Steps Towards Continual Learning in Multivariate Time-Series Anomaly Detection using Variational Autoencoders

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

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

A. Fernández, G. Gómez
Universidad de la República
Uruguay
alicia/ggomez@fing.edu.uy

ABSTRACT

We present *DC-VAE*, an approach to network anomaly detection in multivariate time-series (MTS), using Variational Auto Encoders (VAEs) and Dilated Convolutional Neural Networks (CNN). *DC-VAE* detects anomalies in MTS data through a single model, exploiting temporal and spatial MTS information. We showcase *DC-VAE* in different MTS datasets, and portray its future application in a continual learning framework, exploiting the generative properties of the underlying generative model to deal with continuously evolving data, avoiding catastrophic forgetting. We showcase the functioning of *DC-VAE* in the event of concept drifts, and propose the application of a novel approach to generative-driven continual learning, introducing the Deep Generative Replay model.

CCS CONCEPTS

• **Mathematics of computing → Time series analysis**; • **Networks → Network monitoring**.

KEYWORDS

Variational Autoencoders; Time-Series; Anomaly Detection

ACM Reference Format:
Gastón García González, Pedro Casas, and A. Fernández, G. Gómez. 2022. Steps Towards Continual Learning in Multivariate Time-Series Anomaly Detection using Variational Autoencoders. In *Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22), October 25–27, 2022, Nice, France.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3517745.3563033

1 INTRODUCTION TO *DC-VAE*

Network monitoring data often consists of hundreds or thousands of variables periodically measured and analyzed in the form of time-series, resulting in a complex-to-analyze multivariate time-series (MTS) process. In this paper we introduce *DC-VAE*, a deep-learning based approach to anomaly detection in multivariate time-series, based on popular Variational Auto-Encoders (VAEs) [1, 2, 5]. VAEs are a generative version of classical auto-encoders, with the particularity of having, by conception, a probabilistic manner to describe an observation in the latent space. Thus, rather than training an encoder which outputs a single value describing each latent state attribute, the encoder is formulated to describe a probability distribution for each latent attribute. For a given input, VAEs produce as output prediction not only an expected value, but also the associated standard deviation, corresponding to the distribution the model understands (i.e., has learned) generated the corresponding input. This automatically defines a *normality region* for each independent time-series, which can then be easily exploited for detecting deviations beyond this region. To exploit the temporal dependencies and characteristics of time-series data in a fast and efficient manner, we take a Dilated Convolutional (DC) Neural Network (NN) as the VAE's encoder and decoder architecture. Compared to normal convolutions, dilated convolutions improve time-series modeling by increasing the receptive field of the neural network, reducing computational and memory requirements, and most importantly, enabling training – and detection – on longer-in-the-past temporal sequences.

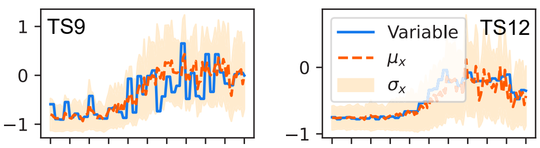

(a) Prediction of TS_9. (b) Prediction of TS_{12}.

Figure 1: Example of time-series analysis through *DC-VAE*. Normal-operation is defined by μ_x and σ_x.

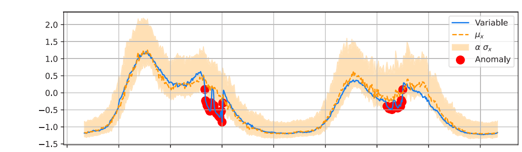

(a) Example of real anomalies in TS_2.

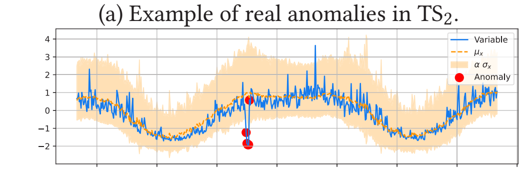

(b) Example of real anomalies in TS_4.

Figure 2: Examples of real anomalies present in the analyzed dataset, and their identification by *DC-VAE*.

2 *DC-VAE* AND CONCEPT DRIFT

Figs. 1 and 2 present the main ideas behind the usage of *DC-VAE* for time-series anomaly detection, in this case portraying the results obtained in the analysis of the TELCO dataset, a proprietary MTS dataset, corresponding to real measurements collected at an

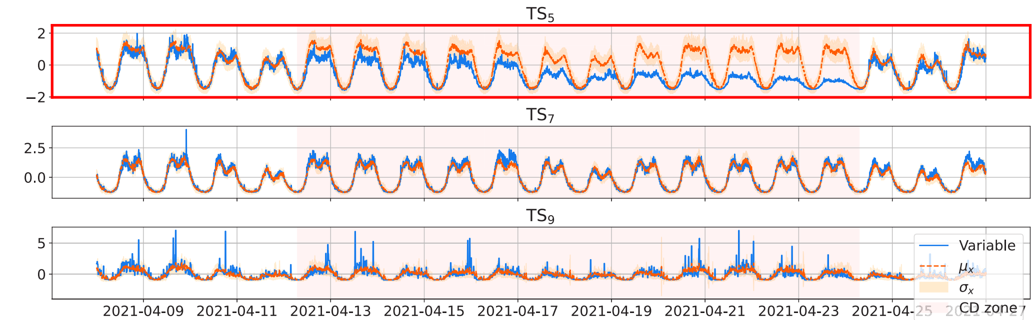

Figure 3: *DC-VAE* response to univariate concept-drift: a gradual linear fall of the values during the day.

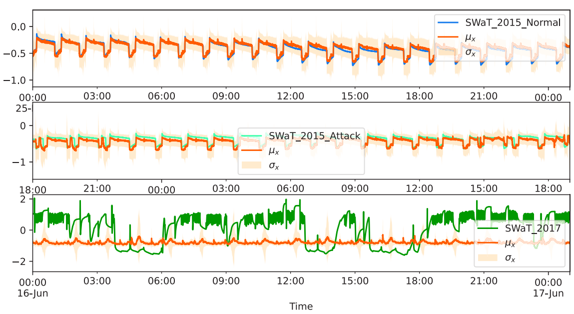

(a) *DC-VAE* under strong concept drift.

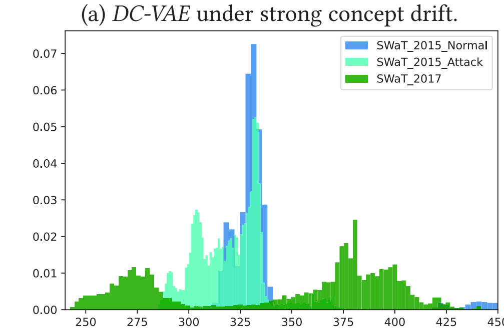

(b) Empirical distributions of subsets S_{2015-N}, S_{2015-A}, and S_{2017}.

Figure 4: Strong subset changes requires re-training.

operation mobile ISP – we are in the process of publicly releasing this dataset to the community. In a nutshell, if the VAE model was trained (mainly) with data describing the normal behavior of the monitored system, then the output for a non-anomalous input $\boldsymbol{x}$ would not deviate from the mean $\boldsymbol{\mu_x}$ more than a specific integer $\boldsymbol{\alpha}$ (calibrated with ground-truth data) times the standard deviation $\boldsymbol{\sigma_x}$. One of the main challenges faced by learning-driven anomaly detection system is their ability to cope with Concept Drift (CD) in the analyzed data – i.e., modifications of the underlying distribution. CD can manifest itself as a shift in the mean, an increase or decrease in the variance, or even as complete data modifications. These CD changes may be related to important trends in the data, requiring proper detection and re-training. Fig. 3 shows an example of *DC-VAE* operation under a CD, where a gradual change in the interval indicated as the CD zone is simulated in TS_5. *DC-VAE* is not capable to track this individual CD, given its multivariate nature. Fig. 4 shows *DC-VAE* under a more drastic CD, in this case considering data from different years (2015 and 2017) from the open SWaT dataset [4] – commonly used for detection of cyber-attacks. Fig. 4(a) shows the tracking of *DC-VAE* in (top) the 2015 normal operation dataset used for training, (middle) the 2015 attack dataset used for testing, and (bottom) the 2017 dataset. *DC-VAE* totally fails to capture the SWaT dataset in 2017, as the underlying distributions of the corresponding data are significantly different, as evidenced in Fig. 4(b).

3 CONTINUAL LEARNING FOR *DC-VAE*

We therefore explore different approaches to cope with the described CDs, in particular exploiting the generative nature of the *DC-VAE* model for continual learning. In a continual learning framework, we assume a continuously evolving stream of data, represented as a sequence of subsets S_i, each characterized by a different underlying distribution. We define a sequence of λ subsets $S_1, \dots, S_\lambda$ sequentially arriving, and assume access to only the data in current subset S_t. We are currently exploring different approaches to tackle this problem, considering two recent models referred to as Deep Generative Replay (DGR) [6] and BooVAE [3]. Given its simplicity and model elegance, we have decided to take DGR as the starting approach to extend *DC-VAE* to the continual learning setup.

The DGR approach uses a *teacher* generative model to generate synthetic data $F_{1\to(t-1)}$ that mimics former training examples in $S_1, \dots, S_{t-1}$. Then, the new *student* model is trained on joint synthetic data F and new data S_t. This approach is conceptually simple, model-agnostic and overcomes catastrophic forgetting, but requires retraining the model while generating the dataset from all previous subsets.

REFERENCES

[1] Ane Blázquez-García, Angel Conde, Usue Mori, and Jose A. Lozano. 2021. A Review on Outlier/Anomaly Detection in Time Series Data. *ACM Comput. Surv.* 54, 3, Article 56 (April 2021), 33 pages.
[2] Diederik P Kingma and Max Welling. 2013. Auto-encoding Variational Bayes. *arXiv preprint arXiv:1312.6114* (2013).
[3] Anna Kuzina, Evgenii Egorov, and Evgeny Burnaev. 2021. BooVAE: Boosting Approach for Continual Learning of VAE. *Advances in Neural Information Processing Systems* 35 (2021).
[4] A. P. Mathur and N. O. Tippenhauer. 2016. SWaT: A Water Treatment Testbed for Research and Training on ICS Security. In *IEEE International Workshop on Cyber-physical Systems for Smart Water Networks (CySWater)*. 31–36.
[5] Guansong Pang, Chunhua Shen, Longbing Cao, and Anton Van Den Hengel. 2021. Deep Learning for Anomaly Detection: A Review. *ACM Comput. Surv.* 54, 2, Article 38 (March 2021), 38 pages.
[6] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. 2017. Continual Learning with Deep Generative Replay. *CoRR* abs/1705.08690 (2017). arXiv:1705.08690