

Deep Generative Replay for Multivariate Time-Series Monitoring with Variational Autoencoders

Gastón García González^{*†}, Pedro Casas^{*}, Alicia Fernández[†]

[†]IIE-FING, Universidad de la República, ^{*}AIT Austrian Institute of Technology

Abstract—Multivariate time-series (MTS) play a crucial role in network monitoring and analysis problems. We explore the usage of generative AI for MTS data modeling, in particular for the sake of *knowledge replay*. Knowledge replay mechanisms help in leveraging past experiences to enhance learning, mitigate forgetting, promote generalization, and enable the transfer of knowledge across different tasks or domains. Using a VAE-based deep architecture for data modeling, we incorporate a Deep Generative Replay (DGR) approach to transfer previously learned latent representations into future learning tasks, enabling continual learning in MTS problems. We study the generative characteristics of VAE-based models on top of a multi-dimensional network monitoring dataset collected from an operational mobile Internet Service Provider (ISP), portraying its usage in the context of DGR learning tasks.

Index Terms—Anomaly Detection, Generative AI, VAE, Multivariate Time-Series, GenDeX

I. INTRODUCTION

Time-series analysis is an essential approach to network monitoring, in particular to profile temporal data behaviors and to detect anomalies in real-time. While time-series based anomaly detection has a long standing literature associated to signal processing techniques [1], modern approaches to time-series anomaly detection based on deep learning technology have flourished in recent years [2]. We have recently introduced *DC-VAE* [3], a deep-learning based approach to *unsupervised anomaly detection* in multivariate time-series (MTS), based on Variational Auto-Encoders (VAEs) [4]. VAEs are a generative version of classical auto-encoders; for a given input, they produce as output prediction not only an expected value, but also the associated standard deviation, corresponding to the distribution the model has learnt. 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 dimension of the input time-series, *DC-VAE* encoder/decoder architecture is based on popular CNNs, using Dilated Convolutions (DCs) [5].

A general problem faced by AI/ML-driven approaches for anomaly detection is their inability to deal with so-called *concept drifts*. Concept drifts correspond to events where the statistical properties of the target variable or the relationships between the input features and the target variable change over time. As such, the patterns and rules that an AI/ML model learned from historical data may no longer hold in the current data, and the model may need to be updated to adapt to the changes. Concept drifts are intrinsically related to

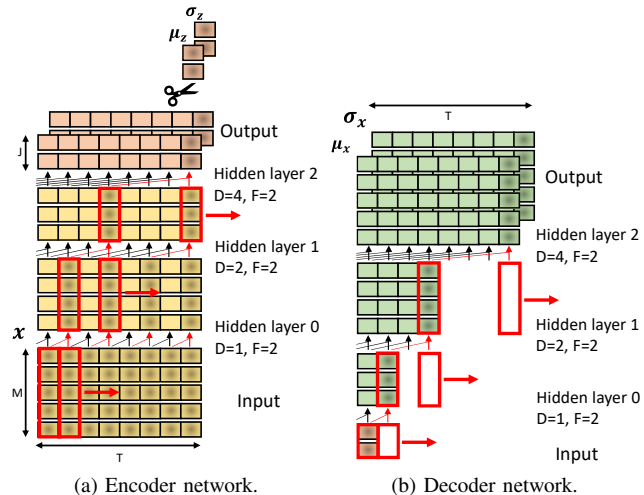


Fig. 1. *DC-VAE* encoder/decoder architecture using causal dilated convolutions, implemented through a stack of 1D convolutional layers.

catastrophic forgetting, which occurs when an AI/ML model trained on a set of tasks or data samples forgets previously learned information when learning new tasks or samples. Both problems require methods to adapt to changing data distributions, by retraining the underlying models.

We resort to a *continual learning* paradigm [6] to address the continual model adaptation and retraining of *DC-VAE*. Continual learning enables a model to learn from a stream of evolving data, without forgetting previously learned knowledge. We leverage the generative AI properties of the underlying VAE model to remember past data. By conception, once the encoder-decoder VAE model has been trained, the decoding function is capable to synthesize new samples mimicking the characteristics of the MTS training datasets, using as input only Gaussian noise.

We combine *DC-VAE* and its generative decoder into GenDeX, an approach to continual learning for anomaly detection in MTS network measurements. In a nutshell, when *DC-VAE* is confronted with concept drifts, or it is applied to a new MTS dataset – e.g., measurements collected at a different network – GenDeX uses the previously trained decoder to synthesize past MTS measurements, and combines them with the new MTS data to retrain the underlying VAE model. GenDeX follows a *Deep Generative Replay* (DGR) [7] paradigm for continual learning, where a generative model produces synthetic data which replays old memories during training, augmenting the heterogeneity and expressiveness of the retraining.




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

In this paper, we study the generative capabilities of *DC-VAE*, describing the underlying architecture and its adaptation to make it operate as a synthetic MTS generator. This paper is a continuation of our initial work on continual learning for anomaly detection in multivariate time-series data [8].

II. *DC-VAE* AND GENDEX

MTS data is generally processed through sliding windows, condensing the information of the most recent T measurements. We define x as a matrix in $\mathbb{R}^{M \times T}$, where M is the number of variables in the MTS process. As depicted in Fig. 1, for a given input x , the trained VAE model produces two different predictions, μ_x and σ_x – matrices in $\mathbb{R}^{M \times T}$, corresponding to the parameterization of the probability distribution which better represents the given input. 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 would not deviate from the mean μ_x more than a specific integer α times the standard deviation σ_x . On the contrary, if the input presents an anomaly, the output would not belong to this normality region.

The main goal of the VAE model is to learn a compressed representation of x in an unsupervised manner. This compressed representation z is referred to as a latent variable, and it is learned by training the VAE to generate data that is similar to the input data. Similar to x , z will also be a sequence of length T , but with a smaller number of dimensions $J < M$, $z \in \mathbb{R}^{J \times T}$. VAEs learn a probabilistic mapping between the input data and its latent variable, which allows to generate new data by sampling from the learned latent variable distribution.

We portray *DC-VAE* in a proprietary MTS dataset, corresponding to real measurements collected at an operation mobile ISP. The TELCO dataset corresponds to twelve different time-series TS_1 to TS_{12} , with a temporal granularity of five minutes per sample, collected and manually labeled for a period of seven months, between January 1 and July 31, 2021. Fig. 2 present *DC-VAE* predictions, using a sliding-window of length $T = 512$ samples, corresponding to roughly two days of past measurements. For each of the displayed time-series TS_i ,

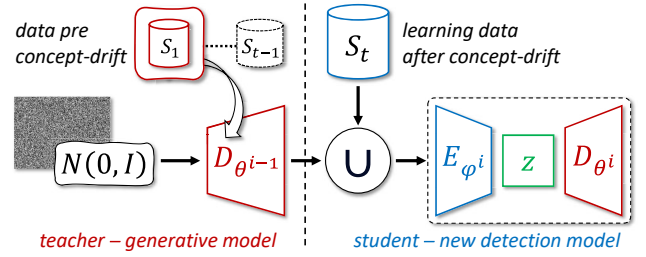


Fig. 3. The GenDeX generative replay approach. At time t , a concept drift significantly modifying the underlying distribution of S_t triggers a model retraining event i .


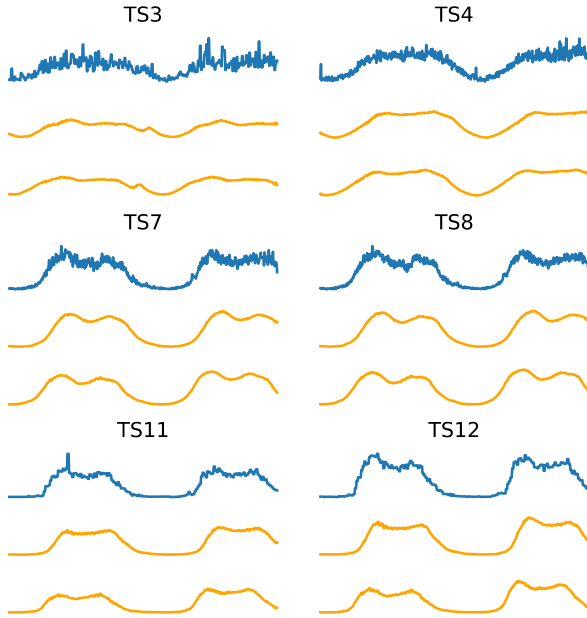


Fig. 4. GenDeX latent space representation. Latent space z with $J = 4$. The colors correspond to the hours of the day.


its real value x_i , along with the outputs of the VAE μ_{x_i} and σ_{x_i} , are reported. The detection of anomalies with different nature is also depicted.

We explore an approach to cope with concept drifts, in particular exploiting the generative nature of the *DC-VAE* model for continual learning. In a continual learning framework, we assume a continually evolving stream of data, represented as a sequence of subsets S_j , each characterized by a specific underlying distribution. We define a sequence of λ_∞ subsets $S_1, \dots, S_{\lambda_\infty}$ sequentially arriving, and assume access to only the data in current subset S_t , with $t \leq \lambda_\infty$. We consider a concept drift occurring at time t , and thus, assume that the underlying distributions of S_1, \dots, S_{t-1} are similar among them, but significantly different from S_t . An initial *DC-VAE* model is trained using S_1 data, which performs accurately till time t . We refer to this model as *DC-VAE* $_0 = \{q_\phi^0, p_\theta^0\} = \{E_\phi^0, D_\theta^0\}$, where E and D represent the encoding and decoding functions, respectively.

GenDeX follows the principles behind DGR to adapt *DC-VAE* $_0$ to the new data S_t , without forgetting the parameterization learned from S_1 , valid for S_1, \dots, S_{t-1} . Fig. 3 explains the GenDeX approach. The decoding function D_θ^0 acts as generator, and it is used to synthesize a new dataset $F_{1 \rightarrow (t-1)}$ out of Gaussian noise, which mimics former training examples in S_1 and its underlying distribution. We say D_θ^0



(a) Synthetic MTS samples generated through GenDeX.



(b) Histograms of generated samples (μ_x).

Fig. 5. For each time-series in TELCO, two examples of time-series generated from noise are depicted. The trend is perfectly captured by the synthetically generated examples. The histograms of samples (μ_x) generated from noise considers the same number of samples as those in the validation set.

acts as the *teacher* model. Then, the new *student* model $DC-VAE_1$ is trained on joint synthetic data F and new data S_t . This approach is model-agnostic and overcomes catastrophic forgetting, as the updated model $DC-VAE_1$ is now capable to handle pre- and post-concept drift data distributions.

III. EXPLORING GENDEX GENERATIVE AI

We assess the generative properties of $DC-VAE$, firstly by analyzing the latent space generated by the encoding function E_ϕ , and then by exploring the generative capabilities of the generative model as represented by the trained decoding function D_θ . The dimension of the latent space in a VAE model is one of the hyper-parameters to define during model evaluation. These dimensions are restricted by the dimensions of the input samples x space, as for the model to only capture the relevant information or energy of the samples, there must be a dimension reduction. By conception and underlying modeling hypothesis, the distribution of the samples z living in the latent space must be a normal distribution with zero mean and an identity covariance matrix. This is enforced during training, by minimizing the standard ELBO loss function [4], consisting of a reconstruction loss (auto-encoding), and a regularization term imposing $z \sim \mathcal{N}(\mu_z, \sigma_z^2)$.

To evaluate the behavior of the encoder E_ϕ , a representation of the latent space is shown for a trained $DC-VAE$ model, using TELCO data. Fig. 4 depicts the resulting latent representation, projecting on each bi-dimensional combination of dimensions $z[i]$, and taking $J = 4$ in this case. Each point corresponds to the projection of a sample from the validation set. with different colors representing a different hour of the day. The distribution of samples in z bi-dimensional projections does

look very close to a zero-one normal distribution. It is certainly centered at zero, and the highest concentration of samples is in the range $[-3, 3]$. If we consider the bi-dimensional latent space $\{z[2], z[3]\}$, we observe how each hour of the day maps to a different angular area in the data distribution. Under this setup, it is enough to feed the decoder D_θ with samples drawn from a zero-one normal distribution to generate synthetic MTS samples out of noise, controlling timing by sampling clockwise.

Fig. 5(a) shows two examples per selected TELCO time-series generated out of noise, along with real time-series included in the original validation set, for two days worth of MTS data. The trend of the time-series is perfectly captured by the synthetically generated examples, with the paramount advantage of these being synthetically generated by D_θ . The MTS process is properly generated, despite having different types of behavior and variability. To evaluate the generative power of GenDeX more broadly, we generate the same number of samples as those in the validation set for each of the selected time-series, and compare them with the real time-series values in the validation set. Fig. 5(b) reports the distribution of the generated and real values, in the form of a histogram. Each pair of distributions strongly overlap, especially for non-spiky values. Time-series TS_3 shows a more variable behavior, which cannot be fully reproduced by the generated baseline, as shown in the corresponding histogram. Recall that we are using GenDeX to track the form and trend of the time-series, by generating μ_x , which cannot capture spiky behaviors. Indeed, we are interested in continually adapting the baseline of the MTS process for anomaly detection, to enable a proper detection of deviations from this baseline.

ACKNOWLEDGMENT

This work has been partially supported by the Austrian FFG ICT-of-the-Future project *DynAISEC – Adaptive AI/ML for Dynamic Cybersecurity Systems* – ID 887504, and by the Uruguayan CSIC project with reference CSIC-I+D-22520220100371UD *Generalization and Domain Adaptation in Time-Series Anomaly Detection*. Gastón García was supported by the ANII scholarship POS-FMV-2020-1-1009239, and by CSIC, under program *Movilidad e Intercambios Académicos 2022*.

REFERENCES

- [1] A. Blázquez-García, A. Conde, U. Mori, and J. A. Lozano, “A review on outlier/anomaly detection in time series data,” *ACM Comput. Surv.*, vol. 54, no. 3, Apr. 2021.
- [2] G. Pang, C. Shen, L. Cao, and A. V. D. Hengel, “Deep learning for anomaly detection: A review,” *ACM Comput. Surv.*, vol. 54, no. 2, Mar. 2021.
- [3] G. García González, S. Martínez Tagliafico, A. Fernández, G. Gómez, J. Acuña, and P. Casas, “DC-VAE, Fine-grained Anomaly Detection in Multivariate Time-Series with Dilated Convolutions and Variational Auto Encoders,” in *2022 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, 2022, pp. 287–293.
- [4] D. P. Kingma and M. Welling, “Auto-encoding Variational Bayes,” *CoRR*, vol. abs/1312.6114, 2013. [Online]. Available: <https://arxiv.org/abs/1312.6114>
- [5] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “Wavenet: A generative model for raw audio,” *arXiv preprint arXiv:1609.03499*, 2016.
- [6] J. Kirkpatrick, R. Pascanu, N. C. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, and R. Hadsell, “Overcoming Catastrophic Forgetting in Neural Networks,” *CoRR*, vol. abs/1612.00796, 2016. [Online]. Available: <http://arxiv.org/abs/1612.00796>
- [7] H. Shin, J. K. Lee, J. Kim, and J. Kim, “Continual Learning with Deep Generative Replay,” *CoRR*, vol. abs/1705.08690, 2017. [Online]. Available: <http://arxiv.org/abs/1705.08690>
- [8] G. G. González, P. Casas, A. Fernández, and G. Gómez, “Steps towards Continual Learning in Multivariate Time-Series Anomaly Detection Using Variational Autoencoders,” in *Proceedings of the 22nd ACM Internet Measurement Conference*, ser. IMC ’22. New York, NY, USA: Association for Computing Machinery, 2022, p. 774–775. [Online]. Available: <https://doi.org/10.1145/3517745.3563033>