

Wavelets, Adapted Waveforms, and De-Noising

Ronald R. Coifman, Ph.D.* Mladen Victor Wickerhauser, Ph.D.†

Abstract

This is a short summary of a talk presented by Wickerhauser at the Frontier Science in EEG Symposium, *Continuous Waveform Analysis*, held 9 October 1993 in New Orleans. We describe some new libraries of waveforms well-adapted to various numerical analysis and signal processing tasks. The main point is that by expanding a signal in a library of waveforms which are well-localized in both time and frequency, one can achieve both understanding of structure and efficiency in computation. We briefly cover the properties of the new “wavelet packet” and “localized trigonometric” libraries. The main focus will be applications of such libraries to the analysis of complicated transient signals: a feature extraction and data compression algorithm for speech signals which uses best-adapted time and frequency decompositions, and an adapted waveform analysis algorithm for removing fish noises from hydrophone recordings. These signals share many of the same properties as EEG traces, but with distinct features that are easier to characterize and detect.

1 Time and frequency analysis

Our goal is to describe tools for adapting methods of analysis to various tasks occurring in harmonic and numerical analysis and signal processing. The main point of this presentation is that by choosing an orthonormal basis, in which space and frequency are suitably localized, one can achieve both understanding of structure and efficiency in computation. In fact we claim that the search for computational efficiency is intimately

*Department of Mathematics, Yale University

†Department of Mathematics, Washington University in St.Louis

related to efficiency in representation (i.e., compression) and to pattern extraction, or structural understanding.

This is traditionally done by choosing appropriate decomposition of the signal into components in the *phase plane*—an abstract 2-dimensional signal representation in which time and frequency are the horizontal and vertical axes, respectively. A waveform is represented by a rectangle in this plane, as seen in Figure 1. Let us call such a rectangle an *information cell*. The position in time and the main frequency can be read from the coordinates of the center of the rectangle. The uncertainty in time and the uncertainty in position are given by the width and height of the rectangle, respectively. Heisenberg’s inequality, or the *uncertainty principle*, implies that the area of such a rectangle can never be less than 1. The amplitude of a waveform can be encoded by darkening the rectangle in proportion to its waveform’s energy.

We will only use phase plane atoms in our analysis; these are waveforms which are so well localized in both time and frequency that the areas of their information cells must be close to 1. They may be represented by information cells of equal area in the phase plane. A basis of such atoms corresponds to a cover of the phase plane by rectangles, and we will depict an orthonormal basis by using disjoint rectangles of exactly equal (unit) area. Of course, only the *Gaussian* function $g(t) = e^{-t^2}$, dilated or modulated, has the minimal information cell area. The other phase plane atoms are not too far off, though, and we will avoid the many restrictions of the Gaussian by relaxing the minimality condition. The only price we will have to pay is that a single atom might in practice require a few of the approximate atoms. Figure 2 shows the output of such an approximate analysis on two Gaussian phase plane atoms.

An orthogonal adapted waveform analysis [Wickerhauser, 1994] approach provides a numerical recipe for the decomposition; the covering partitions are chosen to achieve maximum efficiency with respect to an *information cost function*. Not only does this approach shed light on classical analysis methods, it also suggests new methods of




Figure 1: Idealized information cells in the phase plane.




Figure 2: An actual analysis in the phase plane.

organization and analysis of operators on the signals we have analyzed. These include nonstandard forms for nonlinear operators which better display the interactions among different signal parts, and discrete phase plane approximations for the evaluation of complicated operators. A calculus in compressed variables exists [Wickerhauser, 1992a, Beylkin et al., 1991] exists, making adapted transform methods useful for fast numerical algorithms, for data manipulation and for large scale computation.

2 Example libraries of waveforms

Our goal is to introduce a variety of techniques permitting the mathematician, scientist or engineer to choose the appropriate analysis method in this catalogue of tools and apply it to practical problems. In an orthogonal adapted waveform analysis, the user is provided with a collection of standard libraries of waveforms—called *wavelets*, *wavelet packets*, and *windowed trigonometric waveforms*—which can be combined to fit specific classes of signals. All these functions are phase plane atoms. Examples of such waveforms are displayed in Figure 3.

These libraries are used because they come equipped with fast numerical algorithms, enabling real-time implementation of a variety of analysis and signal processing tasks. These include data compression, parameter extraction for recognition and diagnostics, and fast transformation and manipulation of digital data. The process of analysis is usually done by comparing acquired segments of data with stored waveforms. The numerical comparison algorithm itself is fast and perfectly conditioned, always being a factored sparse orthogonal transformation. The most efficient orthonormal basis for compression of the signal is selected and used to extract and manipulate relevant features.

Consider first the short-time or windowed Fourier transform. Here the basis functions are exponentials, or maybe sines and cosines, which are enveloped so they oscillate only for a short time before going back to zero. They yield a tiling of the phase plane by congruent




Figure 3: Example waveforms: wavelet, wavelet packet, block cosine and smoothly windowed cosine functions.

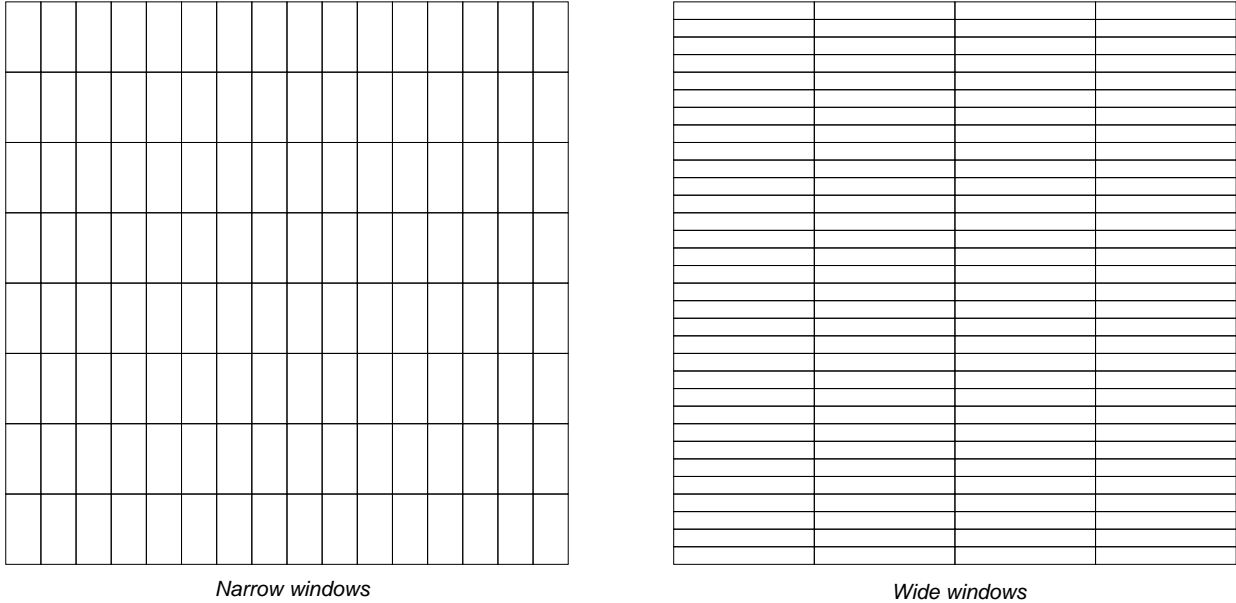


Figure 4: Two windowed Fourier tilings of the phase plane.

rectangles, whose dimensions depend upon the window size. Two choices of window size are shown in Figure 4. It is still necessary to choose a window size appropriate to the analysis. Our measure of quality will be the amount of white space, or negligible waveforms, in the time frequency analysis of a signal. Lots of white space means that most of the components have negligible energy, so that that the signal energy is concentrated into just a few waveforms.

Very short windows are most efficient for sharp impulses, while long windows correspond to information cells which spread energy all over the phase plane, as seen in Figure 5.

Conversely, long windows are more efficient than short ones for nearly continuous tones, as depicted in Figure 6. Hence it can be useful to examine the signal in many window sizes at once and then to choose the *best basis* in the sense of efficient representation.

Remark. The Fourier transform is a rotation by 90 degrees in the phase plane. This is

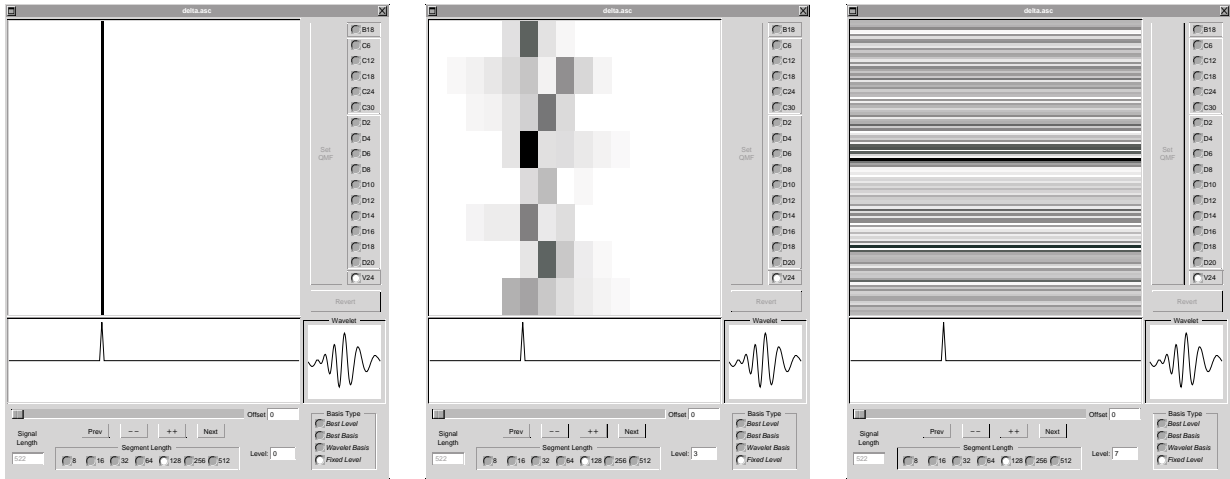


Figure 5: Phase plane analysis of a sharp impulse at increasing window sizes.

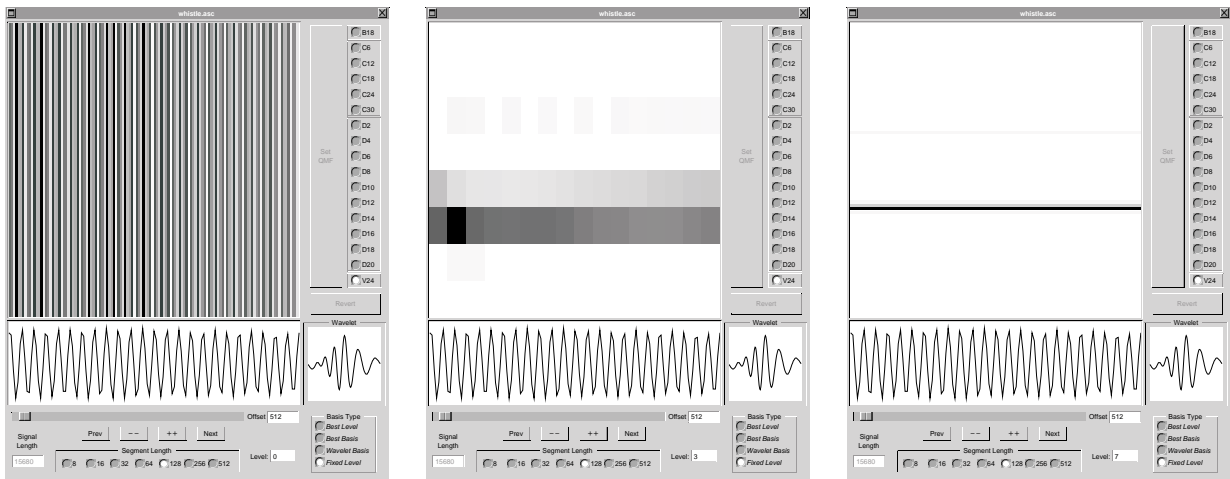


Figure 6: Phase plane analysis of a nearly pure tone at increasing window sizes.




Figure 7: Grid-point and Fourier decompositions of the phase plane.

most evident in Figure 7, which shows the idealized information cells corresponding to grid-point samples (i.e., *Dirac mass* “waveforms”) or sampled pure sine waves as phase plane atoms. It is also possible to apply the *Hermite semigroup transform* or *angular Fourier transform* to obtain information cells which make arbitrary angles with the time and frequency axes.

Wavelet analysis [Daubechies, 1988, Daubechies, 1992] corresponds to windowing frequency space in “octave” windows. Since the information cells have equal area, they cover the phase plane in the manner depicted in Figure 8. A natural extension therefore is provided by allowing all dyadic windows in frequency space and adapted window choice. This sort of analysis is equivalent to wavelet packet analysis.

The *wavelet packet* library is constructed by recursion of the wavelet algorithm. This




Figure 8: Dyadic wavelet tiling of the phase plane.

library contains the wavelet basis, Walsh functions, and smooth versions of Walsh functions called wavelet packets. The wavelet packet analysis algorithms [Coifman et al., 1990, Coifman et al., 1992b, Coifman and Wickerhauser, 1992b] permit us to perform an adapted Fourier windowing directly in the time domain by successive filtering of a function into different frequency bands. The window size selection algorithm, in this context, gives an adapted subband coding algorithm. It should be mentioned that all of these algorithms have higher dimensional generalizations [Wickerhauser, 1992b, Wickerhauser, 1991]; in particular, they can be used to analyze still images and movies in the same way that we are analyzing acoustic signals.

To illustrate a complete analysis in a library, we start with a description of an algorithm to compute the windowed Fourier sine expansion of a function on an interval from the Fourier expansion of its restrictions to the left half and the right half [Auscher et al., 1992, Coifman and Meyer, 1991, Malvar, 1990]. This procedure is depicted in Figure 9. The main idea is that all the signals under the big (reflected) window or “bell”

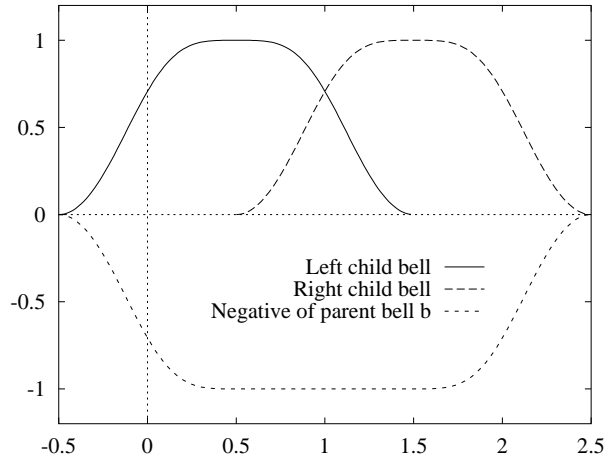


Figure 9: Big window from two small windows.

are combinations of signals under the two smaller windows. We see that in order to compute the Fourier expansion on the large interval, we can start with adjacent pairs of small intervals, combine coefficients to obtain the expansion on their union, and continue until we reach the largest interval at the top level. This scheme is depicted in Figure 10. Along the way we have obtained all dyadic windowed Fourier transforms as intermediate computations. We notice that every disjoint collection of intervals and their orthogonal bases provides us with an orthogonal basis for the union.

3 Choosing the “best basis”

A natural question that arises in connection with the windowed Fourier transform is how to place the windows. The window selection has a big effect on the number of large coefficients in the expansion. So let us now turn to the question of optimizing the windows to obtain an efficient representation of a function.

The *best basis* algorithm fits a phase plane cover to the signal so as to best concentrate the




Figure 10: Computing all windowed Fourier sine transforms at once.

shading into the fewest information cells in the phase plane. This method can use rectangular information cells of all aspect ratios. The *best level* algorithm fits a cover of equal aspect ratio rectangles to the signal, so as to best concentrate the shading. The restriction on the aspect ratios is sometimes useful to avoid undesirable artifacts in partial reconstructions.

We can proceed as follows in our adapted windowed Fourier transform example: we start with the adjacent small intervals and determine the expansion coefficients on each one separately. We then compute the expansion coefficients on the union. Now we can choose that expansion for which the number of coefficients needed to capture 99% of the energy is smallest. Or, we can choose that expansion whose “cost” is smallest: information cost, coding cost, error cost, . . .

We compare the cost of the chosen expansions on two adjacent unions of pairs to the expansion on their union and again pick the best. We continue until we reach an optimal




Figure 11: Segmentation of a French phrase into voiced and unvoiced portions.

distribution of time windows [Coifman and Wickerhauser, 1992a].

This algorithm segments a signal into portions that are individually easy to describe. It can be combined with an appropriate recognition criterion to segment continuous speech into voiced and unvoiced segments [Wesfreid and Wickerhauser, 1992], as shown in Figure 11, or even phonemes.

The adapted wavelet packet decomposition uses the same choice algorithm, only it produces an adapted segmentation of the frequency axis rather than the time axis.

4 Compression

The *compressibility* of a sampled signal is the ratio of the total area of the phase plane (N , for a signal sampled at N points) divided by the total area of the dark information cells (each of area 1). We may automatically analyze signals by expanding them in the best

basis, then drawing the corresponding phase plane representation. The negligible components need not be drawn, as it is not relevant which particular basis is chosen for a subspace containing negligible energy.

As done for the Gaussians in Figure 2, signals can be automatically analyzed in their best wavelet packet bases by a computer program “WPLab” [Rochberg and Wickerhauser, 1992]. The user selects a transform by picking an analyzing quadrature filter from a list of 17 at the right. The “mother wavelet” [Daubechies, 1992] determined by that filter is displayed in the small square window at the lower right, to indicate roughly what the phase plane atoms look like. The signal is plotted in the rectangular window at bottom, and the phase plane representation is drawn in the large main square window. Examples of canonical signals for this analysis are *chirps* (oscillatory signals with increasing modulation), *spikes* (sharp transients), *whistles* (almost periodic functions), and combinations of all three such as human speech. Examples of best basis analyses for real signals are given in Figures 12–20.

5 Adapted waveform “de-noising”

It is also possible to expand a signal in several libraries of waveforms and then to choose the library which best represents it. Or, if no library does particularly well, we can peel off layers of a signal by taking one or a few waveforms out at a time, then re-analyzing the remainders. These are examples of *meta-algorithms* which are used at a high level to choose an appropriate analysis for the given signal.

As an application combining these ideas we now describe an algorithm for *de-noising* or, more precisely, *coherent structure extraction*. This is a difficult and ill-defined problem, not least because what is “noise” is not always well-defined. We chose instead to view an N -sample signal as being noisy or incoherent relative to a basis of waveforms if it does not correlate well with the waveforms of the basis, i.e., if its entropy is of the same order of




Figure 12: Wavelet-packet best-basis analysis of a linear chirp.




Figure 13: Wavelet-packet best-basis analysis of a quadratic chirp.

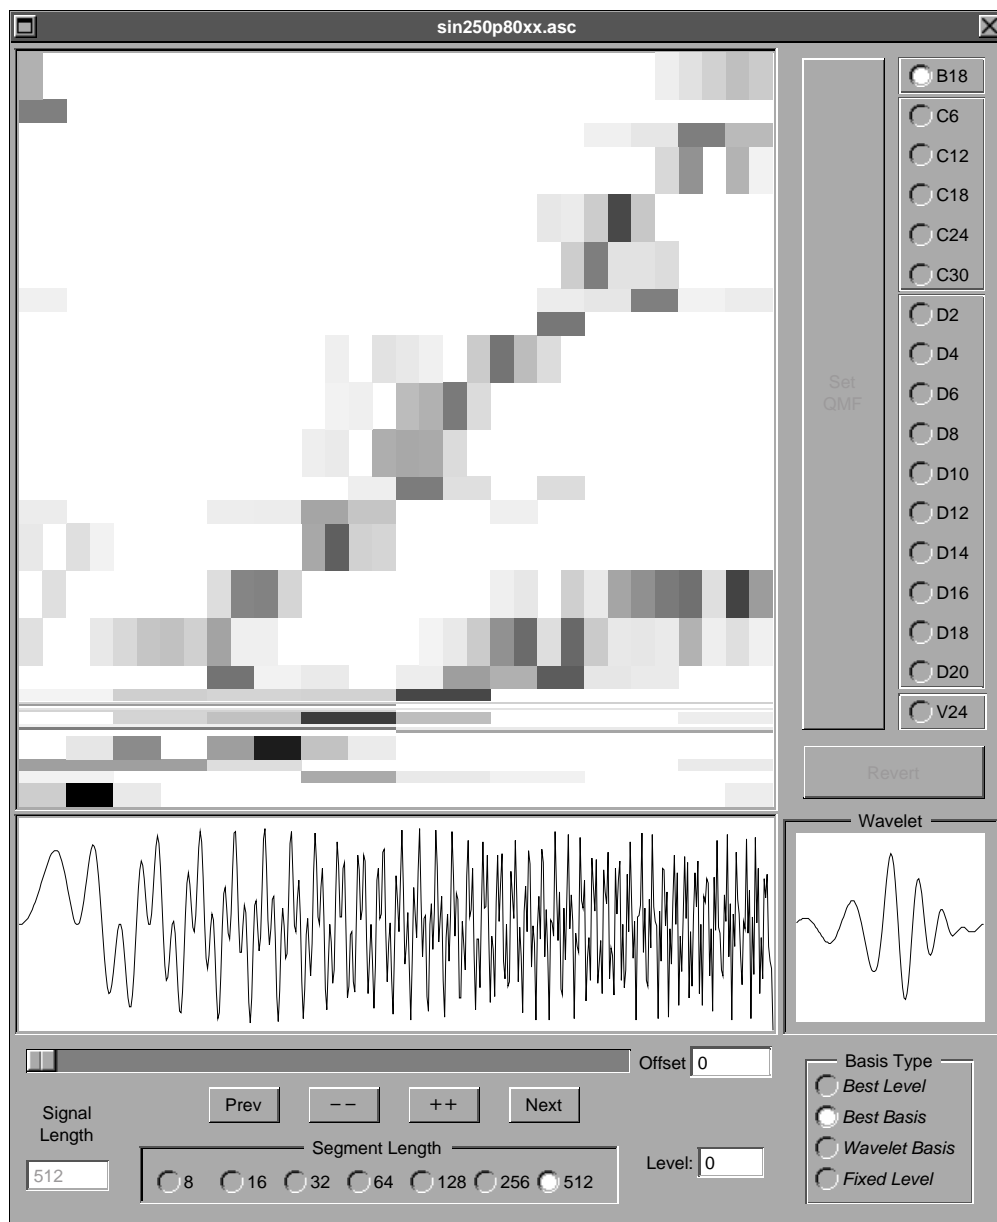


Figure 14: Wavelet-packet best-basis analysis of superposed chirps.




Figure 15: Wavelet-packet best-basis analysis of a sharp transient.




Figure 16: Wavelet-packet best-basis analysis of a decaying transient.




Figure 17: Wavelet-packet best-basis analysis of a narrow wave packet.

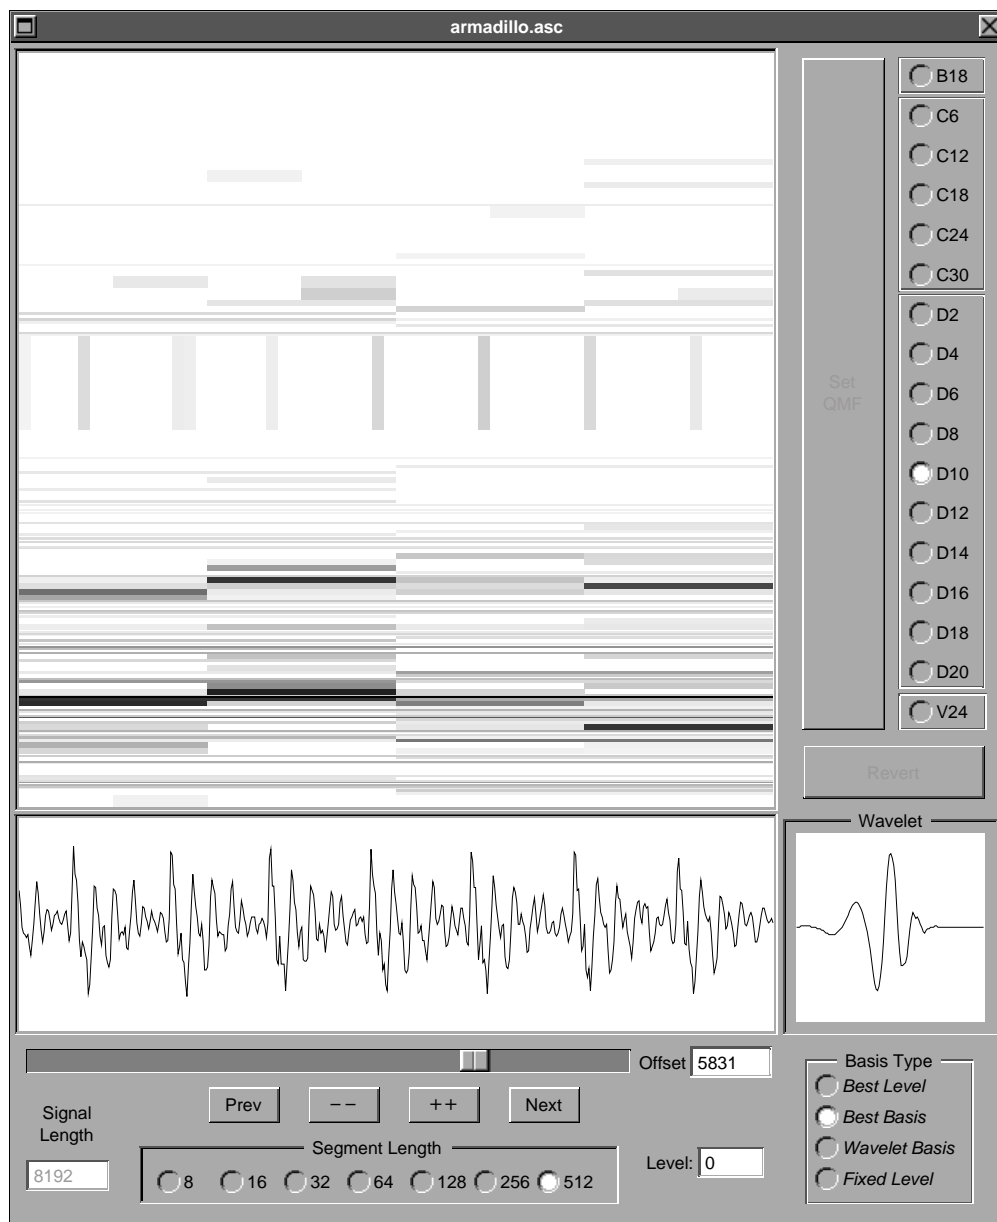


Figure 18: Wavelet-packet best-basis analysis of part of one word.




Figure 19: Wavelet-packet best-basis analysis of part of a phrase.




Figure 20: Wavelet-packet best-basis analysis of part of a poem.

magnitude as

$$\log(N) - \epsilon \tag{1}$$

with small ϵ . From this notion, we are led to the following iterative algorithm based on the several previously-defined libraries of orthonormal bases. We start with a signal f of length N , find the best basis in each library and select from among them the “best” best basis, the one minimizing the cost of representing f . We put the coefficients of f with respect to this basis into decreasing order of amplitude.

The rate at which the coefficients decrease controls the *theoretical dimension* N_0 , which is a number between 1 and N describing how many of the coefficients are significant. We can define N_0 in several ways; the simplest is to count the coefficients with amplitudes above some threshold. Another is to exponentiate the entropy of the coefficient sequence, which matches the criterion in Equation 1.

Theoretical dimension is a kind of information cost. We will say that the signal is *incoherent* if its theoretical dimension is greater than a preset “bankruptcy” threshold $\beta > 0$. The threshold β is chosen to determine if unacceptably bad compression was obtained even with the best choice of waveforms. This condition terminates the iteration when further decompositions gain nothing.

If the signal is not incoherent, then we can pick a fraction $\delta > 0$ and decompose f into $c_1 + r_1$, where c_1 is reconstructed from the δN_0 big coefficients, while r_1 is the remainder reconstructed from the small ones. We proceed by using r_1, r_2, \dots as the signal and iterating the decomposition. The procedure is depicted in Figure 21. We can stop after a fixed number of decompositions, or else we can iterate until we are left with a remainder whose theoretical dimension exceeds β . We then superpose the coherent parts to get the coherent part of the signal. What remains qualifies as *noise* to us, because it cannot be well-represented by any sequence of our adapted waveforms. Thus the adapted waveform de-noising algorithm peels a particular signal into layers; we take as many of the top layers

as we want, assured that the bottom layers are not cost-effective to represent.

The two parameters, β and δ , can be adjusted to match an *a priori* estimate of the signal-to-noise ratio, or can be adjusted by feedback to get the cleanest-looking signal if no noise model is known.

Adapted waveform de-noising is a fast approximate version of the *matching pursuit* procedure described by Mallat [Mallat and Zhang, 1992]. There the waveforms are Gaussians, and just one component is extracted at each iteration. That procedure always produces the best decomposition, at the cost of many more iterations plus more work for each iteration. Mallat's stopping criterion is to test the amplitude ratio of successive extracted amplitudes; this is a method of recognizing remainders which have the statistics of random noise.

As an example, we start with a mechanical rumble masked by the noise of aquatic life, recorded through an underwater microphone. The calculations were performed by the program "denoise" [Majid, 1992, Coifman et al., 1992a], using $\delta = 0.5$ and manually limiting the number of iterations to 4. This application is also described in [Coifman and Wickerhauser, 1992b]; the data was provided through R. R. Coifman. Figure 22 shows the original signal paired with its de-noised version. Note that very little smoothing of the signal has taken place. Figures 23, 24, 25 and 26 respectively show the coherent parts and the remainders of the first 4 iterations. Notice how the total energy in each successive coherent part decreases, while the remainders continue to have roughly the same energy as the original.

Figures 27, 28, 29 and 30 respectively show the successive reconstructions from the coherent parts paired with a plot of the best-basis coefficient amplitudes of the remainders, rearranged into decreasing order. A visual estimate of the theoretical dimension from these plots gives evidence after the fact that little is gained after 4 iterations.

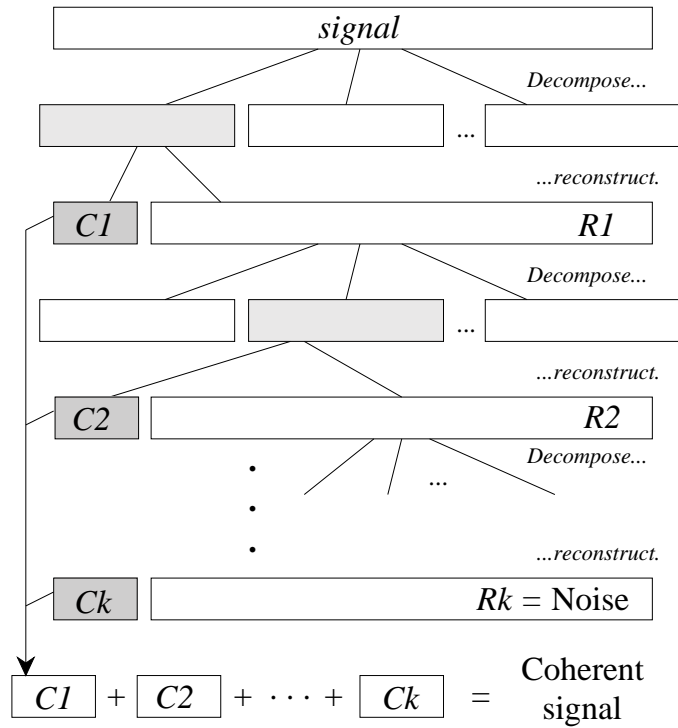


Figure 21: Schematic of adapted waveform de-noising.

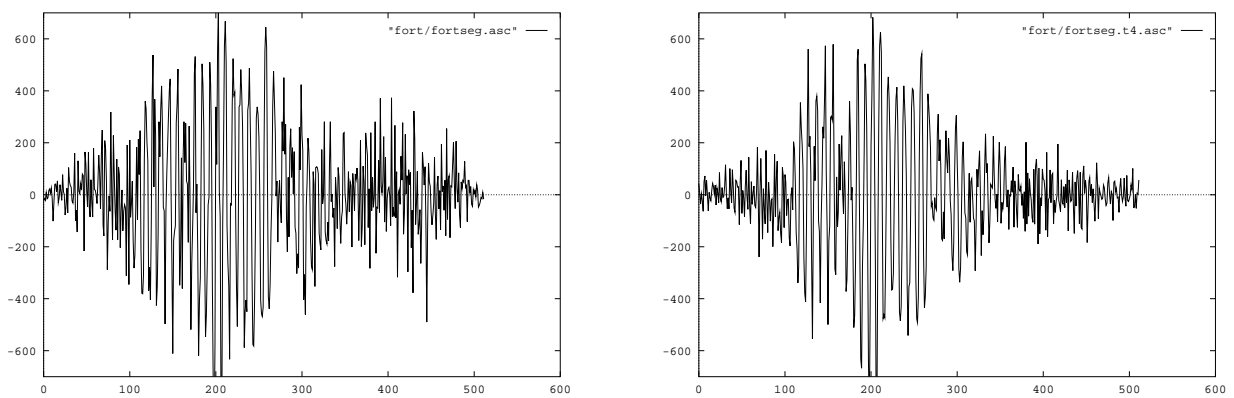


Figure 22: Original and de-noised signals.





Figure 23: First coherent part and first remainder part.





Figure 24: Second coherent part and second remainder part.





Figure 25: Third coherent part and third remainder part.





Figure 26: Fourth coherent part and fourth noise part.





Figure 27: First reconstruction and its sorted remainder coefficients.





Figure 28: Second reconstruction and its sorted remainder coefficients.





Figure 29: Third reconstruction and its sorted remainder coefficients.





Figure 30: Fourth reconstruction and its sorted remainder coefficients.

6 Conclusion

We have analyzed one-dimensional sampled continuous waveforms by decomposing them into building blocks or atoms. The main desirable feature is that the atoms be well-localized both in the time domain and in the frequency domain, in the manner of windowed sines and cosines or Gabor functions. Wavelets and related functions like wavelet packets and local cosines can also be used in such decompositions. They have the added advantage that the resulting expansions are orthogonal or energy-preserving, allowing us to compare and adapt expansions to signals in order to minimize the cost of representation. Such adapted decompositions perform compression and analysis simultaneously. We have designed an idealized graphical presentation of the time-frequency information obtained by such a best-adapted waveform analysis, and from such presentations we can recognize and extract transient features such as parts of speech. Finally, we have shown that the negligibly small components in the analysis may be treated as noise and discarded, and we have designed an iterative algorithm for extracting coherent signals from such noise. Adapted wavelet analyses are practical for realistic signal sizes because the underlying algorithms have low computational complexity.

References

- [Auscher et al., 1992] Auscher, P., Weiss, G., and Wickerhauser, M. V. (1992). Local sine and cosine bases of Coifman and Meyer and the construction of smooth wavelets. In Chui, C. K., editor, *Wavelets—A Tutorial in Theory and Applications*, pages 237–256. Academic Press, Boston. ISBN 0-12-174590-2.
- [Beylkin et al., 1991] Beylkin, G., Coifman, R. R., and Rokhlin, V. (1991). Fast wavelet transforms and numerical algorithms I. *Communications on Pure and Applied Math*, XLIV:141–183.

- [Coifman et al., 1992a] Coifman, R. R., Majid, F., and Wickerhauser, M. V. (1992a). Denoise. Available from Fast Mathematical Algorithms and Hardware Corporation, 1020 Sherman Ave., Hamden, CT 06514 USA.
- [Coifman and Meyer, 1991] Coifman, R. R. and Meyer, Y. (1991). Remarques sur l'analyse de Fourier à fenêtre. *Comptes Rendus de l'Académie des Sciences*, 312:259–261.
- [Coifman et al., 1990] Coifman, R. R., Meyer, Y., Quake, S. R., and Wickerhauser, M. V. (1990). Signal processing and compression with wavelet packets. In [Meyer and Roques, 1992], pages 77–93. ISBN 2-86332-130-7.
- [Coifman et al., 1992b] Coifman, R. R., Meyer, Y., and Wickerhauser, M. V. (1992b). Wavelet analysis and signal processing. In Ruskai, M. B. et al., editors, *Wavelets and Their Applications*, pages 153–178. Jones and Bartlett, Boston. ISBN 0-86720-225-4.
- [Coifman and Wickerhauser, 1992a] Coifman, R. R. and Wickerhauser, M. V. (1992a). Entropy based algorithms for best basis selection. *IEEE Transactions on Information Theory*, 32:712–718.
- [Coifman and Wickerhauser, 1992b] Coifman, R. R. and Wickerhauser, M. V. (1992b). Wavelets and adapted waveform analysis. In Benedetto, J. J. and Frazier, M., editors, *Wavelets: Mathematics and Applications*, Studies in Advanced Mathematics, pages 399–423. CRC Press. ISBN 0-8493-8271-8.
- [Daubechies, 1988] Daubechies, I. (1988). Orthonormal bases of compactly supported wavelets. *Communications on Pure and Applied Mathematics*, XLI:909–996.
- [Daubechies, 1992] Daubechies, I. (1992). *Ten Lectures on Wavelets*, volume 61 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. SIAM Press, Philadelphia, Pennsylvania. ISBN 0-89871-274-2.

- [Majid, 1992] Majid, F. (1992). Applications des paquets d'ondelettes au débruitage du signal. Preprint, Yale University. Rapport d'Option, Ecole Polytechnique.
- [Mallat and Zhang, 1992] Mallat, S. G. and Zhang, Z. (1992). Matching pursuit with time-frequency dictionaries. Preprint, Courant Institute of Mathematical Sciences, New York University, New York.
- [Malvar, 1990] Malvar, H. (1990). Lapped transforms for efficient transform/subband coding. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 38:969–978.
- [Meyer and Roques, 1992] Meyer, Y. and Roques, S., editors (1992). *Progress in Wavelet Analysis and Applications*, Proceedings of the International Conference “Wavelets and Applications”, Toulouse, France. Observatoire Midi-Pyrénées de l'Université Paul Sabatier, Editions Frontieres. ISBN 2-86332-130-7.
- [Rochberg and Wickerhauser, 1992] Rochberg, D. and Wickerhauser, M. V. (1992). WPLab version 3.03 (for NeXT computers). Available by anonymous ftp from [wuarchive, 1991]. Download the file at the path doc/techreports/wustl.edu/math/software/WPLab3.03.tar.Z.
- [Wesfreid and Wickerhauser, 1992] Wesfreid, E. and Wickerhauser, M. V. (1992). Adapted local trigonometric transform and speech processing. Preprint 9236, CEREMADE–University of Paris–Dauphine. Submitted to IEEE Transactions on Signal Processing.
- [Wickerhauser, 1991] Wickerhauser, M. V. (1991). Fast approximate factor analysis. In *Curves and Surfaces in Computer Vision and Graphics II*, volume 1610, pages 23–32, Boston. SPIE. ISBN 0-8194-0747-X.

- [Wickerhauser, 1992a] Wickerhauser, M. V. (1992a). Computation with adapted time-frequency atoms. In [Meyer and Roques, 1992], pages 175–184. ISBN 2-86332-130-7.
- [Wickerhauser, 1992b] Wickerhauser, M. V. (1992b). High-resolution still picture compression. *Digital Signal Processing: a Review Journal*, 2(4):204–226.
- [Wickerhauser, 1994] Wickerhauser, M. V. (1994). *Adapted Wavelet Analysis from Theory to Software*. AK Peters, Ltd., Wellesley, Massachusetts. To appear.
- [wuarchive, 1991] wuarchive (1991). Washington university in st. louis. InterNet Anonymous File Transfer (ftp) Site wuarchive.wustl.edu [128.252.135.4].