

Time-frequency based detection in alpha-stable distributed noise environments

Mark Coates and Bill Fitzgerald

Department of Engineering

University of Cambridge

Cambridge UK CB2 1PZ

mjc40@eng.cam.ac.uk, wjf@eng.cam.ac.uk

Ercan Kuruoğlu*

Xerox Research Centre Europe

Cambridge UK CB2 1AB

kuruoglu@xrce.xerox.com

Abstract

The Gaussian distribution is an inadequate description of certain types of real-life noise that display impulsive characteristics. The α -stable distribution has been suggested as a more appropriate model. Bilinear time-frequency representations (TFRs) and time-scale representations (TSRs) have been used to efficiently implement detectors of non-stationary signals in Gaussian noise. In this paper, we extend their application to the case of impulsive noise, modelled as an isotropic α -stable distribution.

1. Introduction

There is a particular class of detection problems for which bilinear time-frequency or time-scale implementation of receivers is advantageous. Two properties characterise these problems. Firstly, we must be able to express a test statistic in a quadratic form. Secondly, there must be an opportunity to exploit the nature of the TFR/TSR, which arises when the signal to be detected is subjected to an unknown time-frequency or time-scale offset.

Methods for the TFR/TSR based realisation of detection statistics for such problems have been outlined for the case where the noise is modelled as Gaussian [6, 7]. In many detection scenarios, however, particularly those involving atmospheric or underwater acoustic noise, the noise displays impulsive characteristics. There has been widespread interest in the use of α -stable distributions to model noise displaying impulsive characteristics, but the development of receivers using such noise models has been hampered by the lack of a compact analytical form. Near-optimal receivers for the detection of known, deterministic signals using α -stable distributed models have recently been proposed in [3].

*The work described in this paper was undertaken whilst E. Kuruoğlu was with the Signal Processing Laboratory, University of Cambridge.

In this paper, we develop test statistics for the detection of arbitrary nonstationary second-order signals, using a bivariate, isotropic α -stable model for the noise. We consider the problem of detecting such signals in the presence of unknown time-frequency and time-scale offsets, and demonstrate that approximations to locally-optimal test statistics can be expressed in a manner conducive to efficient implementation using TFRs or TSRs. The paper is organised in the following manner. Section 2 reviews a method for approximating α -stable noise models. Section 3 describes the LO detector and near optimal approximation for the case when the signal is subject to no time-frequency or time-scale offset. Section 4 develops test statistics for the case when an unknown offset is applied, and describes TFR/TSR implementations. Section 5 discusses the extension to array-based receivers, and Section 6 examines the performance of the developed receivers.

2. Approximating α -stable Distributions

A bivariate isotropic α -stable distribution can be a good model for narrowband impulsive noise [5]. The α -stable distribution is most readily described by its characteristic function, which, for the special case of the isotropic distribution $f(u, v)$, centred at $(0, 0)$, is:

$$\phi(z_1, z_2) = \exp(-\gamma|\mathbf{z}|^\alpha), \quad (1)$$

where $\mathbf{z} = (z_1, z_2)^T$. The distribution $f(u, v)$ is circularly symmetric: $f(u, v) = h(\sqrt{u^2 + v^2}) = h(r)$.

Kuruoğlu *et al.* [4] have described how sub-Gaussian stable distributions can be analytically represented using a mixture of Gaussians. Any isotropic bivariate $S\alpha S$ random variable $\mathbf{Z} = (u, v)$ centred at zero can be represented as a scale mixture of normal distributions:

$$f_{\mathbf{z}}(\mathbf{z}; \alpha, \gamma) = \frac{1}{(2\pi)^{3/2}(2\gamma)^{4/\alpha}} \times \int_{-\infty}^{\infty} \exp\left(-\frac{(2\gamma)^{2/\alpha} \mathbf{z}^T \mathbf{z}}{2v^2}\right) h(v) v^{-1} dv, \quad (2)$$

where $h(v)$ is the *mixing function*, the form of which was given in [4].

By discretising (2) and normalising, we obtain a finite mixture approximation to the pdf:

$$f_z(\mathbf{z}; \alpha, \gamma) = \sum_{i=1}^M \frac{P(i)}{\sqrt{2\pi\sigma_i}} \exp\left(-\frac{\mathbf{z}^T \mathbf{z}}{2\sigma_i^2}\right). \quad (3)$$

where $P(i)$ are mixture weights, and σ_i variance parameters. For this to be a good approximation, the continuous distribution must be sampled at a large number of points, resulting in an unwieldy, high-order mixture model. It is preferable to sample at fewer locations, but then train the model parameters using the expectation-maximisation (EM) algorithm [4]; in practice, it is observed that good approximations can be achieved with as few as five to ten Gaussians in the mixture. As the α parameter diminishes, and the extent of the tails increases, the required number of Gaussians for an accurate approximation grows.

3. The Locally Optimum Detector

3.1. The Discrete-time Problem

We consider the detection problem involving the choice between the hypotheses:

$$\begin{aligned} H_0 &: \mathbf{x} = \mathbf{w} \\ H_1 &: \mathbf{x} = e^{j\phi} \theta \mathbf{s} + \mathbf{w}. \end{aligned} \quad (4)$$

$\mathbf{x} = (x_1, x_2, \dots, x_n)$ is a vector of n observations, \mathbf{s} is the signal and $\mathbf{w} = (w_1, w_2, \dots, w_n)$ the noise; θ is a positive parameter affecting the amplitude of the received signal, and ϕ a parameter determining a random phase shift.

We assume the following signal and noise properties:

- The noise components w_k are i.i.d. complex random variables with a known common pdf $f_{IQ}(w_{I,k}, w_{Q,k})$, where $w_{I,k}$ and $w_{Q,k}$ are the in-phase and quadrature components. The function f_{IQ} is *circularly symmetric* so that $f_{IQ}(u, v) = h(\sqrt{u^2 + v^2}) = h(r)$.
- The signal samples are complex with a joint pdf $f_s(s_I, s_Q)$, where s_I and s_Q are the in-phase and quadrature components; the signal is zero-mean with a known correlation matrix \mathbf{R}_s .

The amplitude of the received signal is frequently small in comparison to the noise (θ is small); we are then interested in developing a *locally optimum* (LO) detector, which is most powerful for the θ range of interest. Assuming a uniformly-distributed $p(\phi)$, and denoting the amplitude of the k -th sample of the received vector R_k , we can express

the LO detection statistic as [2]:

$$\begin{aligned} \Lambda_{LO}(\mathbf{x}) = & \sum_{k=1}^n \sum_{m=1}^n \frac{h'(R_k)h'(R_m)}{h(R_k)h(R_m)R_kR_m} r_s(m, k) x_k x_m^* \\ & + \sum_{k=1}^n r_s(k, k) \left[\frac{h''(R_k)}{h(R_k)R_k^2} - \left(\frac{h'(R_k)}{h(R_k)R_k} \right)^2 \right] x_k x_k^* \end{aligned}$$

where $r_s(m, k)$ denote elements of the correlation matrix \mathbf{R}_s . Using the notation $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_k x_k y_k^*$:

$$\Lambda_{LO}(\mathbf{x}) = \langle \mathbf{R}_s \mathbf{v}(\mathbf{x}), \mathbf{v}(\mathbf{x}) \rangle + \langle \mathbf{R}_q \mathbf{z}(\mathbf{x}), \mathbf{y} \rangle, \quad (5)$$

where $\mathbf{y} = [\frac{x_1}{R_1}, \dots, \frac{x_n}{R_n}]^T$, $\mathbf{v}(\mathbf{x}) = [g(x_1), \dots, g(x_n)]^T$ and $\mathbf{z}(\mathbf{x}) = [\varepsilon(x_1), \dots, \varepsilon(x_n)]^T$ for:

$$\begin{aligned} g(x_k) &= w(R_k) y_k = \frac{-h'(R_k)}{h(R_k)} \frac{x_k}{R_k} \\ \varepsilon(x_k) &= w'(R_k) y_k \end{aligned} \quad (6)$$

\mathbf{R}_q is the correlation matrix determined by isolating the main diagonal of \mathbf{R}_s .

When we use a finite mixture of Gaussians to approximate the function $h(r)$ describing the noise, the non-linearity $w(r)$ which determines both \mathbf{z} and \mathbf{v} is:

$$w(r) = \frac{-h'(r)}{h(r)} = \frac{\sum_{i=1}^M P(i) \frac{r}{\sigma_i^2} \exp\left(\frac{-r^2}{2\sigma_i^2}\right)}{\sum_{i=1}^M P(i) \frac{1}{\sigma_i} \exp\left(\frac{-r^2}{2\sigma_i^2}\right)} \quad (7)$$

where $P(i)$ and σ_i are the parameters of the mixture.

3.2. The Continuous-time Problem

We now consider the continuous-time equivalent of (4):

$$\begin{aligned} H_0 &: x(t) = w(t) \\ H_1 &: x(t) = e^{j\phi} \theta s(t) + w(t), \end{aligned} \quad (8)$$

where $t \in [0, T]$, the observation period, $w(t)$ is the noise, and $x(t)$ is the observed signal. The signal $s(t)$ is a zero-mean complex signal with (known) correlation function $r_s(t_1, t_2)$. The continuous time LO detection statistic is:

$$\Lambda_{LO}(x) = \langle \mathbf{R}_q z(x), y \rangle + \langle \mathbf{R}_s v(x), v(x) \rangle, \quad (9)$$

where \mathbf{R}_q and \mathbf{R}_s are now correlation operators, $\langle f, g \rangle = \int_t f(t) g^*(t) dt$, and the functions y , z and v are the continuous forms of the vectors in (5). The second-order approximation of the likelihood ratio for the problem (8) is related to the LO test statistic [2]:

$$L(x) = \frac{f_x(x|\theta)}{f_x(x|0)} \approx 1 + \frac{\theta^2}{2} \Lambda_{LO}(x). \quad (10)$$

4. Composite Hypotheses

We now extend the continuous-time detection problem to the case where the signal has been subjected to an unknown time-frequency or time-scale shift. The problem can be stated as a choice between the hypotheses:

$$\begin{aligned} H_0 &: x(t) = w(t) \\ H_1 &: x(t) = e^{j\phi} \theta_s^{(a,b)}(t) + w(t), \end{aligned}$$

where (a, b) denotes a time-shift of a and either a frequency-shift $b = \nu$ or a scale offset $b = c$. The time-frequency or time-scale shift affects only the correlation operators \mathbf{R}_s and \mathbf{R}_q . For a time-frequency shift,

$$\mathbf{R}_s^{(a,\nu)} = \mathbf{R}_s(t_1 - a, t_2 - a) e^{j2\pi\nu(t_1 - t_2)}, \quad (11)$$

and for a time-scale offset,

$$\mathbf{R}_s^{(a,c)} = \frac{1}{c} \mathbf{R}_s((t_1 - a)/c, (t_2 - a)/c). \quad (12)$$

For the fixed shift (a, b) :

$$\Lambda_{\text{LO}}^{(a,b)}(x) = \left\langle \mathbf{R}_q^{(a,b)} z(x), y \right\rangle + \left\langle \mathbf{R}_s^{(a,b)} v(x), v(x) \right\rangle.$$

If the shift parameters are assumed to be deterministic but unknown, we use the generalised likelihood ratio test approach as in [7] to derive the test statistic:

$$\Lambda_{\text{GLO}}(x) = \max_{(a,b)} \Lambda_{\text{LO}}^{(a,b)}(x). \quad (13)$$

If the parameters are assumed to be random with joint pdf $p(a, b)$, an appropriate test statistic is given by

$$T(x) = \int_a \int_b T^{(a,b)}(x) p(a, b) da db, \quad (14)$$

where $T^{(a,b)}(x)$ is the likelihood ratio test statistic for fixed a and b . Using (5) and (10), we have the test statistic:

$$\Lambda_{\text{MLO}}(x) = \int_a \int_b \Lambda_{\text{LO}}^{(a,b)}(x) p(a, b) da db, \quad (15)$$

which is of similar form to the locally optimum statistics derived for the Gaussian noise case in [7].

TFR/TSR implementation

The TFR/TSR implementation of the test statistics is based on the covariance properties of the Weyl symbol $L_{\mathbf{H}}$ of a positive definite linear operator \mathbf{H} , and the Weyl correspondence [1], which states that:

$$\langle \mathbf{H}x, x \rangle = \int_t \int_f L_{\mathbf{H}}(t, f) W_x(t, f) dt df, \quad (16)$$

where $W_x(t, f) = W_{xx}(t, f)$ is the Wigner distribution:

$$W_{xy}(t, f) = \int_{\tau} x(t + \tau/2) y^*(t - \tau/2) e^{-j2\pi f \tau} d\tau. \quad (17)$$

The Weyl symbol displays a time-frequency shift covariance, $L_{\mathbf{R}_s^{(a,\nu)}}(t, f) = L_{\mathbf{R}_s}(t - a, f - \nu)$, and a time-scale shift covariance, $L_{\mathbf{R}_s^{(a,c)}}(t, f) = L_{\mathbf{R}_s}(\frac{t-a}{c}, cf)$. We have $\langle \mathbf{R}_s^{(a,\nu)} x, x \rangle = P_x(a, \nu; \mathbf{R}_s)$, a Cohen's class TFR [1],

$$P_{xy}(a, \nu; \mathbf{R}_s) = \int_u \int_{\xi} W_{xy}(u, \xi) L_{\mathbf{R}_s}(u - a, \xi - \nu) du d\xi$$

and $\langle \mathbf{R}_s^{(a,c)} x, x \rangle = D_x(a, c; \mathbf{R}_s)$, an affine class TSR,

$$D_{xy}(a, c; \mathbf{R}_s) = \int_u \int_{\xi} W_{xy}(u, \xi) L_{\mathbf{R}_s}(\frac{u-a}{c}, c\xi) du d\xi$$

We implement the statistics $\Lambda_{\text{GLO}}(x)$ and $\Lambda_{\text{MLO}}(x)$ using:

$$\Lambda_{\text{LO}}^{(a,b)}(\mathbf{x}) = V_{z(x),y}(a, b; \mathbf{R}_q) + V_{v(x)}(a, b; \mathbf{R}_s)$$

where V is the TFR P or the TSR D depending on the nature of the shift parameters.

A TFR/TSR-based implementation offers a computational saving because of the fast algorithms for evaluating TFRs and TSRs. We consider the TFR case for an N -sample observation, with a possible time-shift of $\pm N/4$ samples and a possible frequency shift of ± 0.25 . If we construct an $N/2 \times N/2$ sample time-frequency lattice covering this range, we must consider $N^2/4$ possible $\mathbf{R}_s^{(a,b)}$. The direct implementation is $O(N^3)$, as it involves $O(N)$ multiplications for each $\mathbf{R}_s^{(a,b)}$. The TFR-based implementation involves the evaluation of a number of spectrograms to calculate $\Lambda^{(a,b)}$ over all relevant (a, b) ; the computational expense is $O(N^2 \log N)$.

5. Array detection

Rao and Jones investigate the use of TFRs and TSRs for implementing array-based receivers under the assumption of Gaussian noise in [6]; they propose structures for the non-coherent, partially-coherent, and coherent detection scenarios. The near-optimal receivers discussed in this paper are readily extended to generate receivers which operate on an array of sensors. In particular, we can derive the near optimal test statistic for the case of coherent receivers [6]:

$$\Lambda_{\text{GLO}}(\mathbf{x}) = \max_{(a,b)} \left[V_{v_M}(a, b; \mathbf{R}_s) + \sum_{i=1}^M V_{z(x_i), y_i}(a, b; \mathbf{R}_q) \right] \quad (18)$$

where $v_M = \sum_{i=1}^M v(x_i)$, a sum over the M available sensors; we assume a known angle of arrival and aligned x_i .

6. Performance

In this section, we illustrate the performance of the proposed detection statistics. We concentrate on the TFR receiver, and use the coherent array formulation of (18) for an array of 4 sensors. We contrast the mixture model time-frequency detector with the time-frequency LO (Gaussian) detector proposed in [6, 7], and a time-frequency implementation of the Cauchy receiver. The LO Gaussian test statistic (for white noise) has the form [6]:

$$\Lambda_G = \max_{(a,b)} [V_{x_M}(a,b; \mathbf{R}_s)]. \quad (19)$$

where $x_M = \sum_{i=1}^M x_i$. The Cauchy test statistic has the same form as the detector presented above, but the non-linearity $w(r)$ of (7) becomes:

$$w(r) = \frac{-h'(r)}{h(r)} = \frac{3r}{(\gamma^2 + r^2)^{3/2}} \quad (20)$$

which arises from a bivariate Cauchy noise model.

We suppose that the signal of interest is characterised by the random process:

$$s(t) = \sum_{k=1}^3 X_k e^{-\beta_k t^2}, \quad (21)$$

where the X_k are i.i.d. ($X_k \sim N(0,1)$) and β_k is a fixed complex number. The signal is thus a weighted sum of three chirps, with the weights determined by the random processes X_k . In our test, we introduce a random time-frequency shift and examine the two noise cases $\alpha = 1.5$ and $\alpha = 1.9$ with the dispersion γ set to 1. Figure 1 displays the detection probability relative to signal strength for a constant false alarm rate of 0.01, and Figure 2 displays the ROC curves for a constant signal strength of $\theta = 3$. The performance of the Gaussian LO detector is poor for the $\alpha = 1.5$ case, and it deteriorates as the noise becomes increasingly impulsive. The mixture model detector performs better than the Cauchy detector for the depicted noise cases; the performance discrepancy diminishes as α approaches 1.

7. Summary

In this paper, we developed near-optimal test statistics for the detection of arbitrary nonstationary second-order signals in impulsive noise, modelled using a bivariate, isotropic α -stable distribution. We considered the extension to the case when the signal to be detected is subjected to an unknown time-frequency or time-scale shift, and showed that approximations to locally-optimal test-statistics can be implemented using TFRs/TSRs. We demonstrated that the performance of the LO linear receiver is poor in even mildly impulsive noise; the alternative detection statistics proposed in this paper offer considerably enhanced performance.

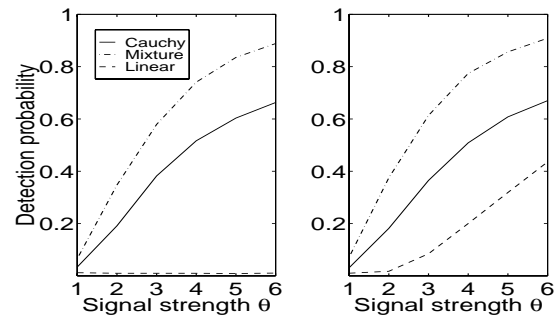


Figure 1. Detection probabilities for constant false-alarm rate = 0.01. The noise is isotropic α -stable with $\gamma = 1.0$; left $\alpha = 1.5$, right $\alpha = 1.9$.

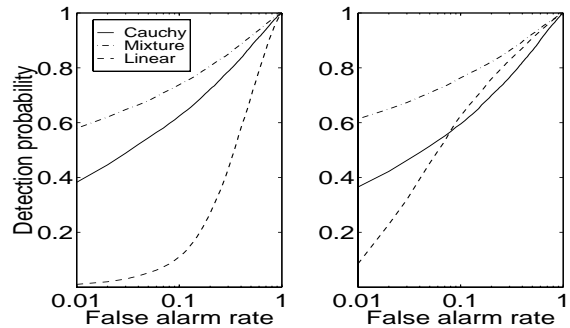


Figure 2. ROC curves for $\theta = 3$. The noise is isotropic α -stable with $\gamma = 1.0$; left $\alpha = 1.5$, right $\alpha = 1.9$.

References

- [1] L. Cohen. *Time-Frequency Analysis*. Prentice-Hall, Englewood Cliffs, N.J., 1995.
- [2] S. Kassam. *Signal Detection in Non-Gaussian Noise*. Springer-Verlag, New York, 1988.
- [3] E. Kuruoğlu, W. Fitzgerald, and P. Rayner. Near optimal detection of signals in impulsive noise modelled with a symmetric α -stable distribution. *IEEE Communications Letters*, 2(10):282–284, Oct. 1998.
- [4] E. Kuruoğlu, C. Molina, and W. Fitzgerald. Approximation of α -stable probability densities using finite mixtures of Gaussians. In *Proc. EUSIPCO'98*, pages 989–992, 1998.
- [5] C. Nikias and M. Shao. *Signal Processing with α -stable Distributions and Applications*. John Wiley & Sons, New York, 1995.
- [6] A. Rao and D. Jones. Nonstationary array signal detection using time-frequency and time-scale representations. In *Proc. 1998 IEEE Int. Conf. Acoust. Speech. Signal Proc.*, pages 1989–1992, 1998.
- [7] A. Sayeed and D. Jones. Optimal detection using bilinear time-frequency and time-scale representations. *IEEE Trans. Signal Processing*, 43:2872–2883, 1995.