APPROXIMATION AND COMPRESSION OF PIECEWISE SMOOTH IMAGES USING A WAVELET/WEDGELET GEOMETRIC MODEL

Justin K. Romberg, Michael B. Wakin, and Richard G. Baraniuk

Department of Electrical and Computer Engineering
Rice University, Houston, TX

ABSTRACT

Inherent to photograph-like images are two types of structures: large smooth regions and geometrically smooth edge contours separating those regions. Over the past years, efficient representations and algorithms have been developed that take advantage of each of these types of structure independently: quadtree models for 2D *wavelets* are well-suited for uniformly smooth images (C^2 everywhere), while quadtree-organized *wedgelet* approximations are appropriate for purely geometrical images (containing nothing but C^2 contours). This paper shows how to *combine* the wavelet and wedgelet representations in order to take advantage of both types of structure simultaneously. We show that the asymptotic approximation and rate-distortion performance of a wavelet-wedgelet representation on piecewise smooth images mirrors the performance of both wavelets (for uniformly smooth images) and wedgelets (for purely geometrical images). We also discuss an efficient algorithm for fitting the wavelet-wedgelet representation to an image; the convenient quadtree structure of the combined representation enables new algorithms such as the recent WSFQ geometric image coder.

1. INTRODUCTION

At the core of image processing lies the problem of characterizing image structure. Building an accurate, tractable mathematical characterization that distinguishes a "real-world, photograph-like" image from an arbitrary set of data is fundamental to any image processing algorithm. There are two particular types of structure that any processing algorithm should exploit: images contain smooth, homogeneous regions (*grayscale regularity*), and these regions are separated by smooth contours (*geometric regularity*).

A vital part of the characterization is the *image representation*. Using an atomic decomposition, we approximate an image $X(s)$ using a linear combination of atoms b_i from a dictionary $\mathcal{D}$

$$\widehat{X}(s) = \sum_i \alpha_i b_i(s), \quad \{b_i\} \subset \mathcal{D} \tag{1}$$

$$X(s) = \widehat{X}(s) + \epsilon(s). \tag{2}$$

We desire a dictionary $\mathcal{D}$ with the following properties:

- Every image of interest is well-approximated ($\|\epsilon(s)\|_2^2$ small) using relatively few terms ($\#\{b_i\} := N$ small). We quantify the approximation power by measuring how fast $\|\epsilon\|_2^2 \to 0$ as $N \to \infty$.
- There is a practical algorithm to find a set of $\{b_i\}$ that yield a good approximation.
- $\mathcal{D}$ has limited complexity; we can efficiently specify (encode) the $\{b_i\}$ used in the approximation.

In this paper, we will introduce a new geometrical representation that is well suited for piecewise smooth images (smooth away from smooth contours).

The wavelet transform [1] has emerged as a powerful representation for image processing. The success of wavelets is due to the fact that they provide a *sparse* representation for 1D smooth signals interrupted by isolated discontinuities [1] (this is an excellent model for "image slices" — 1D cross sections of a 2D image). If $X_{\rm slice}$ is a 1D image slice that is uniformly smooth (C^2 everywhere) and we use an orthonormal wavelet basis for $\mathcal{D}$ in (1), then we can find a set of wavelets $\{b_i\}$ such that $\|\epsilon\|_2^2 \sim N^{-4}$ as $N \to \infty$ (this is also the fastest rate of decay for any orthogonal basis for this class of signals).[1] Amazingly, this rate does not change if we introduce a finite number of discontinuities into $X_{\rm slice}$; the wavelet representation is equally powerful for piecewise smooth and uniformly smooth image slices. Moreover, the $\{b_i\}$ chosen for the piecewise smooth case can be restricted so that they correspond to nodes on a connected binary tree [2]. As a result, the choice of atoms $\{b_i\}$ used to build up the image slice can be efficiently "coded," and we can use fast tree pruning algorithms to find the optimal set of atoms given $X_{\rm slice}$ [3]. Wavelet-based compression algorithms using these ideas can be shown to have optimal asymptotic rate-distortion decay [2, 4].

Instead of a 1D image slice, consider an "image segment" $X_{\rm seg}(s)$ — a 2D local region of an image. If $X_{\rm seg}(s)$ is uniformly smooth (C^2 everywhere), then wavelets coupled with quadtree models and algorithms still achieve the best possible approximation rate — $\|\epsilon\|_2^2 \sim N^{-2}$ for 2D [1]. If, however, $X_{\rm seg}$ is smooth everywhere except along a smooth contour, see Figure 1, the approximation rate slows to $\|\epsilon\|_2^2 \sim N^{-1}$ [5]. That is, unlike the 1D image slice result, adding a discontinuity along a contour to a 2D image segment significantly affects the ability of wavelets to provide a sparse representation. In fact, no matter how smooth the image is away from the contour, or how simple the contour itself is, the approximation rate remains the same. It simply takes too many wavelet basis functions to build up edge contours in images.

Recent research in harmonic analysis has focussed on finding representations for 2D piecewise smooth image segments that mirror the effectiveness of wavelets on 1D image slices (optimal theoretical approximation rate, simple models, and practical processing algorithms) [5–8]. For "horizon class" (or "cartoon") image

Email: jrom@rice.edu, wakin@rice.edu, richb@rice.edu. Web: dsp.rice.edu. This work was supported by the National Science Foundation grants CCR–9973188, ONR grant N00014-02-1-0353, AFOSR grant F49620-01-1-0378, and the Texas Instruments Leadership University Program.

[1] We write $f(k) \sim g(k)$ when there exists a constant C independent of k such that $f(k) \leq Cg(k)$.

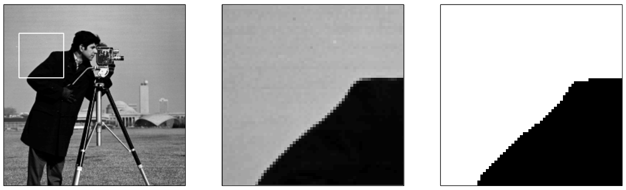

(a) (b) (c)

Fig. 1. *(a) "Cameraman" image, (b) image segment, (c) horizon class image playing the role of H_c in (4).*

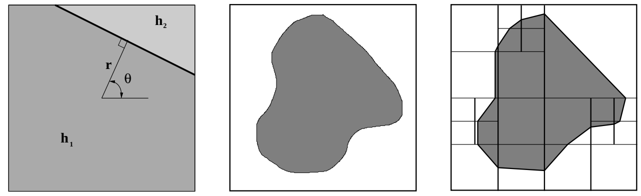

(a) (b) (c)

Fig. 2. *(a) A wedgelet in a dyadic block is parameterized by angle θ, offset r, and heights h_1, h_2. (b) Example of a "simple cartoon image", (c) The wedgelet representation divides the domain of the image into dyadic squares, using a piecewise constant function in each square to approximate the image.*

segments that are constant except along a smooth contour

$$X_{\mathrm{seg}}(s) := X_{\mathrm{seg}}(s_1, s_2) = \begin{cases} 1, & s_2 > c(s_1) \\ 0, & s_2 \le c(s_1) \end{cases} \qquad c(s_1) \in C^2([0,1]) \tag{3}$$

the simplest of these, the wedgelet representation [8, 9], has a success story similar to wavelets in 1D.

A *wedgelet* is a piecewise constant function on a dyadic square S that is discontinuous along a line through S with *orientation* $\ell := (r, \theta)$, see Figure 2(a) for an illustration. A *wedgelet representation* of an image X consists of a dyadic partition of the domain of X along with a wedgelet function in each dyadic square, see Figure 2(c).

Like the 2D wavelet transform, the wedgelet representation can be organized on a quadtree. The nodes of the quadtree define the dyadic partition. Attached to each leaf of the quadtree is a wedgelet that approximates the image over the corresponding dyadic square. Just as we can prune the wavelet quadtree at relatively coarse scales in regions where the image is smooth, we can prune the wedgelet quadtree in regions of a horizon class image where the contour is almost linear. As such, horizon class images can be well approximated using a small number of wedgelets. If we take $\mathcal{D}$ in (1) above to be a suitable wedgelet dictionary, then we can use tree pruning algorithms to select a representation that achieves $\|\epsilon\|_2^2 \sim N^{-2}$. In addition, we can exploit the regularity of the image contours further by incorporating a *multiscale geometry model* into the selection algorithm that favors sets of wedgelets whose orientations "line up" between dyadic blocks [10].

In this paper, we will show that we can achieve the same N^{-2} approximation rate on "C^2/C^2" piecewise smooth image segments

$$X_{\mathrm{seg}}(s) = X_1(s) \cdot H_c + X_2(s) \cdot (1 - H_c)$$
$$X_1(s), X_2(s) \in C^2([0,1]^2) \quad H_c \in \text{Horizon class} \tag{4}$$

using a dictionary composed of wavelets and *wedgeprints* — wedgelets projected onto fine wavelet scales (a concept similar to the wavelet footprints of [11]). The combined dictionary is simple enough that a simple coder based on this dictionary achieves near-optimal asymptotic rate-distortion performance $D(R) \sim (\log R)^2/R^2$. Moreover, given an image segment, we can find a suitable representation using a fast dynamic program. These results serve as a theoretical justification for combined wavelet-wedgelet image coders such as the WSFQ [12].

In addition to improving the approximation and coding rate, the choice of representation carries geometrical side information. The coder will choose to use wedgeprint along contours in the image, making semantic information about the locations edges available in the compressed domain. These "edge maps" could prove useful for classification or performing rapid database searches in the compressed domain.

In Section 2, we present asymptotic approximation and rate-distortion bounds for the wavelet-wedgeprint dictionary. Section 3 briefly discusses how to choose a representation given an image.

2. REPRESENTING PIECEWISE SMOOTH IMAGES WITH WAVELETS AND WEDGEPRINTS

In this section, we outline an argument to show that by using a dictionary composed of orthonormal wavelets and *wedgeprints* (defined below) we can achieve $\|\epsilon\|_2^2 \sim N^{-2}$ for C^2/C^2 image segments. In addition, the structure of the wavelet-wedgeprint dictionary allows us to efficiently encode the atoms $\{b_i\}$ chosen for the approximation, allowing us to translate the approximation result into a rate-distortion bound.

Let be $\psi_{j,k}$ be a set of compactly supported wavelet basis functions, and let $w_{j,k}$ be the corresponding wavelet coefficients of a C^2/C^2 image segment X_{seg}. Choose a maximum (finest) scale J. We will describe a way to prune the wavelet quadtree, using wedgeprint functions instead of wavelets along the contour at the finest scale, to approximate X_{seg} with $\mathrm{Error} \sim 2^{-J}$ using $N \sim 2^{J/2}$ terms.

Wavelet approximation. Simply truncating the wavelet transform at scale J gives us a squared error

$$\mathrm{Error}_{\mathrm{trunc}} \sim 2^{-J}. \tag{5}$$

We will classify the remaining $w_{j,k}$ as type I if the edge contour does not intersect $\psi_{j,k}$, and type II if it does. Since $\psi_{j,k}$ is compactly supported, if $w_{j,k}$ is type I , then all of the descendants of $w_{j,k}$ on the wavelet quadtree are also type I . The regularity of the contour means that at scale j, there will be $\sim 2^j$ wavelet coefficients of type II and $\sim 2^{2j}$ of type I .

The image is uniformly C^2 on the support of a type I $\psi_{j,k}$; as a result, the coefficient magnitudes decay quickly across scale [1]. For type II $w_{j,k}$, the decay is much slower

$$|w_{j,k}|^2 \sim 2^{-6j} \quad \forall\ w_{j,k} \text{ type I} \tag{6}$$
$$|w_{j,k}|^2 \sim 2^{-2j} \quad \forall\ w_{j,k} \text{ type II}\,. \tag{7}$$

We will build the approximation to X_{seg} by keeping all the wavelet coefficients up to scale $J/4$ and then pruning the wavelet quadtree beneath all type I nodes at scales $J/4 \dots J/2$. We will perform $\sim 2^{J/2}$ prunings at scale $J/4$ and $\sim 2^{J/2}$ as we prune

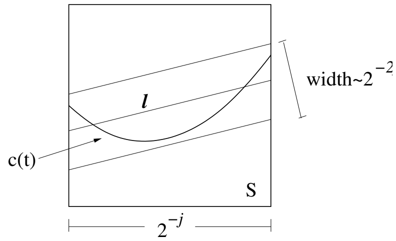

Fig. 3. *Since the contour $c(t)$ is smooth inside of dyadic square S with sidelength 2^{-j} it can be bounded by a strip of width $\sim 2^{-2j}$ around a line ℓ.*

around the contour at scales $J/4+1\ldots J/2$. Each pruning results in an error of $\sim 2^{-3J/2}$, bringing the total error to

$$\text{Error}_{\text{pruneI}} \sim 2^{-J}. \tag{8}$$

The number of type I and type II wavelet coefficients, respectively, retained from scales $0\ldots J/2$ is

$$N_I \sim 2^{J/2} \tag{9}$$

$$N_{II} \sim 2^{J/2}. \tag{10}$$

At scale $J/2$, we are left with $\sim 2^{J/2}$ wavelet coefficients of type II . Simply pruning the wavelet quadtree beneath these nodes would result in an error of $\sim 2^{-J/2}$ using $\sim 2^{J/2}$ coefficients, an approximation rate of N^{-1} rather than N^{-2}. On the other hand, continuing the pruning process down to scale J would result in an error of $\sim 2^{-J}$ using $\sim 2^J$ coefficients; also an N^{-1} approximation.

Wedgeprint approximation. We will remedy this problem by using $\sim 2^{J/2}$ *wedgeprint* functions to approximate the image segments at scale $J/2\ldots J$. To develop the idea, we return briefly to the spatial domain and examine how wedgelets can be used to locally approximate X_{seg} along the contour.

Let S be a square subregion of $[0,1]^2$ with sidelength $K2^{-j}$ through which the contour $c(s_1)$ passes. Since $c(s_1)$ is C^2, there exists a straight line ℓ through S such that c is contained in a strip of width $\sim 2^{-2j}$ around ℓ (see Figure 3) [10]. The image X_{seg} has bounded first derivative, so there exist constants h_1, h_2 such that $|Xseg - h_1| \leq C_{h_1}2^{-j}$ on one side of the strip and $|X_{\text{seg}} - h_2| \leq C_{h_2}2^{-j}$ on the other side. Thus, using a wedgelet $\mathcal{W}(S;\ell,h_1,h_2)$ with orientation ℓ and heights h_1, h_2, we can can approximate X_{seg} on S with error

$$\|X_{\text{seg}}(S) - \mathcal{W}(S;\ell,h_1,h_2)\|_2^2 \sim 2^{-3j}. \tag{11}$$

(The error is $\sim 2^{-4j}$ in the region outside the strip and $\sim 2^{-3j}$ inside the strip.)

Now consider a subtree of wavelet coefficients rooted along the contour at scale $J/2$ and location k. Since the wavelets are compactly supported, the subtree builds up the image in a square S with sidelength $K2^{-J/2}$ (K depends on the size of the support of $\psi_{J/2,k}$). Letting $W^J_{J/2,k}$ be the subspace spanned by all of the basis functions in the subtree beneath node $(J/2,k)$ down to scale J, we define the wedgeprint $\varphi_{J/2,k}(S;\ell)$ to be a projection of a wedgelet onto $W^J_{J/2,k}$:

$$\varphi'_{J/2,k}(S;\ell) = \text{Proj}\left(\mathcal{W}(S;\ell,0,1) \to W^J_{J/2,k}\right) \tag{12}$$

$$\varphi_{J/2,k}(S;\ell) = \frac{\varphi'_{J/2,k}}{\|\varphi'_{J/2,k}\|_2}. \tag{13}$$

By (11), we know that we can use one wedgeprint instead of the $\sim 2^{J/2}$ wavelets in the subtree rooted at $(J/2,k)$ to approximate the projection of X_{seg} onto $W^J_{J/2}$ and pay an error penalty of just $\sim 2^{-3J/2}$. We will use

$$N_\varphi \sim 2^{J/2} \tag{14}$$

total wedgeprints to build up the contour in X_{seg} at fine scales, bringing the wedgeprint error to

$$\text{Error}_\varphi \sim 2^{-J}. \tag{15}$$

Collecting the results, we have used

$$N_{\text{total}} = N_I + N_{II} + N_\varphi \sim 2^{J/2} \tag{16}$$

total wavelets and wedgeprints to approximate the piecewise smooth signal X_{seg} with total squared error

$$\text{Error}_{\text{total}} = \text{Error}_{\text{trunc}} + \text{Error}_{\text{pruneI}} + \text{Error}_\varphi \sim 2^{-J}. \tag{17}$$

To conclude, if we choose $\mathcal{D}$ in (1) to be a combined wavelet-wedgeprint dictionary, we achieve the best-possible asymptotic approximation decay of

$$\|\epsilon\|_2^2 \sim N^{-2} \tag{18}$$

for the class of C^2/C^2 images.

Rate-distortion. In the context of image compression, we are interested more in asymptotic *rate-distortion* performance than approximation decay. An image coder will need to spend bits not only on the expansion coefficients for the atoms chosen from $\mathcal{D}$, but also must encode *which* atoms were chosen. Quantizing the expansion coefficients also introduces and additional source of error.

Fortunately, we can translate (18) into a rate-distortion bound without too much difficulty. If we use $\sim J$ bits to quantize each wavelet coefficient, we will require $R_{\text{wavelet}} \sim J2^{J/2}$ bits overall, while incurring a quantization distortion $D_{\text{wavelet}} \sim 2^{-3J/2}$.

For the wedgeprints, we must quantize the orientation as well as the expansion coefficient. We can limit ℓ to be one of $\sim 2^{2J}$ possibilities (requiring $\sim J$ bits to code) and still have (11) hold. Using another $\sim J$ bits to quantize the $J/2$ wedgeprint coefficients, we use $R_\varphi \sim J2^{J/2}$ total bits while incurring distortion $D_\varphi \sim 2^{-3J/2}$.

The last component to consider is the indexing cost: the number of bits required to specify which wavelets and wedgeprints are being used. Assigning a symbol from {Prune, Wavelet, Wedgeprint} to each node in the quadtree requires less than 2 bits per coefficient. Since all of the basis functions live on a connected tree, we can code all of these symbols using $R_{\text{indexing}} \sim 2^{J/2}$ bits.

In summary, we have an overall rate of

$$R = R_{\text{wavelets}} + R_\varphi + R_{\text{indexing}} \sim J2^{J/2} \tag{19}$$

with distortion

$$D \leq \text{Error}_{\text{total}} + D_{\text{wavelet}} + D_\varphi \sim 2^{-J}. \tag{20}$$

Combining (19) and (20)

$$D(R) \sim \frac{(\log R)^2}{R^2} \tag{21}$$

we see that our simple wavelet-wedgelet coder achieves near optimal asymptotic rate-distortion performance.

3. SELECTING THE WAVELET-WEDGEPRINT REPRESENTATION

We have shown that there exists a configuration of wavelets and wedgeprints that closely approximates a C^2/C^2 image. This section addresses the more practical problem of finding a good wavelet-wedgeprint representation for a given image.

We are able to capture the contours in the image at fine scales using very few wedgeprint functions. Of course, a coding (or other processing) algorithm would not know the locations of the contours *a priori*. The encoder needs to make decisions about where to place the wavelets and wedgeprints dynamically.

Fortunately, the structure of the wavelet-wedgeprint dictionary allows us to formulate and solve an optimization problem to find the best configuration for a given image. For simplicity, we will discuss the approximation problem where the complexity of the representation is simply the number of terms. We refer to [12] for the implementation of an actual image coder.

Given an image and an N, we wish to find the best (smallest error) N-term configuration of wavelets and wedgeprints. Putting the problem in Lagrange form, we wish to solve

$$\min\left(\mathrm{Error} + \lambda N\right) \qquad (22)$$

Since the atoms chosen for the representation lie on a connected tree and we restrict wedgeprints to live on the leaves of this tree, (22) can be solved efficiently with the classical CART dynamic programming algorithm [13]. CART makes a single sweep up the quadtree; at each node, the local cost in error of pruning below this node is weighed against the savings in rate. By passing these decisions upwards through ancestor nodes, we can find the wavelet-wedgeprint tree that globally maximizes (22).

It is worth noting that in the end, the encoder cares little about the actual contour locations in the image. It simply tries to minimize (22) for a given λ. However, the encoder will tend to choose wavelets over smooth regions and wedgeprints over linear contour regions, because the local rate-distortion tradeoff is favorable. The encoder thus naturally adapts to the geometrical structure in the image and does not rely on any pre-processed edge detection.

Figure 4 shows approximations of the "Cameraman" test image using standard wavelet tree pruning and wavelet-wedgeprint tree pruning. Using the wavelet-wedgeprint dictionary, we achieve significant gains both visually (sharper edges and reduced ringing) and in peak signal-to-noise ratio while using 30% fewer terms in the approximation.

We can also tie the orientations of the wedgeprints together along a contour, making them even cheaper to code. A Markov multiscale geometry model [10] can be incorporated into the selection of the representation; the optimal tree pruning with respect to this model is found using the Viterbi algorithm.

4. CONCLUSIONS

We have developed a new framework for "multiscale geometric image processing" that leverages the best of wavelets for representing smooth image regions and wedgelets for representing smooth edge contours. For C^2/C^2 images, the combined wavelet/wedgeprint dictionary provides optimal asymptotic approximation performance, and a simple prototype image coder provides near-optimal rate-distortion performance. While we have used compression to illustrate the effectiveness of our approach, a statistical wavelet/wedgeprint model will enable new algorithms for estimation, detection, classification, segmentation and other statistical image processing tasks. These are currently under investigation.

(a) (b)

Fig. 4. *(a) Wavelet pruning of "Cameraman" (N=7359, PSNR=28.77dB), (b) Wavelet-wedgeprint pruning (N=5158, PSNR=30.27dB).*

5. REFERENCES

[1] S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, San Diego, 2nd edition, 1998.

[2] A. Cohen, W. Dahmen, I. Daubechies, and R. DeVore, "Tree approximation and optimal encoding," *J. Appl. Comp. Harm. Anal.*, vol. 11, pp. 192–226, 2001.

[3] Z. Xiong, K. Ramchandran, and M. T. Orchard, "Space-frequency quantization for wavelet image coding," *IEEE Trans. Image Processing*, vol. 6, no. 5, pp. 677–693, 1997.

[4] A. Cohen, I. Daubechies, O. G. Guleryuz, and M. T. Orchard, "On the importance of combining wavelet-based nonlinear approximation with coding strategies," *IEEE Trans. Inform. Theory*, vol. 48, no. 7, pp. 1895–1921, July 2002.

[5] E. Candés and D. Donoho, "New tight frames of curvelets and optimal representations of objects with C^2 singularities," *Technical Report, Stanford University*, 2002.

[6] E. L. Pennec and S. Mallat, "Image compression with geometrical wavelets," in *IEEE Int. Conf. on Image Proc. – ICIP '01*, Thessaloniki, Greece, Oct. 2001.

[7] R. Shukla, P. L. Dragotti, M. Do, and M. Vetterli, "Rate distortion optimized tree structured compression algorithms," *IEEE Trans. Image Processing*, submitted for publication.

[8] D. Donoho, "Wedgelets: Nearly-minimax estimation of edges," *Annals of Stat.*, vol. 27, pp. 859–897, 1999.

[9] M. N. Do, P. L. Dragotti, R. Shukla, and M. Vetterli, "On the compression of two-dimensional piecewise smooth functions," in *IEEE Int. Conf. on Image Proc. – ICIP '01*, Thessaloniki, Greece, Oct. 2001.

[10] J. K. Romberg, M. B. Wakin, and R. G. Baraniuk, "Wedgelet modeling and algorithms for image processing," *In preparation*, 2003.

[11] P. L. Dragotti and M. Vetterli, "Wavelet footprints: Theory, algorithms, and applications," *IEEE Trans. Signal Processing*, vol. 51, no. 5, May 2003.

[12] M. B. Wakin, J. K. Romberg, H. Choi, and R. G. Baraniuk, "Rate-distortion optimized image compression using wedgelets," in *IEEE Int. Conf. on Image Proc. – ICIP '02*, Rochester, NY, September 2002.

[13] L. Breiman, J. Friedman, R. Olshen, and C. Stone, *Classification and Regression Trees*, Wadsworth, Belmont, CA, 1984.