

Representation and Compression of Multidimensional Piecewise Functions Using *Surflets*

Venkat Chandrasekaran, *Student Member, IEEE*, Michael B. Wakin, *Member, IEEE*, Dror Baron, *Member, IEEE*, and Richard G. Baraniuk, *Fellow, IEEE*

Abstract—We study the representation, approximation, and compression of functions in M dimensions that consist of constant or smooth regions separated by smooth $(M - 1)$ -dimensional discontinuities. Examples include images containing edges, video sequences of moving objects, and seismic data containing geological horizons. For both function classes, we derive the optimal asymptotic approximation and compression rates based on Kolmogorov metric entropy. For piecewise constant functions, we develop a multiresolution predictive coder that achieves the optimal rate–distortion performance; for piecewise smooth functions, our coder has near-optimal rate–distortion performance. Our coder for piecewise constant functions employs *surflets*, a new multiscale geometric tiling consisting of M -dimensional piecewise constant atoms containing polynomial discontinuities. Our coder for piecewise smooth functions uses *surfprints*, which wed *surflets* to wavelets for piecewise smooth approximation. Both of these schemes achieve the optimal asymptotic approximation performance. Key features of our algorithms are that they carefully control the potential growth in surflet parameters at higher smoothness and do not require explicit estimation of the discontinuity. We also extend our results to the corresponding discrete function spaces for sampled data. We provide asymptotic performance results for both discrete function spaces and relate this asymptotic performance to the sampling rate and smoothness orders of the underlying functions and discontinuities. For approximation of discrete data, we propose a new scale-adaptive dictionary that contains few elements at coarse and fine scales, but many elements at medium scales. Simulation results on synthetic signals provide a comparison between surflet-based coders and previously studied approximation schemes based on wedgelets and wavelets.

Index Terms—Compression, discontinuities, metric entropy, multidimensional signals, multiscale representations, nonlinear approximation, rate–distortion, sparse representations, surflets, wavelets.

Manuscript received March 22, 2006; revised April 15, 2008. Current version published December 24, 2008. This work was supported by the NSF under Grant CCF-0431150, ONR under Grant N00014-02-1-0353, AFOSR under Grant FA9550-04-0148, AFRL under Grant FA8650-051850, and the Texas Instruments Leadership University Program. The material in this paper was presented in part the 38th Annual Conference on Information Sciences and Systems Princeton, NJ, March 2004 and the IEEE International Symposium on Information Theory, Chicago, IL, June/July 2004.

V. Chandrasekaran is with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: venkate@mit.edu).

M. B. Wakin is with the Division of Engineering, Colorado School of Mines, Golden, CO 80401 USA (e-mail: mwakin@mines.edu).

D. Baron is with the Department of Electrical Engineering, Technion–Israel Institute of Technology, Technion City, 32000 Haifa, Israel (e-mail: barondror@gmail.com)

R. G. Baraniuk is with the Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005, USA (e-mail: barondror@gmail.com; richb@rice.edu).

Communicated by W. Szpankowski, Associate Editor for Source Coding.

Color versions of Figures 6–8 and 10 in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIT.2008.2008153

I. INTRODUCTION

A. Motivation

Sparse signal representations feature prominently in a broad variety of signal processing applications. The advantages of characterizing signals using just a few elements from a dictionary are clear in the case of compression [3]–[6], but also extend to tasks such as estimation [7]–[10] and classification [11], [12]. The *dimensionality* of signals is an important factor in determining sparse representations. Wavelets provide an appropriate representation for smooth one-dimensional (1-D) signals, achieving the optimal asymptotic representation behavior for this class [13], [14], and they maintain this optimal performance even for smooth 1-D signals containing a finite number of point discontinuities.

Unfortunately, this optimality does not extend completely to two-dimensional (2-D) signals due to the different nature of discontinuities in two dimensions [15], [16]. While smooth signals in two dimensions containing a finite number of *point* singularities are sparsely represented by a wavelet basis, 2-D piecewise smooth signals containing discontinuities along 1-D curves (“edges”) are not represented efficiently by a wavelet basis.

The problem is that the isotropically scaled wavelet basis fails to capture the anisotropic structure of 1-D discontinuities. Due to its global support, the Fourier basis also fails to effectively capture the local nature of these discontinuities. Nonetheless, 2-D piecewise smooth signals containing 1-D discontinuities are worthy candidates for sparse representation. The 1-D discontinuities often carry interesting information, since they signify a boundary between two regions. Edges in images illustrate this point well. Lying along 1-D curves, edges provide fundamental information about the geometry of a scene. Therefore, any signal processing application that relies on a sparse representation requires an efficient tool for representing discontinuities.

A growing awareness of the limitations of traditional bases for representing 2-D signals with 1-D discontinuities has resulted in new multiscale representations. (Multiscale representations offer a number of advantages for signal processing, enabling, for example, predictive and progressive source coding, powerful statistical models, and efficient tree-structured processing.) The resulting solutions fall into two broad categories: tight frames and geometric tilings. Loosely speaking, *tight frames* refer to dictionaries from which approximations are formed using linear combinations of atomic functions, while *geometric tiling* schemes create an adaptive partitioning




Fig. 1. (a) Piecewise constant (“Horizon-class”) functions for dimensions $M = 2$ and $M = 3$. (b) Piecewise smooth function for dimension $M = 2$.

of the signal to which local approximation elements are assigned. The key factor exploited by all these solutions is that the discontinuities often have a coherent geometric structure in one dimension. A primary solution proposed in the class of 2-D tight frames is the curvelet dictionary [16]. The salient feature of curvelets is an anisotropic scaling of the atomic elements, with increased directionality at finer resolution; the resulting dictionary is well-suited to represent 1-D singularities. In the category of geometric tilings, wedgelets [8] are piecewise constant functions defined on 2-D dyadic squares, where a linear edge singularity separates the two constant regions. Tilings of wedgelets at various scales combine to form piecewise linear approximations to the 1-D edges in images. Each of these solutions assumes that the underlying 1-D discontinuity is C^2 -smooth. The class of 2-D signals containing such smooth 1-D discontinuities is often referred to as the *Horizon* function model [8].

Unfortunately, the tight frames and geometric tilings proposed to date also face certain limitations. First, none of these techniques is directly applicable to problems in *arbitrary* higher dimensions; for example, curvelets are not applicable to problems with dimensions beyond 3-D [17], while wedgelets are not applicable to problems with dimensions greater than two. Signals in video (3-D), geophysics (3-D, 4-D), and light-field imaging (4-D, 5-D) [18] frequently contain information-rich discontinuities separating different regions. The confounding aspect, as before, is that these discontinuities often have a well-defined geometric structure in one lower dimension than that of the signal. In video signals, for example, where we can imagine stacking a 2-D image sequence to form a 3-D volume, 3-D regions are separated by smooth, well-structured 2-D discontinuities traced out by the moving boundaries of objects in the video. Therefore, we can model discontinuities in M -dimensional signals as smooth functions of $M - 1$ variables. This is an analogous extension to the *Horizon* function model [8] used for 2-D signals with 1-D smooth discontinuities.

Another limitation of current frame and tiling approaches is that they are intended primarily for signals with underlying C^2 -smooth discontinuities. There exist signals, however, for which the discontinuities are inherently of a higher order of smoothness [19]. While $C^K \subset C^2$ for $K > 2$, dictionaries achieving the optimal performance can be constructed only if the full smoothness of the discontinuities is exploited. Interesting mathematical insights can also be obtained by

considering higher orders of smoothness. Finally, some of the proposed solutions such as wedgelets [8] consider representation of piecewise constant signals in higher dimensions. Real-world multidimensional signals, however, often consist of discontinuities separating smooth (but not constant) regions. This motivates a search for sparse representations for *signals in any dimension* consisting of *regions of arbitrary smoothness* that are separated by *discontinuities in one lower dimension of arbitrary smoothness*.

B. Approximation and Compression

In this paper, we address the problem of approximating and compressing M -dimensional signals that contain a smooth $(M - 1)$ -dimensional discontinuity separating regions in M dimensions¹ (see Fig. 1 for examples in two and three dimensions). The discontinuities in our models have smoothness C^{K_a} in $M - 1$ dimensions. They separate two regions that may be constant (Section I-C1) or C^{K_s} -smooth in M dimensions (Section I-C2). Our *approximation results* characterize the number of atoms required to efficiently represent a signal and are useful for tasks such as estimation [8], [10] and classification [11], [12]. The measure for approximation performance is the asymptotic rate of decay of the distortion between a signal and its approximant as the number of atoms becomes large. *Compression results*, on the other hand, also take into account *which* atoms are used in constructing a sparse approximation and are crucial for communicating an approximation to a signal. Compression performance is assessed by considering the asymptotic rate–distortion behavior, which specifies the rate of decay of the distortion between a signal and its approximant as the number of bits used to encode the approximant becomes large. Since it is impossible to determine for abstract function classes precise approximation or rate–distortion results, we will state our performance results in terms of *metric entropy* [20], which characterizes the order of the number of atoms or bits required to describe a function in the class under consideration.

C. Contributions

We consider two function classes as models for M -dimensional signals containing $(M - 1)$ -dimensional discontinuities.

¹We discuss possible extensions to signals containing multiple discontinuities in Section IV-F.




Fig. 2. Example surflets, designed for (a) $M = 2$, $K_d \in (1, 2]$; (b) $M = 2$, $K_d \in (2, 3]$; (c) $M = 3$, $K_d \in (1, 2]$; (d) $M = 3$, $K_d \in (2, 3]$.

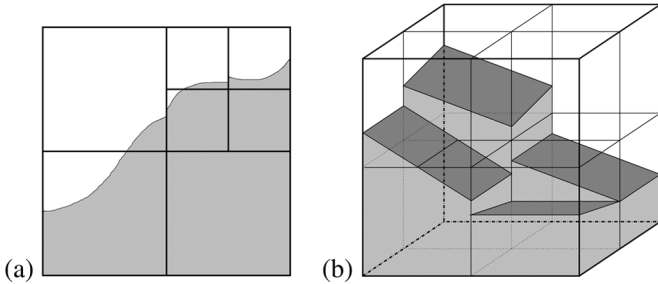


Fig. 3. Example surflet tilings: (a) piecewise cubic with $M = 2$ and (b) piecewise linear with $M = 3$.

This section outlines our contributions toward constructing efficient representations for elements of these classes, including extensions to discrete data.

1) *Piecewise Constant M -Dimensional Functions:* We consider the function class $\mathcal{F}_C(M, K_d)$, consisting of M -dimensional piecewise constant Horizon functions [8] that contain a \mathcal{C}^{K_d} -smooth $(M - 1)$ -dimensional discontinuity separating two *constant* regions. We begin by deriving the optimal asymptotic rates for nonlinear approximation and compression of functions in the class $\mathcal{F}_C(M, K_d)$. These bounds extend the results of Cohen *et al.* [14], Clements [21], and Kolmogorov and Tihomirov [20], which characterize the optimal asymptotic approximation and compression behavior of $(M - 1)$ -dimensional smooth functions.

We introduce a new M -dimensional geometric tiling framework for multiscale representation of functions in $\mathcal{F}_C(M, K_d)$. We propose a dictionary of atoms defined on dyadic hypercubes, where each atom is an M -dimensional piecewise constant function containing an $(M - 1)$ -dimensional polynomial discontinuity. We christen these atoms *surflets* after the fact that they resemble small pieces of surfaces in higher dimensional space (see Fig. 2 for examples in two and three dimensions).

The surflet dictionary is a generalization of the wedgelet dictionary [8] to higher dimensional signals containing discontinuities of arbitrary smoothness (a wedgelet is a surflet with $K_d = 2$ and $M = 2$). We show that tilings of elements drawn from the surflet dictionary can achieve the optimal approximation rate

$$\|f^c - \widehat{f}_N^c\|_{L_2}^2 \approx \left(\frac{1}{N}\right)^{\frac{K_d}{M-1}}$$

for the class $\mathcal{F}_C(M, K_d)$,² where $f^c \in \mathcal{F}_C(M, K_d)$ and \widehat{f}_N^c is the best N -term approximant to f^c . Example 2- and 3-D surflet tilings appear in Fig. 3.

We also propose a tree-structured compression algorithm for M -dimensional functions using surflets and establish that this algorithm achieves the optimal rate–distortion performance

$$\|f^c - \widehat{f}_R^c\|_{L_2}^2 \approx \left(\frac{1}{R}\right)^{\frac{K_d}{M-1}}$$

for the class $\mathcal{F}_C(M, K_d)$. Here, \widehat{f}_R^c is the best approximation to f^c that can be encoded using R bits.

Our approach incorporates the following major features:

- the ability to operate directly on the M -dimensional function, without requiring explicit knowledge (or estimation) of the $(M - 1)$ -dimensional discontinuity;
- the use of multiresolution predictive coding in order to realize significant gains in rate–distortion performance; and
- a technique to quantize and encode higher order polynomial coefficients with lesser precision without a substantial increase in distortion.

Without such a quantization scheme, higher order polynomials would be impractical for representing boundaries smoother than \mathcal{C}^2 , due to an exponential explosion in the number of polynomial parameters and thus the size of the dictionary. A fascinating aspect of our solution is that the size of the surflet dictionary can be reduced tremendously without sacrificing the approximation capability.

2) *Piecewise Smooth M -Dimensional Functions:* We consider the function class $\mathcal{F}_S(M, K_d, K_s)$, consisting of M -dimensional piecewise *smooth* functions that contain a \mathcal{C}^{K_d} -smooth $(M - 1)$ -dimensional discontinuity separating two regions that are \mathcal{C}^{K_s} -smooth in M dimensions. We establish the optimal approximation and rate–distortion bounds for this class.

Despite their ability to efficiently describe a discontinuity, surflets alone are not a suitable representation for functions in $\mathcal{F}_S(M, K_d, K_s)$ because a piecewise *constant* approximation is inefficient on the smooth M -dimensional regions. Conversely, we note that wavelets are well suited for representing these M -dimensional smooth regions, yet they fail to efficiently capture discontinuities. Thus, we propose a dictionary consisting of the following elements:

²We focus here on *asymptotic* performance. We use the notation $f(\alpha) \lesssim g(\alpha)$, or $f(\alpha) = O(g(\alpha))$, if there exists a constant C , possibly large but not dependent on the argument α , such that $f(\alpha) \leq Cg(\alpha)$.




Fig. 4. Comparison of pruned surflet tilings using two surflet dictionaries (see Section VI for more details). (a) Test image with $M = 2$ and $K_d = 3$. (b) The wedgelets from Dictionary 2 can be encoded using 482 bits and yields PSNR of 29.86 dB. (c) The quadratic/wedgelet combination from Dictionary 3 can be encoded using only 275 bits and yields PSNR of 30.19 dB.

- an M -dimensional basis of compactly supported *wavelet* atoms having sufficiently many vanishing moments [22]; and
- a dictionary of M -dimensional *surfprint* atoms. Each surfprint can be viewed as a linear combination of wavelet basis functions. It is derived by projecting a piecewise *polynomial* surflet atom (having two M -dimensional polynomial regions separated by a polynomial discontinuity) onto a wavelet subspace.

The surfprint dictionary is an extension of the wedgeprint dictionary [23] for higher dimensional signals having regions of arbitrary smoothness separated by discontinuities of arbitrary smoothness (a wedgeprint is a surfprint with $K_d = 2$, $K_s = 2$, and $M = 2$). Constructed as projections of piecewise polynomial surflets onto wavelet subspaces, surfprints interface naturally with wavelets for sparse approximation of piecewise smooth M -dimensional functions. We show that the combined wavelet/surfprint dictionary achieves the optimal approximation rate for $\mathcal{F}_S(M, K_d, K_s)$

$$\|f^s - \widehat{f}_N^s\|_{L_2}^2 \lesssim \left(\frac{1}{N}\right)^{\min\left(\frac{K_d}{M-1}, \frac{2K_s}{M}\right)}.$$

Here, $f^s \in \mathcal{F}_S(M, K_d, K_s)$, and \widehat{f}_N^s is the best N -term approximant to f^s . We include a careful treatment of the surfprint polynomial degrees and the number of wavelet vanishing moments required to attain the optimal approximation rates. We also propose a tree-based encoding scheme that comes within a logarithmic factor of the optimal rate–distortion performance for this class

$$\|f^s - \widehat{f}_R^s\|_{L_2}^2 \lesssim \left(\frac{\log R}{R}\right)^{\min\left(\frac{K_d}{M-1}, \frac{2K_s}{M}\right)}.$$

Here, \widehat{f}_R^s is the best approximation to f^s that can be encoded using R bits.

3) *Extensions to Discrete Data:* We also address the problem of representing discrete data obtained by sampling a continuous function from $\mathcal{F}_C(M, K_d)$ or $\mathcal{F}_S(M, K_d, K_s)$. We denote these classes of discretized (or “voxelized”) data by $\widetilde{\mathcal{F}}_C(M, K_d)$ and $\widetilde{\mathcal{F}}_S(M, K_d, K_s)$, respectively, and we allow for different sampling rates in each of the M dimensions.

In order to efficiently represent data from $\widetilde{\mathcal{F}}_C(M, K_d)$, we use a dictionary of discrete surflet atoms derived by voxelizing the surflet atoms of our continuous surflet dictionary. We show that, up to a certain critical scale (which depends on K_d and the sampling rates in each dimension), the approximation results for $\mathcal{F}_C(M, K_d)$ extend to $\widetilde{\mathcal{F}}_C(M, K_d)$. Beyond this scale, however, voxelization effects dominate, and the surflets designed for elements of $\mathcal{F}_C(M, K_d)$ offer unnecessary precision for representing elements of $\widetilde{\mathcal{F}}_C(M, K_d)$. To account for these effects, we propose a new scale-adaptive dictionary that contains few elements at coarse and fine scales, but many elements at medium scales. We also present preliminary simulation results that demonstrate the compression performance of our surflet coding strategies when applied to synthetic signals; see Fig. 4 for an example illustrating the use of quadratic surflets.³

For elements of $\widetilde{\mathcal{F}}_S(M, K_d, K_s)$, we use a dictionary of compactly supported discrete wavelet basis functions. The number of discrete vanishing moments required is the same as that of the continuous wavelet basis applied to members of $\mathcal{F}_S(M, K_d, K_s)$. Discrete surfprint atoms are obtained by projecting discrete surflet atoms onto this discrete wavelet basis. As before, we see a critical scale at which voxelization effects begin to dominate; again these effects can be addressed by the appropriate scale-adaptive surfprint dictionary.

D. Relation to Previous Work

Our work can be viewed as a generalization of wedgelet [8] and wedgeprint [8] representations. Our extensions, however, provide fundamental new insights in the following directions.

- The wedgelet and wedgeprint dictionaries are restricted to 2-D signals, while our proposed representations are relevant in higher dimensions.
- Wedgelets and wedgeprints achieve optimal approximation rates only for functions that are \mathcal{C}^2 -smooth and contain a \mathcal{C}^2 -smooth discontinuity; our results not only show that surflets and surfprints can be used to achieve optimal rates for more general classes, but also highlight the necessary polynomial quantization scheme (a nontrivial extension from wedgelets).

³The peak signal-to-noise ratio (PSNR) we quote is a common measure of distortion that derives from the mean-square error (MSE); assuming a maximum possible signal intensity of I , $\text{PSNR} := 10 \log_{10} \frac{I^2}{\text{MSE}}$.

- In the construction of surfprint atoms, we derive the surfprint polynomial degrees required for optimal performance as a function of K_s , K_d , and M . Such insight cannot be gained from wedgeprints, which are derived simply from the projection of piecewise constant wedgelets.
- We also present a more thorough analysis of discretization effects, including new insights on the multiscale behavior (not revealed by considering wedgelets alone), a new strategy for reducing the surflet dictionary size at fine scales, and the first treatment of wedgeprint/surfprint discretization.

During final preparation of this manuscript we learned of a related generalization of wedgelets to quadratic polynomials [24], and of an extension of directional filter banks to multidimensional signals [25]. Novel components of our work, however, include its broad generality to arbitrary orders of smoothness and arbitrary dimension; the quantization scheme, predictive coding, and analytical results; the scale-adaptive dictionary to deal with voxelization; and the surfprints construction.

E. Paper Organization

In Section II, we define our function models and state the specific goals of our approximation and compression algorithms. We introduce surflets in Section III. In Section IV, we describe our surflet-based representation schemes for elements of $\mathcal{F}_C(M, K_d)$ in detail. In Section V, we present our novel dictionary of wavelets and surfprints for effectively representing elements of $\mathcal{F}_S(M, K_d, K_s)$. Section VI discusses extensions to discrete data from $\mathcal{F}_C(M, K_d)$ and $\widetilde{\mathcal{F}}_S(M, K_d, K_s)$. Section VII summarizes our contributions and insights. The appendices provide additional details and proofs of all the theorems.

II. BACKGROUND

A. Lipschitz Smoothness

A function of D variables has smoothness of order $K > 0$, where $K = r + \alpha$, r is an integer, and $\alpha \in (0, 1]$, if the following criteria are met [20], [21].

- All iterated partial derivatives with respect to the D directions up to order r exist and are continuous.
- All such partial derivatives of order r satisfy a Lipschitz condition of order α (also known as a Hölder condition).⁴

We consider the space of smooth functions whose partial derivatives up to order r are bounded by some constant Ω . We denote the space of such bounded functions with bounded partial derivatives by \mathcal{C}^K , where this notation carries an implicit dependence on Ω . Observe that $r = \lceil K - 1 \rceil$, where $\lceil \cdot \rceil$ denotes rounding up. Also, when K is an integer \mathcal{C}^K includes as a subset the traditional space “ \mathcal{C}^K ” (the class of functions that have $K = r + 1$ continuous partial derivatives).

⁴A function $d \in \text{Lip}(\alpha)$ if $|d(\mathbf{z}_1 + \mathbf{z}_2) - d(\mathbf{z}_1)| \leq C\|\mathbf{z}_2\|^\alpha$ for all D -dimensional vectors $\mathbf{z}_1, \mathbf{z}_2$.

B. Multidimensional Signal Model

Let $\mathbf{x} \in [0, 1]^M$, and let x_i denote its i th element, where boldface characters are used to denote vectors. We denote the first $M - 1$ elements of \mathbf{x} by \mathbf{y} , i.e., $\mathbf{y} = [x_1, x_2, \dots, x_{M-1}] \in [0, 1]^{M-1}$. Let f, g_1 , and g_2 be functions of M variables

$$f, g_1, g_2 : [0, 1]^M \rightarrow \mathbb{R}$$

and let b be a function of $M - 1$ variables

$$b : [0, 1]^{M-1} \rightarrow \mathbb{R}.$$

We define the function f in the following piecewise manner:

$$f(\mathbf{x}) = \begin{cases} g_1(\mathbf{x}), & x_M \geq b(\mathbf{y}) \\ g_2(\mathbf{x}), & x_M < b(\mathbf{y}). \end{cases}$$

1) *Piecewise Constant Model:* We begin by considering the “piecewise constant” case when $g_1 = 1$ and $g_2 = 0$. The $(M - 1)$ -dimensional discontinuity b defines a boundary between two constant regions in M dimensions. The piecewise constant functions f defined in this manner are Horizon-class functions [8]. When $b \in \mathcal{C}^{K_d}$, with $K_d = r_d + \alpha_d$, we denote the resulting space of functions f by $\mathcal{F}_C(M, K_d)$. When $M = 2$, these functions can be interpreted as images containing a \mathcal{C}^{K_d} -smooth 1-D discontinuity that separates a 0-valued region below from a 1-valued region above. For $M = 3$, functions in $\mathcal{F}_C(M, K_d)$ can be interpreted as cubes with a 2-D \mathcal{C}^{K_d} -smooth surface cutting through them, dividing them into two regions—0-valued below the surface and 1-valued above it (see Fig. 1(a) for examples in two and three dimensions).

2) *Piecewise Smooth Model:* Next, we define a model for piecewise smooth functions. For this class of functions, we let $g_1, g_2 \in \mathcal{C}^{K_s}$, with $K_s = r_s + \alpha_s$, and $b \in \mathcal{C}^{K_d}$, with $K_d = r_d + \alpha_d$. The resulting piecewise smooth function f consists of an $(M - 1)$ -dimensional \mathcal{C}^{K_d} -smooth discontinuity that separates two \mathcal{C}^{K_s} -smooth regions in M dimensions (see Fig. 1(b) for an example in two dimensions). We denote the class of such piecewise smooth functions by $\mathcal{F}_S(M, K_d, K_s)$. One can check that both $\mathcal{F}_C(M, K_d)$ and the space of M -dimensional uniformly \mathcal{C}^{K_s} functions are subsets of $\mathcal{F}_S(M, K_d, K_s)$.

C. Approximation and Compression Performance Measures

In this paper, we define dictionaries of atoms from which we construct an approximation \hat{f} to f , which may belong to $\mathcal{F}_C(M, K_d)$ or $\mathcal{F}_S(M, K_d, K_s)$. We analyze the performance of our coding scheme using the L_p distortion measure between the M -dimensional functions f and \hat{f}

$$\|f - \hat{f}\|_{L_p([0,1]^M)}^p = \int_{[0,1]^M} |f - \hat{f}|^p$$

with $p = 2$ (i.e., the standard squared- L_2 distortion measure). We measure the ability of our dictionary of atoms to represent f sparsely by the *asymptotic approximation performance* $\|f - \hat{f}_N\|_{L_2}^2$. Here, N is the number of atomic elements from the dictionary used to construct an approximation to f and \hat{f}_N is the best N -term approximant to f . We consider the rate of decay of $\|f - \hat{f}_N\|_{L_2}^2$ as $N \rightarrow \infty$.

We also present compression algorithms that encode those atoms from the corresponding dictionaries (depending on whether $f \in \mathcal{F}_C(M, K_d)$ or $f \in \mathcal{F}_S(M, K_d, K_s)$) used to construct \hat{f} . We measure the performance of these compression algorithms by the *asymptotic rate–distortion function* $\|f - \hat{f}_R\|_{L_2}^2$, where \hat{f}_R is the best approximation to f that can be encoded using R bits [26]. We consider the behavior of $\|f - \hat{f}_R\|_{L_2}^2$ as $R \rightarrow \infty$.

In [14], Cohen *et al.* establish the optimal approximation rate for D -dimensional C^K -smooth functions d

$$\|d - \widehat{d}_N\|_{L_2}^2 \lesssim \left(\frac{1}{N}\right)^{\frac{2K}{D}}.$$

Similarly, the results of Clements [21] (extending those of Kolmogorov and Tihomirov [20]) regarding metric entropy establish bounds on the optimal achievable asymptotic rate–distortion performance for D -dimensional C^K -smooth functions d

$$\|d - \widehat{d}_R\|_{L_2}^2 \lesssim \left(\frac{1}{R}\right)^{\frac{2K}{D}}.$$

These results, however, are only useful for characterizing optimal *separate* representations for the $(M-1)$ -dimensional discontinuity ($D = M-1, K = K_d, d = b$ in Section II-A) and the M -dimensional smooth regions ($D = M, K = K_s, d = g_1, g_2$ in Section II-A).

We extend these results to nonseparable representations of the M -dimensional function classes $\mathcal{F}_C(M, K_d)$ and $\mathcal{F}_S(M, K_d, K_s)$ in Theorems 1 and 2, respectively.

Theorem 1: The optimal asymptotic approximation performance that can be obtained for all $f^c \in \mathcal{F}_C(M, K_d)$ is given by

$$\|f^c - \widehat{f}_N^c\|_{L_2}^2 \lesssim \left(\frac{1}{N}\right)^{\frac{K_d}{M-1}}.$$

Similarly, the optimal asymptotic compression performance that can be obtained for all $f^c \in \mathcal{F}_C(M, K_d)$ is given by

$$\|f^c - \widehat{f}_R^c\|_{L_2}^2 \lesssim \left(\frac{1}{R}\right)^{\frac{K_d}{M-1}}.$$

Implicit in the proof of the above theorem, which appears in Appendix A, is that any scheme that is optimal for representing and compressing the M -dimensional function $f^c \in \mathcal{F}_C(M, K_d)$ in the squared- L_2 sense is equivalently optimal for the $(M-1)$ -dimensional discontinuity in the L_1 sense. Roughly, the squared- L_2 distance between two Horizon-class functions f_1^c and f_2^c over an M -dimensional domain

$$\mathcal{D} = [\mathcal{D}_b^1, \mathcal{D}_e^1] \times \cdots \times [\mathcal{D}_b^M, \mathcal{D}_e^M]$$

is equal to the L_1 distance over the $(M-1)$ -dimensional subdomain $[\mathcal{D}_b^1, \mathcal{D}_e^1] \times \cdots \times [\mathcal{D}_b^{M-1}, \mathcal{D}_e^{M-1}]$ between the $(M-1)$ -dimensional discontinuities b_1^c and b_2^c in f_1^c and f_2^c , respectively.

More precisely and for future reference, for every \mathbf{y} in the $(M-1)$ -dimensional subdomain of \mathcal{D} , we define the \mathcal{D} -clipping of an $(M-1)$ -dimensional function b as

$$\bar{b}(\mathbf{y}) = \begin{cases} b(\mathbf{y}), & \mathcal{D}_b^M \leq b(\mathbf{y}) \leq \mathcal{D}_e^M \\ \mathcal{D}_e^M, & b(\mathbf{y}) > \mathcal{D}_e^M \\ \mathcal{D}_b^M, & b(\mathbf{y}) < \mathcal{D}_b^M. \end{cases}$$

The \mathcal{D} -active region of b is defined to be

$$\{\mathbf{y} \in [\mathcal{D}_b^1, \mathcal{D}_e^1] \times \cdots \times [\mathcal{D}_b^{M-1}, \mathcal{D}_e^{M-1}] : b(\mathbf{y}) \in [\mathcal{D}_b^M, \mathcal{D}_e^M]\}$$

that subset of the subdomain of \mathcal{D} for which the range of b lies in $[\mathcal{D}_b^M, \mathcal{D}_e^M]$. The \mathcal{D} -clipped L_1 distance between b_1^c and b_2^c is then defined as

$$\overline{L}_1(b_1^c, b_2^c) = \|\bar{b}_1^c - \bar{b}_2^c\|_{L_1([\mathcal{D}_b^1, \mathcal{D}_e^1] \times \cdots \times [\mathcal{D}_b^{M-1}, \mathcal{D}_e^{M-1}])}.$$

One can check that $\|f_1^c - f_2^c\|_{L_2(\mathcal{D})}^2 = \overline{L}_1(b_1^c, b_2^c)$ for any \mathcal{D} .

The following theorem, which is proved in Appendix B, characterizes the optimal achievable asymptotic approximation rate and rate–distortion performance for approximating and encoding elements of the function class $\mathcal{F}_S(M, K_d, K_s)$.

Theorem 2: The optimal asymptotic approximation performance that can be obtained for all $f^s \in \mathcal{F}_S(M, K_d, K_s)$ is given by

$$\|f^s - \widehat{f}_N^s\|_{L_2}^2 \lesssim \left(\frac{1}{N}\right)^{\min\left(\frac{K_d}{M-1}, \frac{2K_s}{M}\right)}.$$

Similarly, the optimal asymptotic compression performance that can be obtained for all $f^s \in \mathcal{F}_S(M, K_d, K_s)$ is given by

$$\|f^s - \widehat{f}_R^s\|_{L_2}^2 \lesssim \left(\frac{1}{R}\right)^{\min\left(\frac{K_d}{M-1}, \frac{2K_s}{M}\right)}.$$

D. “Oracle” Coders and Their Limitations

In order to approximate or compress an arbitrary function $f^c \in \mathcal{F}_C(M, K_d)$, an algorithm is given the function f^c ; we denote its $(M-1)$ -dimensional C_d^K -smooth discontinuity by b^c . As constructed in Section II-B, all of the critical information about f^c is contained in the discontinuity b^c . One would expect any efficient coder to exploit such a fact; methods through which this is achieved may vary.

One can imagine a coder that *explicitly* encodes an approximation \widehat{b}^c to b^c and then constructs a Horizon approximation \widehat{f}^c . Knowledge of b^c could be provided from an external “oracle” [27], or b^c could conceivably be estimated from the provided data f^c . Wavelets provide an efficient method for compressing the $(M-1)$ -dimensional smooth function b^c . Cohen *et al.* [14] describe a tree-structured wavelet coder that can be used to compress b^c with optimal rate–distortion performance in the L_1 sense. It follows that this wavelet coder is optimal (in the squared- L_2 sense) for coding instances of f^c at the optimal rate of Theorem 1. In practice, however, a coder is not provided with explicit information of b^c , and a method for estimating b^c from f^c may be difficult to implement. Estimates for b^c may also be quite sensitive to noise in the data.

A similar strategy could also be employed for $f^s \in \mathcal{F}_S(M, K_d, K_s)$. For such a function, we denote the $(M-1)$ -dimensional C_d^K -smooth discontinuity by b^s and the M -dimensional $C_s^{K_s}$ -smooth regions by g_1^s and g_2^s . Approximations to the discontinuity \widehat{b}^s and the M -dimensional smooth regions \widehat{g}_1^s and \widehat{g}_2^s may be encoded separately and explicitly. This strategy would have disadvantages for the same reasons that were mentioned above. In fact, estimating the discontinuity in this scenario would be much harder.

In this paper, we propose representation schemes and algorithms that approximate f^c and f^s directly in M dimensions. We emphasize that no explicit knowledge of the functions b^c ,

b^s , g_1^s , or g_2^s is required. We prove that surflet-based approximation techniques and encoding algorithms for f^c achieve the optimal decay rates, while our surfprint-based methods for f^s achieve the optimal approximation decay rate and a near-optimal rate–distortion decay rate (within a logarithmic factor of the optimal decay rate of Theorem 2). Although we omit the discussion in this paper, our algorithms can be extended to similar piecewise constant and piecewise smooth function spaces. Our spatially localized approach, for example, allows for changes in the variable along which the discontinuity varies (assumed throughout this paper to be x_M as described in Section II-B).

III. THE SURFLET DICTIONARY

In this section, we introduce a discrete dictionary of M -dimensional atoms called *surflets* that can be used to construct approximations to a function $f^c \in \mathcal{F}_C(M, K_d)$. A surflet is a piecewise constant function defined on an M -dimensional dyadic hypercube, where an $(M-1)$ -dimensional polynomial specifies the discontinuity. Section IV describes compression using surflets.

A. Motivation—Taylor’s Theorem

The surflet atoms are motivated by the following property. If d is a function of D variables in C^K with $K = r + \alpha$, r is a positive integer, and $\alpha \in (0, 1]$, then Taylor’s theorem states that

$$\begin{aligned} d(\mathbf{z} + \mathbf{h}) &= d(\mathbf{z}) + \frac{1}{1!} \sum_{i_1=1}^D d_{z_{i_1}}(\mathbf{z}) h_{i_1} \\ &+ \frac{1}{2!} \sum_{i_1, i_2=1}^D d_{z_{i_1}, z_{i_2}}(\mathbf{z}) h_{i_1} h_{i_2} + \cdots \\ &+ \frac{1}{r!} \sum_{i_1, \dots, i_r=1}^D d_{z_{i_1}, \dots, z_{i_r}}(\mathbf{z}) h_{i_1} \cdots h_{i_r} + O(\|\mathbf{h}\|^K) \end{aligned} \quad (1)$$

where d_{z_1, \dots, z_ℓ} refers to the iterated partial derivatives of d with respect to z_1, \dots, z_ℓ in that order. (Note that there are D^ℓ ℓ th order derivative terms.) Thus, over a small domain, the function d is well approximated using a polynomial of order r (where the polynomial coefficients correspond to the partial derivatives of d evaluated at \mathbf{z}).

Clearly, in the case of f^c , one method for approximating the discontinuity b^c would be to assemble a *piecewise polynomial* approximation, where each polynomial is derived from the local Taylor approximation of b^c (let $D = M - 1$, $K = K_d$, and $d = b^c$ in the above characterization). These piecewise polynomials can be used to assemble a Horizon-class approximation of the function f^c . Surflets provide the M -dimensional framework for constructing such approximations and can be implemented without explicit knowledge of b^c or its derivatives.

B. Definition

A *dyadic hypercube* $X_j \subseteq [0, 1]^M$ at scale $j \in \mathbb{N}$ is a domain that satisfies⁵

$$X_j = [\beta_1 2^{-j}, (\beta_1 + 1) 2^{-j}] \times \cdots \times [\beta_M 2^{-j}, (\beta_M + 1) 2^{-j}]$$

⁵To cover the entire domain $[0, 1]^M$, in the case where $(\beta_i + 1) 2^{-j} = 1$, $i \in \{1, \dots, M\}$, we replace the half-open interval $[\beta_i 2^{-j}, (\beta_i + 1) 2^{-j}]$ with the closed interval $[\beta_i 2^{-j}, (\beta_i + 1) 2^{-j}]$.

with $\beta_1, \beta_2, \dots, \beta_M \in \{0, 1, \dots, 2^j - 1\}$. We explicitly denote the $(M-1)$ -dimensional hypercube *subdomain* of X_j as

$$Y_j = [\beta_1 2^{-j}, (\beta_1 + 1) 2^{-j}] \times \cdots \times [\beta_{M-1} 2^{-j}, (\beta_{M-1} + 1) 2^{-j}]. \quad (2)$$

The *surflet* $s(X_j; p; \cdot)$ is a Horizon-class function over the dyadic hypercube X_j defined through the $(M-1)$ -dimensional polynomial p . For $\mathbf{x} \in X_j$ with corresponding $\mathbf{y} = [x_1, x_2, \dots, x_{M-1}]$

$$s(X_j; p; \mathbf{x}) = \begin{cases} 1, & x_M \geq p(\mathbf{y}) \\ 0, & \text{otherwise} \end{cases}$$

where the polynomial $p(\mathbf{y})$ is defined as

$$\begin{aligned} p(\mathbf{y}) &= p_0 + \sum_{i_1=1}^{M-1} p_{1, i_1} y_{i_1} \\ &+ \sum_{i_1, i_2=1}^{M-1} p_{2, i_1, i_2} y_{i_1} y_{i_2} + \cdots \\ &+ \sum_{i_1, \dots, i_{r_d}=1}^{M-1} p_{r_d, i_1, i_2, \dots, i_{r_d}} y_{i_1} y_{i_2} \cdots y_{i_{r_d}}. \end{aligned}$$

We call the polynomial coefficients $\{p_{\ell, i_1, \dots, i_\ell}\}_{\ell=0}^{r_d}$ the *surflet coefficients*.⁶ We note here that, in some cases, a surflet may be identically 0 or 1 over the entire domain X_j . We sometimes denote a generic surflet by $s(X_j)$, indicating only its region of support.

A surflet $s(X_j)$ approximates the function f^c over the dyadic hypercube X_j . One can cover the entire domain $[0, 1]^M$ with a collection of dyadic hypercubes (possibly at different scales) and use surflets to approximate f^c over each of these smaller domains. For $M = 3$, these surflets tiled together look like piecewise polynomial “surfaces” approximating the discontinuity b^c in the function f^c . Fig. 2 illustrates a collection of surflets with $M = 2$ and $M = 3$.

C. Quantization

We obtain a discrete surflet dictionary $\mathcal{M}(j)$ at scale j by quantizing the set of allowable surflet polynomial coefficients. For $\ell \in \{0, 1, \dots, r_d\}$, the surflet coefficient $p_{\ell, i_1, \dots, i_\ell}$ at scale $j \in \mathbb{N}$ is restricted to values $\{\mu \cdot \Delta_{\ell, j}^{K_d}\}_{\mu \in \mathbb{Z}}$, where the stepsize satisfies

$$\Delta_{\ell, j}^{K_d} = 2^{-(K_d - \ell)j}. \quad (3)$$

The necessary range for μ will depend on the derivative bound Ω (Section II-A). We emphasize that the relevant discrete surflet dictionary $\mathcal{M}(j)$ is *finite* at every scale j .

These quantization stepsizes are carefully chosen to ensure the proper fidelity of surflet approximations without requiring excess bit rate. The key idea is that *higher order terms can be quantized with lesser precision* without increasing the residual error term in the Taylor approximation (1). In fact, Kolmogorov

⁶Because the ordering of the terms $y_{i_1} y_{i_2} \cdots y_{i_\ell}$ in a monomial is irrelevant, only $\binom{\ell + M - 2}{\ell}$ monomial coefficients (not $(M-1)^\ell$) need to be encoded for order ℓ . We preserve the slightly redundant notation for ease of comparison with (1).

and Tihomirov [20] implicitly used this concept to establish the metric entropy for bounded uniformly smooth functions.

IV. REPRESENTATION AND CODING OF PIECEWISE *Constant* FUNCTIONS

A. Overview

We now propose a surflet-based multiresolution geometric tiling approach to approximate and encode an arbitrary function $f^c \in \mathcal{F}_C(M, K_d)$. The tiling is arranged on a 2^M -tree, where each node in the tree at scale j corresponds to a hypercube of side length 2^{-j} . Each node is labeled with a surflet appropriately chosen from $\mathcal{M}(j)$ and is either a leaf node (hypercube) or has 2^M children nodes (children hypercubes that perfectly tile the volume of the parent hypercube). Leaf nodes provide the actual approximation to the function f^c , while interior nodes are useful for predicting and encoding their descendants. This framework enables an *adaptive, multiscale* approximation of f^c —many small surflets can be used at fine scales for complicated regions, while few large surflets will suffice to encode simple regions of f^c (such as those containing all 0 or 1). Fig. 3 shows surflet tiling approximations for $M = 2$ and $M = 3$.

Section IV-B discusses techniques for determining the proper surflet at each node. Section IV-C describes a constructive algorithm for building tree-based surflet approximations. Section IV-D describes the performance of a simple surflet encoder acting only on the leaf nodes. Section IV-E presents a more advanced surflet coder, using a top-down predictive technique to exploit the correlation among surflet coefficients. Finally, Section IV-F discusses extensions of our surflet-based representation schemes to broader function classes.

B. Surflet Selection

Consider a node at scale j that corresponds to a dyadic hypercube X_j , and let Y_j be the $(M - 1)$ -dimensional subdomain of X_j as defined in (2).

We first examine a situation where the coder is provided with explicit information about the discontinuity b^c and its derivatives. In this case, determination of the surflet at the node that corresponds to X_j can proceed as implied by Section III. The coder constructs the Taylor expansion of b^c around any point $\mathbf{y} \in Y_j$ and quantizes the polynomial coefficients (3). We choose

$$\mathbf{y}_{ep} = \left[\left(\beta_1 + \frac{1}{2} \right) 2^{-j}, \left(\beta_2 + \frac{1}{2} \right) 2^{-j}, \dots, \left(\beta_{M-1} + \frac{1}{2} \right) 2^{-j} \right]$$

and call this an *expansion point*. We refer to the resulting surflet as the *quantized Taylor surflet*. From (1), it follows that the squared- L_2 error between f^c and the quantized Taylor surflet approximation $s(X_j)$ (which equals the X_j -clipped L_1 error between b^c and the polynomial defining $s(X_j)$) obeys

$$\begin{aligned} \|f^c - s(X_j)\|_{L_2(X_j)}^2 &= \int_{X_j} (f^c - s(X_j))^2 \\ &= O\left(2^{-j(K_d+M-1)}\right). \end{aligned} \quad (4)$$

However, as discussed in Section II-D, our coder is not provided with explicit information about b^c . Therefore, approximating functions in $\mathcal{F}_C(M, K_d)$ using Taylor surflets is imprac-

tical.⁷ We now define a technique for obtaining a surflet estimate directly from the function f^c . We assume that there exists a method to compute the squared- L_2 error $\|f^c - s(X_j)\|_{L_2(X_j)}^2$ between a given surflet $s(X_j)$ and the function f^c on the dyadic block X_j . In such a case, we can search the finite surflet dictionary $\mathcal{M}(j)$ for the minimizer of this error *without explicit knowledge of b^c* . We refer to the resulting surflet as the *native L_2 -best surflet*. This surflet will necessarily obey (4) as well. Section IV-D discusses the coding implications of using L_2 -best surflets from $\mathcal{M}(j)$. Using native L_2 -best surflets over dyadic blocks X_j achieves near-optimal performance.

As will be made apparent in Section IV-E, in order to achieve optimal performance, a coder must exploit correlations among nearby surflets. Unfortunately, these correlations may be difficult to exploit using native L_2 -best surflets. The problem arises because surflets with small X_j -active regions (Section II-C) may be close in L_2 distance over X_j yet have vastly different underlying polynomial coefficients. (These coefficients are used explicitly in our encoding strategy.)

To resolve this problem, we compute L_2 -best surflet fits to f^c over the L -extension of each dyadic hypercube X_j . That is, if

$$\begin{aligned} X_j &= [\beta_1 2^{-j}, (\beta_1 + 1) 2^{-j}] \times \dots \times [\beta_M 2^{-j}, (\beta_M + 1) 2^{-j}] \\ \text{then the } L\text{-extension of } X_j \text{ is defined to be} \\ X_j^L &= [(\beta_1 - L) 2^{-j}, (\beta_1 + 1 + L) 2^{-j}] \times \dots \\ &\quad \times [(\beta_M - L) 2^{-j}, (\beta_M + 1 + L) 2^{-j}] \end{aligned}$$

where $L > 0$ is an extension factor (designed to expand the domain of analysis and increase correlations between scales).⁸ An *L -extended surflet* is a surflet from $\mathcal{M}(j)$ that is now defined over X_j^L whose polynomial discontinuity has a nonempty X_j -active region. We define the *L -extended surflet dictionary* $\mathcal{R}_L(j)$ to be the set of L -extended surflets from $\mathcal{M}(j)$ plus the all-zero and all-one surflets $s(X_j) = 0$ and $s(X_j) = 1$. An *L -extended L_2 -best surflet* fit to f^c over X_j is then defined to be the L_2 -best surflet to f^c over X_j^L chosen from $\mathcal{R}_L(j)$. Note that even though extended surflets are defined over extended domains X_j^L , they are used to approximate the function only over the associated native domains X_j . Such extended surflet fits (over extended domains) provide sufficient mathematical constraints for a coder to relate nearby surflets, since extended surflets that are close in terms of squared- L_2 distance over X_j^L have similar polynomial coefficients (even if extended surflets have small X_j -active regions, they have large X_j^L -active regions). In Section IV-E, we describe a coder that uses extended surflets from $\mathcal{R}_L(j)$ to achieve optimal performance.

C. Tree-Based Surflet Approximations

The surflet dictionary consists of M -dimensional atoms at various scales. Thus, a 2^M -tree offers a natural topology for arranging the surflets used in an approximation. Specifically, each node at scale j in a 2^M -tree is labeled by a surflet that approximates the corresponding dyadic hypercube region X_j of the function f^c . This surflet can be assigned according to any of the procedures outlined in Section IV-B.

⁷We refer the reader to our technical report [28] for a thorough treatment of Taylor surflet-based approximation of piecewise constant multidimensional functions.

⁸If necessary, each L -extension is truncated to the hypercube $[0, 1]^M$.

Given a method for assigning a surfflet to each tree node, it is also necessary to determine the proper dyadic segmentation for the tree approximation. This can be accomplished using the CART algorithm, which is based on dynamic programming, in a process known as *tree-pruning* [8], [29]. Tree-pruning proceeds from the bottom up, determining whether to prune the tree beneath each node (causing it to become a leaf node). Various criteria exist for making such a decision. In particular, the approximation-theoretic optimal segmentation can be obtained by minimizing the Lagrangian cost $D + \lambda N$ for a penalty term λ . Similarly, the Lagrangian rate–distortion cost $D + \lambda R$ can be used to obtain the optimal rate–distortion segmentation. We note here that the optimal tree approximations computed according to these criteria are *nested*, i.e., if $\lambda_1 > \lambda_2$ the optimal tree for λ_1 is contained inside the optimal tree for λ_2 .

We summarize the construction of a surfflet-based approximation as follows.

Surfflet-based approximation

- **Choose scale:** Choose a maximal scale $J \in \mathbb{Z}$ for the 2^M -tree.
- **Label all nodes:** For each scale $j = 0, 1, \dots, J$, label all nodes at scale j with either a native or an extended L_2 -best surfflet chosen appropriately from either discrete dictionary of surfflets $\mathcal{M}(j)$ or $\mathcal{R}_{\mathcal{L}}(j)$.
- **Prune tree:** Starting at the second-finest scale $j = J - 1$, determine whether each node at scale j should be pruned (according to an appropriate pruning rule). Then proceed up to the root of the tree, i.e., until $j = 0$.

The approximation performance of this algorithm is described in the following theorem, which is proved in Appendix C.

Theorem 3: Using either quantized Taylor surfflets or L_2 -best surfflets (extended or native), a surfflet tree-pruned approximation of an element $f^c \in \mathcal{F}_C(M, K_d)$ achieves the optimal asymptotic approximation rate of Theorem 1

$$\|f^c - \widehat{f}_N^c\|_{L_2}^2 \lesssim \left(\frac{1}{N}\right)^{\frac{K_d}{M-1}}.$$

We discuss computational issues associated with finding best-fit surfflets in Section VI-F, where we also present results from simulations.

D. Leaf Encoding

An initial approach toward surfflet encoding would involve specification of the tree segmentation map (which denotes the location of the leaf nodes) along with the quantized surfflet coefficients at each leaf node. The rate–distortion analysis in Appendix D then yields the following result.

Theorem 4: Using either quantized Taylor surfflets or L_2 -best surfflets (extended or native), a surfflet leaf-encoder applied to an element $f^c \in \mathcal{F}_C(M, K_d)$ achieves the following rate–distortion performance

$$\|f^c - \widehat{f}_R^c\|_{L_2}^2 \lesssim \left(\frac{\log R}{R}\right)^{\frac{K_d}{M-1}}.$$

Comparing with Theorem 1, this simple coder is *near-optimal* in terms of rate–distortion performance. The logarithmic factor

is due to the fact that it requires $O(j)$ bits to encode each surfflet at scale j . In Section IV-E, we propose an alternative coder that requires only a constant number of bits to encode each surfflet.

E. Top-Down Predictive Encoding

Achieving the optimal performance of Theorem 1 requires a more sophisticated coder that can exploit the correlation among nearby surfflets. We now briefly describe a top-down surfflet coder that predicts surfflet parameters from previously encoded values.

Top-down predictive surfflet coder

- **Encode root node:** Encode the best surfflet fit $s([0, 1]^M)$ to the hypercube $[0, 1]^M$. Encode a flag (1-bit) specifying whether this node is interior or a leaf. Set $j \leftarrow 0$.
- **Predict surfflets from parent scale:** For every interior node/hypercube X_j at scale j , partition its domain into 2^M children hypercubes at scale $j + 1$. Compute the polynomial coefficients on each child hypercube X_{j+1} that agree with the encoded parent surfflet $s(X_j^L)$. These serve as “predictions” for the polynomial coefficients at the child.
- **Encode innovations at child nodes:** For each predicted polynomial coefficient, encode the discrepancy with the L -extended surfflet fit $s(X_{j+1}^L)$.
- **Descend tree:** Set $j \leftarrow j + 1$ and repeat until no interior nodes remain.

This top-down predictive coder encodes an entire tree segmentation starting with the root node, and proceeding from the top down. Given an L -extended surfflet $s(X_j^L)$ at an interior node at scale j , we show in Appendix E that the number of possible L -extended surfflets from $\mathcal{R}_{\mathcal{L}}(j)$ that can be used for approximation at scale $j + 1$ is *constant*, independent of the scale j . Thus, given a best fit surfflet at scale 0, a constant number of bits is required to encode each surfflet at subsequent scales. This prediction is possible because L -extended surfflets are defined over L -extended domains, which ensures coherency between the surfflet fits (and polynomial coefficients) at a parent and child node.

We note that predicting L -extended best fit surfflets to dyadic hypercube regions around the borders of $[0, 1]^M$ may not be possible with a constant number of bits when the discontinuity is not completely contained within the dyadic hypercube. However, we make the mild simplifying assumption that the intersections of the discontinuity with the hyperplanes $x_M = 0$ or $x_M = 1$ can be contained within $O(2^{(M-2)j})$ hypercubes at each scale j . Therefore, using $O(K_d j)$ bits to encode such “border” dyadic hypercubes (with the discontinuity intersecting $x_M = 0$ or $x_M = 1$) does not affect the asymptotic rate–distortion performance of the top-down predictive coder. In Appendix E, we prove the following theorem.

Theorem 5: The top-down predictive coder applied to an element $f^c \in \mathcal{F}_C(M, K_d)$ using L -extended L_2 -best surfflets from $\mathcal{R}_{\mathcal{L}}(j)$ achieves the optimal rate–distortion performance of Theorem 1

$$\|f^c - \widehat{f}_R^c\|_{L_2}^2 \lesssim \left(\frac{1}{R}\right)^{\frac{K_d}{M-1}}.$$

Although only the leaf nodes provide the ultimate approximation to the function, the additional information encoded at interior nodes provides the key to efficiently encoding the leaf nodes. In addition, unlike the surflet leaf-encoder of Section IV-C, this top-down approach yields a *progressive* bit stream—the early bits encode a low-resolution (coarse scale) approximation, which is then refined using subsequent bits.

F. Extensions to Broader Function Classes

Our results for classes of functions that contain a single discontinuity can be extended to spaces of signals that contain multiple discontinuities. Functions containing multiple discontinuities that do not intersect can be represented using the surflet-based approximation scheme described in Section IV-C at the optimal asymptotic approximation rate. This is because at a sufficiently high scale, dyadic hypercubes that tile signals containing multiple nonintersecting discontinuities contain at most one discontinuity.

Analysis of the surflet-based approximation scheme of Section IV-C applied to signals containing intersecting discontinuities is more involved. Let $f_{\#}^c$ be an M -dimensional piecewise constant function containing two $(M - 1)$ -dimensional \mathcal{C}^{K_d} -smooth discontinuities that intersect each other (the analysis that follows can easily be extended to allow for more than two intersecting discontinuities). Note that the intersection of $(M - 1)$ -dimensional functions forms an $(M - 2)$ -dimensional manifold. Again, we make the mild simplifying assumption that the intersection of the discontinuities can be contained in $O(2^{(M-2)j})$ hypercubes at each scale j . The following theorem (proved in Appendix F) describes the approximation performance achieved by the scheme in Section IV-C applied to $f_{\#}^c$. A consequence of this theorem is that there exists a smoothness threshold K_d^{th} that defines the boundary between optimal and suboptimal approximation performance.

Theorem 6: Using either quantized Taylor surflets or L_2 -best surflets (extended or native), the approximation scheme of Section IV-C applied to a piecewise constant M -dimensional function $f_{\#}^c$ that contains two intersecting \mathcal{C}_d^K -smooth $(M - 1)$ -dimensional discontinuities achieves performance given by

- $M > 2, K_d \leq \frac{2(M-1)}{M-2}$

$$\|f_{\#}^c - \widehat{f_{\#,N}^c}\|_{L_2}^2 \lesssim \left(\frac{1}{N}\right)^{\frac{K_d}{M-1}}.$$
- $M > 2, K_d > \frac{2(M-1)}{M-2}$

$$\|f_{\#}^c - \widehat{f_{\#,N}^c}\|_{L_2}^2 \lesssim \left(\frac{1}{N}\right)^{\frac{2}{M-2}}.$$
- $M = 2, \text{ any } K_d$

$$\|f_{\#}^c - \widehat{f_{\#,N}^c}\|_{L_2}^2 \lesssim \left(\frac{1}{N}\right)^{\frac{K_d}{M-1}}.$$

Thus, the representation scheme in Section IV-C achieves optimal approximation performance for $M = 2$ even in the presence of intersecting discontinuities, while it achieves optimal performance for $M > 2$ up to a smoothness threshold of

$K_d^{\text{th}} = \frac{2(M-1)}{M-2}$ (for $K_d > K_d^{\text{th}}$, the scheme performs suboptimally: $\|f_{\#}^c - \widehat{f_{\#,N}^c}\|_{L_2}^2 \lesssim \left(\frac{1}{N}\right)^{\frac{K_d^{\text{th}}}{M-1}}$). This performance of the approximation scheme for $M > 2$ is still superior to that of wavelets, which have $K_d^{\text{th, wl}} = 1$. The reason for this difference in performance between the cases $M = 2$ and $M > 2$ is that intersections of discontinuities when $M = 2$ correspond to points,⁹ while intersections in higher dimensions correspond to low-dimensional manifolds. Hence, the number of hypercubes that contain intersections in the two-dimensional case is constant with scale, whereas the number of hypercubes that contain the intersections when $M > 2$ grows exponentially with scale. The analysis above can clearly be extended to prove analogous results for functions containing piecewise \mathcal{C}^{K_d} -smooth discontinuities.

Future work will focus on improving the threshold K_d^{th} for the case $M > 2$. In order to achieve optimal performance for $M > 2$, one may need a dictionary containing regular surflets and specially designed “intersection” surflets that are specifically tailored for intersections. In addition to classes of functions containing multiple intersecting discontinuities, our representation scheme may also be adapted to function spaces where the transient direction of the discontinuity is not fixed to be x_M . This is possible due to the localized nature of our surflet-based approximations. Indeed, we note that with a minor modification, our surflet-based encoder can achieve the near-optimal rate–distortion performance (off by a logarithmic factor) of Section IV-D. This follows from the observation that each node in the surflet tree can be labeled with a best fit surflet from one of M different surflet dictionaries (one for each transient direction). Since the leaf nodes are encoded independently, the rate–distortion performance is given by Theorem 4.

V. REPRESENTATION AND CODING OF PIECEWISE *Smooth* FUNCTIONS

In this section, we extend our coding strategies for piecewise constant functions to encoding an arbitrary element f^s from the class $\mathcal{F}_S(M, K_d, K_s)$ of piecewise smooth functions.

A. Motivation

For a \mathcal{C}^{K_s} -smooth function f in M dimensions, a wavelet basis with sufficient vanishing moments [22] provides approximations at the optimal rate— $\|f - \widehat{f}_N\|_{L_2}^2 \lesssim \left(\frac{1}{N}\right)^{\frac{2K_s}{M}}$. Even if one introduces a finite number of point singularities into the M -dimensional \mathcal{C}^{K_s} -smooth function, wavelet-based approximation schemes still attain the optimal rate. Wavelets succeed in approximating smooth functions because most of the wavelet coefficients have small magnitudes and can thus be neglected. Moreover, an arrangement of wavelet coefficients on the nodes of a tree leads to an interesting consequence: wavelet coefficients used in the approximation of M -dimensional smooth functions are *coherent*—often, if a wavelet coefficient has small magnitude, then its children coefficients also have small magnitude. These properties of the wavelet basis have been exploited in state-of-the-art wavelet-based image coders [4], [5].

⁹Our analysis also applies to “T-junctions” in images, where one edge terminates at its intersection with another.




Fig. 5. Example surflet and the corresponding surfprint. The white box is the dyadic hypercube in which we define the surflet; note that the removal of coarse scale and neighboring wavelets causes the surfprint to appear different from the surflet.

Although wavelets approximate smooth functions well, the wavelet basis is not well-equipped to approximate functions containing higher dimensional manifold discontinuities. Wavelets also do not take advantage of any structure (such as smoothness) that the $(M - 1)$ -dimensional discontinuity might have, and therefore many high-magnitude coefficients are often required to represent discontinuities [16]. Regardless of the smoothness order of the discontinuity, the approximation rate achieved by wavelets remains the same.

Despite this drawback, we desire a wavelet domain solution to approximate $f^s \in \mathcal{F}_S(M, K_d, K_s)$ because most of the function f^s is smooth in M dimensions, except for an $(M - 1)$ -dimensional discontinuity. In order to solve the problem posed by the discontinuity, we propose the addition of *surfprint* atoms to the dictionary of wavelet atoms. A surfprint is a weighted sum of wavelet basis functions derived from the projection of a piecewise polynomial surflet atom (an $(M - 1)$ -dimensional polynomial discontinuity separating two M -dimensional polynomial regions) onto a subspace in the wavelet domain (see Fig. 5 for an example in two dimensions). Surfprints possess all the properties that make surflets well-suited to represent discontinuities. In addition, surfprints coherently model wavelet coefficients that correspond to discontinuities. Thus, we obtain a single unified wavelet-domain framework that is well-equipped to sparsely represent both discontinuities and smooth regions.

The rest of this section is devoted to the definition of surfprints and their use in a wavelet domain framework to represent and encode approximations to elements of $\mathcal{F}_S(M, K_d, K_s)$. We do not discuss the extension of our results to classes of piecewise smooth signals containing multiple intersecting discontinuities, but note that such an analysis would be similar to that described in Section IV-F.

B. Surfprints

Let X_{J_o} be a dyadic hypercube at scale J_o . Let v_1, v_2 be M -dimensional polynomials of degree r_s^{SP} , and let v be an M -dimensional function as follows:

$$v_1, v_2, v : X_{J_o} \rightarrow \mathbb{R}.$$

Let q be an $(M - 1)$ -dimensional polynomial of degree r_d^{SP}

$$q : Y_{J_o} \rightarrow \mathbb{R}.$$

As defined in Section II-B, let $\mathbf{x} \in X_{J_o}$ and let \mathbf{y} denote the first $M - 1$ elements of \mathbf{x} . Let the M -dimensional piecewise polynomial function v be defined as follows:

$$v(\mathbf{x}) = \begin{cases} v_1(\mathbf{x}), & x_M \geq q(\mathbf{y}) \\ v_2(\mathbf{x}), & x_M < q(\mathbf{y}). \end{cases}$$

Next, we describe how this piecewise polynomial function is projected onto a wavelet subspace to obtain a surfprint atom. Let \mathcal{W} be a compactly supported wavelet basis in M dimensions with K_s^{wl} vanishing moments. A surfprint $\text{sp}(v, X_{J_o}, \mathcal{W})$ is a weighted sum of wavelet basis functions with the weights derived by projecting the piecewise polynomial v onto the subtree of basis functions whose idealized supports nest in the hypercube X_{J_o} .

$$\text{sp}(v, X_{J_o}, \mathcal{W}) = \sum_{j \geq J_o, X_j \subseteq X_{J_o}} \langle v, w_{X_j} \rangle w_{X_j} \quad (5)$$

where w_{X_j} represents the wavelet basis function having idealized compact support on the hypercube X_j . (The actual support of w_{X_j} may extend slightly beyond X_j .) The hypercube X_{J_o} thus defines the root node (or coarsest scale) of the surfprint atom.

We propose an approximation scheme in Section V-E, where we use wavelet atoms to represent uniformly smooth regions of f^s and surfprint atoms to represent regions through which the discontinuity passes. Before presenting our approximation scheme, we begin in Section V-C by describing how to choose the surfprint polynomial degrees r_s^{SP} and r_d^{SP} and the number of vanishing moments K_s^{wl} for the wavelet basis.

C. Vanishing Moments and Polynomial Degrees

In general, due to Taylor's theorem, when approximating elements $f^s \in \mathcal{F}_S(M, K_d, K_s)$, the required surfprint polynomial degrees and wavelet vanishing moments are determined by the orders of smoothness K_d and K_s

$$K_s^{\text{wl}} \geq K_s, r_d^{\text{SP}} = \lceil K_d - 1 \rceil \quad \text{and} \quad r_s^{\text{SP}} = \lceil K_s - 1 \rceil.$$

However, the exponent in the expression of Theorem 2 for the optimal approximation rate for $\mathcal{F}_S(M, K_d, K_s)$ indicates that for every (K_d, K_s) , either the $(M - 1)$ -dimensional discontinuity or the M -dimensional smooth region dominates the decay rate. For instance, in two dimensions, the smaller of the two smoothness orders K_d and K_s defines the decay rate.¹⁰ This implies that the surfprint polynomial degrees and/or the number of wavelet vanishing moments can be relaxed (as if either the discontinuity or the smooth regions had a lower smoothness order), without affecting the approximation rate.

Rather than match the surfprint parameters directly to the smoothness orders K_d and K_s , we let K_d^{SP} and K_s^{SP} denote the *operational smoothness orders* to which the surfprint parameters are matched. These operational smoothness orders are selected to ensure the best approximation or rate–distortion performance. The detailed derivations of Appendices G and H yield the following values for the operational smoothness orders.

¹⁰We note also that in the case where the functions g_1 and g_2 , which characterize f^s above and below the discontinuity, have differing orders of smoothness, the smaller smoothness order will determine both the achievable approximation rates and the appropriate approximation strategies.

- **Discontinuity dominates:** In this case, $\frac{K_d}{M-1} < \frac{2K_s}{M}$. We let $K_d^{\text{SP}} = K_d$ and choose $K_s^{\text{SP}} \in [\frac{K_d-1}{2}, K_s]$ and $K_s^{\text{wl}} \in [\frac{K_d M}{2(M-1)}, K_s]$.
- **Smooth regions dominate:** In this case, $\frac{2K_s}{M} < \frac{K_d}{M-1}$. We let $K_s^{\text{wl}} = K_s$, and choose $K_s^{\text{SP}} \in [K_s(1 - \frac{1}{M}) - \frac{1}{2}, K_s]$ and $K_d^{\text{SP}} \in [\frac{2K_s(M-1)}{M}, K_d]$.
- **Both contribute equally:** In this case, $\frac{2K_s}{M} = \frac{K_d}{M-1}$. We let $K_s^{\text{wl}} = K_s$, $K_d^{\text{SP}} = K_d$, and choose $K_s^{\text{SP}} \in [K_s(1 - \frac{1}{M}) - \frac{1}{2}, K_s]$.

The surfprint polynomial degrees are given by

$$r_d^{\text{SP}} = \lceil K_d^{\text{SP}} - 1 \rceil \quad \text{and} \quad r_s^{\text{SP}} = \lceil K_s^{\text{SP}} - 1 \rceil.$$

Therefore, if $\lceil K_d^{\text{SP}} - 1 \rceil < \lceil K_d - 1 \rceil$ and $\lceil K_s^{\text{SP}} - 1 \rceil < \lceil K_s - 1 \rceil$, then the required surfprint polynomial degrees for optimal approximations are lower than what one would naturally expect. Note that even in the scenario where both terms in the exponent of the approximation rate match, one can choose K_s^{SP} slightly smaller than K_s while still attaining the optimal approximation rate of Theorem 2.

D. Quantization

In order to construct a discrete surfprint/wavelet dictionary, we quantize the coefficients of the wavelet and surfprint atoms. The quantization step size ΔK_s^{wl} for the wavelet coefficients depends on the specific parameters of an approximation scheme. We present our prototype approximation scheme and discuss the wavelet coefficient step sizes in Section V-E (see (8) below).

The quantization step size for the surfprint polynomial coefficients of order ℓ at scale j is analogous to the step size used to construct a discrete surflet dictionary (3)

$$\Delta_{\ell,j}^{K_d^{\text{SP}}} = 2^{-(K_d^{\text{SP}} - \ell)j} \quad (6)$$

and

$$\Delta_{\ell,j}^{K_s^{\text{SP}}} = 2^{-(K_s^{\text{SP}} - \ell)j}. \quad (7)$$

As before, the key idea is that higher order polynomial coefficients can be quantized with lesser precision without affecting the error term in the Taylor approximation (1).

E. Surfprint-Based Approximation

We present a tree-based representation scheme using quantized wavelet and surfprint atoms and prove that this scheme achieves the optimal approximation rate for every function $f^s \in \mathcal{F}_S(M, K_d, K_s)$. Let \mathcal{W} be a compactly supported wavelet basis in M dimensions with K_s^{wl} vanishing moments, as defined in Section V-C. Consider the decomposition of f^s into the wavelet basis vectors: $f^s = \sum_j \langle f^s, w_{X_j} \rangle w_{X_j}$. The wavelet coefficients $\langle f^s, w_{X_j} \rangle$ are quantized according to the step size ΔK_s^{wl} defined below. Let these wavelet atoms be arranged on the nodes of a 2^M -tree. We classify the nodes based on the idealized support of the corresponding wavelet basis functions. Nodes whose supports X_j are intersected by the discontinuity b^s are called Type D nodes. All other nodes (over which f^s is smooth) are classified

as Type S. Consider now the following surfprint approximation strategy.¹¹

Surfprint approximation

- **Choose scales and wavelet quantization step size:** Choose a maximal scale $J \in \mathbb{Z}$ and $m, n \in \mathbb{Z}$ such that $\frac{m}{n} = \frac{M}{M-1}$ and both m and n divide J . The quantization step size for wavelet coefficients at all scales j is given by

$$\Delta K_s^{\text{wl}} = 2^{-\frac{j}{m}(K_s^{\text{wl}} + \frac{M}{2})} \quad (8)$$

and thus depends only on the maximal scale J and the parameter m .

- **Prune tree:** Keep all wavelet nodes up to scale $\frac{J}{m}$; from scale $\frac{J}{m}$ to scale $\frac{J}{n}$, prune the tree at all Type S nodes (discarding those wavelet coefficients and their descendant subtrees).
- **Select surfprint atoms:** At scale $\frac{J}{n}$ replace the wavelet atom at each Type D discontinuity node and its descendant subtree (up to depth J) by a quantized surfprint atom chosen appropriately from the dictionary with $J_o = \frac{J}{n}$ in (5):

— M -dimensional polynomials: Choose M -dimensional polynomials v_1 and v_2 of degree $r_s^{\text{SP}} = \lceil K_s^{\text{SP}} - 1 \rceil$. These polynomials should approximate the M -dimensional smooth regions up to an absolute (pointwise) error of $O(2^{-\frac{K_s^{\text{SP}} J}{n}})$. The existence of such polynomials is guaranteed by Taylor's theorem (1) (let $D = M$, $K = K_s^{\text{SP}}$, and $r = r_s^{\text{SP}}$) and the quantization scheme (7).

— $(M-1)$ -dimensional polynomial: Choose an $(M-1)$ -dimensional polynomial q of degree $r_d^{\text{SP}} = \lceil K_d^{\text{SP}} - 1 \rceil$ such that the discontinuity is approximated up to an absolute error of $O(2^{-\frac{K_d^{\text{SP}} J}{n}})$. The existence of such a polynomial is guaranteed by Taylor's theorem (1) (let $D = M-1$, $K = K_d^{\text{SP}}$, and $r = r_d^{\text{SP}}$) and the quantization scheme of (6).

The following theorem summarizes the performance analysis for such surfprint approximations (see Appendix G for the proof).

Theorem 7: A surfprint-based approximation of an element $f^s \in \mathcal{F}_S(M, K_d, K_s)$ as presented above achieves the optimal asymptotic approximation rate of Theorem 2

$$\|f^s - \widehat{f}_N^s\|_{L_2}^2 \lesssim \left(\frac{1}{N}\right)^{\min\left(\frac{K_d}{M-1}, \frac{2K_s}{M}\right)}.$$

An approximation scheme that uses the *best* configuration of N wavelet and surfprint atoms in the L_2 sense would perform at least as well as the scheme suggested above. Hence, surfprint approximation algorithms designed to choose the best N -term approximation (even without explicit knowledge of the discontinuity or the M -dimensional smooth regions) will achieve the optimal approximation rate of Theorem 2.

¹¹The wavelet decomposition actually has $2^M - 1$ distinct directional subbands; we assume here that each is treated identically. Also we assume the scaling coefficient at the coarsest scale $j = 0$ is encoded as side information with negligible cost.

F. Encoding a Surfprint/Wavelet Approximation

We now consider the problem of encoding the tree-based approximation of Section V-E. A simple top-down coding scheme that specifies the pruned tree topology, quantized wavelet coefficients, and surfprint parameters achieves a near-optimal rate–distortion performance (see Appendix H for proof).

Theorem 8: A coding scheme that encodes every element of the surfprint-based approximation of an element $f^s \in \mathcal{F}_S(M, K_d, K_s)$ as presented in Section V-E achieves the near-optimal asymptotic rate–distortion performance (within a logarithmic factor of the optimal performance of Theorem 2)

$$\|f^s - \widehat{f}_R^s\|_{L_2}^2 \lesssim \left(\frac{\log R}{R}\right)^{\min\left(\frac{K_d}{M-1}, \frac{2K_s}{M}\right)}.$$

Repeating the argument of Section V-E, this near-optimal rate–distortion performance serves as an upper bound for an encoding scheme that encodes elements of an L_2 -best approximation. We will discuss the extension of these theoretical results to the approximation of discrete data and related issues in Section VI-C.

VI. EXTENSIONS TO DISCRETE DATA

A. Overview

In this section, we consider the problem of representing discrete data obtained by “voxelizing” (pixelizing in two dimensions) functions from the classes $\mathcal{F}_C(M, K_d)$ and $\mathcal{F}_S(M, K_d, K_s)$. Let f be a continuous M -dimensional function. We discretize f according to a vector $\boldsymbol{\pi} = [2^{\pi_1}, \dots, 2^{\pi_M}] \in \mathbb{Z}^M$, which specifies the number of voxels along each dimension of the discretized M -dimensional function $\widehat{f}_\boldsymbol{\pi}$. Each entry of $\widehat{f}_\boldsymbol{\pi}$ is obtained either by averaging f over an M -dimensional voxel or by sampling f at uniformly spaced intervals. (Because of the smoothness characteristics of $\mathcal{F}_C(M, K_d)$ and $\mathcal{F}_S(M, K_d, K_s)$, both discretization mechanisms provide the same asymptotic performance.) In our analysis, we allow the number of voxels along each dimension to vary in order to provide a framework for analyzing various sampling rates along the different dimensions. Video data, for example, is often sampled differently in the spatial and temporal dimensions. Future research will consider different distortion criteria based on asymmetry in the spatio-temporal response of the human visual system.

For our analysis, we assume that the voxelization vector $\boldsymbol{\pi}$ is fixed and denote the resulting classes of voxelized functions by $\widetilde{\mathcal{F}}_C(M, K_d)$ and $\widetilde{\mathcal{F}}_S(M, K_d, K_s)$. Section VI-B and VI-C describe the sparse representation of elements from $\widetilde{\mathcal{F}}_C(M, K_d)$ and $\widetilde{\mathcal{F}}_S(M, K_d, K_s)$, respectively. In Section VI-D, we discuss the impact of discretization effects on fine scale approximations. In Section VI-E, we provide an overview of practical methods to choose best fit elements from the discrete dictionaries. Finally, we present our simulation results in Section VI-F.

B. Representing and Encoding Elements of $\widetilde{\mathcal{F}}_C(M, K_d)$

Suppose $f^c \in \mathcal{F}_C(M, K_d)$ and let $\widehat{f}_\boldsymbol{\pi}^c \in \widetilde{\mathcal{F}}_C(M, K_d)$ be its discretization. (We view $\widehat{f}_\boldsymbol{\pi}^c$ as a function on the continuous domain $[0, 1]^M$ that is constant over each voxel.) The process

of voxelization affects the ability to approximate elements of $\widetilde{\mathcal{F}}_C(M, K_d)$. At coarse scales, however, much of the intuition for coding $\mathcal{F}_C(M, K_d)$ can be retained. In particular, we can bound the distance from $\widehat{f}_\boldsymbol{\pi}^c$ to f^c . We note that $\widehat{f}_\boldsymbol{\pi}^c$ differs from f^c only over voxels through which b^c passes. Because each voxel has size $2^{-\pi_1} \times 2^{-\pi_2} \dots \times 2^{-\pi_M}$, the number of voxels intersected by b^c is $O(2^{\sum_{i=1}^{M-1} \pi_i} \lceil (\Omega \cdot 2^{-\min(\pi_i)_{i=1}^{M-1}}) / (2^{-\pi_M}) \rceil)$, where Ω is the universal derivative bound (Section II-A). The squared- L_2 distortion incurred on each such voxel (assuming only that the voxelization process is bounded and local) is $O(2^{-(\pi_1 + \dots + \pi_M)})$. Summing over all voxels it follows that the (nonsquared) L_2 distance obeys

$$\|f^c - \widehat{f}_\boldsymbol{\pi}^c\|_{L_2([0,1]^M)} < C_1 \cdot 2^{-(\min \pi_i)/2} \quad (9)$$

where the minimum is taken over all $i \in \{1, \dots, M\}$.

Now we consider the problem of encoding elements of $\widetilde{\mathcal{F}}_C(M, K_d)$. At a particular bit rate R , we know from Theorem 1 that no encoder could represent all elements of $\mathcal{F}_C(M, K_d)$ using R bits and incurring L_2 distortion less than $C_2 \cdot \left(\frac{1}{R}\right)^{\frac{K_d}{2(M-1)}}$. (This lower bound for metric entropy is in effect for R sufficiently large, which we assume to be the case.) Suppose we consider a hypothetical encoder for elements of $\mathcal{F}_C(M, K_d)$ that, using R bits, could represent any element with L_2 distortion of $\widetilde{\mathcal{F}}_C(M, K_d)$ less than some $D_{\text{hyp}}(R)$. This coder could *also* be used as an encoder for elements of $\mathcal{F}_C(M, K_d)$ (by voxelizing each function before encoding). This strategy would yield L_2 distortion no worse than $C_1 \cdot 2^{-(\min \pi_i)/2} + D_{\text{hyp}}(R)$. By applying the metric entropy arguments on $\mathcal{F}_C(M, K_d)$, we have the following constraint on $D_{\text{hyp}}(R)$:

$$C_1 \cdot 2^{-(\min \pi_i)/2} + D_{\text{hyp}}(R) \geq C_2 \cdot \left(\frac{1}{R}\right)^{\frac{K_d}{2(M-1)}}$$

or equivalently

$$D_{\text{hyp}}(R) \geq C_2 \cdot \left(\frac{1}{R}\right)^{\frac{K_d}{2(M-1)}} - C_1 \cdot 2^{-(\min \pi_i)/2}. \quad (10)$$

This inequality helps establish a rate–distortion bound for the class $\widetilde{\mathcal{F}}_C(M, K_d)$. At sufficiently low rates, the first term on the right-hand side (RHS) dominates, and $\widetilde{\mathcal{F}}_C(M, K_d)$ faces similar rate–distortion constraints to $\mathcal{F}_C(M, K_d)$. At high rates, however, the RHS becomes negative, giving little insight into the coding of $\widetilde{\mathcal{F}}_C(M, K_d)$. This breakdown point occurs when $R \sim 2^{(\min \pi_i)(M-1)/K_d}$.

We can, in fact, specify a constructive encoding strategy for $\widetilde{\mathcal{F}}_C(M, K_d)$ that achieves the optimal compression rate up to this breakdown point. We construct a dictionary of discrete surflet atoms by voxelizing the elements of the continuous quantized surflet dictionary. Assuming there exists a technique to find discrete ℓ_2 -best surflet fits to $\widehat{f}_\boldsymbol{\pi}^c$, the tree-based algorithm described in Section IV-C can simply be used to construct an approximation $\widehat{\widehat{f}}_\boldsymbol{\pi}^c$.

Theorem 9: While $R \lesssim 2^{(\min \pi_i)(M-1)/K_d}$, the top-down predictive surflet coder from Section IV-E applied to encode the

approximation $\widehat{f_\pi^c}$ to $\widetilde{f_\pi^c}$ using discrete ℓ_2 -best surflets achieves the rate–distortion performance

$$\left\| \widetilde{f_\pi^c} - \widehat{f_\pi^c} \right\|_{L_2}^2 \lesssim \left(\frac{1}{R} \right)^{\frac{K_d}{M-1}}.$$

As detailed in the proof of this theorem (see Appendix J), the breakdown point occurs when using surflets at a critical scale $J_{\text{vox}} = \frac{\min \pi_i}{K_d}$. Up to this scale, all of the familiar approximation and compression rates hold. Beyond this scale, however, voxelization effects dominate. An interesting corollary to Theorem 9 is that, due to the similarities up to scale J_{vox} , the discrete approximation $\widehat{f_\pi^c}$ itself provides an effective approximation to the function f^c .

Corollary 10: While $R \lesssim 2^{(\min \pi_i)(M-1)/K_d}$, the discrete approximation $\widehat{f_\pi^c}$ provides an approximation to f^c with the following rate–distortion performance:

$$\left\| f^c - \widehat{f_\pi^c} \right\|_{L_2}^2 \lesssim \left(\frac{1}{R} \right)^{\frac{K_d}{M-1}}.$$

The details of the proof appear in Appendix J. While we have provided an effective strategy for encoding elements of $\widetilde{\mathcal{F}}_C(M, K_d)$ at sufficiently low rates (using surflets at scales $j \leq J_{\text{vox}}$), this leaves open the question of how to code $\widetilde{\mathcal{F}}_C(M, K_d)$ at higher rates. Unfortunately, (10) does not offer much insight. In particular, it is not clear whether surflets are an efficient strategy for encoding $\widetilde{\mathcal{F}}_C(M, K_d)$ beyond scale J_{vox} . We revisit this issue in Section VI-D.

C. Representing and Encoding Elements of $\widetilde{\mathcal{F}}_S(M, K_d, K_s)$

Next, let $\widetilde{f_\pi^s}$ be an arbitrary signal belonging to $\widetilde{\mathcal{F}}_S(M, K_d, K_s)$. Similar arguments apply to the voxelization effects for this class. In order to approximate functions in $\widetilde{\mathcal{F}}_S(M, K_d, K_s)$, we use a dictionary of compactly supported discrete wavelet basis functions with K_s^{wl} vanishing moments and discrete surfprint atoms. A discrete surfprint atom is derived by projecting a discrete piecewise polynomial surflet atom onto a subspace of the discrete wavelet basis.

We use the scheme described in Section V-E with

$$\frac{J_{\text{vox}}}{n} = \frac{\min(\pi_i)}{\min(K_d^{\text{SP}}, 2K_s^{\text{SP}} + 1)}$$

to approximate $\widetilde{f_\pi^s}$ by $\widehat{f_\pi^s}$. Using (40), (41), and (42), this scale corresponds to a range of bit rates up to $O(J_{\text{vox}} 2^{(M-1)\frac{J_{\text{vox}}}{n}})$. Within this range, the approximation is encoded as described in Section V-F. The performance of this scheme is evaluated in Appendix K and appears below.

Theorem 11: While $R \lesssim J_{\text{vox}} 2^{(M-1)\frac{J_{\text{vox}}}{n}}$ where

$$J_{\text{vox}} = \frac{n \cdot \min(\pi_i)}{\min(K_d^{\text{SP}}, 2K_s^{\text{SP}} + 1)}$$

the coding scheme from Section V-E applied to encode the approximation $\widehat{f_\pi^s}$ to $\widetilde{f_\pi^s}$ using a discrete wavelet/surfprint dictionary achieves the following near-optimal asymptotic rate–dis-

ortion performance (within a logarithmic factor of the optimal performance of Theorem 2):

$$\left\| \widetilde{f_\pi^s} - \widehat{f_\pi^s} \right\|_{L_2}^2 \lesssim \left(\frac{\log R}{R} \right)^{\min\left(\frac{K_d}{M-1}, \frac{2K_s}{M}\right)}.$$

Again, a corollary follows naturally (see Appendix L for the proof).

Corollary 12: While $R \lesssim J_{\text{vox}} 2^{(M-1)\frac{J_{\text{vox}}}{n}}$, the discrete approximation $\widehat{f_\pi^s}$ provides an approximation to f^s with the following rate–distortion performance:

$$\left\| f^s - \widehat{f_\pi^s} \right\|_{L_2}^2 \lesssim \left(\frac{\log R}{R} \right)^{\min\left(\frac{K_d}{M-1}, \frac{2K_s}{M}\right)}.$$

D. Discretization Effects and Varying Sampling Rates

We have proposed surflet algorithms for discrete data at sufficiently coarse scales. Unfortunately, this leaves open the question of how to represent such data at finer scales. In this section, we discuss one perspective on fine scale approximation that leads to a natural surflet coding strategy.

Consider again the class $\widetilde{\mathcal{F}}_C(M, K_d)$. Section VI-B provided an effective strategy for encoding elements of $\widetilde{\mathcal{F}}_C(M, K_d)$ at sufficiently low rates (using surflets at scales $j \leq J_{\text{vox}} = \frac{\min \pi_i}{K_d}$). Beyond scale J_{vox} , however, the voxelization effects dominate the resolution afforded by surflet approximations. To restore a balance, we suggest a coding strategy for finer scales based on the observation that $\mathcal{F}_C(M, K_d) \subset \mathcal{F}_C(M, K)$ for $K < K_d$. Surflet approximations on the class $\mathcal{F}_C(M, K)$ (tied to the smoothness K) have lower accuracy in general. As a result, $\widetilde{\mathcal{F}}_C(M, K)$ has a higher “breakdown rate” than $\widetilde{\mathcal{F}}_C(M, K_d)$, and discrete surflets tailored for smoothness K will achieve the coding rate $O(R^{-K/(M-1)})$ up to scale $\frac{\min \pi_i}{K}$. While this may not be a worthwhile strategy before scale J_{vox} , it could be useful beyond scale J_{vox} and up to scale $\frac{\min \pi_i}{K}$. In fact, beyond that scale, we can *again* reduce K , obtaining a new breakdown rate and a finer scale to code (using lower order surflets). This gives us a concrete strategy for coding $\widetilde{\mathcal{F}}_C(M, K_d)$ at all scales, although our optimality arguments apply only up to scale J_{vox} . At scale j , we use surflets designed for smoothness $K_j = \min(K_d, \frac{\min(\pi_i)}{j})$, $0 \leq j \leq \min(\pi_i)$. A surflet dictionary constructed using such scale-adaptive smoothness orders consists of relatively few elements at coarse scales (due to the low value of j in the quantization stepsize) and relatively few at fine scales (due to the decrease of K_j), but many elements at medium scales. This agrees with the following intuitive notions.

- The large block sizes at coarse scales do not provide sufficient resolution to warrant large dictionaries for approximation at these scales.
- The relatively small number of voxels in each block at very fine scales also means that a coder does not require large dictionaries in order to approximate blocks at such scales well.
- At medium scales, where the block sizes are small enough to provide good resolution but large enough to contain

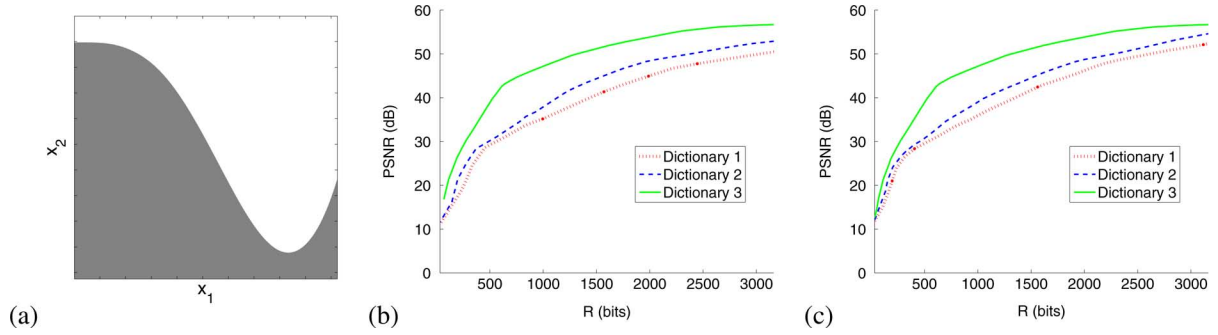


Fig. 6. (a) Test function \tilde{f}_{π}^c . (b) Rate–distortion performance for each dictionary (with the best fixed set of dictionary parameters). (c) Rate–distortion performance for each dictionary (selected using best convex hull in R/D plane over all dictionary parameters).




Fig. 7. (a) Test function \tilde{f}_{π}^c . (b) Rate–distortion performance for each dictionary (selected using best convex hull in R/D plane over all dictionary parameters).

many voxels, the dictionary contains many elements in order to provide good approximations.

Similar strategies can be proposed, of course, for the class $\tilde{\mathcal{F}}_S(M, K_d, K_s)$.

Finally, we note that the interplay between the sampling rate (number of voxels) along the different dimensions and the critical approximation scale J_{vox} can impact the construction of multiscale source coders. As an example of the potential effect of this phenomenon in real-world applications, the sampling rate along the temporal dimension could be the determining factor when designing a surfprint-based video coder because this rate tends to be lower than the sampling rate along the spatial dimensions.

E. Choosing Best Fit Surflets

In order to algorithmically build an approximation \tilde{f}_{π}^c using surflets, it is necessary to first find the ℓ_2 -best surflet fit at each node. Based on the quantization of the surflet dictionary, we must distinguish from among $O(2^{K_d j})$ discrete candidates to find the best surflet fit for each block of voxels at scale j . Fortunately, it is not necessary in practice to actually enumerate each of these possibilities in order to find the best fit. In fact, the same multiscale dependencies that enable efficient top-down predictive coding (see Section IV-E) can be used to facilitate the surflet fitting: given the best fit quantized surflet at a parent node, there are only $O(1)$ possibilities for the best quantized child surflet. Each of these could conceivably be enumerated.

However, it is also possible to avoid a discrete search altogether when finding best fit surflets. One alternative is to formulate the search as a continuous optimization problem, seeking the (nonquantized) surflet coefficients that minimize the ℓ_2 error

from the surflet to the voxelized data. Geometrically, this can be viewed as a search along a manifold (whose dimension matches the number of surflet coefficients) for the closest point to the data point. As described in [30], this search is complicated by the fact that such manifolds are nowhere differentiable. A remedy is proposed to this situation through a coarse-to-fine iterative algorithm that employs successive iterations of Newton’s method on regularized versions of the data cube. The complexity of such an algorithm depends on several factors, such as the accuracy of an initial guess [30]; however, such a starting guess can again be provided from each parent node. In practice, the multiscale Newton algorithm tends to converge in a very small number of iterations (10 or fewer), and it performs significantly faster than a discrete search. We have used it below for the experiments described in Section VI-F. This algorithm can be immediately adapted to find best fit *surfprints* as well.

F. Simulation Results on Synthetic Signals

To demonstrate the potential for coding gains based on surflet representations, we present the following preliminary numerical experiments in two and three dimensions.

1) *2-D Coding*: We start by coding elements of $\tilde{\mathcal{F}}_C(M, K_d)$ with $M = 2$ and $K_d = 3$. We generate 1024×1024 discretized versions of these images (that is, $\pi_1 = \pi_2 = 10$). Our two example images are shown in Figs. 6(a) and 7(a).

On each image we test three types of surflet dictionaries for encoding.

- Dictionary 1 uses wedgelets as implemented in our previous work [23], [31]. In this dictionary we do not use the quantization stepsizes as specified in (3). Rather, we use a quantization step size $\Delta_{\ell, j} \sim 2^{-(1-\ell)j}$. As a result, the

TABLE I
SURFLET DICTIONARY SIZE AT EACH SCALE (USING THE SURFLET PARAMETERS CHOSEN TO GENERATE FIG. 6(B)). OUR SURFLET DICTIONARIES (2 AND 3) ADAPT TO SCALE, AVOIDING UNNECESSARY PRECISION AT COARSE AND FINE SCALES

Scale j	0	1	2	3	4	5	6	7	8	9
Dictionary 1	181631	181631	181631	181631	181631	181631	181631	181631	181631	181631
Dictionary 2	219	4143	62655	987903	987903	248191	62655	15967	4143	1119
Dictionary 3	357	14335	407719	12216207	6264455	248191	62655	15967	4143	1119

quantized wedgelet dictionary has the same cardinality at each scale and is self-similar (simply a dyadic scaling of the dictionary at other scales).

- Dictionary 2 adapts with scale. Following the arguments of Section VI-D, at a given scale j , we use surflets tailored for smoothness

$$K_j = \min\left(2, \frac{\min \pi_i}{j}\right) = \min\left(2, \frac{10}{j}\right).$$

We use surflets of the appropriate polynomial order and quantize the polynomial coefficients analogous to (3); that is, $\Delta_{\ell,j} \sim 2^{-(K_j - \ell)j}$. The limitation $K_j \leq 2$ restricts our surflets to linear polynomials (wedgelets) for comparison with the first dictionary above.

- Dictionary 3 is a surflet dictionary that also adapts with scale. This dictionary is constructed similarly to the second, except that it is tailored to the actual smoothness of the discontinuity in f^c : we set

$$K_j = \min\left(K_d, \frac{\min \pi_i}{j}\right) = \min\left(K_d, \frac{10}{j}\right).$$

This modification allows quadratic surflets to be used at coarse scales $0 \leq j \leq 5$, beyond which K_j again dictates that wedgelets are used.

For each dictionary, we must also specify the range of allowable polynomial coefficients and a constant multiplicative factor on each quantization stepsize. We optimize these parameters through simulation.

Our coding strategy for each dictionary uses a top-down prediction. Based on the prediction from a (previously coded) parent surflet, we partition the set of possible children surflets into two classes for entropy coding. A probability mass of ρ is distributed among the W surflets nearest the predicted surflet (measured using ℓ_2 distance), and a probability mass of $(1 - \rho)$ is distributed among the rest to allow for robust encoding. We optimize the choice of W and ρ experimentally.

To find the discrete ℓ_2 -best fit surflet to a given block, we use a coarse-to-fine iterative algorithm to search for the closest point along the manifold of possible surflets (see Section VI-E). Based on the costs incurred by this coding scheme, we optimize the surflet tree pruning using a Lagrangian tradeoff parameter λ . We repeat the experiment for various values of λ .

Fig. 6(b) shows what we judge to be the best R/D curve for each dictionary (Dictionary 1: dotted curve, 2: dashed curve, and 3: solid curve.) Each curve is generated by sweeping λ but fixing one combination of polynomial parameters/constants. Over all simulations (all polynomial parameters/constants), we also take the convex hull over all points in the R/D plane. The results are plotted in Figs. 6(c) and 7(b).

We see from the figures that Dictionary 2 outperforms Dictionary 1, requiring 0–20% fewer bits for an equivalent distortion

(or improving PSNR by up to 4 dB at a given bit rate). Both dictionaries use wedgelets—we conclude that the coding gain comes from the adaptivity through scale. Table I lists the number of admissible quantized surflets as a function of scale j for each of our three dictionaries.

We also see from the figures that Dictionary 3 often outperforms Dictionary 2, requiring 0–50% fewer bits for an equivalent distortion (or improving PSNR by up to 10 dB at a given bit rate). Both dictionaries adapt to scale—we conclude that the coding gain comes from the quadratic surflets used at coarse scales (which are designed to exploit the actual smoothness $K_d = 3$). Fig. 4 compares two pruned surflet decompositions using Dictionaries 2 and 3. In this case, the quadratic dictionary offers comparable distortion using 40% fewer bits than the wedgelet dictionary.

2) *3-D Coding*: We now describe numerical experiments for coding elements of $\widetilde{\mathcal{F}}_C(M, K_d)$ and $M = 3$. We generate $64 \times 64 \times 64$ discretized versions of these signals (that is, $\pi_i = 6$). Our two example discontinuities b^c are shown in Fig. 8(a) (for which $K_d = 2$) and Fig. 10(a) (for which $K_d = \infty$).

For these simulations, we compare surflet coding (analogous to Dictionary 2 above, with $K_j = \min(2, \frac{6}{j})$) with wavelet coding. Our wavelet coding is based on a 3-D Haar wavelet transform, which we threshold at a particular level (keeping the largest wavelet coefficients). For the purpose of the plots, we assume (optimistically) that each significant wavelet coefficient was coded with zero distortion using only three bits per coefficient. We see from the figures that surflet coding outperforms the wavelet approach for the synthetic examples presented, requiring up to 80% fewer bits than our aggressive wavelet estimate (or improving PSNR by up to 10 dB at a given bit rate). Fig. 9 shows one set of coded results for the synthetic function in Fig. 8; at an equivalent bit rate, we see that surflets offer an improvement in PSNR and also reduce ringing/blocking artifacts compared with wavelets. We also notice from Figs. 8 and 10, however, that at high bit rates the gains diminish relative to wavelets. We believe this is due to small errors made in the surflet estimates at fine scales using our current implementation of the manifold-based technique. Future work will focus on improved surflet estimation algorithms; however using even these suboptimal estimates we *still* see improved performance on synthetic signals across a wide range of bit rates.

VII. CONCLUSION

In this paper, we have studied the representation and compression of piecewise constant and piecewise smooth functions with smooth discontinuities. For both classes of functions, we determined the metric entropy and then provided compression strategies based on multiresolution predictive coding in order to realize significant gains in rate–distortion performance. For

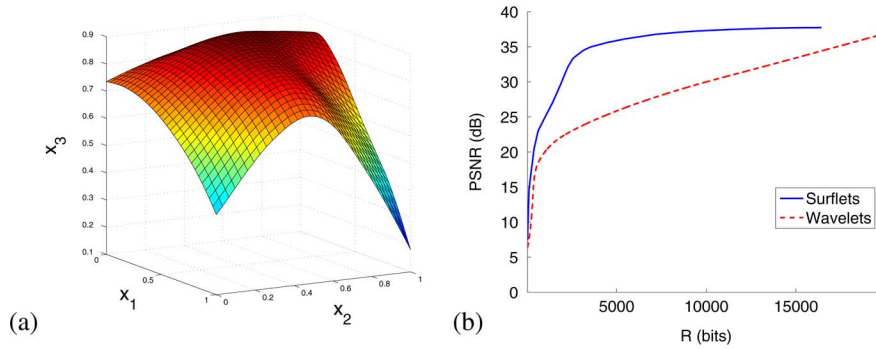


Fig. 8. (a) Horizon b^c used to generate 3-D test function \widetilde{f}_π^c . (b) Rate–distortion performance for surfllet coding compared with wavelet coding.




Fig. 9. Volumetric slices of 3-D coded functions. (a) Original test function \widetilde{f}_π^c from Fig. 8. (b) Surfllet-coded function using 2540 bits; PSNR is 33.22 dB. (c) Wavelet-coded function using approximately 2540 bits; PSNR is 23.08 dB.

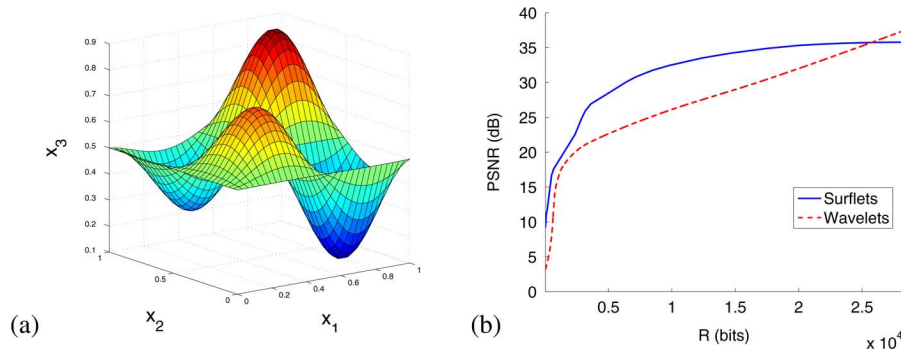


Fig. 10. (a) Horizon b^c used to generate 3-D test function \widetilde{f}_π^c . (b) Rate–distortion performance for surfllet coding compared with wavelet coding.

piecewise constant functions, our surfllet-based compression framework approximates and encodes such functions by assembling piecewise approximations over dyadic hypercubes at different scales, where each surfllet approximant contains a (high-order) polynomial discontinuity that separates two constant regions. This surfllet-based approach achieves the optimal approximation performance and the metric entropy bound. For piecewise smooth functions, we derived surfprints by combining surfllets with wavelets. Our surfprint-based compression framework provides optimal approximation performance and near-optimal rate–distortion performance.

In addition, we extended our results for the continuous signal classes $\mathcal{F}_C(M, K_d)$ and $\mathcal{F}_S(M, K_d, K_s)$ to their corresponding discrete function spaces. We provided asymptotic performance results for both discrete function spaces and related this asymptotic performance to the sampling rate and smoothness orders of the underlying functions and discontinuities. Our simulation results for synthetic 2-D discrete piecewise constant functions demonstrate the coding gains achieved by using higher order

polynomials in the construction of surfllet-based approximations and by defining surfllet dictionaries based on scale-adaptive smoothness orders. Our preliminary 3-D simulation results show that surfllet-based approximations provide improved compression performance for synthetic signals compared to wavelet-based methods over a large range of bit rates.

The insights that we gained, namely, in quantizing higher order terms with lesser precision and using predictive coding to decrease bit rate, can be used to solve more sophisticated signal representation problems. In addition, our methods require knowledge only of the higher dimensional function and not the smooth discontinuity. We believe that the encouraging preliminary simulation results presented in this paper for synthetic signals provide motivation for the future development of surfprint-based coders that are applicable to real-world signals.

Future research will focus on the application of the approximation schemes presented in this paper to statistical estimation of higher dimensional signals containing arbitrary smooth discontinuities given noisy data (extending the piecewise constant

$K_d = 2$, $M = 2$ case treated in [8]). We would like to develop new representation schemes that provide optimal approximation performance for functions containing multiple intersecting discontinuities with high smoothness orders (for $M > 2$). We are also interested in studying practical applications of our coding schemes to the compression of natural images (addressing issues similar to those discussed in [23]), video, and light-field data. Such a study will involve tuning the various parameters related to the quantization schemes, and the construction of the surflet and surfprint dictionaries. Additional work will also improve practical methods for fitting ℓ_2 -best surflets, extending the methods described in [30].

APPENDIX

A. Proof of Theorem 1

Main idea: Let \underline{f}_N^c be an N -term approximant from any representation scheme that provides an approximation to a function $f^c \in \mathcal{F}_C(M, K_d)$ using N terms. We show that we can construct a Horizon-class function from \underline{f}_N^c with the same asymptotic rate as this approximation scheme. As a result, it follows that we only need to consider Horizon-class approximation schemes in establishing a bound on the optimal asymptotic approximation performance for $\mathcal{F}_C(M, K_d)$. (There cannot exist a scheme using non-Horizon approximants that performs better asymptotically.) This connection allows us to directly apply the approximation and metric entropy results pertaining to the $(M - 1)$ -dimensional discontinuity.

Approximation: Let $\mathbf{x} \in [0, 1]^M$ and let \mathbf{y} denote the first $M - 1$ elements of \mathbf{x} (as defined in Section II-B). Define a function \underline{f} such that

$$\underline{f}(\mathbf{x}) = \begin{cases} 1, & \underline{f}_N^c(\mathbf{x}) > 0.5 \\ 0, & \text{otherwise.} \end{cases}$$

Considering the four cases of f^c being 0 or 1 and \underline{f} being 0 or 1, we have

$$\|f^c - \underline{f}\|_{L_2}^2 \leq 4 \cdot \|f^c - \underline{f}_N^c\|_{L_2}^2. \quad (11)$$

Now we construct a Horizon-class function from \underline{f} . Let \widehat{b}_N^c be an $(M - 1)$ -dimensional function defined as

$$\widehat{b}_N^c(\mathbf{y}) = 1 - \int_0^1 \underline{f}(\mathbf{y}, x_M) dx_M.$$

Finally, let \widehat{f}_N^c be a Horizon-class function defined by the $(M - 1)$ -dimensional discontinuity \widehat{b}_N^c

$$\widehat{f}_N^c(\mathbf{x}) = \begin{cases} 1, & x_M \geq \widehat{b}_N^c(\mathbf{y}) \\ 0, & x_M < \widehat{b}_N^c(\mathbf{y}). \end{cases}$$

Again, considering the four cases of f^c being 0 or 1 and \widehat{f}_N^c being 0 or 1, we have

$$\|f^c - \widehat{f}_N^c\|_{L_2}^2 \leq \|f^c - \underline{f}\|_{L_2}^2$$

and using (11) we conclude that

$$\|f^c - \widehat{f}_N^c\|_{L_2}^2 \leq 4 \cdot \|f^c - \underline{f}_N^c\|_{L_2}^2. \quad (12)$$

This result shows that the approximation performance of any scheme that approximates f^c is bounded below by the approximation performance of a corresponding Horizon-class representation scheme.

Because \widehat{f}_N^c is a Horizon-class function

$$\|f^c - \widehat{f}_N^c\|_{L_2}^2 \leq \|b^c - \widehat{b}_N^c\|_{L_1} \quad (13)$$

where b^c is the \mathcal{C}^{K_d} discontinuity in f^c and \widehat{b}_N^c , the implicit estimate to b^c , is the Horizon discontinuity in \widehat{f}_N^c . From the work of Cohen *et al.* [14] regarding optimal approximation rates, the optimal approximation rate for the $(M - 1)$ -dimensional \mathcal{C}^{K_d} class of functions is

$$\|b^c - \widehat{b}_N^c\|_{L_1} \lesssim \left(\frac{1}{N}\right)^{\frac{K_d}{M-1}}. \quad (14)$$

Combining (12)–(14), we have an upper bound on achievable approximation performance for $\mathcal{F}_C(M, K_d)$

$$\|f^c - \widehat{f}_N^c\|_{L_2}^2 \lesssim \left(\frac{1}{N}\right)^{\frac{K_d}{M-1}}. \quad (15)$$

However, (13) is satisfied with equality when both b^c and \widehat{b}_N^c are completely contained inside the unit hypercube (i.e., $0 \leq b^c(\mathbf{y}), \widehat{b}_N^c(\mathbf{y}) \leq 1$), and we also know that (14) provides the *optimal* approximation rate for the $(M - 1)$ -dimensional \mathcal{C}^{K_d} class of functions. Thus, (15) provides the optimal approximation performance that could be achieved for *every* $f^c \in \mathcal{F}_C(M, K_d)$.

Rate distortion: To find the optimal rate–distortion performance for $\mathcal{F}_C(M, K_d)$, a similar argument could be made using the work of Clements [21] (extending Kolmogorov and Tihomirov [20]) regarding metric entropy. It follows from these papers that the optimal asymptotic rate–distortion performance for the $(M - 1)$ -dimensional \mathcal{C}^{K_d} class of functions is

$$\|b^c - \widehat{b}_R^c\|_{L_1} \lesssim \left(\frac{1}{R}\right)^{\frac{K_d}{M-1}}. \quad \square$$

B. Proof of Theorem 2

Let f^s be defined by the $(M - 1)$ -dimensional \mathcal{C}^{K_d} discontinuity b^s separating two M -dimensional \mathcal{C}^{K_s} functions g_1^s, g_2^s . We first establish a lower bound on the optimal approximation rate with respect to the squared- L_2 distortion measure for $\mathcal{F}_S(M, K_d, K_s)$. We note that both the space of M -dimensional uniformly \mathcal{C}^{K_s} -smooth functions and $\mathcal{F}_C(M, K_d)$ are subsets of $\mathcal{F}_S(M, K_d, K_s)$. Cohen *et al.* [14] show that the optimal approximation decay rate for the space of M -dimensional uniformly \mathcal{C}^{K_s} -smooth functions is $(\frac{1}{N})^{\frac{2K_s}{M}}$, while Theorem 1 proves that the optimal approximation rate for $\mathcal{F}_C(M, K_d)$ is $(\frac{1}{N})^{\frac{K_d}{M-1}}$. Therefore, the optimal approximation rate for $\mathcal{F}_S(M, K_d, K_s)$ is bounded below by $(\frac{1}{N})^{\min(\frac{K_d}{M-1}, \frac{2K_s}{M})}$.

We now prove that this lower bound can be achieved, thus establishing the optimal approximation rate for $\mathcal{F}_S(M, K_d, K_s)$. We assume that explicit information about b^s and g_1^s, g_2^s is provided by an external ‘‘oracle.’’ Given such information about the M -dimensional C^{K_s} functions, one could use a wavelet-based approximation scheme [14] to achieve the optimal approximation rate $(\frac{1}{N})^{\frac{2K_s}{M}}$ for such functions. Next, one could use a similar wavelet-based approach in $M - 1$ dimensions to represent b^s with the optimal approximation rate of $(\frac{1}{N})^{\frac{K_d}{M-1}}$. Thus, we have provided the optimal approximation rate for every function $f^s \in \mathcal{F}_S(M, K_d, K_s)$. (The assumption about availability of explicit information about b^s enables to prove the existence of efficient approximations for f^s .) Finally, given the results of Clements [21] (extending Kolmogorov and Tihomirov [20]), the optimal rate–distortion performance for $\mathcal{F}_S(M, K_d, K_s)$ can be derived similarly. \square

C. Proof of Theorem 3

Consider a candidate surflet decomposition grown fully up to level J but pruned back in regions away from the discontinuity to consolidate nodes that are entirely 0- or 1-valued. This surflet decomposition then consists of the following leaf nodes:

- dyadic hypercubes at level J through which the singularity b passes, and which are decorated with a surflet; and
- dyadic hypercubes at various levels through which the singularity b does not pass, and which are all-zero or all-one.

We establish the asymptotic approximation rate for this candidate decomposition. Because this configuration is among the options available to the approximation rate optimized tree-pruning in Section IV-C, this provides an upper bound on the asymptotic approximation rate of the algorithm.

Distortion: First we establish a bound on the distortion in such a decomposition. We assume quantized Taylor surflets for this analysis; this provides an upper bound for the distortion of native L_2 -best surflets as well (since native L_2 -best surflets are chosen from a dictionary that includes the quantized Taylor surflets). In fact, the behavior of the upper bound will also hold for extended L_2 -best surflets, but with slightly larger constants. Let X_J be a dyadic hypercube at level J , and let \mathbf{y}_{ep} be its expansion point. Using Taylor’s theorem (with $D = M - 1, d = b^c, K = K_d, r = r_d, \alpha = \alpha_d$ in (1)), we construct a polynomial approximation of the discontinuity b^c using the Taylor surflet $s^T(X_J; p; \cdot)$ as follows. For each $\mathbf{y} \in Y_J$

$$\begin{aligned} p(\mathbf{y}) &= p(\mathbf{y}_{\text{ep}} + \mathbf{h}) \\ &= [b^c(\mathbf{y}_{\text{ep}}) + c_0 \cdot 2^{-K_d J}] \\ &\quad + \frac{1}{1!} \sum_{i_1=1}^{M-1} [b_{y_{i_1}}^c(\mathbf{y}_{\text{ep}}) + c_{1,i_1} \cdot 2^{-(K_d-1)J}] \cdot h_{i_1} \\ &\quad + \frac{1}{2!} \sum_{i_1, i_2=1}^{M-1} [b_{y_{i_1}, y_{i_2}}^c(\mathbf{y}_{\text{ep}}) \\ &\quad \quad \quad + c_{2,i_1, i_2} \cdot 2^{-(K_d-2)J}] \cdot h_{i_1} h_{i_2} \\ &\quad + \dots \\ &\quad + \frac{1}{r_d!} \sum_{i_1, \dots, i_{r_d}=1}^{M-1} [b_{y_{i_1}, \dots, y_{i_{r_d}}}^c(\mathbf{y}_{\text{ep}}) \\ &\quad \quad \quad + c_{r_d, i_1, \dots, i_{r_d}} \cdot 2^{-\alpha_d J}] \cdot h_{i_1} \dots h_{i_{r_d}} \end{aligned} \quad (16)$$

where each constant $c_{\ell, i_1, \dots, i_{\ell}}$ depends on $b_{y_{i_1}, \dots, y_{i_{\ell}}}^c(\mathbf{y}_{\text{ep}})$ and $|c_{\ell, i_1, \dots, i_{\ell}}| \leq \frac{1}{2}$. The surflet polynomial discontinuity p is constructed by using quantized values of the derivatives of b^c evaluated at \mathbf{y}_{ep} as the polynomial coefficients. The set of all such $s^T(X_J; p; \cdot)$ ’s is precisely $\mathcal{M}(J)$. From Taylor’s theorem (1) and (2), we have that the X_J -clipped L_1 distance between b^c and p is

$$\overline{L_1}(b^c, p) \leq C_3 \cdot 2^{-K_d J - (M-1)J}. \quad (17)$$

Thus, we have that the squared- L_2 error between f^c and $s^T(X_J; p; \cdot)$ over X_J is

$$\|f^c - s^T(X_J; p; \cdot)\|_{L_2(X_J)}^2 \leq C_3 \cdot 2^{-K_d J - (M-1)J}. \quad (18)$$

We construct an approximation \widehat{f}_J^c to f^c at scale J by tiling together all the surflets $s^T(X_J; p; \cdot)$ (where the surflet polynomial p differs from one hypercube to another). Let $N_{j, \text{SL}}$ denote the number of nodes at level j through which the discontinuity b passes. Due to the bounded curvature of b^c , $N_{j, \text{SL}} \lesssim 2^{(M-1)j}$. Therefore, we have that the total distortion is

$$\|f^c - \widehat{f}_J^c\|_{L_2}^2 \leq C_4 \cdot 2^{-K_d J}. \quad (19)$$

Number of terms: Next we establish a bound on the number of surflets required to encode this decomposition (using either quantized Taylor surflets or L_2 -best surflets). We know from the above that due to the bounded curvature of b^c , the number of nodes at level j through which the discontinuity b^c passes is given by $N_{j, \text{SL}} \lesssim 2^{(M-1)j}$. Let $N_{j, \text{ZO}}$ be the number of all-zero and all-one nodes in the pruned decomposition at level j

$$N_{j, \text{ZO}} \leq 2^M \cdot N_{j-1, \text{SL}} \leq 2^M \cdot C_5 \cdot 2^{(M-1)(j-1)} \leq C_6 \cdot 2^{(M-1)j}.$$

Thus, the number of terms required in the approximation is given by

$$N \leq \sum_{j=0}^J (N_{j, \text{SL}} + N_{j, \text{ZO}}) \leq C_7 \cdot 2^{(M-1)J}. \quad (20)$$

Finally, we combine (19) and (20) to obtain the result. \square

D. Proof of Theorem 4

As in Theorem 3, we consider a candidate surflet decomposition grown fully up to level J but pruned back in regions away from the discontinuity. This surflet decomposition then consists of the following nodes:

- leaf nodes at level J , which are decorated with a surflet;
- leaf nodes at various levels, which are all-zero or all-one; and
- internal nodes at various levels, which are decorated with surflets.

Since the leaf nodes are used to construct the approximation, we may use the distortion bound (19) from Theorem 3. This bound holds for both quantized Taylor surflets and L_2 -best surflets.

Number of bits: We establish a bound on the bit rate required to encode this decomposition (using either quantized Taylor surflets or L_2 -best surflets). There are three contributions to the bit rate.

- To encode the structure (topology) of the pruned tree indicating the locations of the leaf nodes, we can use one bit for each node in the tree [32]. Using (20), we have

$$R_1 \leq \sum_{j=0}^J (N_{j,SL} + N_{j,ZO}) \leq C_7 \cdot 2^{(M-1)J}. \quad (21)$$

- For each leaf node that is all-zero or all-one, we can use a constant number of bits to specify the homogeneous nature of the node and the constant value (0 or 1). We have

$$R_2 \leq \sum_{j=0}^J N_{j,ZO} \leq C_8 \cdot 2^{(M-1)J}. \quad (22)$$

- For each leaf node at scale J labeled with a surflet, we must encode the quantized surflet parameters. For a surflet coefficient at scale J of order $\ell \in \{0, \dots, r_d\}$, the number of bits required per coefficient is $O((K_d - \ell)J)$, and the number of such coefficients is $O((M-1)^\ell)$. Hence, the total number of bits required to encode each surflet is $O(J \sum_{\ell=0}^{r_d} (K_d - \ell)(M-1)^\ell) = O(J)$. (Note that our order term describes the scaling with J .) Therefore, we have that

$$R_3 \leq N_{J,SL} \cdot O(J) \leq C_9 \cdot J \cdot 2^{(M-1)J}. \quad (23)$$

Combining (21)–(23), the total bit rate $R(f^c, \widehat{f_j^c})$ required to describe the surflet decomposition satisfies

$$R(f^c, \widehat{f_j^c}) = R_1 + R_2 + R_3 \leq C_{10} \cdot J \cdot 2^{(M-1)J}.$$

We conclude the proof by combining this result with (19). \square

E. Proof of Theorem 5

The proof consists of three steps. First, we show that the L -extended L_2 -best surflet fits to f^c over a hypercube X_j and one of its children X_{j+1} shares a nontrivial common X_{j+1}^L -active region. Second, we use this fact to show that the surflet polynomial coefficients of these L -extended L_2 -best surflet fits are similar in the case where the common X_{j+1}^L -active region is aligned with the center of the hypercube X_{j+1} . Finally, we extend the second step to show that the surflet polynomial coefficients of the L -extended L_2 -best surflet fits are similar regardless of the exact location of the common X_{j+1}^L -active region. We combine these steps to prove that the surflet polynomial coefficients at scale $j+1$ can be encoded using a constant number of bits given the surflet polynomial coefficients at scale j .

Surflet fits over X_j and X_{j+1} share nontrivial active region: Assume that the discontinuity b^c passes through X_{j+1} . (Note that if b^c passes through X_j but not through X_{j+1} , our coder uses an all-0 or all-1 surflet to approximate f^c over X_{j+1} ; checking for such an occurrence does not require explicit knowledge of b^c .) From (17), we have that the L -extended L_2 -best surflet fits $s^L(X_j; p_j; \cdot)$ and $s^L(X_{j+1}; p_{j+1}; \cdot)$ obey the following relations:

$$\overline{L_1}(b^c, p_j) \leq C_3 \cdot 2^{-(K_d+M-1)j} \quad (24)$$

and

$$\overline{L_1}(b^c, p_{j+1}) \leq C_3 \cdot 2^{-(K_d+M-1)(j+1)}. \quad (25)$$

Since b^c passes through X_{j+1} , it is clear that b^c will have a nontrivial $(M-1)$ -dimensional X_{j+1}^L -active region. One can also check that the surflet polynomials p_j and p_{j+1} have similar X_{j+1}^L -active regions. In particular, there exists an $(M-1)$ -dimensional hypercube region A_{j+1} that is contained in the X_{j+1}^L -active regions of b^c , p_j , and p_{j+1} , with side length $\text{h}(A_{j+1}) = C_{11} \cdot 2^{-(j+1)}$ in each dimension with the constant C_{11} independent of j ; otherwise, the bounds (24) and (25) could not hold with b^c passing through X_{j+1} . Hence, $\text{vol}(A_{j+1}) = (C_{11})^{M-1} \cdot 2^{-(M-1)(j+1)}$, where the constant C_{11} depends on the universal derivative bound for functions in $\mathcal{F}_C(M, K_d)$, the extension L , and the dimension of the problem M . We emphasize here that the notion of L -extensions of hypercubes is the key reason that surflet fits at successive scales share such a nontrivial common X_{j+1}^L -active region.

We now restrict the domain of consideration of p_j to the L -extended domain corresponding to the hypercube X_{j+1} at scale $j+1$. We denote the resulting polynomial by p_j^{j+1} . Thus, using (24), (25), and the triangle inequality, the L_1 distance between p_j^{j+1} and p_{j+1} over the $(M-1)$ -dimensional active region A_{j+1} is bounded by

$$\left\| p_j^{j+1} - p_{j+1} \right\|_{L_1(A_{j+1})} \leq C_{12} \cdot 2^{-(K_d+M-1)(j+1)} \quad (26)$$

with the constant C_{12} independent of j .

Surflet polynomial coefficients are similar when A_{j+1} is centered with respect to X_{j+1} : We first assume for simplicity that

$$A_{j+1} = \left[-\frac{C_{11}}{2} \cdot 2^{-(j+1)}, \frac{C_{11}}{2} \cdot 2^{-(j+1)} \right]^{M-1}.$$

In order to relate the similarity (26) between p_j and p_{j+1} in the L_1 sense to their polynomial coefficients, we present the following lemma.

Lemma 1: [33, pp. 72–73] Let $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_D\}$ be a set of linearly independent vectors. Then there exists $C > 0$ such that for any collection of coefficients $\boldsymbol{\gamma} = \{\gamma_i\}_{i=1}^D$

$$|\gamma_i| \leq C \cdot \|\mathbf{v}\|, \quad \forall i \in \{1, \dots, D\}$$

where $\mathbf{v} = \sum_{i=1}^D \gamma_i \mathbf{v}_i$. (For any norm, such a constant exists.)

In order to employ the lemma, we define a monomial basis that contains all monomials of the form $y_1^{i_1} \cdot y_2^{i_2} \cdot \dots \cdot y_{M-1}^{i_{M-1}}$, where $i_1 + i_2 + \dots + i_{M-1} \in \{0, \dots, r_d\}$. We denote by $\mathbf{v}_{i,\ell}$ a monomial basis element of order $i_1 + i_2 + \dots + i_{M-1} = \ell$, with i as an index (i specifies the powers in the basis monomial). At level $j+1$, the domain of each $\mathbf{v}_{i,\ell}$ is restricted to A_{j+1} . We express p_{j+1} and p_j^{j+1} as polynomials in the vector space spanned by the monomial basis

$$p_j^{j+1} = \sum_{i,\ell} a_{i,\ell}^{j+1} \mathbf{v}_{i,\ell}$$

and

$$p_{j+1} = \sum_{i,\ell} a_{i,\ell}^{j+1} \mathbf{v}_{i,\ell}.$$

Now, we define an error vector

$$\mathbf{e}^{j+1} = p_j^{j+1} - p_{j+1}.$$

We define

$$a_{i,\ell}^{j+1} = a_{i,\ell,p_j^{j+1}}^{j+1} - a_{i,\ell,p_{j+1}}^{j+1}$$

and so we have $\mathbf{e}^{j+1} = \sum_{i,\ell} a_{i,\ell}^{j+1} \mathbf{v}_{i,\ell}$. Using (26), over a sub-domain A_{j+1} of volume $(C_{11})^{M-1} \cdot 2^{-(M-1)(j+1)}$

$$\|\mathbf{e}^{j+1}\|_{L_1(A_{j+1})} \leq C_{12} \cdot 2^{-(K_d+M-1)(j+1)}. \quad (27)$$

To complete the proof, we must show that the L_2 -best surflet polynomial coefficients at scale $j+1$ can be encoded using a constant number of bits given the surflet polynomial coefficients at scale j . To do so, note that the quantization bin-size (3) of an ℓ th-order coefficient at scale $j+1$ is $2^{-(K_d-\ell)(j+1)}$. Consequently, it suffices to show that

$$\left| a_{i,\ell}^{j+1} \right| \leq C_{13} \cdot 2^{-(K_d-\ell)(j+1)} \quad (28)$$

with C_{13} independent of j . Such a bound on $|a_{i,\ell}^{j+1}|$ would imply that the L_2 -best surflet coefficients at scale $j+1$ can be encoded using roughly $\log_2(2C_{13} + 1)$ (*constant*) bits, given the L_2 -best surflet coefficients encoded at scale j . We now proceed to establish (28) in the case where

$$A_{j+1} = \left[-\frac{C_{11}}{2} \cdot 2^{-(j+1)}, \frac{C_{11}}{2} \cdot 2^{-(j+1)} \right]^{M-1};$$

in the third part of the proof we will adjust our arguments to establish (28) for arbitrary A_{j+1} .

Since we only have bounds on the L_1 distance between p_{j+1} and p_j^{j+1} over A_{j+1} , we restrict our attention to this shared active region. Normalizing the basis vectors $\mathbf{v}_{i,\ell}$ with respect to the domain A_{j+1} , we have that

$$\mathbf{e}^{A_{j+1}} = \sum_{i,\ell} c_{i,\ell}^{A_{j+1}} \mathbf{w}_{i,\ell}^{A_{j+1}}$$

where

$$\mathbf{w}_{i,\ell}^{A_{j+1}} = \mathbf{v}_{i,\ell} / \|\mathbf{v}_{i,\ell}\|_{L_1(A_{j+1})}$$

$$\|\mathbf{v}_{i,\ell}\|_{L_1(A_{j+1})} = \int_{A_{j+1}} |\mathbf{v}_{i,\ell}|$$

and

$$c_{i,\ell}^{A_{j+1}} = a_{i,\ell}^{j+1} \cdot \|\mathbf{v}_{i,\ell}\|_{L_1(A_{j+1})}.$$

Let $\mathbf{v}_{i,\ell} = y_1^{i_1} \cdot y_2^{i_2} \cdots y_{M-1}^{i_{M-1}}$ be a basis monomial with $i_1 + i_2 + \cdots + i_{M-1} = \ell$. From the definition of the $\|\cdot\|_{L_1(A_{j+1})}$ norm, we have that

$$\begin{aligned} & \|\mathbf{v}_{i,\ell}\|_{L_1(A_{j+1})} \\ &= \int_{\mathbf{y} \in \left[-\frac{C_{11}}{2} \cdot 2^{-(j+1)}, \frac{C_{11}}{2} \cdot 2^{-(j+1)} \right]^{M-1}} \left| y_1^{i_1} \cdots y_{M-1}^{i_{M-1}} \right| d\mathbf{y} \\ &= 2^{-\ell} C_{11}^{\ell+M-1} \frac{2^{-\ell(j+1)} \cdot 2^{-(M-1)(j+1)}}{(i_1+1) \cdots (i_{M-1}+1)}. \end{aligned} \quad (29)$$

Note that C_{11} is independent of j . Because the basis vectors $\mathbf{w}_{i,\ell}^{A_{j+1}}$ are linearly independent over A_{j+1} , we know from the lemma that there exists a $C(j+1)$ such that

$$|c_{i,\ell}^{A_{j+1}}| \leq C(j+1) \cdot \|\mathbf{e}^{A_{j+1}}\|_{L_1(A_{j+1})}. \quad (30)$$

We need to show that $C(j+1)$ is actually independent of j . This would allow us to conclude (28) for $A_{j+1} = \left[-\frac{C_{11}}{2} \cdot 2^{-(j+1)}, \frac{C_{11}}{2} \cdot 2^{-(j+1)} \right]^{M-1}$ because

$$\begin{aligned} \left| a_{i,\ell}^{j+1} \right| &= \left| \frac{c_{i,\ell}^{A_{j+1}}}{\|\mathbf{v}_{i,\ell}\|_{L_1(A_{j+1})}} \right| \\ &= \left| \frac{c_{i,\ell}^{A_{j+1}} \cdot (i_1+1) \cdots (i_{M-1}+1)}{2^{-\ell} \cdot C_{11}^{\ell+M-1} \cdot 2^{-\ell(j+1)} \cdot 2^{-(M-1)(j+1)}} \right| \\ &\leq \left[\frac{C(j+1) \cdot C_{12} \cdot (i_1+1) \cdots (i_{M-1}+1)}{2^{-\ell} \cdot C_{11}^{\ell+M-1}} \right] \\ &\quad \times 2^{-(K_d-\ell)(j+1)} \end{aligned} \quad (31)$$

where we obtain the second equality from (29), and the inequality from (27) and (30). Indeed, if $C(j+1)$ is independent of j , we note from (31) that we could set $C_{13} = [C(j+1) \cdot C_{12} \cdot (i_1+1) \cdots (i_{M-1}+1)] / [2^{-\ell} \cdot C_{11}^{\ell+M-1}]$ in (28).

We let $\xi \in \left[-\frac{C_{11}}{2}, \frac{C_{11}}{2} \right]^{M-1}$ denote the ‘‘relative position’’ within the active hypercube region of a surflet. For *any* level j_1 with $A_{j_1} = \left[-\frac{C_{11}}{2} \cdot 2^{-j_1}, \frac{C_{11}}{2} \cdot 2^{-j_1} \right]^{M-1}$

$$\frac{\mathbf{w}_{i,\ell}^{A_{j+1}}(\xi \cdot 2^{-(j+1)})}{\mathbf{w}_{i',\ell'}^{A_{j+1}}(\xi \cdot 2^{-(j+1)})} = \frac{\mathbf{w}_{i,\ell}^{A_{j_1}}(\xi \cdot 2^{-j_1})}{\mathbf{w}_{i',\ell'}^{A_{j_1}}(\xi \cdot 2^{-j_1})}$$

from (29). Setting $\ell' = 0$, we have that

$$\mathbf{w}_{i,\ell}^{A_{j+1}}(\xi \cdot 2^{-(j+1)}) = 2^{-(M-1)(j_1-j-1)} \cdot \mathbf{w}_{i,\ell}^{A_{j_1}}(\xi \cdot 2^{-j_1}). \quad (32)$$

Thus, we can construct a vector $\mathbf{e}^{A_{j_1}}$ at level j_1 using the *same* coefficients $c_{i,\ell}^{A_{j+1}}$, i.e., $c_{i,\ell}^{A_{j_1}} = c_{i,\ell}^{A_{j+1}}$ so that

$$\begin{aligned} & \|\mathbf{e}^{A_{j+1}}\|_{L_1(A_{j+1})} \\ &= \int_{\mathbf{y} \in A_{j+1}} \left| \sum_{i,\ell} c_{i,\ell}^{A_{j+1}} \mathbf{w}_{i,\ell}^{A_{j+1}}(\mathbf{y}) \right| \cdot d\mathbf{y} \\ & \quad \left(\text{Substituting } \xi_i = y_i \cdot 2^{(j+1)} \right) \\ &= \int_{\xi \in \left[-\frac{C_{11}}{2}, \frac{C_{11}}{2} \right]^{M-1}} \left[\left| \sum_{i,\ell} c_{i,\ell}^{A_{j+1}} \mathbf{w}_{i,\ell}^{A_{j+1}}(\xi \cdot 2^{-(j+1)}) \right| \right. \\ & \quad \left. \times 2^{-(M-1)(j+1)} \cdot d\xi \right] \\ &= \int_{\xi \in \left[-\frac{C_{11}}{2}, \frac{C_{11}}{2} \right]^{M-1}} \left[\left| \sum_{i,\ell} c_{i,\ell}^{A_{j+1}} \mathbf{w}_{i,\ell}^{A_{j_1}}(\xi \cdot 2^{-j_1}) \right| \right. \\ & \quad \left. \times 2^{-(M-1)j_1} \cdot d\xi \right] \\ & \quad \left(\text{Substituting } y_i = \xi_i \cdot 2^{-j_1}; c_{i,\ell}^{A_{j_1}} = c_{i,\ell}^{A_{j+1}} \right) \\ &= \int_{\mathbf{y} \in A_{j_1}} \left| \sum_{i,\ell} c_{i,\ell}^{A_{j_1}} \mathbf{w}_{i,\ell}^{A_{j_1}}(\mathbf{y}) \right| \cdot d\mathbf{y} \\ &= \|\mathbf{e}^{A_{j_1}}\|_{L_1(A_{j_1})} \end{aligned}$$

where we get the third equality from (32). Since the coefficients at the two levels, $j + 1$ and j_1 , are the same, we can set $C_{14} = C(j + 1)$ in (30)

$$\left| c_{i,\ell}^{A_{j_1}} \right| \leq C_{14} \cdot \|e^{A_{j_1}}\|_{L_1(A_{j_1})}.$$

In this manner, one can show that

$$\left| c_{i,\ell}^{A_{j+1}} \right| \leq C_{14} \cdot \|e^{A_{j+1}}\|_{L_1(A_{j+1})} \quad (33)$$

is true for all j (because j_1 is arbitrary), and hence C_{14} is independent of j . Switching back to the original coefficients $a_{i,\ell}^{j+1}$, we have that

$$\left| a_{i,\ell}^{j+1} \right| \leq C_{13} \cdot 2^{-(K_d - \ell)(j+1)} \quad (34)$$

following the logic in (31), with C_{13} independent of j .

Surflet polynomial coefficients are similar independent of the specific location of A_{j+1} : We now suppose $A_{j+1} = [-\frac{C_{11}}{2} \cdot 2^{-(j+1)} + \eta_1, \frac{C_{11}}{2} \cdot 2^{-(j+1)} + \eta_1] \times \dots \times [-\frac{C_{11}}{2} \cdot 2^{-(j+1)} + \eta_{M-1}, \frac{C_{11}}{2} \cdot 2^{-(j+1)} + \eta_{M-1}]$, where η denotes the ‘‘center’’ of A_{j+1} . As in the previous step, let $e^{j+1} = \sum_{i,\ell} a_{i,\ell}^{j+1} \mathbf{v}_{i,\ell}$. Suppose that we transform the basis vectors $\mathbf{v}_{i,\ell}$ to be centered around η rather than around $\mathbf{0}$. The transformed coefficients would then satisfy the bound in (34). To make this point precise, let

$$\mathbf{v}_{i,\ell}^{A_{j+1}}(\mathbf{y}) = \mathbf{v}_{i,\ell}(\mathbf{y} - \eta) = (y_1 - \eta_1)^{i_1} \dots (y_{M-1} - \eta_{M-1})^{i_{M-1}}$$

and let $a_{i,\ell}^{A_{j+1}}$ denote the transformed coefficients of the error vector so that $e^{j+1} = \sum_{i,\ell} a_{i,\ell}^{A_{j+1}} \mathbf{v}_{i,\ell}^{A_{j+1}}$. We have that

$$\left| a_{i,\ell}^{A_{j+1}} \right| \leq C_{13} \cdot 2^{-(K_d - \ell)(j+1)} \quad (35)$$

from (34). In order to complete the proof, we need to show that

$$\left| a_{i,\ell}^{j+1} \right| \leq C_{15} \cdot 2^{-(K_d - \ell)(j+1)} \quad (36)$$

with C_{15} independent of j . Note that the key difference between this step and the result of the previous step of the proof is that A_{j+1} is not necessarily aligned with the center of the hypercube X_{j+1} . Hence, proving (36) is more general than the result in the previous step. Using the binomial theorem, we have the following relationship between the two sets of coefficients:

$$a_{i,\ell}^{j+1} = \sum_{(i',\ell') \in \mathcal{S}(i,\ell)} a_{i',\ell'}^{A_{j+1}} \left(\prod_{k=1}^{M-1} \binom{i'_k}{i'_k - i_k} (-\eta_k)^{i'_k - i_k} \right)$$

where $\mathcal{S}(i,\ell) = \{(i',\ell') : \ell' \geq \ell; i'_1 \geq i_1, \dots, i'_{M-1} \geq i_{M-1}; i'_1 + i'_2 + \dots + i'_{M-1} = \ell'\}$. The outer sum is finite and $|\eta_k| \leq C_{16} \cdot 2^{-(j+1)}$ for all k , thus establishing (36)

$$\begin{aligned} \left| a_{i,\ell}^{j+1} \right| &\leq \sum_{(i',\ell') \in \mathcal{S}(i,\ell)} \left| a_{i',\ell'}^{A_{j+1}} \right| \left(\prod_{k=1}^{M-1} C_{17} \cdot |\eta_k|^{i'_k - i_k} \right) \\ &\leq \sum_{(i',\ell') \in \mathcal{S}(i,\ell)} \left[\left| a_{i',\ell'}^{A_{j+1}} \right| \cdot (C_{17} \cdot C_{16})^{M-1} \right. \\ &\quad \left. \times 2^{-(j+1) \sum_{k=1}^{M-1} i'_k - i_k} \right] \end{aligned}$$

$$\begin{aligned} &\leq \sum_{(i',\ell') \in \mathcal{S}(i,\ell)} \left[C_{13} \cdot 2^{-(K_d - \ell')(j+1)} \right. \\ &\quad \left. \times (C_{17} \cdot C_{16})^{M-1} \cdot 2^{-(j+1)(\ell' - \ell)} \right] \\ &\leq C_{15} \cdot 2^{-(K_d - \ell)(j+1)}. \end{aligned}$$

Here, we use (35) for the third inequality. \square

F. Proof of Theorem 6

We begin by providing a simple performance bound (which we improve upon below) for the approximation scheme described in Section IV-C applied to $f_{\#}^c$. At scale J , the $(M - 2)$ -dimensional intersection manifold passes through $O(2^{(M-2)J})$ hypercubes. Let one such hypercube be denoted by $X_{J,\#}$. The squared- L_2 error in $X_{J,\#}$ due to the use of a surflet (that is ill-suited for representing intersections) can be approximated by the volume of the hypercube and is equal to $O(2^{-MJ})$. Therefore, the total squared- L_2 approximation error in representing the $(M - 2)$ -dimensional intersection manifold is given by $O(2^{-2J})$ (for every M). Comparing this result to the bound in (19), we see that the approximation performance achieved by our surflet-based representation scheme applied to $f_{\#}^c$ is $\|f_{\#}^c - \widehat{f_{\#}^c}\|_{L_2}^2 \lesssim (\frac{1}{N})^{M-1}$. Hence, the representation scheme from Section IV-C applied to $f_{\#}^c$ achieves optimal approximation performance for $K_d = 2$, but suboptimal performance for $K_d > 2$. Thus, the smoothness threshold $K_d^{\text{th}} = 2$. This performance is still better than that achieved by wavelets, which treat the discontinuities as C^1 -smooth functions regardless of any additional structure in the discontinuities, and thus have a smoothness threshold of $K_d^{\text{th,wl}} = 1$ (i.e., the performance achieved by wavelets is $\|f_{\#}^c - \widehat{f_{\#}^c}\|_{L_2}^2 \lesssim (\frac{1}{N})^{M-1}$).

Using more sophisticated analysis, we now improve upon the performance bound described above to show that the approximation scheme has a smoothness threshold K_d^{th} greater than two, thus increasing the range of K_d for which we achieve optimal performance. In order to improve the performance bound, we consider the scenario where the approximation scheme further subdivides those hypercubes at scale J containing intersections. At scale J , the scheme described in Section IV-C uses $N = O(2^{(M-1)J})$ terms to construct an approximation (analyzed in detail in Appendix C). This suggests that an additional $O(2^J)$ mini-hypercubes could be used within each of the $O(2^{(M-2)J})$ hypercubes that contain the $(M - 2)$ -dimensional intersection manifold (again, let one such hypercube be denoted by $X_{J,\#}$), while still maintaining the same order for N . Let the side length of the mini-hypercubes be 2^{-s} . Each of these mini-hypercubes is labeled with a surflet (chosen from a dictionary at scale s). Within each $X_{J,\#}$, the approximation scheme would ideally use the $O(2^J)$ mini-hypercubes to only approximate those regions that contain the intersection manifold rather than to tile the entire hypercube $X_{J,\#}$. Using this idea, we compute the side length, and consequently the smoothness threshold K_d^{th} , as follows.

- $M > 2$: The number of mini-hypercubes used in each $X_{J,\#}$ is $O(2^J)$. We would like all these mini-hypercubes of side length 2^{-s} to approximate regions within $X_{J,\#}$ that contain the intersection manifold. This implies that $2^J \sim 2^{(M-2)(s-J)}$, which results in $s = \lceil \frac{(M-1)J}{M-2} \rceil$. The total

squared- L_2 error due to the use of these mini-hypercubes in each $X_{J,\#}$ is

$$O(2^J \cdot 2^{-Ms}) = O\left(2^J \cdot 2^{-\frac{M(M-1)J}{M-2}}\right).$$

The total squared- L_2 error over all of the $O(2^{(M-2)J})$ $X_{J,\#}$'s due to the mini-hypercubes is given by

$$O\left(2^J \cdot 2^{-\frac{M(M-1)J}{M-2}} \cdot 2^{(M-2)J}\right) = O\left(2^{-\frac{2(M-1)J}{M-2}}\right).$$

Comparing with (19), we have that $K_d^{\text{th}} = \frac{2(M-1)}{M-2}$.

- $M = 2$: In this case, discontinuities intersect at points. Therefore, only a constant number of mini-hypercubes are needed inside each $X_{J,\#}$. As a result, the number of mini-hypercubes that are required to cover the intersection manifold does not grow with scale. Choosing $s = \lceil \frac{K_d J}{M} \rceil$, we see that the total squared- L_2 error due to the use of these mini-hypercubes in each $X_{J,\#}$ is $O(2^{-Ms}) = O(2^{-K_d J})$. The total squared- L_2 error over all of the hypercubes $X_{J,\#}$ (a constant number) due to the mini-hypercubes is also given by $O(2^{-K_d J})$. Comparing with (19), we see that the scheme in Section IV-C achieves optimal approximation performance for every K_d .

Note that the analysis of the approximation scheme as described above requires explicit information about the location of the intersections; however, an approximation scheme based on L_2 -best surflets (with the dictionary containing regular and “mini”-surflets) would not require such explicit information but would still achieve the same performance. \square

G. Proof of Theorem 7

According to the prototype algorithm, there are three sources of error in the approximation—quantizing wavelet coefficients that are kept in the tree, pruning Type S nodes (and their descendant subtrees) from scale $\frac{J}{m}$ to scale $\frac{J}{n}$, and approximating Type D discontinuity nodes at scale $\frac{J}{n}$ (and their descendant subtrees) by surfprint atoms. The terms used in constructing the approximation include Type S wavelet coefficients, Type D wavelet coefficients, and surfprint atoms. We will analyze the approximation error and number of terms separately before calculating the approximation rate of the surfprint-based representation.

Distortion: The sources of distortion contribute in the following manner.

- *Quantizing wavelet coefficients:* There are $N_{1:\frac{J}{m}} = O(2^{\frac{MJ}{m}})$ wavelet nodes up to scale $\frac{J}{m}$. The number of Type D nodes (which are not pruned) at scales $j \in [\frac{J}{m} + 1, \frac{J}{n}]$ is given by $N_{j,D} = O(2^{(M-1)j})$. This implies that the total number of Type D nodes at these scales is bounded by

$$N_{\frac{J}{m}+1:\frac{J}{n},D} = O\left(2^{\frac{J(M-1)}{n}}\right) = O\left(2^{\frac{MJ}{n}}\right)$$

because $\frac{m}{n} = \frac{M}{M-1}$. Using (8), each wavelet coefficient in the approximation is quantized up to resolution $(\Delta K_s^{\text{wl}})^2 = 2^{-\frac{J}{m}(2K_s^{\text{wl}}+M)}$. Therefore, the total quantization distortion is $O(2^{-\frac{2K_s^{\text{wl}}J}{m}})$.

- *Pruning Type S nodes:* First, we consider Type S nodes at scale $\frac{J}{m}$. The magnitude of wavelet coefficients for Type S smooth nodes decays as $O(2^{-(K_s^{\text{wl}}+M/2)j})$ [26]. The squared- L_2 error from a single pruning at scale $\frac{J}{m}$ is given by

$$\sum_{j=\frac{J}{m}}^{\infty} 2^{M(j-\frac{J}{m})} \cdot 2^{-j(2K_s^{\text{wl}}+M)} \lesssim 2^{-\frac{J}{m}(M+2K_s^{\text{wl}})} \quad (37)$$

where a Type S node at scale J/m has $2^{M(j-\frac{J}{m})}$ children nodes at scale j . There are $N_{\frac{J}{m},S} = O(2^{\frac{MJ}{m}})$ Type S nodes at scale $\frac{J}{m}$. Therefore, the total distortion from pruning Type S nodes at scale $\frac{J}{m}$ is $O(2^{-\frac{2K_s^{\text{wl}}J}{m}})$. Second, we consider Type S nodes (not previously pruned) at deep scales greater than $\frac{J}{m}$. The error given by (37) also serves as an upper bound for every Type S pruning from scale $\frac{J}{m} + 1$ to scale $\frac{J}{n}$. For a Type S node at these scales to have not been previously pruned, it must have a Type D parent. Because $N_{\frac{J}{m}+1:\frac{J}{n},D} = O(2^{\frac{MJ}{n}})$, the total error due to pruning is

$$O\left(2^{-\frac{2K_s^{\text{wl}}J}{m}} + 2^{-\frac{J(M+2K_s^{\text{wl}})}{m}} \cdot 2^{\frac{MJ}{m}}\right) = O\left(2^{-\frac{2K_s^{\text{wl}}J}{m}}\right).$$

- *Using surfprint approximations:* The number of Type D discontinuity nodes at scale $\frac{J}{n}$ (approximated by surfprint atoms) is $N_{\frac{J}{n},D} = N_{\frac{J}{n},\text{SP}} = O(2^{\frac{(M-1)J}{n}})$. The error due to each surfprint approximation is given by

$$O\left(2^{-\frac{K_d^{\text{SP}}J}{n}} \cdot 2^{-\frac{(M-1)J}{n}} + 2^{-\frac{(2K_s^{\text{SP}}+M)J}{n}}\right).$$

(This error is bounded by the squared- L_2 error of the quantized piecewise polynomial surflet approximation over each hypercube $X_{J/n}$, extended if necessary to cover the supports of the wavelets.) Therefore, the total error due to surfprint approximations is given by

$$O\left(2^{-\min\left(\frac{K_d^{\text{SP}}J}{n}, \frac{(2K_s^{\text{SP}}+1)J}{n}\right)}\right).$$

Thus, the total squared distortion is given by

$$\|f^s - \widehat{f}_J^s\|_{L_2}^2 = O\left(2^{-\min\left(\frac{2K_s^{\text{wl}}J}{m}, \frac{K_d^{\text{SP}}J}{n}, \frac{(2K_s^{\text{SP}}+1)J}{n}\right)}\right). \quad (38)$$

Number of Terms: The following three types of terms are used in assembling an approximation to $f^s \in \mathcal{F}_S(M, K_d, K_s)$.

- *Coarse-scale wavelets:* The total number wavelets used at coarse scales is

$$N_{1:\frac{J}{m}} = O\left(2^{\frac{MJ}{m}}\right) = O\left(2^{(M-1)\frac{J}{n}}\right).$$

- *Intermediate Type D wavelets:* The number of Type D nodes used at intermediate scales is $N_{\frac{J}{m}+1:\frac{J}{n},D} = O(2^{\frac{(M-1)J}{n}})$.
- *Surfprints:* The total number of surfprints used in the approximation is $N_{\frac{J}{n},\text{SP}} = O(2^{\frac{(M-1)J}{n}})$.

Thus, the total number of terms, N , used in assembling the approximation is

$$\begin{aligned} N &= N_{1:\frac{J}{m}} + N_{\frac{J}{m}+1:\frac{J}{n},D} + N_{\frac{J}{n},SP} = O\left(2^{\frac{(M-1)J}{n}}\right) \\ &= O\left(2^{\frac{MJ}{m}}\right). \end{aligned} \quad (39)$$

Combining (38) and (39), we get the following approximation rate for the performance of our prototype surfprint-based approximation scheme:

$$\left\|f^s - \widehat{f}_N^s\right\|_{L_2}^2 \lesssim \left(\frac{1}{N}\right)^{\min\left(\frac{K_d^{SP}}{M-1}, \frac{2K_s^{wl}}{M}, \frac{2K_s^{SP}+1}{M-1}\right)}.$$

The conditions on the operational smoothness orders K_d^{SP} , K_s^{wl} , and K_s^{SP} (as specified in Section V-C) ensure that the proper terms dominate the decay in this expression and it matches the optimal asymptotic approximation rate of Theorem 2. \square

H. Proof of Theorem 8

Consider an N -term approximation \widehat{f}_N^s to f^s constructed by the scheme in Section V-E. The distortion between \widehat{f}_N^s and f^s is given by (38). We only need to analyze the number of bits required to encode this approximation.

Number of bits: We encode the topology of the tree and the quantized wavelet and surfprint terms.

- To encode the structure of the tree, we use $O(1)$ bits to encode each node in the tree

$$R_1 = O(N) = O\left(2^{\frac{(M-1)J}{n}}\right) = O\left(2^{\frac{MJ}{m}}\right) \quad (40)$$

from (39).

- The number of possible quantization bins for a wavelet coefficient $\langle f^s, w_{X_j} \rangle$ at scale j is given by

$$\text{Bins}(j) = O\left(2^{\frac{j}{m}(K_s^{wl} + \frac{M}{2}) - j}\right)$$

based on the quantization step size (8) and the fact that a wavelet coefficient at scale j near the $(M-1)$ -dimensional discontinuity decays as 2^{-j} [26]. Thus, the number of bits required to encode wavelet coefficients is given by

$$\begin{aligned} R_2 &= \sum_{j=0}^{\frac{J}{m}} (N_{j,S} + N_{j,D}) \log(\text{Bins}(j)) \\ &\quad + \sum_{j=\frac{J}{m}+1}^{\frac{J}{n}} N_{j,D} \log(\text{Bins}(j)) \\ &= O\left(J2^M \frac{J}{m}\right). \end{aligned} \quad (41)$$

- The number of bits required to encode surfprint coefficients at scale $\frac{J}{n}$ is

$$R_3 = N_{\frac{J}{n},SP} O\left(\left(K_d^{SP} + K_s^{SP}\right) \frac{J}{n}\right) = O\left(J2^{(M-1)\frac{J}{n}}\right). \quad (42)$$

Combining (38), (40), (41), and (42), we obtain the desired result. \square

I. Proof of Theorem 9

Let X_J be an M -dimensional dyadic hypercube¹² with $J \leq J_{\text{vox}} := \frac{\min(\pi_i)}{K_d}$. Let \widehat{f}_L^c be the continuous L_2 -best surflet fit to f^c over X_J . We know from (4) that

$$\left\|f^c - \widehat{f}_L^c\right\|_{L_2(X_J)}^2 = O\left(2^{-J(K_d+M-1)}\right). \quad (43)$$

We assume that the values of the discretized function \widetilde{f}_π^c are obtained by averaging f^c over each voxel. The distortion between f^c and \widetilde{f}_π^c over X_J is nonzero only over voxels through which the discontinuity b^c passes. The squared- L_2 distortion over each such voxel is $O(2^{-(\pi_1+\dots+\pi_M)})$. Also, the number of voxels through which b^c passes within $X_J = O(2^{\sum_{i=1}^{M-1}(\pi_i-J)}[\Omega \cdot 2^{-\min(\pi_i)}] / (2^{-\pi_M}))$, where Ω is the universal derivative bound (Section II-A). Thus, we have

$$\begin{aligned} \left\|f^c - \widetilde{f}_\pi^c\right\|_{L_2(X_J)}^2 &= O\left(2^{-J(M-1)-\min(\pi_i)}\right) \\ &= O\left(2^{-J(K_d+M-1)}\right) \end{aligned} \quad (44)$$

where the second equality is due to fact that $J \leq \frac{\min(\pi_i)}{K_d}$. Note that we define \widetilde{f}_π^c as a continuous function (constant over each M -dimensional voxel) in order to compare with f^c . Similarly, one can check that

$$\left\|\widehat{f}_L^c - \widetilde{f}_{L,\pi}^c\right\|_{L_2(X_J)}^2 = O\left(2^{-J(K_d+M-1)}\right) \quad (45)$$

where $\widetilde{f}_{L,\pi}^c$ is the sampled version of \widehat{f}_L^c .

Equations (44) and (45) indicate that at scale J , voxelization effects are comparable to the approximations afforded by surflets. Essentially, then, all of the approximation results for surflets at this scale are preserved when applied to the voxelized function. In particular, combining (43)–(45), we have the following result:

$$\begin{aligned} &\left\|\widetilde{f}_\pi^c - \widetilde{f}_{L,\pi}^c\right\|_{L_2(X_J)}^2 \\ &\leq \left(\left\|f^c - \widetilde{f}_\pi^c\right\|_{L_2(X_J)} + \left\|f^c - \widehat{f}_L^c\right\|_{L_2(X_J)}\right. \\ &\quad \left.+ \left\|\widehat{f}_L^c - \widetilde{f}_{L,\pi}^c\right\|_{L_2(X_J)}\right)^2 \\ &= O\left(2^{-J(K_d+M-1)}\right). \end{aligned}$$

Thus, discrete surflets are as effective on the discrete block as continuous surflets are on the corresponding continuous block (see Appendix C). However, we are only provided with the discrete function \widetilde{f}_π^c and would like to use ℓ_2 -best surflets on dyadic blocks of \widetilde{f}_π^c . Let \widehat{f}_π^c denote the discrete ℓ_2 -best surflet fit to \widetilde{f}_π^c over X_J .

¹²We omit an additive constant that may be added to J_{vox} to ensure a more exact agreement with the voxelization breakdown rate.

By definition, $\|\widetilde{f_{\boldsymbol{\pi}}^c} - \widehat{f_{\boldsymbol{\pi}}^c}\|_{L_2(X_J)}^2 \leq \|\widetilde{f_{\boldsymbol{\pi}}^c} - \widehat{f_{L,\boldsymbol{\pi}}^c}\|_{L_2(X_J)}^2$. Thus, we have that

$$\left\| \widetilde{f_{\boldsymbol{\pi}}^c} - \widehat{f_{\boldsymbol{\pi}}^c} \right\|_{L_2(X_J)}^2 = O\left(2^{-J(K_d+M-1)}\right). \quad (46)$$

It remains to be shown that ℓ_2 -best surflets can be predicted across scales. The proof of this fact is analogous to the Proof of Theorem 5. \square

J. Proof of Corollary 10

Combining (44) and (46) from Appendix I, we have the following result:

$$\begin{aligned} & \left\| f^c - \widehat{f_{\boldsymbol{\pi}}^c} \right\|_{L_2(X_J)}^2 \\ & \leq \left(\left\| f^c - \widetilde{f_{\boldsymbol{\pi}}^c} \right\|_{L_2(X_J)} + \left\| \widetilde{f_{\boldsymbol{\pi}}^c} - \widehat{f_{\boldsymbol{\pi}}^c} \right\|_{L_2(X_J)} \right)^2 \\ & = O\left(2^{-J(K_d+M-1)}\right). \end{aligned}$$

Hence, the quality of the approximation provided by the ℓ_2 -best surflets $\widehat{f_{\boldsymbol{\pi}}^c}$ operating on $\widetilde{f_{\boldsymbol{\pi}}^c}$ to the continuous function f^c , with $\widetilde{f_{\boldsymbol{\pi}}^c}$ considered to be continuous (constant over each M -dimensional voxel), is as good as the approximation performance provided by continuous L_2 -best surflet fits to f^c (see Appendix C). The rest of the proof follows in an analogous manner to the Proof of Theorem 5. \square

K. Proof of Theorem 11

In smooth regions of $\widetilde{f_{\boldsymbol{\pi}}^s}$, discrete wavelet coefficients (corresponding to a discrete wavelet basis with K_s^{wl} vanishing moments applied to $\widetilde{f_{\boldsymbol{\pi}}^s}$) decay as $O(2^{-(K_s^{\text{wl}}+M/2)j})$ [26]. This is the same decay rate that continuous wavelet coefficients obey on smooth regions of f^s (see Appendix G). Therefore, the total error due to the use of discrete wavelets to approximate $\widetilde{f_{\boldsymbol{\pi}}^s}$, the quantization of the corresponding discrete wavelet coefficients, and the pruning of Type S nodes is analogous to the corresponding error in the continuous case analyzed in Appendix G.

What remains to be analyzed is the distortion due to the use of discrete surfprint approximations. This analysis is analogous to the analysis in Appendix I. Consider a Type D node at scale $\frac{J}{n}$ to which a surfprint approximation is assigned. Let the hypercube corresponding to this node be $X_{\frac{J}{n}}$. First, we have the following distortion between f^s and the L_2 -best surfprint fit to f^s over $X_{\frac{J}{n}}$, $\widehat{f_{\text{sp},L}^s}$ (see Appendix G):

$$\begin{aligned} & \left\| f^s - \widehat{f_{\text{sp},L}^s} \right\|_{L_2(X_{\frac{J}{n}})}^2 \\ & = O\left(2^{-\frac{(K_d^{\text{sp}}+M-1)J}{n}} + 2^{-\frac{(2K_s^{\text{sp}}+M)J}{n}}\right) \\ & = O\left(2^{-\frac{(M-1)J}{n}} \cdot 2^{-\min(K_d^{\text{sp}}, 2K_s^{\text{sp}}+1)\frac{J}{n}}\right). \quad (47) \end{aligned}$$

Second, we characterize the error due to voxelization of f^s over $X_{\frac{J}{n}}$. For each voxel in $X_{\frac{J}{n}}$ through which the discontinuity passes, the squared- L_2 error between f^s and $\widetilde{f_{\boldsymbol{\pi}}^s}$ is given by $O(2^{-(\pi_1+\dots+\pi_M)})$. The number of such voxels in

$$X_{\frac{J}{n}} = O\left(2^{\sum_{i=1}^{M-1}(\pi_i - \frac{J}{n})} \lceil (\Omega \cdot 2^{-\min(\pi_i)_{i=1}^{M-1}}) / (2^{-\pi_M}) \rceil \right)$$

where Ω is the universal derivative bound (Section II-A). Discretization over the smooth regions of f^s can be viewed as a construction of 0th-order Taylor approximations (constants) locally over each voxel. Therefore, the squared- L_2 error between f^s and $\widetilde{f_{\boldsymbol{\pi}}^s}$ over a voxel in the smooth region is $O(2^{-2\min(\pi_i)} \cdot 2^{-(\pi_1+\dots+\pi_M)})$. The number of such voxels in $X_{\frac{J}{n}} = O(2^{\sum_{i=1}^M(\pi_i - \frac{J}{n})})$. Thus, the total error between f^s and $\widetilde{f_{\boldsymbol{\pi}}^s}$ is

$$\begin{aligned} & \left\| f^s - \widetilde{f_{\boldsymbol{\pi}}^s} \right\|_{L_2(X_{\frac{J}{n}})}^2 \\ & = O\left(2^{-\frac{(M-1)J}{n} - \min(\pi_i)} + 2^{-\frac{MJ}{n} - 2\min(\pi_i)}\right) \\ & = O\left(2^{-\frac{(M-1)J}{n} - \min(\pi_i)}\right) \\ & = O\left(2^{-\frac{(M-1)J}{n}} \cdot 2^{-\min(K_d^{\text{sp}}, 2K_s^{\text{sp}}+1)\frac{J}{n}}\right). \quad (48) \end{aligned}$$

The last equality is due to the assumption that

$$\frac{J}{n} \leq \frac{\min\{\pi_i\}}{\min\{K_d^{\text{sp}}, 2K_s^{\text{sp}}+1\}}.$$

Similarly, the squared- L_2 error between a continuous surfprint atom and its discrete analog at scale $\frac{J}{n}$ is also given by

$$\begin{aligned} & \left\| \widehat{f_{\text{sp},L}^s} - \widetilde{f_{\text{sp},L,\boldsymbol{\pi}}^s} \right\|_{L_2(X_{\frac{J}{n}})}^2 \\ & = O\left(2^{-\frac{(M-1)J}{n}} \cdot 2^{-\min(K_d^{\text{sp}}, 2K_s^{\text{sp}}+1)\frac{J}{n}}\right). \quad (49) \end{aligned}$$

Combining (47)–(49), we have the following result:

$$\begin{aligned} & \left\| \widetilde{f_{\boldsymbol{\pi}}^s} - \widetilde{f_{\text{sp},L,\boldsymbol{\pi}}^s} \right\|_{L_2(X_{\frac{J}{n}})}^2 \\ & \leq \left(\left\| f^s - \widetilde{f_{\boldsymbol{\pi}}^s} \right\|_{L_2(X_{\frac{J}{n}})} + \left\| f^s - \widehat{f_{\text{sp},L}^s} \right\|_{L_2(X_{\frac{J}{n}})} \right. \\ & \quad \left. + \left\| \widehat{f_{\text{sp},L}^s} - \widetilde{f_{\text{sp},L,\boldsymbol{\pi}}^s} \right\|_{L_2(X_{\frac{J}{n}})} \right)^2 \\ & = O\left(2^{-\frac{(M-1)J}{n}} \cdot 2^{-\min(K_d^{\text{sp}}, 2K_s^{\text{sp}}+1)\frac{J}{n}}\right). \quad (50) \end{aligned}$$

There are $O(2^{\frac{(M-1)J}{n}})$ Type D nodes at scale $\frac{J}{n}$. From (50), we have that the total error due to discrete surfprint approximations is given by $O(2^{-\min(K_d^{\text{sp}}, 2K_s^{\text{sp}}+1)\frac{J}{n}})$. Similar to the argument

made in Appendix I, we have by definition that the discrete surfprint approximation $\widetilde{f_{sp,\pi}^s}$, constructed using the discrete ℓ_2 -best surfprint atoms, satisfies

$$\|\widetilde{f_{\pi}^s} - \widetilde{f_{sp,\pi}^s}\|_{L_2(X_J)}^2 \leq \|\widetilde{f_{\pi}^s} - f_{sp,L,\pi}^s\|_{L_2(X_J)}^2.$$

Thus

$$\left\| \widetilde{f_{\pi}^s} - \widetilde{f_{sp,\pi}^s} \right\|_{L_2\left(X_{\frac{J}{n}}\right)}^2 = O\left(2^{-\frac{(M-1)J}{n}} \cdot 2^{-\min(K_d^{SP}, 2K_s^{SP}+1)\frac{J}{n}}\right).$$

Combining this result with the arguments made in the first paragraph of this proof, the squared- L_2 error between the discrete function $\widetilde{f_{\pi}^s}$ and the discrete wavelet/surfprint approximation $\widetilde{f_{\pi}^s}$ (this approximation now represents a composite wavelet/surfprint approximation; the previous paragraph only analyzed surfprint approximations) is

$$\left\| \widetilde{f_{\pi}^s} - \widehat{f_{\pi}^s} \right\|_{L_2}^2 = O\left(2^{-\min\left(\frac{2K_s^{w1}J}{m}, \frac{K_d^{SP}J}{n}, \frac{(2K_s^{SP}+1)J}{n}\right)}\right). \quad (51)$$

The rest of this proof is analogous to the Proof of Theorem 8 in Appendix H. \square

L. Proof of Corollary 12

This proof is similar to the proof in Appendix J. We begin by extending the bound provided by (45). This bound holds for Type D hypercubes at scale $\frac{J}{n}$. The total number of such Type D hypercubes at scale $\frac{J}{n}$ is given by $O(2^{(M-1)\frac{J}{n}})$. Following the logic preceding (48), we have that the squared- L_2 error between f^s and $\widetilde{f_{\pi}^s}$ over a Type S hypercube at scale j is given by $O(2^{-Mj-2\min(\pi_i)})$, and the total number of such Type S hypercubes at scale j is given by $O(2^{Mj})$. Combining these arguments with (51) from Appendix K, we have the following result:

$$\begin{aligned} \left\| f^s - \widehat{f_{\pi}^s} \right\|_{L_2}^2 &\leq \left(\left\| f^s - \widetilde{f_{\pi}^s} \right\|_{L_2} + \left\| \widetilde{f_{\pi}^s} - \widehat{f_{\pi}^s} \right\|_{L_2} \right)^2 \\ &= O\left(2^{-\min\left(\frac{2K_s^{w1}J}{m}, \frac{K_d^{SP}J}{n}, \frac{(2K_s^{SP}+1)J}{n}\right)}\right). \end{aligned}$$

The rest of the proof follows in an analogous manner to the proof of Theorem 8. \square

ACKNOWLEDGMENT

The authors wish to thank Justin Romberg, Ron DeVore, Hyeokho Choi, Mark Embree, and Rebecca Willett for enlightening discussions; Peter Jones for inspiring conversations on β -numbers; and Emmanuel Candès and David Donoho for helpful questions.

REFERENCES

[1] V. Chandrasekaran, M. B. Wakin, D. Baron, and R. Baraniuk, "Compression of higher dimensional functions containing smooth discontinuities," in *Proc. Conf. Information Sciences and Systems—CISS 2004*, Princeton, NJ, Mar. 2004.

[2] V. Chandrasekaran, M. B. Wakin, D. Baron, and R. Baraniuk, "Surflets: A sparse representation for multidimensional functions containing smooth discontinuities," in *Proc. IEEE Int. Symp. Information Theory (ISIT2004)*, Chicago, IL, Jun./Jul. 2004, p. 563.

[3] R. A. DeVore, B. Jawerth, and B. J. Lucier, "Image compression through wavelet transform coding," *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 719–746, Mar. 1992.

[4] J. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3445–3462, Dec. 1993.

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

[6] 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. Inf. Theory*, vol. 48, no. 7, pp. 1895–1921, Jul. 2002.

[7] D. L. Donoho, "Denosing by soft-thresholding," *IEEE Trans. Inf. Theory*, vol. 41, no. 3, pp. 613–627, May 1995.

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

[9] H. Krim, D. Tucker, S. Mallat, and D. Donoho, "On denoising and best signal representation," *IEEE Trans. Inf. Theory*, vol. 45, no. 7, pp. 2225–2238, Nov. 1999.

[10] R. Willett and R. Nowak, "Platelets: A multiscale approach for recovering edges and surfaces in photon-limited medical imaging," *IEEE Trans. Med. Imag.*, vol. 22, no. 3, pp. 332–350, Mar. 2003.

[11] E. Mammen and A. B. Tsybakov, "Smooth discrimination analysis," *Ann. Statist.*, vol. 27, pp. 1808–1829, 1999.

[12] A. B. Tsybakov, "Optimal aggregation of classifiers in statistical learning," *Ann. Statist.*, vol. 32, no. 1, pp. 135–166, 2004.

[13] I. Daubechies, "The wavelet transform, time-frequency localization and signal analysis," *IEEE Trans. Inf. Theory*, vol. 36, no. 5, pp. 961–1005, Sep. 1990.

[14] 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.

[15] J. Kovacevic and M. Vetterli, "Nonseparable multidimensional perfect reconstruction filter banks and wavelet bases for R^n ," *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pt. 2, pp. 533–555, Mar. 1992.

[16] E. J. Candès and D. L. Donoho, A. Cohen, C. Rabut, and L. L. Schumaker, Eds., "Curvelets—A surprisingly effective nonadaptive representation for objects with edges," in *Curve and Surface Fitting*. Nashville, TN: Vanderbilt Univ. Press, 1999.

[17] E. Candès, L. Demanet, D. L. Donoho, and L. Ying, "Fast discrete curvelet transforms," *Multiscale Model. Simul.*, vol. 5, pp. 861–899, 2005.

[18] M. Levoy and P. Hanrahan, "Light field rendering," in *SIGGRAPH'96: Proc. 23rd Annu. Conf. Computer Graphics and Interactive Techniques*, New Orleans, LA, 1996, pp. 31–42.

[19] J. VanderWerf, N. Kourjanskaia, S. Shladover, H. Krishnan, and M. Miller, "Modeling the effects of driver control assistance systems on traffic," in *Proc. National Research Council Transportation Research Board 80th Annu. Meeting*, Washington, DC, Jan. 2001.

[20] A. N. Kolmogorov and V. M. Tihomirov, " ϵ -entropy and ϵ -capacity of sets in functional spaces," *Amer. Math. Soc. Transl. (Ser. 2)*, vol. 17, pp. 277–364, 1961.

[21] G. F. Clements, "Entropies of several sets of real valued functions," *Pacific J. Math.*, vol. 13, pp. 1085–1095, 1963.

[22] S. Mallat, *A Wavelet Tour of Signal Processing*, 2nd ed. San Diego: Academic, 1999.

[23] M. B. Wakin, J. K. Romberg, H. Choi, and R. G. Baraniuk, "Wavelet-domain approximation and compression of piecewise smooth images," *IEEE Trans. Image Process.*, vol. 15, no. 5, pp. 1071–1087, May 2006.

[24] A. Lisowska, "Effective coding of images with the use of geometrical wavelets," in *Proc. Decision Support Systems Conf.* (in Polish), Zakopane, Poland, 2003.

[25] Y. Lu and M. N. Do, "Multidimensional directional filter banks and surfacelets," *IEEE Trans. Image Process.*, vol. 16, no. 4, pp. 918–931, Apr. 2007.

[26] D. L. Donoho, M. Vetterli, R. A. DeVore, and I. Daubechies, "Data compression and harmonic analysis," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2435–2476, Oct. 1998.

- [27] M. N. Do, P. L. Dragotti, R. Shukla, and M. Vetterli, "On the compression of two-dimensional piecewise smooth functions," in *IEEE Int. Conf. Image Processing—ICIP'01*, Thessaloniki, Greece, Oct. 2001, pp. 14–17.
- [28] V. Chandrasekaran, M. Wakin, D. Baron, and R. Baraniuk, "Compressing Piecewise Smooth Multidimensional Functions Using Surflets: Rate-Distortion Analysis Rice Univ. ECE Dep., Houston, TX, Tech. Rep., Mar. 2004.
- [29] J. Romberg, M. Wakin, and R. Baraniuk, "Multiscale Geometric Image Processing," in *Proc. SPIE Visual Communications and Image Processing Conf.*, Lugano, Switzerland, Jul. 2003, pp. 1265–1272.
- [30] M. B. Wakin, D. L. Donoho, H. Choi, and R. G. Baraniuk, "The multiscale structure of non-differentiable image manifolds," in *Proc. Wavelets XI at SPIE Optics and Photonics*, San Diego, Aug. 2005.
- [31] J. K. Romberg, M. B. Wakin, and R. G. Baraniuk, "Multiscale wedgelet image analysis: Fast decompositions and modeling," in *Proc. IEEE Int. Conf. Image Processing—ICIP'02*, Rochester, NY, 2002, pp. 585–588.
- [32] F. M. J. Willems, Y. M. Shtarkov, and T. J. Tjalkens, "The context-tree weighting method: Basic properties," *IEEE Trans. Inf. Theory*, vol. 41, no. 3, pp. 653–664, May 1995.
- [33] E. Kreyszig, *Introduction to Functional Analysis with Applications*. New York: Wiley, 1978.

Venkat Chandrasekaran (S'03) received the B.S. degree in electrical engineering, the B.A. degree in mathematics from Rice University, Houston, TX, in 2005, and the S.M. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 2007, where he is currently working towards the Ph.D. degree with the Stochastic Systems Group.

His research interests include statistical signal processing, machine learning, and computational harmonic analysis.

Michael B. Wakin (S'01–M'06) received the B.S. degree in electrical engineering and the B.A. degree in mathematics in 2000 (*summa cum laude*), the M.S. degree in electrical engineering in 2002, and the Ph.D. degree in electrical engineering in 2007, all from Rice University, Houston, TX.

He was an NSF Mathematical Sciences Postdoctoral Research Fellow at the California Institute of Technology, Pasadena, from 2006 to 2007 and an Assistant Professor at the University of Michigan in Ann Arbor from 2007 to 2008. He is now an Assistant Professor in the Division of Engineering at the Colorado School of Mines, Golden. His research interests include sparse, geometric, and manifold-based models for signal and image processing, approximation, compression, compressive sensing, and dimensionality reduction.

Dror Baron (S'99–M'03) received the B.Sc. (*summa cum laude*) and M.Sc. degrees from the Technion–Israel Institute of Technology, Haifa, Israel, in 1997 and 1999, and the Ph.D. degree from the University of Illinois at Urbana-Champaign in 2003, all in electrical engineering.

From 1997 to 1999, he worked at Witcom Ltd. in modem design. From 1999 to 2003, he was a research assistant at the University of Illinois at Urbana-Champaign, where he was also a Visiting Assistant Professor in 2003. From 2003 to

2006, he was a Postdoctoral Research Associate in the Department of Electrical and Computer Engineering at Rice University, Houston, TX. From 2007 to 2008, he was a quantitative financial analyst with Menta Capital. He is now a visiting scientist in the Department of electrical Engineering at Technion–Israel Institute of Technology, Haifa. His research interests include information theory and signal processing.

Dr. Baron was a recipient of the 2002 M. E. Van Valkenburg Graduate Research Award, and received honorable mention at the Robert Bohrer Memorial Student Workshop in April 2002, both at the University of Illinois. He also participated from 1994 to 1997 in the Program for Outstanding Students, comprising the top 0.5% of undergraduates at the Technion.

Richard G. Baraniuk (S'85–M'93–SM'98–F'02) received the B.Sc. degree in 1987 from the University of Manitoba, Winnipeg, Canada, the M.Sc. degree in 1988 from the University of Wisconsin-Madison, and the Ph.D. degree in 1992 from the University of Illinois at Urbana-Champaign, all in electrical engineering.

After spending 1992–1993 with the Signal Processing Laboratory of Ecole Normale Supérieure, in Lyon, France, he joined Rice University, Houston, TX, where he is currently the Victor E. Cameron Professor of Electrical and Computer Engineering. He spent sabbaticals at Ecole Nationale Supérieure de Télécommunications in Paris, France, in 2001 and Ecole Fédérale Polytechnique de Lausanne, Switzerland, in 2002. His research interests lie in the area of signal and image processing. In 1999, he launched Connexions, a nonprofit publishing project that invites authors, educators, and learners worldwide to "create, rip, mix, and burn" free textbooks, courses, and learning materials from a global open-access repository.

Dr. Baraniuk has been a Guest Editor of special issues for the *IEEE Signal Processing Magazine* on "Signal Processing and Networks" in 2002 and "Compressive Sampling" in 2008 and for the PROCEEDINGS OF THE IEEE on "Educational Technology" in 2008. He is currently an Associate Editor for the *ACM Transactions on Sensor Networks* and *Applied and Computational Harmonic Analysis*. He was Co-Technical Program Chair for the 2007 IEEE Statistical Signal Processing Workshop. He received a NATO postdoctoral fellowship from NSERC in 1992, the National Young Investigator award from the National Science Foundation in 1994, a Young Investigator Award from the Office of Naval Research in 1995, the Rosenbaum Fellowship from the Isaac Newton Institute of Cambridge University in 1998, the C. Holmes MacDonald National Outstanding Teaching Award from Eta Kappa Nu in 1999, the Charles Duncan Junior Faculty Achievement Award from Rice in 2000, the University of Illinois ECE Young Alumni Achievement Award in 2000, the George R. Brown Award for Superior Teaching at Rice in 2001, 2003, and 2006, the Hershel M. Rich Invention Award from Rice in 2007, the Wavelet Pioneer Award from SPIE in 2008, and the Internet Pioneer Award from the Berkman Center for Internet and Society at Harvard Law School in 2008. He was selected as one of *Edutopia Magazine's* Daring Dozen educators in 2007. Connexions received the Tech Museum Laureate Award from the Tech Museum of Innovation in 2006. His work with Kevin Kelly on the Rice single-pixel compressive camera was selected by *MIT Technology Review Magazine* as a TR10 Top 10 Emerging Technology in 2007. He was coauthor on a paper with Matthew Crouse and Robert Nowak that won the IEEE Signal Processing Society Junior Paper Award in 2001 and another with Vinay Ribeiro and Rolf Riedi that won the Passive and Active Measurement (PAM) Workshop Best Student Paper Award in 2003. He was elected a Fellow of the IEEE in 2001 and a Plus Member of AAA in 1986.