

WAVELET SOFT-THRESHOLDING OF TIME-FREQUENCY REPRESENTATIONS

Richard G. Baraniuk*

Department of Electrical and Computer Engineering
Rice University
P.O. Box 1982, Houston, TX 77251-1892, USA

ABSTRACT

The high noise sensitivity of the Wigner distribution makes smoothing a necessity for producing readable time-frequency images of noise corrupted signals. Since linear smoothing suppresses noise at the expense of considerable smearing of the signal components, we explore two nonlinear denoising techniques based on soft-thresholding in an orthonormal basis representation. Soft-thresholding provides considerable noise reduction without greatly impairing the time-frequency resolution of the denoised distribution.

1. INTRODUCTION

Time-frequency representations (TFRs), which map one-dimensional signals into two-dimensional images indicating their joint time-frequency energy content [1], have proven indispensable in a wide range of applications, including speech, music, acoustics, biology, radar, sonar, and geophysics. The most popular TFR is undoubtedly the short-time Fourier transform

$$Z_s(t, f) = \int s(\tau) w^*(\tau - t) e^{-j2\pi f\tau} d\tau, \quad (1)$$

which computes a local Fourier analysis by sliding a window function w along the signal s . (Here, t and f represent time and frequency, respectively.) The squared magnitude of this transform is known as the spectrogram.

While the spectrogram is a powerful time-frequency analysis tool, the window function w limits its resolution in time-frequency. Since good time resolution requires a narrow window but good frequency resolution requires a wide window, high resolution simultaneously in both directions is unattainable with this transform. Figure 1 shows a spectrogram TFR of a test signal composed of two Gaussian components modulated to different center frequencies.

This major limitation of the spectrogram prompted the development of bilinear TFRs that attempt to match the window function to the signal. A primary example is the Wigner distribution

$$W_s(t, f) = \int s\left(t + \frac{\tau}{2}\right) s^*\left(t - \frac{\tau}{2}\right) e^{-j2\pi f\tau} d\tau, \quad (2)$$

*This work was supported by the Texas Advanced Technology Program, grant no. TX-ATP 003604-002.

which can be interpreted as a scaled short-time Fourier transform computed using the time-reversed signal as window. Because of this window matching, the Wigner distribution possesses excellent time-frequency resolution. Figure 2 illustrates the Wigner TFR of the same test signal.¹

Unfortunately, the bilinear nature of the Wigner distribution results in a high noise sensitivity, limiting its application to only virtually noise-free signals. Figure 3, for example, shows the Wigner distribution of the test signal immersed in 6dB SNR additive white Gaussian noise. For applications involving noise corrupted signals, some amount of smoothing must be applied to the Wigner distribution in order to obtain readable time-frequency images. Current methods smooth via a linear convolution with a kernel function ϕ . The resulting TFRs belong to Cohen's class of bilinear TFRs [1] and can be written as

$$C_s(t, f) = \iint W_s(u, v) \phi(t - u, f - v) du dv. \quad (3)$$

While lowpass kernels smooth and suppress noise in the Wigner distribution, they also smear the signal components that we wish to view at full resolution. For example, Figures 4 and 5 illustrate the spectrogram ($\phi =$ the Wigner distribution of the analysis window w) and the Choi-Williams distribution ($\phi =$ the Fourier transform in θ, τ of the function $\exp(-\theta^2\tau^2/\sigma)$ [1]) of the noisy test signal. While the noise has been largely suppressed in both images, the two Gaussian components have been blurred.

Taking the Fourier transform of (3) yields an alternate interpretation of Cohen's class TFRs as weighting the narrow-band ambiguity function A_s of the signal by a smoothing function Φ

$$W_s \xrightarrow{2-d FT} A_s \xrightarrow{\text{weight}} \Phi A_s \xrightarrow{2-d FT^{-1}} C_s \quad (4)$$

The functions A_s and Φ correspond to the Fourier transforms (FTs) of the Wigner distribution and smoothing kernel ϕ , respectively.

Since all linearly smoothed TFRs share a common noise reduction vs. smearing tradeoff, in this paper we will explore

¹Besides increasing the time-frequency resolution of the Wigner distribution, its nonlinearity also introduces artifacts called cross-components [1] that appear midway between each pair of true signal components.

two nonlinear denoising techniques. The *soft-thresholding* algorithms we will consider can provide considerable noise reduction without greatly impairing the time-frequency resolution of the TFR. Wavelet soft-thresholding algorithms — multiresolution denoising techniques applicable to a wide range of signals and images — construct nonlinear estimates of signals or images embedded in additive white Gaussian noise using a simple three-step procedure [2]: (1) compute the wavelet transform of the data; (2) translate (“soft-threshold”) the wavelet coefficients towards zero by a set threshold value; (3) invert the modified wavelet coefficients to obtain the final estimate.

When applied to the Wigner distribution (which can be interpreted as a two-dimensional image), wavelet soft-thresholding corresponds to nonlinear, scalar processing of the coefficients of this distribution in a wavelet basis representation

$$W_s \xrightarrow{2-d \text{ WT}} B_s \xrightarrow{\text{threshold}} \Gamma_\gamma(B_s) \xrightarrow{2-d \text{ WT}^{-1}} D_s. \quad (5)$$

Here B_s represents the wavelet transform (WT) coefficients of W_s , Γ_γ represents soft-thresholding with threshold γ , and D_s represents the denoised TFR. In contrast, Cohen’s class TFRs result from linear scalar processing of the coefficients of the Wigner distribution in the sinusoidal Fourier basis representation (see (4)). Figure 6 illustrates a wavelet soft-thresholded TFR for the same noisy test signal utilized for Figures 3–5. Unlike the spectrogram and Choi-Williams distribution, this TFR offers reduced noise levels without degraded resolution.

After a brief review of wavelet soft-thresholding in Section 2, we discuss its application to TFRs in Section 3. Since wavelet processing of the Wigner distribution sacrifices some of its desirable properties, in Section 4, we introduce soft-thresholding of the ambiguity function representation from (4). We close in Section 5 with some preliminary conclusions. While tantalizing, we will find that since the Wigner distribution of a noisy signal does not conform to the standard additive white Gaussian noise model, the application (or misapplication!) of soft-thresholding techniques to time-frequency analysis remains as ad hoc as previous nonlinear schemes such as Wigner distribution thresholding, median filtering, and so on.

2. WAVELET SOFT-THRESHOLDING

The wavelet transform of a one-dimensional continuous-time signal s is defined as

$$B_s(m, k) = 2^{-k/2} \int s(t) \psi^*(2^{-k}t - m) dt. \quad (6)$$

When the dilates and translates of the wavelet function ψ form an orthonormal basis, we have the signal representation, or inverse wavelet transform

$$s(t) = \sum_{m,k} B_s(m, k) 2^{-k/2} \psi(2^{-k}t - m). \quad (7)$$

Roughly speaking, the wavelet transform of a smooth signal is concentrated in a relatively small number of wavelet coefficients. On the other hand, the transform of a white noise signal spreads out over all coefficients.

The wavelet thresholding concept arose from combining these two observations with the conventional wisdom that simple thresholding performs well as a data recovery technique whenever the data lies above the noise floor. Wavelet thresholding addresses the following data recovery problem (stated in one dimension for simplicity): Recover the smooth, discrete-time signal $s(i)$, $i = 1, \dots, N$, given the corrupted observations $s(i) + n(i)$, where $n(i)$ is a white Gaussian sequence of zero mean and variance σ^2 . The algorithm of Donoho and Johnstone [2] runs as follows:

1. Compute the wavelet transform of $s + n$ using a discrete-time, finite-data analog to (6) (an interval adapted filterbank).
2. Translate all wavelet coefficients $B_{s+n}(m, k)$ towards zero by the amount $\gamma = \sqrt{2 \log(N)/N} \sigma$.
3. Invert the thresholded coefficients using the discrete-time, finite-data analog to (7).

A multidimensional wavelet transform [3] extends this procedure to image and other data in higher dimensions.

3. WAVELET SOFT-THRESHOLDING THE WIGNER DISTRIBUTION

In addition to being straightforward and intuitively reasonable, wavelet soft-thresholding possesses two remarkable properties [2], both potentially useful for TFR denoising. First, with high probability, the data estimate is at least as smooth as the desired noise-free data. Thus, given a smooth set of Wigner distribution signal components² embedded in noise, wavelet denoising should not introduce artifacts that could be interpreted as new components. Second, the estimate achieves almost the minimax mean-square error over every one of a wide variety of smoothness measures, including many where linear estimators do not and cannot achieve the minimax value. Thus, nonlinear denoising of the Wigner distribution should offer higher performance than linear smoothing. Simulations support this intuition; Figure 6 illustrates a wavelet soft-thresholded Wigner distribution for the same noisy test signal utilized for Figures 3–5.

Unfortunately, it appears difficult to go beyond simulations for justifying wavelet soft-thresholding in this context, because our data recovery model does not match that for which the algorithm was developed. In particular, the Wigner distribution of the signal $s + n$, given by

$$W_{s+n} = W_s + W_n + 2 \operatorname{Re} W_{s,n} \quad (8)$$

²The signal components in a Wigner distribution time-frequency image are always at least absolutely continuous, due to the integral in (2).

where the last term involves the cross-Wigner distribution

$$W_{s,n}(t, f) = \int s\left(t + \frac{\tau}{2}\right) n^*\left(t - \frac{\tau}{2}\right) e^{-j2\pi f\tau} d\tau, \quad (9)$$

corresponds to data W_s plus interference $W_n + 2 \operatorname{Re} W_{s,n}$. This interference is anything but Gaussian and white: $W_{s,n}$ is Gaussian, yet highly correlated and signal-dependent, while W_n is neither Gaussian nor uncorrelated. Further complicating matters, note that since $W_{s,n}$ has variance proportional to $\|s\|^2 \sigma^2$ and W_n has variance proportional to σ^4 , one term will dominate depending on the particular value of SNR.

Nevertheless, as very little is known about the probability density of W_n , ad hoc methods such as thresholding must suffice until a more complete theory for Wigner distribution estimation can be derived.³ Some progress has been made with the stationary power spectrum [5]; a similar approach might prove useful here and result in an explicit formula for the threshold γ . At present, we take γ either as a free parameter or adjust it automatically to optimize some measure of TFR performance as in [6,7].⁴

The denoising provided by wavelet soft-thresholding comes at some expense in terms of the desirable mathematical properties of the Wigner distribution [1]. First, due to the nonlinearity of the processing, the energy preservation and marginal properties fail to hold true. Second, the time-frequency shift covariance property is lost: Because the discrete wavelet transform is not covariant to shifts, a time-frequency shift in the Wigner distribution will result in a different thresholding pattern and thus a slightly different denoised TFR. Third, with separable wavelet processing, the rotation covariance property of the Wigner distribution abandons us as well (although it should be noted that nonadaptive linear smoothing cannot retain rotation covariance either).

4. SOFT-THRESHOLDING THE AMBIGUITY FUNCTION

The wavelet transform proves so useful as a soft-thresholding basis transformation, because wavelets form unconditional bases for an incredible variety of signal spaces, including most of those related to smoothness [2]. Sinusoids are more limited in their utility for soft-thresholding, because they do not form unconditional bases for most of these spaces. Nevertheless, in light of the loss of desirable TFR properties mentioned in the previous section, it appears reasonable to consider also the Fourier basis for soft-thresholding the Wigner distribution.

³Obviously, denoising the signal and then computing the Wigner distribution of the result avoids any problems with the additive white Gaussian noise requirement. However, signal components modulated to lower frequencies are "smoother" than their counterparts at higher frequencies, and hence they are preferentially treated by the wavelet transform before thresholding. This favoritism is demonstrated in Figure 7, where we show the Wigner distribution of the denoised test signal. For modulated signals, the Wilson bases [4] probably represent a better alternative to wavelets for soft-thresholding.

⁴Even when the additive white Gaussian noise model is valid, determination of the noise power σ can be very difficult in practice, usually requiring some experimentation or optimization.

The resulting scheme fits in the framework of (4), but with a now signal-dependent kernel Φ_s that soft-thresholds the ambiguity function of the signal

$$W_s \xrightarrow{2-d \text{ FT}} A_s \xrightarrow{\text{threshold}} \Gamma_\gamma(A_s) \xrightarrow{2-d \text{ FT}^{-1}} E_s. \quad (10)$$

Note that E_s belongs to Cohen's class and is time-frequency shift covariant; additional constraints can be imposed on the thresholding to ensure that it satisfies other properties such as energy preservation and marginals, if desired. Furthermore, the Rényi information measures [7] can be utilized to optimize the threshold value. Figure 8 illustrates a TFR arising from soft-thresholding the ambiguity function of the noisy test signal; it closely resembles the wavelet-denoised TFR of Figure 6. Less ad hoc approaches to adaptive kernel design are detailed in [6,8].⁵

5. CONCLUSIONS

While our results are preliminary and admittedly somewhat ad hoc, nonlinear smoothing techniques have potential for providing time-frequency analyses with Wigner-like resolution down to low SNRs. The hallmarks of the wavelet soft-thresholding technique — simplicity, use of information across scales, smoothness preservation, and near optimality for additive white Gaussian noise — remain tantalizing, but more work is required in order to justify its application to TFRs. It is likely that a detailed analysis of the correlated, nonGaussian interference will inspire modifications to the algorithm, with a corresponding performance increase.

Finally, we note that soft-thresholded representations of time t and scale a (related to the continuous wavelet transform and the scalogram [3]) can be obtained simply by processing the reparameterized Wigner distribution $W_s(t, f_0/a)$.

REFERENCES

1. L. Cohen, "Time-frequency distributions — A review," *Proc. IEEE*, vol. 77, no. 7, 1989.
2. D. Donoho, "De-noising by soft-thresholding," preprint.
3. O. Rioul and M. Vetterli, "Wavelets and signal processing," *IEEE Signal Proc. Mag.*, October, 1991.
4. I. Daubechies, S. Jaffard, J. Journé, "A simple Wilson basis with exponential decay," *SIAM J. Math. Anal.*, vol. 22, no. 2, 1991.
5. P. Moulin, "Wavelet thresholding techniques for power spectrum estimation," *IEEE Trans. Signal Proc.*, preprint.
6. R. Baraniuk and D. Jones, "A signal-dependent time-frequency representation: Optimal kernel design," *IEEE Trans. Signal Proc.*, vol. 41, no. 4, 1993.
7. P. Flandrin, R. Baraniuk, O. Michel, "Time-frequency complexity and information," *Proc. ICASSP '94*.
8. A. Sayeed and D. Jones, "Optimal kernels for Wigner-Ville spectral estimation," *Proc. ICASSP '94*.

⁵In fact, the algorithm (10) is closely related to a related version of the "1/0" optimal kernel design from [6].

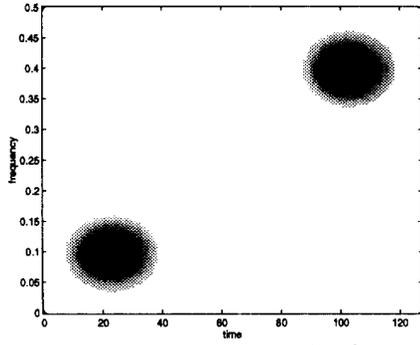


Fig. 1: Spectrogram of the test signal.

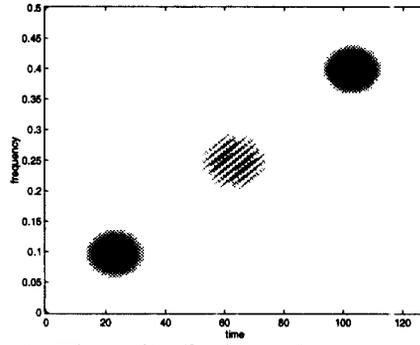


Fig. 2: Wigner distribution of the test signal.

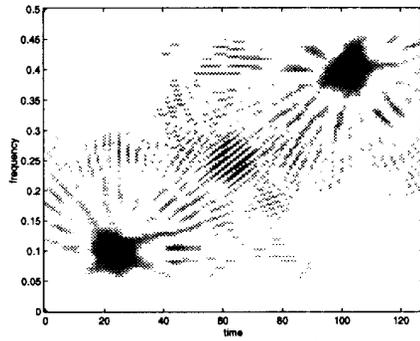


Fig. 3: Wigner distribution of noisy test signal.

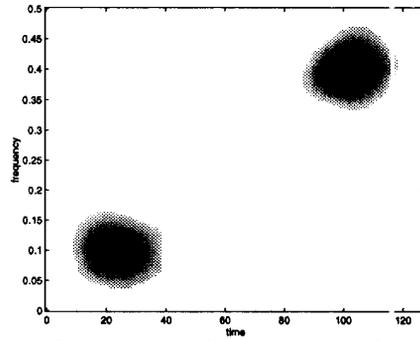


Fig. 4: Spectrogram of noisy test signal.

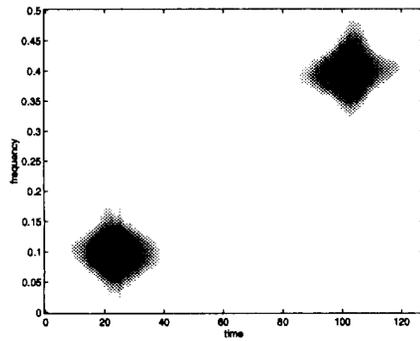


Fig. 5: Choi-Williams distribution.

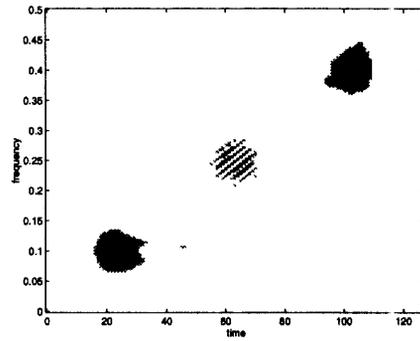


Fig. 6: Wavelet soft-thresholded TFR.

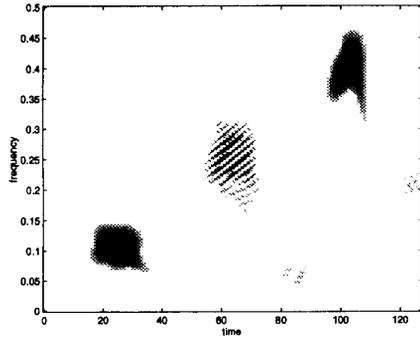


Fig. 7: Wigner distribution of wavelet soft-thresholded noisy test signal.

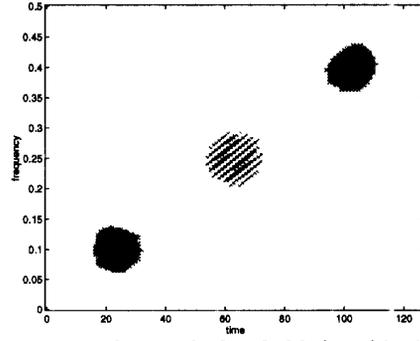


Fig. 8: TFR from soft-thresholded ambiguity function.