Network Anomaly Detection with Net-GAN, a Generative Adversarial Network for Analysis of Multivariate Time-Series

Gastón García González
Universidad de la República & AIT
gastong@fing.edu.uy

Pedro Casas
AIT Austrian Institute of Technology
pedro.casas@ait.ac.at

Alicia Fenández
Universidad de la República
alicia@fing.edu.uy

Gabriel Gómez
Universidad de la República
ggomez@fing.edu.uy

ABSTRACT

We introduce *Net-GAN*, a novel approach to network anomaly detection in time-series, using recurrent neural networks (RNNs) and generative adversarial networks (GAN). Different from the state of the art, which traditionally focuses on univariate measurements, Net-GAN detects anomalies in multivariate time-series, exploiting temporal dependencies through RNNs. Net-GAN discovers the underlying distribution of the baseline, multivariate data, without making any assumptions on its nature, offering a powerful approach to detect anomalies in complex, difficult to model network monitoring data. We present preliminary detection results in different monitoring scenarios, including anomaly detection in sensor data, and intrusion detection in network measurements.

CCS CONCEPTS

• **Computing methodologies** → **Anomaly detection**; **Machine learning algorithms**;

KEYWORDS

Anomaly Detection; Multivariate Time-Series; Generative Models; GAN; LSTM

ACM Reference Format:
Gastón García González, Pedro Casas, Alicia Fenández, and Gabriel Gómez. 2020. Network Anomaly Detection with Net-GAN, a Generative Adversarial Network for Analysis of Multivariate Time-Series. In *ACM Special Interest Group on Data Communication (SIGCOMM '20 Demos and Posters), August 10–14, 2020, Virtual Event, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3405837.3411393

1 INTRODUCTION

Network monitoring data generally consists of hundreds or thousands of counters periodically collected in the form of time-series, resulting in a complex-to-analyze multivariate time-series process (MTS). In particular, detecting anomalies in such multivariate, temporal data is challenging. Without loss of generality, we refer to the MTS as a set of n, non-iid time series sampled at the same rate, referred to as $x_t = \{x_t(1), x_t(2), \ldots, x_t(n)\} \in \mathbb{R}^n$. Current approaches to anomaly detection tackle this challenge by either focusing on univariate time-series analysis – running an independent detector for each time-series $x_t(i)$, or by considering multidimensional input data $x \in \mathbb{R}^n$ at each time t, neglecting the temporal aspects of the MTS. To improve the state of affairs we propose Net-GAN, a novel unsupervised approach to anomaly detection in MTS data, based on Recurrent Neural Networks (RNNs), trained through a Generative Adversarial Networks framework (GAN) [2].

The usage of generative models for semi-supervised anomaly detection helps to solve two major problems faced in this specific field: the high imbalance between normal operation and anomaly instances, as well as the lack of labeled instances for learning and validation purposes. Generative models such as Variational Auto-Encoders (VAEs) or Generative Adversarial Networks (GANs) are powerful approaches to learn the underlying distributions of data samples, in a purely data-driven, model-agnostic manner. Such models can be used in the practice to construct better baselines (i.e., profiles for normal operation) for the anomaly detection task, improving the identification of instances which deviate from this baseline. Examples of VAEs and GANs for anomaly detection are presented in [6] and [7], respectively. Most of previous work in this direction treats data as temporally independent samples, loosing the information provided by causality and temporal correlation.

To capture the temporal correlations characterizing an MTS, we adapt the original GAN model proposed in [2], replacing the multilayer perceptrons by recursive, LSTM networks for both generator and discriminator models. The input data is therefore sequences of multi-dimensional measurements, of length T: $\{x_{t-T}, ..., x_t\}$. Net-GAN is inspired by previous work on GANs for time-series synthesizing and anomaly detection [1, 3, 4].

2 THE NET-GAN APPROACH

Fig. 1 depicts the Net-GAN architecture and both the model training and anomaly detection procedures. In the ***training phase*** (left), the generator G draws synthetic sample sequences $G(z)$ from Gaussian noise – the latent space Z, with the objective of deceiving the discriminator D, which in turn learns to determine whether training samples are real or derived from the generative distribution. The classification result proposed by D is additionally fed back to G, serving as a reinforcement loop to guide the generation process. As both G and D compete to achieve their adversarial tasks, synthetic samples become more and more "realistic", and the discriminator becomes robust to noise, improving the detection of non-conforming

SIGCOMM '20 Demos and Posters, August 10–14, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-8048-5/20/08…$15.00
https://doi.org/10.1145/3405837.3411393

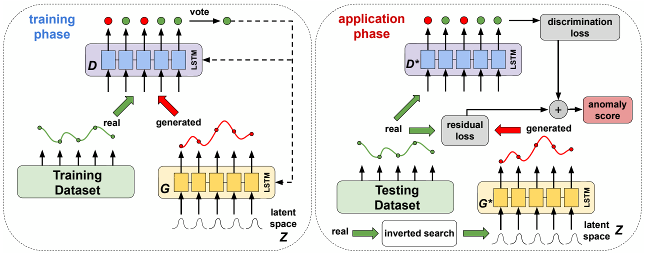

Figure 1: Net-GAN architecture and its application.

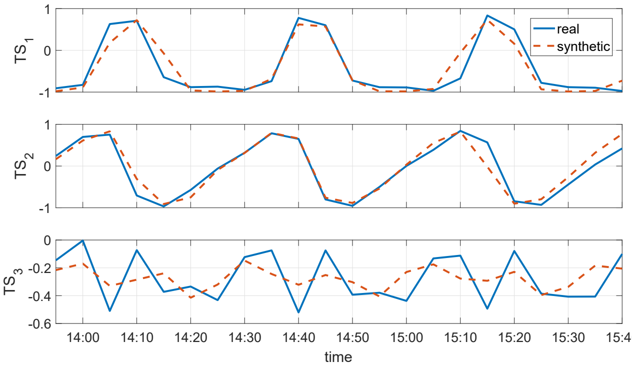

Figure 2: Synthetically generated time-series.

(i.e., out of the baseline) samples. In the ***application phase*** (right), the trained discriminator D^* acts naturally as an anomaly detector, detecting deviations from the baseline, through a *discrimination loss* function. The trained generator G^* is also used to improve detection performance, serving as baseline generation; by doing an inverse search in the latent space – for example, by constructing an inverse model for the generator [7], we find the sample $z \in Z$ which generates the closest sample to the tested one, producing a *residual loss*. Finally, both the discrimination and the residual loss functions are combined into an *anomaly score*, which is compared to a calibrated threshold to take the final decision.

3 PRELIMINARY EVALUATION RESULTS

We evaluate Net-GAN's anomaly detection performance on two different publicly available datasets, here referred to as CPS [3] and SYN-NET [5], respectively. The CPS dataset consists of synthetically generated attacks targeting industrial control systems, in particular a safe water treatment plant. It includes sensor measurements for 51 different physical properties related to the plant and the water treatment process. In total, 946.722 samples are collected with a 1-second resolution, over 11 days. The SYN-NET dataset is a synthetically generated dataset for network intrusion detection, including normal operation traffic generated by a group of 25 users (e.g., HTTP/HTTPS browsing, FTP file transfer, SSH and mail, etc.), with controlled attacks over-imposed, of very different nature. In particular, we test Net-GAN for the detection of botnet traffic (0.2% of total flows), DDoS attacks (4.3%), port scan activity (16.6%), and infiltration activity (only 36 flows). SYN-NET consists of more than a million flows – 83%/17% benign/malign traffic.

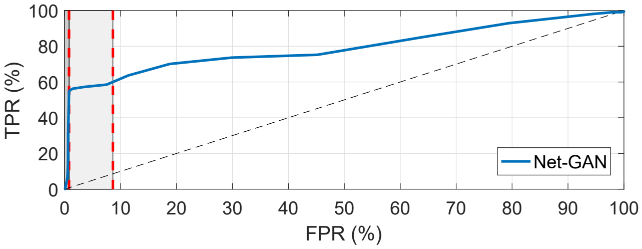

Figure 3: Detection of anomalies in CPS data.

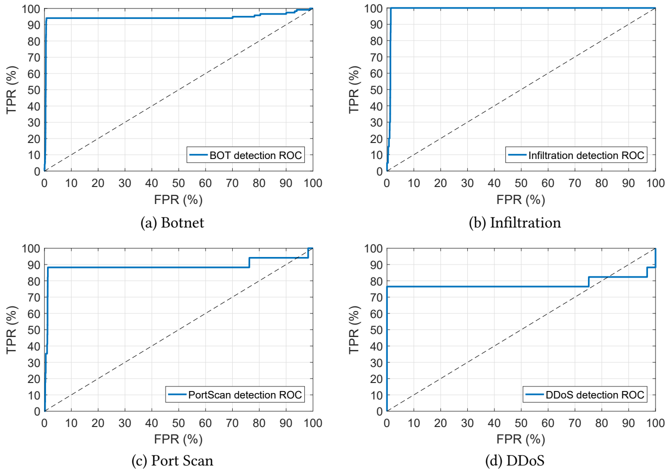

(a) Botnet

(b) Infiltration

(c) Port Scan

(d) DDoS

Figure 4: Detection of attacks in SYN-NET data

To show the generation/synthesization capabilities of Net-GAN, Fig. 2 depicts some of the (min-max normalized) time series generated by the trained generator G^*, along with the corresponding real time-series. To reduce noise, samples are aggregated in time-windows of 10'. Time-series of higher magnitude are better reconstructed (TS_1 and TS_2), whereas noisy ones – such as TS_3, are more difficult to track.

Fig. 3 reports the detection performance achieved by Net-GAN in the CPS dataset. For simplicity, we only use D^* as detector. While results are preliminary, Net-GAN detects 56% of the attacks with a FPR below 1%. Finally, Fig. 4 reports the detection performance in the SYN-NET dataset, for the four considered network attacks. About 93%, 100%, 89%, and 78% of the attacks are detected with a FPR below 1%, for botnet, infiltration, port scan, and DDoS, respectively, showing promising results. Being purely data-driven, we are currently working with bigger network measurement datasets to improve Net-GAN generation capabilities.

ACKNOWLEDGMENTS

This work has been partially supported by the ANII-FMV project FMV_1_2019_1_155850 "Anomaly detection with continual and streaming machine learning on big data telecommunications networks", and by Telefónica. Gastón García was supported by a ANII No. POS_NAC_2018_1_152279 scholarship and by CSIC-UdelaR

REFERENCES

[1] Cristóbal Esteban, Stephanie L. Hyland, and Gunnar Rätsch. 2017. Real-valued (Medical) Time Series Generation with Recurrent Conditional GANs. *CoRR* abs/1706.02633 (2017). arXiv:1706.02633 http://arxiv.org/abs/1706.02633

[2] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2 (NIPS'14)*. MIT Press, Cambridge, MA, USA, 2672–2680.

[3] Dan Li, Dacheng Chen, Baihong Jin, Lei Shi, Jonathan Goh, and See-Kiong Ng. 2019. MAD-GAN: Multivariate Anomaly Detection for Time Series Data with Generative Adversarial Networks. In *Artificial Neural Networks and Machine Learning - ICANN 2019: Text and Time Series - 28th International Conference on Artificial Neural Networks, Munich, Germany, September 17-19, 2019, Proceedings, Part IV (Lecture Notes in Computer Science)*, Igor V. Tetko, Vera Kurková, Pavel Karpov, and Fabian J. Theis (Eds.), Vol. 11730. Springer, 703–716. https://doi.org/10.1007/978-3-030-30490-4_56

[4] Federico Di Mattia, Paolo Galeone, Michele De Simoni, and Emanuele Ghelfi. 2019. A Survey on GANs for Anomaly Detection. *CoRR* abs/1906.11632 (2019). arXiv:1906.11632 http://arxiv.org/abs/1906.11632

[5] Iman Sharafaldin., Arash Habibi Lashkari., and Ali A. Ghorbani. 2018. Toward Generating a New Intrusion Detection Dataset and Intrusion Traffic Characterization. In *Proceedings of the 4th International Conference on Information Systems Security and Privacy - Volume 1: ICISSP,*. INSTICC, SciTePress, 108–116. https://doi.org/10.5220/0006639801080116

[6] S. Zavrak and M. Äřskefiyeli. 2020. Anomaly-Based Intrusion Detection From Network Flow Features Using Variational Autoencoder. *IEEE Access* 8 (2020), 108346–108358.

[7] Houssam Zenati, Chuan Sheng Foo, Bruno Lecouat, Gaurav Manek, and Vijay Ramaseshan Chandrasekhar. 2018. Efficient GAN-Based Anomaly Detection. *CoRR* abs/1802.06222 (2018). arXiv:1802.06222 http://arxiv.org/abs/1802.06222