

Anomaly Detection of Time Series With Smoothness-Inducing Sequential Variational Auto-Encoder

Longyuan Li, Junchi Yan[✉], *Member, IEEE*, Haiyang Wang, and Yaohui Jin, *Member, IEEE*

Abstract—Deep generative models have demonstrated their effectiveness in learning latent representation and modeling complex dependencies of time series. In this article, we present a smoothness-inducing sequential variational auto-encoder (VAE) (SISVAE) model for the robust estimation and anomaly detection of multidimensional time series. Our model is based on VAE, and its backbone is fulfilled by a recurrent neural network to capture latent temporal structures of time series for both the generative model and the inference model. Specifically, our model parameterizes mean and variance for each time-stamp with flexible neural networks, resulting in a nonstationary model that can work without the assumption of constant noise as commonly made by existing Markov models. However, such flexibility may cause the model fragile to anomalies. To achieve robust density estimation which can also benefit detection tasks, we propose a smoothness-inducing prior over possible estimations. The proposed prior works as a regularizer that places penalty at nonsmooth reconstructions. Our model is learned efficiently with a novel stochastic gradient variational Bayes estimator. In particular, we study two decision criteria for anomaly detection: reconstruction probability and reconstruction error. We show the effectiveness of our model on both synthetic data sets and public real-world benchmarks.

Index Terms—Anomaly detection, deep generative model, recurrent neural network, time series, variational auto-encoder (VAE).

I. INTRODUCTION

TIME series anomaly detection is an important and challenging problem, and it has been studied over decades across wide application domains, including intelligence transportation [1], health care [2], KPI monitoring [3], web intrusion detection [4], environmental monitoring [5], fault diagnosis [6], and malware detection [7]. Such a broad

application is also reflected in the diversity of formulations and models to time series anomaly detection. In contrast to anomaly detection on static data, one widely accepted idea is that the temporal continuity plays a key role in all formulations and unusual changes in the time-series data are detected and used to model anomalies. For multidimensional correlated time-series data, the correlations make the problem further complex, which highlights the need for robust and flexible models.

Consider a matrix $\mathbf{X} \in \mathbb{R}^{M \times T}$ for a time-series corpus that consists of M correlated streams with T time steps for each stream. On the one hand, there are generally three scenarios of time series anomaly detection depending on the availability of training data and anomaly labels: supervised, semi-supervised, and unsupervised [8]. Anomaly detection tasks without *clean*¹ data or anomaly labels also referred to as unsupervised detection. Unsupervised methods have their advantages that no anomaly labels is needed, which is cost-saving, and they can enjoy better flexibility when the anomaly pattern drifts over time. On the other hand, time series anomaly detection can be, in general, divided into two settings: i) subsequence or whole sequence level anomaly whereby a subsequence $\mathbf{x}_{m,t_1:t_2}$ is labeled as an anomaly; ii) point level anomaly for which a measurement $x_{m,t}$ at time t in sequence m is treated as an anomaly. The former scenario is usually more challenging than the latter. For more comprehensive surveys, readers are referred to [9]–[11].

Technique and methodology to address the above two settings are rather different, and we focus on the popular point-wise anomaly detection for multidimensional time series in this article, which has received wide attention in the literature [5], [12], [13]. In some cases, sequence or subsequence level anomaly detection can be generalized from the point level model. In addition, many practical applications require point level detection results as abnormal signal can quickly disappear in the time series especially when the sampling rate is low. We leave it for future work for generalizing our model to more macro-level anomaly detection tasks. On the other hand, as a lack of labeled anomalies necessitates the use of unsupervised approaches, we focus on unsupervised detection model for its practical applicability. This also aligns with the rising trend of adopting unsupervised learning algorithms for anomaly detection such as one-class SVM [14], GMM [15], tDCGAN [7], VAE [3], and VRNN [16].

¹We use "clean" to describe time-series data that contain only normal patterns.

Manuscript received March 24, 2019; revised October 10, 2019 and February 10, 2020; accepted March 5, 2020. Date of publication April 13, 2020; date of current version March 1, 2021. This work was supported in part by the China Major State Research Development Program under Grant 2018YFC0830400, Grant 2018AAA0100704, and Grant 2016YFB1001003, in part by the NSFC under Grant 61972250 and Grant U19B2035, and in part by STCSM under Grant 18DZ1112300. (Corresponding authors: Junchi Yan; Yaohui Jin.)

Longyuan Li, Haiyang Wang, and Yaohui Jin are with the State Key Laboratory of Advanced Optical Communication System and Network, MoE Key Laboratory of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, Shanghai 200240, China (e-mail: jeffli@sjtu.edu.cn; 0103050180@sjtu.edu.cn; jinyh@sjtu.edu.cn).

Junchi Yan is with the Department of Computer Science and Engineering, MoE Key Laboratory of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, Shanghai 200240, China (e-mail: yanjunchi@sjtu.edu.cn).

Color versions of one or more of the figures in this article are available online at <https://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNNLS.2020.2980749

Different from general time-series modeling aiming to fit with the data curve, anomaly detection often need to recover “normal” patterns robustly in the presence of anomalies. This is challenging since the model is optimized to fit all data points. Previous works [17], [18] show that inductive bias is crucial for generalization when learning deep generative models. Motivated by previous works, we seek to take a Bayesian approach. Specifically, we devise a temporal smoothness prior which can be coherently incorporated in the sequential variational auto-encoder (VAE) learning procedure, under the well-established probabilistic framework. As such, the learning of latent space can be more focused on the “normal” pattern rather than blindly minimizing the reconstruction error between the raw signal and the generated one by the latent variables. Based on the learned generation model, the anomalies can be detected via a certain criterion by comparison of the reconstructed signal and the raw data point.

II. RELATED WORK

This article is devoted to unsupervised point level anomaly detection for multidimensional time series. We review mostly related work by roughly dividing them into two groups: deterministic models and probabilistic methods.

A. Deterministic Models

From the optimization perspective, classical approaches often introduce additional regularization term to account for *temporal smoothness*. The extra regularization enables tradeoff between *fitness* and *smoothness*, making the model less sensitive to anomalies associated with the training data [19], [20]. Representative works include time series de-trending [21], and its extension for multidimensional correlated time series [22]. These methods tend to learn the general trend of time series. However, they often pay less attention to the stochastic nature of time series, as such the modeling capabilities are degraded when the data are noisy and anomaly contaminated.

B. Probabilistic Models

In these methods, the core idea is to learn the marginal likelihood $p_\theta(\mathbf{x})$ of the data-generation process. Then the anomalies are detected by thresholding reconstruction probability, using the techniques, for example, Hidden Markov Models [23], [24] and Matrix Factorizations [25], [26]. However, generative models are often computationally restricted to a simple linear model with conjugate probability distributions due to the intractable integral of marginal likelihood $p_\theta(\mathbf{x}) = \int p_\theta(\mathbf{z}) p_\theta(\mathbf{x}|\mathbf{z}) d\mathbf{z}$. Moreover, limited by the potential intractability, the noise is often assumed constant, making them unsuitable for nonstationary time series.

Recently, the work of Kingma and Welling [27] proposes the Stochastic Gradient Variational Bayes (SGVB) estimator and Auto-encoded Variational Bayes (AEVB) algorithm. One popular technique under this framework is rVAE. By using neural network for parameterizing a flexible inference model $q_\phi(\mathbf{z}|\mathbf{x})$ to approximate posterior $p_\theta(\mathbf{z}|\mathbf{x})$, along with the reparameterization trick, the lower bound \mathcal{L} of marginal likelihood is differentiable. In consequence, the model can be effectively

trained. To a certain extent, the AEVB algorithm liberates the limitations when devising complex probabilistic generative models, especially for deep generative models.

One step further, by taking advantage of the AEVB algorithm, recent studies have introduced deep generative models for anomaly detection. The work [28] uses VAE [27] for anomaly detection on nonsequential data sets. However, vanilla VAE makes i.i.d. assumption among data points $\mathbf{x} = \{\mathbf{x}^{(i)}\}_{i=1}^N$, which is unrealistic for real-world time-series data. To solve this problem, the very recent approach [3] further extends vanilla VAE by preprocessing the raw data with sliding windows, and the windows are assumed to obey i.i.d. Although this approach mitigates the artifact, the temporal structure is still not well modeled. Alternatively, researchers seek to use a deterministic recurrent neural network (RNN) [29] to produce hidden states, then the latent variable is conditioned on the RNN hidden states and the previous latent variables, such that the temporal structure is captured. The resulting models include VRNN [16] and STORN [30], which can be viewed as sequential versions of VAE. However, these models essentially hardly consider the presence of anomalies in their objective functions.

In this article, we propose a smoothness-inducing sequential VAE (SISVAE) model for learning multidimensional correlated time-series data. In particular, we are focused on point level anomaly detection which can be naturally handled by our model, where the overall procedure is illustrated in Fig. 1.

The technical highlights of this article are as follows.

- 1) We devise a Bayesian technique to incorporate the smoothness prior to the learning of a deep generative model for multidimensional time series anomaly detection. Such a design, to our best knowledge, has not been studied in the literature and it inherits merits for the efficiency of the classical optimization model and the uncertainty modeling capability by the deep generative models. Specifically, the proposed variational smoothness regularizer encourages both smooth mean and variance transitions over time. The resulting objective, with the network can be learned by standard backpropagation.
- 2) Different from Markov models that depend on constant noise assumption, our model parameterizes mean and variance individually for each time-stamp using neural networks. This enables the dynamic anomaly detection thresholds adaption according to estimated data noise. This feature is important as the variance of time series may vary over time.
- 3) We perform extensive empirical evaluations on both synthetic data sets and several real-world benchmarks, showing that our model consistently outperforms the state-of-the-art competing models. We also study two decision criteria for performing anomaly detection tasks on a trained SISVAE model: reconstruction probability and reconstruction error.

III. PRELIMINARIES

Before focussing on the technical details of our proposed model, some preliminaries are described as background to facilitate the presentation of our main method, and the glossary for the terms used in this article is shown in Table I.




Fig. 1. Procedure illustration of the proposed model. (a) Input multidimensional (here 4-dim) time series for anomaly detection. (b) Time-varying probability density function estimation using the proposed model—shaded region denote variance levels. (c) Anomalies (red circle) detection by thresholding (a) multidimensional time series with anomalies. (b) Density estimation (shaded area). (c) Anomaly detection (red circle).

TABLE I
GLOSSARY FOR THE TERMS USED IN THIS ARTICLE AND THEIR INTERPRETATION

Terms	Symbol	Meaning
latent variable model	$p_{\theta}(\mathbf{x}, \mathbf{z}) = p_{\theta}(\mathbf{x} \mathbf{z})p(\mathbf{z})$	Description of the data generating process involving unobserved random variable \mathbf{z} .
AEVB algorithm	$\min_{\phi} D_{KL}[q_{\phi}(\mathbf{z} \mathbf{x}) p_{\theta}(\mathbf{z} \mathbf{x})]$	Use a pair of <i>encoder</i> and <i>decoder</i> to learn a complex deep latent variable model.
generative model	$p_{\theta}(\mathbf{x} \mathbf{z})$	A probabilistic mapping from latent variable \mathbf{z} to observation \mathbf{x} , also called <i>decoder</i> .
inference model	$q_{\phi}(\mathbf{z} \mathbf{x})$	An approximation to the intractable true posterior $p_{\theta}(\mathbf{z} \mathbf{x})$, also called <i>encoder</i> .
reparameterization trick	$\mathbf{z} = g_{\phi}(\epsilon, \mathbf{x}),$ $\epsilon \sim p(\epsilon)$	Using a deterministic variable to express random variable $\mathbf{z} = g_{\phi}(\epsilon, \mathbf{x})$, where ϵ is an auxiliary variable with independent marginal $p(\epsilon)$.
latent state	\mathbf{z}_t	latent variable associate with observation \mathbf{x}_t .
latent temporal dependencies	$p(\mathbf{z}_{1:T})$	Latent variables are not independent but exhibits some kinds of structure, such as 1st-order Markov chain.
trending prior	$p(\mathbf{z}_t)$	A time-varying prior distribution for latent state.
non-linear transitions	$p(\mathbf{z}_t \mathbf{z}_{<t})$	Non-linear relationship between latent space and past latent states. It is linear for Kalman filter and Hidden Markov Models.
non-linear emission probability	$p(\mathbf{x}_t \mathbf{z}_t)$	Nonlinear relationship between latent space and observation. It is linear for Kalman filter and Hidden Markov Models.
smoothness prior	$p(\mathbf{f})$	A generic contextual constraint on the true signal based on human knowledge is the smoothness.
non-stationary noise	ϵ_t	The noise model varies over time, in many currently used models, only the mean of Gaussian is modeled with dynamics, while the variance is often set to be constant
anomaly score	\mathbf{A}	Score assigned by the model for each observation, higher score indicates higher probability of anomaly.

A. Problem Formulation

Due to various reasons, such as faulty sensors or extreme events, the real-world time-series data are often contaminated with anomalies. Consider we have time-series data $\mathbf{x}_{1:T} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T)$ contaminated with anomalies, and each $\mathbf{x}_t \in \mathbb{R}^M$ (indicating T -dimensional time series) is modeled by

$$\mathbf{x}_t = \mathbf{f}(t) + \epsilon_t + \mathbf{I}_t \odot \mathbf{e}_t \quad (1)$$

where $\mathbf{f}(t)$ is the true signal, ϵ_t is nonstationary noise of observations, \mathbf{I} is a $M \times T$ binary matrix, which is the indicator matrix of anomalies. The value $\mathbf{I}_{m,t} = 1$ when there is an anomaly at observation $\mathbf{x}_{m,t}$, and \mathbf{e}_t is the anomaly vector, and \odot is element-wise multiplication. Our goal is to recover the true signal $\mathbf{f}(t)$ and detect anomaly \mathbf{I} from $\mathbf{x}_{1:T}$.

B. Smoothness Prior Modeling

For the unsupervised learning problem, to recover true signal $\mathbf{f}(t)$ from anomaly-contaminated data, it is important to impose inductive bias or prior knowledge to the model as suggested in [17] and [18]. For time-series modeling, we first introduce smoothness prior modeling of time series as described in [31]. Consider the time series $y_{1:T}$ (for convenience, we deal with the univariate case), it is assumed to

consist of the sum of a function f and observation noise ϵ

$$y_t = f(t) + \epsilon_t \quad (2)$$

where $f \in \mathcal{F}$ is an unknown smooth function, and $\epsilon \sim \mathcal{N}(0, \sigma^2)$. The problem is to estimate f from the observations. A common approach is to use a parametric model such as polynomial regression or neural network to approximate f . The quality of estimation is dependent upon the appropriateness of the assumed model class. The idea is to balance a tradeoff of goodness-of-fit to the data, and a smoothness criterion. The solution is to minimize the objective with an appropriately chosen smoothness trade-off parameter λ

$$\min_f \sum \{(y_t - f(t))^2 + \lambda (\nabla^k f(t))^2\} \quad (3)$$

where ∇^k is a k th order differential constraint on the solution f , with $\nabla f(t) = f_t - f_{t-1}$, $\nabla^2 = \nabla(\nabla f(t))$, etc.

The solution can be solved from a Bayesian perspective: consider a homoscedastic (variance is constant over time) time series $y_t \sim \mathcal{N}(f(t), \sigma^2)$ with constant variance σ^2 . The solution becomes maximizing

$$\mathcal{L}(f) = \exp\left(-\frac{1}{2\sigma^2} \sum_{t=1}^T (y_t - f(t))^2\right) \exp\left(-\frac{\lambda}{2\sigma^2} \sum_{t=1}^T (\nabla^k f(t))^2\right). \quad (4)$$

Its Bayesian interpretation can be written by

$$\pi(f|y, \lambda, \sigma^2, k) \propto p(y|\sigma^2, f)\pi(f|\lambda, \sigma^2, k) \quad (5)$$

where $\pi(f|\lambda, \sigma^2, k)$ is the prior distribution of f , and $p(y|\sigma^2, f)$ is data distribution conditioned on σ^2 and $\pi(f|y, \lambda, \sigma^2, k)$ is the posterior of f .

C. Deep Generative Models and SGVB Estimator

Consider a data set $\mathbf{x} = \{\mathbf{x}^{(i)}\}_{i=1}^N$ consisting of N i.i.d. data points $\mathbf{x}^{(i)}$. From a latent variable modeling perspective, we assume the data are generated by some random processes that capture the variations in the observed variables \mathbf{x} . Thus, we have the integral of marginal likelihood

$$p_\theta(\mathbf{x}) = \int p_\theta(\mathbf{z})p_\theta(\mathbf{x}|\mathbf{z})d\mathbf{z} \quad (6)$$

where θ is generative model parameters, and $p_\theta(\mathbf{z})$ is the prior distribution over latent variable \mathbf{z} , and $p_\theta(\mathbf{x}|\mathbf{z})$ is the conditional distribution that generates data \mathbf{x} . VAE [27] models the conditional probability $p_\theta(\mathbf{x}|\mathbf{z})$ with a neural network as a flexible approximator. However, introducing a complex nonlinear mapping for the generation process from \mathbf{z} to \mathbf{x} results in intractable inference of posterior $p_\theta(\mathbf{z}|\mathbf{x})$. Hence, VAE uses a variational approximation $q_\phi(\mathbf{z}|\mathbf{x})$ of the posterior, that is, the evidence lower bound (ELBO) of the marginal likelihood as written below

$$\mathcal{L}(\theta, \phi; \mathbf{x}) = -D_{\text{KL}}(q_\phi(\mathbf{z}|\mathbf{x})\|p_\theta(\mathbf{z})) + \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] \quad (7)$$

where the first R.H.S. term is the Kullback–Leibler divergence between two probability distributions Q and P , and ϕ denotes variational parameters. The generative model and inference model are trained jointly by maximizing the ELBO with respect to their parameters. The KL-divergence $D_{\text{KL}}(q_\phi(\mathbf{z}|\mathbf{x})\|p_\theta(\mathbf{z}))$ of (7) can often be integrated analytically. However, the gradient of the second term is problematic because latent variable $\mathbf{z} \sim q_\phi(\mathbf{z}|\mathbf{x})$ is stochastic which blocks backpropagation. The Stochastic Gradient Variational Bayes (SGVB) [32] expresses the random variable \mathbf{z} as a deterministic variable $\mathbf{z} = g_\phi(\boldsymbol{\epsilon}, \mathbf{x})$, where $\boldsymbol{\epsilon}$ is an auxiliary variable with independent marginal distribution $p(\boldsymbol{\epsilon})$. Then we have the SGVB estimator $\tilde{\mathcal{L}}(\theta, \phi; \mathbf{x}) \simeq \mathcal{L}(\theta, \phi; \mathbf{x})$ of the lower bound

$$\tilde{\mathcal{L}}(\theta, \phi; \mathbf{x}) = -D_{\text{KL}}(q_\phi(\mathbf{z}|\mathbf{x})\|p_\theta) + \log p_\theta(\mathbf{x}|\mathbf{z}) \quad (8)$$

where $\mathbf{z} = g_\phi(\boldsymbol{\epsilon}, \mathbf{x})$ and $\boldsymbol{\epsilon} \sim p(\boldsymbol{\epsilon})$.

Learning is performed by optimizing the SVGB estimator with backpropagation. VAE [27] uses centered isotropic multivariate Gaussian $p_\theta(\mathbf{z}) = \mathcal{N}(\mathbf{z}; \mathbf{0}, \mathbf{I})$ as prior of latent variable \mathbf{z} . The variational approximate posterior $q_\phi(\mathbf{z}|\mathbf{x})$ is a multivariate Gaussian with diagonal covariance matrix: $\mathcal{N}(\boldsymbol{\mu}, \text{diag}(\boldsymbol{\sigma}^2))$, and its mean $\boldsymbol{\mu}$ and variance $\boldsymbol{\sigma}^2$ are parameterized by a neural network. For $p_\theta(\mathbf{x}|\mathbf{z})$, a Bernoulli distribution for binary data or multivariate Gaussian distribution (as adopted in our model) for continuous data is often used.

IV. PROPOSED MODEL

We present the proposed SISVAE model for multidimensional time series point anomaly detection. We first describe how the sequential VAE can be employed for time series density estimation, which involves a generative model and inference model to learn the latent variable for data points. Then we devise a smooth prior under the Bayesian framework which leads to a smoothness-induced sequential VAE model. We show the model can be used to detect the anomaly either by absolute reconstruction loss or reconstruction probability.

A. Sequential VAE as Density Estimator of Time Series

For time-series modeling, many prior publications have extended VAE in a sequential manner [16], [33]–[35]. The main idea is to use a deterministic Recurrent Neural Network (RNN) as a backbone, represented by the sequence of RNN hidden states. At each timestamp, the RNN hidden state is conditioned on the previous observation \mathbf{x}_{t-1} and stochastic latent variable \mathbf{z}_{t-1} , such that the model is able to capture sequential features from time-series data.

Let $\mathbf{x}_{1:T} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T)$ denote a multidimensional time series, such as observation data collected by multiple sensors of T time slots. Our goal is to learn the data distribution of the training time series $p(\mathbf{x}_{1:T})$. The model is composed of two parts, the generative model and the inference model, as illustrated in Fig. 2.

1) *Generative Model:* We specify the generative model $p_\theta(\mathbf{x}_{1:T}|\mathbf{z}_{1:T})$ of the following factorization:

$$p_\theta(\mathbf{x}_{\leq T}, \mathbf{z}_{\leq T}) = \prod_{t=1}^T p_\theta(\mathbf{x}_t|\mathbf{z}_{\leq t}, \mathbf{x}_{< t})p_\theta(\mathbf{z}_t|\mathbf{x}_{< t}, \mathbf{z}_{< t}). \quad (9)$$

The first factor of R.H.S of (9) $p_\theta(\mathbf{x}_t|\mathbf{z}_{\leq t}, \mathbf{x}_{< t})$ is the conditional likelihood of a data point given all previous data and latent states. We notate $\hat{\mathbf{x}}_t$ as the random variable for the reconstructed data distribution of observation \mathbf{x}_t , we have

$$p(\hat{\mathbf{x}}_t|\mathbf{z}_t) \sim \mathcal{N}(\boldsymbol{\mu}_{x,t}, \text{diag}(\boldsymbol{\sigma}_{x,t}^2)) \quad (10)$$

where $[\boldsymbol{\mu}_{x,t}, \boldsymbol{\sigma}_{x,t}] = \varphi_\theta^{\text{dec}}(\varphi_\theta^{\text{z}}(\mathbf{z}_t), \mathbf{h}_{t-1})$

where $\boldsymbol{\mu}_{x,t}$ and $\boldsymbol{\sigma}_{x,t}$ denote the sufficient statistic parameters of the generated data distribution, and θ is the parameter set of the mapping function φ^{dec} . In this multivariate Gaussian case, they denote mean and variance, respectively. The hidden states are updated with the recurrence equation using gated recurrent unit. The second factor $p_\theta(\mathbf{z}_t|\mathbf{x}_{< t}, \mathbf{z}_{< t})$ is the prior of latent state, depending on previous RNN hidden state \mathbf{h}_{t-1}

$$\mathbf{z}_t \sim \mathcal{N}(\boldsymbol{\mu}_{0,t}, \text{diag}(\boldsymbol{\sigma}_{0,t}^2)) \quad (11)$$

where $[\boldsymbol{\mu}_{0,t}, \boldsymbol{\sigma}_{0,t}] = \varphi_\theta^{\text{prior}}(\mathbf{h}_{t-1})$

where $\boldsymbol{\mu}_{0,t}$ and $\boldsymbol{\sigma}_{0,t}$ are sufficient statistic parameters of the prior. By introducing temporal structure into the prior distribution of the latent variable, the representation power is enhanced. The hidden states are updated with the recurrence equation using gated recurrent unit (GRU) [36]

$$\mathbf{h}_t = f_\theta(\varphi_\theta^{\text{z}}(\mathbf{z}_t), \mathbf{h}_{t-1}). \quad (12)$$




Fig. 2. Graphical illustration of the proposed SISVAE model. Circle nodes: random variables. Diamond node: deterministic variables. Gray nodes: observable. White nodes: unobservable. Black solid lines: backbone recurrent neural network, as shown in (12). (a) Network structure of inference model $q_\phi(\mathbf{z}|\mathbf{x})$, red solid lines are the inference step as shown in (16), and black dotted lines are the transition prior as shown in (11). (b) Network structure of generative model $p_\theta(\mathbf{x}|\mathbf{z})$, blue solid lines are generation step in (10). (a) Inference model $q_\phi(\mathbf{z}|\mathbf{x})$. (b) Generative model $p_\theta(\mathbf{x}|\mathbf{z})$.

The GRU transition equations of our recurrence model are

$$\begin{aligned} \mathbf{y}_t &= [\varphi_\theta^z(\mathbf{z}_t)] \\ \mathbf{r}_t &= \sigma(\mathbf{W}_r \mathbf{y}_t + \mathbf{U}_r \mathbf{h}_{t-1} + \mathbf{b}_r) \\ \mathbf{s}_t &= \sigma(\mathbf{W}_s \mathbf{y}_t + \mathbf{U}_s \mathbf{h}_{t-1} + \mathbf{b}_s) \\ \tilde{\mathbf{h}}_t &= \tanh(\mathbf{W}_h + \mathbf{r}_t(\mathbf{U}_h \mathbf{h}_{t-1}) + \mathbf{b}_h) \\ \mathbf{h}_t &= (1 - \mathbf{s}_t) \tilde{\mathbf{h}}_t + \mathbf{s}_t \mathbf{h}_{t-1} \end{aligned} \quad (13)$$

where \mathbf{W}_* , \mathbf{U}_* , \mathbf{b}_* are model parameters of GRU in SISVAE model, and σ is the sigmoid function. In (10)–(12), $\varphi_\theta^{\text{prior}}$ and $\varphi_\theta^{\text{dec}}$ are realized by networks. $\varphi_\theta^x(\mathbf{x}_t)$ and $\varphi_\theta^z(\mathbf{z}_t)$ are feature extractors of observed data \mathbf{x}_t and latent variable \mathbf{z}_t respectively, which are also realized by networks.

2) *Inference Model*: We use variational inference to learn an approximate posterior over latent variables conditioned on the given data, which means we need to construct an approximating distribution q_ϕ parameterized by ϕ . We train the generative model using auto-encoding variable Bayes (AEVB) algorithm [27]

$$\max_{\theta, \phi} \mathbb{E}_{p(\mathbf{x}_{1:T})} \left[\mathbb{E}_{q_\phi} \left[\log \frac{p_\theta(\mathbf{x}_{1:T}, \mathbf{z}_{1:T})}{q_\phi(\mathbf{z}_{1:T}|\mathbf{x}_{1:T})} \right] \right]. \quad (14)$$

We construct a q_ϕ distribution of the following factorization:

$$q_\phi(\mathbf{z}_{\leq T}|\mathbf{x}_{\leq T}) = \prod_{t=1}^T q_\phi(\mathbf{z}_t|\mathbf{x}_{\leq t}, \mathbf{z}_{<t}). \quad (15)$$

The approximate posterior of latent variable \mathbf{z}_t depends on \mathbf{x}_t and \mathbf{h}_{t-1}

$$\begin{aligned} p(\mathbf{z}_t|\mathbf{x}_t) &\sim \mathcal{N}(\boldsymbol{\mu}_{z,t}, \text{diag}(\boldsymbol{\sigma}_{z,t}^2)) \\ \text{where } [\boldsymbol{\mu}_{z,t}, \boldsymbol{\sigma}_{z,t}] &= \varphi_\phi^{\text{enc}}(\varphi_\theta^x(\mathbf{x}_t), \mathbf{h}_{t-1}) \end{aligned} \quad (16)$$

where $\boldsymbol{\mu}_{z,t}$, $\boldsymbol{\sigma}_{z,t}$ denote the mean and variance of the approximate posterior, respectively, and $\varphi_\phi^{\text{enc}}$ is approximated by a neural network.

B. Anomaly Detection With Sequential VAE

As discussed before, we seek to devise suitable inductive bias for anomaly detection of time series in the sequential VAE framework. To this end, we begin with the factorization of the existing loss function. Specifically, we substitute the

numerator and denominator in (14) with (9) and (15) respectively, the objective of sequential VAE becomes

$$\tilde{\mathcal{L}}_{\text{SAE}} = \sum_{t=1}^T \left\{ \underbrace{D_{\text{KL}}(q_\phi(\mathbf{z}_t|\mathbf{x}_{\leq t}, \mathbf{z}_{<t}) \| p_\theta(\mathbf{z}_t|\mathbf{x}_{<t}, \mathbf{z}_{<t}))}_{\text{inference loss}} + \underbrace{\log p_\theta(\mathbf{x}_t|\mathbf{z}_{\leq t}, \mathbf{x}_{<t})}_{\text{reconstruction loss}} \right\} \quad (17)$$

The first term is the loss of inference model, which can be viewed as a regularizer, encouraging the model to learn disentangled feature by putting a prior over the latent variable [37]. The second term is the reconstruction loss of the generative model. Suppose $x_{m,t} = f_m(t) + \epsilon_{m,t} + e_{m,t}$ is an anomaly, and its associated reconstruction distribution is $\hat{x}_{m,t} = \mathcal{N}(\mu_{m,t}, \sigma_{m,t})$. The generative model tries to reconstruct the contaminated $x_{m,t}$, rather than the true signal $f_m(t)$, otherwise the reconstruct loss, which is the negative log likelihood $-p_\theta(x_{m,t}|\hat{x}_{m,t})$ will be high. To minimize the objective loss function from the data with anomalies, the model will either learn a biased mean $\mu_{m,t}$, or an over-estimated variance $\sigma_{m,t}$. As a result, the sequential VAE can be vulnerable to anomalies.

To account for such a bias caused by the anomalies, one natural idea is to induce smoothness prior to the generative model. Recall in Section III-B, the smoothness criterion refers to the accumulative difference of signal $\sum_{t=1}^T \nabla f_t$. One might be tempted to simply add such a smoothness regularization to the objective. However, since the sequential VAE is under the probabilistic framework, such a forced combination is not mathematically coherent. However, the classical approach treat observation as a deterministic variable, however, in sequential VAE, we treat each observation \mathbf{x}_t as a random variable, and we estimate the probability distribution for each observation $p(\mathbf{x}_t) \sim \mathcal{N}(\boldsymbol{\mu}_t, \text{diag}(\boldsymbol{\sigma}_t^2))$, such that the classical regularizer failed to work in a probabilistic framework, which calls for a more comprehensive method.

We extend the smoothness regularization to a probabilistic framework, by following the common assumption that the probability density function over time would vary smoothly over time. Among all possible reconstructed series $\hat{\mathbf{x}}_{1:T} \in \mathcal{X}$, the true signal should have a smooth transition of distribution.

To construct a valid regularizer, we need to measure the smoothness of a time-varying probability distribution.

For two consecutive data points of a single time series: let $x_{m,t-1}$ and $x_{m,t}$ be associated with reconstructions $p_\theta(\hat{x}_{m,t-1})$ and $p_\theta(\hat{x}_{m,t})$, respectively. Let $d(\cdot, \cdot)$ to be a distance metric between two distributions, we propose accumulative transition cost of time-varying probability distributions

$$\mathcal{L}_{\text{smooth}} = \sum_{t=1}^T \sum_{m=1}^M d(p_\theta(\hat{x}_{m,t-1}), p_\theta(\hat{x}_{m,t})). \quad (18)$$

In this article, we choose KL-divergence as the distance metric between distributions. For isotropic Gaussian case in this article, the transition cost of two normal distributions is

$$\begin{aligned} D_{\text{KL}}(\mathcal{N}(\mu_{m,t-1}, \sigma_{m,t-1}) || \mathcal{N}(\mu_{m,t}, \sigma_{m,t})) \\ = \log \frac{\sigma_{m,t}}{\sigma_{m,t-1}} + \frac{\sigma_{m,t-1}^2 + (\mu_{m,t-1} - \mu_{m,t})^2}{2\sigma_{m,t}^2} - \frac{1}{2} \end{aligned} \quad (19)$$

The final objective involves: 1) encoding loss for inference model; 2) expected reconstruction error for generative model; and 3) the proposed variational smoothness regularizer. We apply stochastic gradient variational Bayes (SGVB) estimator to approximate the expected reconstruction error $\mathbb{E}_{q_\phi(\mathbf{z}_{\leq T} | \mathbf{x}_{\leq T})}[\log p_\theta(\mathbf{x}_t | \mathbf{z}_{\leq t}, \mathbf{x}_{< t})]$. The learning objective is

$$\begin{aligned} \mathcal{L}_{\text{SISVAE}} \\ = \sum_{t=1}^T \left\{ \underbrace{-D_{\text{KL}}(q_\phi(\mathbf{z}_t | \mathbf{x}_{\leq t}, \mathbf{z}_{< t}) || p_\theta(\mathbf{z}_t | \mathbf{x}_{< t}, \mathbf{z}_{< t}))}_{\text{inference loss}} \right. \\ \quad \underbrace{+ \log p_\theta(\mathbf{x}_t | \mathbf{z}_{\leq t}, \mathbf{x}_{< t})}_{\text{reconstruction loss}} \\ \quad \left. + \lambda \underbrace{D_{\text{KL}}(p_\theta(\hat{\mathbf{x}}_{t-1} | \mathbf{z}_{\leq t-1}, \mathbf{x}_{< t-1}) || p_\theta(\hat{\mathbf{x}}_t | \mathbf{z}_{\leq t}, \mathbf{x}_{< t}))}_{\text{smoothness loss}} \right\} \end{aligned} \quad (20)$$

where $\mathbf{z}_t = \boldsymbol{\mu}_{z,t} + \boldsymbol{\sigma}_{z,t} \odot \boldsymbol{\epsilon}$ and $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$, and λ is a smoothness regularization hyperparameter. We set $p_\theta(\hat{\mathbf{x}}_t | \mathbf{z}_{\leq t-1}, \mathbf{x}_{< t-1}) = p_\theta(\hat{\mathbf{x}}_t | \mathbf{z}_{\leq t}, \mathbf{x}_{< t})$ when $t = 1$.

C. Training

1) *Preparing Data Set*: Ideally, we could train SISVAE with a whole time series as input, and the model outputs the whole reconstruction. However, in many real applications, the time series is typically very long. The objective function (20) of our model is summed sequentially over time, such that computing the gradient of it requires backpropagation over time [38]. Research studies [39] observe that RNNs are hard to train when the sequence length is over 200, and the model may face the problem of vanishing gradients and exploding gradients, which is also undesirable for anomaly detection tasks.

To solve this problem, we divide the long time series into short chunks, using sliding windows technique. We apply a sliding windows to the time series, which slides over multiple time series synchronously. The sliding windows is

Algorithm 1 Auto-Encoded Variational Bayes Algorithm for Learning SISVAE for Time Series

Input: Time series matrix $\mathbf{X} \in \mathbb{R}^{M \times T}$

Input: λ for smoothness, W and s for sliding windows

Split time series into chunks $\mathcal{D} = \{\mathcal{X}_d\}_{1:D}$, $\mathcal{X} \in \mathbb{R}^{M \times W}$

Randomly initialize θ, ϕ ;

Initialize $\mathbf{h}_0 = \mathbf{0}$;

while not converged do

 Sample a mini-batch of P chunks from data set \mathcal{D} ;

for each chunk in parallel do

for $t = 1$ **to** W **do**

$\mathbf{x}_t = \mathcal{X}_{:,t}$;

$\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$ //random sampling;

$[\boldsymbol{\mu}_{z,t}, \boldsymbol{\sigma}_{z,t}] = \phi_\phi^{\text{enc}}(\phi_\theta^{\text{x}}(\mathbf{x}_t), \mathbf{h}_{t-1})$ //encoding;

$\mathbf{z}_t = \boldsymbol{\mu}_{z,t} + \boldsymbol{\sigma}_{z,t} \odot \boldsymbol{\epsilon}$ // reparam trick;

$[\boldsymbol{\mu}_{0,t}, \boldsymbol{\sigma}_{0,t}] = \phi_\theta^{\text{prior}}(\mathbf{h}_{t-1})$ //prior;

$[\boldsymbol{\mu}_{x,t}, \boldsymbol{\sigma}_{x,t}] = \phi_\theta^{\text{dec}}(\phi_\theta^{\text{z}}(\mathbf{z}_t), \mathbf{h}_{t-1})$ //decoding;

$\mathbf{h}_t = f_\theta(\phi_\theta^{\text{z}}(\mathbf{z}_t), \mathbf{h}_{t-1})$ //recurrence;

end

end

 Compute gradient $\nabla_\theta \mathcal{L}$ and $\nabla_\phi \mathcal{L}$ with $\mathbf{z}_{\leq W}$;

 Update parameters θ and ϕ ;

end

return θ, ϕ

controlled by two parameters: window size W and step size s . Specifically, for each data set $\mathbf{X} \in \mathbb{R}^{M \times T}$, we have

$$\begin{aligned} \mathcal{X}_1 &:= \mathbf{X}_{M,1:1+W} \\ \mathcal{X}_2 &:= \mathbf{X}_{M,1+s:1+s+W} \\ &\vdots \\ \mathcal{X}_d &:= \mathbf{X}_{M,1+(d-1)*s:1+(d-1)*s+W} \end{aligned} \quad (21)$$

where $\mathcal{X} \in \mathbb{R}^{M \times W}$, each \mathcal{X} is a short time-series chunk with length W , putting together, we have

$$\mathcal{D} = \{\mathcal{X}_1, \dots, \mathcal{X}_d, \dots, \mathcal{X}_D\} \quad (22)$$

where \mathcal{D} is the whole data set. Each item is a short time-series chunk. The learning procedure is depicted in Algorithm 1.

D. Computing Anomaly Scores

The next step is to compute anomaly score for each data point $\mathbf{x}_{m,t}$. For a trained model, input a chunk of time series $\mathcal{X} = \{\mathbf{x}_t\}_{1:W}$, each \mathbf{x} is an M -dimensional vector, the model first encode it into $\mathbf{z}_{1:W}$, then decode to a sequence of random variables $\hat{\mathcal{X}} = \{\mathcal{N}(\boldsymbol{\mu}_t, \text{diag}(\boldsymbol{\sigma}_t))\}_{1:W}$. The whole encoding and decoding process is called *probabilistic reconstruction*.

Recall that in (1) of Section III-A, we use \mathbf{I} to indicate the indicator matrix of anomaly labels. In this step, we use \mathbf{A} to indicate the matrix of anomaly scores.

One natural choice is to use absolute reconstruction error, like anomaly detection with deterministic auto-encoders [40]. The anomaly score $e_{m,t}$ for observation $x_{m,t}$ is computed as

$$e_{m,t} = \|\mathbf{x}_{m,t} - \boldsymbol{\mu}_{m,t}\| \quad (23)$$




Fig. 3. When the variance of data varies over time, absolute reconstruction error may fail to detect some anomalies.

where $\mu_{m,t}$ is the m th dimension of reconstructed mean μ_t . Intuitively, a large reconstruction error indicates anomaly behavior of time series, which is a common assumption for many reconstruction-based anomaly detection algorithms. The underlying assumption of the absolute reconstruction error is that the variance is stationary, which means that the variance is constant over time. However, in many real-world data, the variance is nonstationary, and it varies over time. We illustrate this using an example as shown in Fig. 3. We generate some time-series data with the Gaussian process as the true signal, denoted as the real blue line centering scatters, then we add some Gaussian noise $\mathcal{N}(0, \sigma_t)$, where σ_t increases over time. The black dots are the noisy observation, and blue transparent regions increasing the level of noise. Then, we add two points, point A and point B, they both deviate from the true signal with distance 1. We see that, point A is more likely to be an anomaly point, while point B is more likely to be a normal point located at some noisy period. In this case, if we use absolute reconstruction error, we may fail to detect anomaly point A, or otherwise report normal point B as an anomaly.

Due to the nonstationary nature of real-world time series, we need a robust method for computing anomaly scores. Previous VAE-based anomaly detection methods have introduced reconstruction probability based anomaly score, here we devise a reconstruction probability for sequential auto-encoders. Specifically, the probability density of input \mathbf{x} is

$$p_{\theta}(\mathbf{x}) = \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})}[p_{\theta}(\mathbf{x}|\mathbf{z})]. \quad (24)$$

Draw L samples from $q_{\phi}(\mathbf{z}|\mathbf{x})$, and the anomaly score is negative of average of L probability. For numerical stability, we often work with log probability, and we have

$$\log p_{\theta}(\mathbf{x}) = \frac{1}{L} \sum_{l=1}^L p_{\theta}(\mathbf{x}_t | \mu_x^{(l)}, \sigma_x^{(l)}). \quad (25)$$

However, for sequential VAE, this can be more challenging. Because \mathbf{z} is not independent but has sequential dependencies. We choose to adopt sequential Monte Carlo (SMC) [41], which is widely used in state-space models, such as Bayesian filters [42], and dynamical systems [43].

We conclude the SMC for computing anomaly scores and detecting anomalies in Algorithm 2. For the full multivariate time-series matrix, $\mathbf{X} \in \mathbb{R}^{M \times T}$, we slice them into nonoverlapped chunks and reconstruct each of them individually.

V. EXPERIMENTS

We present the experiment for robust modeling and anomaly detection for multidimensional correlated time series,

Algorithm 2 SMC for Computing Anomaly Score and Detecting Anomalies

Input: Trained model θ, ϕ , threshold α , time series chunk $\mathcal{X} \in \mathbb{R}^{M \times W}$, # of SMC iterations L .

Output: Anomaly score \mathbf{A} , detected anomaly \mathbf{I}_{α}

Initialize $\mathbf{A} = \mathbf{0}$; **for** $l = 1$ **to** L **do**

for $t = 1$ **to** W **do**

$\mathbf{x}_t = \mathcal{X}_{:,t}$;

$\epsilon \sim \mathcal{N}(0, \mathbf{I})$ // random sampling;

$[\mu_{z,t}, \sigma_{z,t}] = \phi_{\phi}^{\text{enc}}(\phi_{\theta}^{\text{x}}(\mathbf{x}_t), \mathbf{h}_{t-1})$ // encoding;

$\mathbf{z}_t = \mu_{z,t} + \sigma_{z,t} \odot \epsilon$ // reparam trick;

$[\mu_{x,t}, \sigma_{x,t}] = \phi_{\theta}^{\text{dec}}(\phi_{\theta}^{\text{z}}(\mathbf{z}_t), \mathbf{h}_{t-1})$ // decoding;

$\mathbf{h}_t = f_{\theta}(\phi_{\theta}^{\text{z}}(\mathbf{z}_t), \mathbf{h}_{t-1})$ // recurrence;

$\mathbf{A}_{:,t} += -\frac{1}{L} \log p(\mathbf{x}_t | \mu_{x,t}, \sigma_{x,t})$;

end

end

$\mathbf{I}_{\alpha} = \mathbf{A} > \alpha$; // element-wise comparison;

return $\mathbf{A}, \mathbf{I}_{\alpha}$;

compared with peer VAE methods. Recall that in this article we deal with point level anomaly detection without supervision.

A. Protocols and Settings

1) *Performance Metrics*: For time-series data $\mathbf{X} \in \mathbb{R}^{M \times T}$, the model outputs anomaly score matrix $\mathbf{A} \in \mathbb{R}^{M \times T}$, and each item $A_{m,t}$ is the likelihood of data point $x_{m,t}$ being an anomaly. For detection, we set a threshold α to obtain the final anomalies label indicator matrix $\mathbf{I}_{\alpha} := \mathbf{A} > \alpha$. Then we compute precision, recall, true-positive rate (TPR), false-positive rate (FPR) at time series point level.

The tradeoff is needed as the detection characteristics vary by different threshold α . When the threshold is tight, the precision will be higher and recall will be lower. We compute the area under ROC (AUROC), the area under precision recall curve (AUPRC) and the best F1-score. AUROC quantifies the FPR and TPR with different α , and AUPRC quantifies the precision and recalls with different α . Note that all three metrics are independent of specific threshold.

There are inherent differences between AUROC and AUPRC. For anomaly detection tasks where class (anomaly or normal data points) distributions are heavy imbalances, the AUROC metric is less insensitive to false-positives than the AUPRC metric. As discussed in [44] and [45], the AUPRC score is more discriminative than the AUROC score. Previous work may choose one of the metrics to evaluate the performance. To comprehensively evaluate the model's performance, we report both AUROC and AUPRC scores, but we mainly focus on analysis AUPRC in this article.

2) *Peer Methods*: Recent deep learning-based methods are compared, including Donut and STORN.

1) **Donut [3]**. The model proposed very recently is based on VAE, which takes a fixed length sliding windows of a univariate time series as input. To capture “normal patterns” of data, the modified ELBO technique is devised to leverage generative model $p_{\theta}(\mathbf{x}|\mathbf{z})$ and latent

TABLE II

SUMMARY OF DIFFERENT TIME SERIES ANOMALY DETECTION MODELS. NOTE THE TRIANGLE FOR DONUT DENOTES THE OBJECTIVE ROBUSTNESS IS INCORPORATED INDIRECTLY BY CONSIDERING THE LABELS OF ANOMALIES IN THE OBJECTIVE WHILE OUR METHOD FULFILLS THIS PURPOSE BY USING THE (SMOOTH) TRENDING PRIOR

	Traditional Models		Deep Latent Variable Models		
	ARMA	LDS	Donut	STORN	SISVAE-p
latent temporal dependencies	–	✓	✗	✓	✓
non-linear emission probability	–	✓	✓	✓	✓
non-linear latent transitions	–	✗	✗	✓	✓
trending prior	–	✓	✗	✗	✓
multi-dimensional	✗	✓	✗	✓	✓
non-stationary variance	✗	✗	✓	✓	✓
long sequence	✗	✗	✓	✗	✓
objective robustness	–	✗	▽	✗	✓

regularization $p_{\theta}(\mathbf{z})$ of VAE, such that the model is encouraged to learn only normal patterns.

- 2) **STORN** [30]. As a sequential version of VAE, the model is based on stochastic recurrent neural network (STORN). Different from SISVAE, the sequences of latent variables \mathbf{z}_t are generated independently. The model uses lower bound output as the anomaly score.

We also include three traditional statistical models.

- 1) *History Average (HA)*. We use absolute deviation from history average value as the anomaly score. We include this model in comparison to show how a very simple model performs within different data sets.
- 2) *Auto-regressive Moving Average (ARMA)* [46]. ARMA is a classic time series analysis model [47]. We use absolute fitting residuals as the anomaly score.
- 3) *Linear Dynamical System (LDS)*. This model is also known as *Kalman filter (KMF)*. Different from deep latent variable-based state-space models, the transition function and emission probability are linear (see Table II).

We further include some variants of our model to study the contribution and behavior of each component in our model.

- 1) *SISVAE-p*. Our recommended model, which uses the sliding windows of a time-series matrix as input, similar to Donut. This model is more applicable for long sequences.
- 2) *SISVAE-e*. Different from SISVAE-p, it takes absolute reconstruction error as an anomaly score. We use this model to show the performance gap by different detection criteria.
- 3) *SISVAE-0*. To show the effect of our smoothness regularizer, we let SISVAE-0 to be SISVAE-p for $\lambda = 0$.
- 4) *VAE-s*. To test the contribution of the proposed smoothness prior technique, we add the regularization to vanilla VAE as a comparison named VAE-s. The model can be viewed as our SISVAE model without a latent temporal structure. We use this model to show the effectiveness of introducing a latent temporal structure.

In particular, we compare the characteristics of peer methods in Table II, whereby both traditional methods, that is, HA, ARMA, KMF, and deep latent model, that is, Donut, STORN, SISVAE, and SISVAE-mb are included.

Implementation Details: For preprocessing, we set window size W to be 120 for each data set. The sequential VAE-based models, STORN, and SISVAE have a similar network architecture, which is defined by two varying hyperparameters, h_dim for dimension of a hidden layer, and z_dim for dimension of latent variable \mathbf{z} . We set h_dim to be 200, and z_dim to be 40 for all data sets. We empirically set the smoothness regularization hyperparameter $\lambda = 0.5$ for SISVAE during training process. The encoder, decoder, and RNN of models use the same hidden dimension. The three models are implemented using PyTorch [48] 0.4, and trained on a single Nvidia RTX 2080ti GPU. We train the three models 200 epochs using Adam [49] with learning rate 0.001. Note for Donut, the M-ELBO technique proposed by Donut may require some anomaly labels to learn ‘normal’ patterns. For a fair comparison, we provide no anomaly label to the model, and the experiments are conducted under fully unsupervised setting. We use the code open sourced by the authors. We set h_dim to be 200, and z_dim to be 40, same as our model. The model is fundamentally designated for univariate time series, and it is nontrivial to extend it to the multidimensional case. We test two settings: 1) train a model for each sequence individually, and compute anomaly score individually; 2) flatten the multidimensional time-series data after preprocessing to accommodate the model, the anomaly score is computed together. We find that the second protocol exhibits better performance. This may be because the normal patterns are common along different sequences. With more sequences as input, the model can be less influenced by anomalies. We adopt the second protocol for Donut in the experiments. For all deep latent variable models, we use negative reconstruction probability as anomaly score output. For all three traditional statistical models, we use the same preprocessing method. For history average (HA), the anomaly score is absolute observation value since all series have zero mean. For ARMA and LDS, we use absolute reconstruction error as anomaly score.

B. Experiments on Synthetic Data

In this section, we aim to gain insights into the behavior of our proposed SISVAE model. We investigate the anomaly detection performance of our model for different proportions of anomalies. Furthermore, we are interested in the anomaly detection performance under different settings of smoothness regularization hyper-parameter λ .

Referring to the synthetic data generation protocol in [24], we generate 100 time series sequences of length 200 from a multioutput Gaussian process with a multiple output kernel [50], implemented by a python package named GPy [51]. Sequences are correlated through a random sampled positive-definite coregionalization matrix, and we add random non-stationary Gaussian noise on the 100-dimensional sequence of length 200. The location of anomalies is simulated by a binary mask with Bernoulli distribution of probability $\{0.5\%, 1\%, 2\%, 5\%, 10\%\}$ for different fractions of anomalies. We insert Poisson noise where the binary mask is equal to one as anomalies, and the intensity λ_i of each anomaly is proportional




Fig. 4. Results of controlled experiment on synthetic data sets. (a) Anomaly detection performance for various proportions of anomalies. (b) Anomaly detection performance of SISVAE with different values of regularization hyperparameter λ . (c) AUPRC score over training iterations with different regularizers (KL-regularized is our proposed regularizer and the mean-regularized means only the mean is regularized as done in classical methods). (d) Convergence property of reconstruction term and recognition term of objective function, top: unregularized model, bottom: KL-regularized model. Here, SISVAE denotes the reconstruction probability-based model, the result for reconstruction error-based method performs worse hence is omitted here. (a) Sensitivity to anomaly ratio, (b) sensitivity to regularizer, (c) effects of regularizer, and (d) convergence of losses.

to the amplitude of inserted place x_t . The plus and minus sign is randomly sampled from Bernoulli distribution with $p = 0.5$.

1) *Effects of Anomaly Proportion*: Model performance variation with respect to the proportion of anomalies is shown in Fig. 4(a). Since all the models are unsupervised, they are inevitably influenced by the anomalies, because a full reconstruction of both normal and anomaly data is encouraged in the objective function. The proposed variational smoothness regularizer places penalties at consecutive outputs of the generative model that violates the temporal smoothness assumption. As observed from the result, as the anomaly proportion grows, the AUPRC scores of two unregularized model SISVAE-0 and STORN drop significantly, whereas the regularized models' performance almost keeps maintained. SISVAE-0, the proposed SISVAE model shows superior performance over competing methods in all settings of anomaly proportions.

2) *Effects of Regularization Hyper-Parameter*: The variational smoothness hyper-parameter λ in the objective by (17) determines the penalty of nonsmooth output of the generative model. We analyze the detection performance of the SISVAE model under different settings of hyperparameters. We train the model with $\lambda = \{0.01, 0.05, 0.1, 0.2, 0.5, 0.8, 1.0\}$ using synthetic data set with 2% anomaly proportion. Fig. 4(b) shows the change of AUPRC performance w.r.t λ , where the model achieves the best performance when λ is set to be 0.5.

3) *Convergence Study*: We analyze the contribution of our regularizer to convergence property. On the synthetic 2% anomaly proportion data set, we train three models: 1) SISVAE-p without regularization ($\lambda = 0$); 2) classical smoothness regularizer of deterministic model where only the mean is regularized which is formulated as: $\lambda \sum_{t=1}^T \sum_{m=1}^M (\nabla_t^2 \hat{x}_{m,t})^2$; 3) proposed KL regularizer. We record the AUPRC score over training. As shown in Fig. 4(c), at the first 100 iterations, all the three models are trained very fast, and the detection performance increases stably. This suggests that the models are learning the major "normal" patterns of the data, whereas the anomalies are temporally neglected. However, after 100 iterations, the performance of the unregularized and mean-regularized models fluctuates severely between 0.45 and 0.85, and our KL-regularized model (i.e. SISVAE-p) also fluctuates but in a relatively small interval. We think the fluctuation is caused by a contradiction between generative (decoder) model and recognition (encoder) model: the generative model is trying hard to reconstruct every

TABLE III
STATISTICS OF THE USED DATA SETS

Dataset	A1	A2	A3	A4	μ PMU
# sequences	65	100	100	100	36
# length	1420	1421	1680	1680	1080
# points	92300	142100	168000	168000	38800
# anomalies	1462	426	943	1044	240
# proportion	1.58%	0.32%	0.56%	0.62%	0.62%

data point perfectly, while the inference model aims to map the observed anomaly contaminated data points into a simpler low-dimensional latent space. The decoder and encoder are playing a seesaw battle in the training procedure, thus causing a sawtooth-shape AUPRC score. Under the probabilistic framework, classical mean-regularizer which is in its nature deterministic, cannot work well. Its convergence curve is akin to the unregularized model. In contrast, the proposed variational smoothness regularizer weakens the reconstruction objective, and guides the decoder to generate sequences that obey temporal smoothness assumption, preventing the model from overfitting "anomaly" patterns. As shown in Fig. 4(d), the reconstruction loss and inference loss of the regularized model converge stably, while the training of unregularized model is unstable. Hence, the KL-regularized SISVAE model reaches a higher AUPRC score.

C. Experiments on Real-World Data

Data sets. This experiment involves the following data sets, we conclude statistics of data sets in Table III.

- 1) *Yahoo's S5 Webscope Data Set*: This data set is a part of Yahoo's Webscope program.² It consists of 367 time series. This data set consists of four different classes A1/A2/A3/A4 with 65/100/100/100 time series, respectively. Class A1 has real data collected from computational systems, which records real production traffic of Yahoo's website, and the anomalies are labeled by human experts manually. Classes A2, A3, and A4 contain synthetic data, whereby anomalies are constructed and are inserted at random positions.
- 2) *Open μ PMU Power Network Data Set*: This data set is collected from a power distribution equipped in Alameda, CA, USA, with advanced smart meters

²<https://webscope.sandbox.yahoo.com/>

TABLE IV
PERFORMANCE ON REAL-WORLD PUBLIC DATA SETS. SISVAE-E IS OMITTED FOR ITS WORSE PERFORMANCE THAN SISVAE-P

Dataset	Metric	HA	ARMA	LDS	Donut	STORN	SISVAE-0	VAE-s	SISVAE-e	SISVAE-p
A1	AUROC	0.616	0.845	0.664	0.825	0.867	0.862	0.886	0.865	0.876
	AUPRC	0.276	0.305	0.099	0.178	0.166	0.184	0.185	0.300	0.369
	F1	0.337	0.351	0.140	0.215	0.240	0.254	0.254	0.358	0.401
A2	AUROC	0.956	0.658	0.952	0.999	0.926	0.974	0.993	0.989	0.981
	AUPRC	0.745	0.233	0.401	0.623	0.435	0.608	0.536	0.731	0.817
	F1	0.774	0.457	0.480	0.691	0.515	0.588	0.611	0.709	0.789
A3	AUROC	0.724	0.486	0.989	0.990	0.917	0.912	0.907	0.951	0.993
	AUPRC	0.252	0.005	0.835	0.753	0.732	0.746	0.713	0.865	0.954
	F1	0.331	0.011	0.831	0.798	0.757	0.791	0.763	0.880	0.951
A4	AUROC	0.715	0.487	0.979	0.991	0.850	0.834	0.836	0.831	0.924
	AUPRC	0.107	0.005	0.737	0.598	0.438	0.521	0.525	0.521	0.774
	F1	0.165	0.010	0.745	0.704	0.478	0.582	0.573	0.582	0.820
μ PMU	AUROC	0.629	0.587	0.805	0.714	0.684	0.628	0.629	0.630	0.814
	AUPRC	0.084	0.023	0.303	0.312	0.185	0.174	0.089	0.153	0.310
	F1	0.225	0.049	0.356	0.389	0.279	0.247	0.228	0.244	0.367

called phasor measurement units (μ PMU)³ [52]. There are three μ PMU devices, each recording three-phase voltage, current magnitude, and phase angle. Anomalies caused by voltage disturbance are labeled by human experts. We collect measurement data and anomaly labels during the whole day of June 3 with 20 s resolution. We use 36-dimensional time series of length 1080.

Preprocessing In real applications, different time series may have different scales, for example, temperature and air pressure sensors have different scales. So, we follow the common practice to normalize each sequence to standard normal distribution.

1) *Threshold-Based Anomaly Detection*: First, we conduct a threshold-based anomaly detection experiment. For each model, we collect anomaly score matrix A for each model, then we compute AUROC, AUPRC, and F1-score for each possible threshold α . We select the highest F1-score the model could achieve as the final F1-score. We compare all the peer methods in this experiment, the results of which are shown in Table IV. SISVAE-p outperforms other models in most data sets and metrics. For SISVAE-p and SISVAE-e, we find that our proposed SMC-based reconstruction probability for computing anomaly score notably improves the performance. Among all the five data sets, we find SISVAE-p and SISVAE-e perform similarly on A2 data set. We further find that the A2 data set is, in general, stationary regarding variance, which means that modeling the time-varying variance is important for anomaly detection. For traditional statistical models, we find that the simple method HA shows decent performance in some data sets, however, the potential of HA is very limited. We are surprised that ARMA failed on A3, A4, and μ PMU data set. Our further study uncovers that these three data sets are highly nonstationary with both mean and variance, which violate the basic stationary assumption of ARMA. In contrast, we find LDS performs competitively on A3, A4, and μ PMU compared to the best model, however, it fails on data set A1. This suggests that linear transition and linear emission probability

of LDS may not be able to capture the complex nonlinear relationship. For VAE-s, the proposed smoothness prior regularizer works even without a temporal structure model. In our analysis, the regularizer constructs inductive bias of the true signal, leading to robustness to the unknown anomaly.

As discussed, we find models' anomaly detection performance varies across data sets, such as LDS and ARMA. To make the overall performance comparison more intuitionistic, we compute the rank of models for each data set and performance metric, and compute the mean and standard deviation of ranks for each model. The plots of rankings are shown in Fig. 5 showing our model outperforms other models notably.

2) *Model Characteristics*: The detection threshold α is often set according to specific needs. In order to investigate the detailed detection characteristic, we plot the precision-recall curve (PRC), and receiver operating characteristic curve (ROC), as shown in Figs. 6 and 7. Compared to LDS and STORN, SISVAE-p outperforms at most thresholds. For precision and recall curve in Yahoo A2, A3, and A4 data set, note that Donuts performs strangely, where the precision and recall curves do not monotonically decrease as expected, but increase first and then decrease. We find this is because Donut places very large anomaly scores to some normal points, while the other three models do not have this problem. This means modeling latent temporal dependencies is important for time series anomaly detection tasks.

3) *Ranking-Based Anomaly Detection*: Due to a limited budget, sometimes we can only deal with a limited number of detected anomalies hence the precision need be high. In this context, precision@K is a useful metric, which is widely used in information retrieval [53]. For each model, we select data points that are assigned with top K highest anomaly score, and precision@K is computed by V/K , where V is the number of true anomalies in K predictions. We examine models with $K = 10, 50, 200$, and the results are shown in Fig. 8. We can see that, when K is 10, STORN, LDS, and SISVAE-p perform well. However, Donut fails on A2, A3, and A4 data set. The reason is that Donut places a very high anomaly score on some normal data points. Overall, SISVAE-p's precision

³<https://powerdata.lbl.gov>

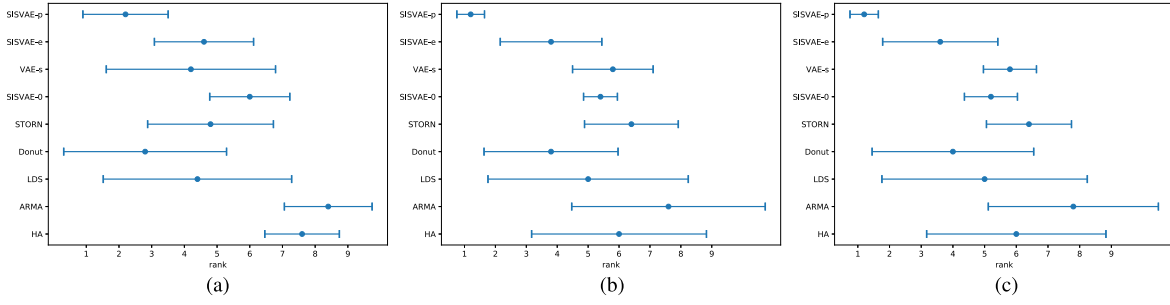


Fig. 5. Ranking of models with respect to three evaluation metrics. Bar ranges indicate the standard variations centered with their ranking mean. (a) AUROC. (b) AUPRC. (c) F1.




Fig. 6. Visualization of precision and recall curve of SISVAE-p, Donut, LDS, and STORN on real-world data sets. (a) Yahoo-A1. (b) Yahoo-A2. (c) Yahoo-A3. (d) Yahoo-A4. (e) μ PMU.

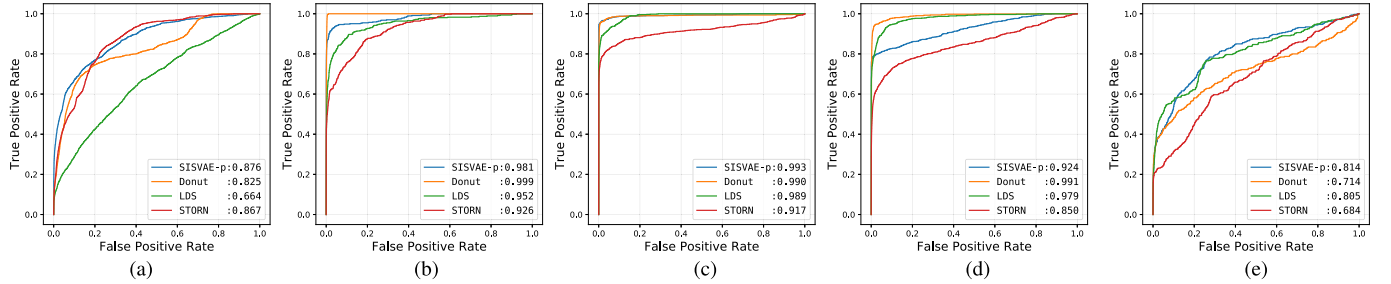


Fig. 7. Visualization of ROC of SISVAE-p, Donut, LDS, and STORN on real-world data sets. (a) Yahoo-A1. (b) Yahoo-A2. (c) Yahoo-A3. (d) Yahoo-A4. (e) μ PMU.

keeps stable as K increases. This means SISVAE-p makes more precise detection when the number of detection trials is limited.

4) *Tradeoff Between Recall and Precision*: Despite the fact that SISVAE-p is better than baselines in most cases, however, the preference of detection metric also affects the choice of model. It is important to know when to use which model. The cost of validating a detected anomaly can be expensive that calls for model's high precision. For large-scale power grid system, it is required that the precision need be strictly above 0.95. In this case, SISVAE-p outperforms other methods in recall. In contrast, for fraud detection, recall can be more important. If we set a hard bar for recall to 0.95, we find Donut is also competitive, it outperforms SISVAE-p in A2 and A4 data set w.r.t. precision.

D. Case Study

We compare the three deep generative models, SISVAE-p, STORN, and Donut with real-world examples from Yahoo-A1 data set. For each model, we choose to set threshold α when

F1-score is maximized. The data set consists of 65 sequences each with length 1420. In what follows, we show some representative time series fragments for discussion.

1) *False Alarms on Normal Time Series*: Fig. 9(a) shows a normal sequence without any anomaly. Since the training data are contaminated with anomalies, the model can be influenced by anomalies and fails to capture normal patterns such that false alarms can be reported. Fig. 9(b) illustrates the density estimation of the three models, and Fig. 9(c) shows the anomaly score generated by three models. Note STORN reports three false alarms, while SISVAE-p and Donut do not. Moreover, we find the anomaly score generated by SISVAE-p is stably low, while Donut is more fluctuating given noisy data. The results demonstrate that SISVAE-p is robust to contaminated data for unsupervised training.

2) *Detection of Point-Level Anomalies*: Fig. 10(a) includes a sequence from Yahoo-A1 data set with an anomaly at around time point 300. This type of anomaly is often called *additive outlier* [54] or *extreme value* [55]. As shown in Fig. 10(b), density estimation of STORN and Donut are corrupted by the




Fig. 8. Visualization of precision@K for ranking-based detection tasks on real-world data sets. (a) Yahoo-A1. (b) Yahoo-A2. (c) Yahoo-A3. (d) Yahoo-A4. (e) μ PMU.

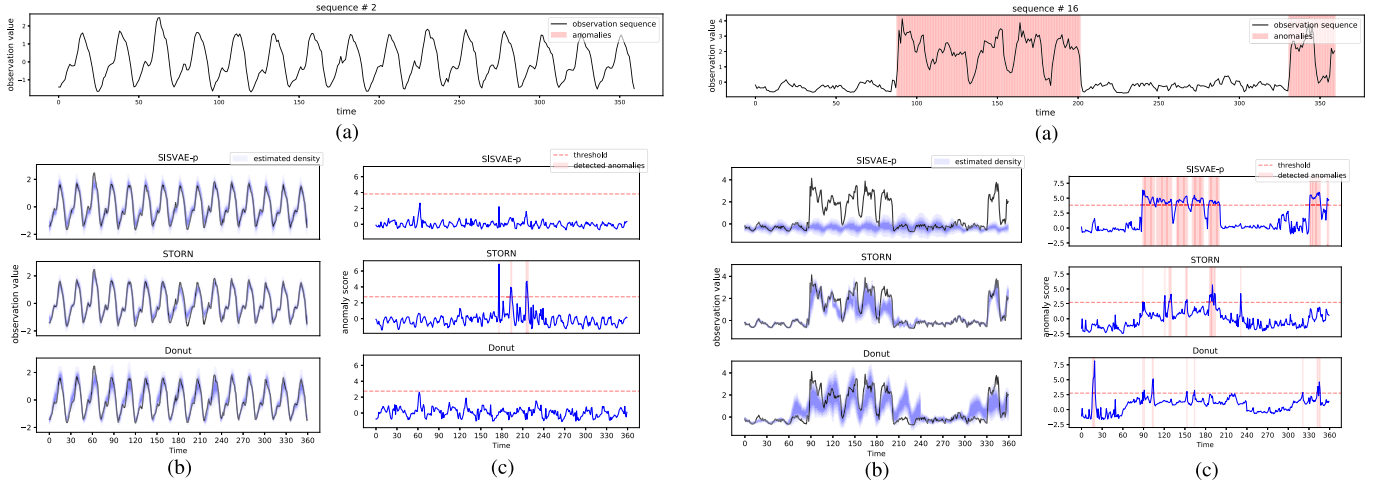


Fig. 9. Example for time series without anomaly. Note false alarm by STORN. (a) Studied dimension of the 65-dim time series labeled with ground-truth anomaly windows [stretched twice compared to (b) and (c)]. (b) Density estimation. (c) Anomaly detection score.

Fig. 11. Example for time-series sample with window-sized anomaly. Note the missed anomaly regions by STORN and Donut (a) studied dimension of the 65-dim time series labeled with ground-truth anomaly windows [stretched twice compared to (b) and (c)]. (b) Density estimation. (c) Anomaly detection score.

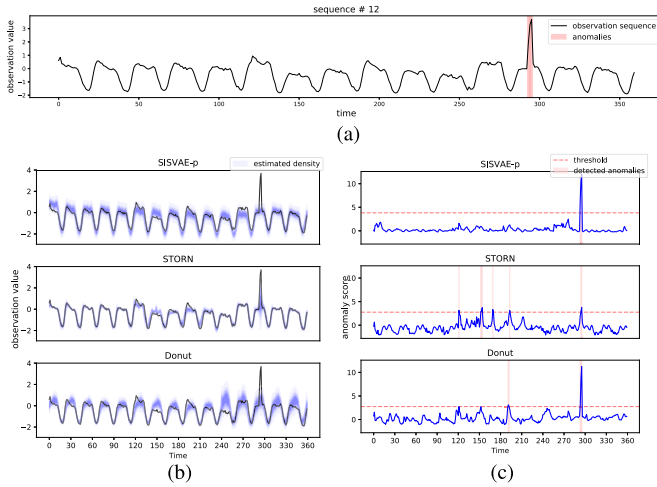


Fig. 10. Example for time-series sample with point-level anomaly. Note the false alarms by STORN and Donut (a) studied dimension of the 65-dim time series labeled with ground truth anomaly windows [stretched twice compared to (b) and (c)]. (b) Density estimation. (c) Anomaly detection score.

point anomaly, such that they place much probability density to the anomaly point. In contrast, SISVAE-p successfully recovers the true underlying density. Fig. 10(c) shows the anomaly score generated by the three models, and we can see that all the models identify the anomaly point. However, STORN and Donut report some false alarms. Moreover, the anomaly score

of SISVAE-p is more smooth than those of other models, suggesting that SISVAE-p is more robust to the moderate fluctuation of normal data. Overall, SISVAE is capable of detecting point-level anomalies robustly.

3) *Detection of Subsequence Level Anomalies:* In some cases, the anomalies are in the presence of subsequences, such as the anomaly cardiac cycle in ECG data [56]. Fig. 11(a) shows a sequence with two anomaly subsequences, whereby the red shaded region covers the true anomaly. Fig. 11(b) shows the density estimation, whereby Donut and STORN are notably influenced by the anomaly subsequence. In contrast, SISVAE-p estimates a smooth probability density over time, which is more likely to be the true density. Fig. 11(c) shows SISVAE-p successfully detecting most of the anomaly subsequences, whereas STORN and Donut suffer from false-negative and false-positive. The results demonstrate the effectiveness of the proposed smoothness prior regularizer, helping recover the true density and improve detection performance.

E. Further Comparison With Donut

1) *Simulation Study on Uni-Variate Time Series:* As discussed in Section V-A, Donut is designated for univariate time series. In order to investigate the detection performance in 1D case, we conduct simulation study on data generated from Mackey-Glass equation [57], which serves a benchmark




Fig. 12. Generated Mackey-Glass time-series data and inserted anomalies.

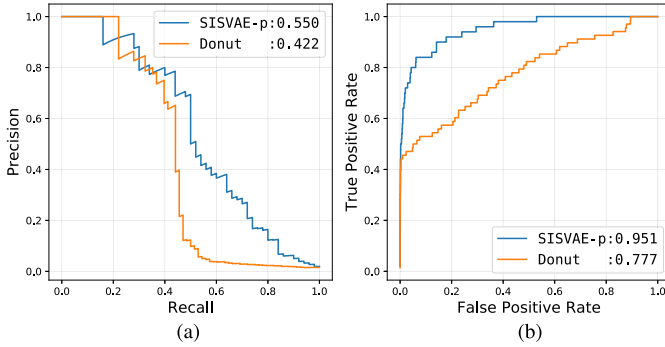


Fig. 13. Comparison of anomaly detection performance between SISVAE-p and Donut on Mackey-Glass time-series data set. (a) Precision and recall curve. (b) ROC.

for nonlinear models since the true underlying function is nonlinear [58]. The time series is generated as

$$\frac{dx(t)}{dt} = \frac{\alpha x(t - \tau)}{1 + x^\beta(t - \tau)} - \gamma x, \quad \gamma, \beta, n > 0 \quad (26)$$

with $\gamma = 0.1, \beta = 10, \alpha = 0.2$. We run the Mackey-Glass equation 5000 iterations, then we collect a univariate time series with length 5000. Two types of anomalies are added:

- 1) *Point Level Anomaly*: We insert point anomalies at 0.3% time step randomly, each point anomaly is a combination of: 1) Poisson noise with $\lambda = 1$; 2) Gaussian noise with zero mean and unit variance; 3) 0.2 fixed bias.
- 2) *Subsequence Level Anomaly*: Two subsequences anomalies are inserted, starting at a random time step, with random length sampled from even distribution of range $10 \sim 28$. For each anomaly subsequence, a true signal is replaced by a sequence sampled from a Gaussian process with RBF kernel of unit variance, and 0.3 length-scale.

The generated Mackey-Glass time series is shown in Fig. 12. To illustrate the generated anomalies, we zoomed-in view three fragments out of the whole time series.

In this experiment, we find that Donut shows better performance when z_dim is lower. We enumerate Donut's performance on $z_dim = \{2, \dots, 10\}$, and we find Donut performs best when $z_dim = 5$. For SISVAE-p, we find the performance is relatively stable when we change z_dim . Hence, we use the same setup as discussed in Section V-A. The ROC and PRC are shown in Fig. 13. In conclusion, benefiting from the sequential network structure of SISVAE, and the proposed smoothness prior regularizer, SISVAE-p exhibits better anomaly detection performance for univariate time series.


VI. CONCLUSION

We have presented a deep generative model-based method, that is, SISVAE (and the variants) for robust estimation and anomaly detection for multidimensional correlated time series. To capture “normal” patterns from anomaly-contaminated time series, we propose a novel variational smoothness regularizer, which provides robustness by placing a penalty on nonsmooth output of the generative model. Furthermore, we discuss two decision criteria for anomaly detection: reconstruction probability and reconstruction error. Systemic experiments are conducted on both synthetic and real-world data sets, which demonstrate that the proposed method outperforms state-of-the-art unsupervised anomaly detection models for multidimensional time series. For future work, we are interested in exploring heavy-tailed distribution such as Student- t distribution as the output of the generative model.

REFERENCES

- [1] X. Li, Z. Li, J. Han, and J.-G. Lee, “Temporal outlier detection in vehicle traffic data,” in *Proc. IEEE 25th Int. Conf. Data Eng.*, Mar. 2009, pp. 1319–1322.
- [2] M. Hauskrecht, I. Batal, M. Valko, S. Visweswaran, G. F. Cooper, and G. Clermont, “Outlier detection for patient monitoring and alerting,” *J. Biomed. Informat.*, vol. 46, no. 1, pp. 47–55, Feb. 2013.
- [3] H. Xu *et al.*, “Unsupervised anomaly detection via variational auto-encoder for seasonal KPIs in Web applications,” in *Proc. World Wide Web Conf. World Wide Web (WWW)*, 2018, pp. 187–196.
- [4] V. Jyothsna, V. R. Prasad, and K. M. Prasad, “A review of anomaly based intrusion detection systems,” *Int. J. Comput. Appl.*, vol. 28, no. 7, pp. 26–35, 2011.
- [5] D. J. Hill and B. S. Minsker, “Anomaly detection in streaming environmental sensor data: A data-driven modeling approach,” *Environ. Model. Softw.*, vol. 25, no. 9, pp. 1014–1022, Sep. 2010.
- [6] J. C. Yan, C. H. Tian, J. Huang, and F. Albertao, “Incremental dictionary learning for fault detection with applications to oil pipeline leakage detection,” *Electron. Lett.*, vol. 47, no. 21, pp. 1198–1199, 2011.
- [7] J.-Y. Kim, S.-J. Bu, and S.-B. Cho, “Zero-day malware detection using transferred generative adversarial networks based on deep autoencoders,” *Inf. Sci.*, vols. 460–461, pp. 83–102, Sep. 2018.
- [8] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly detection: A survey,” *ACM Comput. Surv.*, vol. 41, no. 3, p. 15, 2009.
- [9] C. C. Aggarwal, “Outlier analysis,” in *Data Mining*. Cham, Switzerland: Springer, 2015, pp. 237–263.
- [10] M. Gupta, J. Gao, C. C. Aggarwal, and J. Han, “Outlier detection for temporal data: A survey,” *IEEE Trans. Knowl. Data Eng.*, vol. 26, no. 9, pp. 2250–2267, Sep. 2014.
- [11] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly detection for discrete sequences: A survey,” *IEEE Trans. Knowl. Data Eng.*, vol. 24, no. 5, pp. 823–839, May 2012.
- [12] R. S. Tsay, “Outliers in multivariate time series,” *Biometrika*, vol. 87, no. 4, pp. 789–804, Dec. 2000.
- [13] S. Basu and M. Meckesheimer, “Automatic outlier detection for time series: An application to sensor data,” *Knowl. Inf. Syst.*, vol. 11, no. 2, pp. 137–154, Feb. 2007.
- [14] S. M. Erfani, S. Rajasegarar, S. Karunasekera, and C. Leckie, “High-dimensional and large-scale anomaly detection using a linear one-class SVM with deep learning,” *Pattern Recognit.*, vol. 58, pp. 121–134, Oct. 2016.
- [15] R. Laxhammar, G. Falkman, and E. Sviestins, “Anomaly detection in sea traffic—a comparison of the Gaussian mixture model and the kernel density estimator,” in *Proc. 12th Int. Conf. Inf. Fusion (FUSION)*, Jul. 2009, pp. 756–763.
- [16] J. Chung, K. Kastner, L. Dinh, K. Goel, A. C. Courville, and Y. Bengio, “A recurrent latent variable model for sequential data,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 2980–2988.
- [17] F. Locatello *et al.*, “Challenging common assumptions in the unsupervised learning of disentangled representations,” in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 4114–4124.


- [18] S. Zhao, H. Ren, A. Yuan, J. Song, N. Goodman, and S. Ermon, "Bias and generalization in deep generative models: An empirical study," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 10792–10801.
- [19] C. Shang, X. Huang, J. A. K. Suykens, and D. Huang, "Enhancing dynamic soft sensors based on DPLS: A temporal smoothness regularization approach," *J. Process Control*, vol. 28, pp. 17–26, Apr. 2015.
- [20] Y. Cai, H. Tong, W. Fan, and P. Ji, "Fast mining of a network of coevolving time series," in *Proc. SIAM Int. Conf. Data Mining*, Philadelphia, PA, USA: SIAM, Jun. 2015, pp. 298–306.
- [21] W. Enders, *Applied Econometric Time Series*. Hoboken, NJ, USA: Wiley, 2008.
- [22] Y. Zhou, H. Zou, R. Arghandeh, W. Gu, and C. J. Spanos, "Non-parametric outliers detection in multiple time series a case study: Power grid data analysis," in *Proc. AAAI*, 2018, pp. 4605–4612.
- [23] Y. Cai, H. Tong, W. Fan, P. Ji, and Q. He, "Facets: Fast comprehensive mining of coevolving high-order time series," in *Proc. 21st ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2015, pp. 79–88.
- [24] N. Görmitz, M. Braun, and M. Kloft, "Hidden Markov anomaly detection," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 1833–1842.
- [25] L. Li, J. McCann, N. S. Pollard, and C. Faloutsos, "DynaMMo: Mining and summarizing of coevolving sequences with missing values," in *Proc. 15th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining (KDD)*, 2009, pp. 507–516.
- [26] L. Xiong, X. Chen, and J. Schneider, "Direct robust matrix factorization for anomaly detection," in *Proc. IEEE 11th Int. Conf. Data Mining*, Dec. 2011, pp. 844–853.
- [27] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," 2013, *arXiv:1312.6114*. [Online]. Available: <http://arxiv.org/abs/1312.6114>
- [28] J. An and S. Cho, "Variational autoencoder based anomaly detection using reconstruction probability," *Special Lecture IE*, vol. 2, pp. 1–18, Dec. 2015.
- [29] J. L. Elman, "Finding structure in time," *Cognit. Sci.*, vol. 14, no. 2, pp. 179–211, 1990.
- [30] M. Soelch, J. Bayer, M. Ludersdorfer, and P. van der Smagt, "Variational inference for on-line anomaly detection in high-dimensional time series," 2016, *arXiv:1602.07109*. [Online]. Available: <http://arxiv.org/abs/1602.07109>
- [31] G. Kitagawa and W. Gersch, *Smoothness Priors Analysis of Time Series*, vol. 116. Springer, 2012.
- [32] D. P. Kingma and M. Welling, "Stochastic gradient VB and the variational auto-encoder," in *Proc. 2nd Int. Conf. Learn. Represent. (ICLR)*, vol. 19, 2014.
- [33] O. Fabius and J. R. van Amersfoort, "Variational recurrent auto-encoders," 2014, *arXiv:1412.6581*. [Online]. Available: <http://arxiv.org/abs/1412.6581>
- [34] J. Bayer and C. Osendorfer, "Learning stochastic recurrent networks," 2014, *arXiv:1411.7610*. [Online]. Available: <http://arxiv.org/abs/1411.7610>
- [35] M. Fraccaro, S. K. Sønderby, U. Paquet, and O. Winther, "Sequential neural models with stochastic layers," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 2199–2207.
- [36] K. Cho *et al.*, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," 2014, *arXiv:1406.1078*. [Online]. Available: <http://arxiv.org/abs/1406.1078>
- [37] I. Higgins *et al.*, "Beta-VAE: Learning basic visual concepts with a constrained variational framework," in *Proc. Int. Conf. Learn. Represent.*, 2017, p. 6.
- [38] P. J. Werbos, "Backpropagation through time: What it does and how to do it," *Proc. IEEE*, vol. 78, no. 10, pp. 1550–1560, Oct. 1990.
- [39] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2013, pp. 1310–1318.
- [40] C. Zhou and R. C. Paffenroth, "Anomaly detection with robust deep autoencoders," in *Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining (KDD)*, 2017, pp. 665–674.
- [41] A. Doucet, N. De Freitas, and N. Gordon, "An introduction to sequential Monte Carlo methods," in *Sequential Monte Carlo Methods in Practice*. New York, NY, USA: Springer, 2001, pp. 3–14.
- [42] A. Doucet, S. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statist. Comput.*, vol. 10, no. 3, pp. 197–208, 2000.
- [43] J. S. Liu and R. Chen, "Sequential Monte Carlo methods for dynamic systems," *J. Amer. Stat. Assoc.*, vol. 93, no. 443, pp. 1032–1044, Aug. 1998.
- [44] J. Davis and M. Goadrich, "The relationship between precision-recall and ROC curves," in *Proc. 23rd Int. Conf. Mach. Learn. (ICML)*, 2006, pp. 233–240.
- [45] T. Saito and M. Rehmsmeier, "The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets," *PLoS ONE*, vol. 10, no. 3, 2015, Art. no. e0118432.
- [46] P. Galeano, D. Peña, and R. S. Tsay, "Outlier detection in multivariate time series by projection pursuit," *J. Amer. Stat. Assoc.*, vol. 101, no. 474, pp. 654–669, Jun. 2006.
- [47] J. D. Hamilton, *Time Series Analysis*, vol. 2. Princeton, NJ, USA: Princeton Univ. Press, 1994.
- [48] A. Paszke *et al.*, "Pytorch: An imperative style, high-performance deep learning library," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 8024–8035.
- [49] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*. [Online]. Available: <http://arxiv.org/abs/1412.6980>
- [50] M. A. Álvarez, L. Rosasco, and N. D. Lawrence, "Kernels for vector-valued functions: A review," *Found. Trends Mach. Learn.*, vol. 4, no. 3, pp. 195–266, 2012.
- [51] GPy. (2012). *GPy: A Gaussian Process Framework in Python*. Accessed: 2012. [Online]. Available: <https://github.com/SheffieldML/GPy>
- [52] E. M. Stewart, A. Liao, and C. Roberts, "Open μ PMU: A real world reference distribution micro-phaser measurement unit data set for research and application development," *IEEE Power Eng. Lett.*, Oct. 2016. [Online]. Available: <https://gig.lbl.gov/publications/open-pmu-real-world-reference>
- [53] R. Baeza-Yates *et al.*, *Modern Information Retrieval*. New York, NY, USA: Addison-Wesley, 2011.
- [54] P. Burrage and A. M. Robert Taylor, "Additive outlier detection via extreme-value theory," *J. Time Ser. Anal.*, vol. 27, no. 5, pp. 685–701, Sep. 2006.
- [55] A. Siffer, P.-A. Fouque, A. Termier, and C. Largouet, "Anomaly detection in streams with extreme value theory," in *Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining (KDD)*, 2017, pp. 1067–1075.
- [56] M. C. Chuah and F. Fu, "ECG anomaly detection via time series analysis," in *Proc. Int. Symp. Parallel Distrib. Process. Appl.* Berlin, Germany: Springer, 2007, pp. 123–135.
- [57] M. Mackey and L. Glass, "Oscillation and chaos in physiological control systems," *Science*, vol. 197, no. 4300, pp. 287–289, Jul. 1977.
- [58] H. Kantz and T. Schreiber, *Nonlinear Time Series Analysis*, vol. 7. Cambridge, U.K.: Cambridge Univ. Press, 2004.



Longyuan Li received the B.E. degree in electronic engineering from the Huazhong University of Science and Technology, Wuhan, China, in 2013. He is currently pursuing the Ph.D. degree with the Department of Electronic Engineering, Shanghai Jiao Tong University, Shanghai, China.

His research interests include machine learning and data mining. He focuses on probabilistic models and Bayesian nonparametric models, for sequential data, such as multidimensional time series. He is also interested in deep probabilistic graphical models,

and uncertainty estimation.



Junchi Yan (Member, IEEE) received the Ph.D. degree in electrical engineering from Shanghai Jiao Tong University, Shanghai, China, in 2015.

He has been a Research Staff Member and a Principal Scientist with IBM Research, Shanghai, where he started his career, since April 2011. He is currently an Associate Professor with the Department of Computer Science and Engineering, Shanghai Jiao Tong University. His research interests are machine learning and computer vision.

Dr. Yan serves as the Area Chair for ICPR 2020, CVPR 2021, Senior PC for CIKM 2019, and an Associate Editor for IEEE ACCESS.



Haiyang Wang received the B.E. and Ph.D. degrees in electrical engineering from Shanghai Jiao Tong University, Shanghai, China, in 2013 and 2019, respectively.

His research area includes machine learning and data mining, especially massive spatial-temporal data mining and intelligent transportation.



Yaohui Jin (Member, IEEE) was a Technical Staff Member with Bell Labs Research China, Shanghai, China. He joined Shanghai Jiao Tong University, Shanghai, in 2002, where he is currently a Professor with the State Key Laboratory of Advanced Optical Communication Systems and Networks and the Deputy Director of Network and Information Center. He is the Founder of OMNILab, Shanghai, which is an open innovation lab focusing on data analysis. He has authored or coauthored more than 100 technical articles in leading conferences and journals. He is

the owner of over ten patents. In 2014, OMNILab won the champion of CCF national big data challenge among nearly 1000 teams, and was the champion of the Shanghai open data innovation and creation competition. He has served over ten technical committees. He enthuses public service and science popularization, actively promotes crowd engaged innovation, and interdisciplinary collaboration. His research interests include civic engagement and open innovation, cloud computing network architecture, and streaming data analysis.