

Time-Series Event Prediction with Evolutionary State Graph

Wenjie Hu^{†*}, Yang Yang[†], Ziqiang Cheng[†], Carl Yang[‡], Xiang Ren[§]

[†]College of Computer Science and Technology, Zhejiang University, Hangzhou, China

^{*}Alibaba Cloud, Alibaba Group, Hangzhou, China

[‡]Emory University, Atlanta, GA, USA

[§]University of Southern California, Los Angeles, CA, USA

*dulin.hwj@alibaba-inc.com, †{aston2une, yangya, petecheng}@zju.edu.cn, ‡j.carlyang@emory.edu, §xiangren@usc.edu

ABSTRACT

The accurate and interpretable prediction of future events in time-series data often requires the capturing of representative patterns (or referred to as *states*) underpinning the observed data. To this end, most existing studies focus on the representation and recognition of states, but ignore the *changing transitional relations* among them. In this paper, we present *evolutionary state graph*, a dynamic graph structure designed to systematically represent the evolving relations (edges) among states (nodes) along time. We conduct analysis on the dynamic graphs constructed from the time-series data and show that changes on the graph structures (e.g., edges connecting certain state nodes) can inform the occurrences of events (i.e., time-series fluctuation). Inspired by this, we propose a novel graph neural network model, *Evolutionary State Graph Network (EvoNet)*, to encode the evolutionary state graph for accurate and interpretable time-series event prediction. Specifically, *Evolutionary State Graph Network* models both the node-level (state-to-state) and graph-level (segment-to-segment) propagation, and captures the node-graph (state-to-segment) interactions over time. Experimental results based on five real-world datasets show that our approach not only achieves clear improvements compared with 11 baselines, but also provides more insights towards explaining the results of event predictions.

KEYWORDS

Time series prediction, evolutionary state graph, graph networks

ACM Reference Format:

Wenjie Hu^{†*}, Yang Yang[†], Ziqiang Cheng[†], Carl Yang[‡], Xiang Ren[§]. 2021. Time-Series Event Prediction with Evolutionary State Graph. In *Proceedings of the Fourteenth ACM International Conference on Web Search and Data Mining (WSDM '21)*, March 8–12, 2021, Virtual Event, Israel. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3437963.3441827>

¹This work was done when the first author was studying in Zhejiang University. Corresponding author: Yang Yang, yangya@zju.edu.cn

²The code and data are publicly released at <https://github.com/zjunet/EvoNet>.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WSDM '21, March 8–12, 2021, Virtual Event, Israel

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8297-7/21/03...\$15.00

<https://doi.org/10.1145/3437963.3441827>




Figure 1: time-series can be segmented and recognized with several states (e.g., 1-4). Based on this, we construct the evolutionary state graph, where each node indicates a state and the edges represent their transitional relations across adjacent segments. Upon this, we develop EvoNet to further capture significant modes for effective time-series event prediction.

1 INTRODUCTION

The prediction of future events (e.g., anomalies) in time-series data has been an important task for temporal data mining [1, 12, 27, 28]. One common approach is latent state machines. For example, HMM [34], RNN [5] and their variants [11, 19] use series of latent representations to encode temporal data. However, such black-box encoding does not directly capture representative patterns (or referred to as “states”) that carry physical meanings in practice, such as walk or run in the observations from fitness-tracking devices. While these methods sometimes can obtain strong results, they are still sensitive to noises [37], provide poor interpretability, and are hard to debug when things go wrong. For this reason, many recent studies focus on discretizing time-series and finding the underlying states, with methods such as sequence clustering [17, 39], dictionaries (e.g. SAX [24, 37], BoP [25]) and shapelets [26, 35]. While effectively handling noises and providing better interpretability, they only recognize the states but *ignore the potential effects of relations among them*.

To jointly model the states and their relations, recent studies have started to explore the usage of *graph structures*, such as GCN-LSTM [27] and Time2Graph [9]. However, GCN-LSTM requires an explicit graph as input (e.g., in-app action graph), which is difficult to directly get from general time-series data. Time2Graph uses shapelets to discover states and relations, but it only computes a single static graph over the whole timeline, despite the fact that the state relations might change over time (e.g., node-level dynamics and graph-level migration, cf. Section 3.1 for details). To the best of

our knowledge, no existing studies have successfully captured and modeled the time-varying relations among the time-series states.

In this work, we observe that time-series are often affected by the joint influence of different states, and in particular, the *change of relations among states*. For example, in the sequential observations from fitness-tracking devices, *stopping exercise* from an *intense run* may cause the *fainting* event, while the monitoring data will look normal if one *stops exercise* from *jogging*; from online shopping records, a sudden interest change from *electronics* to *cosmetics* might be more suspicious than a smooth one from *cosmetics* to *fashion*.

Motivated by such observations, we propose a novel framework for time-series event prediction, by constructing and modeling a dynamic graph structure as shown in Figure 1. Following existing studies [9, 25, 26, 37], we model time-series based on the underlying states. However, to preserve more information from the original time-series data, we model each *time-series segment* as belonging to multiple states with different *recognition weights*, and leverage a directed graph to model the *transitional relations* among states between adjacent segments. Since the graph evolves along the time-series, we refer to it as an *evolutionary state graph*. Our empirical observations find that: 1) time-series evolution can be translated into different levels of graph dynamics; 2) when an event occurs, the time-series fluctuation can be expressed as the migration of graph structure, in particular, the dynamics of some edges connecting certain states (Section 3.1).

Despite the insights provided by our empirical observations, there still remains the challenge of how to quantitatively leverage the evolutionary state graph to improve the performance of time-series event prediction. Existing GNN models only consider a static graph or node-level dynamics [9, 23, 27, 31], which cannot be directly used for learning with our evolutionary state graph. In light of this, we propose a novel GNN model, *Evolutionary State Graph Network* (EvoNet), to further model the graph-level propagation and node-graph interactions with a temporal attention mechanism. The learned representations are then fed into an end-to-end model for time-series event prediction (Section 3.2).

To validate the effectiveness of EvoNet, we conduct experiments on five real-world datasets. Our experimental results demonstrate the superiority of EvoNet over 11 state-of-the-art baselines on time-series event prediction (Section 4.4). We further conduct comprehensive ablation and hyper-parameter studies to validate the effectiveness of our proposed method (Section 4.5). Finally, we demonstrate the insights towards prediction explanation by visualizing EvoNet and its evolutionary state graph (Section 4.6).

The main contributions of this work are summarized as follows:

- Through real-world data analysis, we find the time-varying relations among states important for time-series event prediction.
- We propose the evolutionary state graph to capture the dynamic relations among states, and develop EvoNet to improve the performance of event prediction based on such graphs.
- We conduct extensive experiments on five datasets to demonstrate that our method can both make more accurate predictions, and provide more insight towards explaining them.

2 BACKGROUND AND PROBLEM

Time-series event prediction. We consider the task of predicting future events in a given time-series sequence, following similar definition in previous work [1, 12, 27, 28]. Each time-series sequence with T chronologically paired segments can be represented as

$$\langle \mathbf{X}_{1:T}, \mathbf{Y}_{1:T} \rangle = \{(\mathbf{X}_1, \mathbf{Y}_1), (\mathbf{X}_2, \mathbf{Y}_2), \dots, (\mathbf{X}_T, \mathbf{Y}_T)\},$$

where $\mathbf{X}_t \in \mathbb{R}^{r \times d}$ and $\mathbf{Y}_t \in \mathbb{Z}$ denote a *time-series segment* [2] and the *observed event* in the corresponding time (e.g., anomalies), respectively. Each segment \mathbf{X}_t is a contiguous subsequence, i.e., $\mathbf{X}_t = \{\mathbf{x}_1, \dots, \mathbf{x}_\tau\}$, where $\mathbf{x}_i \in \mathbb{R}^d$ is a d -dimensional observation at the i -th time unit; segment length τ is a hyper-parameter which indicates certain physical meanings (e.g. 24 hours). If a time-series sequence can be divided by T segments of equal length τ , we then have $\langle \mathbf{X}_{1:T}, \mathbf{Y}_{1:T} \rangle = \{(\{\mathbf{x}_{\tau \times t+1}, \dots, \mathbf{x}_{\tau \times t+\tau}\}, \mathbf{Y}_t)_{0 \leq t < T}\}$. In this work, we aim to predict the future event \mathbf{Y}_{T+1} via discovering time-series states behind $\langle \mathbf{X}_{1:T}, \mathbf{Y}_{1:T} \rangle$ and modeling their dynamic relations.

State. A *state* v is a segment that indicates a representative pattern in the time-series sequence, denoted as $\Theta_v \in \mathbb{R}^{r \times d}$. In our study, we adopt existing methods (e.g., Symbolic Aggregate Approximation [24, 37], Bag of Patterns [25], Shapelets [26, 35], sequence clustering [17, 39]) for recognizing interpretable states from time-series data (e.g., symbolic values, shapes or clusters), which are shown to be effective in handling noises and providing good interpretability.

Segment-to-state representation. Once the states have been recognized, one can then models each time-series segment \mathbf{X}_t as a composition of states—i.e., quantify the *recognition weight* of each state for a segment to characterize the segment-state associations. Formally, given a segment \mathbf{X}_t and a state Θ_v , the recognition weight $P(\Theta_v|\mathbf{X}_t)$ is a measurement of similarity, defined as follows.

$$P(\Theta_v|\mathbf{X}_t) = \frac{\max([\mathcal{D}(\mathbf{X}_t, \Theta_v)]_{v \in \mathcal{V}}) - \mathcal{D}(\mathbf{X}_t, \Theta_v)}{\max([\mathcal{D}(\mathbf{X}_t, \Theta_v)]_{v \in \mathcal{V}}) - \min([\mathcal{D}(\mathbf{X}_t, \Theta_v)]_{v \in \mathcal{V}})}, \quad (1)$$

where $\mathcal{D}(\mathbf{X}_t, \Theta_v)$ can be formalized as the Euclidean Distance or other distances based on different time-series representation and state recognition methods. The smaller this distance, the higher the weight $P(\Theta_v|\mathbf{X}_t)$.

3 EVONET FRAMEWORK

In this section, we present a novel framework for time-series event prediction. We name the proposed framework *Evolutionary State Graph Network* (EvoNet), as it transforms the time-series into a dynamic graph based on the states and recognition weights, and constructs a GNN-based neural network to capture significant correlations and improve the ability of event prediction.

3.1 Evolutionary State Graph

Inspired by existing models introduced in Section 2, we aim to leverage the underlying states for effective and interpretable modeling of time-series. A straightforward approach is to regard a time series as a sequence of the *most likely states* (for each segment in the sequence), and then model their sequential dependencies [1, 17?]. However, one segment may not belong only to a single state; rather it should be recognized as multiple states with different weights. To this end, one can adopt a multiscale recurrent network (MRNN) [32] to model a multidimensional sequence of state weights, but this method does not highlight the *transitions among the states*, which may essentially determine whether an event occurs. Therefore, in

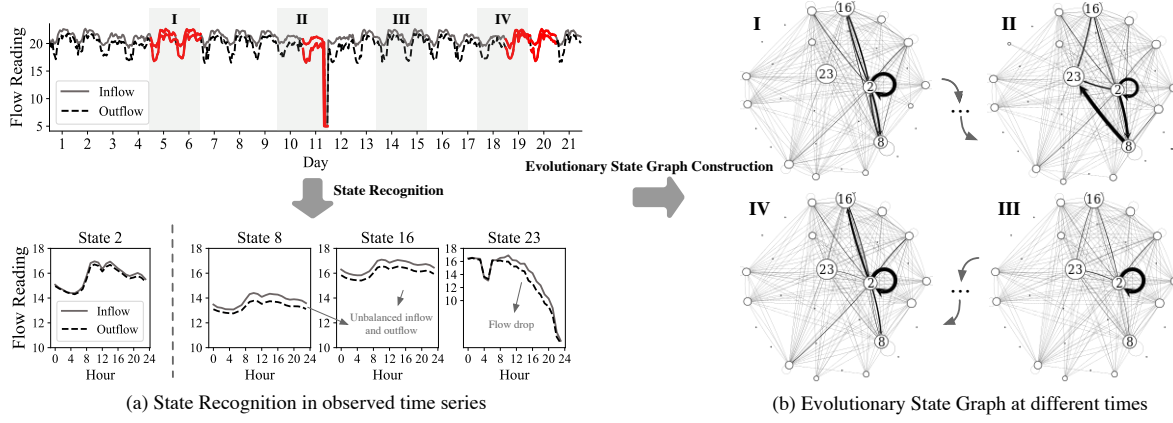


Figure 2: Example of an evolutionary state graph constructed to predict network anomalies. (a)-(b) present a case in the *NetFlow* dataset, such that the hourly inflow and outflow are recorded by the network monitor, while the red line indicates an anomaly occurred in this day. (a) visualizes four states recognized by EvoNet, while (b) presents the evolutionary state graph in four different intervals (I, II, III, IV marked in (a)). Each node in the graph indicates a state, while edge’s thickness indicates the weights of relations. The thicker the edge, the greater the weight. We can see that the evolutionary state graph can help to derive insights to analyze time-series, such as the fact that an unbalanced inflow and outflow or flow drop will lead to anomalies (e.g., state transitions: #2→#16, #8→#23).

this work we propose a novel dynamic graph structure to describe the relations among the states and explore how the dynamic shifts of states can reveal time-series evolution.

Evolutionary state graph. We define the *evolutionary state graph* as a sequence of weighted-directed graphs $\langle G^{(1:T)} \rangle$. Specifically, each graph is formulated as $G^{(t)} = \{\mathcal{V}, \mathcal{E}^{(t)}, \mathcal{M}^{(t)}\}$ to represent the transitions from the states of segments X_{t-1} to those of X_t . Each node in the graph indicates a state v ; each edge $e_{(v,v')} \in \mathcal{E}^{(t)}$ represents the *transitional relation* (or *relation* in short) from v to v' , along with the *transition weight* $m_{(v,v')}^{(t)} \in \mathcal{M}^{(t)}$. Assuming the state weights observed for each segment to be independent, the transition weights are computed by

$$m_{(v,v')}^{(t)} = \mathbf{P}(\Theta_{(v,v')} | X_{(t-1,t)}) = \mathbf{P}(\Theta_v | X_{t-1}) \times \mathbf{P}(\Theta_{v'} | X_t), \quad (2)$$

which is the joint weight that X_{t-1} is recognized to the state v , while X_t is recognized to the state v' .

Compared with existing time-series representations based on states, our evolutionary state graph preserves more information from the original data along the timeline through the modeling of multiple states in each segment and their changing transitional relations. It allows the subsequent model to be more powerful and provide richer interpretations in its predictions, while inheriting from state-based representations the robustness towards noises.

Real-world example and analysis of evolutionary state graph.

To demonstrate how the evolutionary state graph reveals the evolution of time-series and helps the prediction of events, we conduct an observational study on the *Netflow* dataset (cf. Section 4.1 for details of the dataset). As we can see from the case shown in Figure 2, when an anomaly event occurs, the state transitions (#2→#16) and (#2→#8) are more frequent at time I; similarly, the state transitions (#2→#8) and (#8→#23) are obvious at time II. These transitions reveal that the unbalanced inflow and outflow (state 8 and state 16), or flow drop (state 23), will cause anomalies of network devices. At time III, no anomaly occurs during this period. We can see that states primarily stay in #2. There is then an anomaly in the

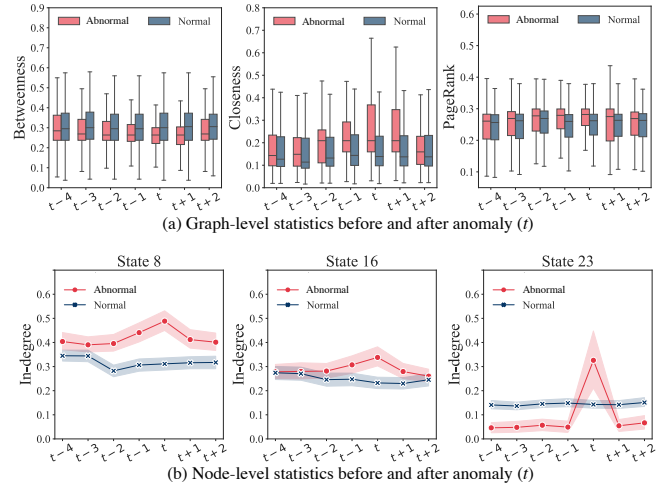


Figure 3: The statistics of the constructed evolutionary state graph in Figure 2. (a)-(b) present some statistics based on the graph-level measurements (e.g., betweenness [7], closeness [29], pagerank [30]) and node-level ones (in-degree) before and after anomaly t .

next immediate moment at time IV. Accordingly, we can see a clear increase of the state transition #2→#16.

In light of the observations, we conduct statistic analysis related to the constructed evolutionary state graph based on the abnormal samples (an anomaly occurs at time t) and normal samples (no anomaly occurs). The distributions of the different graph-level and node-level measurements at different times (before and after anomaly t) are visualized in Figure 3. From the figure, we can clearly see that when an anomaly occurs, the abnormal graph (red bar) tends to be denser; i.e., the betweenness scores gets lower, while the closeness scores gets higher. Figure 3b presents three typical states and compares their in-degree before and after anomaly t . We can see that the in-degrees of state 8 and 16, indicating the unbalanced inflow and outflow, gradually increase before t ; this illustrates that the network gradually becomes abnormal. The in-degree of state

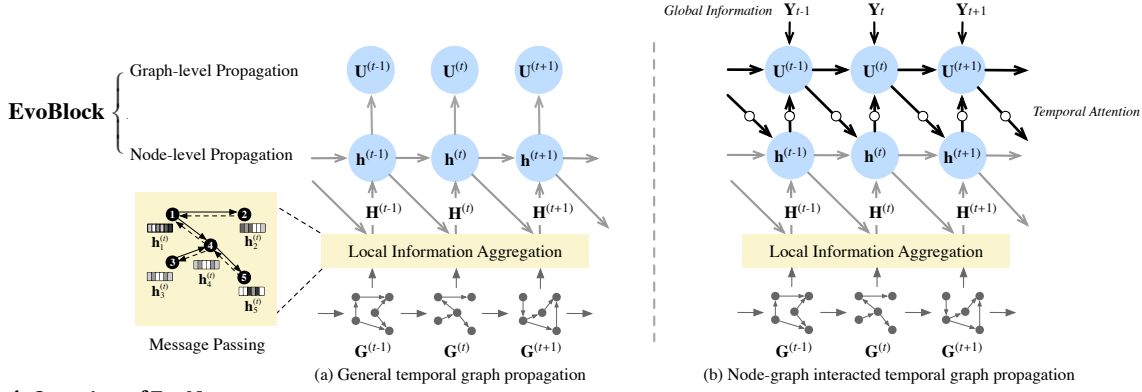


Figure 4: Overview of EvoNet. After building the evolutionary state graph on the raw time series, EvoNet conducts *local information aggregation* and *temporal graph propagation*. For local information aggregation, each node in graph $G^{(t)}$ has a feature vector $\mathbf{h}_v^{(t)}$. The solid edges indicate the passing messages, while the dashed edges indicate the feedback ones. The graph-level patterns $\mathbf{U}^{(t)}$ and node-level features $\mathbf{h}^{(t)}$ are then propagated by the recurrent *EvoBlock*, based on the aggregated intermediate representation $\mathbf{H}^{(t)}$: a) The general architecture for *EvoBlock* on the evolutionary state graph, where $\mathbf{U}^{(t)}$ is pooled from each $\mathbf{h}_v^{(t)}$, $v \in \mathcal{V}$; b) The architecture of EvoNet, where graph-level and node-level propagation influence each other, based on the temporal attention mechanism.

23 suddenly increases, indicating that the flow drop is an unexpected event. When no anomaly occurs, we can see that the normal evolutionary state graph (blue lines) generally remains unchanged.

Through the example, we show how the transformation of time-series into evolutionary state graphs allows us to capture the relations between states and their evolution. Meanwhile, we also learn that the graph-level and node-level evolutions can reveal different contextual information related to the time-series events: the node-level evolution reveals the states' skips when events occur, while the graph-level evolution presents the time-series migration. Intuitively, we shall capture these two levels of information simultaneously when learning with evolutionary state graphs.

3.2 Evolutionary State Graph Network

Overview. Motivated by Section 3.1, unlike most existing works [1, 26, 37] which model the independent effects of each state, we develop EvoNet to capture the following two types of information through the leverage of our evolutionary state graph:

- *Local structural influence*: the same state v will cause different observations when v is transmitted from different states. In other words, the relations among states matter. For example, *stopping exercise* from an *intense run* may cause fainting, while the monitoring data will look healthier if one *stops exercise* from *jogging*.
- *Temporal influence*: previous transitions of states will influence the current observed data. For example, (*intense run* \rightarrow \dots \rightarrow *stopping exercise*) and (*jogging* \rightarrow *jogging* \rightarrow \dots \rightarrow *stopping exercise*) lead to different fitness effects.

The above two types of influence can be naturally represented by the evolutionary state graph: the local structural influence is primarily determined by local-pairwise relations among nodes in each graph, while the temporal influence is determined by how relations evolve over different graphs. Inspired by Graph Neural Networks (GNN) [3], we model both the structural and temporal influences of evolutionary state graph by designing two mechanisms: *local information aggregation* and *temporal graph propagation*.

Figure 4 illustrates the overall structure of EvoNet. Given the observations $\langle \mathbf{X}_{1:T}, \mathbf{Y}_{1:T} \rangle$, we first recognize states for each segment

\mathbf{X}_t and construct the evolutionary state graph $\langle \mathbf{G}^{(1:T)} \rangle$. Next, we define a representation vector $\mathbf{h}_v^{(t)} \in \mathbb{R}^{|\mathbf{h}|}$ for each node v in graph $\mathbf{G}^{(t)}$ to encode v 's node-level patterns, and define a representation vector $\mathbf{U}^{(t)} \in \mathbb{R}^{|\mathbf{U}|}$ for $\mathbf{G}^{(t)}$ to encode the graph-level information. Based on this, EvoNet aggregates local structural information by means of message passing, and further incorporates temporal information using the recurrent *EvoBlock*. EvoNet then applies the learned representations (\mathbf{h}, \mathbf{U}) towards the prediction task.

Local information aggregation. In order to aggregate the local structural information in each $\mathbf{G}^{(t)}$, EvoNet aims to make two linked nodes share similar representations. To achieve this, we let each node representation $\mathbf{h}_v^{(t)}$ in $\mathbf{G}^{(t)}$ aggregate the *messages* of its neighbors, and thus compute its new representation vector. Initially, we let $\mathbf{h}_v^{(0)} = \Theta_v$. Recall that Θ_v is obtained from the state recognition on all segments, which records the time-series information of state v . Then, following the message-passing neural network (MPNN) [16] directly, we have the following aggregation scheme:

$$\mathbf{H}_v^{(t)} = \sum_{v' \in N(v)} \mathcal{F}_{\text{MP}}(\mathbf{h}_{v'}^{(t-1)}, e_{(v,v')}^{(t)}) \quad (3)$$

where $\mathbf{H}_v^{(t)}$ is the intermediate representation of node v following aggregation, which combines the messages from all neighbors $N(v)$ in the graph $\mathbf{G}^{(t)}$. The message function $\mathcal{F}_{\text{MP}}(\cdot, \cdot)$ can be implemented by many existing neural networks, such as GGNN[23]:

$$\mathcal{F}_{\text{MP}}(\mathbf{h}_{v'}^{(t-1)}, e_{(v,v')}^{(t)}) = W_{\text{MP}} \cdot \left[m_{(v,v')}^{(t)} \times \mathbf{h}_{v'}^{(t-1)} \right] + b_{\text{MP}} \quad (4)$$

where $m_{(v,v')}^{(t)} \times \mathbf{h}_{v'}^{(t-1)}$ is the passing message, while W_{MP} and b_{MP} are the learnable parameters, indicating the passing weight and bias. We also have other implementations for \mathcal{F}_{MP} , such as pooling, GCN [13], GraphSAGE [18], GAT [38], etc. Herein, we serve \mathcal{F}_{MP} as interchangeable modules in EvoNet and conduct experiments in Section 4.5 to analyze the effectiveness of different implementations.

Temporal graph propagation. In addition to aggregating the local structural information, previous transitions also influence current representations. Moreover, when events occur, the modes of the graph-level and node-level evolution will change (Section 3.1). Intuitively, we should capture these two kinds of temporal information

simultaneously. To achieve this, we design a recurrent block, named *EvoBlock*, to capture the evolving information in the evolutionary state graph. *EvoBlock* combines the local aggregated representation $\mathbf{H}_v^{(t)}$ and the past representation $(\mathbf{h}_v^{(t-1)}, \mathbf{U}^{(t-1)})$, formulated as

$$\mathbf{h}_v^{(t)}, \mathbf{U}^{(t)} := \mathcal{F}_{\text{recur}}(\mathbf{H}_v^{(t)}, \mathbf{h}_v^{(t-1)}, \mathbf{U}^{(t-1)}) \quad \text{for } v \in \mathcal{V} \quad (5)$$

where $\mathcal{F}_{\text{recur}}$ indicates a recurrent function that allows us to incorporate information from the previous timestamp in order to update current representations. When there are few messages from other nodes, i.e., $m_{(v',v)}^{(t)} \rightarrow 0$, $(\mathbf{h}_v^{(t)}, \mathbf{U}^{(t)})$ will be more influenced by the previous $(\mathbf{h}_v^{(t-1)}, \mathbf{U}^{(t-1)})$. Otherwise, the messages will influence current representations more.

As shown in Figure 4a, most existing works implement $\mathcal{F}_{\text{recur}}$ using simple recurrent neural networks on node-level propagation (e.g., GGSNN [23] adopts GRU [11], GCN-LSTM [27] adopts LSTM [19], etc.). For the graph-level propagation $\mathbf{U}^{(t)}$, these methods simply pool the node-level representations, i.e., $\mathbf{U}^{(t)} = \sum_{v \in \mathcal{V}} \mathbf{h}_v^{(t)}$. However, in our empirical observations (Section 3.1), both the graph and nodes in the evolutionary state graph will present different temporal information when events occur. In order to improve the ability of event prediction, $\mathcal{F}_{\text{recur}}$ shall consider the contextual information of previous events $\mathbf{Y}_{1:T}$ when modeling the graph-level propagation, and then influence the node-level representations via the node-graph interactions. Accordingly, events are generally scattered in the timeline; thus, we propose a temporal attention mechanism for capturing significant temporal information in node-graph interactions. More specifically, as shown in Figure 4b, we have

$$\mathcal{F}_{\text{recur}}(\mathbf{H}_v^{(t)}, \mathbf{h}_v^{(t-1)}, \mathbf{U}^{(t-1)}) = \begin{cases} \mathbf{h}_v^{(t)} = \Phi_h(\mathbf{h}_v^{(t-1)}, \mathbf{H}_v^{(t)} \oplus \alpha_t \mathbf{U}^{(t-1)}) \\ \mathbf{U}^{(t)} = \Phi_U(\mathbf{U}^{(t-1)}, \mathbf{Y}_t \oplus \alpha_t \sum_{v \in \mathcal{V}} \mathbf{h}_v^{(t)}) \\ \alpha_t = \text{softmax}\left(W_\alpha \left(\mathbf{U}^{(t-1)} \oplus \sum_{v \in \mathcal{V}} \mathbf{H}_v^{(t)}\right)\right) \end{cases} \quad (6)$$

where “ \oplus ” indicates the concatenation operator. The current node-level representation is computed using the function $\Phi_h(\cdot, \cdot)$, based on the past representations $(\mathbf{h}_v^{(t-1)}, \mathbf{U}^{(t-1)})$ and current aggregations $\mathbf{H}_v^{(t)}$, while the current graph-level representation is computed by $\Phi_U(\cdot, \cdot)$ based on the past $\mathbf{U}^{(t-1)}$ and current event \mathbf{Y}_t , as well as all node representations $\mathbf{h}_v^{(t)}$. The attention score α_t re-weights the node-graph interaction of the t -th temporal step, which is computed based on the concatenated patterns of $\mathbf{U}^{(t-1)}$ and all aggregations $\sum_{v \in \mathcal{V}} \mathbf{H}_v^{(t)}$, under the learnable weight W_α . We use the softmax function to normalize α_t during different time steps.

Recurrent function $\Phi_*(\cdot, \cdot)$ smooths the two inputted vectors of each temporal step, and can be implemented using many existing approaches. Herein, we provide an example of $\Phi_h(\cdot, \cdot)$ implemented by LSTM. Formally, we have

$$\begin{aligned} \Phi_h(\mathbf{h}_v^{(t-1)}, \mathbf{H}_v^{(t)} \oplus \alpha_t \mathbf{U}^{(t-1)}) = & \begin{cases} \mathbf{t}^{(t)} = \mathbf{H}_v^{(t)} \oplus \alpha_t \mathbf{U}^{(t-1)} \\ \mathbf{F}^{(t)} = \sigma(W_F \cdot [\mathbf{h}_v^{(t-1)}, \mathbf{t}^{(t)}] + b_F) & \mathbf{I}^{(t)} = \sigma(W_I \cdot [\mathbf{h}_v^{(t-1)}, \mathbf{t}^{(t)}] + b_I) \\ \mathbf{C}^{(t)} = \mathbf{F}^{(t)} \odot \mathbf{C}^{(t-1)} + \mathbf{I}^{(t)} \odot \tanh(W_C \cdot [\mathbf{h}_v^{(t-1)}, \mathbf{t}^{(t)}] + b_C) \\ \mathbf{O}^{(t)} = \sigma(W_O \cdot [\mathbf{h}_v^{(t-1)}, \mathbf{t}^{(t)}] + b_O) & \mathbf{h}_v^{(t)} = \mathbf{O}^{(t)} \odot \tanh(\mathbf{C}^{(t)}) \end{cases} \quad (7) \end{aligned}$$

where $\mathbf{F}^{(t)}$, $\mathbf{I}^{(t)}$ and $\mathbf{O}^{(t)}$ are forget gate, input gate and output gate respectively, while σ is a sigmoid activation function. The current node vectors are updated by receiving their own previous memory and current memory. In our experiments, we compare the performance of different methods for *EvoBlock* (Table 2).

End-to-End Model Learning. Thus, the representations $\mathbf{h}_v^{(t)}$ and $\mathbf{U}^{(t)}$ capture both the node-level and graph-level information respectively until the t -th temporal step, which can then be applied to predict the next event \mathbf{Y}_{t+1} . More specifically, we encode the current evolutionary state graph $\mathbf{G}^{(t)}$ into representation $\mathbf{h}_G^{(t)}$ based on the concatenated features of all $\mathbf{h}_v^{(t)}$ and $\mathbf{U}^{(t)}$, which can be formulated as

$$\mathbf{h}_G^{(t)} = \mathcal{F}_{\text{fc}}\left(\mathbf{U}^{(t)} \oplus \sum_{v \in \mathcal{V}} \mathbf{h}_v^{(t)}\right) \quad (8)$$

where \mathcal{F}_{fc} acts as a fully connected layer. We then learn a classifier, such as a neural network or XGBoost [8], which takes $\mathbf{h}_G^{(t)}$ as input and estimates the probability of the next event, $\mathbf{P}(\mathbf{Y}_{t+1} | \mathbf{h}_G^{(t)})$. To learn the parameters θ of the proposed *EvoNet* and classifier, we employ an end-to-end framework, based on the Adam optimization algorithm [22] to minimize the cross-entropy loss \mathcal{L} as follows:

$$\mathcal{L} = - \sum \hat{\mathbf{Y}}_{t+1} \log \mathbf{P}(\mathbf{Y}_{t+1} | \mathbf{h}_G^{(t)}) + (1 - \hat{\mathbf{Y}}_{t+1}) \log (1 - \mathbf{P}(\mathbf{Y}_{t+1} | \mathbf{h}_G^{(t)})) \quad (9)$$

where $\hat{\mathbf{Y}}_{t+1} \in \{0, 1\}$ is the ground truth that indicating whether a future event will occur. The procedure of state recognition and graph propagation are carried out step by step: we first recognize the states and construct an evolutionary state graph, then conduct the evolutionary state graph propagation to model the time-series. $|\mathcal{V}|$ -node graphs are constructed in T segments, such that the time complexity of each iteration is $O(T \times |\mathcal{V}|^2)$.

4 EXPERIMENTS

We apply our method to the prediction of upcoming events in time-series data, and aim to answer the following three questions:

- **Q1:** How does *EvoNet* perform on the time-series prediction task, compared with other baselines from the state-of-the-art?
- **Q2:** How does the proposed *EvoBlock* effectively bridge the graph-level and node-level information over time?
- **Q3:** How do different configurations, e.g., state number, segmentation length, implementation of state recognition and message passing, influence the performance?

4.1 Datasets

We employ five real-world datasets to conduct our experiments, including two public ones (*DJIA30* and *WebTraffic*) from Kaggle¹, and another three (*NetFlow*, *ClockErr* and *AbServe*) provided by China Telecom², State Grid³ and Alibaba Cloud⁴, respectively. Table 1 presents the overall dataset statistics.

DJIA 30 Stock Time Series (DJIA30). This dataset comes from Kaggle. It contains around 15K daily readings, each of which records four observations on a trading day: three kinds of trade price and a trade number. The task is to predict abnormal price volatility

¹An online community of data scientists and machine learners.

²A major mobile service provider in China.

³A major electric power company in China.

⁴The largest cloud service provider in Asia.

Table 1: Dataset statistics

Dataset	DJIA30	WebTraffic	NetFlow	ClockErr	AbServe
#(samples)	15,540	2,992,184	238,000	6,879,834	12,224
positive ratio(%)	19.5	28.2	8.6	14.9	1.5

(variance greater than 1.0) in the next week (five trading days) based on the most recent records from the past year (50 weeks). In total, we identify around 12K normal cases and 3K abnormal ones.

Web Traffic Time Series Forecasting (WebTraffic). This dataset comes from Kaggle. It contains around 3M daily readings, each of which records the number of views for a specific Wikipedia article. The task is to predict whether there will be a rapid growth (curve slope greater than 1.0) in the next month (30 days) based on the most recent records from the past 12 months. In total, we identify around 900K positive cases (rapid growth) and 2M negative ones.

Information Networks Supervision (NetFlow). This dataset is provided by China Telecom. It consists around 238K hourly readings, each of which records the hourly in- and out-flow of network devices. When an abnormal flow goes through the device ports, an alarm will be recorded. Our goal is to predict future anomalies (next day) based on records from the past 15 days. In total, we identify around 200K normal cases and 20K abnormal ones.

Watt-hour Meter Clock Error (ClockErr). This dataset is provided by the State Grid of China. It consists of around 6M weekly readings, each of which records the deviation time and delay of watt-hour meters. When the deviation time exceeds 120, the meter is marked as abnormal. Our goal is to predict anomalies in the next month based on records from the past 12 months. In total, we identify around 5M normal cases and 1M abnormal ones.

Abnormal Server Response (AbServe). This dataset is provided by Alibaba Cloud. It consists of around 12K server monitoring series, each of which records the minutely readings of different metrics (e.g., CPU, disk, memory, etc.). When a server fails to respond, the log will record the anomaly. Our goal is to predict anomalies in next 5 minutes based on records from the previous one hour. In total, we identify 11.8K normal cases and 0.2K abnormal ones.

4.2 Baseline Methods

We compare our proposed EvoNet with several groups of baselines:

Feature-based models. Several popular feature-based algorithms have been proposed for time-series analysis. In this paper, we choose some typical algorithms to compare with our model: Bag of Patterns (*BoP*) [25], Vector Space Model using SAX (*SAX-VSM*) [37] and Fast Shapelet (*FS*) [35]. These methods capture different state representations, which serve as features for event predictions.

Sequential models. Another typical group of algorithms interpret the time-series as a new sequence of states, and model their sequential dependencies. In this paper, we use several famous frameworks as baselines: switching-time-series model (*S-HMM*) [1] models the Markov dependencies of state sequences; multiscale recurrent neural network (*MRNN*) [32] takes the concatenated multi-source sequences ($X_t \oplus Y_t$) as input and learns one latent representation for prediction; hierarchical recurrent neural network (*HRNN*) [10] captures more correlations between X_t and Y_t , which conducts the same mechanism as Evoblock.

Graph-based models. Recently, many GNN-based works are proposed to model the (dynamic) graphs. In this paper, we choose several state-of-the-art algorithms as baselines to model the evolutionary state graph, and conduct the same approaches for event prediction as EvoNet: gated graph neural network (*GGNN*) [23] initializes the node vector $\mathbf{h}^{(0)}$ using a one-hot vector of the corresponding state; it conducts GGNN [23] for local message passing and only adopts a GRU structure [11] for node-level propagation. *GCN-LSTM* [27] uses states' patterns Θ to initialize the node vector $\mathbf{h}^{(0)}$; it conducts GCN [13] for local message passing and LSTM structure [19] for node-level propagation. *EvolveGCN* [31] is a dynamic graph neural network that builds a multi-layer framework to combine RNN and GCN; it also focuses on node-level propagation. *ST-MGCN* [15] is a spatiotemporal multi-graph convolution network, in which Y_t serve as contextual information for propagation. It directly fuses the contextual information into node-level representations rather than learning graph-level representations and modeling the node-graph interactions. *Time2Graph* [9] adopts shapelet to extract states; it aggregates the graphs at different times as a static graph and conduct DeepWalk [33] to learn graph's representations, which then serve as features for event predictions.

EvoNet variants. We also compare EvoNet with its derivatives by modifying some key components to see how they fare: 1) we sample the most possible state sequence (i.e., each segment is recognized with highest state weight) for each time-series, and directly use LSTM to model the new sequence without building and modeling the evolutionary state graph, denoted as *EvoNet w/o G*; 2) we build evolutionary state graph for time-series but model it without conducting temporal attention mechanism, denoted as *EvoNet w/o A*; 3) we conduct complete EvoNet for time-series modeling, denoted as *EvoNet*. Herein, EvoNet uses the state patterns Θ to initialize node vector $\mathbf{h}^{(0)}$ and conducts graph-level and node-level propagation for $\langle \mathbf{G}^{(1:T)} \rangle$. We implement state recognition and local message passing using Kmeans [21] and GGNN [23] respectively. We will study how different implementations influence the performance later in Section 4.5.

4.3 Implementation details

We conduct experiments on the five real-world datasets. We split the train/test set by 0.8 at the time line, such that preceding segments are used for training and the following ones are used for testing. We also split 10% samples from train set as validation set in order to avoid overfitting. We run all experiments on a single GPU with a batch size of 1000, and train our models for 100 iterations in total, starting with a learning rate of 0.001 and reducing it by a factor of 10 at every 20 iterations.

4.4 Performance Comparison

We compare the performance of EvoNet and other baselines in order to answer Q1, and also conduct ablation studies to answer Q2. For the binary event prediction tasks, we use F1 score and AUC as our evaluation metrics, due to the unbalanced positive ratio. All reports are the average results of five times repeated experiments, along with their standard deviations (see details in Table 2).

1. Feature-based models vs. others. We observe that all feature-based methods perform poorly, because they only capture the states

Table 2: Comparison of prediction performance on five real-world datasets (%). The bold text indicates the best performance among all methods, while the underline text indicates the second-best performance.

Datasets Models		DJIA30		WebTraffic		NetFlow		ClockErr		AbServe	
		F1-score	AUC	F1-score	AUC	F1-score	AUC	F1-score	AUC	F1-score	AUC
Feature-based models	BoP [25]	24.92±0.40	50.92±0.19	44.31±0.33	66.87±0.09	54.01±0.89	81.36±0.45	60.01±0.49	85.20±0.38	42.59±0.60	70.22±0.37
	FS [35]	24.38±0.97	50.55±0.42	43.89±0.76	66.96±0.23	52.84±1.63	79.21±0.69	58.34±0.83	84.32±0.71	46.95±0.91	72.04±0.56
	SAX-VSM [37]	26.06±0.45	51.42±0.20	44.66±0.49	67.63±0.15	61.11±1.44	83.95±0.71	62.44±0.65	85.97±0.64	47.98±0.75	73.88±0.49
Sequential models	S-HMM [1]	25.20 ±0.48	51.14±0.20	43.09±0.41	66.54±0.12	58.05±0.87	81.89±0.49	59.55±0.60	84.99±0.61	48.71±0.60	73.65±0.38
	MRNN [32]	21.20±0.42	49.39±0.19	44.43±0.57	67.51±0.17	69.15±0.93	85.11±0.49	60.95±0.87	85.06±0.76	47.08±0.69	72.21±0.46
	HRNN [10]	26.43±0.87	52.66±0.29	45.79±0.82	68.27±0.26	72.42±1.25	91.19±0.57	61.14±1.19	85.38±0.83	50.93±0.78	78.13±0.51
Graphical models	GGSNN [23]	23.72±0.91	51.56±0.31	43.30±1.25	67.14±0.38	72.92±1.54	90.38±0.68	64.96±1.13	86.81±0.84	48.79±0.83	74.08±0.50
	GCN-LSTM [27]	25.76±0.85	52.66±0.30	45.67±0.90	68.15±0.29	75.05±1.38	91.43±0.60	65.65±1.04	87.03±0.78	50.95±0.80	78.14±0.50
	EvolveGCN [31]	26.16±1.24	53.01±0.55	45.90±1.58	68.38±0.41	75.21±2.47	91.56±1.08	65.82±1.92	87.17±1.29	50.63±1.37	78.01±0.98
	ST-MGCN [15]	26.93±0.97	53.39±0.39	45.96±0.91	68.74±0.27	77.79±1.40	91.95±0.64	66.61±1.11	87.78±0.83	<u>51.21±0.85</u>	<u>78.33±0.52</u>
	Time2Graph [9]	26.50±0.91	53.28±0.39	<u>46.03±1.12</u>	<u>68.74±0.43</u>	76.94±1.83	91.61±0.64	67.01±1.46	88.00±1.23	50.50±1.02	77.87±0.98
Our models	EvoNet w/o G	25.81±0.80	52.67±0.33	45.66±0.85	68.45±0.38	74.92±1.42	91.40±0.63	65.71±0.99	87.10±0.80	50.89±0.80	78.10±0.50
	EvoNet w/o A	<u>29.11±0.83</u>	<u>54.47±0.37</u>	45.95±0.91	68.55±0.25	<u>79.37±1.43</u>	<u>92.45±0.66</u>	69.21±1.17	89.92±0.80	51.20±0.81	78.10±0.50
	EvoNet	30.47±0.93	55.07±0.39	47.02±0.95	69.03±0.27	80.25±1.43	92.67±0.65	<u>68.62±1.21</u>	<u>89.76±0.82</u>	53.44±0.87	79.97±0.52

as features but ignore the influence of relations. *FS* is unstable relatively and *SAX-VSM* outperforms another two methods. We note that other models capture the relations and outperform feature-based ones, demonstrating the significance of relation modeling.

2. Sequential models vs. graphical models. We compare the sequential models with graphical models in order to present the effectiveness of different methodologies for relation modeling. Most graph networks outperform *MRNN* and *S-HMM*, illustrating that modeling the dynamic relations of states is more significant compared to modeling their sequential dependencies. *HRNN* effectively improves the performance and even beats some graph neural models on the DJIA30 datasets, which suggest that we should try to capture the multi-level correlations in the temporal modeling.

3. Effectiveness of temporal modeling. For the temporal modeling on the evolutionary state graph, different graph neural models adopt different mechanisms. Due to the monotonous information expression of one-hot annotations, *GGSNN* is not as good as the latter methods. Accordingly, *GCN-LSTM* utilizes more state information and performs better. *EvolveGCN*, which builds multi-layer deep networks to combine RNN and GCN, is unstable. *Time2Graph* models the aggregated static graph and ignores the temporal dependencies, which dose not outperform *ST-MGCN* and *our models*. As we expect, our proposed *EvoNet* model conducts the node-graph interaction during the temporal graph propagation, making it more suitable for the temporal modeling of the evolutionary state graph.

4. Ablation study on propagation mechanism. As shown in Table 2 (*Our models*), we attempt to validate the effectiveness of the proposed *EvoBlock*. We can see that, due to simple modeling on state sequence, *EvoNet w/o G* performs poorly. When we build and model the evolutionary state graph for time-series (*EvoNet w/o A*), the performances are improved with the information of node-graph interaction. The temporal attention mechanism can capture the significant correlations during the temporal propagation, meaning that it outperforms other implementations as expected (*EvoNet*). We present several cases in Section 4.6 to support this conclusion.

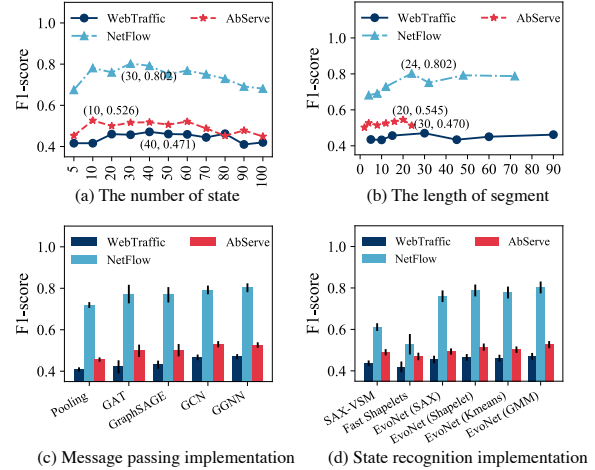


Figure 5: The impact of the different parameters. (a)-(d) present comparisons of state number $|\mathcal{V}|$, segment length τ , implementation of message passing and state recognition, respectively, over three datasets (WebTraffic, NetFlow and AbServe) in Section 4.1.

4.5 Parameter Analysis

We examine the sensitivities of four important parameters to answer Q3: state number $|\mathcal{V}|$, segment length τ , implementation of message passing and state recognition. Due to space limitations, we present the results based on only three datasets in Figure 5. We test $|\mathcal{V}|$ with values from 5 to 100 with interval 10, and test τ with different lengths that are smaller or greater than the period length of the event. We compare the pooling method, GAT [38], GraphSAGE [18], GCN [13] and GGNN [23] for message passing, as well as SAX word [37], Shapelets [26], Kmeans [21] and GMM [6] for state recognition. The F1-score is used as a metric to compare these parameters across the datasets.

1. Sensitivities of state number $|\mathcal{V}|$. As shown in Figure 5a, prediction performance curves differ depending on the dataset, illustrating that the state number $|\mathcal{V}|$ is sensitive to the data owns

patterns. Moreover, the performance is not bound to improve as $|V|$ increases, suggesting that $|V|$ is an empirically determined parameter and is unsuitable for large values.

2. Sensitivities of segment length τ . Another sensitive parameter is the segment length τ , the variation of which may change the temporal scale of event Y_t , and thus the positive ratio of ground truth. We can see the performances in Figure 5b do not vary significantly, meaning that it can be an empirical parameter that is generally determined by the realistic demand (e.g. an acceptable temporal scale of anomaly detection, etc.)

3. Implementation of message passing. Figure 5c presents the comparisons for different implementations of message passing. We can observe that GAT and GraphSAGE perform poorly and are unstable due to their full attention or sampling operation, which is unsuitable for the small-scale graph. The performances of GGNN and GCN are similar, and both outperform the pooling method.

4. Implementation of state recognition. As shown in Figure 5d, we test different implementations of state recognition, and further compare them with some feature-based baselines (i.e., SAX-VSM [37] and Fast Shapelets [35]). We can see that EvoNet can clearly improve the performance of SAX-VSM and Fast Shapelets when models the relations. Moreover, the implementations of cluster methods and shapelet outperform the SAX word; this is because each SAX word is simply a symbolic value representing state, while other representations are a vector describing state patterns, which provide more information for modeling the evolutionary state graph.

4.6 Case Studies

In this section, we apply our EvoNet method to a real-world anomaly prediction scenario in Alibaba Cloud⁵, enabling us to demonstrate how this method can be used to find meaningful relational clues to explain its results. As described in Section 4.1, the minutely time-series of server monitor are segmented by the interval $\tau = 5$ (empirical length). In order to present clearly, we cluster 10 states for constructing evolutionary state graph and conduct EvoNet for anomaly prediction. We visualize the results including several states and the evolutionary state graph at different times. The temporal attention scores learned by EvoNet are also visualized to validate its effectiveness. All results are presented in Figure 6.

1. Effectiveness of temporal attention mechanism. As shown in Figure 6(a)-(b), we adopt heat map to visualize the attention scores α_t learned by Eq 6 at different times. We can see that the attention scores successfully highlight the positions of anomalies in (a) (i.e., the positions near 13:00 and 20:00), which demonstrate that the temporal attention mechanism is useful for EvoNet to capture significant temporal information.

2. Interpretability of evolutionary state graph. We then explore how an evolutionary state graph can be used to find meaningful insights that can explain anomaly event. As shown in Figure 6(a), we mark two intervals, I and II, to visualize evolutionary state graph and explore some meaningful insights. The results are shown in Figure 6(c)-(d). We can see that there is a major transition $\#2 \rightarrow \#4 \rightarrow \#0$

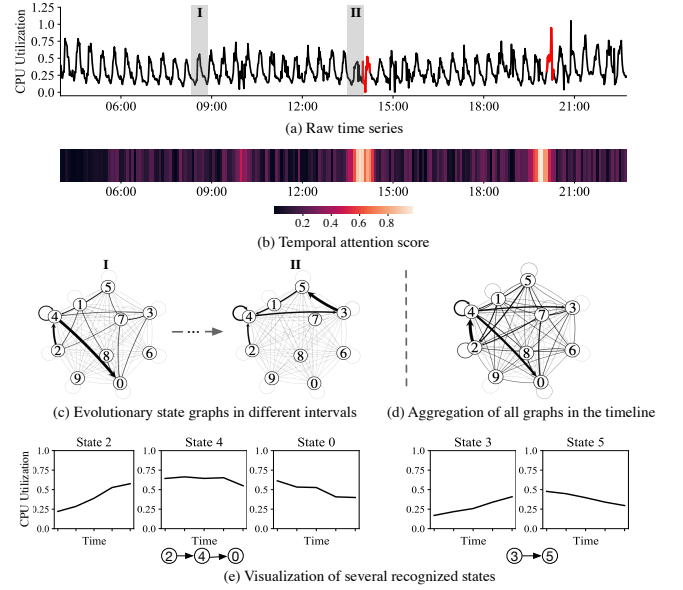


Figure 6: A case of EvoNet conducted for anomaly analysis of cloud service. (a) visualizes the raw time-series of CPU utilization. The red lines indicates anomaly events during different intervals ($\tau = 5$). We set 10 states for constructing evolutionary state graph and conducting EvoNet. Heat map in (b) presents the attention score α_t of EvoNet at different temporal steps. (c)-(d) visualize the evolutionary state graph in two different intervals (I, II marked in (a)) and the aggregated graph in the whole timeline, respectively. (e) presents the raw segments corresponding to the five different states.

(i.e., thick edges) in the graph of I, while $\#3 \rightarrow \#5$ is a major transition in the graph of II. Note that there is an anomaly occurring immediately after interval II. When we aggregate all evolutionary state graphs in the timeline (Figure 6(d)), we can find that the transition $\#2 \rightarrow \#4 \rightarrow \#0$ is the major path in the graph, while $\#3 \rightarrow \#5$ is a rare path (i.e., thin edges). These observations indicate that the transition $\#3 \rightarrow \#5$ occurred in interval II is abnormal, which is consistent with the anomaly of cloud service. As shown in Figure 6(e), we present the average curve of segments with different states. We can see that the transition $\#2 \rightarrow \#4 \rightarrow \#0$ indicates a process of service, i.e., CPU utilization rises from 0.25 to 0.75 and drops after maintaining a period. On the contrary, transition $\#3 \rightarrow \#5$ indicates that CPU utilization rises to 0.5 and then drops immediately. These observations demonstrate that this anomaly may be caused by the CPU's fault.

5 RELATED WORK

time-series modeling. time-series modeling aims to capture the representative patterns underpinning observed data. One important trend here is sequential modeling, such as HMM [34], RNN [5] and their variants [11, 19, 20, 39]. They define one latent representation to capture all the patterns by modeling the sequential dependencies, rather than distinguishing different states. Another trend is mining discretized sequential patterns, such as switch time-series models [1] and dictionaries [24, 25, 37]. They model the time-series by capturing different states of segments independently, but ignore the influence from their relations. Some works have applied *graph structure* into the relation modeling of time-series states [9, 17, 27],

⁵Our method has been deployed by SLS, Alibaba Cloud, the largest log service provider in Asia, acting as a common function.

which aims to represent different segments, rather than capturing the dynamics. To the best of our knowledge, no existing studies have successfully modeled the time-varying relations among states.

Graph neural networks. Models in the graph family [3, 13, 14, 18, 23, 38, 40] have been applied to many real-world scenarios, including learning the dynamics of physical systems [4, 36], predicting traffic on roads [15], etc.. These studies present the effectiveness of GNNs for modeling structural information. Recently, some works have attempted to model dynamic graphs using GNNs [15, 27, 31], although they focus primarily on the explicit graphical structure. To the best of our knowledge, no existing studies have successfully modeled dynamic relations in non-graphical data, such as time-series.

6 CONCLUSIONS

In this paper, we study the problem of how relations among states reflect the evolution of temporal data. We propose a novel representation, the evolutionary state graph, to present the time-varying relations among time-series states. In order to capture these effective patterns for downstream tasks, we further propose a GNN-based model, EvoNet, to conduct dynamic graph modeling. As for the validation of EvoNet’s effectiveness, we conduct extensive experiments on five real-world datasets. Experimental results demonstrate that our model clearly outperforms 11 state-of-the-art benchmark methods. Based on this, we can find some meaningful relations among the states that allow us to understand temporal data.

Acknowledgments. Yang Yang’s work is supported by NSFC (61702447), the National Key Research and Development Project of China (No. 2018AAA0101900), the Fundamental Research Funds for the Central Universities, and research funding from the State Grid Corporation of China. Xiang Ren’s work is supported by the DARPA MCS program under Contract No. N660011924033 with the United States Office Of Naval Research and NSF SMA 18-29268.

REFERENCES

- [1] Pierre Ailliot and Valerie Monbet. 2012. Markov-switching autoregressive models for wind time series. *Environmental Modelling and Software* 30 (2012), 92–101.
- [2] Anthony J Bagnall, Jason Lines, Aaron Bostrom, James Large, and Eamonn J Keogh. 2017. The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *DMKD* 31, 3 (2017), 606–660.
- [3] Peter Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchezgonzalez, Vinicius Flores Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. 2018. Relational inductive biases, deep learning, and graph networks. *arXiv: Learning* (2018).
- [4] Peter Battaglia, Razvan Pascanu, Matthew Lai, Danilo Jimenez Rezende, and Koray Kavukcuoglu. 2016. Interaction networks for learning about objects, relations and physics. *NeurIPS* (2016), 4509–4517.
- [5] Yoshua Bengio, Patrice Y Simard, and Paolo Frasconi. 1994. Learning long-term dependencies with gradient descent is difficult. *TNNLS* 5, 2 (1994), 157–166.
- [6] Philippe Loic Marie Bouttefroy, Abdesslem Bouzerdoum, Son Lam Phung, and Azeddine Beghdadi. 2010. On the analysis of background subtraction techniques using Gaussian Mixture Models. *ICASSP* (2010), 4042–4045.
- [7] Ulrik Brandes. 2001. A Faster Algorithm for Betweenness Centrality. *Mathematical Sociology* 25, 2 (2001), 163–177.
- [8] Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A Scalable Tree Boosting System. *SIGKDD* (2016), 785–794.
- [9] Ziqiang Cheng, Yang Yang, Wei Wang, Wenjie Hu, Yueting Zhuang, and Guojie Song. 2020. Time2Graph: Revisiting Time Series Modeling with Dynamic Shapelets. *AAAI* (2020), 3617–3624.
- [10] Junyoung Chung, Sungjin Ahn, and Yoshua Bengio. 2017. Hierarchical Multiscale Recurrent Neural Networks. *ICLR* (2017).
- [11] Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. 2015. Gated Feedback Recurrent Neural Networks. *ICML* (2015), 2067–2075.
- [12] Nan Du, Hanjun Dai, Rakshit Trivedi, Utkarsh Upadhyay, Manuel Gomezrodriguez, and Le Song. 2016. Recurrent Marked Temporal Point Processes: Embedding Event History to Vector. *SIGKDD* (2016), 1555–1564.
- [13] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. 2015. Convolutional networks on graphs for learning molecular fingerprints. In *NeurIPS*. 2224–2232.
- [14] Rui Feng, Yang Yang, Wenjie Hu, Fei Wu, and Yueting Zhuang. 2018. Representation Learning for Scale-free Networks. *AAAI* (2018), 282–289.
- [15] Xu Geng, Yaguang Li, Leye Wang, Lingyu Zhang, Jieping Ye, Yan Liu, and Qiang Yang. 2019. Spatiotemporal Multi-Graph Convolution Network for Ride-hailing Demand Forecasting. *AAAI* 33 (2019), 3656–3663.
- [16] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. 2017. Neural Message Passing for Quantum Chemistry. *ICML* (2017), 1263–1272.
- [17] David Hallac, Sagar Vare, Stephen P Boyd, and Jure Leskovec. 2017. Toeplitz Inverse Covariance-Based Clustering of Multivariate Time Series Data. *SIGKDD* (2017), 215–223.
- [18] William L Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. *NeurIPS* (2017).
- [19] Sepp Hochreiter and Jurgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation* (1997), 1735–1780.
- [20] Wenjie Hu, Yang Yang, Jianbo Wang, Xuanwen Huang, and Ziqiang Cheng. 2020. Understanding Electricity-Theft Behavior via Multi-Source Data. *WWW* (2020), 2264–2274.
- [21] Tapas Kanungo, David M Mount, Nathan S Netanyahu, Christine D Piatko, Ruth Silverman, and Angela Y Wu. 2002. An efficient k-means clustering algorithm: analysis and implementation. *TPAMI* 24, 7 (2002), 881–892.
- [22] Diederik P Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. *ICLR* (2015).
- [23] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard S Zemel. 2016. Gated Graph Sequence Neural Networks. *ICLR* (2016).
- [24] Jessica Lin, Eamonn J Keogh, Li Wei, and Stefano Lonardi. 2007. Experiencing SAX: a novel symbolic representation of time series. *DMKD* 15, 2 (2007), 107–144.
- [25] Jessica Lin, Rohan Khade, and Yuan Li. 2012. Rotation-invariant similarity in time series using bag-of-patterns representation. *IJIS* (2012), 287–315.
- [26] Jason Lines, Luke M Davis, Jon Hills, and Anthony Bagnall. 2012. A shapelet transform for time series classification. In *SIGKDD*. ACM, 289–297.
- [27] Yozen Liu, Xiaolin Shi, Lucas Pierce, and Xiang Ren. 2019. Characterizing and Forecasting User Engagement with In-app Action Graph: A Case Study of Snapchat. *SIGKDD* (2019), 2023–2031.
- [28] Yue Ning, Sathappan Muthiah, Huzefa Rangwala, and Naren Ramakrishnan. 2016. Modeling Precursors for Event Forecasting via Nested Multi-Instance Learning. *SIGKDD* (2016), 1095–1104.
- [29] Tore Opsahl, Filip Agneessens, and John Skvoretz. 2010. Node centrality in weighted networks: Generalizing degree and shortest paths. *Social Networks* 32, 3 (2010), 245–251.
- [30] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. The PageRank Citation Ranking: Bringing Order to the Web. *WWW* (1999), 161–172.
- [31] Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kanezashi, Tim Kaler, and Charles E Leisersen. 2020. EvolveGCN: Evolving Graph Convolutional Networks for Dynamic Graphs. *AAAI* (2020).
- [32] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2013. On the difficulty of training recurrent neural networks. *ICML* (2013), 1310–1318.
- [33] Bryan Perozzi, Rami Alrfou, and Steven Skiena. 2014. DeepWalk: online learning of social representations. *SIGKDD* (2014), 701–710.
- [34] L R Rabiner and Biinghwang Juang. 1986. An introduction to hidden Markov models. *IEEE Assp Magazine* 3, 1 (1986), 4–16.
- [35] Thanawin Rakthanmanon and Eamonn Keogh. 2013. Fast shapelets: A scalable algorithm for discovering time series shapelets. *ICDM* (2013), 668–676.
- [36] Alvaro Sanchez, Nicolas Heess, Jost Tobias Springenberg, Josh Merel, Raia Hadsell, Martin A Riedmiller, and Peter Battaglia. 2018. Graph Networks as Learnable Physics Engines for Inference and Control. *ICML* (2018), 4467–4476.
- [37] Pavel Senin and Sergey Malinchik. 2013. SAX-VSM: Interpretable Time Series Classification Using SAX and Vector Space Model. *ICDM* (2013), 1175–1180.
- [38] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2018. Graph Attention Networks. *ICLR* (2018).
- [39] Yun Yang and Jianmin Jiang. 2014. HMM-based hybrid meta-clustering ensemble for temporal data. *KBS* (2014), 299–310.
- [40] Lekui Zhou, Yang Yang, Xiang Ren, Fei Wu, and Yueting Zhuang. 2018. Dynamic Network Embedding by Modeling Triadic Closure Process. *AAAI* (2018), 571–578.