leq C2^{-2j},$$

for some C dependent on f but independent of j .

This means that for certain normal subdivision has the same approximation rate as classical, regular subdivision. This follows from Lemmata 8 and 9: first, we are sure that the length of all subinter-




Fig. 15. Although asymptotically normal polylines on smooth data perform as well as regular subdivision on the x -coordinate, initial convergence of this smooth function may be as slow as the worst case analysis of the piecewise linear function in Figure 12.

vals converge to zero, and because they do so, the second lemma guarantees that all subintervals have length of order 2^{-j} . The result then follows from the fact that linear interpolation converges quadratically if f has a bounded second derivative.

Remark Although the asymptotic convergence of normal polylines shows the same rate as the convergence for regular subdivision on the x -coordinate, the initial convergence could be substantially slower. This happens if a smooth function is well approximated by a sharp cusp as in Figure 15. The initial convergence behaves as described in the worst case analysis of Section 4. From practical point of view, it is crucial to find initial points on the ‘sharpest’ cusps in the function. In a sense, these points with highest second derivative bear the essential information of the smooth curve. In the following sections, we consider smooth horizons in 2-d, i.e., smooth line singularities. The analysis in this section has illustrated that placing initial points near highly curved parts of such a line singularity is an interesting choice.

6 Analysis II: Normal Meshes for 2-d Piecewise Constant Functions

6.1 Horizon class images

This section analyses the adaptivity of the normal offset scheme for so-called *horizon class* images [10], that is, functions of the form:

$$f(x, y) = 1_{\{y > H(x)\}},$$

where we take for the horizon $H(x)$ a C^2 -smooth function, as in Figure 1. The notation 1_A means that $f(x, y)$ is the indicator function on A : $f(x, y) = 1$ if $(x, y) \in A$ and 0 otherwise. In 1-d, this horizon corresponds to a step function.




Fig. 16. (a) Putting limits to how far one can search in the normal direction. This is controlled by the reference domain. The dotted lines indicate the boundaries in which piercing points are accepted. (b) A 1-d version of the procedure followed when no normal piercing point is found in the acceptance region

6.2 2-d topological problems near singularities

As discussed above, adjacency handling in 2-d is non-trivial. Careful procedure design is necessary to deal with topological exceptions which otherwise may slow down the approximation convergence. This section discusses a simple algorithm which may not be optimal with respect to smoothness. It does preserve, however, the singularity approximation potential of a normal mesh scheme.

Adjacency in 2-d is handled by triangulation. The question arises how to proceed at the lines of singularity. At first sight, it might look natural to have vertical triangles between three vertices on the horizon. These triangles constitute a piecewise linear approximation of a vertical, curved surface. Subsequent refinements may cause flipping and mutually crossing triangles when the edges of these vertical triangles are being subdivided: the new piercing points all lie on the horizon (the data singularity), but not necessarily in a consistent way. The reason is that the projection of vertical triangles onto the domain space has area to zero. Hence, this function domain can no longer serve as a parameter or reference space to control the refinement process.

In order to avoid flipping triangles, the parameter domain (in our case domain of the function, i.e., a subset of \mathbb{R}^2) could put limits on how far the algorithm can look in the normal direction. The projection of the piercing point onto the parameter domain should be such that the refined triangular grid shows no overlaps (crossing edges) in this reference domain. If such a piercing point cannot be found in the normal direction, the new point should be taken as far as possible in the normal direction. This approach at least partly saves the benefits from a normal offset pointing towards the location of the horizon. A correction coefficient in vertical direction is then needed to find the corresponding function value. Figure 16(b) has the 1-d version of this case: starting from a midpoint m we are not allowed to find piercing points further than n (which in 1-d obviously does not make any sense), so we continue to look in the vertical direction, leading to point v .

6.3 Polyline approximation of the horizon

Working with vertical triangles on singularities poses at least two problems:

- (1) *Practically* lots of narrow, arbitrarily oriented triangles appear near the data discontinuity.
- (2) *Theoretically* all the complicated exception handling may and probably will destroy the 1-d fast convergence. Because of their intractability, vertical triangles are excluded, thereby giving up their intrinsic good approximation of horizons.

Since triangulation is not a goal as such, an alternative is to provide two function values, a lower and upper value, in each vertex. The goal is now a gradual improvement of these values as an approximation of the true discontinuity, just as in the 1-d case. We consider three types of edges in the triangulation: the first class contains edges that connect two horizon vertices. The second class connects a horizon with a non-horizon vertex, whereas the third class does not interfere with the horizon. Subdivision of edges in the first class should lead to a new horizon vertex. To maintain control over this process, we look for a piercing point in a direction normal to the vertical plane containing these two horizon vertices. This means that the slope of the edge connecting them is not taken into account.

Subdivision of edges in the second class (connecting a horizon point with a non-horizon vertex) should *not* lead to a new horizon vertex, but rather to a better approximation away from the horizon. We look for a piercing point which lies *in* the vertical plane containing that edge and the slope of that edge is the *only* degree of freedom to be filled in when computing the normal direction. Either this new vertex does or does not coincide with the existing end point on the horizon. If it does not, no special action has to be taken: we just insert the newly found vertex. If it does, we update one of the two function values in the existing point, just as in the 1-d case. We also insert a new vertex at the midpoint of the edge for two reasons. First, for simplicity, we want to subdivide every triangle into four children. Second, this avoids long and skinny triangles perpendicular to the horizon. We store the vertical offset between the function value in this vertex and its prediction on the subdivided edge as a detail coefficient, together with a label indicating that this is not a normal but vertical offset. We do not store the normal offset, since the label already tells the reconstruction algorithm that it should go as far as possible, up to the previously introduced vertex.

Edges in the third class do not need a special treatment. Details and an overview of this algorithm follow in Sections 6.5 and 6.6.

6.4 Straight horizons

Before we proceed with algorithmic details, let us take a moment to study the behaviour of this algorithm on straight line horizons. As soon as two points on the horizon have been detected, the only error comes from triangles with one edge on the horizon. By construction, the area of such a triangle is always less than a quarter of its parent: the parent has four children and the refinement




Fig. 17. Normal mesh approximation in 2-d: the shaded area corresponds to the dominant term in the approximation error.

points are always closer to the horizon than half of the total edge length. As a consequence, the total area of the error zone is reduced by a factor more than two in every step. Moreover, all edges that link a horizon point with a non-horizon point inherit the 1-d property that the error height H_j after j refinements has the same order of magnitude as the edge length S_j :

$$H_j = \mathcal{O}(S_j). \quad (4)$$

As a consequence, the total squared L_2 -error satisfies:

$$\|f - f_j\|^2 = \mathcal{O}(2^{-3j}), \quad (5)$$

to be compared to

$$\|f - f_j\|^2 = \mathcal{O}(2^{-j}) \quad (6)$$

for a wavelet approximation. The gain lies in the simultaneous reduction of error width and height.

6.5 Smooth horizons

If the horizon is a smooth curve (that is, at least differentiable with respect to local coordinates), then the algorithm aims at a polyline approximation of this curve. We then obtain the situation in Figure 17.

All triangles shown in this picture are in contact with the singularity and therefore they are not entirely flat in 3-d. Nevertheless, thanks to the fast 1-d error decay, these triangles rapidly approach the horizontal surface which they are supposed to approximate. In the shaded area however, between the horizon and its polyline approximation, the approximation of the 2-d function f is essentially incorrect.




Fig. 18. (a) Triangles near the polyline approximation (ac) of the horizon need special treatment in order to keep track of the location of the horizon. Refinement of the polyline approximation should lead to a closer approximation. Without special action, the true horizon may fall outside control of the approximating edge. (b) Therefore we give absolute priority to refinement of the approximating edge.

This term turns out to be dominant, so it is crucial to keep it under control. In order to do so, the algorithm must keep track of *where* the horizon is (i.e., its geometry) at all times and without (too much) overhead.

At the analysis (decomposition), singularities can be detected when finding piercing points: if the piercing surface is vertical (in practical applications: when it is “very steep”), we have a horizon point. At the reconstruction, the only way we can detect horizon points from the data themselves is by finding the vertices with two different function values. A newly found horizon point however has only one function value in its first step of existence. If a horizon vertex follows from the refinement of an edge in the approximating polyline, this can be monitored in the reconstruction phase as well.

Since triangle edges near the horizon are subdivided in a special way, it is important to:

- (1) ensure that refinement of the horizon approximation leads to new horizon points as often as possible
- (2) detect new horizon points immediately if they do not originate directly from the horizon approximation.

A refinement of an edge connecting two horizon points should lead to new horizon point, if possible. The procedure as described in Section 6.2, and in Figure 16, does not always let the new point go far enough. Figure 18(a) illustrates what may happen near sharp curvatures. Since the search field within a triangle is divided into three regions for each edge, the real edge may fall outside the control of its polyline approximation (ac). Therefore, we modify the procedure, such that in triangles containing the polyline approximation, the polyline edge has absolute priority. This is shown in Figure 18(b). Also, retriangulation after refinement is different: normally, the new vertices are connected with each other and with the end points of the edge that they subdivide. In Figure 18(b), however, the edge mn would have intersections with edges ap and cp . Therefore, we modify the retriangulation algorithm for triangles in which a new horizon vertex is found on a normal search direction: if necessary, the new horizon vertex is connected to its opposite, old




Fig. 19. First possible configuration where horizon vertices are found outside the refinement process of the polyline approximation of the horizon. At the reconstruction, this configuration is detected by the fact that the horizon approximating edge (ac) is not subdivided in a normal direction.




Fig. 20. Second configuration with horizon vertices not following from refinement of the polyline approximation (ac) of the horizon. Vertex b was probably an unforeseen, additional horizon vertex in the previous step, but only now this has to be taken into account.

vertex (leading to edge bp in Figure 18(b)).

Figure 18(b) also illustrates the second issue: connecting a new horizon vertex with its neighbours may result in two edges crossing the horizon. This poses no problem for edges that are part of the polyline approximation, but other edges crossing the horizon, like pn and pm in the figure, cause new horizon points which are not a result of the polyline refinement procedure. It is acceptable for the algorithm to store the vertex numbers of the first points of a newly found horizon. From then on, the algorithm should be able to reconstruct the horizon from the coefficients only.

To this end, we use an additional bit attached to each coefficient indicating the interpretation of that coefficient. If an edge connects two horizon points, the normal search direction was taken horizontal, as discussed before. If no intersection point is found, this means that the horizon has at least one intersection with at least one of the two remaining edges. So, if the interpretation bit tells the reconstruction algorithm that the corresponding coefficient is an exception, we are in one out of three situations, depicted in Figures 19, 20, and 21. A new horizon point may also show up while the present polyline approximation can still be subdivided in a normal direction. See Figure 22. The next step however immediately leads to the situation in Figure 20.

The two labels of the remaining edges in this triangle indicate in which situation we are: is the third vertex (b) at the opposite side of the horizon, and if so, do we have just one new horizon vertex (as




Fig. 21. Third configuration with horizon vertices not following from refinement of the polyline approximation (ac) of the horizon.




Fig. 22. Fourth configuration with horizon vertices not following from refinement of the polyline approximation of the horizon. Unlike the first three, this one cannot be detected from the refinement action on the horizon approximating edge. Nevertheless, subdivision immediately leads to the situation in Figure 20.

in Figure 19), or has the horizon intersections with both edges. If there is just one intersection (first case), this must be on the longer edge. Indeed, otherwise the polyline approximation of the horizon would have been refined in this step (vertex p in Figure 19 would have been on the horizon.)

The second possibility is Figure 20, where the two other edges should be refined as polyline approximations of the horizon. Note that in Figure 20, a new exception for the next step is generated promptly in vertex q .

The last case, as depicted in Figure 21 is the most complicated one. The approximating polyline segment belongs to two triangles. If these triangles both lie on the same side of the horizon, this means we are in the situation of Figure 21. A simple examination of the actual function values in neighbours of vertex b reveals which triangle should receive a special treatment: all edges of neighbouring triangles that contain b now have an intersection with the horizon.

In all three cases, to make sure that the algorithm locates the intersection immediately, we search along the edge itself instead of going normal. The detail coefficient is now a *tangential* offset. Figure 23 shows what happens in 1-d. Note that this tangential search is fast in *locating* the singularity. Edges that are marked as not having an intersection with the horizon are subdivided in a vertical direction. The question arises whether this combination of tangential offsets for edges and vertical offsets elsewhere could serve as the algorithm as such. This however would be poor for further




Fig. 23. *Tangential refinement. Good for fast location of singularities, not appropriate for further convergence.*

convergence near the singularity. The use of tangential offsets is limited to these exceptional cases.

6.6 Overview of the subdivision algorithm

Before stating and proving the main result, we summarize the subsequent steps of the subdivision procedure:

- (1) (a) The input of the *analysis (decomposition)* is a horizon class function and an initial, coarse scale grid. That initial grid consists typically of four corner points constituting the rectangular working domain. If this grid contains points on the horizon, this information has to be stated explicitly at the input. Each vertex in the initial grid has an x - and y -coordinate and two function values f_1 and f_2 . These two values are equal, unless the vertex lies on the horizon.
- (b) The input of the *synthesis (reconstruction)* algorithm consists of the initial grid and the normal mesh coefficients together with one additional bit for every coefficient with additional information on the interpretation of that coefficient (see below). We refer to this bit as the interpretation bit.
- (c) The objective of both analysis and synthesis is to (re-)construct an approximation of the horizon class function. This approximation includes a polyline approximation of the horizon. This polyline consists of segments which coincide with triangle sides of an edge-adaptive triangulation. The construction of this multiscale triangulation is part of the output, as well as one bit for every vertex indicating whether or not this vertex lies on the horizon. This bit is referred to as the horizon bit. It is created at the moment a vertex is inserted. The analysis has also the normal mesh coefficients and corresponding interpretation bits as output, as well as a list of vertices that are on a newly detected horizon.
- (2) For the construction of the approximation, repeat the following steps:
 - (a) First scan all vertices that have previously been identified as lying on the horizon curve. For all edges connecting two such horizon vertices, compute the midpoint and define the normal search direction as orthogonal to the triangle edge and parallel to the $x - y$ -plane. Find the piercing point of this search direction with the horizon.
 - (i) If this piercing point cannot be found within one of the triangles adjacent to the edge, we are in one of the situations of Figures 19, 20, or 21. Find a piercing point in vertical direction and label the corresponding interpretation bit as ‘vertical’.

- (ii) In order to find out in which of three possible situations (Figures 19, 20, or 21) we are, subdivide the longest triangle side, using tangential subdivision. If this succeeds, we are in the situation of Figure 19. Label the interpretation bit as ‘vertical’. Given the history, the reconstruction algorithm is able to interpret ‘vertical’ as ‘tangential’. Label the shortest triangle side as ‘normal’.
 - (iii) If the tangential refinement ends up in the existing opposite vertex (point \mathbf{b} in the figures), we are in the situation of Figures 20, or 21. If point \mathbf{b} had been detected as the first vertex of a newly found horizon, this is the situation of Figure 20. Remove this point from the list of newly found horizons. Label the refinements of both remaining triangle sides as ‘normal’. Refinement proceeds in the standard way for edges connecting two horizon points. Otherwise, label both refinements as ‘vertical’ and refine accordingly.
- (b) Next, refine all other edges.
- (i) If an edge connects a horizon point with a non-horizon point, define the normal search direction to lie in the vertical plane containing that edge, as explained in Section 6.3.
 - (ii) Otherwise compute the normal search direction, based on the function values in vertices of the triangles adjacent to the edge under consideration.
 - (iii) Find the piercing point of this search direction. If this piercing point lies outside the acceptance region as defined in Figure 16, continue the search in vertical direction. If the objective function is pierced in a point of discontinuity, add the newly found point at the end of the list of newly found horizons.

6.7 Convergence behaviour

The previous sections presented an algorithm for normal mesh subdivision on 2-d functions. This section proves that this algorithm has an optimal approximation rate. This convergence analysis is made difficult by the geometrical and topological configurations of triangulation and singularities, and the corresponding exceptions. The subsequent analysis starts from three assumptions, listed below. The first two are about the smoothness of the singularity, the third one is about the triangulation. The idea behind this third assumption is that at a given point, the triangulation is sufficiently fine so that we can replace the complicated subdivision procedure from the previous sections by a straightforward procedure, which is easy to analyze.

Assumption 1 *For ease of notations, we assume that the horizon can be described as a smooth function $y = H(x)$.*

Obviously, since normal offsets have no natural restriction towards functions, any curve with the same smoothness will work, even if it cannot be written in an explicit function form.

Assumption 2 *By a smooth horizon, we mean a twice differentiable function $y = H(x)$ with a bounded second derivative.*

In practice, this means we exclude curves with infinitely sharp cusps.

A bounded second derivative implies that the horizon curve has a bounded curvature. Indeed, the *curvature* κ in a point on a curve in two or more dimensions is defined as

$$\kappa = \frac{d\phi}{ds}, \tag{7}$$

where ϕ stands for the tangential angle of the given curve in a certain point and s is the arc length in this point. For planar curves that can be written as a function $y = H(x)$, the curvature equals

$$\kappa = \frac{\frac{d\phi}{dx}}{\frac{ds}{dx}} = \frac{\frac{d}{dx} \left(\arctan \frac{dy}{dx} \right)}{\sqrt{1 + y'(x)}} = \frac{y''(x)}{\left(1 + [y'(x)]^2\right)^{3/2}}.$$

It follows immediately that $|\kappa| \leq |y''(x)|$. The condition of bounded second derivatives could easily be relaxed to bounded curvature. It only makes the calculations a bit more complicated.

Assumption 3 *We assume that the normal refinement procedure with the exception handling from Section 6.5 creates an initial mesh in which the horizon is approximated by a polyline. Let h be the polyline resolution, i.e., the length of the longest segment of the polyline. We assume that h is sufficiently small. Also, we assume that the altitude of triangles on the polyline is sufficiently large and that none of these triangles has an edge tangential to the horizon. What we mean by ‘sufficiently large’ and ‘sufficiently small’ will be clarified in the proof of the subsequent theorem.*

Alternatively, we could assume that the initial mesh satisfies this condition, no matter how it was created.

We now state the main result of this paper:

Theorem 11 *Given a horizon class function f defined on $[0, 1] \times [0, 1]$, or on any compact subset of \mathbb{R}^2 .*

then under the Assumptions 1, 2, and 3, the normal mesh approximation f_n with n non-zero offsets converges in L_2 at a rate of

$$\|f - f_n\| = \mathcal{O}(n^{-1}),$$

Proof of Theorem 11:

The proof consists of two parts. The first part is a construction, including a proof that the construction is possible, the second part is the asymptotic analysis of this construction. We assume that after j refinement steps, using the procedure described in Section 6.5, the triangulation is in a shape satisfying the assumptions. For the proof, we concentrate on what happens with further refinement of one segment of the polyline. We introduce local coordinates, such that the segment coincides with the $y = 0$, and the two vertices are at $(0, 0)$ and $(h_j, 0)$, see Figure 24. $h_j \leq h$ is now

the resolution of this polyline segment. For the further refinement, we want to guarantee that the horizon approximation behaves as a normal polyline approximation of a 1-d smooth curve, without interference from the triangles adjacent to this polyline approximation. In order to avoid such interference, we introduce a ‘forbidden zone’, near the horizon, where no triangle edge refinement can take place, except for triangle edges that constitute segments of the polyline approximation of the horizon. In order to be able to construct such a forbidden zone, we must be sure that there is enough space in all subsequent refinements. The forbidden zone itself will avoid triangles with two small angles on the horizon. We also want to avoid triangles with arbitrarily obtuse angles on the horizon (i.e., angles arbitrarily close to π). This is guaranteed by the following lemma:

Lemma 12 *Given is a horizon curve with bounded curvature, i.e., in all points the curvature κ as defined in (7), satisfies $|\kappa| \leq K$, for some positive K . Given an initial polyline approximation of this horizon curve, embedded in a triangulation where the polyline resolution h_0 satisfies $h_0 \leq c \cdot 1/K$, with $c < 1$. Call $\alpha_{0,k}$ the angles in the initial triangulation (i.e., the scale level 0 triangulation) adjacent to a segment of the polyline approximation. (The index k runs over the total number of angles, say $k = 0, \dots, n - 1$.) Consider a normal refinement scheme where the approximation of the horizon is refined as a normal polyline. Edges connecting a horizon vertex with a non-horizon vertex are refined as described in Section 6.3. It then holds for all j that*

$$\max_k \alpha_{j,k} \leq \alpha_{\max} := \max(\max_k \alpha_{0,k}, \pi/2) + (\pi/2) \cdot \arcsin(Kh_0/2). \quad (8)$$

The proof of this lemma can be found in an appendix. This lemma is based on an assumption of bounded curvature. The reason is that the proof goes by induction where every step begins with the introduction of a local coordinate system. The subsequent coordinate systems are translations *and* rotations of each other. Curvature is independent of these operations, whereas the second derivative is not. On the other hand, if the initial horizon has bounded second derivative, it also has bounded curvature and if the polyline resolution satisfies the conditions of the Lemma, all local coordinate systems constructed as in Figure 24 allow to define $M = \sup_{[0,h]} |H''(x)|$.

As elaborated in the proof of Lemma 9, it follows immediately that $\sup_{[0,h_j]} |H'(x)| \leq Mh_j/2$ and $\sup_{[0,h_j]} |H(x)| \leq Mh_j^2/8$.

We are now ready to define the ‘forbidden zone’ near the horizon which replaces all exception handling procedures described in Section 6.5. This forbidden area is a strip adjacent to the segment of the polyline approximation of the horizon. The border of this forbidden area is parallel to the polyline segment, and the y -value F_j should accumulate to catch the following effects:

- (1) The maximum horizon value: $Mh_j^2/8$.
- (2) The maximum slope over a distance of $Mh_j/2$ (from $h_j/2$ to one of the end points)
- (3) A triangle with an obtuse angle in one of its horizon vertices.

Let $C_t = \tan(\alpha_{\max} - \pi/2)$, then we take

$$F_j = \frac{3Mh_j^2}{8} \cdot \frac{1}{1 - C_t M h_j/2}.$$




Fig. 24. Triangles containing the horizon cannot have arbitrarily obtuse angles in horizon vertices. The bold line is the horizon curve, the full lines are adjacent triangles at two successive scales. Dashed lines are constructions for the proof of Lemma 12.

The second factor corrects for possible obtuse angles. As described in Section 6.3, edges connecting a horizon vertex with a non-horizon vertex are subdivided in a 1-d fashion. If the new point has x and y coordinates on the existing vertex, we proceed as before: the approximation value in that vertex is updated, and we also insert a new vertex at the middle point of that edge (vertical offset). If the normal offset however points towards a vertex on the edge in the forbidden zone, i.e., a point with y -coordinate below F_j , we stop at the forbidden border and from there proceed vertically. The edges connecting the new vertex with its neighbours cannot possibly have an intersection with the horizon.

Before proceeding and using this exception handling rule, we must be sure that the construction of a forbidden zone is possible in each successive step. First of all, the altitude of the initial triangle must be at least twice the width of the forbidden strip. Second, this must also be the case in all subsequent steps. The worst case scenario in step j corresponds to a new vertex right on the border of the forbidden strip, see Figure 25. The condition is $F_{j+1} \leq \tau_{j+1,1}/2$, where $\tau_{j+1,1}$ is the altitude of the child triangle. It is clear that this altitude satisfies $\tau_{j+1,1} \geq F_j - Mh_j^2/8$. Because, by Lemma 9, $\lim_{h_j \rightarrow 0} h_{j+1}/h_j = 1/2$, we can, for sufficiently small h_j choose a constant C arbitrarily close to 1, such that $h_{j+1} \leq Ch_j/2$ and we know for certain that $C \leq 2$, so we can write:

$$\begin{aligned}
 F_{j+1} &= \frac{3Mh_{j+1}^2}{8} \frac{1}{1 - C_t M h_{j+1}/2} \\
 &\leq \frac{3Mh_{j+1}^2}{8} \frac{1}{1 - C_t M h_j/2} \\
 &\leq \frac{C^2 3Mh_j^2}{4 \cdot 8} \frac{1}{1 - C_t M h_j/2} \\
 &= \frac{3C^2 M h_j^2}{4 \cdot 8} \frac{1}{1 - C_t M h_j/2}
 \end{aligned}$$




Fig. 25. At sufficiently fine scale, it is possible to replace all exception handling with one simple rule, i.e., the construction of a forbidden area, in which no new vertex is allowed. This area does not contain the present polyline segment, so refinements on existing vertices is still allowed. This forbidden area prevents all exceptions in further steps, and makes the asymptotic analysis easier.

$$\begin{aligned}
&= \frac{3C^2}{4} \left(\frac{3Mh_j^2}{8} - \frac{Mh_j^2}{8} \right) \frac{1}{2} \frac{1}{1 - C_t M h_j / 2} \\
&\leq \frac{3C^2}{4} \left(\frac{3Mh_j^2}{8} \frac{1}{1 - C_t M h_j / 2} - \frac{Mh_j^2}{8} \right) \frac{1}{2} \\
&= \frac{3C^2}{4} \left(F_j - \frac{Mh_j^2}{8} \right) \frac{1}{2} \\
&= \frac{3C^2}{4} \frac{\tau_{j+1,1}}{2}
\end{aligned}$$

The construction is therefore possible as soon as $3C^2/4 \leq 1$, which is the case for h_j sufficiently small.

Figure 17 illustrates that two effects determine the convergence rate of the normal offset approximation:

- (1) the polyline approximation of the edge, and
- (2) the approximation near that polyline.

We call f_j the approximation of f after j refinement steps (i.e., at level j) and $f_j^{[1]}$ the horizon class function where the horizon is the polyline approximation after j refinement steps of the horizon in f . While f_j has non-flat triangles near the polyline approximation of the horizon, all triangles in $f_j^{[1]}$ are flat. f_j can be considered as an approximation of $f_j^{[1]}$, which in its turn is an approximation of f . We then have for the L_2 -error norm:

$$\|f - f_j\|^2 \leq \|f - f_j^{[1]}\|^2 + \|f_j^{[1]} - f_j\|^2.$$




Fig. 26. One refinement at the convex side of the horizon.

The actual proof (the second part) now consists of three analyses:

- (1) The asymptotic behaviour of the polyline approximation of the horizon.

Since the forbidden area prevents any exception, the polyline approximation of the horizon reduces to what happens in 1-d with a normal polyline approximation of a smooth curve. If we call H_j the polyline approximation of horizon H , then according to Corollary 10 we have $|H_j(x) - H(x)| \leq C2^{-2j}$. As a consequence,

$$\|f - f_j^{[1]}\|^2 = \mathcal{O}(2^{-2j}).$$

- (2) Triangles lying on the convex side of the horizon, with two vertices on the horizon. The analysis for these triangles starts from what we know about straight horizons, discussed in Section 6.4. In Figure 26, we know by construction that the area $|\Delta pma|$ of triangle Δpma satisfies $|\Delta pma| \leq |\Delta abc|/4$. The correction for the non-straight horizon leads to $|\Delta pna| \leq |\Delta abc|(1 + M2^{-2j})/4$. If we denote by A_j the total area of triangles with two vertices on the horizon, we have that $A_{j+1} \leq A_j(1 + M2^{-2j})/2$. Since the correction vanishes as j increases, we can write

$$A_j = \mathcal{O}\left((2 - \varepsilon)^{-j}\right),$$

for arbitrary positive ε .

- (3) Triangles lying on the concave side of the horizon, or triangles near an inflection point, have a forbidden area. This zone cannot change the fact that the total area A_j of triangles with two vertices on the horizon satisfies $A_{j+1} \leq A_j/2$.

Just as in Expression (4) for the straight horizon case, the error height H_j in a vertex on the horizon after j refinements has the same order of magnitude as the length S_j of the edge connecting this horizon vertex with a non-horizon vertex. This allows to conclude that

$$\|f_j - f_j^{[1]}\|^2 = \mathcal{O}\left((2 - \varepsilon)^{-3j}\right)$$

The factor ε is due to the convex side only.

Putting these elements together, we obtain a convergence rate of:

$$\|f - f_j\|^2 = \mathcal{O}(2^{-2j}).$$

The number n of non-zero coefficients needed for the approximation f_j at level j , satisfies:

$$n = \mathcal{O}(2^j),$$

from which we conclude that the L_2 -error of a nonlinear n -term approximation f_n decays as $\|f - f_n\| = \mathcal{O}(n^{-1})$. \square

For reconstruction, the procedure requires storage of one coefficient for each new vertex plus a label telling whether this offset is normal or vertical. We also have to label horizon vertices that are not found by subdivision of an edge connecting two horizon vertices. The number of these ‘‘initial’’ horizon points is of course neglectible.

7 Numerical Results

We first run the algorithm on an artificial test image, depicted in Figure 27(a). This image consists of four areas with constant grey values, separated by sharp, edges with bounded curvature (lines and circular arcs). Figure 27(b) shows a normal mesh approximation after 4 refinement steps, and Figure 27(c) displays the triangulation superposed to the original image. These images clearly illustrate that the triangulation is adaptive, leading to sharp approximations of the edges. An approximation with a non-adaptive triangulation is depicted in Figures 27(d) and 27(e). Figure 27(f) shows an approximation with the 5 coarsest levels of a wavelet transform. As can be expected, the approximation of the edges is blurred and blocky. This approximation involves more detail coefficients than the normal mesh approximation in Figure 27(b). For the wavelets, we used the biorthogonal wavelets of Cohen, Daubechies, and Feauveau with two vanishing moments. The lifted implementation of this wavelet transform is the simplest scheme with a linear prediction operator. It is therefore close to the prediction operator in the normal offset approach.

We now apply the normal offset decomposition to two examples of (digital) images. These images are discrete sets of pixels. In order to find normal piercing points, we need to interpolate these pixel matrices. A trivial triangular mesh allows for a piecewise planar interpolation in each point. As a consequence, there is no real discontinuity, only steep transitions. Many edges in images are blurred over several pixels anyway. The special actions to deal with real discontinuities are therefore unnecessary in this practical example.

Figure 28 compares a six level wavelet reconstruction with a six level normal offset reconstruction. The peppers image in Figure 29 is more dominated by large, smooth areas, separated by long, smooth edges. In this Figure, we compare a five level wavelet reconstruction with a five level normal offset reconstruction. Again, the adaptive triangulation leads to sharper edge reconstruction.




Fig. 27. An artificial test example: (a) image with four areas of constant grey value, separated by edges with bounded curvature. (b) Normal mesh approximation after 4 refinement steps. (c) Corresponding adaptive mesh, superposed to the original image. (d) Approximation by lifted wavelet decomposition on non-adaptive mesh. (e) The corresponding mesh. (f) Approximation using 5 coarsest levels of a classical tensor-product (i.e., separable) 2-d wavelets.

For small structures, such as the eyes in the photograph, or for texture, wavelets perform better in filling up the details. This suggests that a combination of both approaches might be interesting in practical applications. The application of normal meshes for real images is subject of current research. The results in this paper should therefore be seen as provisional.

8 Discussion and Conclusions

We have introduced an adaptive multiscale triangulation scheme for images based on a normal mesh decomposition. The scheme outperforms wavelet approximation thanks to the combined efforts of:

- adaptivity of the normal mesh approach. Normal mesh coefficients carry both information on location (or geometry) and discontinuity of the edge.



(a)




(b)

Fig. 28. Normal offsets (b), compared to classical wavelets (a). CDF 2,2 wavelets, i.e., linear prediction followed by two-taps update (i.e., update with two coefficients). Normal offsets are directed towards the edges. Triangulation allows for sharper edge reconstruction.



(a)



(b)

Fig. 29. A second test example, similar conclusions as in Figure 28.

- better approximation of edge contours when using triangulations instead of blocky tensor product wavelets.

Normal meshes could be used for image modeling, compression, and processing. However, algorithms will have to take into account that the decomposition is highly nonlinear.

Normal offsets are the key to *adaptive* triangulation of 2-d data sets. These data may contain line singularities, posing substantial problems to any tensor product based decomposition. In a normal offset decomposition the multiscale detail coefficients carry information on the location of the line singularities. For horizon class images, this leads to a $\mathcal{O}(n^{-1})$ approximation. The procedure is highly nonlinear. Topological exceptions need to be dealt with carefully.

Current research concentrates on the applicability of the normal offset concept on real images:

- (1) In practice, a good initial mesh seems to have crucial impact on the performance. The same idea of normal search can be used to select a limited number of crucial, coarsest scale samples (pixels).
- (2) The nonlinear character of the decomposition itself makes it harder to analyse the effect of removing or modifying a given coefficient. An analysis in 1-d is possible, e.g. in L_1 . In 2-d, the topological exceptions complicate the whole thing: changing a single coefficient may influence the topology on the following, finer grids.
- (3) The error of a normal mesh approximation in 2-d is completely dominated by the error of this piecewise linear approximation of the geometry information in the edge. This observation suggests that “curved” triangles have the potential of catching the geometry information even better.
- (4) This observation also explains why nonlinear approximation, for compression, is a non-trivial task. Thresholding or tree structured coefficient selection has to deal with the topological aspects.
- (5) In practice, images are obtained as samples on a square grid, so using normal meshes is equivalent to a *remeshing* operation. A second inverse remesh would be necessary to display a normal mesh approximation using a conventional display or printer. This makes things more complicated in practice.

Appendix A: Proof of Lemma 7

First note that the midpoint of the lobe containing $\pi/3$ satisfies:

$$m_n = \pi/3 + \Delta_n,$$

with

$$\Delta_n = (-1)^{n+1} \frac{\pi}{12 \cdot 2^n}.$$

The end points are the midpoints from the previous two steps. If n is even, then $a_n = m_{n-2}$ and $b_n = m_{n-1}$. If n is odd, we have $a_n = m_{n-1}$ and $b_n = m_{n-2}$. The rest of the proof concentrates on the n even case. The odd case is completely similar.

Second, the average log error reduction for angles in the $\pi/3$ -lobe can be simplified to:

$$\overline{\eta}_n(\alpha) = \frac{1}{n} \log \left(\prod_{k=0}^{n-1} \frac{1}{\tan(2^k \alpha)} \cdot \prod_{k=1}^n \frac{1}{\tan(2^k \alpha)} \right)$$

and its derivative equals:

$$\overline{\eta}_n'(\alpha) = -\frac{1}{n} \sum_{k=1}^n \left(\frac{2^k}{\sin(2^k \alpha)} + \frac{2^{k+1}}{\sin(2^{k+1} \alpha)} \right).$$

We show that for all even n :

- (1) $\overline{\eta}_n'(m_n) > 0$,
- (2) $\overline{\eta}_n'(\alpha) - \widehat{\eta}_n'(\alpha)$ has no zero,

which is sufficient to prove the lemma. Note that $\widehat{\eta}_n'(\alpha) = -\overline{\eta}_n'(2m_n - \alpha)$.

We proceed by induction. The main difficulty is that the interval on which we consider $\overline{\eta}_n$ changes too, so let us set:

$$x = 2^{n-2}(\alpha - m_n),$$

and

$$\begin{aligned} f_n(x) &:= n \overline{\eta}_n'(2^{-n+2}x + m_n) \\ &= -\sum_{k=1}^n \left(\frac{2^k}{\sin(2^k(2^{-n+2}x + m_n))} + \frac{2^{k+1}}{\sin(2^{k+1}(2^{-n+2}x + m_n))} \right) \\ &= -\sum_{i=0}^{n-1} 2^{n-i} \left(\frac{1}{\sin(2^{2-i}x + 2^{n-i}\frac{\pi}{3} + 2^{n-i}\Delta_n)} + \frac{2}{\sin(2^{3-i}x + 2^{n-i+1}\frac{\pi}{3} + 2^{n-i+1}\Delta_n)} \right) \\ &= -\sum_{i=0}^{n-1} 2^{n-i} \left(\frac{1}{\sin(2^{2-i}x + 2^{n-i}\frac{\pi}{3} + 2^{2-i}\Delta_2)} + \frac{2}{\sin(2^{3-i}x + 2^{n-i+1}\frac{\pi}{3} + 2^{3-i}\Delta_2)} \right) \end{aligned}$$

Note that $\sin(2^{m+2}\pi/3 + u) = \sin(2^m\pi/3 + u)$, so the first $n-1$ terms in the expression for $f_{n+2}(x)$ are equal to $4f_n(x)$ and hence

$$f_{n+2}(x) = 4f_n(x) - 4 \left(\frac{1}{\sin(2^{2-n}x + 4\frac{\pi}{3} + 2^{2-n}\Delta_2)} + \frac{2}{\sin(2^{3-n}x + 8\frac{\pi}{3} + 2^{3-n}\Delta_2)} \right)$$

$$-2 \left(\frac{1}{\sin(2^{1-n}x + 2\frac{\pi}{3} + 2^{1-n}\Delta_2)} + \frac{2}{\sin(2^{2-n}x + 4\frac{\pi}{3} + 2^{2-n}\Delta_2)} \right)$$

If we define

$$H(u) = -4 \left(\frac{1}{\sin(4u + 4\frac{\pi}{3})} + \frac{2}{\sin(8u + 8\frac{\pi}{3})} \right) - 2 \left(\frac{1}{\sin(2u + 2\frac{\pi}{3})} + \frac{2}{\sin(4u + 4\frac{\pi}{3})} \right)$$

we have

$$f_{n+2}(x) = 4f_n(x) + H \left(\frac{x + \Delta_2}{2^n} \right),$$

where $f_0(x) = 0$.

This allows to fill in the values we need:

- (1) $\overline{\eta}'_n(m_n) = f_n(0) = 4f_{n-2}(0) + H \left(\frac{\Delta_2}{2^{n-2}} \right)$. So, $f_2(0) = H(\Delta_2) \approx 1.4889$. Since $H(u)$ is monotonously decreasing on $[\Delta_2, 0]$, we can construct a minorant for $f_n(0)$:

$$\hat{f}_n = 4\hat{f}_{n-2} + H(0).$$

It is easy to prove that all elements in \hat{f}_n are positive, and hence so is $f_n(0)$. This proves that $\overline{\eta}'_n(m_n) > 0$.

- (2) Next, we compare the derivative of the mean log decay with its mirror:

$$\begin{aligned} \overline{\eta}'_n(\alpha) - \widehat{\eta}'_n(\alpha) &= f_n(x) + f_n(-x) = (f_{n-2}(x) + f_{n-2}(-x)) \\ &\quad + \left(H \left(\frac{x + \Delta_2}{2^{n-2}} \right) + H \left(\frac{-x + \Delta_2}{2^{n-2}} \right) \right) \end{aligned}$$

We want to prove that this never has a zero. It can be proven that the first two expressions, for $n = 2$ and $n = 4$ are always positive with a minimum in $x = 0$. In order to extend this to arbitrary n , we first note that for even $n > 4$

$$H \left(\frac{x + \Delta_2}{2^{n-2}} \right) + H \left(\frac{-x + \Delta_2}{2^{n-2}} \right)$$

reaches its (negative) minimum in the end points $x = \pm\pi/16$. Second, this minimum is bounded by $2H(0)$. (All this requires plots and/or intensive calculations.) This allows to construct a positive minorant sequence for $f_n(x) + f_n(-x)$, from which it follows that $\overline{\eta}'_n(\alpha) - \widehat{\eta}'_n(\alpha)$ is strictly positive on the entire lobe.

This concludes the proof of Lemma 7. Following the argument in the last paragraphs of Section 3, this also concludes the proof of Theorem 3.

Appendix B: Proof of Lemma 12

Consider a detail of the initial setting at level j , sketched in Figure 24. The polyline segment \mathbf{ab} is a local approximation of the horizon curve. This curve is a function $y = H(x)$ in the local coordinate system.

After refinement, the angles in the vertices from the previous (coarser) scale, \mathbf{a} and \mathbf{b} equal $\alpha_{j+1, \ell(j,k)} = \alpha_{j,k} \pm \theta_{j,k}$, where the index function $\ell(j, k)$ maps the angle at scale j , location k to its corresponding angle at scale $j + 1$. The offset angle $\theta_{j,k}$, indicated in Figure 24 is maximized by maximizing the normal offset, i.e., the distance between the normal subdivision point \mathbf{n} and its predicting midpoint \mathbf{m} . Given the maximum curvature, K , the curve leading to the largest offset, is a circular arc with radius $R = 1/K$. By basic goniometry, the corresponding offset is then $R - \sqrt{R^2 - h_j^2/4} \approx Kh_j^2/8$, and the maximum angle offset

$$\theta_{j,\max} := \max_k \theta_{j,k} = (1/2) \cdot \arcsin(h_j/2R) \leq \pi h_j/8R = K\pi h_j/8. \quad (9)$$

The angles in the newly inserted vertex \mathbf{n} are maximized if the new points \mathbf{p} , \mathbf{q} , \mathbf{r} , and \mathbf{s} , coincide with the midpoints of the corresponding edges (thereby giving zero normal offsets). Indeed, by construction, these refinement points cannot be further away from the horizon. From Figure 24, it follows that

$$\widehat{\mathbf{anp}} \leq \widehat{\mathbf{aip}} = \widehat{\mathbf{ajc}} = \widehat{\mathbf{abc}} + \theta.$$

For the angle $\widehat{\mathbf{anr}}$, the analysis falls apart into two cases. If the angle $\widehat{\mathbf{arm}} \leq \pi/2$, then the circle passing through \mathbf{arm} has its center below \mathbf{am} , hence the line segment \mathbf{mn} lies entirely outside that circle. As a consequence, the angle $\widehat{\mathbf{anr}}$ is smaller than $\widehat{\mathbf{amr}}$, since both are subtended by the same chord, \mathbf{ar} , but $\widehat{\mathbf{amr}}$ is inscribed in the circle, while $\widehat{\mathbf{anr}}$ lies outside the circle. This first case implies the case that $\widehat{\mathbf{amr}}$ is obtuse. In other words, if the angle $\widehat{\mathbf{abd}} = \widehat{\mathbf{amr}} \geq \pi/2$, its offspring $\widehat{\mathbf{anr}}$ is smaller. If $\widehat{\mathbf{arm}} > \pi/2$, the angle $\widehat{\mathbf{anr}}$ can be arbitrarily close to $\pi/2 - \theta$, namely if $\widehat{\mathbf{mar}}$ tends to zero while $\widehat{\mathbf{amr}}$ is kept constant.

All together, we have that

$$\max_k \alpha_{j+1,k} \leq \max\left(\pi/2, \max_k(\alpha_{j,k} + \theta_{j,k})\right). \quad (10)$$

By induction, it follows that

$$\max_k \alpha_{j,k} \leq \max\left(\max_k \alpha_{0,k}, \pi/2\right) + \sum_{i=1}^j \max_k \theta_{i,k}. \quad (11)$$

Using equation (9), this becomes:

$$\max_k \alpha_{j,k} \leq \max(\max_k \alpha_{0,k}, \pi/2) + (K\pi/8) \cdot \sum_{i=1}^j h_i. \quad (12)$$

A straightforward induction argument proves that h_i is always smaller than the arc length of the circle segment with radius R , subtended by a chord of length $h_0/2^i$. These majorant arc lengths constitute a geometrical series themselves, so the summed arc lengths are bounded by twice the initial arc length of a circle segment with radius R , subtended by a chord of length h_0 . As a consequence, we may write:

$$\begin{aligned} \max_k \alpha_{j,k} &\leq \max(\max_k \alpha_{0,k}, \pi/2) + (K\pi/8) \cdot 2 \cdot 2R \arcsin(h_0/2R) \\ &= \max(\max_k \alpha_{0,k}, \pi/2) + (\pi/2) \arcsin(Kh_0/2). \end{aligned}$$

□

Acknowledgments

We would like to thank the two referees for their very useful and constructive comments. Maarten Jansen wishes to thank also his colleagues Frank Redig and Remco van der Hofstad for interesting discussions on this paper.

References

- [1] G. Al-Regib, Y. Altunbasak, and Mersereau R. M. Hierarchical motion estimation with content-based meshes. *IEEE Transactions on Circuits and Systems for Video Technology*, 13(10):1000 – 1005, 2003.
- [2] E. J. Candès and D. L. Donoho. Recovering edges in ill-posed inverse problems optimality of curvelet frames. *Annals of Statistics*, 30:784–842, 2000.
- [3] E. J. Candès and D. L. Donoho. New tight frames of curvelets and optimal representations of objects with piecewise c -2 singularities. *Comm. Pure Appl. Math.*, 57(2):219–266, 2004.
- [4] I. Daubechies. *Ten Lectures on Wavelets*. CBMS-NSF Regional Conf. Series in Appl. Math., Vol. 61. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1992.
- [5] I. Daubechies, O. Runborg, and W. Sweldens. Normal polyline approximation. *Constructive Approximation*, 20(3):399–463, May 2004.
- [6] S. Dekel and D. Leviatan. Adaptive multivariate approximation using binary space partitions and geometric wavelets. *SIAM Journal on Numerical Analysis*, To Appear:–, 2005.

- [7] R. A. DeVore, B. B. Jawerth, and V. Popov. Compression of wavelet decompositions. *Amer. J. Math.*, 114:737–785, 1992.
- [8] R. A. DeVore, B. Jawerth, and B. J. Lucier. Image compression through wavelet transform coding. *IEEE Transactions on Information Theory*, 38(2):719–746, 1992.
- [9] M. N. Do and M. Vetterli. Contourlets. In G. Welland, editor, *Beyond Wavelets*, pages 83–107. Academic Press, 2003.
- [10] D. L. Donoho. Wedgelets: Nearly minimax estimation of edges. *Annals of Statistics*, 27(3):859–897, 1999.
- [11] D. L. Donoho and X. Huo. Beamlets and multiscale image processing. Technical report, Department of Statistics, Stanford University, 2001.
- [12] D. L. Donoho and I. M. Johnstone. Ideal spatial adaptation via wavelet shrinkage. *Biometrika*, 81:425–455, 1994.
- [13] D. L. Donoho, M. Vetterli, R. A. DeVore, and I. Daubechies. Data compression and harmonic analysis. *IEEE Transactions on Information Theory*, 44(6):2435–2476, October 1998.
- [14] N. Dyn, D. Levin, and S. Rippa. Data dependent triangulations for piecewise linear interpolation. *IMA Journal of Numerical Analysis*, 10:137–154, 1990.
- [15] N. Dyn and S. Rippa. Data-dependent triangulations for scattered data interpolation and finite-element approximation. *Appl. Num. Math.*, 12:89–105, 1993.
- [16] R. M. Figueras i Ventura, P. Vanderghenst, and P. Frossard. Low rate and scalable image coding with redundant representations. *IEEE Transactions on Image Processing*, Submitted:–, 2005.
- [17] I. Guskov, K. Vidimee, W. Sweldens, and P. Schröder. Normal meshes. In *SIGGRAPH 2000 Conference Proceedings*, 2000.
- [18] E. Le Pennec and S. Mallat. Sparse geometrical image representations with bandelets. *IEEE Transactions on Image Processing*, 2003.
- [19] S. Mallat. Wavelets, approximation, and compression. *IEEE Signal Processing Magazine*, 18(5):59–73, September 2001.
- [20] S. G. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 525 B Street, Suite 1900, San Diego, CA, 92101-4495, USA, 1998.
- [21] I. Niven. *Irrational Numbers*, volume 11 of *Carus Mathematics Monographs*. Mathematical Association of America, 1956.
- [22] S. Rippa. Long and thin triangles can be good for linear interpolation. *SIAM J. Numer. Anal.*, 29(1):257–270, 1992.
- [23] J. Romberg, M. Wakin, and R. Baraniuk. Approximation and compression of piecewise smooth images using a wavelet/wedgelet geometric model. In *Proc. IEEE Int. Conf. on Image Proc. — ICIP '03*, 2003.
- [24] R. Shukla, P. L. Dragotti, M. Do, and M. Vetterli. Rate distortion optimized tree structured compression algorithms. *IEEE Transactions on Image Processing*, To Appear:–, 2005.

- [25] W. Sweldens and P. Schröder. Building your own wavelets at home. In *Wavelets in Computer Graphics*, ACM SIGGRAPH Course Notes, pages 15–87. ACM, 1996.
- [26] P van Beek, A. M. Tekalp, N. Zhuang, I. Celasun, and M. H. Xia. Hierarchical 2-d mesh representation, tracking, and compression for object-based video. *IEEE Transactions on Circuits and Systems for Video Technology*, 9(2):353–369, March 1999.
- [27] R. M. Willett and R. D. Nowak. Platelets: a multiscale approach for recovering edges and surfaces in photon-limited medical imaging. *IEEE Transactions on Medical Imaging*, 22(3):332–350, March 2003.