

Leveraging a Probabilistic PCA model to Understand the Multivariate Statistical Network Monitoring Framework for Network Security Anomaly Detection

Fernando Pérez-Bueno, Luz García, Gabriel Maciá-Fernández, Rafael Molina, *Senior Member, IEEE*



Abstract—Network anomaly detection is a very relevant research area nowadays, especially due to its multiple applications in the field of network security. The boost of new models based on variational autoencoders and generative adversarial networks has motivated a reevaluation of traditional techniques for anomaly detection. It is, however, essential to be able to understand these new models from the perspective of the experience attained from years of evaluating network security data for anomaly detection. In this paper, we revisit anomaly detection techniques based on PCA from a probabilistic generative model point of view, and contribute a mathematical model that relates them. Specifically, we start with the probabilistic PCA model and explain its connection to the Multivariate Statistical Network Monitoring (MSNM) framework. MSNM was recently successfully proposed as a means of incorporating industrial process anomaly detection experience into the field of networking. We have evaluated the mathematical model using two different datasets. The first, a synthetic dataset created to better understand the analysis proposed, and the second, UGR'16, is a specifically designed real-traffic dataset for network security anomaly detection. We have drawn conclusions that we consider to be useful when applying generative models to network security detection.

Index Terms—Anomaly Detection; PPCA; Generative Models; Network Security.

1 INTRODUCTION

Network security constitutes mainly a research and development focus nowadays, with a forecasted market of \$170.4 billion in 2022 according to Gartner [1], and a constant flow of worrying news concerning security incidents, like data breaches, denial of service, data exfiltration, privacy issues, advanced persistent threats, or even government cyberwar issues [2].

Many security technologies have been developed in recent years to deal with these relevant problems. The in-depth security design paradigm advocates the use of different layers of security protection to deal with such

problems. Among these layers, the detection of network security incidents is a crucial part of developing effective response measures when attacks occur.

There are two main approaches to the detection of network security incidents. One, *signature detection* works with rules defined by experts that identify known attacks. This approach works efficiently on known attacks but lacks the flexibility to detect unseen attack patterns or multi-stage attacks. An alternative approach, known as *anomaly detection*, based on building a normality model and detecting deviations from it, has gained popularity. However, the main limitation of anomaly detection technologies is related to the appearance of false positives or negatives, due to the lack of accuracy in the normality models obtained.

In this context, many different approaches to anomaly detection have been proposed in the literature [3]. Recent research in the deep learning area is generating high expectations with regard to the promising results that could be obtained in the field of network anomaly detection [4] [5] [6]. Preliminary results on the use of generative models, like variational autoencoders or generative adversarial networks, show high performance. There is, however, a concern that the internals of these models are not completely understood and that they might only show good results for specific datasets. In summary, there is a need to comprehensively understand the evolution from widely used models to these new deep generative models.

A family of models formed by those based on PCA are currently being successfully applied in the network anomaly detection problem. The first proposal to use PCA was that of Lakhina et al. [7] [8] in 2004. A training data matrix was built with different features extracted from network flows. The latent PCA model was then extracted from it and subsequently used to evaluate new network samples to decide on their abnormality.

Later in 2016, Camacho et al. [9] proposed the use of a framework based on PCA called MSNM (*Multivariate Statistical Network Monitoring*). MSNM essentially adapts a methodology known as MSPC (*Multivariate Statistical Process Control*), which has been extensively and successfully used in the field of anomaly detection in industrial process

- Fernando Pérez-Bueno (corresponding author) and Rafael Molina are with Dpto. de Ciencias de la Computación e I. A., Universidad de Granada, Spain. E-mail: fpb@ugr.es; rms@decsai.ugr.es
- Luz García and Gabriel Maciá-Fernández are with Dpto. de Teoría de la Señal, Telemática y Comunicaciones, Universidad de Granada, Spain. E-mail: luzgm@ugr.es; gmacia@ugr.es

control to deal with the specificities of the network anomaly detection problem. In brief, the framework suggests the construction of a PCA model¹ for normal traffic patterns, and the use of two statistics, termed D and Q , which are thresholded in order to determine if an anomaly occurs. While the statistic Q is similar to the one used by Lakhina, the use of D , a Hotelling’s T2 statistic [10], is a novelty that also captures possible deviations from the latent model, rather than only in the observation model (Q statistic).

Despite the effectiveness of MSNM in anomaly detection (see the comparative analysis in [11]), two problems appear in the application of this model. First, there is a lack of understanding on how D and Q statistics behave in different scenarios and why. This implies that obtaining good results in certain scenarios could lead to false conclusions being drawn. Second, there is a need to relate these PCA based anomaly detection models to more novel approaches like generative models, e.g., VAEs, which are currently being used profusely in the anomaly detection area. Understanding this relationship allows well-known lessons from PCAs to be applied to generative models.

A better way to understand these models from a generative perspective was provided by Tipping and Bishop [12], in which the PCA model is derived from a probabilistic PCA (PPCA) model when the variance of the latent distribution becomes zero. In addition, PPCA is at the basis of the definition of VAEs [13].

Starting from the PPCA model (a generative model that explains PCA), the focus of this paper is to derive an analytical model that elucidates why the use of the two MSNM statistics, D and Q , leads to effective anomaly detection and how this is understood in the framework of generative models. Note that the aim of this paper is not to contribute with a novel approach derived from MSNM, but rather to relate this model with generative models (analytically and empirically).

As it will be shown, while the Q statistic measures the quality of the reconstruction model, the D statistic will represent a regularization term in the generative model. As previously discussed, these conclusions should help to understand how to soundly employ generative models, like VAEs and GANs, for network anomaly detection. As an example, some of the current proposals for anomaly detection using VAEs, e.g., [14], use only the reconstruction model, thus discarding the contributions of the regularization term that could be relevant, as we will show in what follows.

In summary, the contribution of this paper is three-fold: *i*) A PPCA model is leveraged to understand the MSNM framework from a generative point of view. We then develop a mathematical framework to explain, from a probabilistic point of view, the meaning of the Q and D statistics. *ii*) Using the generative model, we show how some limitations of the MSNM model are circumvented. Specifically, the authors of [15] propose the use of a combined weighted statistic for both Q and D (called the t-Score). Then, we obtain a probabilistic interpretation for this weighted combination with PPCA-MSNM. *iii*) We test the generative model on both a synthetic dataset and a real

network traffic dataset to show how detection results stay coherent with both models.

The rest of this paper is structured as follows. In Section 2, we present related work and explain the contributions of this paper in more detail. Then, we provide the basics for MSNM and PPCA in Sections 3 and 4, respectively. The mathematical model that connects MSNM and PPCA is proposed in Section 5. This proposal is validated with the experiments shown in Section 6. Finally, conclusions are drawn in Section 7.

2 RELATED WORK

Many statistical strategies for anomaly detection on the basis of Lakhina’s approximation [7] [8] have been profusely proposed. The benefits of PCA’s unsupervised nature have motivated the appearance of a wide range of work, like the PCA-based traffic matrix estimation of [16], the network anomography proposed by [17], or the combination of distributed tracking and in-network PCA-based anomaly detection of [18] among many others. Limitations of these models, like the high sensitivity to calibration settings, ineffective detection of large anomalies or difficulties to capture temporal correlations have been reported [19]. In addition, the proposals to solve these limitations also use different frameworks. Robust PCA [20], and its variation [21] [22], or the Karhunen-Loève expansion used by [23], are examples of the achieved progress.

In 2016, Camacho *et al.* proposed the use of the Multivariate Statistical Network Monitoring (MSNM) framework [9] [15] as an improvement to previous PCA proposals. In essence, MSNM is an adaptation from a sibling framework traditionally used in the field of industrial process control, known as MSPC (Multivariate Statistical Process Control) [24] [25] [26]. In order to face the particularities of the networking field, MSNM adapted the MSPC methodology to introduce new data pre-processing strategies and processing steps, like the deparsing of network traces [11]. In addition, MSNM research has focused on the evaluation of its implementation in real networks, the optimization of its parameters with semi-supervised models, enabling big-data processing, its application to hierarchical architectures for issuing privacy and traffic reduction [9], enhancing visualization of network anomalies or supporting authentication systems [27].

The use of deep generative models in the field of anomaly detection, is currently a hot research topic due to good performance achieved by the use of deep learning techniques. Many authors have followed this approach using variational autoencoders (VAEs) as a natural evolution of PCA in the frame of reconstruction approaches to detect anomalies. Despite the fact that image and video processing research was the initial promoter of these models [6], they are also being extensively evaluated in the field of network anomaly detection. A recent survey of applications and techniques can be found in [4], yet, it is surprising to see that in many of these studies, like [14], where they consider the time gradient effect, or in the conditional VAE implemented in [28], or in others like [29] or [30], anomaly scores are evaluated by only considering the reconstruction error provided by the VAE. That is, the regularization term present in the

1. The use of linear PCA techniques is justified by the advantages in the diagnosing of anomalies in posterior phases of an incident lifecycle.

VAE marginal likelihood is not taken into account. Very few proposals, like [31], use both the regularization and the error reconstruction terms of the marginal likelihood to evaluate anomaly scores.

As will be shown in the rest of this paper, not using the regularization term in VAEs is similar to the well-known limitation in Lakhina's PCA approaches that do not consider the latent variable space deviations for anomaly detection. Thus, although evaluating the impact of the use of both terms (regularization and reconstruction) in VAEs is out of the scope of this paper, the goal here is to show the connection between these terms in a PPCA generative model (precursor of VAE) and MSNM detection statistics.

3 MSNM FOR ANOMALY DETECTION

Multivariate Statistical Network Monitoring (MSNM) [9] transfers the theory of Statistical Processes Control, which has been used for a long time in industrial applications, to network traffic analysis. Its goal is to jointly analyze several interrelated variables to differentiate common from special causes of variation called anomalies. The approach consists of five steps [11], which are explained in what follows.

First, raw network traffic data from different sources are parsed and transformed into a set of quantitative features, often using the feature-as-a-counter approach. Examples of such features are the number of times an event takes place, the count of a given word in a log, the number of times a given event takes place in a given time window, or the number of traffic flows with a given destination port in a NetFlow. The selection of the specific features and their parsing step require an effective comprehension of the data.

Second, all features are fused for the multivariate analysis, maintaining a common sampling rate. As a result, traffic flow matrices of N observations featured in M -dimensional vectors $N \times M$ are ready to be analyzed.

The *third* and main step of MSNM is the anomaly detection, which is based on PCA. The well-known technique is applied to the mean-centered and auto-scaled M -dimensional dataset of N observations ($N \times M$), and projected into a subspace of range $P < M$ that maximizes the variance. To do so, original features are transformed into Principal Components (PC) using the eigenvectors of the covariance matrix $\mathbf{X}^T \mathbf{X} / N$. As a result, a residual matrix \mathbf{E} is generated as the differential error between projections and real samples. The transformation follows Eq. (1), with \mathbf{T} ($N \times P$) and \mathbf{V} ($M \times P$) being the score and loading matrices, respectively :

$$\mathbf{X} = \mathbf{T} \cdot \mathbf{V}^T + \mathbf{E} \quad (1)$$

Based on such transformation, MSNM proposes the usage of two complementary statistics extracted from the PCA analysis:

- The Q -statistic, also called *Squared Prediction Error*, comprises the residuals in the n -th row of \mathbf{E} for a given observation x_n , following expression (2). As mentioned in Section 1, Q evaluates the reconstruction error of the projection used:

$$Q_n = e_n e_n^t \quad (2)$$

- The D -statistic, or *Hotelling's T2 statistic*, is computed by applying Eq. (3) to PCA scores. As mentioned in Section I, D represents a regularization term that rates how close the observation is to the data prior distribution. For a given observation x_n , and as t_n is the score vector in the n -th row of \mathbf{T} :

$$D_n = t_n \Lambda^{-1} t_n^t \quad (3)$$

Intuitively, Q measures the capability of the model to recover a certain point in the data, while the regularization term D measures the similarity of the latent representation with respect to those in the calibration data. As each term focuses on different domains, they are able to capture different types of anomalies.

It is well known that the Q -statistic has a high anomaly detection capability, and its usage together with the D -statistic is a key feature of the MSNM approach that offers attractive improvements [9]. Once D and Q are calculated, they are used to model the normal operating conditions for the calibration of the MSNM system. In order to do so, *upper control limits (UCL)* for a given significance level are defined for both Q , termed UCL_Q , and D , termed UCL_D . Diverse combinations of the two statistics can be used to provide a final expression for the anomaly evaluation. The authors of [15] propose the following weighting of Q and D to generate the anomaly score for a given observation x_n :

$$t - Score_n = \frac{P \cdot D_n}{M \cdot UCL_D} + \frac{(M - P) \cdot Q_n}{M \cdot UCL_Q} \quad (4)$$

Once the anomaly is detected, the *fourth* step of MSNM approximation is the pre-diagnosis. Features associated with the anomalies are identified in order to make an initial guess on their root causes. Contribution plots or others tools like oMEDA [9] are commonly used to identify such features.

Finally, the *fifth* step consists of de-parsing the information pointed out during the detection (anomaly time-stamps) and pre-diagnosis (anomaly related features) phases. As a final result, raw information about the anomaly is extracted from specific logs or network traces.

In what follows, we will introduce PPCA as a generative model framework (described in Section 4) in order to derive the expressions for Q and D statistics, as well as the interpretation of Eq. (4) in Section 5.

4 PPCA FOR ANOMALY DETECTION

Probabilistic PCA (PPCA) provides a method to calculate the principal subspace of a set of data vectors using a generative point of view and a maximum-likelihood framework. In order to be able to understand MSNM from a generative model perspective, it is important to first understand how PPCA is related to PCA.

This section provides a description of this connection, following the presentation of probabilistic PCA provided in [32]. While PCA is based on a deterministic linear projection of the data on a lower dimensional subspace, PPCA is a linear-Gaussian framework that considers a latent distribution for the data. Therefore a whole distribution of possible latent candidates is available for each observed data point. PPCA includes the measurement of deviations in the latent space, achieved in MSNM with the addition of the D

statistic, among other advantages (see [12] for a complete list).

Let \mathbf{X} denote the $N \times M$ network traffic matrix whose i -th row, \mathbf{x}_i^T , corresponds to the i -th observed instance, $i = 1, \dots, N$. That is, $\mathbf{X}^T = [\mathbf{x}_1, \dots, \mathbf{x}_N]$. We assume that each column of \mathbf{X}^T has been centered and normalized by its standard deviation. We also assume that calibrated observations are used. For each instance (network traffic observation) an explicit latent variable \mathbf{z} with P components is introduced. As we will see, it corresponds to components in a principal-component subspace. Next, a Gaussian prior distribution $p(\mathbf{z})$ over the latent variable and a Gaussian observation distribution $p(\mathbf{x}|\mathbf{z})$ are introduced. Specifically:

$$p(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (5)$$

$$p(\mathbf{x}|\mathbf{z}) = \mathcal{N}(\mathbf{x}|\mathbf{W}\mathbf{z}, \sigma^2\mathbf{I}) \quad (6)$$

where \mathbf{W} is a $M \times P$ matrix whose columns span a linear (the principal-component) subspace within the data space, and the scale σ^2 governs the variance of the conditional distribution. Notice that, in what follows, we will omit the dependence of $p(\mathbf{x}|\mathbf{z})$ on \mathbf{W} and σ^2 for simplicity.

To estimate the values of the parameters \mathbf{W} and σ^2 , we use maximum likelihood, and the marginal distribution $p(\mathbf{x})$ is required. It can be easily calculated because the prior and the observation models in Eqs. (5) and (6) are both Gaussian. It follows that:

$$p(\mathbf{x}) = \int p(\mathbf{x}|\mathbf{z})p(\mathbf{z})d\mathbf{z} = \mathcal{N}(\mathbf{x}|\mathbf{0}, \mathbf{C}), \quad (7)$$

where

$$\mathbf{C} = \mathbf{W}\mathbf{W}^T + \sigma^2\mathbf{I}. \quad (8)$$

The above likelihood requires the calculation of \mathbf{C} which may consume a lot of computational resources. This can be alleviated when N is larger than P (the dimension of the principal component subspace) and by utilizing the matrix inversion identity

$$\mathbf{C}^{-1} = \sigma^{-2}(\mathbf{I} - \mathbf{W}\mathbf{M}^{-1}\mathbf{W}^T), \quad (9)$$

where

$$\mathbf{M} = \mathbf{W}^T\mathbf{W} + \sigma^2\mathbf{I}. \quad (10)$$

Using our normalized and calibrated observations, the maximum likelihood estimates are calculated by solving

$$\begin{aligned} \mathbf{W}_{\text{ML}}, \sigma_{\text{ML}}^2 &= \arg \max_{\mathbf{W}, \sigma^2} \sum_{n=1}^N \ln p(\mathbf{x}_n|\mathbf{W}, \sigma^2) \\ &= \arg \max_{\mathbf{W}, \sigma^2} -\frac{NP}{2} \ln(2\pi) - \frac{N}{2} \ln |\mathbf{C}| - \frac{1}{2} \sum_{n=1}^N \mathbf{x}_n^T \mathbf{C}^{-1} \mathbf{x}_n \\ &= \arg \max_{\mathbf{W}, \sigma^2} -\frac{N}{2} \{P \ln(2\pi) + \ln |\mathbf{C}| + \text{tr}(\mathbf{C}^{-1}\mathbf{S})\} \end{aligned} \quad (11)$$

where \mathbf{S} is the data normalized and calibrated covariance matrix.

Maximizing the above equation with respect to \mathbf{W} and σ^2 is not easy. However, it can be shown —see [32]— that

$$\mathbf{W}_{\text{ML}} = \mathbf{U}(\mathbf{L} - \sigma^2\mathbf{I})^{1/2}\mathbf{R} \quad (12)$$

where \mathbf{U} is a $M \times P$ matrix whose columns correspond to the P eigenvectors associated with the P largest eigenvalues of the data normalized and calibrated covariance matrix

\mathbf{S} , $\lambda_1, \dots, \lambda_P$. \mathbf{L} is a diagonal matrix with diagonal values these eigenvalues, and \mathbf{R} is an $P \times P$ orthonormal matrix that represents any rotation. Note that σ^2 is constrained to be smaller than the lowest eigenvalue λ_P (the minimum element in the diagonal matrix \mathbf{L} , thus avoiding a negative square root in Eq. (12)). Thus, $\sigma^2 \in [0, \lambda_P)$. Furthermore, the maximum likelihood for σ^2 is

$$\sigma_{\text{ML}}^2 = \frac{1}{M-P} \sum_{i=P+1}^M \lambda_i \quad (13)$$

Note that with the estimated \mathbf{W}_{ML} and σ_{ML}^2 we can calculate, for a given normalized sample \mathbf{x} , the quantity

$$\ln p(\mathbf{x}|\mathbf{W}_{\text{ML}}, \sigma_{\text{ML}}^2) = -\frac{M}{2} \ln(2\pi) - \frac{1}{2} \ln |\mathbf{C}| - \frac{1}{2} \mathbf{x}^T \mathbf{C}^{-1} \mathbf{x} \quad (14)$$

and use it to decide whether \mathbf{x} is an anomaly (we have omitted the dependency of \mathbf{C} on \mathbf{W}_{ML} and σ_{ML}^2 for simplicity). The whole process for anomaly detection in PPCA is represented in algorithm 1. First, the data need to be mean-centered around zero and scaled as PCA and PPCA are sensitive to feature scaling. Then, the parameters of the model are estimated using Eq. (13) and Eq. (12) on trusted calibration data. Once the parameters are fixed, new data can be checked for anomalies.

Algorithm 1 PPCA for anomaly detection.

Require: Centered and normalized observations \mathbf{X} , predefined threshold thr .

Obtain σ_{ML}^2 using Eq. (13).

Obtain \mathbf{W}_{ML} using Eq. (12).

For a new sample \mathbf{x}^{new} , decide whether it is an anomaly by thresholding $\ln p(\mathbf{x}^{new}|\mathbf{W}_{\text{ML}}, \sigma_{\text{ML}}^2)$ defined in Eq. (14), that is if,

$$\frac{1}{2} \mathbf{x}^T \mathbf{C}^{-1} \mathbf{x} > thr. \quad (15)$$

Notice that Algorithm 1 can be used even with a value σ^2 that is different from σ_{ML}^2 . The only constraint on σ^2 is that it has to be smaller than λ_P , see Eq. (12).

Using $\frac{1}{2} \mathbf{x}^T \mathbf{C}^{-1} \mathbf{x} > thr$ to detect anomalies does not provide an insight into the roles played by the prior and the observation models in the detection process. This is crucial since it will allow us to relate PPCA and MSNM, as we will see in the following section.

5 RELATING PPCA TO MSNM

5.1 Revisiting PPCA for anomaly detection

In order to relate PPCA and MSNM we will evaluate other variance values apart from σ_{ML}^2 , in the expressions for PPCA. Thus, for clarity, we will make the formulations using a generic variance δ (smaller than λ_P). We will use the Laplace approximation [32] described in Appendix A as a starting point for the analysis. Given a marginal probability $p(\mathbf{x}|\mathbf{W}, \delta)$, with \mathbf{W} and δ being its score matrix and variance, and following Eq. (33), it follows that:

$$\ln p(\mathbf{x}|\mathbf{W}, \delta) = -\frac{1}{2} (\hat{\mathbf{z}}^T \hat{\mathbf{z}} + \frac{1}{\delta} \|\mathbf{x} - \mathbf{W}\hat{\mathbf{z}}\|^2) + \text{const.} \quad (16)$$

where $\hat{\mathbf{z}}$ is the mode of the posterior distribution $p(\mathbf{z}|\mathbf{x})$.

Therefore, leaving constants aside for the sake of clarity, we can equalise the thresholding expression of $\ln \mathbf{p}(x|\mathbf{W}, \delta)$ for PPCA —Eq. (14)— to its Laplace approximation:

$$\frac{1}{2} \mathbf{x}^T \mathbf{C}^{-1} \mathbf{x} = \frac{1}{2} (\hat{\mathbf{z}}^T \hat{\mathbf{z}} + \frac{1}{\delta} \| \mathbf{x} - \mathbf{W} \hat{\mathbf{z}} \|^2) \quad (17)$$

In order to calculate the term $\hat{\mathbf{z}}$ in Eq. (17), the closed expression for $\mathbf{p}(z|x)$ derived in [32] can be used:

$$\mathbf{p}(z|\mathbf{x}) = \mathcal{N}(z|\mathbf{M}^{-1} \mathbf{W}^T \mathbf{x}, \delta \mathbf{M}^{-1}) \quad (18)$$

and since for this Gaussian distribution the mode and the mean coincide, it follows that:

$$\hat{\mathbf{z}} = \mathbf{M}^{-1} \mathbf{W}^T \mathbf{x} \quad (19)$$

As seen in Eq. (12) (Section 4), \mathbf{W} depends on the variance $\delta \in [0, \lambda_P)$:

$$\mathbf{W}(\delta) = \mathbf{U}(\mathbf{L} - \delta \mathbf{I})^{1/2}, \quad (20)$$

where for the sake of simplicity, $\mathbf{R} = \mathbf{I}$ is used in (12). Introducing Eq. (20) in Eq. (10), and taking into account that $\mathbf{U}^T \mathbf{U} = \mathbf{I}$, it follows that:

$$\mathbf{M} = \mathbf{W}(\delta)^T \mathbf{W}(\delta) + \delta \mathbf{I} = (\mathbf{L} - \delta \mathbf{I}) + \delta \mathbf{I} = \mathbf{L} \quad (21)$$

Using Eqs. (20) and (21) in Eq. (19):

$$\hat{\mathbf{z}}(\delta) = \mathbf{M}^{-1} \mathbf{W}(\delta)^T \mathbf{x} = \mathbf{L}^{-1} (\mathbf{L} - \delta \mathbf{I})^{1/2} \mathbf{U}^T \mathbf{x}. \quad (22)$$

In short, the whole process for anomaly detection with this revisited PPCA is summarized in Algorithm 2. Note that it is equivalent to Algorithm 1, and allows arbitrary variance values to be used. The parameters of the model are estimated on trusted calibration data before checking new data for anomalies.

Algorithm 2 PPCA for anomaly detection (Revisited)

Require: Centered and normalized observations \mathbf{X} , predefined threshold thr .

Set a certain variance $\delta \in [0, \lambda_P)$.

Obtain $\mathbf{W}(\delta)$ using Eq. (20).

Obtain $\mathbf{M}(\delta)$ using Eq. (21).

Obtain $\hat{\mathbf{z}}(\delta)$ using Eq. (22).

For a new sample \mathbf{x}^{new} , decide whether it is an anomaly by thresholding $\ln \mathbf{p}(\mathbf{x}^{new}|\mathbf{W}(\delta), \delta)$ defined in Eq. (16), that is if,

$$\frac{1}{2} (\hat{\mathbf{z}}^T(\delta) \hat{\mathbf{z}}(\delta) + \frac{1}{\delta} \| \mathbf{x} - \mathbf{W}(\delta) \hat{\mathbf{z}}(\delta) \|^2) > thr \quad (23)$$

5.2 Analysis of the variance: connecting PPCA and MSNM

Recall that Algorithm 2 can be used with a value of $\delta \in [0, \lambda_P)$ which is different from σ_{ML}^2 . We can observe the connection between PPCA and MSNM when studying the range of values that this variance δ might take. To do so, let us analyze the expression of $\ln \mathbf{p}(x|W, \delta)$ in Eq. (16). Leaving aside constant terms, Eq. (16) is formed using two terms, the first of which considers the influence of the latent space in the probability calculation, and somehow acts as a *regularization term*:

$$\hat{\mathbf{z}}(\delta)^T \hat{\mathbf{z}}(\delta) \quad (24)$$

while the second evaluates the difference between a sample \mathbf{x} and its reconstruction from the latent space. Thus, it plays the role of a *reconstruction error*:

$$\| \mathbf{x} - \mathbf{W} \hat{\mathbf{z}}(\delta) \|^2 \quad (25)$$

Note that the reconstruction error contribution to the sample's probability is weighted by a factor $1/\delta$. At this point, we separate the influence of δ in this weighting factor from that of δ in $\mathbf{W}(\delta)$ and $\hat{\mathbf{z}}(\delta)$. Thus, we let the weighting factor take a fixed value $1/\alpha$ while δ keeps taking possible values in $[0, \lambda_P)$. Thus, a function $f_\alpha(\delta)$ can be defined to study the behavior of $\ln \mathbf{p}(x|W, \delta)$ for different values $\delta \in [0, \lambda_P)$:

$$f_\alpha(\delta) = \frac{1}{2} (\hat{\mathbf{z}}^T(\delta) \hat{\mathbf{z}}(\delta) + \frac{1}{\alpha} \| \mathbf{x} - \mathbf{W}(\delta) \hat{\mathbf{z}}(\delta) \|^2) \quad (26)$$

It can be seen that —see Appendix B and Figure 1— $f_\alpha(\delta)$ in Eq. (26) has 3 properties:

- (i) $f_\alpha(\delta)$ is convex,
- (ii) its minimum value is achieved at $\delta = \alpha/2$,
- (iii) $f_\alpha(0) = f_\alpha(\alpha)$.

Based on the fact that MSNM uses PCA for its modelling and PPCA converges to PCA when $\delta = 0$ — see [32]—, we have explored the value of $f_\alpha(\delta)$ for $\delta = 0$:

$$f_\alpha(0) = \frac{1}{2} (\mathbf{x}^T \mathbf{U} \mathbf{L}^{-1} \mathbf{U}^T \mathbf{x} + \frac{1}{\alpha} \| \mathbf{x} - \mathbf{U} \mathbf{U}^T \mathbf{x} \|^2) \quad (27)$$

Eq. (27) is exactly the quantity used by MSNM to detect anomalies since:

- $\mathbf{x}^T \mathbf{U} \mathbf{L}^{-1} \mathbf{U}^T \mathbf{x}$, the regularization part of the probability, matches Eq. (3) that defines the D -statistic.
- $\| \mathbf{x} - \mathbf{U} \mathbf{U}^T \mathbf{x} \|^2$, the reconstruction error part of the probability, matches Eq. (2) that defines the Q -statistic.
- Both terms, regularization and reconstruction error, contribute with a different weighting: *i)* $1/\sigma_{ML}^2$ in PPCA, and *ii)* the empirical value $(M - P/M)$ —see Eq. (4)— in MSNM.

The results obtained show that the conditions that make MSNM and PPCA coincide are:

$$\begin{aligned} \alpha &= \sigma_{ML}^2 \\ \delta &= \sigma_{ML}^2 \text{ or } \delta = 0 \end{aligned} \quad (28)$$

The behavior of the function $f_\alpha(\delta)$ is shown in Figure 1.

Notice also that in this *revisited version of PPCA* we could use a value of δ that is different from zero and σ_{ML}^2 . In that case, for a given anomaly probability threshold, the anomaly detection model would accept less observations (lower probability) if $\delta \in [0, \alpha]$ and more if $\alpha < \delta < \lambda_P$.

To end Section 5, note that the conclusions obtained for PPCA are applicable to other generative models based on PPCA, like Variational Autoencoders (VAEs). In VAEs, the calculation of $\mathbf{p}(z|\mathbf{x})$ in a closed form is not possible and, thus, we must rely on the so called Evidence Lower Bound (ELBO), which comprises two different terms (as in PPCA):

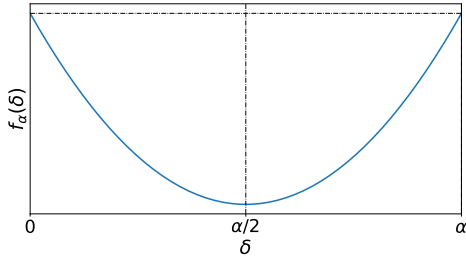


Figure 1. Analysis of $f_\alpha(\delta)$.

a reconstruction error and a regularization term. Our conclusions for PPCA lead to the same recommendations for VAEs, *i.e.*, both terms should be used in the calculation of anomaly scores. We consider this to be a relevant contribution of this paper, as we can still find examples in the state-of-the-art [14] [28] [29] [30] that only use the error reconstruction term for anomaly detection.

6 EXPERIMENTAL VALIDATION

6.1 Datasets Description

To experimentally validate the theory introduced in the previous sections, we will make use of two datasets: one with synthetic data, and another with real network traffic information, as described below.

6.1.1 Synthetic dataset

The synthetic dataset is intended to provide a 2-dimensional plottable scenario that provides an easy interpretation of the anomalies and behaviour of the different models. Using the PCA observation model given by $\mathbf{x} = \mathbf{W}\mathbf{z} + \epsilon$, we obtain N samples of \mathbf{z} from a Gaussian distribution with zero mean and unit variance. Then we multiply \mathbf{z} by an arbitrary $\mathbf{W} = [0.707, 0.707]^T$ and add ϵ sampled from $\mathcal{N}(\mathbf{0}, 0.1 \cdot \mathbf{I})$ to obtain N bidimensional samples of \mathbf{x} . The samples will follow a Gaussian distribution along the selected \mathbf{W} . Using this method, we generate 1000 samples of clean data to calibrate our models. For the testing set, we mix 1000 new clean data samples with two different types of anomalies. The first type of anomaly is sampled from a distribution that is different from that of the above described samples. Specifically, we choose a multivariate Gaussian with zero mean and a diagonal variance matrix $5 \times \mathbf{I}$. Anomaly type 1 is not designed according to the linear model. The second type of anomaly is generated in the latent space, we sample z_{anom} using a Gaussian with mean 5 and unit variance. The values of z_{anom} are then randomly multiplied by -1 following a Bernoulli (0.5) distribution. This procedure provides anomalies that follow the generative model but whose latent distribution is different from that of the calibration data. For each type of anomaly mentioned, 100 data points are introduced. See Figure 2 for a visual representation of the data and anomalies.

6.1.2 UGR'16 dataset

The UGR'16 dataset [33] was designed for the evaluation of cyclostationary-based network IDSs. It contains real anonymized NetFlow traces captured over several months

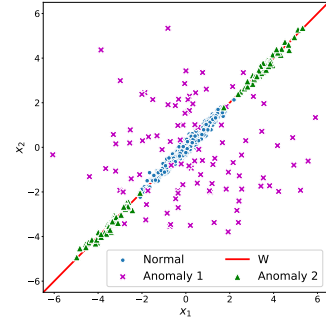


Figure 2. Synthetic dataset test bench. Normal data is presented in blue circles, Anomaly type 1 is presented in fuchsia with \times -shaped markers, while the green triangles refer to Anomaly type 2. The arbitrary line \mathbf{W} used to create the data is plotted in red.

in a tier-3 ISP. The traces include legitimate traffic from many virtualized services in the cloud, like web servers with proprietary and standard configurations, and other hosted services like DNS, FTP, mail servers, etc. In addition, a set of malicious virtual machines was configured in the network to generate attack traffic. The traffic is captured from two border routers of the ISP network. Thus, the dataset includes both legitimate traffic and realistic attack scenarios, all of them labeled.

UGR'16 divides the data in two differentiated sets: calibration and test. The calibration set contains real background traffic data gathered from March to June in 2016 (4 months). The test set mixes real background traffic and synthetically generated malicious traffic, gathered from July to August 2016. Although the data was captured in 2016, it was published in 2018 and is still considered of interest in recent work [34], [35].

To train our models, we use the data gathered during working days in May, where less anomalies were detected after data obtained was analyzed [33], and no synthetically generated attacks were introduced. The test set uses data gathered on those working days when synthetically generated attacks were interlaced within background traffic. The attack types that were synthetically generated include:

- Low-rate Denial of Service (DoS): TCP SYN packets are sent to the victims port 80 using 1280-bit packets and a rate of 100 packets/s. The rate is not high enough to avoid the normal operation of the network being affected.
- Port scanning: A continuous SYN scan of common ports of victims. Two variants are implemented for this attack: Scan11 (one-to-one scan attack) and Scan44 (Four-to-four scan attack).
- Botnet traffic: Obtained from the execution of the malware known as Neris in a controlled environment (See [36] for details about the malware and [33] for details about its injection in the data.).

The test set also included labels for a real UDP Scan campaign that was identified in the background traffic.

The labels are provided as a list of timestamps (in mins) of when the attacks were executed. Table 1 summarizes the traces in the train and test sets. The NetFlow logs cannot be directly used to feed PCA-based anomaly detection systems

Table 1

Summary of data flows from UGR'16 used to train and test the models

Flow type	Calibration	Test
Background traffic	31680	8714
DoS	NA	299
Scan44	NA	65
Scan11	NA	66
Nerisbotnet	NA	488
UDPscan	NA	9
SSHscan	NA	9
Spam	NA	3616

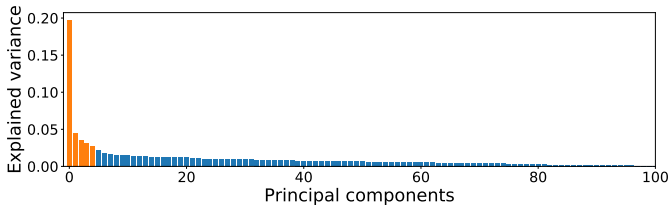


Figure 3. Explained variance per component in the UGR'16 dataset.

[8]. Thus, following [11], we use the FCParse² tool to extract 143 quantitative features from the NetFlow logs. Each feature counts the number of times that a given event takes place during each minute, *e.g.*, number of flows with a given destination port, number of flows with an accumulated payload size greater than a threshold, etc. Features were manually defined in [11] from domain knowledge using regular expressions. Therefore, they represent information that experts would use to manually identify anomalies in the data. In this paper, we focus on differentiating the minutes labeled as anomalies from those labeled as background data. We study each type of anomaly/attack separately. The anomaly scores are calculated for the whole test set at once. Then, classification metrics are calculated for each anomaly type against background data only (binary classification)

The evaluation of the variance captured by the different principal components of PCA on the calibration set (see Figure 3), shows that the principal components 1-5 are the most relevant and capture a higher percentage of variance than the rest. The model with 5 P explains 33.65% of the total variance.

6.2 Experiments

This section focuses on describing three different experiments: *i)* Anomaly detection using the MSNM and PPCA models with a common threshold to show that detection results from both models coincide. *ii)* Anomaly detection using only either the regularization or the reconstruction error term separately, to analyze their performance for different types of anomalies. *iii)* Finally, we assess the correctness of the weighting values for both regularization and reconstruction error terms, for MSNM and PPCA models.

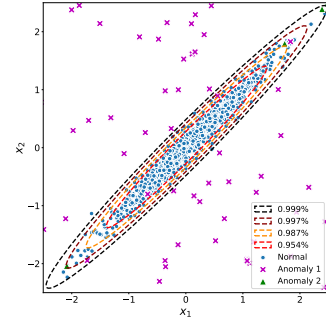
6.2.1 MSNM and PPCA equivalence

To experimentally demonstrate the MSNM and PPCA equivalency shown in Section 5, we perform anomaly detection on the datasets presented in Section 6.1 using both

Table 2

Performance metrics for the synthetic dataset using different threshold values selected according to different confidence levels on the calibration set.

Confidence level	95%	96%	97%	98%	99%
Threshold χ^2	5.99	6.43	7.01	7.82	9.21
Accuracy	0.9925	0.9933	0.9925	0.9866	0.9808
False Alarm Ratio	0.002	0.001	0.001	0.001	0

Figure 4. Detail of the synthetic test set shown in Figure 2. The detection area obtained with the threshold for different confidence levels α on the calibration data is shown in dashed lines.

MSNM and PPCA models. The anomaly score is calculated using Eq. (26) with the condition in Eq. (28), *i.e.*, $\alpha = \sigma_{ML}^2$ and $\delta = 0$ or $\delta = \sigma_{ML}^2$ for the MSNM and PPCA models, respectively. When a threshold value is needed, it should be chosen in accordance with the values in the calibration set, according to our confidence in the absence of malicious traffic in the training data. In our case, the threshold is the 99-percentile of the anomaly scores of the calibration set for each model. The threshold can be modified if the confidence in the calibration data is reduced.

Note also that once σ_{ML}^2 and \mathbf{W}_{ML} have been calculated, Eq. (14) can be used to set a non-experimental value. Since $p(x|\mathbf{W}_{ML}, \sigma_{ML}^2)$ is a normal distribution we can select a threshold δ such as

$$\int_{\mathbf{x}^T \mathbf{C}^{-1} \mathbf{x} \leq \delta} p(x|\mathbf{W}_{ML}, \sigma_{ML}^2) dx \geq \alpha \quad (29)$$

for a confidence level α . We can take into account that $\mathbf{x}^T \mathbf{C}^{-1} \mathbf{x} \sim \chi^2(M)$ and select for this distribution a threshold δ with a given confidence level α . This theoretical approach is less used in practice.

In Figure 4 we can see how the detection area changes according to the confidence level α . Table 2 includes the threshold value, accuracy and False Alarm Ratio using the theoretical bound $\mathbf{x}^T \mathbf{C}^{-1} \mathbf{x} \leq \delta$.

Figure 5 shows the parallelism between MSNM and PPCA models in the synthetic test set. The same data points are identified as anomalies by both models, obtaining identical ROC curves and AUC values. The accuracy (0.9858) and the False Alarm Ratio (0.012%) were also the same values for both models. The equivalence of both models is also tested on the UGR'16 dataset. Although the anomaly score is calculated for all testing data at once, we calculate the ROC curve and AUC for each attack type against the background traffic. Table 3 includes identical results for

2. Available at: <https://github.com/josecamachop/FCParser>

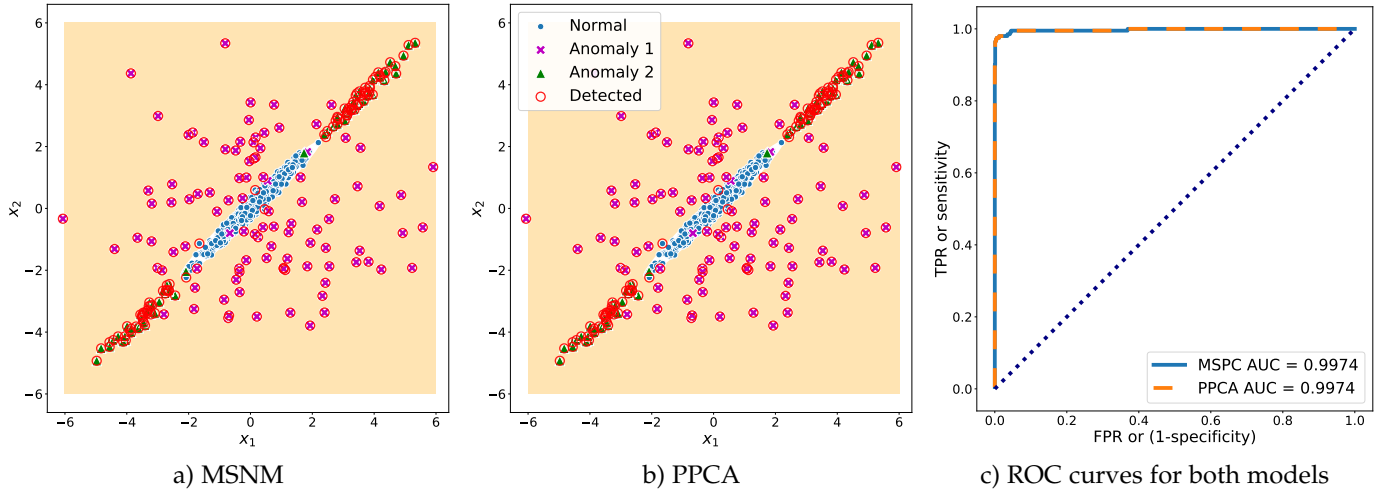


Figure 5. Anomaly detection on the synthetic dataset for both a) MSNM and b) PPCA models. Blue circles identify normal data, fuchsia \times s and green triangles identify different types of anomalies. Red circumferences mark the points detected as anomalies. The orange shadows cover the detection area in each case. c) ROC curves for both models.

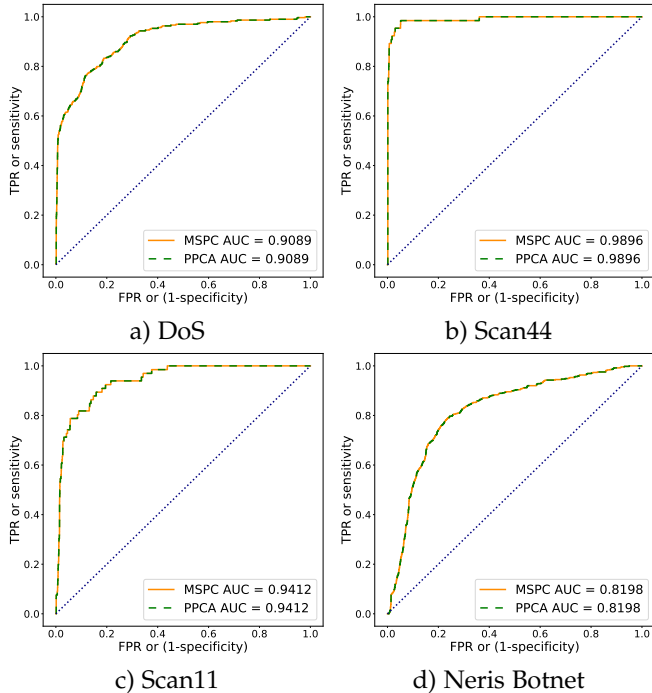


Figure 6. ROC curves for different attacks on UGR'16 using MSNM and PPCA with $P=3$.

both models when detecting different kind of attacks, using several numbers of principal components, $P = [1, 2, \dots, 5]$. For further details, Figure 6 shows the ROC curves obtained for both models in different attacks using $P = 3$.

6.2.2 Using the regularization and reconstruction error terms to detect different types of anomalies

As previously indicated, both terms in Eq. (26) are able to identify different behaviour in the data. To quickly gain an overview of how both terms work, Figure 7 depicts the anomalies in the synthetic dataset that are detected using a single PPCA/MSNM term only. Note that when using a single term, the threshold is the 99th-percentile

of the calibration set for that term. When we only use the reconstruction error term —Eq. (25)—, we can see in Figure 7(a) how we are unable to identify anomalies that have been generated following the linear model. Note that this case is equivalent to using a simple PCA analysis. Using the regularization term only —Eq. (24)—, see Figure 7(b), results in a similar problem. This term is calculated using only the latent $\hat{z}(\delta)$, therefore all data points whose latent value is within the distribution of the normal data remains undetected. To effectively detect all anomalies, it is clearly necessary to use both terms.

When it comes to complex networking data, such as those in UGR'16, it is difficult to predict which anomalies will be correctly detected by the two terms in the PPCA/MSNM model. Table 4 shows that different attacks are captured differently by both terms. Some attacks (DoS, Scan44) are captured correctly by both terms. Neris Botnet traffic can be identified using the reconstruction error term, but cannot usually be identified with the regularization term.

The results for Scan11 and UDP Scan attacks show that the ability of the model to identify anomalies using the error reconstruction or regularization terms depends on the number of principal components used. We have included Fig 8 as an example of where the discriminative power of the regularization term increases with P , even outperforming the error reconstruction term performance. Note that the error reconstruction term is not severely affected by the number of P because the explained variance of the model is low, and therefore the error reconstruction for most anomalies remains high. As the number of P increases, the latent space is able to capture more features of the data, also allowing anomalies in the latent domain to better identified.

6.2.3 Weighting reconstruction error and regularization terms

The previous experiments have shown that both terms in Eq. (26) are relevant and should be used. Furthermore, a single term is not usually able to correctly identify all types

Table 3
AUC Values for the attacks on UGR'16 using MSNM and PPCA

Attack type	1 P		2 P		3 P		4 P		5 P	
	MSNM	PPCA	MSNM	PPCA	MSNM	PPCA	MSNM	PPCA	MSNM	PPCA
DoS	0.9118	0.9118	0.9089	0.9089	0.9089	0.9089	0.9097	0.9097	0.9091	0.9091
Scan44	0.9903	0.9903	0.9902	0.9902	0.9896	0.9896	0.9882	0.9882	0.9880	0.9880
Scan11	0.9384	0.9384	0.9390	0.9390	0.9412	0.9412	0.9318	0.9318	0.9303	0.9303
Nerisbotnet	0.8204	0.8204	0.8211	0.8211	0.8198	0.8198	0.8201	0.8201	0.8203	0.8203
UDPscan	0.7826	0.7826	0.7844	0.7844	0.7707	0.7707	0.7707	0.7707	0.7727	0.7727
SSHscan	0.5593	0.5593	0.5624	0.5624	0.5569	0.5569	0.5588	0.5588	0.5614	0.5614
Spam	0.4669	0.4669	0.4610	0.4610	0.4588	0.4588	0.4585	0.4585	0.4512	0.4512

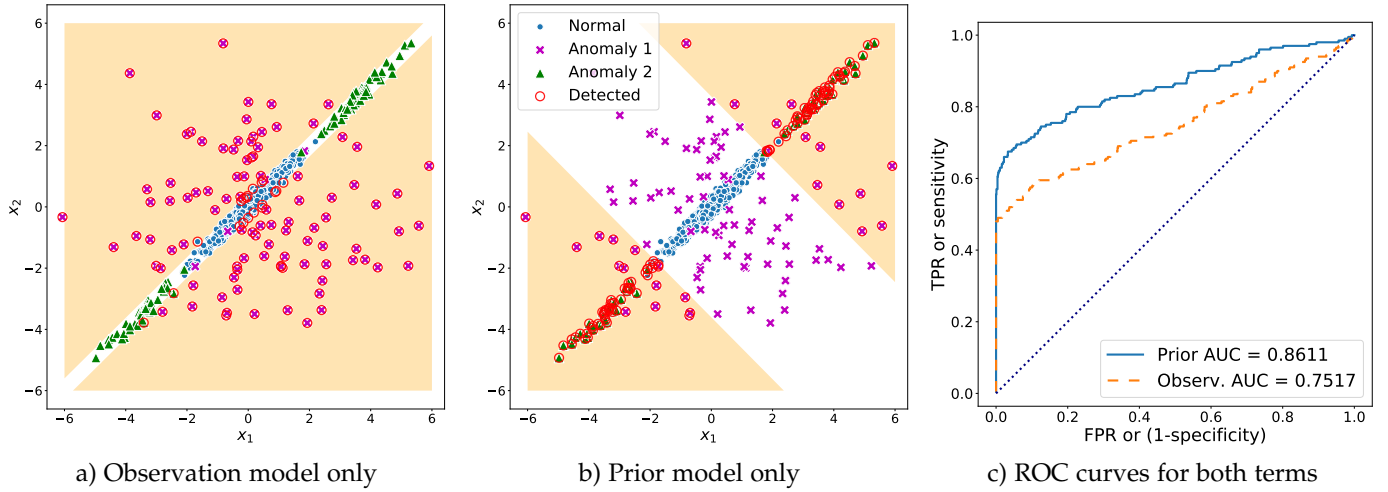


Figure 7. Detected anomalies in the synthetic dataset using only a single term. Blue circles identify normal data, fuchsia xs and green triangles identify different types of anomalies. Red circumferences mark the points detected as anomalies. The orange shadows cover the detection area in each case.

Table 4
AUC Values for the attacks on UGR'16 using divergence (Prior) and error (Observ.) terms separately.

Attack type	1 P		2 P		3 P		4 P		5 P	
	Prior	Observ.	Prior	Observ.	Prior	Observ.	Prior	Observ.	Prior	Observ.
DoS	0.5242	0.9118	0.9191	0.9085	0.9178	0.9085	0.9050	0.9095	0.9152	0.9087
Scan44	0.8747	0.9903	0.8483	0.9902	0.9968	0.9895	0.9960	0.9878	0.9957	0.9875
Scan11	0.4966	0.9385	0.3771	0.9391	0.7714	0.9416	0.9691	0.9304	0.9691	0.9285
Nerisbotnet	0.4358	0.8207	0.3686	0.8214	0.4741	0.8199	0.4480	0.8203	0.4926	0.8204
UDPscan	0.6618	0.7817	0.7084	0.7838	0.8886	0.7682	0.8659	0.7674	0.8372	0.7698
Mean	0.5986	0.8886	0.6443	0.8886	0.8097	0.8855	0.8368	0.8831	0.8419	0.8830

of anomalies. The forthcoming question is how to combine them to obtain a single anomaly score.

On the one hand, the work in [15] suggested the use of the significance levels UCL_Q and UCL_D and a weighting parameter calculated as P/M (see Section 3). On the other hand, PPCA provides a probabilistic interpretation for the weighting parameter, combining both terms according to their contribution to the marginal likelihood —see Eq. (17)—. In this section, we compare the MSNM approach in Eq. (4) with the PPCA approach in Eq. (17).

In the synthetic dataset, both strategies obtain a similar result in AUC: 0.9974 for PPCA and 0.9973 for MSNM. Note that in both cases, the combination of both terms obtains a better AUC than the values reported in the previous section for each term separately.

The difference between the results from PPCA and

MSNM is more clearly shown in Table 7, where the UGR'16 dataset is analyzed. Even when the MSNM approach results in slightly better AUC values for some combinations of attacks and number of P, the PPCA AUC score yields a better mean value in each case. While the differences are subtle, they indicate that PPCA is better when it comes to combining information from both regularization and reconstruction error terms. Looking at Table 7, we can observe that PPCA obtains better values than MSNM in those attacks where the regularization term is more informative. In addition, note that while a single term might be better at recognising a given attack type, the mean AUC values are always higher than those obtained using only a single term.

Table 5
Accuracy Values for the attacks on UGR'16 using divergence (Prior) and error (Observ.) terms separately.

Attack type	1 P		2 P		3 P		4 P		5 P	
	Prior	Observ.	Prior	Observ.	Prior	Observ.	Prior	Observ.	Prior	Observ.
Dos	0.9237	0.8946	0.9655	0.8939	0.9273	0.8932	0.8939	0.8943	0.9287	0.8922
scan44	0.9507	0.9011	0.9713	0.9008	0.9320	0.9004	0.8979	0.9015	0.9347	0.8991
scan11	0.9460	0.8999	0.9667	0.8995	0.9300	0.8992	0.8974	0.9001	0.9337	0.8977
nerisbotnet	0.9022	0.8801	0.9223	0.8798	0.8850	0.8795	0.8541	0.8806	0.8880	0.8788
anomaly-udpscan	0.9520	0.9001	0.9728	0.8998	0.9312	0.8995	0.8969	0.9005	0.9340	0.8981
mean	0.9349	0.8952	0.9597	0.8948	0.9211	0.8943	0.8880	0.8954	0.9238	0.8932

Table 6
Mean false alarm ratio on UGR'16 using divergence (Prior) and error (Observ.) terms separately.

1 P		2 P		3 P		4 P		5 P	
Prior	Observ.	Prior	Observ.	Prior	Observ.	Prior	Observ.	Prior	Observ.
0.0474	0.0995	0.0267	0.0998	0.0684	0.1002	0.1027	0.0992	0.0656	0.1016

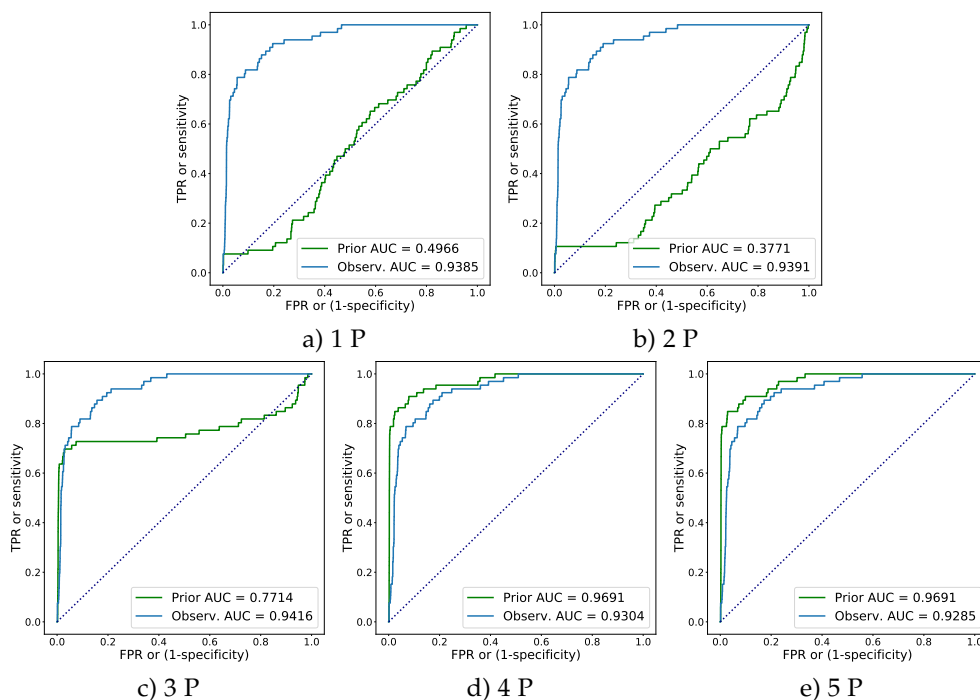


Figure 8. Evolution of the ROC curves of both divergence and error terms on the Scan11 attack for different latent space sizes.

7 CONCLUSIONS

In this paper we have analyzed the use of the generative model known as probabilistic PCA (PPCA) for detection of anomalies in network security. Specifically, we have provided a detailed mathematical model that connects MSNM, a well-known framework for network anomaly detection, and PPCA. The generative PPCA model provides a probabilistic point of view to explain the MSNM framework and the meaning of its principal elements: The use of Q and D statistics, which are derived as an error reconstruction term and a regularization term, respectively, in the PPCA formulation.

Understanding the role of both terms in the anomaly detection process is a key step towards the correct use of generative models in this security research field. Specifi-

cally, a direct application is that of correctly using other generative models like VAEs and GANs. In a review of research works that use these models, we note that while the error reconstruction term is widely used due to its high anomaly detection capability, the regularization term is often forgotten or discarded. We have theoretically and experimentally assessed that both terms are relevant and capture complementary information. This implies that they should be used together for a robust anomaly detection.

In addition, the PPCA generative framework provides a combination of both terms that considers their contribution to the marginal distribution $p(\mathbf{x})$ with a probabilistic interpretation, thus offering a non-empirical solution for the weighting parameter required by MSNM.

Although the linear PPCA generative model is easy

Table 7
AUC Values for the attacks on UGR'16 using different number of PCs and models: MSNM —Eq. (4)— and PPCA —Eq. (17)—.

Attack type	1 P		2 P		3 P		4 P		5 P	
	MSNM	PPCA	MSNM	PPCA	MSNM	PPCA	MSNM	PPCA	MSNM	PPCA
DoS	0.9118	0.9118	0.9085	0.9089	0.9085	0.9089	0.9095	0.9097	0.9087	0.9091
Scan44	0.9903	0.9903	0.9902	0.9902	0.9895	0.9896	0.9878	0.9882	0.9875	0.9880
Scan11	0.9385	0.9384	0.9391	0.9390	0.9415	0.9412	0.9305	0.9318	0.9286	0.9303
Nerisbotnet	0.8207	0.8204	0.8214	0.8211	0.8199	0.8198	0.8203	0.8201	0.8204	0.8203
UDPscan	0.7817	0.7826	0.7838	0.7844	0.7683	0.7707	0.7674	0.7707	0.7700	0.7727
Mean	0.8886	0.8887	0.8886	0.8887	0.8855	0.8860	0.8831	0.8841	0.8830	0.8841

Table 8
Accuracy Values for the attacks on UGR'16 using different number of PCs and models: MSNM —Eq. (4)— and PPCA —Eq. (17)—.

	MSNM	PPCA	MSNM	PPCA	MSNM	PPCA	MSNM	PPCA	MSNM	PPCA
Dos	0.8946	0.8947	0.8939	0.8946	0.8932	0.9655	0.8944	0.8947	0.8922	0.8932
scan44	0.9011	0.9012	0.9008	0.9011	0.9004	0.9713	0.9015	0.9017	0.8991	0.9004
scan11	0.8999	0.9000	0.8995	0.8999	0.8992	0.9667	0.9001	0.9003	0.8977	0.8992
nerisbotnet	0.8801	0.8802	0.8798	0.8801	0.8794	0.9223	0.8806	0.8807	0.8788	0.8795
anomaly-udpscan	0.9001	0.9003	0.8998	0.9001	0.8995	0.9728	0.9005	0.9007	0.8981	0.8995
mean	0.8952	0.8953	0.8948	0.8952	0.8943	0.9597	0.8954	0.8956	0.8932	0.8943

Table 9
Mean false alarm ratio on UGR'16 using different number of PCs and models: MSNM —Eq. (4)— and PPCA —Eq. (17)—.

1 P		2 P		3 P		4 P		5 P	
MSNM	PPCA	MSNM	PPCA	MSNM	PPCA	MSNM	PPCA	MSNM	PPCA
0.0995	0.0994	0.0998	0.1000	0.1002	0.1000	0.0992	0.0989	0.1016	0.1018

to understand and helps obtain the above conclusions, its simplicity limits its generalization to non-linear data. Non-linearity is often present in real traffic data and would be difficult to capture and detect. Also, PPCA inherits some of the disadvantages from PCA, which is quite common when working with latent space models. Even with a linear model, the latent combination of the original features is not easy to interpret. So, as shown in Figure 8, the choice of the latent space size P is critical for the detection of some attacks. While the use of linear detection models is useful in the later diagnosis of network incidents, this paper intends to establish a first step for more complex generative approaches to network anomaly detection. Further research on the combination of both error reconstruction and regularization terms for those models is also needed. Finally, the choice of the threshold is a well-known problem for anomaly detection. The calibration set provides a benchmark to choose the decision boundary, but it should be determined according to the confidence in the calibration data. The Gaussian form of $p(x|\mathbf{W}_{ML}, \sigma_{ML}^2)$ allows us to choose the confidence-based threshold. Although the threshold might also be experimentally determined with the use of testing sets, this approach relies on known attacks, and therefore does not provide information about the optimal threshold for new or unknown attacks. Finally, system requirements usually determine the threshold to use in industry applications. A high cost of undetected false negatives might induce the use of a lower threshold that will produce a higher number of false positives. When combining both decision terms, the use of a combined or separated threshold for each term is

also an open field for future research.

ACKNOWLEDGMENT

This work was sponsored in part by the Agencia Estatal de Investigación under project PID2019-105142RB-C22/AEI/10.13039/501100011033, Spanish MINECO (Ministerio de Economía y Competitividad) project TIN2017-83494-R and FEDER/Junta de Andalucía-Consejería de Transformación Económica, Industria y Universidades/project A-TIC-215-UGR18. The work by Fernando Pérez-Bueno was sponsored by Ministerio de Economía, Industria y Competitividad under FPI contract BES-2017-081584.

The authors would like to thank Daniel Cortés Troya for his collaboration in the early stages of this work.

APPENDIX

.1 Laplace approximation

It follows that:

$$p(\mathbf{x}) = \int p(\mathbf{z})p(\mathbf{x}|\mathbf{z})d\mathbf{z} = \int \exp[\ln p(\mathbf{z})p(\mathbf{x}|\mathbf{z})]d\mathbf{z} \quad (30)$$

where $f(\mathbf{z}) = \ln(p(\mathbf{z})p(\mathbf{x}|\mathbf{z}))$ is a quadratic function that can be expanded around the maximum a posteriori (MAP)

$$\hat{\mathbf{z}} = \arg \max_{\mathbf{z}} f(\mathbf{z}) = \arg \max_{\mathbf{z}} \ln(p(\mathbf{z})p(\mathbf{x}|\mathbf{z})) \quad (31)$$

to obtain

$$p(\mathbf{z})p(\mathbf{x}|\mathbf{z}) = \exp[\ln p(\mathbf{z})p(\mathbf{x}|\mathbf{z})] \quad (32)$$

$$\propto \exp[f(\hat{\mathbf{z}}) - \frac{1}{2}(\mathbf{z} - \hat{\mathbf{z}})^T(\mathbf{I} + \frac{1}{\sigma^2}\mathbf{W}^T\mathbf{W})(\mathbf{z} - \hat{\mathbf{z}})],$$

which produces, from Eq. (30), $\ln p(\mathbf{x}) = f(\hat{\mathbf{z}}) + \text{const}$, and so

$$\ln p(\mathbf{x}|\mathbf{W}, \boldsymbol{\sigma}^2) = -\frac{1}{2}(\hat{\mathbf{z}}^T \hat{\mathbf{z}} + \frac{1}{\sigma^2} \|\mathbf{x} - \mathbf{W}\hat{\mathbf{z}}\|^2) + \text{const.} \quad (33)$$

2 Analysis of $f_\alpha(\delta)$

Let us define

$$f_\alpha(\delta) = \frac{1}{2}(\hat{\mathbf{z}}^T(\delta)\hat{\mathbf{z}}(\delta) + \frac{1}{\alpha} \|\mathbf{x} - \mathbf{W}(\delta)\hat{\mathbf{z}}(\delta)\|^2), \quad (34)$$

and study its properties.

We observe that this function depends on \mathbf{x} but we do not make this dependency explicit in order to simplify the notation. Note also that we will end up studying the basic properties of the associated quadratic form.

Theorem 1. *Let us consider $f_\alpha(\delta)$ defined in Eq. (34) with $\alpha < \lambda_P$ and $\delta \in [0, \lambda_P)$. Then*

- $f_\alpha(\delta)$ is a convex function on δ with minimum at $\alpha/2$,
- and $f_\alpha(0) = f_\alpha(\alpha)$.

Proof. Therefore,

$$\begin{aligned} f_\alpha(\delta) &= \frac{1}{2}(\mathbf{x}^T \mathbf{U}(\mathbf{L} - \delta \mathbf{I})\mathbf{L}^{-2} \mathbf{U}^T \mathbf{x} + \frac{1}{\alpha} \|\mathbf{x} - \mathbf{U}(\mathbf{L} - \delta \mathbf{I})\mathbf{L}^{-1} \mathbf{U}^T \mathbf{x}\|^2) \\ &= \frac{1}{2}(\mathbf{x}^T \mathbf{U}(\mathbf{L} - \delta \mathbf{I})\mathbf{L}^{-2} \mathbf{U}^T \mathbf{x} \\ &+ \frac{1}{\alpha}(\mathbf{x}^T \mathbf{x} + \mathbf{x}^T \mathbf{U}(\mathbf{L} - \delta \mathbf{I})^2 \mathbf{L}^{-2} \mathbf{U}^T \mathbf{x} - 2\mathbf{x}^T \mathbf{U}(\mathbf{L} - \delta \mathbf{I})\mathbf{L}^{-1} \mathbf{U}^T \mathbf{x})) \end{aligned}$$

and

$$\begin{aligned} f'_\alpha(\delta) &= \frac{1}{2}(-\mathbf{x}^T \mathbf{U} \mathbf{L}^{-2} \mathbf{U}^T \mathbf{x} \\ &+ \frac{1}{\alpha}(-2\mathbf{x}^T \mathbf{U}(\mathbf{L} - \delta \mathbf{I})\mathbf{L}^{-2} \mathbf{U}^T \mathbf{x} + 2\mathbf{x}^T \mathbf{U} \mathbf{L}^{-1} \mathbf{U}^T \mathbf{x})). \end{aligned} \quad (35)$$

Furthermore, we have the following identity for the sum of the matrices involved in $f'_\alpha(\delta)$,

$$-\mathbf{L}^{-2} + \frac{2}{\alpha}(-(\mathbf{L} - \delta \mathbf{I})\mathbf{L}^{-2} + \mathbf{L}^{-1}) = \mathbf{L}^{-2}[-\mathbf{I} + \frac{2}{\alpha}\delta \mathbf{I}] \quad (36)$$

from which we can see that the sum is the zero matrix iff $\delta = \alpha/2$.

We also note that

$$f''_\alpha(\delta) = \frac{1}{2\alpha} \mathbf{x}^T \mathbf{U} \mathbf{L}^{-2} \mathbf{U}^T \mathbf{x} \geq 0 \quad (37)$$

So, $f_\alpha(\delta)$ is a convex quadratic function on δ whose minimum is achieved at $\delta = \alpha/2$.

Furthermore, using the Taylor expansion around the minimum it follows that

$$f_\alpha(\delta) = f_\alpha\left(\frac{\alpha}{2}\right) + \frac{1}{2}\left(\delta - \frac{\alpha}{2}\right)^2 f''_\alpha\left(\frac{\alpha}{2}\right) \quad (38)$$

and so $f_\alpha(0) = f_\alpha(\alpha)$.

In summary, $f'_\alpha(\delta)$ is convex, its minimum value is achieved at $\delta = \alpha/2$, and $f_\alpha(0) = f_\alpha(\alpha)$. \square

REFERENCES

- [1] R. Contu, D. Kish, C. Canales, S. Deshpande, E. Kim, and D. Gardner, "Forecast analysis: Information security, worldwide, 2q18 update," [urlhttps://www.gartner.com/en/documents/3889055](https://www.gartner.com/en/documents/3889055), 2018.
- [2] New York Times, "Blackout Hits Iran Nuclear Site in What Appears to Be Israeli Sabotage," <https://www.nytimes.com/2021/04/11/world/middleeast/iran-nuclear-natanz.html>, 2021.
- [3] A. B. Nassif, M. A. Talib, Q. Nasir, and F. M. Dakalbab, "Machine learning for anomaly detection: A systematic review," *IEEE Access*, vol. 9, pp. 78 658–78 700, 2021.
- [4] D. Kwon, H. Kim, J. Kim, S. C. Suh, I. Kim, and K. J. Kim, "A survey of deep learning-based network anomaly detection," *Cluster Computing*, vol. 22, no. 1, pp. 949–961, 2019.
- [5] R. Chalapathy and S. Chawla, "Deep learning for anomaly detection: A survey," 2019.
- [6] L. Ruff, J. R. Kauffmann, R. A. Vandermeulen, G. Montavon, W. Samek, M. Kloft, T. G. Dietterich, and K.-R. Müller, "A unifying review of deep and shallow anomaly detection," *Proceedings of the IEEE*, 2021.
- [7] A. Lakhina, M. Crovella, and C. Diot, "Diagnosing network-wide traffic anomalies," *ACM SIGCOMM computer communication review*, vol. 34, no. 4, pp. 219–230, 2004.
- [8] —, "Characterization of network-wide anomalies in traffic flows," in *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, 2004, pp. 201–206.
- [9] G. Maciá-Fernández, J. Camacho, P. García-Teodoro, and R. A. Rodríguez-Gómez, "Hierarchical PCA-based multivariate statistical network monitoring for anomaly detection," in *2016 IEEE international workshop on information forensics and security (WIFS)*. IEEE, 2016, pp. 1–6.
- [10] H. Hotelling, "Multivariate quality control," *Techniques of Statistical Analysis*, 1947. [Online]. Available: <https://ci.nii.ac.jp/naid/10021322508/en/>
- [11] J. Camacho, J. M. García-Jiménez, N. M. Fuentes-García, and G. Maciá-Fernández, "Multivariate big data analysis for intrusion detection: 5 steps from the haystack to the needle," *Computers & Security*, vol. 87, p. 101603, 2019.
- [12] M. E. Tipping and C. M. Bishop, "Probabilistic principal component analysis," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 61, no. 3, pp. 611–622, 1999.
- [13] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," 2014.
- [14] H. Xu, W. Chen, N. Zhao, Z. Li, J. Bu, Z. Li, Y. Liu, Y. Zhao, D. Pei, Y. Feng, J. Chen, Z. Wang, and H. Qiao, "Unsupervised Anomaly Detection via Variational Auto-Encoder for seasonal KPIs in Web Applications," in *Proc. of the 2018 World Wide Web Conference*, 2018, pp. 187–196.
- [15] J. Camacho, P. García-Teodoro, and G. Maciá-Fernández, "Traffic monitoring and diagnosis with multivariate statistical network monitoring: a case study," in *2017 IEEE security and privacy workshops (SPW)*. IEEE, 2017, pp. 241–246.
- [16] A. Soule, A. Lakhina, N. Taft, K. Papagiannaki, K. Salamantian, M. C. Antonio Nucci, and C. Diot, "Traffic matrices: balancing measurements, inference and modeling," vol. 33, no. 1. USENIX Association, Berkeley, CA, USA, 2005, pp. 331–344.
- [17] Y. Zhang, Z. Ge, A. Greenberg, and M. Roughan, "Network anomography," in *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement (IMC'05)*. USENIX Association, Berkeley, CA, USA, 2005, pp. 317–330.
- [18] L. Huang, X. Nguyen, M. Garofalakis, M. Jordan, A. Joseph, and N. Taft, "In-network pcas and anomaly detection," in *Proceedings of Neural Information Processing Systems (NIPS)*. NIPS Foundation, 2006.
- [19] H. Ringberg, A. Soule, J. Rexford, and C. Diot, "Sensitivity of PCA for traffic anomaly detection," *SIGMETRICS. Perform Eval. Rev.*, vol. 35, no. 1, pp. 109–120, 2007.
- [20] B. Rubinstein, B. Nelson, L. Huang, A. Joseph, S. hon Lau, R. Satish, N. Taft, , and J. Tygar, "Antidote: understanding and defending against poisoning of anomaly detectors," 2009.
- [21] Z. Wang, K. Hu, K. Xu, Y. Baolin, and X. Dong, "Structural analysis of network traffic matrix via relaxed principal component pursuit," *Computer Networks*, vol. 56, pp. 2049–2067, 2012.
- [22] C. Pascoal, M. De Oliveira, R. Valadas, P. Filzmoser, P. Salvador, and A. Pacheco, "Robust feature selection and robust pca for internet traffic anomaly detection," 2012, pp. 1755–1763.

- [23] D. Brauckhoff, K. Salamantian, and M. May, "Applying PCA for Traffic Anomaly Detection: Problems and Solutions," in *Proceedings of IEEE INFOCOM*, 2009, pp. 2886–2870.
- [24] T. Kourti and J. F. MacGregor, "Multivariate SPC methods for process and product monitoring," *Journal of Quality Technology*, vol. 28, no. 4, 1996.
- [25] A. Ferrer, "Latent Structures-Based Multivariate Statistical Process Control: A Paradigm Shift," *Quality Engineering*, vol. 26, pp. 72–91, Jan. 2014, publisher: Taylor & Francis _eprint: <https://doi.org/10.1080/08982112.2013.846093>.
- [26] J. M. González-Martínez, A. Ferrer, and J. A. Westerhuis, "Real-time synchronization of batch trajectories for on-line multivariate statistical process control using Dynamic Time Warping," *Chemometrics and Intelligent Laboratory Systems*, vol. 105, pp. 195–206, 2011.
- [27] S. Soufiane, R. Magán-Carrión, I. Medina-Bulo, and H. Bouden, "Preserving authentication and availability security services through multivariate statistical network monitoring," *Journal of Information Security and Applications*, vol. 58, p. 102785, 2021.
- [28] M. López-Martin, B. Carro, A. Sanchez-Esguevillas, and J. Lloret, "Conditional variational autoencoder for prediction and feature recovery applied to intrusion detection in iot," *Sensors*, vol. 17, p. 1967, 2017.
- [29] J. An and S. Cho, "Variational autoencoder based anomaly detection using reconstruction probability," *Technical Report*, pp. 1–18, 2015.
- [30] S. Zavrak and M. Iskefiyelo, "Anomaly-based intrusion detection from network flow features," *IEEE Access*, vol. 8, pp. 108 346–108 358, 2020.
- [31] M. Sölch, J. Bayer, M. Ludersdorfer, and P. van der Smagt, "Variational inference for on-line anomaly detection in high-dimensional time series," in *International Conference on Machine Learning anomaly detection workshop*. IEEE, 2016.
- [32] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Berlin, Heidelberg: Springer-Verlag, 2006.
- [33] G. Maciá-Fernández, J. Camacho, R. Magán-Carrión, P. García-Teodoro, and R. Therón, "UGR'16: A new dataset for the evaluation of cyclostationarity-based network IDSs," *Computers & Security*, vol. 73, pp. 411–424, 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167404817302353>
- [34] M. Ring, S. Wunderlich, D. Scheuring, D. Landes, and A. Hotho, "A survey of network-based intrusion detection data sets," *Computers & Security*, vol. 86, pp. 147–167, 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S016740481930118X>
- [35] M. Catillo, A. Pecchia, M. Rak, and U. Villano, "Demystifying the role of public intrusion datasets: A replication study of dos network traffic data," *Computers & Security*, vol. 108, p. 102341, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167404821001656>
- [36] S. García, M. Grill, J. Stiborek, and A. Zunino, "An empirical comparison of botnet detection methods," *Computers & Security*, vol. 45, pp. 100–123, 2014. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167404814000923>