IMPROVED POWER-LAW DETECTION OF TRANSIENTS

Zhen Wang and Peter Willett

U-157, University of Connecticut
Storrs, CT 06269
e-mail: willett@engr.uconn.edu

ABSTRACT

Recently, a power-law statistic operating on DFT data has emerged as a basis for a remarkably robust detector of transient signals having unknown structure, location and strength. In this paper we offer a number of improvements to the original power-law detector. Specifically, the power-law detector requires that its data be pre-normalized and spectrally white; a CFAR and self-whitening version is developed and analyzed. Further, it is noted that transient signals tend to be contiguous both in temporal and frequency senses, and consequently new power-law detectors in the frequency and the wavelet domains are given. The resulting detectors offer exceptional performance and are extremely easy to implement. There are no parameters to tune, and they may be considered "plug-in" solutions to the transient detection problem.

1. INTRODUCTION

There has been significant recent attention to Nuttall's power-law detector [3, 4] due to its simple implementation and good performance. The test is based on the following formulation. Under the signal-absent hypothesis (H_0) – that the time-domain data is complex white Gaussian noise – pre-processing by the magnitude-square DFT yields independent and identically distributed (iid) exponential random variates. Under the signal- present hypothesis (H_1), the DFT observations are no longer a homogeneous population of exponentials; Nuttall's basic assumption is that there are two positive exponential populations:

$$\begin{aligned} \mathbf{H}_0: \quad f(\mathbf{X}) &= \prod_{k=1}^{N} \frac{1}{\mu_0} e^{-X_k/\mu_0} \qquad (1) \\ \mathbf{H}_1: \quad f(\mathbf{X}) &= \prod_{k \notin S} \frac{1}{\mu_0} e^{-X_k/\mu_0} \prod_{k \in S} \frac{1}{\mu_1} e^{-X_k/\mu_1} \end{aligned}$$

This research was supported by NUWC through contract N66604-99-1-5021, and by ONR through contract N00014-98-1-0049.

where N is the total number of FFT bins, $\mathbf{X}$ are the magnitude-squared FFT bins and S indicates a subset with size M. It is assumed that M signal-present bins are uniformly distributed among the N FFT bins.

Nuttall developed power-law statistics [3] as an approximation to the optimal detector, and these have the form:

$$T(\mathbf{X}) = \sum_{k=1}^{N} {X_k}^{\nu} \qquad (2)$$

where ν is an adjustable exponent. Through extensive computational work it has been found that a good compromise value for ν is 2.5 when information about M is completely unavailable. As an extension of power-law to unknown noise level (μ_0) cases, a constant false-alarm rate (CFAR) version was introduced [5]:

$$T_{cpl}(\mathbf{X}) = \frac{\sum_{k=1}^{N} {X_k}^{\nu}}{(\sum_{k=1}^{N} X_k)^{\nu}} \qquad (3)$$

Clearly, T_{cpl} is not affected by a scale factor.

The statistic (2) is remarkably good; this paper has been written because:

1. The statistic (2) is designed with white noise of known power in mind; the fact is that the performance of (3) is disappointing in white noise, while for colored noise it has very little appeal.

2. The statistic (2) is essentially optimal [3] given its frequency-domain model of (1) when there is nothing whatever known about the signal-bearing set S. There is some tendency for real transient signals to aggregate their energy in a band.

3. There is no reason why a DFT must be the pre-processing step: we investigate the extension that a transform other than the DFT be used.

Basically, the power-law detector is as yet neither a plug-in solution nor is it as good as it can be, and we offer some remedy here.

2. THE NEW CFAR POWER-LAW DETECTOR

The focus of this section is to detect transients buried in colored noise with unknown but stationary spectrum. We write in a matrix a block of NL time domain observations as $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_L)$, where $\mathbf{x}_i$ is a column vector of dimension N whose k^{th} element is the time sample of index $(i-1)L+k$. We immediately transform each column to its magnitude-squared frequency domain equivalent $\mathbf{X}_i$, and record $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2, ..., \mathbf{X}_L)$.

Following ideas similar to those frequently used in radar CFAR processing (e.g. [2]) we define the new power-law statistic as

$$T_{fc}(\mathbf{X}) = \sum_{j=1}^{N} z_j^{\nu}, \text{ where } z_j = \frac{X_{jL}}{\frac{1}{L-1}\sum_{i=1}^{L-1} X_{ji}} \quad (4)$$

where ν is a real exponent.

The best value for the power ν in (4) is in general strongly dependent on M, the number of signal-present bins – this is not at all desirable, since our goal is to find a detection structure which does not depend on knowledge of such signal qualities. Now, given a statistic T, its output SNR can be expressed as

$$SNR_T = \frac{(E(T|\mathbf{H}_1) - E(T|\mathbf{H}_0))^2}{var(T|\mathbf{H}_0)} \quad (5)$$

We define the SNR loss as

$$ISL(\nu, M) = St(\nu, M) - St_{opt}(M)$$

where $St_{opt}(M) = \sqrt{SNR.M}$ indicates the required input aggregate SNR for the optimal statistic at a particular bandwidth M. The ISL measures the input aggregate SNR which is sacrificed through use of a fixed exponent ν, as compared to the best possible exponent ν for that M or the corresponding optimal statistic. In [7] it is found from examination of the ISL that $\nu = 2.5$ is the best choice, a result corroborating that in [3]. We apply the input SNR-loss analysis to the detector in (4), to select ν. Example results are shown in figure 1. Generally, $1.5 < \nu < 2$ is a good choice when information of M is completely unknown.

3. DETECTORS BASED ON CONTIGUITY IN THE FREQUENCY DOMAIN

There is often a tendency for transient signal energy to aggregate itself in the frequency domain; that is, to be at least somewhat bandlimited. We thus modify Nuttall's assumption that the M signal-present bins are uniformly and independently distributed amongst the record of N.

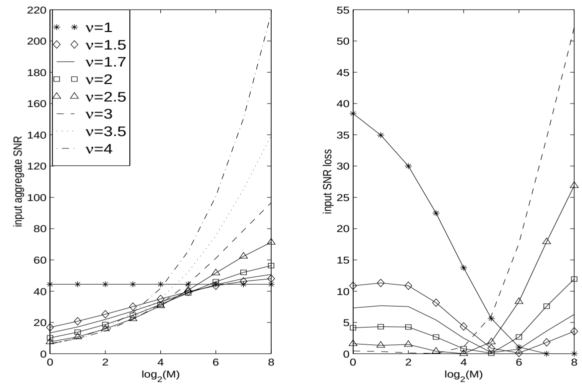

Fig. 1. SNR for CFAR power-law statistics, with settings the output $SNR = 6$, $N = 256$. Left figure: SNR for different ν; right figure: the input SNR-loss for different ν.

3.1. The Prenormalized Case

New random variables are obtained by combining two contiguous frequency bins, and we define new power-law detectors

$$T_{f2}(\mathbf{U}) = \sum_{j=2}^{N} {U_j}^{\nu} = \sum_{j=1}^{N} (X_{j-1} + X_j)^{\nu} \quad (6)$$

where $\{X_j\}$ and N have same meanings as in (1). The statistic of (6) is easily extended as

$$T_{f3}(\mathbf{U}) = \sum_{j=3}^{N} {U_j}^{\nu} = \sum_{j=1}^{N} (X_{j-2} + X_{j-1} + X_j)^{\nu} \quad (7)$$

to the case of three contiguous bins, and further extension is straightforward.

3.2. The Self-Normalizing Case

A similar combining process was adopted in the colored noise case by letting $U_{ji} = X_{j-1,i} + X_{j,i}$. This combining approach results in modified model and generates new CFAR power-law detector in the frequency domain as

$$T_{fc2}(\mathbf{U}) = \sum_{j=2}^{N} \left(\frac{U_{jL}}{\frac{1}{L-1}\sum_{i=1}^{L-1} U_{ji}} \right)^{\nu} \quad (8)$$

The similar detector T_{fc3} combines 3 contiguous bins.

4. DETECTORS IN THE WAVELET DOMAIN

For time-domain observations, the DFT transforms a pure "time description" into a pure "frequency description" and thus clearly cannot take advantage of time contiguity. The wavelet transform (WT) finds a good compromise. The original work of Nuttall explored only the case that the pre-processing transformation was the DFT – the extension to other transforms, especially the wavelet transform, is natural. The Haar wavelet [6] is explored due to the fact that a statistic which assumes as little as possible about the transient to be detected is preferable. Analogous to the frequency domain detectors T (the original power-law detector), T_{f3}, T_{cf} and T_{cf3}, we define T_w, T_{w3}, T_{cw} and T_{cw3}. These are essentially the same detectors, with the exception that the pre-processing transform is in a multi-resolution decomposition. Combinations of wavelet coefficients are according to the tree structure of the filter-bank (see [7] for details); due to this, combining two wavelet coefficients (i.e. T_{w2} and T_{cw2}) is inappropriate.

5. PERFORMANCE COMPARISON

5.1. Prenormalized Data

The detection performance of the improved detectors in the frequency and the wavelet domains are compared to the power-law in [3] with power $\nu = 2.5$. We set $N = 256$ and $S_t = 100$ where S_t is the aggregate SNR of transients. The results are shown in figure 2, where the transient signal was created by passing white Gaussian noise through a bandpass FIR filter (the number of signal-energy-containing FFT bins $M \approx 25$). From figures 2 and 3, it is clear that combining 2 or 3 contiguous FFT or wavelet bins together does indeed improve detection performance.

5.2. Self-Normalizing Case.

Results for the self-normalizing case are shown in figure 4 for different values of M, where colored ambient noise is created by passing white Gaussian noise through an FIR filter. The transient, of duration 50, has the same PSD as the noise. The exponent $\nu = 1.5$ is used for all power-law detectors. We further plot P_d vs. aggregate SNR in figures 5 with $M \approx 10$. Clearly, for each ν, combining contiguous FFT or wavelet bins will improve the performance over S_t.

6. SUMMARY

In [3] Nuttall derived and justified a new and easy-to-implement statistic for the detection of short-duration (transient) signals: the sum of magnitude-square DFT outputs from a block of N time domain data, each raised to a power typically in the range 2-3. This test has been found to be very effective indeed.

detector	provenance	pre-proc.	comb.	CFAR
T	Nuttall	DFT	1	no
T_{cpl}	Nuttall	DFT	1	partial
T_{f2}	new	DFT	2	no
T_{f3}	new	DFT	3	no
T_{fc}	new	DFT	1	yes
T_{fc2}	new	DFT	2	yes
T_{fc3}	new	DFT	3	yes
T_w	new	wavelet	1	no
T_{w3}	new	wavelet	3	no
T_{wc}	new	wavelet	1	yes
T_{wc3}	new	wavelet	3	yes

Table 1. Categorization of various transient detectors discussed in this paper.

The power-law detector is almost a plug-in transient detector for all purposes, but not quite: pre-whitened and pre-normalized data is required. We have thus extended the power-law detector to be self-normalizing by raising to an exponent not the DFT data directly, but instead the power in each DFT bin relative to the average power in previous DFTs. In this case a somewhat smaller exponent, in the range 1.5-2, should be used. It has been noted that there is a tendency among real transient signals to be bandlimited to some degree, and hence a combined-bin power-law detector is proposed. It was additionally noted that the power-law dogma of pre-processing via the DFT is open to challenge, and indeed a power-law processor operating on (Haar) wavelets is developed, made self-normalizing, and augmented to use combined bins (since transient signals most transient signals are aggregated not just in frequency, but also in time/scale).

We give a taxonomy in table 1. Our overall conclusion is that although all of these tests work well, the combined/wavelet power-law detector is perhaps the finest of all.

7. REFERENCES

[1] M. Basseville and I. Nikiforov, Detection of Abrupt Changes, Prentice Hall, 1993.

[2] P. Gandhi and S. Kassam, “Analysis of CFAR Processors in Nonhomogeneous Background”, IEEE Transactions on Aerospace and Electronic Systems, pp. 427-444, July 1988.

[3] A. Nuttall, “Detection Performance of Power-Law Processors for Random Signals of Unknown Location, Structure, Extent, and Strength”, NUWC-NPT Technical Report 10,751, Sep., 1994.

[4] A. Nuttall, “Near-Optimum Detection Performance of Power-Law Processors for Random Signals of Unknown Location, Structure, Extent, and

Arbitrary Strengths", NUWC-NPT Technical Report 11,123, April 1996.

[5] A. Nuttall, "Performance of Power-Law Processors with Normalization for Random Signals of Unknown Structure", NUWC-NPT Technical Report 10,760, May 1997.

[6] G. Strang and T. Nguyen, "Wavelets and Filter Banks", Wellesley-Cambridge Press, Wellesley, MA, 1996.

[7] Z. Wang and P. Willett, "All-Purpose and Plug-In Power-Law Detectors for Transient Signals", submitted to IEEE Transactions on Signal Processing, August 2000.

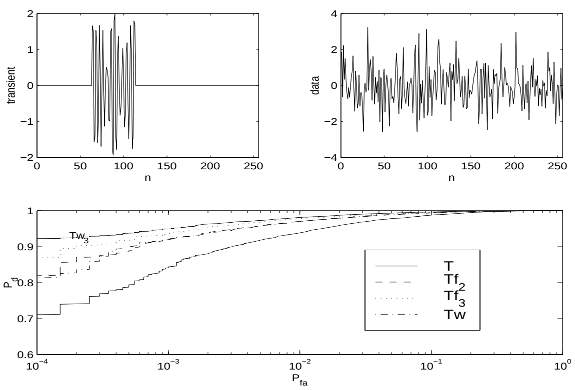

Fig. 2. Receiver operating characteristic (ROC) of new power-law statistics in the frequency and the wavelet domains versus that of Nuttall's power-law detector with $\nu = 2.5$. The upper plots illustrate an example of the transient signal.

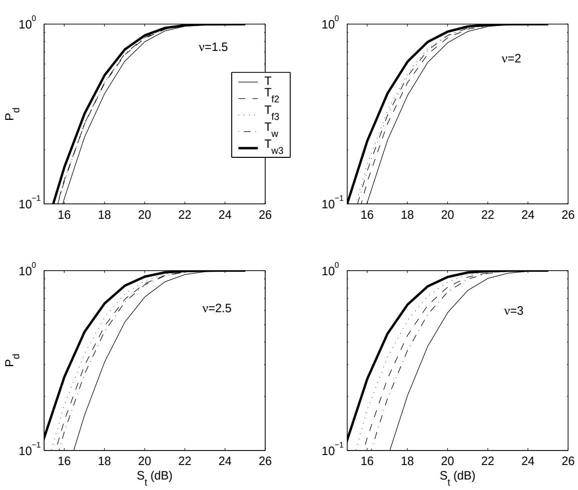

Fig. 3. Similar to figure 2, Detection performances of new power-law statistics in the frequency and the wavelet domains; S_t denotes the aggregate SNR.

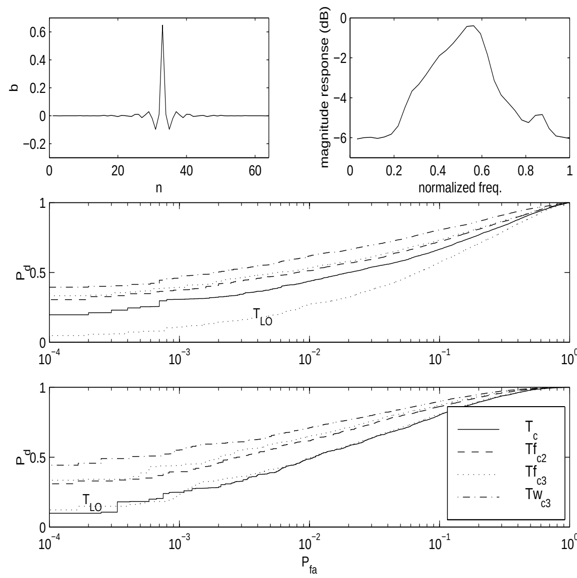

Fig. 4. Receiver operating characteristic (ROC) of statistics in the frequency and the wavelet domains for transient detection in colored noise. T_{LO} is the locally-optimum detector (see [7]).

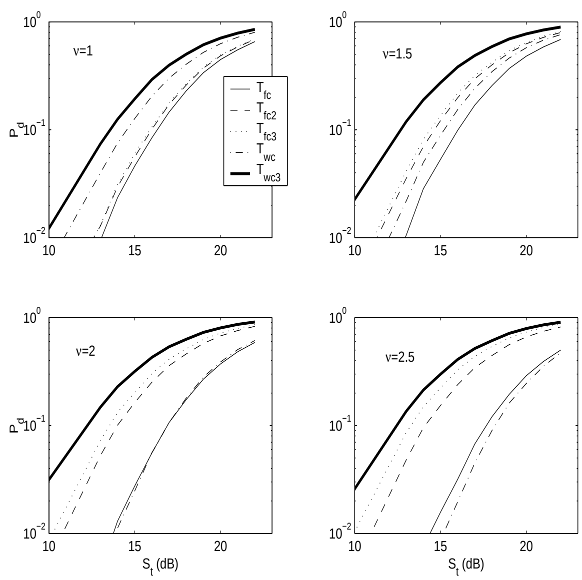

Fig. 5. Detection performance of power-law detectors in the frequency and the wavelet domains for transient detection in colored noise with different ν. Here S_t means the aggregate SNR, $N = 256$ and $M = 10$. The same noise model as in figure 4 is applied, and the same terms as in figure 4 are used.