
Decoupling Local and Global Representations of Time Series

Sana Tonekaboni¹
University of Toronto

Chun-Liang Li
Google

Sercan Ö. Arık
Google

Anna Goldenberg
University of Toronto

Tomas Pfister
Google

Abstract

Real-world time series data are often generated from several sources of variation. Learning representations that capture the factors contributing to this variability enables a better understanding of the data via its underlying generative process and improves performance on downstream machine learning tasks. This paper proposes a novel generative approach for learning representations for the global and local factors of variation in time series. The local representation of each sample models non-stationarity over time with a stochastic process prior, and the global representation of the sample encodes the time-independent characteristics. To encourage decoupling between the representations, we introduce counterfactual regularization that minimizes the mutual information between the two variables. In experiments, we demonstrate successful recovery of the true local and global variability factors on simulated data, and show that representations learned using our method yield superior performance on downstream tasks on real-world datasets. We believe that the proposed way of defining representations is beneficial for data modelling and yields better insights into the complexity of real-world data.¹

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

⁰¹Work done while at Google

¹The implementation of our framework and the experiments can be found at: https://github.com/googleinterns/local_global_ts_representation

1 Introduction

Learning high-quality representations for complex data such as time series helps better understanding the underlying generative process of the data (Ghahramani, 2015) and has a substantial impact on the performance of downstream machine learning (ML) tasks (Bengio et al., 2013). Condensing information into expressive lower-dimensional representations can improve interpretability (Bai et al., 2018) and lead to better generalization and domain adaptation (Chen et al., 2012). One challenge for downstream ML tasks for time-series data is that collecting labels can be expensive, as most real-world tasks may involve human expert domain knowledge. Also, the underlying state of the data may evolve, rendering human understanding of the data more difficult. Unsupervised representation learning would be of extortionate value in such cases, providing a powerful tool to uncover the underlying state and summarize the complex data into informative general-purpose representations.

Observational time series are often generated from different underlying factors of variation that can be identified through the representations. Some of these factors present the global attributes of a sample and are unique to each individual, while others are descriptive of the variations of the underlying state over time. For instance, consider a medical diagnosis application: the observed patient physiological signals are a product of the individual’s attributes such as gender, age and pre-existing medical conditions, and the medical treatment they received over time. These hidden factors could influence the time series in different ways. Uncovering and decoupling these factors is a fundamental challenge in representation learning that will lead to substantial improvement in data understanding as well as the performance of ML models that utilize the learned representations.

In this paper, we introduce an unsupervised representation learning method for time series that decouples

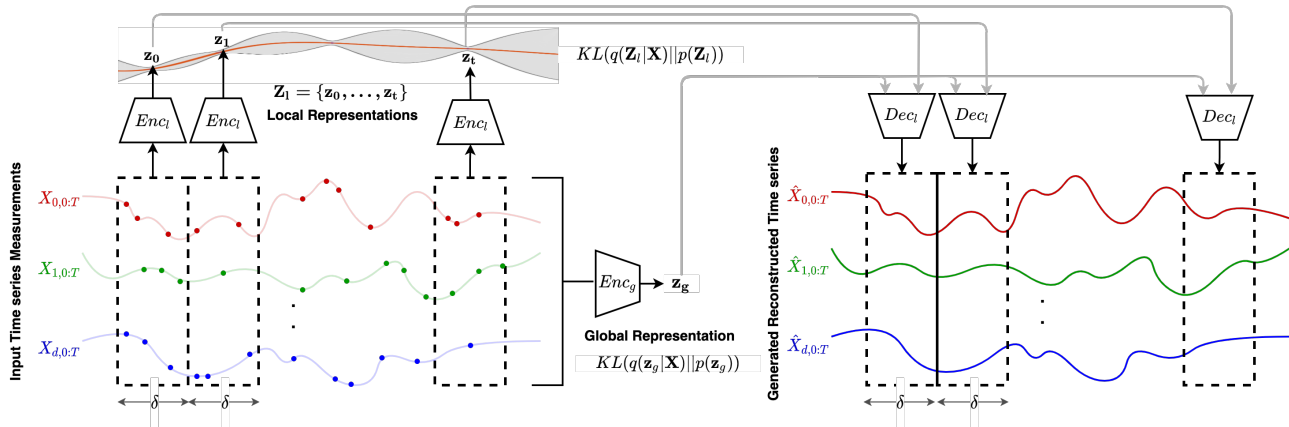


Figure 1: Overview of our method for learning global and local representations. The proposed method learns the distribution of the local representations of windows of samples over time using the local encoder Enc_l . The global encoder Enc_g models the posterior of the global representation, and the decoder Dec generates the time series using samples from the posterior of the local and global representations.

the global and local properties of the signal into separate representations. Our method is based on generative modelling and assumes that each time series sample is generated from two underlying factors: a global and local representation. The global representations are unique to each sample and represent static time-independent characteristics. The local representations encode the dynamic underlying state of windows of the time series sample as they evolve over time, taking into account the non-stationarity of the samples. We use variational approximation to model the posterior of the two sets of representations and model the temporal dynamics using a prior sampled from a Gaussian Process (GP). To ensure that the two sets of representations are decoupled and they model distinct data characteristics, we introduce a counterfactual regularization with the goal of mutual information minimization to disentangle the representations. Fig. 1 overviews our framework. Decoupling local and global representations using our proposed method provides the following important benefits:

1. Exploiting global patterns in the data and coupling them with local variations serves as an inductive bias and improves the downstream tasks from time series **forecasting** performance to **classification** accuracy.
2. By disentangling the factors of variation, we can represent the underlying information in a more **efficient** way by encoding the necessary information for target tasks with more compact representations.
3. Knowing the various factors of variation in generation of a time series helps identify and disentangle

the underlying explanatory factors and results in more **interpretable** representations.

4. Having this information also allows one or the other representations to be used flexibly for downstream tasks depending on which property is appropriate for a specific use case. For instance, global representations can help us **identify sub-groups** with similar properties.
5. Using our approach, we can **generate counterfactual instances** of data with controlled factors of variation, utilizing the global and local representations. This enables generating different manifestations of events like disease progression in individuals with different characteristics.

2 Related Work

The performance of many ML models relies heavily on the quality of the data representations (Bengio and LeCun, 2007). This applies to all data types, but it is vital for complex ones such as time series that can be high-dimensional, high-frequency and non-stationary (Yang and Wu, 2006; Långkvist et al., 2014). Due to the difficulties in labelling time-series data, unsupervised approaches are often preferred in such settings. They include different categories of methods that are based on reconstruction (Yuan et al., 2019; Fortuin et al., 2018, 2020; Chorowski et al., 2019), clustering (Ma et al., 2019; Lei et al., 2019), contrastive objectives (Oord et al., 2018; Franceschi et al., 2019; Hyvarinen et al., 2019; Tonekaboni et al., 2020; Hyvarinen and Morioka, 2016), and others.

While all the methods above have been shown to successfully encode the informative parts of the signal in

a low dimensional representation, improving the interpretability of the representations is still an active area of research. An effort in this direction is disentangling the dimensions of the encoding, which enables representing different factors of variation in independent dimensions. Earlier approaches learn disentangled representations using supervised data (Yang et al., 2015; Tenenbaum and Freeman, 2000), while more recent methods provide unsupervised solutions to tackle this problem using adversarial training (Chen et al., 2016; Kim and Mnih, 2018; Kumar et al., 2018) or regularization of the data distribution (Dezfouli et al., 2019). However, associating these factors with interpretable notions in the data domain remains challenging. A different line of work focuses on decoupling global and local representations of samples into separate representations (or dimensions of representation). This idea has been explored for visual data to separate the factors of variation associated with the labels from other sources of variability. Mathieu et al. (2016) use a conditional generative model with adversarial regularization, while Ma et al. (2020) learn decoupled representations of global and local information of images relying on empirical characteristics of VAE and flow models.

For time series, VAE-based methods have been used to disentangle dynamic and static factors of representations for video and speech data. FHVAE (Hsu et al., 2017) uses a hierarchical VAE model to learn *sequence-dependent variables* (corresponding to the speaker factor) and *sequence-independent variables* (corresponding to the linguistic factors) in modeling speech, but it does not explicitly encourage disentanglement. DSVAE (Yingzhen and Mandt, 2018) formalizes the idea of disentangled representation by explicitly factorizing the posterior to model the static and dynamic factors. Other methods have been proposed to augment DSVAE by explicitly enforcing disentanglement in their objective function. S3VAE (Zhu et al., 2020) introduces additional loss terms that encourage the invariance of the learnt representations by permuting the frame and leveraging external supervision for the motion labels. An additional regularization is imposed by minimizing the mutual information between the two sets of representations. Similar regularizations are used in C-DSVAE (Bai et al., 2021) using contrastive estimation of mutual information. Note that for all the above methods, the focus is on the generated samples and not the quality nor interpretability of the representations. Therefore, the dynamic local representations are not designed to summarize the information of the time series samples over time. Other efforts for disentangling global and local representations are designed to improve specific downstream modeling tasks. For instance, Sen et al. (2019) leverage both local and global patterns to improve forecasting based on matrix

factorization techniques. Similar ideas have also been introduced to improve forecasting performance (Wang et al., 2019; Nguyen and Quanz, 2021). Schulam and Saria (2015) learn population and sub-population parameters to model personalized disease trajectories.

3 Method

In this section, we present the notations used throughout this paper, followed by the problem definition and description of the method.

3.1 Notation

Let $\mathbf{X}^{(i)} \in \mathbf{R}^{d \times T}$ be a multivariate time series sample ($i \in N$), with d input features and T measurements over time. Each time series sample is generated from two latent variables \mathbf{z}_g and \mathbf{Z}_l . The global representation $\mathbf{z}_g^{(i)} \in \mathbf{R}^{d_g}$ is a vector of size d_g that represents the global properties of sample i . The local representation of sample i , $\mathbf{Z}_l^{(i)} = [\mathbf{z}_0^{(i)}; \mathbf{z}_1^{(i)}; \dots; \mathbf{z}_t^{(i)}]$, is composed of a set of vector representations of non-overlapping windows $[\mathbf{X}_0^{(i)}; \mathbf{X}_1^{(i)}; \dots; \mathbf{X}_t^{(i)}]$ of time series $\mathbf{X}^{(i)}$. Each $\mathbf{z}_t^{(i)} \in \mathbf{R}^{d_l}$ is the representation for a window of time series $\mathbf{X}_t^{(i)} \in \mathbf{R}^{d \times \delta}$ of length δ that encodes information for all features within the window. The windows split the sample into $\lceil \frac{T}{\delta} \rceil$ consecutive parts as shown in Fig. 1. The size of the global representations and the local representation is determined by d_g and d_l , respectively. In order to handle missing measurements, each sample has a mask channel $M^{(i)} \in \mathbf{R}^{d \times T}$ to indicate which data points are measured and which ones are missing. Irregularly-sampled time series can also be converted into regularly-sampled signals with the additional mask channel to indicate the measurements. For simplicity, we drop the sample index i in the rest of the paper.

Our probabilistic data generation mechanism assumption is shown in Fig. 2. We model the conditional likelihood distribution of the data as follows: $\mathbf{X} \sim p(\mathbf{X}|\mathbf{Z}_l, \mathbf{z}_g)$. \mathbf{X}_t in Fig. 2 represents a window of time series, and \mathbf{z}_t is the local representation for that window. The local representation of the windows can change over time as the underlying state of the time series changes. The dependencies of these local representations are modeled using a prior sampled from $\mathcal{GP}(m(t), k(t, t'))$. The global representation \mathbf{z}_g is the same for all windows within a sample, and its prior is modelled as a Normal Gaussian distribution.

3.2 Modeling Distributions

As parts of the learning algorithm, we use variational approximations to model the following three distribu-




Figure 2: Graphical model of the generative process. Each window of time series \mathbf{X}_t is generated from the global representation of the samples \mathbf{z}_g and the local representation of the window \mathbf{z}_t . The sample \mathbf{X} is composed of the series of consecutive windows \mathbf{X}_t , and the pset of all local representations $\mathbf{Z}_l = [\mathbf{z}_0; \mathbf{z}_1; \dots; \mathbf{z}_t]$ are sampled from a Gaussian Process prior.

tions:

1. The conditional likelihood distribution of the time series sample $p(\mathbf{X} | \mathbf{Z}_l, \mathbf{z}_g)$ conditioned on the local and global representations. We approximate this using a **Decoder** model ($Dec(\mathbf{Z}_l, \mathbf{z}_g)$).
2. The posterior over the local representations $q(\mathbf{Z}_l | \mathbf{X})$. We approximate this using the **Local Encoder** ($Enc_l(\mathbf{X})$). This encoder slices the time series into consecutive windows and approximates the joint distribution of all the local representations over time.
3. The posterior distribution of the global representations $q(\mathbf{z}_g | \mathbf{X})$. The **Global Encoder** ($Enc_g(\mathbf{X})$) approximates the parameters of this conditional distribution. The encoder can input the entire time series, or any part of it, for estimating the global representations. This is particularly useful in the presence of missing data and allows using the part of the signal with fewer missing observations. To ensure robustness and to encourage the global representation to be constant throughout a sample, the encoder is trained on random sub-windows of the sample.

The local representation should model the temporal behaviour and the underlying time-varying states. The local encoder learns the representation of consecutive windows independently; however, we cannot assume that these representations are *i.i.d.* Casale et al. (2018) show that for video data, by accounting for temporal covariance between the representations, we can learn more general and informative encodings. Similar to ideas presented in Casale et al. (2018) and Fortuin

et al. (2020), we impose the temporal dependencies between the local representations using a prior sampled from a Gaussian Process (GP). GP is a non-parametric Bayesian method well-suited for modelling temporal data that helps with robust modelling even in the presence of uncertainty (Roberts et al., 2013). We model each dimension of the local representations (indexed by $j \in d_l$) independently over time, using unique GPs with different kernel functions. The intuition behind this choice is to decompose the latent representations into dimensions with unique temporal behaviours that are characterized by the covariance structure. E.g., a dimension with a periodic kernel can model the seasonality of the underlying state of the signal. This is an important property since the local representations should model the non-stationarity in the samples. The local encoder $Enc_l(\mathbf{X})$ approximates the posterior of each dimension using a multivariate Gaussian distribution as shown in Eq. 1.

$$q(Z_{0:T,j} | \mathbf{X}) = \mathcal{N}(Z_{0:T,j}; \mu_j, \Sigma_j). \quad (1)$$

Following Fortuin et al. (2020), the precision matrix of the covariance Σ_j^{-1} , is parameterized as a product of bi-diagonal matrices (Eq. 2), where B_j is an upper triangular band matrix that guarantees positive definiteness and symmetry in Σ_j^{-1} . With this sparse estimation of the precision matrix, sampling from the posterior becomes more efficient (Mallik, 2001) and the estimated covariance can still be dense and model long-range dependencies in time.

$$\Sigma_j^{-1} = B_j^T B_j, \{B_j\}_{tt'} = \begin{cases} b_{tt'}^j & \text{if } t' \in \{t, t+1\}, \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

With the parametric approximations for the distributions, we can use the *ELBO* objective in Eq. 3 to train the models.

$$\begin{aligned} \mathcal{L} = & -\mathbb{E}_{\mathbf{Z}_l, \mathbf{z}_g} [\log(p(\mathbf{X} | \mathbf{Z}_l, \mathbf{z}_g))] \\ & + \beta [D_{\text{KL}}(q(\mathbf{Z}_l | \mathbf{X}) \| p(\mathbf{Z}_l)) \\ & + D_{\text{KL}}(q(\mathbf{z}_g | \mathbf{X}) \| p(\mathbf{z}_g))] \end{aligned} \quad (3)$$

The log-likelihood term ensures the generated signals are realistic, and the D_{KL} divergence terms minimize the distance between the estimated distributions and their priors. As mentioned, the prior over the global representations $p(\mathbf{z}_g)$ is assumed to be a standard Gaussian ($\mathcal{N}(0, 1)$) and the prior over the local representations $p(\mathbf{Z}_l)$ is a zero-mean GP prior defined over time ($\mathcal{GP}(0, k_j(t, t'))$). We assume different kernel functions and parameters (i.e. the length scale) for

different dimensions of the latent representations to model various dynamics at multiple time scales. Our framework is compatible with many kernel structures, including but not limited to RBF, Cauchy, and periodic. A list of kernel functions used in our experiments is presented in the Appendix 6.2. Note, the negative log-likelihood is only estimated for the observed measurements to account for the missing values.

Eq. 3, however, does not guarantee all the properties that we expect from the representations. Nothing would prevent all information to flow through the local representation \mathbf{Z}_l , which has a higher encoding capacity. This means that the model can easily converge to a solution where all information about the global behaviour of the signal is encoded in the local representations. As a result, \mathbf{z}_g would become random noise ignored by the decoder and the local representation would no longer represent underlying states, independent of sample variabilities. To address this, we introduce the counterfactual regularization term of the loss, described in the next section.

3.3 Counterfactual Regularization

The issue of information flowing through one set of representations can be prevented if we have labels available for the global sources of variations (Reed et al., 2015; Klys et al., 2018). However, in practice, they are rarely available, and the underlying factors are often unknown. We propose a counterfactual regularization, encouraging \mathbf{z}_g to be informative and the global behaviours to be only encoded in the global representation. This means that variation in the local representation should not change the global identity of the time series and vice versa. For each sample $\mathbf{X}^{(i)}$ during training, we generate a counterfactual sample \mathbf{X}^* with the local representation of $\mathbf{X}^{(i)}$ ($\mathbf{Z}_l^* = \mathbf{Z}_l^{(i)}$), and a random global representations sampled from the prior ($\mathbf{z}_g^* \sim p(\mathbf{z}_g)$), as explained in Fig. 3.

Ideally, this generated counterfactual sample \mathbf{X}^* would have no signs of the global properties of $\mathbf{X}^{(i)}$. If the two representations are independent, $\mathbf{Z}_l^{(i)}$ should not contain any information about $\mathbf{z}_g^{(i)}$; therefore, $\mathbf{z}_g^{(i)}$ should have low likelihood under the estimated posterior distribution of the global representations, conditioned on the counterfactual sample $q(\mathbf{z}_g|\mathbf{X}^*)$. Using the global encoder, we estimate this posterior distribution ($Enc_g(\mathbf{X}^*)$) and encourage the likelihood ratio for \mathbf{z}_g to \mathbf{z}_g^* to be low. Our proposed regularization term is as follows:

$$\mathcal{L}_{reg} = \mathbb{E}_{\mathbf{z}_g, \mathbf{Z}_l} \frac{q(\mathbf{z}_g|\mathbf{X}^*)}{q(\mathbf{z}_g^*|\mathbf{X}^*)}, \quad (4)$$




Figure 3: The counterfactual sample \mathbf{X}^* generation process for the regularization term.

As a result, the final objective becomes:

$$\begin{aligned} \mathcal{L} = & -\mathbb{E}_{\mathbf{Z}_l, \mathbf{z}_g} [\log(p(\mathbf{X}|\mathbf{Z}_l, \mathbf{z}_g))] \\ & + \beta [D_{\text{KL}}(q(\mathbf{Z}_l|\mathbf{X})\|p(\mathbf{Z}_l))] \\ & + D_{\text{KL}}(q(\mathbf{z}_g|\mathbf{X})\|\mathcal{N}(0, 1))] + \lambda \mathcal{L}_{reg}. \end{aligned} \quad (5)$$

Counterfactual regularization and disentanglement. An essential role of counterfactual regularization is encouraging independence implicitly. One way to achieve independence between global and local variables is through minimizing the mutual information $I(\mathbf{z}_g, \mathbf{Z}_l)$ between the two variables. Following Moyer et al. (2018), the mutual information between the two sets of representations can be decomposed as:

$$\begin{aligned} I(\mathbf{z}_g; \mathbf{Z}_l) &= I(\mathbf{z}_g; \mathbf{X}) - I(\mathbf{z}_g; \mathbf{X}|\mathbf{Z}_l) \\ &= I(\mathbf{z}_g; \mathbf{X}) + H(\mathbf{z}_g|\mathbf{X}) - H(\mathbf{z}_g|\mathbf{Z}_l) \\ &= I(\mathbf{z}_g; \mathbf{X}) + H(\mathbf{z}_g|\mathbf{X}) + \mathbb{E}_{\mathbf{z}_g, \mathbf{Z}_l} \log p(\mathbf{z}_g|\mathbf{Z}_l). \end{aligned} \quad (6)$$

The first two terms measure the information captured by the global representation \mathbf{z}_g , which is also considered in the variational autoencoder objective (Eq. 3). Minimizing $I(\mathbf{z}_g; \mathbf{Z}_l)$ can therefore be done by minimizing $\mathbb{E}_{\mathbf{z}_g, \mathbf{Z}_l} \log p(\mathbf{z}_g|\mathbf{Z}_l)$. As we do not have access to the distribution $p(\mathbf{z}_g|\mathbf{Z}_l)$, existing works (Cheng et al., 2020) use an additional network to construct the variational approximation. Instead of introducing an additional network to approximate $p(\mathbf{z}_g|\mathbf{Z}_l)$, which increases training complexity and computation time, we reuse the global encoder for counterfactual regularization for this approximation as follows. Since \mathbf{z}_g^* is sampled from the prior distribution independent of \mathbf{Z}_l , then $q(\mathbf{z}_g|\mathbf{Z}_l) = q(\mathbf{z}_g|\mathbf{Z}_l, \mathbf{z}_g^*) \simeq q(\mathbf{z}_g|\mathbf{X}^*)$ if the decoder $\mathbf{Z}_l, \mathbf{z}_g^* \rightarrow \mathbf{X}^*$ preserves the information. This implies that minimizing (4) implicitly minimizes $I(\mathbf{z}_g; \mathbf{Z}_l)$ and encourages decoupling between the representations.

4 Experiments

Evaluating unsupervised representation learning is challenging due to the lack of well-defined labels for the underlying representations and factors of variation. However, the representations’ generalizability and informativeness can be assessed on different downstream tasks. We present a number of experiments to evaluate the performance of our method in comparison to the following benchmarks: (i) Variational Auto Encoder (VAE) as a standard unsupervised representation learning framework, (ii) GP-VAE (Fortuin et al., 2020), (iii) C-DSVAE (Bai et al., 2021), which separates dynamic and static representations², and (iv) a model trained supervised for the task.

All baselines are trained to learn representations for consecutive windows in the time series sample. For consistency and to allow comparison of performance results, the encoder and decoder architectures of different baselines are kept the same. One of the benefits of decoupling local and global representations is that we can condense the representation into fewer dimensions, especially since the global representation is unique throughout the sample. We have chosen different encoding sizes for our baselines to reflect this advantage better. The performance of all models is compared across multiple evaluation tests using two different time series dataset:

1. Physionet ICU Dataset (Goldberger et al., 2000): A medical time-series dataset of records from 12,000 adult ICU stays in the Intensive Care Unit. The temporal measurements consist of physiological signals and different lab measurements. For each recording, there are general descriptors of the patient (age, type of ICU admission, etc.) as well as labels indicating in-hospital mortality. We use such descriptors as approximates for the global properties of the signals.
2. UCI Beijing Multi-site Air Quality Dataset (Zhang et al., 2017): This dataset includes hourly measurements of multiple air pollutants from 12 nationally-controlled monitoring sites, collected over four years. The measurements are also matched with meteorological data from the nearest weather station. We partition the data such that each time series sample is the pollutant reading for parts of a particular month of the year.

To fully demonstrate the capability of our method, we focus on time series data with non-stationarity. Unlike

²The implementation of this method for time series data doesn’t have the augmentations proposed for video setting because such transformations (cropping, color distortion, etc.) are not defined for time series.

most classification tasks where the time series is windowed to represent samples of different classes, here we are interested in long time series where the behaviour of the signals might change over time. More information on the experiment datasets, cohort selection and processing, are provided in the Appendix 6.1.

4.1 Improving Downstream Prediction Tasks

General representations should capture the important information in the data and can therefore be leveraged for downstream prediction tasks by training simple predictors on them. This approach is commonly used for evaluating the quality of representations (Oord et al., 2018; Franceschi et al., 2019; Fortuin et al., 2020). For the Physionet dataset, we consider the mortality prediction task. As a supervised baseline, we train an end-to-end model that directly uses the time series measurements to predict the risk of in-hospital mortality. For all other baselines, a simple Recurrent Neural Network (RNN) is trained to predict the risk using the representations over time. For the Air Quality dataset, the defined task is to estimate the average daily rain. Similarly, a simple RNN model is used for training this predictor using the representations. For our approach and C-DSVAE where the global and local representations are encoded separately, the global representations are concatenated with the output of the RNN model to estimate the downstream task. Appendix 6.2 provides more details on the architecture of the models used in our experiments.

Table 1 shows the performance of all models on the prediction tasks. For better comparison, we have baselines with different representation dimensionality. The results show that our method outperforms others for the ICU mortality prediction task, with even fewer representation dimensions, and comes second to C-DSVAE for daily rain estimation. For Physionet, we even perform closely to a fully-supervised model. GP-VAE performs better than the regular VAE, as it properly models the correlation between representations of samples over time. As we increase the dimensionality of the representations, this model improves albeit with higher complexity and less interpretability in the encodings. By decoupling the representations, our method achieves superior performance with smaller dimensionality. Lastly, we demonstrate that the counterfactual regularization substantially improves the performance of our method (as shown by the performance results of our method with no regularization).

Model	Dimensions	ICU Mortality Prediction		Average Daily Rain Estimation
		AUPRC	AUROC	Mean Absolute Error
Our method	$8 \times \text{steps} + 8$	0.365 ± 0.092	0.752 ± 0.011	1.824 ± 0.001
Our method - no reg.	$8 \times \text{steps} + 8$	0.238 ± 0.026	0.672 ± 0.010	1.825 ± 0.001
GP-VAE	$8 \times \text{steps}$	0.266 ± 0.034	0.662 ± 0.036	1.824 ± 0.001
GP-VAE	$16 \times \text{steps}$	0.282 ± 0.086	0.699 ± 0.018	1.826 ± 0.001
VAE	$8 \times \text{steps}$	0.157 ± 0.053	0.564 ± 0.044	1.831 ± 0.005
VAE	$16 \times \text{steps}$	0.118 ± 0.001	0.491 ± 0.037	1.840 ± 0.012
C-DSVAE	$8 \times \text{steps} + 8$	0.158 ± 0.005	0.565 ± 0.007	1.806 ± 0.012
Supervised	NA	0.446 ± 0.036	0.802 ± 0.043	0.079 ± 0.001

Table 1: Performance of different representation learning methods on downstream predictive tasks.

Model	Dimensions	Air Quality	Physionet
		Accuracy	Accuracy
Our method	8	57.93 ± 3.53	46.98 ± 3.04
Our method - no reg.	8	38.35 ± 2.67	32.54 ± 0.00
GP-VAE	8	36.73 ± 1.40	42.47 ± 2.02
GP-VAE	16	33.57 ± 1.50	44.67 ± 0.50
VAE	8	27.17 ± 0.03	34.71 ± 0.23
VAE	16	31.20 ± 0.33	35.92 ± 0.38
C-DSVAE	8	47.07 ± 1.20	32.54 ± 0.00
Supervised	NA	62.43 ± 0.54	62.00 ± 2.10

Table 2: Performance of different representation learning methods on identifying similar subgroups of data.

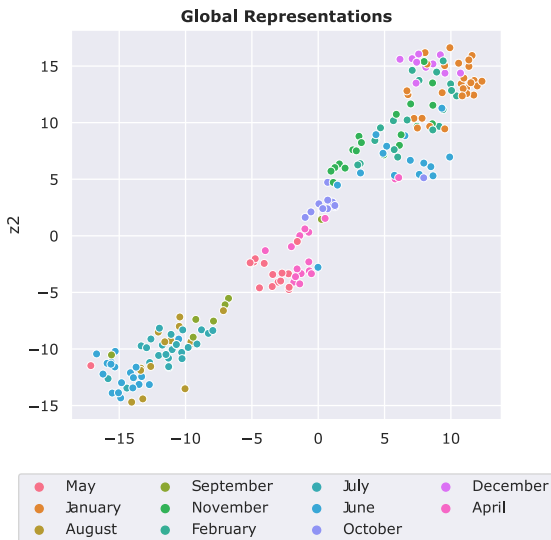


Figure 4: t-SNE visualization of the representations of the Air Quality dataset. Each data point is the global representation of a time series sample and the color indicates the month of year of that measurement.

4.2 Subgroup Identification

In many cases, we are interested in identifying or clustering samples with similar global properties invariant to other factors that can influence the time series trajectory. The global representations provide the information that allows us to identify such subgroups of the data. This is an important application, and as an example of the benefits, earlier work shows that in many applications, performing cluster-specific modelling or prediction can improve the overall performance of ML models (Giordano et al., 2019; Bouveyron and Jacques, 2011). This experiment evaluates how well our global representations identify clusters of similar samples. For the Physionet dataset, we choose the ICU unit type as a proxy label for identifying subgroups because patients that are admitted to these different units have some underlying similarities. For the Air Quality dataset, we are interested in identifying the month of the year that each recording belongs to.

We use a simple MLP classifier to predict the subgroup of each sample using the learned representations. Our method and C-DSVAE define global representations separately, and that is what we use for the training.

Model	Dimensions	Air Quality		Physionet	
		NLL	MSE	NLL	MSE
Our method	$8 \times \text{steps}+8$	1.445 \pm 0.052	0.609 \pm 0.026	2.609 \pm 0.032	1.183 \pm 0.016
Our method - no reg.	$8 \times \text{steps}+8$	3.866 \pm 0.344	0.906 \pm 0.040	3.250 \pm 0.145	1.323 \pm 0.048
GP-VAE	$8 \times \text{steps}$	1.908 \pm 0.114	0.743 \pm 0.048	2.626 \pm 0.063	1.197 \pm 0.031
GP-VAE	$16 \times \text{steps}$	3.080 \pm 0.281	0.941 \pm 0.065	2.833 \pm 0.069	1.295 \pm 0.035

Table 3: Performance of different representation learning methods used for time series forecasting.

For the other baselines, we randomly select the representation of one window for this task as the representations encode both the local and global properties of the sample. Table 2 summarizes the classification performance of all baselines. Our results support the claim that the global representations capture the global characteristics of each sample. We can therefore identify samples with similar characteristics better than other baselines by only using these representations regardless of the changes over time. Fig. 4 also shows the 2-dimensional projection of the global representations of the Air Quality dataset, learned using our method. We can observe the clear distinction between the global properties of measurements from the warmer months of the year and the ones around winter and fall. Even within these categories, the representations seem to cluster the months together.

4.3 More Accurate Forecasting

Forecasting high-dimensional time series is another important application in many domains, like retail and finance. Prior works observe that exploiting global patterns and coupling them with local calibration help prediction performance on many datasets (Sen et al., 2019). Our proposed method, designed in a similar vein for representation learning, can therefore be well-suited to improve forecasting performance. More importantly, since we model the local representations over time using a GP, the conditional distribution over the future local representations can be estimated by conditioning over the observed historical representations (Williams and Rasmussen, 2006); As shown in Eq. 7. Here, Z_i^* represent the local representations, and t^* corresponds to the time steps in the future.

$$\begin{aligned}
Z_i^*|Z_i, t, t^* &\sim \mathcal{N}(\mu(Z_i^*), \text{Cov}(Z_i^*)) \\
\text{Cov}(Z_i^*) &= \mathbf{K}(Z_i^*, Z_i^*) - \mathbf{K}(Z_i^*, Z_i)\mathbf{K}(Z_i, Z_i)^{-1}\mathbf{K}(Z_i, Z_i^*) \\
\mu(Z_i^*) &= \mathbf{K}(Z_i^*, Z_i)\mathbf{K}(Z_i, Z_i)^{-1}Z_i
\end{aligned} \tag{7}$$

We evaluate the performance of our method and the GP-VAE baselines for the time series forecasting, as shown by the results in Table 3. The objective is

to predict two windows, equivalent to 28 observations or 2-days measurements for the Air Quality data and eight measurements in Physionet, using the full history of the signal. While deep learning models are prone to overfitting with increasing forecast horizon especially in non-stationary settings, our probabilistic method achieves good performance in such scenarios by estimating the expected trajectory of the signal using predictions for future local representations. Additional plots showing the forecasting performance are provided in the Appendix 6.3.

4.4 Learning Disentangled Representations

Many applications require generative models that are capable of synthesizing new instances where certain key factors of variation are held constant and other varied (Mathieu et al., 2016). For example, in medical diagnosis, one may wish to model a treatment trajectory on different individuals. Decoupling the global representations from local variations enables such control over the generative process. We have curated a simulated dataset to test this functionality as the evaluation of generated counterfactual samples is difficult on real-world data without the knowledge of the ground truth. This dataset consists of time series samples composed of seasonality, trend, and noise. There are four types of signals in the dataset, with two global classes that determine the trend and intercept and two local classes that determine the seasonal component’s frequency. More information on this dataset is provided in the Appendix 6.1.

Using our approach, we learn one-dimensional local and global representations for each sample. We then generate new samples, controlling the local or global behaviours, and observe the effect of changes on the generated time-series. Fig. 5 shows an example of this experiment. On the top left, we show a sample from the dataset. Over the multiple rows, we generate samples by gradually changing the global representation while keeping the local behaviour the same. We see that this changes the trend and intercept of the signal while maintaining the frequency. This is expected given our knowledge of the generative process of this




Figure 5: Examples of counterfactual signal generation in simulated dataset. The top left plot shows the original time-series in orange. Each row shows generated samples with the same underlying local representations as the original signal, with variable \mathbf{z}_g . Each column represent different \mathbf{Z}_l for the different \mathbf{z}_g values.

data. Then, we keep \mathbf{z}_g constant for each sample and generate samples by changing the local representation in the second column. The change in local representations manifests as a change in the frequency, confirming our hypothesis. Note that some of the samples generated (Fig. 5 middle row) are the ones never seen by the model, but still, their properties are correctly inferred. This experiment shows that we can learn the true underlying factors of data and shows that our generative approach can synthesize realistic samples by changing certain underlying factors of variation.

5 Conclusion

This paper introduces a generative approach for learning global and local representations for time series. We demonstrate the benefits of decoupling these representations for improving downstream performance, efficiency and a better understanding of the underlying generative factors of the data. Decoupling the underlying factors of variation in time series brings us one step closer to a better understanding of the generative process of complex temporal data. As a future direction, we would like to investigate how to associate the representation with concrete notions in the data space for improved interpretability and explore potential applications of counterfactual generation in real-world data.

References

- Bai, J., Wang, W., and Gomes, C. P. (2021). Contrastively disentangled sequential variational autoencoder. *NeurIPS*, 34.
- Bai, T., Zhang, S., Egleston, B. L., and Vucetic, S. (2018). Interpretable representation learning for healthcare via capturing disease progression through time. In *KDD*.
- Bengio, Y., Courville, A., and Vincent, P. (2013). Representation learning: A review and new perspectives. *PAMI*, 35(8):1798–1828.
- Bengio, Y. and LeCun, Y. (2007). Scaling learning algorithms towards AI. In *Large Scale Kernel Machines*. MIT Press.
- Bouveyron, C. and Jacques, J. (2011). Model-based clustering of time series in group-specific functional subspaces. *Advances in Data Analysis and Classification*, 5(4):281–300.
- Casale, F. P., Dalca, A. V., Saglietti, L., Listgarten, J., and Fusi, N. (2018). Gaussian process prior variational autoencoders. In *NeurIPS*.
- Chen, M., Xu, Z., Weinberger, K. Q., and Sha, F. (2012). Marginalized denoising autoencoders for domain adaptation. In *ICML*.
- Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., and Abbeel, P. (2016). Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *NIPS*, pages 2180–2188.
- Cheng, P., Hao, W., Dai, S., Liu, J., Gan, Z., and

- Carin, L. (2020). Club: A contrastive log-ratio upper bound of mutual information. In *ICML*.
- Chorowski, J., Weiss, R. J., Bengio, S., and van den Oord, A. (2019). Unsupervised speech representation learning using wavenet autoencoders. *IEEE Trans audio, speech, and language processing*, 27(12):2041–2053.
- Dezfouli, A., Ashtiani, H., Ghattas, O., Nock, R., Dayan, P., and Ong, C. S. (2019). Disentangled behavioral representations. In *NeurIPS*.
- Fortuin, V., Baranchuk, D., Rätsch, G., and Mandt, S. (2020). Gp-vae: Deep probabilistic time series imputation. In *AISTATS*.
- Fortuin, V., Hüser, M., Locatello, F., Strathmann, H., and Rätsch, G. (2018). Som-vae: Interpretable discrete representation learning on time series. In *ICLR*.
- Franceschi, J.-Y., Dieuleveut, A., and Jaggi, M. (2019). Unsupervised scalable representation learning for multivariate time series. *NeurIPS*.
- Ghahramani, Z. (2015). Probabilistic machine learning and artificial intelligence. *Nature*, 521(7553):452–459.
- Giordano, G., Haberman, S., and Russolillo, M. (2019). Coherent modeling of mortality patterns for age-specific subgroups. *Decisions in Economics and Finance*, 42(1):189–204.
- Goldberger, A. L., Amaral, L. A., Glass, L., Hausdorff, J. M., Ivanov, P. C., Mark, R. G., Mietus, J. E., Moody, G. B., Peng, C.-K., and Stanley, H. E. (2000). Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals. *Circulation*, 101(23):e215–e220.
- Hsu, W.-N., Zhang, Y., and Glass, J. (2017). Unsupervised learning of disentangled and interpretable representations from sequential data. *arXiv preprint arXiv:1709.07902*.
- Hyvarinen, A. and Morioka, H. (2016). Unsupervised feature extraction by time-contrastive learning and nonlinear ica. *NIPS*.
- Hyvarinen, A., Sasaki, H., and Turner, R. (2019). Non-linear ica using auxiliary variables and generalized contrastive learning. In *AISTATS*.
- Johnson, A. E., Dunkley, N., Mayaud, L., Tsanas, A., Kramer, A. A., and Clifford, G. D. (2012). Patient specific predictions in the intensive care unit using a bayesian ensemble. In *Computing in Cardiology (CinC), 2012*, pages 249–252. IEEE.
- Kim, H. and Mnih, A. (2018). Disentangling by factorising. In *ICML*.
- Klys, J., Snell, J., and Zemel, R. (2018). Learning latent subspaces in variational autoencoders. *arXiv preprint arXiv:1812.06190*.
- Kumar, A., Sattigeri, P., and Balakrishnan, A. (2018). Variational inference of disentangled latent concepts from unlabeled observations. In *ICLR*.
- Längkvist, M., Karlsson, L., and Loutfi, A. (2014). A review of unsupervised feature learning and deep learning for time-series modeling. *Pattern Recognition Letters*, 42:11–24.
- Lei, Q., Yi, J., Vaculin, R., Wu, L., and Dhillon, I. S. (2019). Similarity preserving representation learning for time series clustering. In *IJCAI*, volume 19, pages 2845–2851.
- Ma, Q., Zheng, J., Li, S., and Cottrell, G. W. (2019). Learning representations for time series clustering. *NeurIPS*.
- Ma, X., Kong, X., Zhang, S., and Hovy, E. H. (2020). Decoupling global and local representations via invertible generative flows. In *ICLR*.
- Mallik, R. K. (2001). The inverse of a tridiagonal matrix. *Linear Algebra and its Applications*, 325(1-3):109–139.
- Mathieu, M. F., Zhao, J. J., Zhao, J., Ramesh, A., Sprechmann, P., and LeCun, Y. (2016). Disentangling factors of variation in deep representation using adversarial training. *NIPS*.
- Moyer, D., Gao, S., Brekelmans, R., Galstyan, A., and Ver Steeg, G. (2018). Invariant representations without adversarial training. *NIPS*, 31:9084–9093.
- Nguyen, N. and Quanz, B. (2021). Temporal latent auto-encoder: A method for probabilistic multivariate time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 9117–9125.
- Oord, A. v. d., Li, Y., and Vinyals, O. (2018). Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Reed, S. E., Zhang, Y., Zhang, Y., and Lee, H. (2015). Deep visual analogy-making. *NIPS*.
- Roberts, S., Osborne, M., Ebden, M., Reece, S., Gibson, N., and Aigrain, S. (2013). Gaussian processes for time-series modelling. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1984):20110550.
- Schulam, P. and Saria, S. (2015). A framework for individualizing predictions of disease trajectories by exploiting multi-resolution structure. *NIPS*.
- Sen, R., Yu, H.-F., and Dhillon, I. (2019). Think globally, act locally: A deep neural network approach to high-dimensional time series forecasting. *arXiv preprint arXiv:1905.03806*.

- Tenenbaum, J. B. and Freeman, W. T. (2000). Separating style and content with bilinear models. *Neural computation*, 12(6):1247–1283.
- Tonekaboni, S., Eytan, D., and Goldenberg, A. (2020). Unsupervised representation learning for time series with temporal neighborhood coding. In *ICLR*.
- Wang, Y., Smola, A., Maddix, D., Gasthaus, J., Foster, D., and Januschowski, T. (2019). Deep factors for forecasting. In *International conference on machine learning*, pages 6607–6617. PMLR.
- Williams, C. K. and Rasmussen, C. E. (2006). *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA.
- Yang, J., Reed, S. E., Yang, M.-H., and Lee, H. (2015). Weakly-supervised disentangling with recurrent transformations for 3d view synthesis. In *NIPS*.
- Yang, Q. and Wu, X. (2006). 10 challenging problems in data mining research. *International Journal of Information Technology & Decision Making*, 5(04):597–604.
- Yingzhen, L. and Mandt, S. (2018). Disentangled sequential autoencoder. In *International Conference on Machine Learning*, pages 5670–5679. PMLR.
- Yuan, Y., Xun, G., Suo, Q., Jia, K., and Zhang, A. (2019). Wave2vec: Deep representation learning for clinical temporal data. *Neurocomputing*, 324:31–42.
- Zhang, S., Guo, B., Dong, A., He, J., Xu, Z., and Chen, S. X. (2017). Cautionary tales on air-quality improvement in beijing. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 473(2205):20170457.
- Zhu, Y., Min, M. R., Kadav, A., and Graf, H. P. (2020). S3vae: Self-supervised sequential vae for representation disentanglement and data generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6538–6547.

6 SUPPLEMENTARY MATERIAL

6.1 Datasets

This section provides additional details about each of the datasets used in our experiments.

6.1.1 Physionet Dataset

This dataset was made available as part of the PhysioNet Computing in Cardiology Challenge 2012³ with the objective of predicting mortality in ICU patients. The data consists of records from 12,000 ICU stays. All patients were adults who were admitted for a wide variety of reasons to cardiac, medical, surgical, and trauma ICUs. For cohort selection, out of the entire set of features that include multiple lab measurements over time and physiological signals collected at the bedside, we have selected the ones with less than 60 percent missing observations over the entire dataset. The length of all samples is restricted to be between 40 to 80 measurements. Some of the time series processing steps are borrowed from Johnson et al. (2012)⁴ (without any feature extraction for signals).

For all the experiments, we only use the time series measures for training the unsupervised representation learning frameworks, without the static information about each patient. We use the general descriptors as proxy labels for evaluating global representations throughout the experiments.

6.1.2 Air Quality Dataset

The Beijing multi-site Air Quality dataset⁵ includes hourly air pollutants data from 12 nationally-controlled air-quality monitoring sites, collected from the Beijing Municipal Environmental Monitoring Center. The meteorological data in each air-quality site are matched with the nearest weather station from the China Meteorological Administration. The time period is from March 1st, 2013, to February 28th, 2017. We create our dataset by dividing this time series into samples from different stations and of different months of the year. For our experiments, we use the pollutant measurements as our time series, the daily rain amount as a proxy for local behaviour and the station and month of each sample as the global characteristics.

	$Z_l = 1$	$Z_l = 2$
$\mathbf{z}_g = 1$	$\gamma = 0.05$ $a = 1.8$ $c = -1.5$ $b = 40$ $\alpha = 0.5$	$\gamma = 0.05$ $a = 0.8$ $c = -1.5$ $b = 20$ $\alpha = 0.8$
$\mathbf{z}_g = 2$	$\gamma = -0.05$ $a = 1.8$ $c = +1.5$ $b = 40$ $\alpha = 0.5$	$\gamma = -0.05$ $a = 0.8$ $c = +1.5$ $b = 20$ $\alpha = 0.8$

Table 4: Parameters of the simulated dataset

6.1.3 Simulated Dataset

We have created the simulated dataset to assess the quality of the counterfactual generated samples, as the ground-truth generative process of the data is known to us. The dataset consists of a total of 500 samples, with 100 measurements over time. Each sample is a one-dimensional time series and can be decomposed as follows:

$$\mathbf{X}_{(t)} = \alpha \times (\gamma t + a \sin(\frac{bt}{2\pi}) + c) + \epsilon \quad (8)$$

Where $\epsilon \sim \mathcal{N}(0, 0.1)$ is the Gaussian noise, and the rest of the parameters are determined by the global and local representation class of the sample. There are two possible classes defined for each of the local and global representations, resulting in overall four different types of time series samples in the cohort. Table 4 describes the underlying generative process of each sample based on the underlying global and local variable.

6.2 Experiment Setup

For reproducibility purposes, the implementation code for all the experiments is made available. We have also provided additional information about each experiment in this section.

6.2.1 Training our model

For training our method on all the different datasets, we have chosen the list of hyper-parameters, reported in Table 5. The window and the representation sizes are a design choice and depend on the properties of the data. The rest of the parameters are selected based on the model performance on the validation set.

6.2.2 Downstream task

The downstream task test evaluates the generalizability and usability of the learned representations. All baseline methods learn a representation for each window of time series. In addition to the representations over time, our method learns a single representation

³<https://physionet.org/content/challenge-2012/1.0.0/>

⁴<https://github.com/alistairewj/challenge2012>

⁵<https://archive.ics.uci.edu/ml/datasets/Beijing+Multi-Site+Air-Quality+Data>

Parameter	Physionet	Air Quality	Simulation
λ	1.0	0.5	2.0
β	0.1	0.1	0.1
Window size (δ)	4	24	10
Global representation size (d_g)	8	8	1
Local representation size (d_l)	8	8	1
Optimizer	Adam	Adam	Adam
Learning rate	0.001	0.001	0.01
Prior Kernels	RBF, Cauchy	RBF, Cauchy	RBF
Prior Kernel scales	2, 1, 0.5, 0.25	2, 1, 0.5, 0.25	1

Table 5: List of the selected parameters for training our method on different dataset

vector for the global variable. For the mortality prediction task, we use a single layer RNN, followed by a fully connected layer to estimate the risk of mortality. To integrate the global representation, our method concatenates the final hidden state of the RNN with the global representation vector. The designated task for the Air Quality dataset is the prediction of average daily rain. Each local representation encodes a window of 24 samples, equivalent to a day of measurements. A 2-layer MLP uses the local representations concatenated with the global representation to predict this value. A big challenge with the rain prediction is that it is a highly imbalanced regression problem, where on almost 90% of the days, there is no rain, and for the rest, the amount is highly variable. To mitigate this, we use a weighted mean absolute error loss for training, and report the mean squared error as a performance metric to ensure consistency among results.

6.2.3 Subgroup classification

For this test, the goal is to identify similar subgroups using the global representations. For baselines where the two representations are not separated, we randomly select one of the local representations over time to use for the prediction. The global properties are expected to be encoded in the same representation vector in these baselines. Using a two layer MLP and the learned representations, we classify samples into the subgroups defined by out proxy labels (type of ICU for Physionet and month of year for the Air quality dataset).

6.3 Supplementary Plots




Figure 6: Risk of mortality estimation over time, based on the local and global representations (As part of the downstream task test for Physionet data).

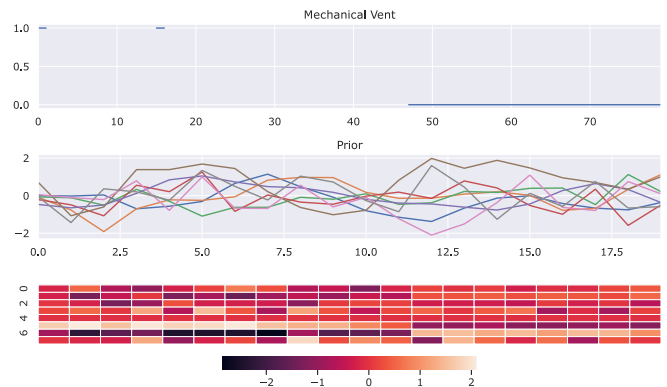


Figure 7: Exploratory analysis of local representations in Physionet data. The bottom heatmap demonstrates the 8 dimensional local representation of windows over time (x-axis). The middle bar shows the priors over all the different dimensions and finally the plot on the top shows the indicator of the mechanical ventilation. In the representation patterns, we can see signals indicating whether a patient is ventilated or not. Note that the ventilation information is never provided as an input to any of the encoders.

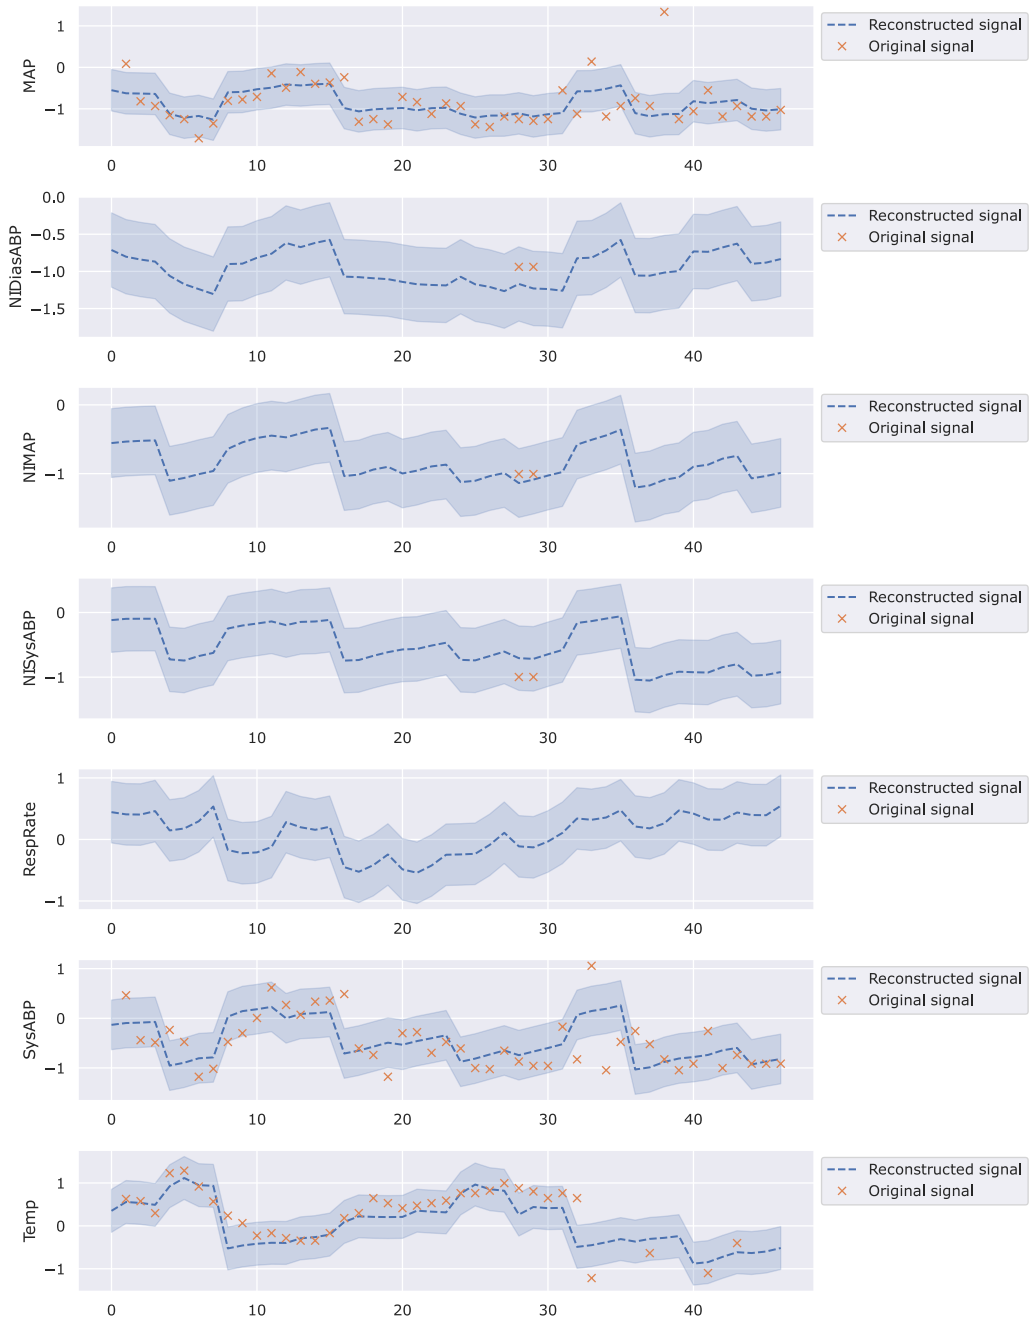


Figure 8: Reconstructed of a Physionet sample with high missing rate




Figure 9: Forecasting pollutant measurements for the Air Quality dataset. The measurements before the vertical red line are observed (shown by the green line), and the forecasted measurements for 2 windows of time series are shown with the blue line, with the green plot demonstrating the expected prediction. The shaded regions indicate one standard deviation around the estimated distribution.