

Authenticity Analysis of Wavelet Approximations in Visualization

Pak Chung Wong
pcw@cs.unh.edu

R. Daniel Bergeron
rdb@cs.unh.edu

Department of Computer Science, University of New Hampshire
Durham, New Hampshire 03824, USA

Abstract

Wavelet transforms include data decompositions and reconstructions. This paper is concerned with the authenticity issues of the data decomposition, particularly for data visualization. A total of six datasets are used to clarify the approximation characteristics of compactly supported orthogonal wavelets. We present an error tracking mechanism, which uses the available wavelet resources to measure the quality of the wavelet approximations.

1 Introduction

Since the beginning of the 80's, there has been an explosion in research in the area of wavelet theory and its applications [3]. Major research achievements during the first decade included the Morlet wavelet, the Meyer wavelet, the Battle Lemarié wavelet, the Daubechies wavelet, and Mallat's multiresolution representation [10]. The last has become the *de facto* algorithm for most, if not all, wavelet applications in computer graphics and visualization for the last several years. Some of these efforts include volume rendering [12, 13], radiosity and textures [7, 5], volume morphing [8], progressive transmission [16], spacetime control [9], painting [1], curves [4], and shape models [17].

We are interested in using wavelets to provide a progressive refinement environment for scientific data visualization [19]. Such an environment should allow a scientist to visualize massive amounts of data at a coarse resolution to identify areas that warrant investigation at finer resolutions. To be effective the coarse resolutions must provide an accurate visual representation of the finer resolutions. In particular we would like to ensure that *important* patterns in the high resolution data are evident in the visualization of the low resolution data. Unfortunately, the definition of *what is important* is both application and task-dependent. On the other hand, we can develop a common evaluation and representation of the accuracy of a coarse resolution based on the loss of information between the coarse resolution and the fine resolution representations. We call this the *authenticity* of the representation.

The goal of this paper is to investigate the authenticity issues of the wavelet approximations. Common approaches of error measures for lossy reconstructions are discussed.

We define the meaning of *function pattern*, which we use to measure the quality of the approximations. Six test datasets, each with different characteristics, are used to clarify facts and pitfalls of orthogonal wavelets. An error tracking mechanism, which does not require any post-transformation computation, is presented with real life examples.

2 Surfing the little waves

This section sketches basic wavelet theory as it relates to our application. The reader is referred to [15, 3, 2] for more details. For novices, we recommend [18].

2.1 Scaling functions and wavelets

A function $\phi(x)$ is *refinable* if it can be represented by integer translates of its dilation, $\phi(2x)$. For example, in Figure 1, function $f(x)$ is *piecewise constant* on the unit




Figure 1: Function $f(x)$ is refinable and can be represented by its dilation $f(2x)$ and translate of its dilation $f(2x - 1)$.

interval; it can be represented by its dilation $f(2x)$, together with the translate of its dilation $f(2x - 1)$, which are piecewise constant on half unit intervals. Mathematically, $\phi(x)$ can be represented by integer translates of its own dilations such that

$$\phi(x) = \sum_k c_k \phi(2x - k) \quad (1)$$

where c_k is constant and $k \in \mathbb{Z}$. In wavelet literature, ϕ is known as the *scaling* function. The box function f in Figure 1 is indeed the scaling function of the Haar wavelet.

Given $\phi(x)$, a linear space $V_j \subset L^2(\mathbb{R})$ can be obtained by *spanning* the integer translates of its dilations. ($L^2(\mathbb{R})$ denotes the Hilbert space of measurable, square-integrable one-dimensional functions.) The linear space is formally defined as

$$V_j = \text{span}(\phi(2^j x - k)_{k \in \mathbb{Z}}). \quad (2)$$

Since ϕ is refinable, $V_j \subset V_{j+1}$ for $j \in \mathbb{Z}^+$. These two infinite spaces are related by an orthogonal projection. If,

for example, V_j is spanned by a finite basis with vectors x_1 and x_2 , a projection of function $f \in V_{j+1}$ onto space V_j can be depicted as shown in Figure 2. In an infinite dimension




Figure 2: A projection of function $f(x) \in V_{j+1}$ onto V_j .

space, the shaded region represents V_j of Equation (2). If we define W_j to be the *orthogonal complement* of V_j with respect to V_{j+1} (i.e., $V_j \perp W_j$), then

$$V_j \oplus W_j = V_{j+1}. \quad (3)$$

The space W_j is spanned by the integer translates of the dilations of a function ψ known as the *wavelet*. A function $f \in V_{j+1}$ can be scaled to a coarse resolution of V_j by losing details, which are captured in W_j . If Equation (3) is applied iteratively, we get

$$\begin{aligned} V_j &= W_{j-1} \oplus V_{j-1} = \dots \\ &= W_{j-1} \oplus W_{j-2} \oplus \dots \oplus W_0 \oplus V_0. \end{aligned} \quad (4)$$

Therefore we can represent function $f \in V_j$ by a set of integer translates of the dilated wavelets, also known as a *wavelet basis*. This iteration process is the foundation of Mallat’s multiresolution wavelet analysis [10].

2.2 Multiresolution analysis

In practice, we can only measure or record scientific data in discrete values. A wavelet implementation on discrete data is given by Mallat [10]. Given a one-dimensional dataset with N items at resolution j , an application of an orthogonal wavelet decomposition generates $\frac{N}{2}$ coefficients of low frequency *approximations*, and $\frac{N}{2}$ coefficients of high frequency *details* at resolution $j - 1$, as shown in Figure 3. The coefficients correspond to the bases in V_{j-1}




Figure 3: Wavelet decomposition on a one-dimensional data.

and W_{j-1} discussed in Section 2.1. From Equation (4), we know that this operation can be applied iteratively on the approximations to get increasingly coarse data, for as long as the size of the approximations $\geq 2p$, where p is the number of *vanishing moments* of the wavelet. (We discuss vanishing moment in more detail below.)

Wavelet transforms are invertible. If both the approximations and the details of any one resolution are available,

it is possible to have a *lossless* reconstruction of the approximations of the next finer resolution.

3 General error measure

Applications of wavelets are developing rapidly in both graphics and visualization research. Many of them rely on the vanishing moment property of the wavelet transform, in which a substantial amount of wavelet detail is ignored because of its relatively insignificant values. The process is usually followed by a reverse transform operation, which produces a *lossy* reconstruction. This is by far the most popular application of wavelets.

A common way to measure the effectiveness of a lossy operation is by *visual comparison* based on feature classification by a human being on a small number of features such as the peak value and fundamental frequency. This *subjective* method assumes medium-sized data which can be displayed in a single picture.

Traditionally, applications which require higher error accuracy might use more *objective* (quantitative) measures such as the root-mean-square error, given by

$$\sqrt{\frac{1}{N} \sum_{i=1}^N (f_i - f'_i)^2} \quad (5)$$

where f_i and f'_i are the original and the reconstructed data respectively. Unfortunately, however, this method requires both functions to be the same size, which is not the case when comparing a function with its coarse-resolution reconstructions. Furthermore, it is not always the case that objective measures are better than subjective ones.

In our application [19] we combine objective and subjective error measures in an interactive environment to explore data. Orthogonal wavelets provide fine to coarse resolutions to browse terabyte sized scientific data and search for information. The wavelet approximations give an overview of the very large data set at a lower resolution, and the details of each resolution allow the reconstruction of the data in higher resolutions once interesting targets are located. Obviously, the zoomed data has a much smaller size than the original, so conventional error measures such as the root-mean-square method cannot be used. Instead of computing an error measure based on comparing the approximation coefficients to the input data, we compute one based on the detail coefficients.

4 Authenticity analysis

We would like to have some mechanism for validating that a lower resolution representation of a signal is an authentic approximation. Fortunately, by using a wavelet representation, the energy loss due to space projection can be obtained from the wavelet details of each resolution. Once we define a measure of the energy loss, we can use

that measure for both analysis and visualization of the error. Defining an error measure that is consistent throughout the wavelet coefficient hierarchy is complicated by the change in coefficient scale and the different numbers of coefficients at each level of the hierarchy. For our initial experimentation, we have chosen an error metric in which the total error due to projection from spaces V_{j+1} to V_j is given by

$$totalError_j = \sum_k |w_{jk}| \quad (6)$$

where w_{jk} is the wavelet detail at resolution j . Even this simple measure shows positive visual results as shown in Section 5, but more research is needed to determine how it compares to other potential measures.

It is known that some wavelets lose more energy than others during decomposition. This brings us to the discussion of vanishing moments.

4.1 Vanishing moments

Our discussion is restricted to compactly supported orthogonal wavelets. Others such as the Morlet wavelet and the Meyer wavelet have infinite support on the whole real line because they use sinusoids as the building blocks. They are C^∞ .

The accuracy of piecewise wavelet approximation can be characterized by the number of vanishing moments p of the wavelet. Strang [15] describes this as how well the polynomials $1, x, x^2, \dots, x^{p-1}$ are reproduced by the approximation. For a wavelet ψ with p vanishing moments,

$$\int \psi(x)x^m dx = 0 \quad (7)$$

where $m = 0, \dots, p-1$. In wavelet literature, the value of $2p$ is usually used as a subscript to identify a wavelet. This comes from the fact that the number of coefficients of a wavelet filter with p vanishing moments is equal to $2p$. Figure 4 shows H_2 (Haar wavelet with $p = 1$), D_4 , D_{12} , and




Figure 4: Top: H_2 , D_4 . Bottom: D_{12} , D_{20} .

D_{20} (Daubechies wavelets with $p = 2, 6,$ and 10 .) It is also true that a compactly supported wavelet with p vanishing moments is p times continuously differentiable, i.e., C^p .

4.2 Function pattern

In this paper, the *pattern* of a function is loosely defined by its *shape* or *look* rather than its mathematical definition. Since most graphic displays are normalized in our application [19], the number of a function's *extreme values* (i.e., local maximum and minimum) and their *relative locations* are far more important than the absolute values. Following the same philosophy, we pay more attention to the *frequency* or the *period* of a periodic function than its *amplitude*.

4.3 Approximation characteristics

In this section, compactly supported orthogonal wavelets with $p = 1$ to 10 are applied to sinusoid-based test data with different patterns. Due to limited space, we only plot the results generated by wavelets H_2 and D_{20} , which represent the lowest and the highest number of vanishing moments in our discussion. These examples are used to clarify the approximation characteristics of orthogonal wavelets.

4.3.1 Effect of vanishing moment on merging of data

A higher number of vanishing moments implies more coefficients in a wavelet filter matrix, which means that the implementation requires that more numbers be multiplied and added (merged) together during decomposition. However, the features in the data are not necessarily merged any faster. This is illustrated by a sinusoid with $2^9 = 512$ discrete samples shown in Figure 5. Figure 6 depicts the




Figure 5: Dataset #1 with 512 items.




Figure 6: Wavelet coefficients of dataset #1 with 512 items.

first resolution of the decomposition. The left half is the approximation and the right is the detail. The tiny wavelet detail indicates good approximation for both wavelets.

Figure 7 shows the same signal represented with $2^8 = 256$ wavelet coefficients. Both wavelets show strong approximations and weak details, with D_{20} having smaller details.

After two more resolutions, the size of the wavelet coefficients is down to $2^6 = 64$, as shown in Figure 8. H_2




Figure 7: Wavelet coefficients of dataset #1 with 256 items.




Figure 8: Wavelet coefficients of dataset #1 with 64 items.

continues to lose more energy than D_{20} . Both wavelets, however, manage to maintain the basic pattern (four peaks) of the original function.

Although more numbers are multiplied with higher p values during decomposition, the merging process is dominated by the offset of the data during each multiplication, which shifts two positions at a time for orthogonal wavelets.

4.3.2 Energy loss versus resolution

Each wavelet decomposition introduces energy loss (error) into the approximation. By examining Figures 6–8, we see that H_2 has significantly larger energy loss at each resolution. However, the amount of energy loss depends not only on the wavelet itself, but also on the *smoothness* of the data, which may change during each resolution. This is discussed in the following section.

4.3.3 Energy loss versus data smoothness

The rate of decay of wavelet decomposition is governed by the number of vanishing moments of the wavelet. In general, as illustrated in our previous example, smooth functions like D_{20} approximate functions better than H_2 . However, this is not always true.

We use a second sinusoid with a higher frequency to demonstrate. The discrete function size is $2^9 = 512$. There are a total of 21 peaks as depicted in Figure 9.




Figure 9: Dataset #2 with 512 items.

Figure 10 shows the wavelet coefficients after two resolutions with 256 discrete items. Because of the smoothness of the data, the very small energy loss of D_{20} hardly shows




Figure 10: Wavelet coefficients of dataset #2 with 256 items.

up in the figure. Both H_2 and D_{20} retain all 21 peaks even though H_2 shows larger details.

The next resolution is depicted in Figure 11. It has 128




Figure 11: Wavelet coefficients of dataset #2 with 128 items.

discrete wavelet coefficients: 64 approximations and 64 details. All 21 spikes stay in the approximation, which is rather non-smooth by now.

Disaster hits when the size of the function reaches $2^6 = 64$. The *Nyquist* limit (the lower bound to retain all the spikes) is past. As a result, features merge and large details are created. For the first time in our examples, D_{20} has larger details than H_2 as shown in Figure 12.




Figure 12: Wavelet coefficients of dataset #2 with 64 items.

H_2 does not approximate functions accurately because it only has one vanishing moment. However, it does not tend rapidly to zero at finer levels, and the accuracy also depends on the smoothness of the function. When the function is smooth, more vanishing moments lead to smaller wavelet details. On the other hand, more vanishing moments also lead to more large wavelet details when the function is non-smooth, so D_{20} is not always a better choice over H_2 .

The phenomenon shown in Figures 10–12 is repeated for the rest of the resolutions. The first three resolutions of this example are marked by their inactivity or repose. Most of the major features of the approximations stay the same, i.e., the details are very small. That creates a state of *quiescence*. It also, however, establishes a *horizon* beyond which disaster lurks. In Figure 12, half of the spikes in the approximation are gone, and the details (the energy loss) are actually much larger than the approximations.

Definition In a progressive refinement environment, a *quiescence* is defined as a state of inactivity in which most of the distinctiveness of the refining target stays.

One of the major issues of our application [19] is the detection of quiescence. Only the lowest resolution representation needs to be maintained from a set of resolutions that are part of a single quiescent state.

4.3.4 Energy loss versus data value

Energy loss of a wavelet decomposition depends on the data values. Higher data values imply more energy loss, and vice versa. This is illustrated by a sinusoid whose amplitude decreases continuously as depicted in Figure 13.




Figure 13: Dataset #3 with 512 items.

After two resolutions, the wavelet coefficients (with 256 discrete values) of dataset #3 is shown in Figure 14. The




Figure 14: Wavelet coefficients of #3 with 256 items. values of the details increase as the data values increase.

4.3.5 Energy loss versus data frequency

Energy loss of a wavelet decomposition depends on the data frequency. Higher data frequency implies more energy loss, and vice versa. This occurs because the data gets less smooth as frequency increases. A *chirp*, whose frequency increases continuously, illustrates our statement as shown in Figure 15.




Figure 15: Dataset #4 with 1,024 items.

Figure 16 shows the wavelet coefficients after one resolution of wavelet decomposition. Both wavelets indicate




Figure 16: Wavelet coefficients of dataset #4 with 1024 items.

that the higher the frequency, the larger the energy loss. At the same time, they also keep the patterns intact.

4.3.6 Pitfall

From Equation (4), we would expect that the wavelet decomposition (with floating point computations) will continue until the size of the approximation reaches $2p$. However, certain patterns, even with high energy content, can vanish suddenly and prematurely.

The sinusoid in Figure 17 is created with an integer




Figure 17: Dataset #5 with 512 items.

value of cycles within the 512 discrete values. After three resolutions of decomposition, both H_2 and D_{20} shown in Figure 18 have perfect zigzag patterns. In fact, the pattern




Figure 18: Wavelet coefficients of dataset #5 with 64 items.

interleaves high-low discrete values with the same amount of energy but reversed directions. Then the energy vanishes suddenly as shown in Figure 19 when a high energy spike




Figure 19: Wavelet coefficients of dataset #5 with 32 items.

zeros out its neighbor, a low energy data item with exactly the same absolute value.

4.4 Error tracking and results

We do not claim that our sinusoid-based examples above simulate real life scientific data. They are presented because they show the special characteristics of orthogonal wavelets. Now we present results of non-sinusoid functions. Since wavelet transforms are well adapted to respond locally to rapid changes in function values [3], orthogonal wavelets are often used as edge detectors [10, 11]. Energy loss during wavelet decomposition usually implies edge (i.e., pattern) changes of the approximations. Small changes produce little wavelet details, which can largely be ignored. Large changes such as aliasing, however, produce significant wavelet details. These details are good indicators of the quality of the wavelet approximation of each resolution.

A function with 512 discrete values is created to demonstrate the idea of using wavelet details to measure the authenticity of wavelet approximations, as described at the beginning of section 4. The function has seven special features including: 1) two discrete steps, 2) a portion of a sinusoid, 3) two steep slopes ($f(x) = x^3$ and its mirror image), 4) a sharp spike, 5) some fluctuating signals, 6) a large flat block, and 7) a small block. Each of them is identified with its feature number as depicted in Figure 20.




Figure 20: Dataset #6 with 512 items.

When D_4 is applied to this function, it produces 256 discrete approximations followed by 256 details, as shown in Figure 21. In the graph, the wavelet details corresponding




Figure 21: Wavelet coefficients of dataset #6 with 512 items.

to the original features are marked with the same numbers. Starting from the left of the details, we see two sharp spikes which indicate the two sudden slope changes of feature 1. The energy loss is due to the fact that these two sharp edges are smoothed out. The details of feature 2 are almost invisible, reflecting the high approximation power of D_4 on sinusoids. The two steep slopes of feature 3 create only small spikes, which also indicate good approximations. The isolated spike of feature 4 is preserved with some loss. The fluctuating signals of feature 5 produce the largest amount of

loss, indicating major pattern changes. Each of the following two blocks (features 6 and 7) produce two consecutive spikes, indicating the two sharp edges of the blocks. The next two resolutions are shown in Figures 22 and 23.




Figure 22: Wavelet coefficients of dataset #6 with 256 items.




Figure 23: Wavelet coefficients of dataset #6 with 128 items.

The *maxima* and the *totals* of the absolute values of wavelet details of the first four resolutions are listed in Figure 24. The *summation* of wavelet details shows the

N	$\max w_i $	$\sum w_i $
256	0.9124	6.2810
128	1.5074	10.7591
64	1.1257	9.1903
32	1.2708	8.8527

Figure 24: A summary of D_4 on dataset #6.

overall approximation quality, while the *maximum* indicates the worst error. We notice from the graphs that D_4 keeps the function pattern intact. This is largely reflected by the correlated numbers in the table.

These numbers, however, fail to show the energy loss of a particular feature such as the fluctuating signals, which produce the largest details. Figure 25 is a summary of

N	$\sum_{\text{feature}} w_i $						
	1	2	3	4	5	6	7
256	0.708	0.014	0.477	0.177	3.763	0.966	0.177
128	1.557	0.081	0.754	0.225	6.177	1.641	0.325
64	~2.560		1.432	0.099	3.790	0.813	0.498
32	~7.177					~1.676	

Figure 25: A summary (by feature) of D_4 on dataset #6.

wavelet details based on the different features. In general, the fluctuating signals (feature 5) have the biggest pattern changes, which explain why they have the largest details. Feature 1 has one of the larger losses. The two steps are totally smoothed out after the second resolution.

Figure 26 lists the results of D_{20} applied to dataset #2

N	$max w_i $	$\sum w_i $
256	0.3718	0.9792
128	0.4690	1.3124
64	1.2059	13.2433
32	3.8953	72.6761
16	2.1756	8.5699
8	0.7576	1.6806

Figure 26: A summary of D_{20} on dataset #4.

as presented in Section 4.3. The resolution printed in bold shows the occurrence of aliasing as well as the end of the first quiescence covering the previous three resolutions. A second quiescence starts right after this resolution.

5 Visualization of error

We briefly describe a wavelet-based multiresolution visualization system [19], and use it to display approximations of scientific data and its corresponding error. The system, as depicted in Figure 27, supports progressive refinement




Figure 27: A wavelet-based visualization tool.

data analysis with resolution as fine as one data item per pixel. Ten orthogonal wavelets with vanishing moments from $p = 1$ to 10 are provided. The system keeps track of the accumulated data loss as well as data loss from individual resolutions during wavelet transforms. This error can be displayed alongside the compressed data. Input data can be in *CDF* [14] or system-defined binary formats. A variety of colormaps are available.

Figures 27 and 28 show a 1D data set extracted from the CD-ROM `USA_NASA_DDF_ISTP_KP_0003` recorded from the spacecraft GEOTAIL of the ISTP [6] project. The data set contains electron average energy data recorded every 64 seconds around the earth for the first three months of 1994. It has a total of $2^{17} = 131,072$ integers.

D_4 is used to generate a total of 7 resolutions, from the 0^{th} resolution with $2^{10} = 1,024$ to the 6^{th} resolution with $2^{16} = 65,536$ items. The display is a one-dimensional line plot with colors indicating the accumulated data loss (error)

of each item. A rainbow colormap shown in Figure 28a is used for the display.

We start from the coarsest (0^{th}) resolution with 1,024 items, an approximation of the original 131,072 items. From Section 4.3, we know that the energy loss of a wavelet transform depends on a number of factors including the data values and the smoothness of the data. Six features are chosen from the data, as indicated in Figure 28b, to illustrate these ideas as well as the importance of the approximation error display.

Feature 1 contains highly fluctuating data with some of the highest data values. Both of them contribute to the very high energy loss of the approximation. This is reflected by the darker colors (green/blue) of the error display in Figures 28b and 28c.

In terms of data values, features 4, 5 and 6 are more or less close to each other. However, when we look at the smoothness of the data, we notice that feature 6 is smoother than the other two. It implies that feature 6 has the lowest energy loss. This is accurately reflected by the color of the error displayed in Figures 28b. Both features 4 and 5 have green spikes while feature 6 is light orange.

The value of the error representation is particularly evident when looking at features 2 and 3. They have very similar values, spreads, and shapes. Feature 3, however, has more green (higher value) spikes than feature 2. These two features (marked by the dotted rectangles) are zoomed to a finer (4^{th}) resolution, as shown in Figures 28d and 28e. This finer resolution reveals that feature 2 is indeed very smooth data, while feature 3 is relatively non-smooth. Figure 28e shows that feature 3 has multiple spikes spread across the area. These narrow spikes fade away during the downsampling process because the *Nyquist* limit (to hold all the spikes) is reached. Their errors are clearly reflected in our display. By using the color to represent accumulated error, we are able to identify areas of the coarsest resolution representation that warrant investigation at finer resolutions.

6 Conclusions and future work

We illustrate the approximation characteristics of compactly supported orthogonal wavelets. The idea of using wavelet details to measure the authenticity of wavelet approximation is presented and results from test data and real scientific data are discussed.

Our immediate goal is to extend this work to functions of higher dimensions. This paper is part of our on-going efforts on very large data visualization using wavelets [19].

Acknowledgements

This work has been supported in part by the National Science Foundation under grant IRI-9117153. The ISTP data is courtesy of Dr. Terry Onsager of the Institute for the Study of Earth, Oceans, and Space at the University of New Hampshire.

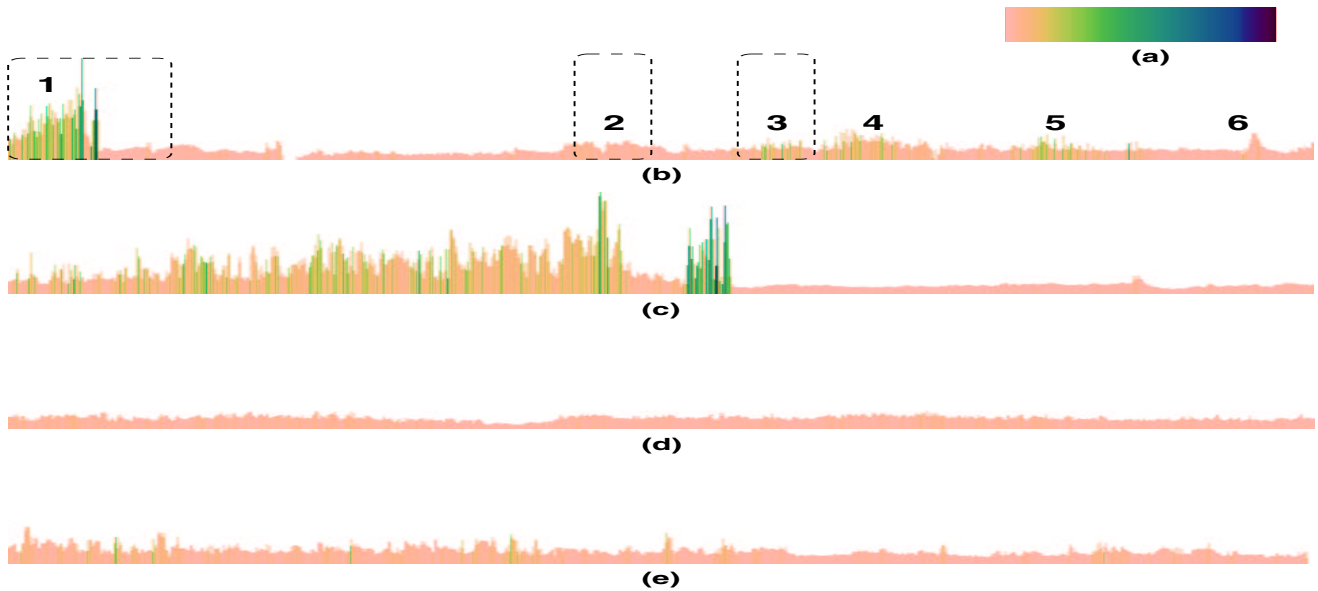


Figure 28: a) Rainbow colormap. b) The coarsest approximation of the ISTP average energy data is displayed at the 0^{th} resolution. Interesting features are identified with numbers. Dotted rectangles are zooming windows. The color indicates the accumulated approximation error. c) Feature 1 at the 3^{rd} resolution. d) Feature 2 at the 4^{th} resolution. e) Feature 3 at the 4^{th} resolution.

References

- [1] Deborah F. Berman, Jason T. Bartell, and David H. Salesin. Multiresolution painting and compositing. In *SIGGRAPH 94 Conference Proceedings*, pages 85–90, August 1994.
- [2] Charles K. Chui. *An Introduction to Wavelets*. Wavelet Analysis and its Applications – Volume 1. Academic Press, 1992.
- [3] Ingrid Daubechies. *Ten Lectures on Wavelets*. SIAM, Philadelphia, Pennsylvania, 1992.
- [4] Adam Finkelstein and David H. Salesin. Multiresolution curves. In *SIGGRAPH 94 Conference Proceedings*, pages 261–268, August 1994.
- [5] Reid Gershbein, Peter Schröder, and Pat Hanrahan. Textures and radiosity: Controlling emission and reflection with texture maps. In *SIGGRAPH 94 Conference Proceedings*, pages 51–58, August 1994.
- [6] Goddard Space Flight Center, National Aeronautics and Space Administration, Greenbelt, Maryland. *International Solar–Terrestrial Physics (ISTP) Key Parameter Generation Software (KPGS) Standards and Conventions – Version 1.1*, December 1992.
- [7] Steven J. Gortler, Peter Schröder, Michael F. Cohen, and Pat Hanrahan. Wavelet radiosity. In *SIGGRAPH 93 Conference Proceedings*, pages 221–230, August 1993.
- [8] Taosong He, Sidney Wang, and Arie Kaufman. Wavelet-based volume morphing. In *Proceedings of IEEE Visualization '94*, pages 85–92, Washington, DC, October 1994.
- [9] Zicheng Liu, Steven J. Gortler, and Michael F. Cohen. Hierarchical spacetime control. In *SIGGRAPH 94 Conference Proceedings*, pages 35–42, August 1994.
- [10] Stephane G. Mallat. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7):674–693, 1989.
- [11] Stephane G. Mallat and Sifen Zhong. Characterization of signals from multiscale edges. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(7):710–732, 1992.
- [12] Shigeru Muraki. Application and rendering of volume data using wavelet transforms. In *Proceedings IEEE Visualization '92*, pages 21–28, Boston, MA, October 1992.
- [13] Shigeru Muraki. Volume data and wavelet transforms. *IEEE Computer Graphics and Applications*, 13(4):50–56, 1993.
- [14] National Space Science Data Center, Greenbelt, Maryland. *CDF User's Guide, Version 2.4*, February 1994.
- [15] Gilbert Strang. Wavelets and dilation equations: A brief introduction. *SIAM Review*, 31(4):614–627, December 1989.
- [16] Hai Tao and Robert J. Moorhead. Progressive transmission of scientific data using biorthogonal wavelet transform. In *Proceedings of IEEE Visualization '94*, pages 93–99, Washington, DC, October 1994.
- [17] B. C Vemuri and A. Radisavljevic. Multiresolution stochastic hybrid shape models with fractal priors. *IEEE Transactions on Graphics*, 13(2):177–207, April 1994.
- [18] Pak Chung Wong and R. Daniel Bergeron. A child's garden of wavelet transforms. Technical report, Department of Computer Science, University of New Hampshire, New Hampshire, 1994.
- [19] Pak Chung Wong and R. Daniel Bergeron. Hierarchical representation of very large data sets for visualization using wavelets. In *Proceedings of Dagstuhl Workshop on Scientific Visualization*. IEEE CS Press, 1995. To be published.