

Time2Graph: Revisiting Time Series Modeling with Dynamic Shapelets

Ziqiang Cheng[†], Yang Yang^{†*}, Wei Wang[‡], Wenjie Hu[†], Yueting Zhuang[†], Guojie Song[§]

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

[‡]State Grid Huzhou Power Supply Co. Ltd, China

[§]Key Laboratory of Machine Perception, Ministry of Education, Peking University

Abstract

Time series modeling has attracted extensive research efforts; however, achieving both reliable efficiency and interpretability from a unified model still remains a challenging problem. Among the literature, *shapelets* offer interpretable and explanatory insights in the classification tasks, while most existing works ignore the differing representative power at different time slices, as well as (more importantly) the evolution pattern of shapelets. In this paper, we propose to extract *time-aware shapelets* by designing a two-level timing factor. Moreover, we define and construct the *shapelet evolution graph*, which captures how shapelets evolve over time and can be incorporated into the time series embeddings by graph embedding algorithms. To validate whether the representations obtained in this way can be applied effectively in various scenarios, we conduct experiments based on three public time series datasets, and two real-world datasets from different domains. Experimental results clearly show the improvements achieved by our approach compared with 16 state-of-the-art baselines.

1 Introduction

Time series modeling aims to discover the temporal relationships within chronologically arranged data. The key issue here is how to extract the representative features of a time series. A large part of previous frameworks range from classical feature engineering and representation learning to deep learning based models. While these methods have achieved good performance (Malhotra et al. 2016; Johnson et al. 2016), they have also been subject to criticism regarding their lack of interpretability. On the other hand, *shapelets*, the time series subsequences that are representative of a class (Ye and Keogh 2011), can offer directly interpretable and explanatory insights in the classification scenario, and shapelet-based models have proven to be promising in various practical domains (Ye and Keogh 2009; Xing, Pei, and Yu 2012; Lines et al. 2012; Rakthanmanon and Keogh 2013; Grabocka et al. 2014; Hills et al. 2014; Bostrom and Bagnall 2017).

Existing efforts have mainly considered shapelets as static. However, in the real world, shapelets are often dynamic,

which is reflected in two respects. First, the same shapelet appearing at different time slices may have a range of different impacts. For instance, in the scenario of detecting electricity theft, low electricity consumption in summer or winter is more suspicious than it is in spring, as refrigeration or heating equipments costs more electrical power. Second, determining the ways in which shapelets evolve is vital to a full understanding of a time series. In fact, shapelets with small values at a particular time can hardly distinguish an electricity thief from a normal user who indeed regularly consumes a low level of electricity. An alternative method would involve identifying users who once had high electricity consumption shapelets but suddenly consumes very few electrical power for a while. In other words, an important clue here is how shapelets evolve over time. We refer to the subsequences of a time series that are able to reflect its representativeness at different time slices as *time-aware shapelets*.

There are several challenges involved in modeling the dynamics of shapelets. First, how can time-aware shapelets be defined, and then, be extracted? Traditional algorithms generate a set of static time series subsequences as candidates, then select subsequences with the most discriminatory power according to certain criterion (Lines et al. 2012). However, dynamic shapelets have not been well defined yet, and the criterion used before are often not differentiable if parameterized timing weights are added. To the best of our knowledge, the question of how time-aware shapelets might be extracted remains open. Second, how can the evolutions of shapelets be captured? One possible solution would be to explore the *shapelet sequences*, which consist of ordered shapelets assigned to each *segment* of the time series. However, as each segment may be assigned by several shapelets with different probabilities, enumerating all possible shapelet sequences is quite time-consuming, where the time complexity is $\mathcal{O}(m^K)$ (m is the number of segments and K is the number of shapelets). A more reliable solution would be to construct a transition matrix (graph) whereby each element (edge) represents the probability of a shapelet appearing after another, but how to construct and analyze such matrix (graph) reasonably is nontrivial. Third, how can the evolution patterns of shapelets for modeling the given time series be utilized? Inspired by the recent success of

*Corresponding author: Yang Yang, yangya@zju.edu.cn
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

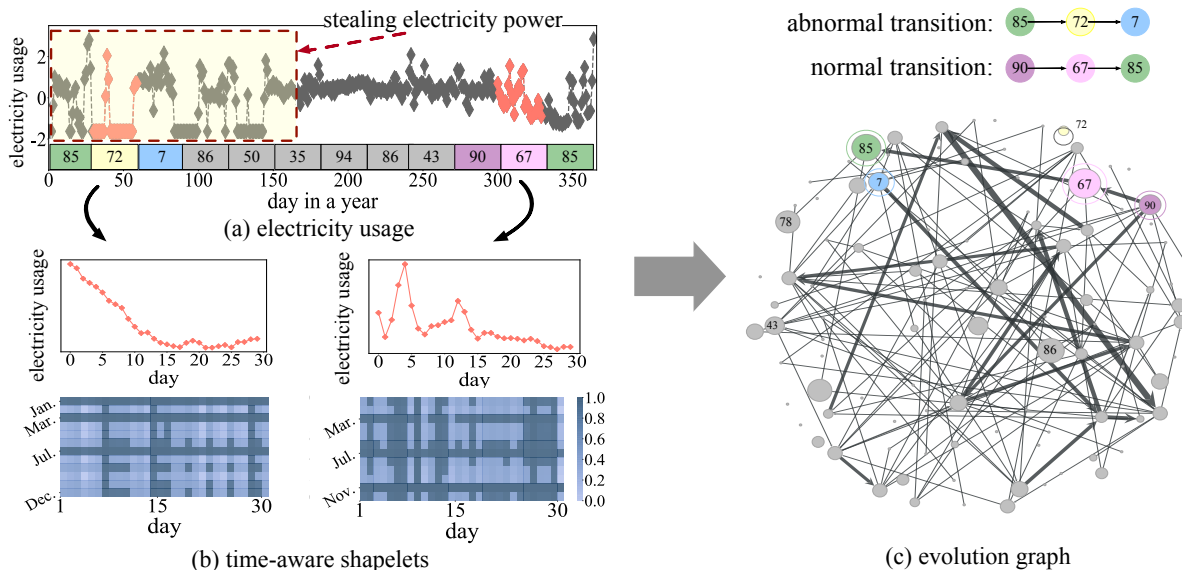


Figure 1: Illustration of (a) electricity consumption data, (b) two shapelets and corresponding timing factors extracted from the observed time series, and (c) the shapelet evolution graph. In shapelet #72, there is an abnormal peak at the beginning, and then a continuous drop until the end, while its timing factor is highlighted mainly in Jan, Mar, and Jun. As for the normal shapelet #67, weights in the timing factor may just reflect the frequency of its occurrence. In figure (c), the size of each vertex is proportional to its weighted in-degree, and the width of the edge is proportional to its betweenness.

representation learning (Perozzi, Al-Rfou, and Skiena 2014; Grover and Leskovec 2016), embedding the evolution patterns of shapelets into a latent feature space Σ can be effective. Then it seems feasible to represent the original time series by using a vector extended from Σ . Designing a reasonable representation learning algorithm is the final challenge.

To address the abovementioned challenges, in this paper, we propose a novel approach to learn the representations of a time series by extracting time-aware shapelets and constructing a shapelet evolution graph. We first define a two-level timing factor to quantify the discriminatory power of shapelets at different time, then construct a graph to represent the evolution patterns of shapelets. Fig. 1 shows an example from real-world electricity consumption record data: Fig. 1a demonstrates the one-year electricity usage of a user who has stolen electrical power from January to May while using electrical power normally in the remaining months. We assign each month the most representative shapelet at that time and present the shapelets #72 and #67, along with their timing factors in Figure 1b, where dark areas indicate that the corresponding shapelet is more discriminative relative to light areas. The shapelet evolution graph is presented in Fig. 1c, illustrating how a shapelet would transfer from one to another *in a normal case*: for the normal electricity consumption record, there is a clear path for its shapelet transition (#90 → #67 → #85) in the graph. For the abnormal data, however, the path (#85 → #72 → #7) does not exist, indicating that the connectivity of the shapelet transition path provides an evidential basis for detecting an abnormal time series. Finally, we translate the problem of learning representations of shapelets and time series into a graph embedding problem.

We summarize our contributions to the field as follows: 1) We propose the concept of time-aware shapelets and design a

learning algorithm to extract them; 2) We construct a shapelet evolution graph and translate the problem of representation learning for shapelets and time series into graph embedding; and 3) We validate the effectiveness of our approach based on three public and two real-world datasets. Experimental results show that our approach achieves notably better performance when compared with 16 state-of-the-art baselines.

2 Preliminaries

A time series set $T = \{t_1, \dots, t_{|T|}\}$, where each t contains n chronologically arranged elements, i.e., $t = \{x_1, \dots, x_n\}$. A *segment* s of t is a contiguous subsequence, i.e., $s = \{x_i, \dots, x_j\}$. If t can be divided by m segments of equal length l , then we have $t = \{\{x_{l*k+1}, \dots, x_{l*k+l}\}, 0 \leq k \leq m-1\}$. To measure the dissimilarity of sequences, we denote the distance between two segments s_i and s_j as $d(s_i, s_j)$, where $d(\cdot, \cdot)$ can be intuitively formalized as the *Euclidean Distance* (ED). But ED cannot deal with varied sequence length and timing shifts, and in the context of time series modeling, time warping techniques are often used to address such problems. The central idea of time warping is to find an appropriate alignment for the given pair of sequences, where an alignment is defined as

Definition 1 Alignment. Given two segments s_i and s_j with length l_i and l_j respectively, an alignment $\mathbf{a} = (\mathbf{a}_1, \mathbf{a}_2)$ is a pair of two index sequences of length p , satisfying that

$$\begin{aligned}
 1 &\leq \mathbf{a}_k(1) \leq \dots \leq \mathbf{a}_k(p) = l_k, \\
 \mathbf{a}_k(n+1) - \mathbf{a}_k(n) &\leq 1, \\
 &\text{for } k = i, j, \text{ and } 1 \leq n \leq p-1
 \end{aligned} \tag{1}$$

We denote all possible alignments for two segments s_i and s_j as $\mathcal{A}(s_i, s_j)$, then one popular time-warping based measurement, *DTW* (Dynamic Time Warping), can be illustrated

as Eq. (2), where $\tau(s_i, s_j | \mathbf{a})$ is a predefined dissimilarity for two sequences under the alignment \mathbf{a} (Müller 2007). We refer the alignment achieving the minimum in Eq. (2) as \mathbf{a}^* .

$$d_{DTW}(s_i, s_j) = \min_{\mathbf{a} \in \mathcal{A}(s_i, s_j)} \tau(s_i, s_j | \mathbf{a}) \quad (2)$$

We can further measure the dissimilarity between a segment s and a time series $\mathbf{t} = \{s_1 \cdots s_m\}$. Inspired by the literature that we often say a segment s is close to \mathbf{t} if there exists some segment s' in \mathbf{t} between which the distance of s is rather small, we define the distance between s and \mathbf{t} as

$$D(s, \mathbf{t}) = \min_{1 \leq k \leq m} d(s, s_k) \quad (3)$$

Based on these definitions, previous work have proposed novel methods to extract typical subsequences, i.e., shapelets, to distinguish the *representative power* of segments:

Definition 2 Shapelet. A shapelet v is a segment that is representative of a certain class. More precisely, it can separate T into two smaller sets, one that is close to v and another far from v by some specific criteria, such that for a time series classification task, positive and negative samples can be put into different groups. The criteria can be formalized as

$$\mathcal{L} = -g(S_{pos}(v, T), S_{neg}(v, T)) \quad (4)$$

\mathcal{L} measures the dissimilarity between positive and negative samples towards the shapelet v . $S_*(v, T)$ denotes the set of distances with respect to a specific group T_* , i.e., positive or negative class; the function g takes two finite sets as input, returns a scalar value to indicate how far these two sets are, and it can be *information gain* (Ye and Keogh 2011), or some dissimilarity measurements on sets, i.e., *KL divergence*.

3 Time2Graph Framework

In this section, we present a novel representation learning algorithm for time series modeling. We name the proposed framework as *Time2Graph*, as it transforms time series to a graph with shapelets and their transitions. We extract time-aware shapelets from a large pool of candidates (Sec. 3.1), then construct the *Shapelet Evolution Graph* to capture the correlations between shapelets (Sec. 3.2), and finally learn the time series representation vectors (Sec. 3.3) by concatenating segment embeddings which is composed by shapelet embeddings obtained from the *Shapelet Evolution Graph*.

3.1 Time-Aware Shapelet Extraction

Existing works often ignore that subsequences may have various representative powers at different time. For example, low consumption of electrical power in spring is normal, whereas it is a strong signal for identifying abnormal users in summer, since high temperatures often lead to more electricity usage. Therefore, we consider time-aware shapelets in this paper.

We define two factors for quantitatively measuring the timing effects of shapelets at different levels. Specifically, we introduce the *local factor* \mathbf{w}_n to denote the inner importance of the n -th element of a particular shapelet, then the distance between a shapelet v and a segment s is redefined as

$$\begin{aligned} \hat{d}(v, s | \mathbf{w}) &= \tau(v, s | \mathbf{a}^*, \mathbf{w}) \\ &= \left(\sum_{k=1}^p \mathbf{w}_{\mathbf{a}_1^*(k)} \cdot (\mathbf{v}_{\mathbf{a}_1^*(k)} - \mathbf{s}_{\mathbf{a}_2^*(k)})^2 \right)^{\frac{1}{2}} \end{aligned} \quad (5)$$

where \mathbf{a}^* refers to the best alignment for *DTW* which we have discussed in Eq. (2). The intuitive explanation of Eq. (5) is to project the weight \mathbf{w} onto the *DTW* alignment path. On the other hand, at a *global level*, we aim to measure the timing effects across segments on the discriminatory power of shapelets. It is inspired from the intuition that shapelets may represent totally different meaning at different time steps, and it is straightforward to measure such deviations by adding segment-level weights. Formally, we set a *global factor* \mathbf{u}_m to capture the cross-segments influence, then the distance between a shapelet v and a time series \mathbf{t} can be rewritten as

$$\hat{D}(v, \mathbf{t} | \mathbf{w}, \mathbf{u}) = \min_{1 \leq k \leq m} \mathbf{u}_k \cdot \hat{d}(v, s_k | \mathbf{w}) \quad (6)$$

where \mathbf{t} is divided into m segments, i.e., $\mathbf{t} = \{s_1, \cdots, s_m\}$. Eq. (6) denotes the two-level time-aware distance between a shapelet v and a time series \mathbf{t} , and the parameters \mathbf{w}, \mathbf{u} associated with each specific shapelet can be learned separately under some proper criteria. Given a classification task, we establish a supervised learning method to select the most important time-aware shapelets and learn corresponding timing factors \mathbf{w}_i and \mathbf{u}_i for each shapelet v_i . In particular, we have a pool of segments as shapelet candidates that selected from all subsequences, and a set of time series T with labels. We only consider binary classification here, to which it is straightforward to extend multinomial classification. For each shapelet candidate v , we modify Eq. (4) as

$$\hat{\mathcal{L}} = -g(S_{pos}(v, T), S_{neg}(v, T)) + \lambda \|\mathbf{w}\| + \epsilon \|\mathbf{u}\| \quad (7)$$

where λ and ϵ are the hyperparameters of penalties, and the differentiable function $g(\cdot, \cdot)$ measures the distances between distributions of two finite sets. In practice, we assume that the given sets both follow some particular distributions, e.g., Gaussian Distribution. We can easily estimate the distribution parameters by closed-form solutions, and then the gradients of g can be derived from those differentiable parameters. After learning the timing factors from shapelet candidates, we select the top K shapelets with minimal loss in Eq. (7).

3.2 Shapelet Evolution Graph

After obtaining shapelets, many works use BoP (Baydogan et al. 2013) or similar methods to represent the time series, but these algorithms ignore the correlations between shapelets. Here, *correlations* consist of the co-occurrence along with the occurrence order. To capture such relationship, we propose the concept of *Shapelet Evolution Graph* as follows:

Definition 3 Shapelet Evolution Graph. It is a directed and weighted graph $G = (V, E)$ in which V consists of K vertices, each denoting a shapelet, and each directed edge $e_{ij} \in E$ is associated with a weight w_{ij} , indicating the occurrence probability of shapelet $v_i \in V$ followed by another shapelet $v_j \in V$ in the same time series.

Graph Construction (Algorithm 1). We first assign each segment s_i of each time series to several shapelets that have the closest distances to s_i according to the time-aware dissimilarity. Then a problem naturally rises as how far the distance would be considered as *closest*? One simple but effective solution is to predefine a threshold δ such that distances less

than δ are treated as close, and in practice, we can determine δ by experimental statistics on the training dataset. For convenience, we denote those shapelets assigned to segment s_i as $v_{i,*}$, and say that $v_{i,j}$ is the j -th assignment of s_i . For the purpose of measuring how reasonable our assignment is, we standardize assignment probability $p_{i,j}$ as

$$p_{i,j} = \frac{\max(\hat{d}_{i,*}(v_{i,*}, s_i)) - \hat{d}_{i,j}(v_{i,j}, s_i)}{\max(\hat{d}_{i,*}(v_{i,*}, s_i)) - \min(\hat{d}_{i,*}(v_{i,*}, s_i))} \quad (8)$$

where $\hat{d}_{i,*}(v_{i,*}, s_i) = \mathbf{u}_*[i] * \hat{d}(v_{i,*}, s_i | \mathbf{w}_*)$ (Eq. (5)), with the constraint that $\hat{d}_{i,*} \leq \delta$. So the shapelets set $v_{i,*}$ is assigned to segment s_i with probability $p_{i,*}$, and $v_{i+1,*}$ is assigned to s_{i+1} with probability $p_{i+1,*}$ respectively, where s_i is followed by s_{i+1} in one time series. Then, for each pair (j, k) , we create a weighted edge from shapelet $v_{i,j}$ to $v_{i+1,k}$ with weight $p_{i,j} \cdot p_{i+1,k}$ and merge all duplicated edges as one by summing up their weights. Finally, we normalize the edge weights sourced from each node as 1, which naturally interprets the edge weight between each pair of nodes, i.e., v_i and v_j into the conditional probability $P(v_j | v_i)$ that shapelet v_i being transformed into v_j in an adjacent time step.

Algorithm 1: Shapelet Evolution Graph Construction

Input:

time series set $T = \{t_1, \dots, t_{|T|}\}$,
 K shapelets $\{v_1 \dots v_K\}$, distance threshold δ

Output:

Shapelet evolution graph G

- 1: Initialize the graph G with K vertices
- 2: **for all** segment s_i of t in T **do**
- 3: **for all** shapelet v_j where $\hat{d}(v_j, s | \mathbf{w}_j) * \mathbf{u}_j[i] \leq \delta$ **do**
- 4: Assign v_j to s_i with probability defined in Eq. (8)
- 5: **end for**
- 6: **end for**
- 7: **for all** adjacent segment pair (s_i, s_{i+1}) of t in T **do**
- 8: **for all** assigned shapelet pair $(v_{i,j}, v_{i+1,k})$ **do**
- 9: Add directed edge $e_{j,k}$ with weight $p_{i,j} * p_{i+1,k}$
- 10: **end for**
- 11: **end for**
- 12: Normalize the edge weights for each vertex
- 13: **return** G

3.3 Representation Learning

Finally, we learn the representations for both the shapelets and the given time series by using the shapelet evolution graph constructed as above. We first employ an existing graph embedding algorithm (DeepWalk (Perozzi, Al-Rfou, and Skiena 2014)) to obtain vertex (shapelet) representation vectors $\mu \in \mathbb{R}^B$, where B is the embedding size (latent dimension). In our case, a path in G intuitively reflects possible transitions between shapelets.

Next, for a time series $t = \{s_1 \dots s_m\}$ with corresponding assigned shapelets $\{v_{1,*} \dots v_{m,*}\}$ and assignment probabilities $\{p_{1,*}, \dots, p_{m,*}\}$, we retrieve each shapelet $v_{i,j}$'s representation vector $\mu(v_{i,j})$ multiplied by assignment probability $p_{i,j}$, and sum them over each segment. If there exists

some segment that we cannot assign any shapelets to it by applying the predefined distance threshold, the embeddings of this segment would be left as empty. It is reasonable since shapelet embeddings are always non-zero which is guaranteed by graph embedding models (more precisely, shapelet embeddings are bound to be normalized), so segments without shapelet assignments are clearly distinguished by empty values. By far, we get the segment representations, and finally concatenate all those m segment embedding vectors to obtain the representation vector Φ for time series t as follows:

$$\Phi_i = \left(\sum_j p_{i,j} \cdot \mu(v_{i,j}) \right), 1 \leq i \leq m \quad (9)$$

The representation vector of the time series can then be applied as features for various time series classification tasks, by the way of feeding the embedding features into an outer classifier. See the formulation of the representation learning framework in Algorithm 2, and the details of the downstream classification tasks are introduced in Sec. 4.1.

Algorithm 2: Time Series Embedding Framework

Input:

time series $t = \{s_1 \dots s_m\}$, K shapelets $\{v_1 \dots v_K\}$
shapelet evolution graph G , graph embedding size B

Output:

time series embedding vector Φ

- 1: Embeds graph G as μ^B with representation-size as B
- 2: Initialize Φ as an empty vector.
- 3: **for all** segment s_i in t , $1 \leq i \leq m$ **do**
- 4: $temp \leftarrow$ zero vectors with size B
- 5: **for all** shapelet $v_{i,j}$ that assigned to s_i **do**
- 6: $temp += p_{i,j} * \mu^B(v_{i,j})$, where shapelet assignment can be referred from Algorithm 1
- 7: **end for**
- 8: $\Phi.append(temp)$
- 9: **end for**
- 10: **return** Φ

4 Experiments

4.1 Experimental Setup

We use three public datasets, *Earthquakes* (EQS), *WormsT-woClass* (WTC) and *Strawberry* (STB) from the *UCR Time Series Archive* (Dau et al. 2018), along with two real-world datasets, *Electricity Consumption Records* (ECR) from State Grid of China and *Network Traffic Flow* (NTF) from *China Telecom.*, to validate our proposed model. Table 1 shows the overall statistics of those five datasets:

Dataset \ Metric	EQS	WTC	STB	ECR	NTF
#(time series)	461	258	983	60,872	5,950
positive ratio(%)	25.3	42.2	64.3	2.3	6.4

Table 1: Overall statistics of 5 datasets in the experiments.

We briefly introduce the two real-world datasets as follows: **Electricity Consumption Records (ECR)**. This dataset is

Methods \ Datasets	public dataset			real-world dataset					
	EQS	WTC	STB	Prec	ECR	F_1	Prec	NTF	F_1
	Accuracy				Recall			Recall	
NN-ED	68.22	62.41	95.60	18.71	10.48	13.44	37.71	46.35	41.59
NN-DTW	70.31	68.16	95.53	15.52	18.15	16.73	33.20	43.75	37.75
NN-WDTW	69.50	67.74	95.44	15.52	18.15	16.73	35.29	46.86	40.27
NN-CID	69.41	69.56	95.51	18.18	13.71	15.63	32.56	43.75	37.33
DDDTW	70.79	70.92	95.60	18.78	13.71	15.85	30.48	42.71	35.58
<i>XGBoost</i> (origin)	74.82	62.34	95.92	38.36	19.48	25.86	71.43	17.86	28.57
<i>XGBoost</i> (feature)	75.54	64.94	97.03*	56.82	16.23	25.25	80.00*	21.43	33.80
BoP	74.80	74.42*	96.45	14.86	4.44	6.83	43.40	47.92	45.55
TSF	74.67	68.51	96.27	26.32	2.02	3.75	57.52	33.85	42.62
EE	73.50	71.74	95.88	10.18	33.47	15.62	42.98	27.08	33.23
SAXVSM	73.76	72.10	96.97	21.59	42.74	28.69	30.19	50.00	37.65
LS	74.22	73.57	92.49	0.00	0.00	0.00	0.00	0.00	0.00
FS	74.66	70.58	91.66	10.45	79.84*	18.48	63.55	35.42	45.49
LPS	66.78	74.26	96.35	17.00	24.19	19.97	24.17	30.21	26.85
LSTM	74.82	42.86	63.84	13.64	31.86	19.11	7.22	16.67	10.08
VAE	71.22	62.34	71.35	19.02	14.11	16.20	59.79	30.21	40.14
Shapelet-Seq	75.53	55.84	78.10	14.37	66.94	23.66	18.45	61.98*	28.44
<i>Time2Graph</i> -static	76.98	70.13	95.68	33.81	29.22	31.36	80.00*	28.57	42.11
<i>Time2Graph</i>	79.14*	72.73	96.76	30.10*	40.26	34.44*	71.52*	56.25	62.97*

Table 2: Comparison of classification performance on the public and real-world datasets (%). Results of the *Time2Graph* model are bold, and a star (*) means the best performance among all methods.

provided by the State Grid Corporation of China and contains the daily electricity consumption records (K·Wh) of 60,872 users over the span of one year (2017). For every user, it records the daily total electricity usage, on-peak usage, and off-peak usage. Some users may take unauthorized actions concerning the electricity meter or power supply lines to cut costs (i.e., electricity theft), and there are a total of 1,433 (2.3%) users who have been manually confirmed as having stolen electrical power. Given users and their electricity consumption record, the task is to determine which users have stolen electrical power in the past year.

Network Traffic Flow (NTF). This dataset is provided by China Telecom, the major mobile telecommunications service provider in China. It consists of 5,950 network traffic series, each of which describes the hourly inflow and outflow of different servers, from April 6th 2017 to May 15th 2017. When an abnormal flow goes through server ports and some process is suddenly dead, an alarm state is recorded by the operating system (objective ground-truth); there are 383 (6.4%) servers with abnormal flow series. The goal is to use the daily network traffic data to detect whether there are abnormal flows during a period.

Besides, we select several *UCR* datasets from many candidates by the following reasons that: 1) to maintain the consistency of evaluation metrics between the real-world and public datasets, we only consider binary-label ones in *UCR*; 2) we have to make sure that there are enough training cases because we need sufficient samples to capture the normal transitions between shapelets (many binary-label datasets in *UCR* only have less than 100 training samples), and 3) we omit all datasets categorized as “image”, because the pro-

posed intuition (timing factor, shapelet evolutions) may not be appropriate for time series transformed from images. After filtering based on the abovementioned criterion, and due to space limitation, we only present those three in the paper.

We compare our proposed *Time2Graph* model with several groups of the state-of-the-art baselines:

Distance-based Models. Previous work has stated that in most time series classification tasks, 1-NN-based methods are hard to beat (Wang et al. 2013; Bagnall et al. 2017). As for the distance metric applied in 1-NN, we use Euclidean Distance (ED), Dynamic Time Warping (DTW), Weighed DTW (WDTW) (Jeong, Jeong, and Omiaomu 2011), Complexity-Invariant Distance (CID) (Batista et al. 2014) and Derivative DTW (DDDTW) (Górecki and Łuczak 2013) as candidates.

Feature-based Models. We first extract some statistical features (the mean, standard deviation, etc.), or just take the raw time series as input (*XGBoost* (feature/origin)), and use the same outer classifier as which *Time2Graph* uses (*xgboost*) to validate the effectiveness of the representation learning framework. Besides, several popular feature-based algorithms have been proposed for time series classification tasks. In this paper, we choose some typical algorithms to compare with our model: Bag-of-Patterns (BoP) (Lin, Khade, and Li 2012), Time Series Forest (TSF) (Deng et al. 2013), Elastic Ensembles (EE) (Lines and Bagnall 2015), and Vector Space Model using SAX (SAXVSM) (Senin and Malinchik 2013).

Shapelet-based Models. Another typical group of algorithms extracts *shapelets* to capture the intrinsic features of the original time series data. In this paper, we use several famous shapelet-based frameworks as baselines: Learn Time Series Shapelets (LS) (Grabocka et al. 2014), Fast Shapelets

(FS) (Rakthanmanon and Keogh 2013) and Learned Pattern Similarity (LPS) (Baydogan and Runger 2016).

Deep learning models. We consider two commonly-used deep models, LSTM and VAE, due to their efficacy in feature-representation tasks and processing time series data.

Time2Graph variants. We also compare *Time2Graph* model with its derivatives by modifying some key components to see how they fare: a) We sample the most possible shapelet sequence (i.e., each segment is assigned with highest assignment probability) for each time series, and use *LSTM* to conduct end-to-end classifications, denoted as *Shapelet-Seq*; b) We learn shapelets without considering timing factors, and embed them in the same way of *Time2Graph*, and refer this method as *Time2Graph-static*.

We choose *XGBoost* (Chen and Guestrin 2016) as the outer classifier, and use 5-fold nested cross-validation to conduct fine-tuning on hyper-parameters. As for the reproducibility, the source codes, along with the implementation details, parameter settings and documentations, can be found on the project homepage: <https://petecheng.github.io/Time2Graph>.

4.2 Comparison Results

Table 2 shows the comparison results for classification tasks. All three public datasets from *UCR Archive* uniformly use accuracy as evaluation metric, and for those two real-world datasets, which are both very imbalanced, we show the prediction precision, recall and F1 score.

We conclude from Table 2 that the *Time2Graph* model achieves competitive performance on the three public datasets. Specifically, on the *EQS* dataset, *Time2Graph* achieves the highest accuracy (79.14%), while on the other two datasets, *Time2Graph* also beats most of the baseline methods. When it comes to the real-world datasets, it is clear that *Time2Graph* significantly outperforms all baselines in the F1 score metric. Even though some baseline methods achieve higher precision or recall, all of them seem to encounter biases on positive or negative samples.

We next compare *Time2Graph* with its variance. As mentioned in Sec. 1, *Shapelet-Seq* model suffers from the size of possible sequences, and when we only sample sequences with the highest probability, its performance fails to match several baselines, since a substantial amount of information is lost during sequence sub-sampling. The performance incrementation from *Time2Graph-Static* to *Time2Graph* demonstrates the predictive power brought by time-aware shapelets, and the additional interpretability and insights derived from timing factors are shown in Sec. 4.4. In summary, the *Time2Graph* model is better at finding effective patterns, as well as capturing evolutionary characteristics in time series.

4.3 Parameter Analysis




Figure 2: Parameter analysis.

We examine the sensitivities of three important hyperparameters: number of selected shapelets K , graph embedding size B and segment length l . Due to space limitations, we only present the results for the public datasets, which are shown in Fig. 2. From the results, we see that K should be large enough to capture a sufficient number of shapelets; while when K is too large, it will bring in less representative shapelets as noise (Fig. 2a). Another parameter that should be tuned is the segment length l . We can see from Fig. 2c that it achieves the best results when l is 24, which is exactly the same as the number of hours in a day. It seems not to be a coincidence that, in *ECR* dataset, the best segment length is 30, i.e., the number of days in a month, while the optimal choice for *NTF* is 24, again the number of hours in a day. We may conclude that the segment length l should be carefully selected based on the physical characteristics of the original data source. As for the embedding size, we see that accuracy improves in most cases when it is increasing (Fig. 2b), whereas the efficiency will be sacrificed. It is also difficult to train the outer classifiers for features with dimensions that are too large; accordingly, an appropriate graph embedding size is necessary for better performance.

4.4 Case Study of Time-Aware Shapelets

In the following two sections, we conduct several case studies to explore the interpretability of our proposed model, and we use *ECR* dataset as the example since much domain knowledge is available here from experts.

The first question is, do the shapelets we extracted indeed have varying levels of discriminatory power? As shown in Fig. 4a, the training loss grows much slower at the right end, and the KL divergence of distributions of distances between positive and negative samples towards the top (ranked 1-50) shapelets on the test set is statistically significantly larger than that for the bottom (ranked 51-100) shapelets ($p = 7.7 * 10^{-6}$). This reflects the effectiveness of the selected shapelets to some extent. To rigorously check the differences in shapelet variance, we show each shapelet’s mean value and standard deviation (std) in Fig. 4b. Again, the std of top shapelets are statistically significantly larger than those of bottom ones ($p = 7.5 * 10^{-3}$), while the mean values across shapelets exhibit very little difference; this suggests that typical patterns intend to be unstable. And to make further illustration, we compare the top-1 shapelet extracted by *LS* (a popular baseline) and *Time2Graph* in Fig. 4c, d. The scale and trends of these two shapelets differ a lot, and Fig. 4d provides additional information towards time-aware shapelets in *T2G*: this specific shapelet matters in spring and summer (from month-level timing factor), and weights more at the peak of time series data during the month (from day-level timing factor). Such clue is the distinct advantage of our proposed model on the interpretability.

4.5 Case Study of the Shapelet Evolution Graph

We finally conduct experiments to construct shapelet evolution graphs for different time steps in order to see how the graphs change and how the nodes evolve. Fig. 3 shows two graphs, one for January and another for July. In January, shapelet #45 has large in/out degrees, and its corresponding




Figure 3: Shapelet evolution graphs at different times. The positions of vertices in the two graphs are the same. The vertex size is proportional to its weighted in-degree, the same as the edge width to its betweenness.




Figure 4: Shapelet analysis. (a) shows the train and test loss; (b) shows the mean and std of shapelets; (c) compares the top (rank-1) shapelet between *Learn Shapelets* (*LS*) and *time2graph* (*T2G*), and (d) visualizes the timing factors of the *T2G*-shapelet in (c).

timing factor is highlighted in January and February (dark areas). It indicates that shapelet #45 is likely to be a common pattern at the beginning of a year. As for July, shapelet #45 is no longer as important as it was in January. Meanwhile, shapelet #42, which is almost an isolated point in January, becomes very important in July. Although we do not explicitly take seasonal information into consideration when constructing shapelet evolution graphs, the inclusion of the timing factors means that they are already incorporated into the process of the graph generation.

5 Related Work

Time series modeling have attracted extensive research over a wide range of fields, such as image alignment (Peng et al. 2014), speech recognition (Shimodaira et al. 2002), etc. One important technique here is Dynamic Time Warping (DTW) (Müller 2007), which aims to find the proper distance measurement between time series data, and a wide range of applications (Jeong, Jeong, and Omiaomu 2011; Górecki

and Łuczak 2013) have been proposed based on this measure.

Traditional time series classification models try to extract efficient features from original data and develop a well-trained classifier, such as BoP (Lin, Khade, and Li 2012), TSF (Deng et al. 2013), EE (Lines and Bagnall 2015), etc. However, the major challenge is that there are no explicit features in sequences (Xing, Pei, and Keogh 2010), so much research has focused on time series embedding (Bagnall et al. 2017): Algorithms based on DTW and traditional embedding techniques (Hayashi, Mizuhara, and Suematsu) aim to project original time series data into feature-vector space; Symbolic representations (Lin et al. 2007; 2003; Schäfer 2015) transform time series using symbols such as characters in a given alphabet; Shapelet discovery-based models (Ye and Keogh 2011; Lines et al. 2012; Rakthanmanon and Keogh 2013; Hou, Kwok, and Zurada 2016; Baydogan and Runger 2016), from another perspective, try to find typical subsequences based on certain criteria such as information gain. Another relevant work to this paper is graph embedding, and one popular category (Goyal and Ferrara 2018) lies in the random walk-based methods, such as DeepWalk (Perozzi, Al-Rfou, and Skiena 2014) and node2vec (Grover and Leskovec 2016). Besides, dynamic graph embedding which aims to model the network dynamics and evolutions over time (Zhu et al. 2016; Zhou et al. 2018) also attracts great interests in recent years.

6 Conclusion and Discussion

In this paper, we proposed a novel framework *Time2Graph* to learn time-aware shapelets for time series representations. Moreover, to capture the co-occurrence and peer influence between shapelets, we put forward the idea of considering a time series as a graph, in which the nodes refer to shapelets, and weighted edges denote transitions between shapelets with varying probabilities. By conducting experiments on three public datasets from *UCR Archive* and two real-world datasets, we demonstrate the effectiveness and interpretability of our proposed model.

Acknowledgments. The work is supported by NSFC (61702447), a research funding from Chinese Knowledge Center of Engineering Science and Technology (CKCEST), and the Fundamental Research Funds for the Central Universities.

References

- Bagnall, A.; Lines, J.; Bostrom, A.; Large, J.; and Keogh, E. 2017. The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *DMKD*. 31(3):606–660.
- Batista, G. E.; Keogh, E. J.; Tataw, O. M.; and De Souza, V. M. 2014. Cid: an efficient complexity-invariant distance for time series. *DMKD*. 28(3):634–669.
- Baydogan, M. G., and Runger, G. 2016. Time series representation and similarity based on local autopatterns. *DMKD*. 30(2):476–509.
- Baydogan, M.; Gokce, R.; George, T.; and Eugene. 2013. A bag-of-features framework to classify time series. *T-PAMI*. 35(11):2796–2802.
- Bostrom, A., and Bagnall, A. 2017. Binary shapelet transform for multiclass time series classification. In *TLSD-KCS'17*. 24–46.
- Chen, T., and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *KDD'16*.
- Dau, H. A.; Keogh, E.; Kamgar, K.; Yeh, C.-C. M.; Zhu, Y.; Gharghabi, S.; Ratanamahatana, C. A.; Yanping; Hu, B.; Begum, N.; Bagnall, A.; Mueen, A.; and Batista, G. 2018. The ucr time series classification archive. https://www.cs.ucr.edu/~eamonn/time_series_data_2018/.
- Deng, H.; Runger, G.; Tuv, E.; and Vladimir, M. 2013. A time series forest for classification and feature extraction. *Information Sciences* 239:142–153.
- Górecki, T., and Łuczak, M. 2013. Using derivatives in time series classification. *DMKD*. 26(2):310–331.
- Goyal, P., and Ferrara, E. 2018. Graph embedding techniques, applications, and performance: A survey. *KBS*. 151:78–94.
- Grabocka, J.; Schilling, N.; Wistuba, M.; and Schmidt-Thieme, L. 2014. Learning time-series shapelets. In *KDD'14*, 392–401. ACM.
- Grover, A., and Leskovec, J. 2016. node2vec: Scalable feature learning for networks. In *KDD'16*, 855–864.
- Hayashi, A.; Mizuhara, Y.; and Suematsu, N. Embedding time series data for classification. In *MLDM'05 workshop*.
- Hills, J.; Lines, J.; Baranauskas, E.; Mapp, J.; and Bagnall, A. 2014. Classification of time series by shapelet transformation. *DMKD*. 28(4):851–881.
- Hou, L.; Kwok, J.; and Zurada, J. 2016. Efficient learning of timeseries shapelets. In *AAAI'16*.
- Jeong, Y.-S.; Jeong, M. K.; and Omiaomu, O. A. 2011. Weighted dynamic time warping for time series classification. *Pattern Recognition*. 44(9):2231–2240.
- Johnson, M.; Duvenaud, D. K.; Wiltchko, A.; Adams, R. P.; and Datta, S. R. 2016. Composing graphical models with neural networks for structured representations and fast inference. In *NIPS'16*, 2946–2954.
- Lin, J.; Keogh, E.; Lonardi, S.; and Chiu, B. 2003. A symbolic representation of time series, with implications for streaming algorithms. In *DMKD'03 Workshop*.
- Lin, J.; Keogh, E.; Wei, L.; and Lonardi, S. 2007. Experiencing sax: a novel symbolic representation of time series. *DMKD*. 15(2):107–144.
- Lin, J.; Khade, R.; and Li, Y. 2012. Rotation-invariant similarity in time series using bag-of-patterns representation. *JIS*. 39(2):287–315.
- Lines, J., and Bagnall, A. 2015. Time series classification with ensembles of elastic distance measures. *DMKD*. 29(3):565–592.
- Lines, J.; Davis, L. M.; Hills, J.; and Bagnall, A. 2012. A shapelet transform for time series classification. In *KDD'12*, 289–297.
- Malhotra, P.; Ramakrishnan, A.; Anand, G.; Vig, L.; Agarwal, P.; and Shroff, G. 2016. Lstm-based encoder-decoder for multi-sensor anomaly detection. *arXiv preprint arXiv:1607.00148*.
- Müller, M. 2007. Dynamic time warping. *IRMM'07* 69–84.
- Peng, X.; Huang, J.; Hu, Q.; Zhang, S.; and Metaxas, D. N. 2014. Head pose estimation by instance parameterization. In *ICPR'14*, 1800–1805.
- Perozzi, B.; Al-Rfou, R.; and Skiena, S. 2014. Deepwalk: Online learning of social representations. In *KDD*, 701–710.
- Rakthanmanon, T., and Keogh, E. 2013. Fast shapelets: A scalable algorithm for discovering time series shapelets. In *SDM*, 668–676.
- Schäfer, P. 2015. The boss is concerned with time series classification in the presence of noise. *DMKD*. 29(6):1505–1530.
- Senin, P., and Malinchik, S. 2013. Sax-vsm: Interpretable time series classification using sax and vector space model. In *ICDM'13*, 1175–1180.
- Shimodaira, H.; Noma, K.-i.; Nakai, M.; and Sagayama, S. 2002. Dynamic time-alignment kernel in support vector machine. In *NIPS'02*, 921–928.
- Wang, X.; Mueen, A.; Ding, H.; Trajcevski, G.; Scheuermann, P.; and Keogh, E. 2013. Experimental comparison of representation methods and distance measures for time series data. *DMKD*. 26(2):275–309.
- Xing, Z.; Pei, J.; and Keogh, E. 2010. A brief survey on sequence classification. *ACM SIGKDD Explorations Newsletter*. 12(1):40–48.
- Xing, Z.; Pei, J.; and Yu, P. S. 2012. Early classification on time series. *KIS*. 31(1):105–127.
- Ye, L., and Keogh, E. J. 2009. Time series shapelets: a new primitive for data mining. In *KDD'09*, 947–956.
- Ye, L., and Keogh, E. 2011. Time series shapelets: a novel technique that allows accurate, interpretable and fast classification. *DMKD*. 22(1):149–182.
- Zhou, L.-k.; Yang, Y.; Ren, X.; Wu, F.; and Zhuang, Y. 2018. Dynamic network embedding by modeling triadic closure process. In *AAAI'18*.
- Zhu, L.; Guo, D.; Yin, J.; Ver Steeg, G.; and Galstyan, A. 2016. Scalable temporal latent space inference for link prediction in dynamic social networks. *IEEE Transactions on Knowledge and Data Engineering* 28(10):2765–2777.