

Towards Foundation Auto-Encoders for Time-Series Anomaly Detection

Gastón García González, (†) Pedro Casas, Emilio Martínez, Alicia Fernández
IIE-FING, Universidad de la República; (†) AIT Austrian Institute of Technology
gastong@fing.edu.uy, pedro.casas@ait.ac.at, emartinez@fing.edu.uy, alicia@fing.edu.uy

Abstract

We investigate a novel approach to time-series modeling, inspired by the successes of large pretrained foundation models. We introduce FAE (Foundation Auto-Encoders), a foundation generative-AI model for anomaly detection in time-series data, based on Variational Auto-Encoders (VAEs). By foundation, we mean a model pretrained on massive amounts of time-series data which can learn complex temporal patterns useful for accurate modeling, forecasting, and detection of anomalies on previously unseen datasets. FAE leverages VAEs and Dilated Convolutional Neural Networks (DCNNs) to build a generic model for univariate time-series modeling, which could eventually perform properly in out-of-the-box, zero-shot anomaly detection applications. We introduce the main concepts of FAE, and present preliminary results in different multi-dimensional time-series datasets from various domains, including a real dataset from an operational mobile ISP, and the well known KDD 2021 Anomaly Detection dataset.

Keywords

Time-Series Data, Anomaly Detection, VAE, Foundation Models

1 Introduction

While the literature offers a plethora of traditional statistical models for anomaly detection in time-series data, they often struggle with the non-stationary, non-linear, and noisy nature of real data, leading to suboptimal predictions. In recent years, there has been a surge in the adoption of modern deep learning-based approaches for time-series anomaly detection [28, 31], owing to their ability to handle complex dependencies and generate realistic data sequences. Generative AI methodologies, in particular, have gained attention for their performance in time-series modeling [11, 25, 37].

In this paper, we focus on devising a Generative AI model capable of matching or even surpassing the performance of conventional time-series modeling methods without the need for training on the specific target dataset - a concept known as Zero-Shot Learning (ZSL). ZSL is a problem setup in deep learning where, at test time, a learner observes samples from classes which were not observed during training, and needs to predict the class that they belong to. We therefore investigate if a model pretrained on multiple time-series data can learn temporal patterns useful for accurate forecasting on previously unseen time-series. For doing so, we use as starting point the DC-VAE model [13], a deep-learning-based, unsupervised, and multivariate approach to real-time anomaly detection in MTS data, based on popular Variational Auto-Encoders (VAEs) [23].

VAEs are generative AI models that learn the underlying distribution of the data and can generate new samples from this distribution. In the context of time-series data, VAEs can capture latent representations of temporal patterns and generate sequences that exhibit

similar characteristics, making them powerful for generalization and ZSL. VAEs learn a low-dimensional latent space representation of the input data, which captures the underlying structure of the data in a compressed form. By learning meaningful representations, VAEs can generalize well to unseen data points that lie within the same distribution as the training data, supporting generalization to new instances. As generative models, VAEs can generate new samples from the learned latent space distribution, potentially enabling ZSL, as the model can produce samples that belong to unseen classes or categories without explicitly training on them. By sampling from the latent space, VAEs can generate diverse and realistic data points even for classes not present in the training set.

We introduce and investigate FAE (Foundation Auto-Encoders), a foundation generative-AI model for anomaly detection in time-series data, based on VAEs. FAE uses DC-VAE’s network architecture [13], originally designed for multivariate anomaly detection. In particular, it leverages VAEs and Dilated Convolutional Neural Networks (DCNNs) to build a generic model for *univariate time-series modeling*, which could eventually perform properly in out-of-the-box, zero-shot anomaly detection applications. The reasons for moving from multivariate to univariate time-series analysis are twofold: from an architectural point of view, we become independent of the spatial dimensionality of a MTS dataset - i.e., we fix the spatial input dimensionality to one - and can therefore apply exactly the same architecture without any modifications; from an analytics perspective, while a univariate model is at a disadvantage compared to a multivariate model - i.e., it loses access to cross-correlational information, which we have shown might be critical for better data modeling [13, 14] - the univariate version puts the focus exclusively on the temporal behavior of the data, which is exactly the target of the generalization we are looking for - we want a model that generalizes across the temporal dimension and not necessarily the spatial one, which varies among different problems.

We introduce the main concepts and ideas behind FAE, along with its network architecture, and present preliminary results in the analysis of two different datasets: (1) a multi-dimensional network monitoring dataset - TELCO [15], collected from an operational mobile Internet Service Provider (ISP), recently released openly to the research community; and (2) a large-span dataset of 250 multi-domain time series released for the KDD2021 yearly data contest [22]. The remainder of the paper is organized as follows: Section 2 presents an overview of the related work. In Section 3 we describe the FAE model and its underlying network architecture. Section 4 reports the preliminary results obtained with FAE on the analysis of time-series from TELCO and KDD2021, with a particular focus on generalization and zero-shot modeling. Discussion on the potential of FAE, along with its limitations, is presented in Section 5. Finally, Section 6 concludes the paper.




Figure 1: FAE’s encoder/decoder architecture using causal dilated convolutions, implemented through a stack of 1D convolutional layers.

2 Related Work

The diversity of data characteristics and types of anomalies results in a lack of universal anomaly detection models. Modern approaches to time-series anomaly detection based on deep learning technology have flourished in recent years [28, 31]. Due to their data-driven nature and achieved performance in multiple domains, generative models such as VAEs and Generative Adversarial Networks (GANs) [19] have gained relevance in the anomaly detection field [5, 9, 11, 17, 26, 38, 39].

VAEs [8, 23, 24] represent a powerful and widely-used class of models to learn complex data distributions. Unlike GANs, a potential limitation of VAEs is the prior assumption that latent sample representations are independent and identically distributed. While this is the most common assumption followed in the literature, there is ongoing research on the benefits of accounting for covariances between samples in time to improve model performance [3, 10, 18, 29]. For example, while the original work [23] assumes that the prior over the parameters and latent variables are centered isotropic Gaussian and the true posteriors are approximately Gaussian with approximately diagonal covariance, [10] proposes an approximation capturing temporal correlations, by considering a Gaussian process prior in the latent space.

Modeling data sequences through a combination of variational inference and deep learning architectures has been vastly researched in other domains in recent years, mostly by extending VAEs to Recurrent Neural Networks (RNNs), with architectures such as STORN [1], VRNN [6], and Bi-LSTM [32] among others. Convolutional layers with dilation have been also incorporated into some of these approaches [25, 36], allowing to speed up the training process based on the possibilities of parallelization offered by these architectures. One of these approaches using Dilated Convolutional Neural Networks as the encoder-decoder architecture for VAEs is our DC-VAE model [12, 13].

Transformer-based models [33] are gaining popularity in recent years for time-series analysis, given their remarkable performance in large-scale settings, such as long sequence time-series forecasting (LSTF). LSTF requires capturing long-range dependencies between input and output efficiently. Earlier examples include the TFT interpretable model [27] and the MQTransformer model [4]. The Informer model [40] introduced Transformers for long sequence forecasting through sparse self-attention mechanisms. This concept

has since been further refined through various forms of inductive bias and attention mechanisms in models like the Autoformer [34] and the FEDformer [41].

Finally, there is a recent surge in papers targeting the conception of foundation models for time-series data, capable of generating accurate predictions for diverse datasets not seen during training. The underlying concept of these models is to rely on highly expressive, large-scale architectures which are trained on millions or billions of time-series data points, coming from very diverse domains and having high heterogeneity in terms of temporal behaviors and characteristics. TimeGPT-1 [16], PromptCast [35], LLMTime [20], TimesFM [7], Lag-Llama [30], and Time-LLM [21] are all examples of novel foundation models for time-series forecasting, which target a ZSL application. FAE follows exactly this concept, but using a much smaller and simpler architecture. While this adds certain limitations in terms of expressiveness and therefore generalization capabilities, it also opens the door to the exploration of other venues, such as the combined utilization of smaller foundation models in the form of ensembles, in combination with domain detection strategies.

3 FAE Model and Network Architecture

Time-series are generally processed through sliding windows, condensing the information of the most recent T measurements. We define X as the input vector in $\mathbb{R}^{1 \times T}$. For a given input X , the trained VAE model produces two different predictions, μ_X and σ_X – vectors in $\mathbb{R}^{1 \times T}$, 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. 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.

Figure 1 depicts the encoder/decoder architecture used in FAE, which is an adaptation of DC-VAE’s architecture, for the case of univariate time-series analysis. The FAE model functions as a univariate model trained on various series within a system simultaneously, treating them as distinct classes of series. Similar to the original DC-VAE version, FAE allows for monitoring of all time-series within a MTS process using a single model, albeit analyzing one time-series at a time. The architecture, based on dilated convolutional neural networks (DCNNs), is capable to exploit the temporal dependence of values for longer sequences. The main difference with DC-VAE is that the new architecture has to accommodate univariate input samples $X \in \mathbb{R}^{1 \times T}$, rather than multivariate ones. To maintain the concept of compression – i.e., the dimension of the latent space Z has to be lower than the input dimension of X – the latent space in FAE reduces dimensionality along the temporal

dimension; in DC-VAE, the dimensionality reduction operates in the spatial dimension.

Using DCNNs forces to keep the sequence length T at the output of each hidden layer – referred to as H – thus $H \in \mathbb{R}^{U \times T}$, where U is the number of filters in the layer. As shown in Figure 1, the dilation of each layer at the encoder increases exponentially as the network deepens, ensuring that each time t of the output at the final layer has information from all the previous times of X up to t , i.e., $[X_0, X_1, \dots, X_t]$. This means that the last sample of the output at position $T - 1$ contains information from the entire input sequence X . Then, two layers in parallel with J filters of size one are applied at the output of the last hidden layer, bringing the output to a J -dimensional latent space. The output of the encoder results from keeping only the values at time $T - 1$ at the output of these filters, resulting in vectors $\mu_Z, \sigma_Z \in \mathbb{R}^{J \times 1}$, which define the distribution in the latent space of the inputs X , where $J < T$. Using the reparameterization trick [23], the latent vector $Z \in \mathbb{R}^{J \times 1}$ is generated, corresponding to the encoding of observation X , and is then fed into the decoder. The decoder remains the same as in DC-VAE, which is symmetric with respect to the encoder. However, given that its input requires a sequence of T values and not a single one, the input to the decoder is generated by repeating the vector Z for a total of T times, obtaining an input sample $\in \mathbb{R}^{J \times T}$. As a result, FAE's decoder extracts information from the same latent vector Z for each time t , to generate the output parameters $\mu_X, \sigma_X \in \mathbb{R}^{1 \times T}$, which are used to evaluate deviations from the input observation X . If the FAE model was trained (mainly) with data describing the normal behavior of the analyzed time-series, then the value of a non-anomalous sample X_t at time t would not deviate from the predicted mean μ_{X_t} more than a specific integer α times the standard deviation σ_{X_t} . On the contrary, if the sample is anomalous, it would not belong to the region determined by the predicted mean and standard deviation.

In terms of size of the architecture, an interesting characteristic of FAE is that its structure and number of layers is defined by the length T of the sliding window. In particular, the number of hidden layers N and the length of filters F are related through the dilation factor $d = F^h$ of the DCNNs, which grows exponentially with the layer depth $n \in [0, N - 1]$. Subsequently, N is the minimum value that verifies: $T \leq 2 * F^{N-1}$. In the architectural example (cf. Figure 1), the window length is $T = 8$ and the filter length is $F = 2$, and the target is achieved by taking $N = 3$ hidden layers. This direct relationship between T and the network architecture has a strong practical impact, making it easy to construct the encoder/decoder based on the desired temporal-depth of the analysis. As a final reference of architectural complexity, for a relatively small FAE architecture, using $T = 256$ samples (less than one day of samples, at a 5' sampling-rate), the network exposes roughly half a million free parameters to train. Training FAE with a sufficiently large and heterogeneous training set, comprising multiple time-series of different characteristics, enhances its capability to generalize to unseen data and eventually to different domains.

Table 1: Grid of hyperparameters used in the model calibration.

Hyperparameter	Grid Search Ranges	Best
T - sequence length	$\{128 - 512\}$, step=32	256
J - latent dimension	$\{16 - T/4\}$, step=16	48
γ - learning rate	$\{1e^{-5} - 5e^{-4}\}$	$6e^{-5}$
m - mini-batch size	$\{16 - 96\}$, step=16	32
U - number of filters	$\{16 - 128\}$, step=16	128

4 FAE Time-Series Prediction in the Practice

4.1 Application of FAE in TELCO

We experiment with FAE in the analysis of an open MTS dataset arising from the monitoring of an operational mobile ISP, consisting of time-series with different structural properties. Referred to as the TELCO dataset [15], this *large-scale* – about 750 thousand samples, *long time-span* – seven months' worth of measurements (January 1st to July 31st, 2021) collected at a five-minutes scale, *multi-dimensional* – twelve different time-series, network monitoring dataset includes ground-truth labels for anomalous events at each individual time-series, manually labeled by the experts of the network operation center (NOC) managing the mobile ISP. The twelve time-series are typical data monitored in a mobile ISP, including volume of data traffic, number of SMS, number and amount of prepaid data transfer fees, number and cost of calls, etc.

In this paper we focus on a more qualitative analysis of FAE's performance, focusing on its ability to properly track and reconstruct the different TELCO time-series, and the KDD2021 time-series in section 4.6. Figure 2 depicts a one-month example from the complete TELCO MTS dataset. Different time-series expose different behaviors, e.g., some of them are noisier (TS₃ and TS₉), others have lower dynamic ranges (TS₁), and some others show a smoother evolution (TS₂). All time-series exhibit daily seasonality, but some behave differently on weekends compared to workdays, while others show monthly trends either ascending or descending.

4.2 Hyperparameter Search and Training

One of the most important aspects when working with deep learning models is the search of model and training hyperparameters, along with the subsequent training of the model. Table 1 shows the grid used for the hyperparameter search, as well as the best values (smallest validation loss), identified by Tree-structured Parzen Estimator (TPE) search [2]. In total, 50 attempts were tested on the grid. In the table, T corresponds to the sequence length, and J is the dimensionality of the latent space. Training hyperparameters include the learning rate γ and the mini-batch size m . Finally, U is the number of filters for each hidden convolutional layer, which together with the number of layers and the input and output dimensions define the size of the architecture in terms of the number of trainable parameters p . Considering the five minutes sampling rate of the time-series, the selected sequence length of $T = 256$ samples corresponds to a time window of 21hs and 20 minutes. The exact number of trainable parameters in the identified architecture is $p = 483.840$.

We split the full, 7-months dataset in three independent, time-ordered sub-sets, using measurements from January to March for model training (3 months), April for model validation (1 month),




Figure 2: TELCO time-series, for one month worth of data (March 2021), sampled at a five minutes rate.




Figure 3: Predictions with fully-trained FAE and multivariate DC-VAE in two days of TELCO samples.

and May to July for testing purposes (3 months). One of the disadvantages of the FAE model as compared to the multivariate DC-VAE model is the training time, and hence the time required for hyperparameter search. The FAE model requires approximately five times more training time than DC-VAE.

4.3 FAE Modeling Performance

We evaluate the prediction performance of FAE in samples from the testing set, considering training on the full three months of data, for the 12 time-series, i.e., more than 300.000 samples. Figure 3(a) depicts the resulting predictions μ_{X_t} and σ_{X_t} for two days of testing samples X_t from June 2021, for four representative time-series, including TS₁, TS₄, TS₈, and TS₁₂. To add more variability, we consider a working day (Friday 4th) and a weekend day (Saturday 5th). FAE can properly track different types of behavior in the time-series, including the strong seasonal daily component, but also the operation during workdays and weekends, clearly visible in TS₁₂. Interesting to note is how different periods of time-series variability




Figure 4: Latent space representation – temporal evolution.

result in more or less tight normal-operation regions estimated by FAE, as defined by σ_{X_t} . Figure 3(b) depicts the predictions obtained by the former DC-VAE multivariate model in the same four time-series, using a different time period in April – in this case from the validation test – from Friday 16th till Saturday 17th. Results are similar, but in particular for TS₁₂, DC-VAE can better capture the drop observed on Saturday evening, exploiting the strong spatial correlation observed on Saturday between TS₁₂, TS₁₁, and both TS₁ and TS₂ (cf. Figure 2). Nevertheless, note that FAE predictions are slightly better than DC-VAE's for TS₈ and TS₁₂ on Friday.

To better understand the modeling capabilities of FAE, we focus on the analysis of the latent space Z , for the different time-series and the different times of the analysis. Recall that the latent space set for




Figure 5: Latent space representation, (a) per different time-series (TS₁, TS₄, TS₈, TS₁₂), and (b,c) specifically for TS₁ and TS₄ in a temporal basis, considering workdays (purple) and weekends (yellow).




Figure 6: Latent space representation for TS₁₂, in a daily basis – from day 1 in purple to day 31 in yellow, for a full month.

F AE in this analysis is $J = 48$; to easily visualize Z as a two- or three-dimensional space, we apply standard PCA analysis, and study the top-two and top-three principal components $Z_{PCi}, \dots, i=1,2,3$. Figure 4 shows the latent representation of each sample X_t for a single day, for all the 12 time-series, depicting the encoded samples in colors, each color representing a different three-hour period of the day. Plotting the first two principal components shows that each hour period maps to a certain position in the latent space, and interestingly, the direction mirrors the progression of hours on a clock, ordered continuously by hour of the day.

Figure 5(a) depicts now the first three principal components, displaying the previously analyzed four time-series TS₁, TS₄, TS₈, and TS₁₂, independently. FAE maps each time-series to a different region of the latent space, showing it can properly differentiate among different time-series characteristics. Closely located samples from different time-series exhibit similar behaviors in the time-series space. For example, time-series TS₁ and TS₁₂ are closer to the center of the latent space, and both exhibit a similar behavior (cf. Figure 2), with marked differences between weekends and workdays. On the

other hand, time-series TS₄ and TS₈ have a similar temporal behavior without marked variations between workdays and weekends, and are located together and farther from the center of Z .

The difference between time-series for workdays and weekends is further explored in Figures 5(b,c), where workday samples are displayed in purple color, and weekends in yellow, for (b) time-series TS₁ and (c) time-series TS₄. The plots include two spheres with radius one and two, reflecting the expected Gaussian distribution of the latent space. The difference between workdays and weekends are clear for TS₁, with weekends clustering closer to the center (smaller dynamic range, cf. Figure 2), and workdays located closer at the sphere borders (bigger dynamic range). This workdays/weekends difference is not observed for TS₄. Note in Figure 2 how time-series TS₁₁ and TS₁₂ exhibit a downtrend behavior during the month, which can also be observed in the latent-space. Figure 6 depicts the latent representation of time-series TS₁₂, where days of the month are differentiated by color, from day 1 in purple to day 31 in yellow. As the month goes by, the representation in the latent space moves from the outside borders closer towards the center.

Finally, Figure 7 reports the prediction results, along with flagged anomalies, for both DC-VAE (left) and FAE (right) for the complete 12 time-series in TELCO. Modeling behavior and detection performance are comparable between both models.

To wrap-up these preliminary evaluations, we observe how FAE can properly capture and differentiate among the different temporal behaviors present in the time-series used for training, suggesting a sufficiently expressive model and architecture to model a large and heterogeneous dataset of time-series. In addition, the visual analysis of the latent representations in FAE evidences how VAEs – despite their generative nature – are rather transparent in their operation and behavior, making interpretation and analysis simpler and more human-friendly. This is indeed a strong advantage of VAEs as a powerful yet explainable generative AI model, as compared to modern generative AI approaches, which operate in a more black-box manner.




Figure 7: FAE vs DC-VAE in TELCO.




Figure 8: FAE vs DC-VAE in TELCO2.

4.4 FAE vs DC-VAE in TELCO2

What happens when we modify the application domain? In Figure 8, we apply the pretrained DC-VAE and FAE models in a modified version of TELCO, corresponding to newer data captured in 2024, with modifications in some of the time-series. We refer to this new dataset as TELCO2. On the left side, we depict the modeling and detection performance of DC-VAE, and FAE on the rightside. Both models are pretrained in TELCO, and then applied on TELCO2. In this scenario, it is clear to see the flexibility and advantages of FAE over a multivariate version of the model, which would need re-calibration to cope with the slight yet present variations in the

domain of the dataset. In particular, the spatial modification of correlations and time-series' characteristics make it challenging for DC-VAE to properly adjust to this domain modification.

4.5 Zero-shot Modeling Behavior

We investigate now the performance of FAE in a zero-shot setting, testing the model for time-series not seen at training time. We focus the analysis on TS_{12} , due to its combined seasonality and particular temporal trend, as well as its strong correlated behavior to TS_{11} (cf. Figure 2). We train FAE on three different training datasets from TELCO, using the same 3/1/3-months temporal split for training/validation/testing, but considering a different number of time-series TS_i . The first model uses all 12 time-series – we refer to it as *full-FAE*; the second model considers a zero-shot setting for TS_{12} , with a training dataset which includes time-series TS_1 to TS_{11} , leaving out all samples from TS_{12} ; given the strong temporal correlation between TS_{12} to TS_{11} , we also train a third model leaving out all samples from TS_{11} and TS_{12} , i.e., training on time-series TS_1 to TS_{10} .

The full FAE model mimics a situation where we pretrain with a sufficiently large and heterogeneous dataset which covers the statistical behavior of the downstream data – i.e., a model that *has seen it all*. The other two models mimic two different levels of zero-shot learning: the former represents a pure zero-shot setting for TS_{12} , where the pretrained model has nevertheless observed a similar statistical behavior in a different time-series, i.e., TS_{11} – in particular, it has seen both the seasonality and the monthly trend behaviors; the latter represents a more challenging setting, where the pretrained model has not seen the monthly trend behavior, which is not present in TS_1 to TS_{10} .

Figure 9 presents the prediction performance of the three models, when applied to two weeks of TS_{12} samples, from May 5 to May 19, 2021. In Figure 9(a), the modeling performance for full-FAE is optimal, as it can properly track the different behaviors and patterns in the time-series, similarly to Figure 3(a). A similar performance is observed in Figure 9(b) for the second model, which learns the characteristics of TS_{12} at training time, from TS_{11} . Not surprisingly, the performance of the third model in Figure 9(c) is significantly worse than for the other two models, given the lack of a similar temporal pattern in the training data. To some extent, there is an identification with the patterns observed in time-series TS_1 – note how the daily sharp peaks are exacerbated – which is coherent with their close representations in the latent space (cf. Figure 5(a)). Nevertheless, it somehow manages to capture and track the monthly downtrend, even without previous evidence of it.

4.6 Application of FAE in KDD2021

In this last evaluation, we extend the application of FAE to the KDD2021 dataset, consisting of 250 different time-series arising from many different domains. The dataset is part of the larger UCR Time Series Classification Archive (https://www.cs.ucr.edu/~eamonn/time_series_data_2018/). Within the dataset, each time-series comes in two parts, a train partition and a test partition, which are generally of different sizes. For consistency with TELCO, we take 12 different time series from KDD2021 and apply FAE on




Figure 9: Zero-shot modeling experimentation, predicting TS₁₂ for two weeks in the testing dataset (May 2021). (a) FAE is trained on the full, 12 time-series training set – modeling performance is optimal. (b) FAE is trained on 11 time-series, leaving out TS₁₂ – performance remains almost unchanged. (c) FAE is trained on 10 time-series, leaving out TS₁₁ and TS₁₂ – modeling performance is impacted.

them, training on the train partition and reporting testing results. We refer to this dataset as KDD2021-12.

Figure 10 depicts the predictions obtained with FAE in KDD2021-12. Tracking performance is highly accurate, in particular when considering the anomaly detection application which FAE targets, where the main goal is not to fully reconstruct the output, but rather to provide a mean and variance values where samples are expected to be contained. A couple of observations are worth considering: first, there is a much higher heterogeneity in KDD2021-12 as compared to TELCO, yet FAE manages to properly track the expected normal behavior of the time-series. Second, for some of the KDD2021-12 time series – e.g., TS₁₀, there are long periods of inactivity (constant value) followed by a very spiky behavior; in those cases, note how FAE provides a proper prediction contained in the forecasted variance, which covers the actual time-series values. Finally, Figure 11 depicts the resulting latent space for the KDD2021-12 time-series, which evidences a more complex mix to model, evidenced by clouds of dots from different time-series mixed together; nevertheless, there is a clear structure in the data – clusters with single time-series or with time-series with similar behavior, which again confirms the principled concepts employed by FAE when it comes to the application of VAEs to understand the modeling approach.

5 Discussion and Limitations of FAE

Selecting VAEs for our foundation model exploration has both benefits and limitations, which we briefly discuss next. On the benefits side, we have shown how easy it is to explore and interpret the functioning of the encoding and the behavior of the encoded time-series in the latent space, making the model transparent and easy to tame, particularly for training. While VAEs may struggle with capturing long-range dependencies in the data, we have shown how the integration of DCNNs as part of the encoding/decoding




Figure 10: FAE applied in 12 time-series from KDD2021.

networks enables tracking multiple different temporal behaviors in the time-series, from seasonality to long-term trends.

FAE shows potential to be a strong foundation model for time-series analysis, but so far, we have only trained and tested the model with a rather small number of datasets from a limited number of different domains, and thus still require further assessment. Indeed, thorough testing and validation would be necessary to assess FAE’s performance and generalization capabilities in different scenarios. Ultimately, the effectiveness of FAE as a foundation model would depend on its performance in various real-world scenarios and its ability to generalize to different datasets and different domains.

While powerful, VAEs may have limitations in terms of expressiveness when tasked with learning and mapping a large-scale number of highly heterogeneous time-series data. VAEs operate under the assumption of a latent variable space with a simple distribution, which may not always capture the intricate and diverse characteristics of more complex and diverse time-series data.

We note that the performance of FAE in generalization and zero-shot learning tasks can be affected by factors such as the complexity of the data, the dimensionality of the latent space, the choice of the encoder/decoder architecture, and the quality and diversity of the training data. Additionally, while VAEs can capture global structure in the data distribution, they may not always capture fine-grained details or handle complex data distributions as effectively as other generative models. Finally, in terms of scalability, FAE’s performance may degrade with extremely large datasets or highly heterogeneous data, as the model complexity may need to increase significantly.

6 Concluding Remarks

We have introduced FAE, a novel approach for time-series modeling, motivated by the performance realized by large pretrained foundation models in different domains. FAE targets the detection of anomalies in univariate time-series data, leveraging VAEs and DCNNs to pretrain on large-scale, heterogeneous time-series datasets, potentially enabling to properly model and track a baseline for normal operation, even on unseen datasets.




Figure 11: FAE’s latent space for the 12 KDD2021 time-series.

The preliminary assessment of FAE’s performance has shown promising results. In particular, we have provided evidence of FAE’s capabilities to capture and distinguish various temporal behaviors within the training time-series, demonstrating a promising capacity to model large and heterogeneous datasets effectively. The interpretability of FAE’s latent representations showcased VAEs’ transparency in operation, facilitating simpler analysis and interpretation compared to black-box generative AI models.

Our exploration extended to the zero-shot learning scenario, where FAE’s performance on unseen time-series was assessed. We tested FAE in three settings, ranging from optimal modeling performance to more challenging scenarios. While FAE performs properly in tracking different behaviors and patterns in the time-series, even in the absence of previous evidence, there is room for further improvement, especially in capturing previously unseen temporal trends.

These initial findings underscore FAE’s potential as a feasible foundation model for time-series analysis. However, it is essential to note that our evaluation was limited to a rather small number of datasets from a limited number of different domains. As part of our ongoing work, we are focusing on comprehensively testing FAE’s performance and generalization capabilities across diverse scenarios, using much larger and heterogeneous time-series datasets for training, considering more heterogeneous domains.

Acknowledgment

This work has been partially supported by the FWF Austrian Science Fund, Project I-6653 *GRAPHS4SEC*, 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*.

References

- [1] Justin Bayer and Christian Osendorfer. 2014. Learning stochastic recurrent networks. *arXiv preprint arXiv:1411.7610* (2014).
- [2] James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. 2011. Algorithms for Hyper-parameter Optimization. *Advances in neural information processing systems* 24 (2011).
- [3] Francesco Paolo Casale, Adrian V. Dalca, Luca Saglietti, Jennifer Listgarten, and Nicoló Fusi. 2018. Gaussian Process Prior Variational Autoencoders. In *Advances*

- in *Neural Information Processing Systems*.
- [4] Kevin C. Chen, Lee Dicker, Carson Eisenach, and Dhruv Madeka. 2022. MQTransformer: Multi-horizon Forecasts with Context Dependent Attention and Optimal Bregman Volatility. In *KDD 2022 Workshop on Mining and Learning from Time Series – Deep Forecasting: Models, Interpretability, and Applications*.
- [5] Run-Qing Chen, Guang-Hui Shi, Wanlei Zhao, and Chang-Hui Liang. 2021. A joint model for IT operation series prediction and anomaly detection. *Neurocomputing* 448 (2021), 130–139.
- [6] Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville, and Yoshua Bengio. 2015. A recurrent latent variable model for sequential data. *Advances in Neural Information Processing Systems* 28 (2015).
- [7] Abhimanyu Das, Weihao Kong, Rajat Sen, and Yichen Zhou. 2024. A Decoder-only Foundation Model for Time-series Forecasting. *arXiv:2310.10688 [cs.CL]*
- [8] Carl Doersch. 2016. Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908* (2016).
- [9] Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. 2016. Adversarial feature learning. *arXiv preprint arXiv:1605.09782* (2016).
- [10] Vincent Fortuin, Dmitry Baranchuk, Gunnar Rätsch, and Stephan Mandt. 2020. GP-VAE: Deep Probabilistic Time Series Imputation. In *International conference on artificial intelligence and statistics*. PMLR, 1651–1661.
- [11] Gastón García González, Pedro Casas, Alicia Fernández, and Gabriel Gómez. 2021. On the Usage of Generative Models for Network Anomaly Detection in Multivariate Time-Series. *SIGMETRICS Perform. Eval. Rev.* 48, 4 (may 2021), 49–52. <https://doi.org/10.1145/3466826.3466843>
- [12] Gastón García González, Sergio Martínez Tagliafico, Alicia Fernández, Gabriel Gómez, José Acuña, and Pedro Casas. 2022. 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)*. 287–293.
- [13] Gastón García González, Sergio Martínez Tagliafico, Alicia Fernández, Gabriel Gómez, José Acuña, and Pedro Casas. 2023. One Model to Find Them All – Deep Learning for Multivariate Time-Series Anomaly Detection in Mobile Network Data. *IEEE Transactions on Network and Service Management* (2023), 1–1. <https://doi.org/10.1109/TNSM.2023.3340146>
- [14] Gastón García González, Pedro Casas, and Alicia Fernández. 2023. Fake it till you Detect it: Continual Anomaly Detection in Multivariate Time-Series using Generative AI. In *2023 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. 558–566. <https://doi.org/10.1109/EuroSPW59978.2023.00068>
- [15] Gastón García González, Sergio Martínez Tagliafico, Alicia Fernández, Gabriel Gómez, José Acuña, and Pedro Casas. 2023. TELCO – a new Multivariate Time-Series Dataset for Anomaly Detection in Mobile Networks. <https://doi.org/10.21227/skpg-0539>
- [16] Azul Garza and Max Mergenthaler-Canseco. 2023. TimeGPT-1. *arXiv:2310.03589 [cs.LG]*
- [17] Alexander Geiger, Dongyu Liu, Sarah Alnegheimish, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni. 2020. TadGAN: Time series anomaly detection using generative adversarial networks. In *2020 IEEE International Conference on Big Data (Big Data)*. IEEE, 33–43.
- [18] Laurent Girin, Fanny Roche, Thomas Hueber, and Simon Leglaive. 2019. Notes on the use of variational autoencoders for speech and audio spectrogram modeling. In *DAFx 2019-22nd International Conference on Digital Audio Effects*. 1–8.
- [19] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, Vol. 27.
- [20] Nate Gruver, Marc Finzi, Shikai Qiu, and Andrew Gordon Wilson. 2023. Large Language Models Are Zero-Shot Time Series Forecasters. *arXiv:2310.07820 [cs.LG]*
- [21] Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Y Zhang, Xiaoming Shi, Pin-Yu Chen, Yuxuan Liang, Yuan-Fang Li, Shirui Pan, and Qingsong Wen. 2024. Time-LLM: Time series forecasting by reprogramming large language models. In *International Conference on Learning Representations (ICLR)*.
- [22] E. Keogh, T. Dutta Roy, U. Naik, and A. Agrawal. 2021. Multi-dataset Time-Series Anomaly Detection Competition, SIGKDD 2021. <https://competitions.hexagonml.com/practice/competition/39/>
- [23] Diederik P Kingma and Max Welling. 2013. Auto-encoding Variational Bayes. *CoRR* abs/1312.6114 (2013). [arXiv:1312.6114](https://arxiv.org/abs/1312.6114)
- [24] Diederik P Kingma and Max Welling. 2019. An introduction to variational autoencoders. *arXiv preprint arXiv:1906.02691* (2019).
- [25] Guokun Lai, Bohan Li, Guoqing Zheng, and Yiming Yang. 2018. Stochastic WaveNet: A Generative Latent Variable Model for Sequential Data. *arXiv preprint arXiv:1806.06116* (2018).
- [26] 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 *International Conference on Artificial Neural Networks*. Springer, 703–716.
- [27] Bryan Lim, Sercan Arik, Nicolas Loeff, and Tomas Pfister. 2021. Temporal Fusion Transformers for Interpretable Multi-horizon Time Series Forecasting. *International Journal of Forecasting* 37, 4 (2021), 1748–1764.

- [28] 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.
- [29] Siddharth Ramchandran, Gleb Tikhonov, Kalle Kujanpää, Miika Koskinen, and Harri Lähdesmäki. 2021. Longitudinal variational autoencoder. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 3898–3906.
- [30] Kashif Rasul, Arjun Ashok, Andrew Robert Williams, Hena Ghonia, Rishika Bhagwatkar, Arian Khorasani, Mohammad Javad Darvishi Bayazi, George Adamopoulos, Roland Riachi, Nadhir Hassen, Marin Biloš, Sahil Garg, Anderson Schneider, Nicolas Chapados, Alexandre Drouin, Valentina Zantedeschi, Yuriy Nevmyvaka, and Irina Rish. 2024. Lag-Llama: Towards Foundation Models for Probabilistic Time Series Forecasting. [arXiv:2310.08278](https://arxiv.org/abs/2310.08278) [cs.LG]
- [31] Sebastian Schmidl, Phillip Wenig, and Thorsten Papenbrock. 2022. Anomaly detection in Time Series: a Comprehensive Evaluation. *Proc. VLDB Endow.* 15, 9 (may 2022), 1779–1797. <https://doi.org/10.14778/3538598.3538602>
- [32] Samira Shabanian, Devansh Arpit, Adam Trischler, and Yoshua Bengio. 2017. Variational bi-LSTMs. *arXiv preprint arXiv:1711.05717* (2017).
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, Vol. 30.
- [34] Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. 2021. Autoformer: Decomposition Transformers with Auto-Correlation for Long-Term Series Forecasting. *CoRR abs/2106.13008* (2021). [arXiv:2106.13008](https://arxiv.org/abs/2106.13008) <https://arxiv.org/abs/2106.13008>
- [35] Hao Xue and Flora D. Salim. 2023. PromptCast: A New Prompt-Based Learning Paradigm for Time Series Forecasting. *IEEE Transactions on Knowledge and Data Engineering* (2023), 1–14. <https://doi.org/10.1109/TKDE.2023.3342137>
- [36] Zichao Yang, Zhiting Hu, Ruslan Salakhutdinov, and Taylor Berg-Kirkpatrick. 2017. Improved Variational Autoencoders for Text Modeling using Dilated Convolutions. In *International conference on machine learning*. PMLR, 3881–3890.
- [37] Jinsung Yoon, Daniel Jarrett, and Mihaela van der Schaar. 2019. Time-series Generative Adversarial Networks. In *Advances in Neural Information Processing Systems*, Vol. 32.
- [38] Sultan Zavrak and Murat Iskefiyeli. 2020. Anomaly-based intrusion detection from network flow features using variational autoencoder. *IEEE Access* 8 (2020), 108346–108358.
- [39] Houssam Zenati, Chuan Sheng Foo, Bruno Lecouat, Gaurav Manek, and Vijay Ramaseshan Chandrasekhar. 2018. Efficient GAN-based anomaly detection. *arXiv preprint arXiv:1802.06222* (2018).
- [40] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. 2021. Informer: Beyond Efficient Transformer for Long Sequence Time-Series Forecasting. In *Thirty-Fifth AAAI Conference on Artificial Intelligence*. AAAI Press, 11106–11115. <https://doi.org/10.1609/AAAI.V35I12.17325>
- [41] Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. 2022. FEDformer: Frequency Enhanced Decomposed Transformer for Long-term Series Forecasting. *CoRR abs/2201.12740* (2022). [arXiv:2201.12740](https://arxiv.org/abs/2201.12740) <https://arxiv.org/abs/2201.12740>