UNSUPERVISED ANOMALY DETECTION FOR CONTAINER CLOUD VIA BILSTM-BASED VARIATIONAL AUTO-ENCODER

22-27 MAY 2022
ICASSP 2022
SINGAPORE

Paper #2451

Yulong Wang*†, Xingshu Chen*, Qixu Wang*, Run Yang†, Bangzhou Xin†
Sichuan University*
Institute of Computer Application, China Academy of Engineering Physics†

Introduction

Motivation

- The appearance of Linux container technology has profoundly changed the development and deployment of multi-tier distributed applications.
- The imperfect system resource isolation features and the kernel-sharing mechanism will introduce significant security risks to the container-based cloud.

Problem Formulation

Definition 1. ***Sequential Behaviours*** $\mathcal{S} = \{id, s, period\}$ *is defined as collected system call sequences for container id during a specific period.* $s = \langle s_1, s_2, \ldots, s_L \rangle : \forall k, 1 \leq k \leq L$ *and* $0 \leq s_k \leq N$, *where* L *represents the length of the captured sequence,* N *is the number of system call types, and* s_k *is the system call number with a unique integer.*

Problem 1. *Given a Sequential Behaviour* $\mathcal{S}$ *from a container, we aim to construct a deep generative network (i.e., real-time unsupervised anomaly detector)* $\mathcal{G}$ *consisting of an encoder* $\mathcal{E}$ *and a decoder* $\mathcal{D}$. *The latent variable* z *is learned from the visible input* x *by* $\mathcal{E}(x)$, *and the reconstructed sequence* $\hat{x}$ *is generated by* $\mathcal{D}(z)$. *The anomaly score function* $f : loss(x, \hat{x}) \mapsto \delta$ *can be quantified by the loss between reconstructed sequence and the input sequence.*

Contribution

- A robust and real-time unsupervised anomaly detection system is proposed termed *Pudding* for container cloud using system call sequences.
- Our method is applicable to the environment without a large number of labeled samples and has a strong perception of the unforeseen normal patterns.
- ***Pudding*** can easily be integrated into the container cloud platform as a security service without any hardware or kernel modifications while ensuring the transparency of container services.

Proposed Method

- The ***Pudding*** integrates data collecting, preprocessing, generative modeling, and anomaly detection modules in a non-intrusive manner. The overview of our proposed system:

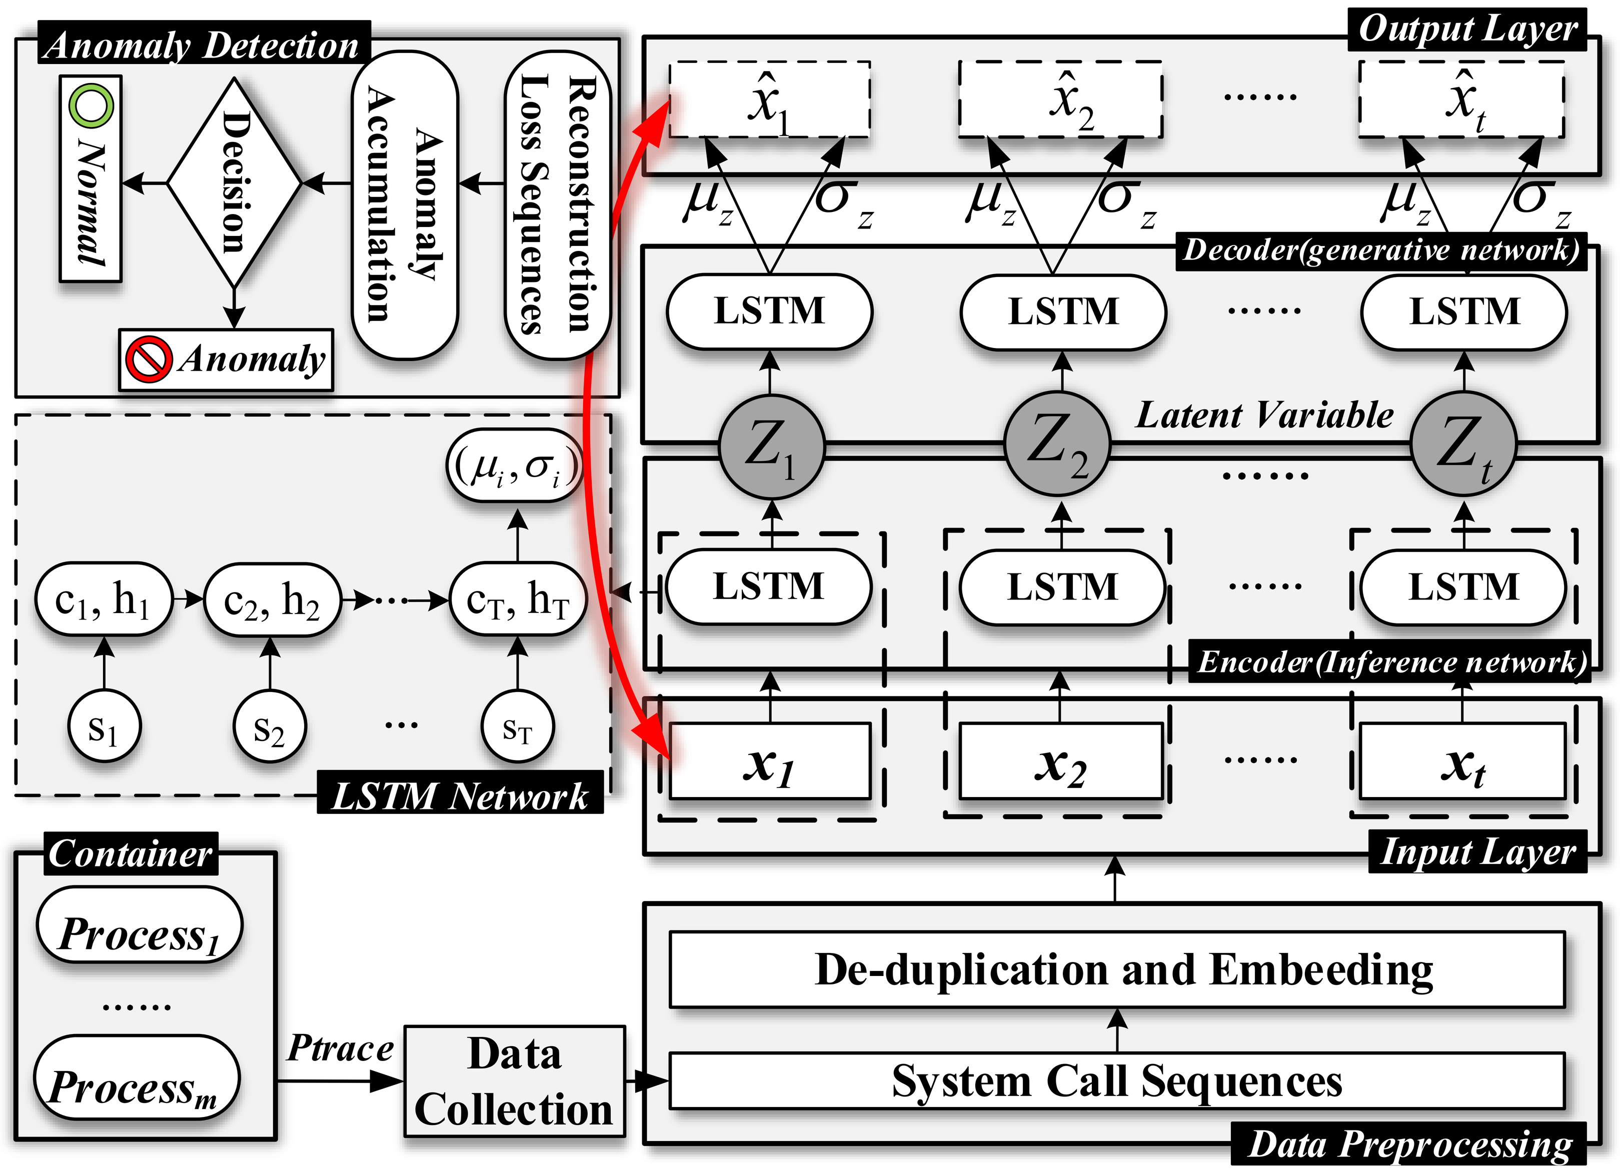

- **Data Collection**: The data collection module is implemented by a Linux system call termed ***ptrace***, which automatically senses the creation, operation, and extinction of the container in real-time.
- **Generative Network Structure:** The deep generative network G is constructed by *BiLSTM*-based variational auto-encoder, where one *BiLSTM* is utilized as an inferential network (encoder) to estimate the underlying probability distribution of the latent variable z, another *BiLSTM* is utilized as the generative network (decoder) to sample the reconstructed output $\hat{z}$ from the conditional probability distribution $p(x_t|z_t)$.
- **Anomaly Detection**: he anomaly detector is directly connected to the generation network $\mathcal{g}$ and receives the reconstruction probability served as our anomaly detection scores, which will be applied to make the final decision subsequently.

$$\mathcal{F}_{score}(x_i) = E_{z \sim q_\phi(z|x_i)}\left[\log p_\theta(z \mid x)\right] \approx \frac{1}{M}\sum_{m=1}^{M} \log p_\theta\left(x_i \mid z^{(m)}\right)$$

Experiments

Dataset

- The UNM public dataset and the system call sequences collected in real container environment.
- The dataset we built contains a total of 1,315,815 time-series data, of which anomalies account for 0.63%.

Results Analysis

- The trend of temporal reconstruction probabilities for normal and abnormal sequences, respectively.

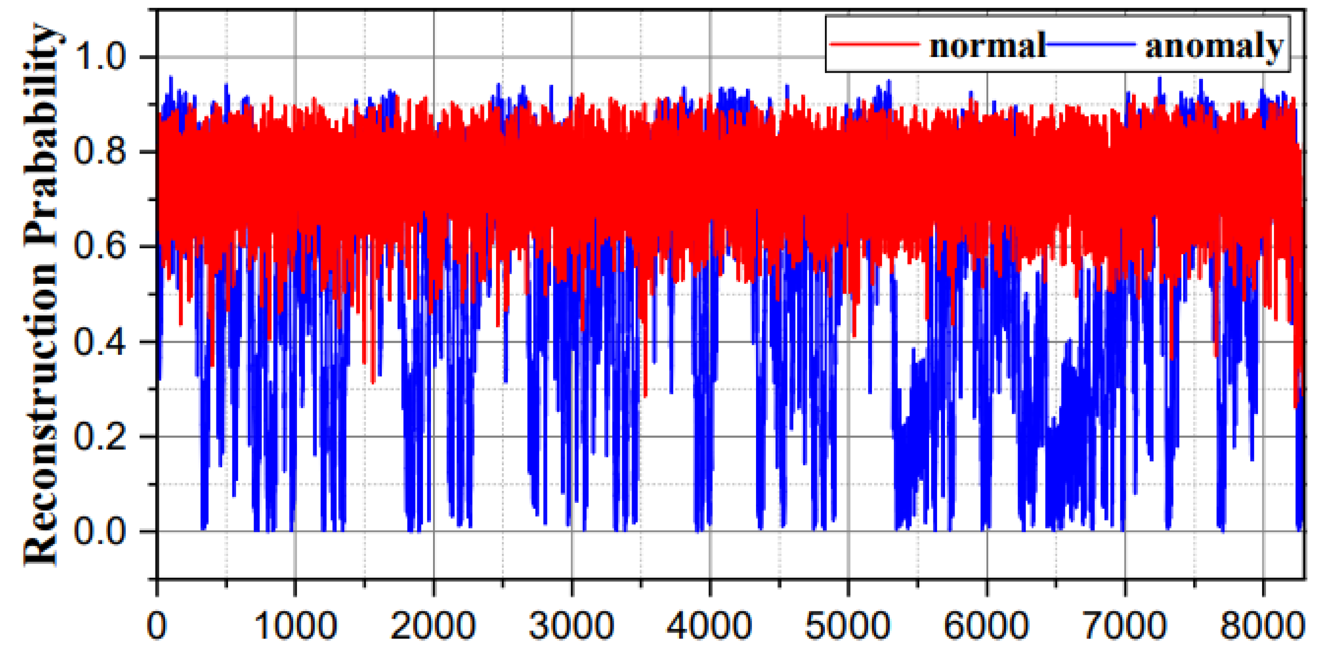

- Detection Performance comparison of different baseline approaches under various evaluation metrics.

Approaches	Evaluation Metrics			
	ACC/%	PRE/%	REC/%	F1/%
LOF [18]	79.27	70.89	99.35	82.74
One-class SVM [19]	72.99	69.34	82.23	75.27
Isolation Forests [18]	72.84	64.89	99.52	78.56
VAE-base [13]	77.94	71.32	93.49	80.91
Ours	**90.01**	**84.59**	**97.87**	**90.75**

Conclusion

Our proposed method leverages the generative characteristics of VAE to learn the robust representations of normal patterns by reconstruction probabilities while being sensitive to long-term dependencies. Our evaluations on real-world datasets show that our method achieves excellent detection performance without introducing much performance overhead to the container cloud.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant U19A2081, 61872059, 61502085