SIGNAL DETECTION IN PARA COMPLEX NORMAL NOISE

Yonatan Woodbridge[(1)], *Gal Elidan*[(1),(2)] *and Ami Wiesel*[(1)]

(1) Hebrew University of Jerusalem (2) Google Inc.

ABSTRACT

In this paper we address target detection in correlated non-Gaussian noise. We introduce a powerful class of multivariate complex valued distribution that allows us to specify flexible non-Gaussian marginals, as well as correlation between the variables, while preserving circular symmetry. For noise belonging to this class, we study the fundamental problem of signal detection under different settings, and develop the needed (generalized) likelihood ratio tests. We also consider the problem of estimation of the noise parameters, and derive the maximum likelihood formulations. We compare the performance of the proposed methods using numerical simulations on synthetic data, and demonstrate the importance of using both correlations and non-Gaussiantiy.

***Index Terms*—** Circular symmetry, non-Gaussian, copula, detection.

1. INTRODUCTION

Detection of a signal corrupted by noise is a fundamental problem in statistical signal processing [1, 2]. The basic formulation assumes a known complex signal with a possibly unknown scaling factor, and with various models for the noise distribution. These range from expressive models with independent marginals [2, Chapter 10] to correlated models that are Gaussians [2, Chapter 4], as well as hybrids that only allow for limited control of the correlations [3, 4] (see details below). Our goal is to provide a unified and flexible framework for detection in non-Gaussian yet correlated noise. Towards this, we introduce a novel and powerful class of noise distributions, and derive the associated detectors and estimators.

As a branch of electrical engineering, statistical signal processing mostly deals with complex valued random variables which are useful in modeling electromagnetic waves. A common assumption is circular symmetry (CS), i.e. that the distribution is invariant to a constant phase rotation [5]. In the scalar case, this simply means that the amplitude and phase are statistically independent, and that the phase is uniformly distributed on the unit circle. The multivariate case is more subtle conceptually though some intuition can be obtained by examining the second order statistics, where CS is equivalent to properness, and dictates a zero mean and a specific covariance structure; the multivariate scenario is also more difficult in practice, from the constructive modeling perspective.

Consequently, the two common multivariate CS models are the independently and identically distributed non-Gaussian (IIDNG), which works with independent scalars, and the multivariate complex normal (CN), which is characterized by second order statistics. IIDNG models are limited in that they cannot capture dependencies between the variables. They can be colored using a linear transformation [6, 7], but this destroys their marginal properties. CN models, on the other hand, allow us to control correlation, but assume Gaussian marginals and thus cannot account for heavy tailed signals. CNs can be extended to complex elliptical distributions (CEDs) [3, 4], which allow for arbitrary marginals but limited correlation control. For example, CEDs do not include IIDNG as a special case. They are constructed by multiplying a CN vector with a positive random scalar that changes their marginals. Thus, even if the original CN has independent and identically distributed (IID) variables, these elements will be statistically dependent after the multiplication. Nonetheless, CEDs are commonly used in signal detection, e.g., [8, 9, 10].

The main contribution in this paper is a new class of multivariate distributions which allows for flexible parameterized marginals, control over correlations, and also guarantees multivariate CS. We denote this class as *Para Complex Normal* (PCN), and construct it as follows. We begin with a CN vector and apply an identical element-wise parameterized operation on the amplitudes of its elements. This transformation provides the three promised properties, and in fact is the only component-wise transformation that results in a multivariate CS distribution. PCN is directly related to the recent nonparanormal (NPN) proposal, where a similar transformation was non-parametric [11]. PCN is also related to copulas, and in particular to the class of Gaussian Copulas [12, 13, 14, 15, 16]. However, these frameworks deal with real-valued random variables with no notion of CS, and are thus designed for completely different applications.

Armed with the modeling power of PCN, we consider classical signal detection problems. We begin with detection of a known signal in PCN noise with known statistics, and derive the associated likelihood ratio test (LRT). Next, we extend the setting to the case of a known signal scaled by an unknown scalar, and derive the associated Generalized likelihood ratio test (GLRT). This test requires an expensive line search for estimating the scaling. Thus, we propose an alternative computationally efficient heuristic which shows reasonable performance. Finally, we address the issue of unknown PCN statistics. Following [17, 18], we assume a secondary dataset with i.i.d. target-less PCN samples, and derive the maximum likelihood (ML) estimator of the distribution parameters. The computational complexity of our numerical solution depends heavily on the number of parameters in the PCN transformation (typically only 2-3 parameters), but its dependence on the dimension of the vector is negligible. We demonstrate the performance of the different estimators and detectors via numerical simulations. In particular, we show the importance of both diverging from the Gaussian assumption, as well as allowing for correlations.

2. THE PCN DISTRIBUTION

In this section we introduce a class of complex-valued multivariate distributions which: 1) allows for flexible marginals; 2) has a correlation structure; 3) is CS. As discussed, existing frameworks satisfy only subsets of the these three properties, and our goal is specifically to fill this gap. We start by recalling the definition of a complex multivariate CS distribution [19]:

Definition 1. *Let $\mathbf{z} \in \mathbf{C}^p$ be a complex random vector of dimension p. We say that $\mathbf{z}$ is (multivariate) CS if $\mathbf{z}e^{j\phi}$ has the same distribution as $\mathbf{z}$ for any real ϕ.*

Probably the most common CS setting is the CN distribution denoted by $\mathcal{CN}(\mathbf{0}, \mathbf{Q})$, where the set of complex random variable is jointly Gaussian, and parameterized by a correlation matrix $\mathbf{Q}$ (note that a mean of zero follows from CS). Our proposal is to begin with this building block and modify its marginals. This leads to the following PCN definition:

Definition 2. *Let $g_{\boldsymbol{\theta}} : \mathbf{R}_+ \to \mathbf{R}_+$ be an invertible and differentiable function parametrized by the vector $\boldsymbol{\theta}$. Let $\mathbf{Q}$ be a positive definite correlation matrix with unit diagonal elements. Then, a complex random vector $\mathbf{w} = [w_1, \cdots, w_p]^T$ has a ParaComplexNormal (PCN) distribution denoted by*

$$\mathbf{w} \sim \mathcal{PCN}_g(\boldsymbol{\theta}, \mathbf{Q}).$$

if

$$\mathbf{z} = g_{\boldsymbol{\theta}}(|\mathbf{w}|)\, e^{j\angle \mathbf{w}} \sim \mathcal{CN}(\mathbf{0}, \mathbf{Q}), \tag{1}$$

where $g_{\boldsymbol{\theta}}(\cdot)$ is applied element-wise, i.e.,

$$z_k = g_{\boldsymbol{\theta}}(|w_k|)\, e^{j\angle w_k}, \quad k = 1, \cdots, p.$$

An obvious consequence of this definition is that it is straightforward to generate synthetic PCN samples. For this purpose, one simply generates a standard CN sample and applies $g_{\boldsymbol{\theta}}^{-1}(\cdot)$ to each marginal amplitudes. The result is clearly PCN.

As promised, the PCN distribution satisfies the required properties, as formalized in the following lemmas.

Lemma 1. *If $\mathbf{w}$ has a PCN distribution then it is multivariate CS.*

Proof. Since $\mathbf{z}$ is CS, $\mathbf{z}e^{j\phi} \sim \mathbf{z}$. Applying the inverse function $g_{\boldsymbol{\theta}}^{-1}$ to both sides yields $g_{\boldsymbol{\theta}}^{-1}(|\mathbf{z}|)\, e^{j(\angle \mathbf{z}+\phi)} \sim g_{\boldsymbol{\theta}}^{-1}(|\mathbf{z}|)\, e^{j\angle \mathbf{z}}$. For the left hand side, we have $g_{\boldsymbol{\theta}}^{-1}(|\mathbf{z}|)\, e^{j(\angle \mathbf{z}+\phi)} = \mathbf{w}e^{j(\angle \mathbf{z}+\phi)} = \mathbf{w}e^{j\phi}$. For the right hand side, we have $g_{\boldsymbol{\theta}}^{-1}(|\mathbf{z}|)\, e^{j\angle \mathbf{z}} = \mathbf{w}$ directly. Thus, $\mathbf{w} \sim \mathbf{w}e^{j\phi}$. □

Interestingly, under weak conditions (i.e., non-degeneracy), the converse to this lemma is also true: to ensure that PCN is CS using marginal transformations, we must use a transformation that modifies only the amplitude. We formally provide and prove the result in the journal version [20].

Lemma 2. *The PCN can be designed to have arbitrary (yet identical) CS marginal distributions. Specifically, to obtain CS marginals with magnitude CDF*

$$F_{\boldsymbol{\theta}}(\omega) = P_{\boldsymbol{\theta}}(|w_k| \leq \omega)$$

we choose

$$g_{\boldsymbol{\theta}}(|w_k|) = F_{Ray}^{-1} \circ F_{\boldsymbol{\theta}}(|w_k|)$$

where $F_{Ray}(x) = 1 - e^{-x^2}$ is the Rayleigh CDF

Proof. Having directly constructed the desired marginal CDF, we simply need to show that our choice of $g_{\boldsymbol{\theta}}()$ results in a CN distribution. CS marginals are characterized by uniform phase which is statistically independent of the amplitude, and $g_{\boldsymbol{\theta}}()$ only involves the magnitude. Thus, the phase and (transformed) magnitude remain independent, and the phase distribution is unchanged. $F_{\boldsymbol{\theta}}$ is always uniformely distributed (for any continuous random variable) and applying the inverse Rayleigh CDF to it results in a Rayleigh distributed magnitude. Together with the uniform and independent phase, the distribution is CN (see, for instance, [4]. There the squared magnitude is considered, which is chi-square distributed). □

The previous lemmas show that the PCN distribution is multivariate CS and allows for a flexible choice of the marginal magnitudes. The PCN class also allows flexible control on the correlation structure via the choice of $\mathbf{Q}$. For example, by choosing $\mathbf{Q} = \mathbf{I}$, we obtain an IIDNG distribution, and by choosing $\mathbf{Q} = \mathbf{1}\mathbf{1}^T$ we obtain a vector with identical elements. More details on the relation between the correlation matrix of a PCN vector and the matrix $\mathbf{Q}$ are provided in the journal version [20].

The density of $\mathbf{w} \sim PCN(0, \mathbf{Q})$ can also be derived, as a function of $g_{\boldsymbol{\theta}}()$, and will be useful in the following sections:

Lemma 3. *The PDF of PCN is*

$$f_{PCN}(\mathbf{w}) = \frac{1}{\pi^p |\mathbf{Q}|} e^{-\mathbf{z}(\mathbf{w})^H \mathbf{Q}^{-1} \mathbf{z}(\mathbf{w})} \prod_{k=1}^{p} \frac{\dot{g}_{\boldsymbol{\theta}}(|w_k|)\, g_{\boldsymbol{\theta}}(|w_k|)}{|w_k|} \tag{2}$$

where $\mathbf{z}(\mathbf{w}) = g_{\boldsymbol{\theta}}(|\mathbf{w}|)\, e^{j\angle \mathbf{w}}$, as defined in (1). The marginal pdfs are identical and equal to

$$f(w_k) = \frac{1}{\pi} e^{-|g_{\boldsymbol{\theta}}(w_k)|^2} \frac{\dot{g}_{\boldsymbol{\theta}}(|w_k|)\, g_{\boldsymbol{\theta}}(|w_k|)}{|w_k|}$$

and have a uniform phase independent of the magnitudes.

Proof. Due to space limitations, we only give a sketch of the proof technique and leave the details to [20]. We use a real-valued representation $\mathbf{w}' = \left[w_1^R, w_1^I, ..., w_p^R, w_p^I\right]^T$ and $\mathbf{z}' = \left[z_1^R, z_1^I, ..., z_p^R, z_p^I\right]^T$. Transformation from $\mathbf{w}'$ to $\mathbf{z}'$ is done by:

$$z_k^R = \frac{g_{\boldsymbol{\theta}}(|w_k|)}{|w_k|} w_k^R, \quad z_k^I = \frac{g_{\boldsymbol{\theta}}(|w_k|)}{|w_k|} w_k^I. \tag{3}$$

Recall that, from the transformation theorem for random variables,

$$f(\mathbf{w}') = f(\mathbf{z}'(\mathbf{w}'))\, |\mathbf{J}|, \quad \mathrm{J}_{k,\ell} = \frac{\partial z_k'}{\partial w_\ell'}, \quad k, \ell = 1, ..., 2p,$$

where $\mathbf{J}$ is the Jacobian matrix of the transformation. From the ordering of elements in $\mathbf{w}'$, $\mathbf{J}$ is block diagonal and each block is easy to compute (details omitted). Combining this with (3) yields the above PDF. □

Example: We conclude this section by illustrating PCN via a specific example: PCN with complex Generalized Gaussian marginals (GGPCN). These marginals are known to be useful in characterizing signals in many applications, e.g. [3, 4, 10, 21]. The PDF of the marginals is

$$f(w_k) = \frac{1}{2\pi} \frac{\beta}{\sigma^{\frac{4}{\beta}} \Gamma\left(\frac{2}{\beta}\right)} e^{-\frac{1}{\sigma^2}|w_k|^\beta},$$

and is parameterized by two parameters $\boldsymbol{\theta} = \left[\, \beta \;\, \sigma^2 \,\right]^T$, that relate to shape and scale, respectively. Two well known special cases are $\beta = 2$ which coincides with the Gaussian distribution, and $\beta = 1$ which coincides with Laplacian distribution. The distribution of the GG magnitudes $|w_k|$ is less known, but their powers $|w_k|^\beta$ are simply Gamma distributed. Their CDFs are goverened by the incomplete gamma function which is available in most scientific toolboxes.

Thus, the function $g_{\boldsymbol{\theta}}$, its inverse and its derivative can be easily implemented. For example, a MATLAB pseudocode is

$$\begin{aligned} g_{\boldsymbol{\theta}}(x) &= \sqrt{-\log\left(1-\text{gammainc}\left(\frac{x^{\beta}}{2\sigma^2},\frac{2}{\beta}\right)\right)} \\ g_{\boldsymbol{\theta}}^{-1}(x) &= \left(2\sigma^2\text{gammaincinv}\left(1-e^{-x^2},\frac{2}{\beta}\right)\right)^{\frac{1}{\beta}} \\ \dot{g}_{\boldsymbol{\theta}}(x) &= \frac{\beta x}{2\Gamma\left(\frac{2}{\beta}\right)\left(2\sigma^2\right)^{\frac{2}{\beta}}}e^{-\frac{x^{\beta}}{2\sigma^2}}\frac{1}{g_{\boldsymbol{\theta}}(x)}e^{g_{\boldsymbol{\theta}}^2(x)}. \end{aligned}$$

□

3. SIGNAL DETECTION

Having defined the multivariate CS flexible PCN distribution, we now consider the fundamental signal detection problems in the face of PCN noise. In this section we assume full knowledge of the PCN distribution, and consider the unknown noise parameters that requires estimation in the next.

3.1. Known Signal

We start with the simplest scenario where the observations are modeled as

$$\mathbf{y} = A_0\mathbf{s} + \mathbf{w}, \tag{4}$$

where A_0 is a real known deterministic parameter, $\mathbf{s}$ is a known steering vector of dimension p, and $\mathbf{w} \sim \mathcal{PCN}_g(\boldsymbol{\theta},\mathbf{Q})$ with known $g_{\boldsymbol{\theta}}()$, $\boldsymbol{\theta}$ and $\mathbf{Q}$. This formulation generalizes two classical textbook detection problems. When $g_{\boldsymbol{\theta}}(|w|) = |w|$, the model reduces to a signal corrupted by colored Gaussian noise, e.g., [2, chapter 4]. When $\mathbf{Q} = \mathbf{I}$, the model reduces to a signal corrupted by independent and identically distributed non-Gaussian noise, e.g., [2, chapter 10].

Our goal is to decide whether a signal is present or not. This leads to a simple hypothesis test:

$$\begin{aligned} \mathcal{H}_0: &\quad A = 0 \\ \mathcal{H}_1: &\quad A = A_0. \end{aligned}$$

We use the likelihood ratio test (LRT) to solve the decision problem. This test is optimal in a sense that it is the most powerful for a given significance level, according to Neyman-Pearson lemma [2, chapter 3]. The LRT is defined as

$$T(\mathbf{y}) = \frac{f(\mathbf{y};A_0,\mathcal{H}_1)}{f(\mathbf{y};0,\mathcal{H}_0)} = \frac{f_{\text{PCN}}(\mathbf{y}-A_0\mathbf{s})}{f_{\text{PCN}}(\mathbf{y})} > \tau, \tag{5}$$

where $f_{\text{PCN}}(\cdot)$ is defined in (2), and τ is chosen to fit a given false alarm probability

$$\text{Prob}\left(T(\mathbf{y}) > \tau|\mathcal{H}_0\right) = \alpha.$$

In this basic formulation, all the parameters are known and the distribution of the test is, in principle, straight forward. Nonetheless, computation of the threshold can be challenging since the distribution involves non-Gaussian random variables and a non-linear transformation, preventing a closed form solution.

We propose a numerical approach via simulations; generate K IID $\mathbf{w}$'s according to the PCN distribution, and define the empirical CDF of the test as $\hat{F}_T$. Then, we set $\tau = \hat{F}_T^{-1}(1-\alpha)$. We emphasize that this procedure, similarly to the computation of distribution tables, can be performed off-line before observing the vector $\mathbf{y}$.

3.2. Unknown Signal Amplitude

We now consider a more realistic setting in which the amplitude of the signal in unknown, e.g., as in [2, chapter 7]. We still assume full knowledge of the PCN noise distribution. Using the same model of (4), this formulation leads to a composite hypothesis test

$$\begin{aligned} \mathcal{H}_0: &\quad A = 0 \\ \mathcal{H}_1: &\quad A \neq 0. \end{aligned} \tag{6}$$

Since A is unknown, we will use a generalized likelihood ratio test (GLRT) which replaces the unknown A_0 in (5) with its ML estimate $\hat{A}$. This requires the solution of the following line-search

$$\hat{A}_{\text{ML}} = \arg\max_{A\neq 0} f_{\text{PCN}}\left(\mathbf{y}-\hat{A}\mathbf{s}\right).$$

In applications where this search is too computationally expensive, we propose to use a Gaussian approximation and estimate the amplitude using a simple weighted least squares (WLS) approach:

$$\hat{A}_{\text{WLS}} = \frac{\Re\left(\mathbf{s}^H\mathbf{Q}^{-1}\mathbf{y}\right)}{\mathbf{s}^H\mathbf{Q}^{-1}\mathbf{s}}.$$

This estimate is also asymptotically consistent, but not efficient, in the statistical estimation sense. Thus, it provides a reasonable trade off between accuracy and computational complexity.

As before, the threshold τ is chosen to fit a given false alarm probability. Importantly, even though A is missing, all the $\mathcal{H}_0$ parameters are known, and the distribution of the signal-free test can be computed via simulations. However, in contrast to the known amplitude case, this procedure cannot be performed off-line since the test itself depends on $\hat{A}_{\text{ML}}$ or $\hat{A}_{\text{WLS}}$, both of which are functions of $\mathbf{y}$. Thus, we can only perform these computations after the observations are available.

4. PARAMETER ESTIMATION

We now consider the scenario where the noise PCN distribution itself is unknown. In this case, it is common to first estimate the parameters of the noise distribution, and then perform detection assuming known (estimated) noise characteristics. To do so, we adopt a standard approach [17, 18] and assume that, in addition to the primary data $\mathbf{y}$ used for detection, we have a set of secondary data with only noise originating from the same noise distribution

$$\mathbf{w}_k \sim \mathcal{PCN}(\boldsymbol{\theta},\mathbf{Q}), \quad k = 1,\cdots N,$$

where $\mathbf{y}$ and $\mathbf{w}_k$ are all statistically independent of each other. We will use this secondary dataset to estimate $\{\boldsymbol{\theta},\mathbf{Q}\}$, which will then be used in the detection as if they were the true parameters of the PCN distribution. In what follows, we explain how to perform the estimation phase.

We propose to estimate $\{\boldsymbol{\theta},\mathbf{Q}\}$ given $\{\mathbf{w}_k\}_{k=1}^N$ using the maximum likelihood approach. The method is known to be asymptotically unbiased and efficient in the sense of minimizing the mean squared error among all unbiased estimators. The log-likelihood of the observations, denoted by $L\left(\{\mathbf{w}_k\}_{k=1}^N;\boldsymbol{\theta},\mathbf{Q}\right)$, is given by the logarithm of the PDF in (2). It can be shown that for any $\boldsymbol{\theta}$, the maximum likelihood with respect to $\mathbf{Q}$ is attained by the standard sample correlation matrix of the transformed data

$$\left[\hat{\mathbf{Q}}\left(\{\mathbf{w}_k\}_{k=1}^N;\boldsymbol{\theta}\right)\right]_{ij} = \frac{\mathbf{S}_{ij}(\boldsymbol{\theta})}{\sqrt{\mathbf{S}_{ii}(\boldsymbol{\theta})}\sqrt{\mathbf{S}_{jj}(\boldsymbol{\theta})}},$$

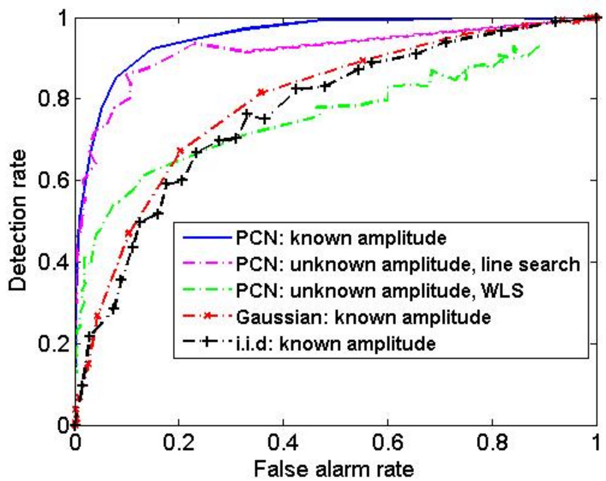

Fig. 1. Detection ROC curves with known PCN noise distribution

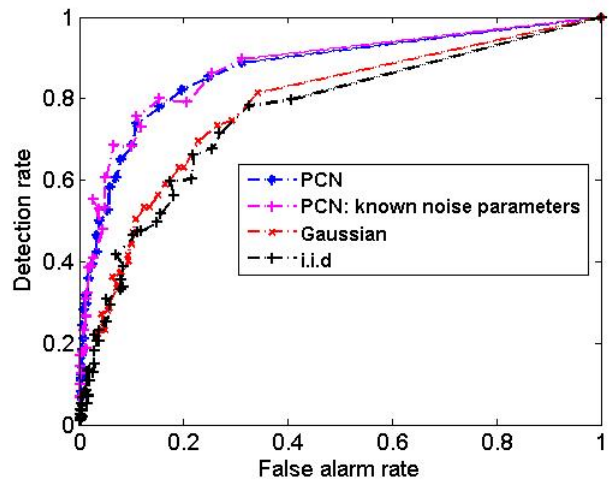

Fig. 2. Detection ROC curves for unknown PCN noise distribution.

where

$$\mathbf{S}(\boldsymbol{\theta}) = \frac{1}{N}\sum_{k=1}^{N} \mathbf{z}(\mathbf{w}_k)\,\mathbf{z}^H(\mathbf{w}_k).$$

Recall that $\mathbf{z}(\mathbf{w}_k) = g_{\boldsymbol{\theta}}(|\mathbf{w}_k|)\, e^{j\angle w_k}$. Plugging this back into the log-likelihood, we obtain an objective parameterized only by $\boldsymbol{\theta}$:

$$\widetilde{L}\left(\{\mathbf{w}_k\}_{k=1}^{N};\boldsymbol{\theta}\right) = L\left(\{\mathbf{w}_k\}_{k=1}^{N};\boldsymbol{\theta},\hat{\mathbf{Q}}\left(\{\mathbf{w}_k\}_{k=1}^{N};\boldsymbol{\theta}\right)\right).$$

We can now optimize over this low dimensional objective using a simple grid search. For example, with GG-PCN marginals, we only need to search over the values of two parameters, namely σ and β.

5. SIMULATIONS

In this section, we demonstrate the merit of the PCN noise representation using numerical simulations. More detailed and exhaustive experiments will appear in the journal version [20].

We start with signal detection in PCN noise, where the noise distribution is known. The signal itself is a vector of dimension $p = 20$, where all entries are identical and equal to $1 + j$, i.e $\mathbf{s} = (1 + j)\mathbf{1}$. The shape noise parameter is $\beta = 1$ (Laplace), while the scale parameter is $\sigma = 1$. The correlation matrix has random entries with strong correlation, i.e most of the off diagnoal entries have real and imagianry parts that are close to one. The true amplitude is $A_0 = 1$. Figure 1 shows the receiver operating characteristic (ROC) averaged over 50 independent trials for our PCN detection test under three different settings: PCN LRT with known amplitude; PCN GLRT with ML estimate of the amplitude; PCN GLRT with WLS estimate of the amplitude. We also compare to the two standard baselines: CN LRT with known amplitude (i.e., PCN with no transformation); standard IIDNG LRT with known amplitude (i.e. PCN with $\mathbf{Q} = \mathbf{I}$). The advantage of our LRT PCN detector when the signal is known is quite substantial even at false alarm rates that are far from zero. Further, our GLRT detector is also consistently better than the baselines, despite being at the disadvantage where the signal amplitude is unknown. Finally, although the performance of our WLS detector degrades without knowledge of the amplitude and with the Gaussian approximation, we are still on part with the baselines that have full knowledge of the signal.

Recall that to compute the threshold of our detector, we use a numerical method. Thus, to complement the ROC curves, we also examine the robustness of this computation. Using the same setting, we set a false alarm probability of $\alpha = 0.05$, and compute the threshold using $K = 100$ trials. We then test the accuracy of the threshold for noise variance values ranging from 0 to 15. As desired, the mean false alarm rate for each noise variance is extremely close to $\alpha = 0.05$, while the standard deviation is in the $[0.02, 0.03]$ range, even with the small number of computations.

We now consider detection in the more challenging scenario where the parameters of the PCN noise distribution are unknown. As discussed, in this scenario we first estimate the parameters from a secondary dataset with only noise, and then use these estimates for detection, as if they were the true parameters. Figure 2 shows the ROC curves for this scenario where the true noise distribution parameters are as before. The signal now is a vector with dimension of $p = 10$, and all entries are identical and equal to $1+j$. As before, our PCN approach is compared to the CN and IID models. 200 samples were use for the secondary dataset, whereas 100 samples were used for detection and the computation of the ROC curves. For reference, we also show the ROC curve for the case where the parameters of the noise distribution are known. Clearly, estimation noise does not noticeably influence the performance of our PCN detector, and the advantage over the baselines is evident throughout much of the false alarm rate range.

6. CONCLUSION

We presented PCN, a novel and powerful class of multivariate complex CS distributions, and developed detection tests in the face of noise belonging to this class. Importantly, the PCN noise distribution family allows for arbitrary CS marginals, as well as flexible control of the correlation. To the best of our knowledge, ours is the first work to consider detection with this combined flexibility in the complex CS setting that is central to signal processing. Complete proofs and additional evaluation are provided in the journal version [20].

7. REFERENCES

[1] H. V. Poor, *An introduction to signal detection and estimation*. Springer Science & Business Media, 2013.

[2] S. M. Kay, “Fundamentals of statistical signal processing, vol. ii: Detection theory,” *Signal Processing. Upper Saddle River, NJ: Prentice Hall*, 1998.

[3] E. Ollila and V. Koivunen, “Generalized complex elliptical distributions,” in *Sensor Array and Multichannel Signal Processing Workshop Proceedings, 2004*, pp. 460–464, IEEE, 2004.

[4] E. Ollila, J. Eriksson, and V. Koivunen, “Complex elliptically symmetric random variables generation, characterization, and circularity tests,” *Signal Processing, IEEE Transactions on*, vol. 59, no. 1, pp. 58–69, 2011.

[5] A. Lapidoth, *A foundation in digital communication*. Cambridge University Press, 2009.

[6] S. M. Kay, “Asymptotically optimal detection in incompletely characterized non-gaussian noise,” *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 37, no. 5, pp. 627–633, 1989.

[7] D. Sengupta and S. Kay, “Efficient estimation of parameters for non-gaussian autoregressive processes,” *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 37, no. 6, pp. 785–794, 1989.

[8] F. Gini and M. Greco, “Covariance matrix estimation for cfar detection in correlated heavy tailed clutter,” *Signal Processing*, vol. 82, no. 12, pp. 1847–1859, 2002.

[9] A. Wiesel, “Geodesic convexity and covariance estimation,” *Signal Processing, IEEE Transactions on*, vol. 60, no. 12, pp. 6182–6189, 2012.

[10] T. Zhang, A. Wiesel, and M. S. Greco, “Multivariate generalized gaussian distribution: Convexity and graphical models,” *Signal Processing, IEEE Transactions on*, vol. 61, no. 16, pp. 4141–4148, 2013.

[11] H. Liu, J. Lafferty, and L. Wasserman, “The nonparanormal: Semiparametric estimation of high dimensional undirected graphs,” *The Journal of Machine Learning Research*, vol. 10, pp. 2295–2328, 2009.

[12] A. Soriano, L. Vergara, J. Moragues, and R. Miralles, “Unknown signal detection by one-class detector based on gaussian copula,” *Signal Processing*, vol. 96, pp. 315–320, 2014.

[13] S. G. Iyengar, P. K. Varshney, and T. Damarla, “A parametric copula-based framework for hypothesis testing using heterogeneous data,” *Signal Processing, IEEE Transactions on*, vol. 59, no. 5, pp. 2308–2319, 2011.

[14] A. Sundaresan, P. K. Varshney, and N. S. Rao, “Copula-based fusion of correlated decisions,” *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 47, no. 1, pp. 454–471, 2011.

[15] H. Yu, J. Dauwels, and X. Wang, “Copula gaussian graphical models with hidden variables,” in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*, pp. 2177–2180, IEEE, 2012.

[16] A. Subramanian, A. Sundaresan, and P. Varshney, “Detection of dependent heavy-tailed signals,” 2015.

[17] E. J. Kelly, “An adaptive detection algorithm,” *Aerospace and Electronic Systems, IEEE Transactions on*, no. 2, pp. 115–127, 1986.

[18] I. S. Reed, J. D. Mallett, and L. E. Brennan, “Rapid convergence rate in adaptive arrays,” *Aerospace and Electronic Systems, IEEE Transactions on*, no. 6, pp. 853–863, 1974.

[19] R. G. Gallager, “Circularly-symmetric gaussian random vectors,” *preprint*, pp. 1–9, 2008.

[20] Y. Woodbridge, G. Elidan, and A. Wiesel, “Para complex normal distributions,” *in preparation*, 2015.

[21] M. Novey, T. Adali, and A. Roy, “A complex generalized gaussian distributioncharacterization, generation, and estimation,” *Signal Processing, IEEE Transactions on*, vol. 58, no. 3, pp. 1427–1433, 2010.