# Towards Similarity-Aware Time-Series Classification

Daochen Zha*    Kwei-Herng Lai*    Kaixiong Zhou*    Xia Hu*

**Abstract**

We study time-series classification (TSC), a fundamental task of time-series data mining. Prior work has approached TSC from two major directions: (1) similarity-based methods that classify time-series based on the nearest neighbors, and (2) deep learning models that directly learn the representations for classification in a data-driven manner. Motivated by the different working mechanisms within these two research lines, we aim to connect them in such a way as to jointly model time-series similarities and learn the representations. This is a challenging task because it is unclear how we should efficiently leverage similarity information. To tackle the challenge, we propose **Sim**ilarity-Aware **T**ime-**S**eries **C**lassification (SimTSC), a conceptually simple and general framework that models similarity information with graph neural networks (GNNs). Specifically, we formulate TSC as a node classification problem in graphs, where the nodes correspond to time-series, and the links correspond to pair-wise similarities. We further design a graph construction strategy and a batch training algorithm with negative sampling to improve training efficiency. We instantiate SimTSC with ResNet as the backbone and Dynamic Time Warping (DTW) as the similarity measure. Extensive experiments on the full UCR datasets and several multivariate datasets demonstrate the effectiveness of incorporating similarity information into deep learning models in both supervised and semi-supervised settings. Our code is available at `https://github.com/daochenzha/SimTSC`.

## 1 Introduction

Time-series classification (TSC) is a fundamental task of time-series data mining. Given a collection of time-series with the attached labels, TSC aims to train a classifier to classify unseen time-series. With the increasing amount of temporal data available, TSC has broad applications, such as human activity recognition, health care, and cyber security [28].

Many research efforts have been devoted to TSC. Similarity-based (distance-based) methods are widely used [1]. The main idea is to combine a $k$-NN clas-

---
*Department of Computer Science, Rice University. {daochen.zha,khlai,Kaixiong.Zhou,Xia.Hu}@rice.edu

Figure 1: Average ranks ($\downarrow$) of ResNet and DTW on the full 128 UCR datasets, where different numbers of labels per class is given (see Section 4.1 for more details).

sifier with a similarity measure for classification. Dynamic Time Warping (DTW) [23, 26], which calculates the optimal match between two time-series, is one of the most popular similarity measures. It is shown that DTW plus a 1-NN classifier can achieve reasonably good accuracy [27]. However, the similarity is often obtained in an unsupervised fashion followed by a simple $k$-NN classifier, which could be sub-optimal. Another promising research line is deep learning. Without any crafting in feature engineering, deep learning methods perform end-to-end training on the raw time-series and learn the representations. Recent studies suggest that convolutional layers, such as ResNet and Fully Convolutional Networks (FCN), significantly outperform DTW and achieve competitive performance to the state-of-the-art TSC algorithms on the UCR benchmarks [30, 9].

While deep learning methods are simple and effective, they highly rely on the supervision of training labels for automatic representation learning; they thus often fall short when very few labels are given. Figure 1 compares the average ranks of ResNet and DTW plus 1-NN on the full 128 UCR datasets [7] with different numbers of provided training labels per class. While ResNet dominates DTW with sufficient training labels, it delivers unsatisfactory performance with very few labels. In contrast, DTW classifies the time-series by reasoning with pair-wise similarities instead of directly learning time-series representations, which could be less sensitive to the number of provided labels. Motivated by the different working mechanisms within these two research lines, we explore the possibility of connecting them in such a way as to jointly model time-series similarities and learn the representations.

However, it is non-trivial to achieve this goal due to the following challenges. First, it is unclear how we can incorporate similarity information into representation learning. The commonly used architectures, such as CNN and LSTM, cannot model similarity. While some papers have explored neural networks for time-series similarity learning [2, 12], they learn the similarity in the first place and then simply apply the learned similarity to classification, which still relies on $k$-NN. Second, even though we can enable similarity in deep learning models, it is challenging to balance similarity information and the original representation learning. Incorporating too much similarity information may lead to indistinguishable representations with poor accuracy.

To address these challenges, we propose **Sim**ilarity-Aware **T**ime-**S**eries **C**lassification (SimTSC), a conceptually simple and general framework for incorporating similarity information into deep learning models. Motivated by graph neural networks (GNNs) in modeling node relationships [8, 13, 17], we reformulate TSC as a node classification problem. Specifically, we correspond each time-series to a node in a graph and each pair-wise similarity to a link between nodes. A graph convolution layer is then applied on the top of a backbone (e.g., ResNet) to jointly perform feature extraction and model time-series similarities. We instantiate SimTSC with ResNet and DTW, which are the representative deep learning model and similarity measure, respectively. Extensive experiments suggest that this simple design improves accuracy, particularly with very few labels. We make the following contributions:

- Explore the possibility of connecting the research efforts of similarity-based methods and deep learning models for TSC.

- Propose SimTSC, a simple and general framework that can combine any similarity measures with any deep learning models from the view of graphs.

- Design a graph construction strategy that focuses on the top neighbors for efficient aggregation of graph convolution. We also develop a batch training algorithm with negative sampling to enable the training of SimTSC on large datasets.

- Instantiate SimTSC with ResNet and DTW. We conduct extensive experiments on the full 128 UCR datasets and several multivariate datasets. SimTSC outperforms ResNet, DTW, and the state-of-the-art supervised and semi-supervised deep models significantly, particularly when very few labels are given. We also present extensive hyperparameter studies and ablations.

Table 1: Main symbols and definitions.

| Symbol | Definition |
|--------|-----------|
| $x \in \mathbb{R}$ | A real-value in a time-series |
| $\mathbf{x} \in \mathbb{R}^T$ | A univariate time-series with length $T$ |
| $\mathbf{X} \in \mathbb{R}^{M \times T}$ | An $M$-dimensional time-series with length $T$ |
| $\mathcal{X} \in \mathbb{R}^{N \times M \times T}$ | A 3D matrix consisting of $N$ multivariate time-series |
| $\mathcal{X}^{\text{train}}$ | A set of training time-series |
| $\mathbf{y}^{\text{train}}$ | The labels corresponding to $\mathcal{X}^{\text{train}}$ |
| $\mathcal{X}^{\text{test}}$ | A set of testing time-series |
| $\mathbf{y}^{\text{test}}$ | The labels corresponding to $\mathcal{X}^{\text{test}}$ |
| $\mathcal{X}^{\text{unlabeled}}$ | A set of unlabeled time-series |
| $d(\mathbf{X}_1, \mathbf{X}_2)$ | The similarity (distance) of two time-series |
| $\mathbf{D} \in \mathbb{R}^{N \times N}$ | The similarity (distance) matrix |
| $\tilde{\mathbf{A}} \in \mathbb{R}^{N \times N}$ | The normalized adjacency matrix in graph |
| $\tilde{\mathbf{X}} \in \mathbb{R}^{N \times M}$ | The attribute information matrix in graph |
| $\alpha$ | A scaling factor |
| $K$ | The number of neighbors for each node |

## 2 Preliminaries

We start with a problem description and then provide a background of time-series similarity measure, deep learning for TSC, and graph neural networks.

**2.1 Problem Statement** We use lowercase alphabet, e.g., $x \in \mathbb{R}$, to represent a scalar value, lowercase boldface letter, e.g., $\mathbf{x} = [x_1, x_2, ..., x_T] \in \mathbb{R}^T$, to denote a vector of length $T$, uppercase boldface alphabet, e.g., $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_M] \in \mathbb{R}^{M \times T}$, to denote a matrix consisting of $M$ vectors, where each vector can have a different length, and calligraphic font, e.g., $\mathcal{X} = [\mathbf{X}_1, \mathbf{X}_2, ..., \mathbf{X}_N] \in \mathbb{R}^{N \times M \times T}$, to denote a 3D matrix. We summarize the main symbols in Table 1. We first give formal definitions of time-series.

DEFINITION 1. (UNIVARIATE TIME-SERIES) *A univariate time-series $\mathbf{x}$ of length $T$ is represented as a vector $[x_1, x_2, ..., x_T]$.*

DEFINITION 2. (MULTIVARIATE TIME-SERIES) *An $M$-dimensional time-series $\mathbf{X}$ consists of $M$ univariate time-series $[\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_M]$.*

Without loss of generality, we unify the above definitions by expanding dimension for univariate time-series. Specifically, we regard a univariate time-series $\mathbf{x}$ as a 1-dimensional time-series $\mathbf{X} \in \mathbb{R}^{1 \times T}$, that is, a univariate time-series is a special case of multivariate time-series with $M = 1$. We will use the unified notation $\mathbf{X}$ to represent a time-series throughout the paper.

We formally describe the problem of TSC. Given some testing time-series $\mathcal{X}^{\text{test}} = [\mathbf{X}_1, \mathbf{X}_2, ..., \mathbf{X}_{N^{\text{test}}}]$ and the labels $\mathbf{y}^{\text{test}} = [y_1, y_2, ..., y_{N^{\text{test}}}]$, where $N^{\text{test}}$ is the number of testing time-series, we aim to train a classifier that can predict the labels based on $\mathcal{X}^{\text{test}}$ under one of the following settings:

- **Supervised setting**: The classifier is trained based on a training time-series dataset $\mathcal{X}^{\text{train}} = [\mathbf{X}_1, \mathbf{X}_2, ..., \mathbf{X}_{N^{\text{train}}}]$ and its corresponding labels $\mathbf{y}^{\text{train}} = [y_1, y_2, ..., y_{N^{\text{train}}}]$, where $N^{\text{train}}$ is the number of training time-series.

- **Inductive semi-supervised setting**: In addition to $\mathcal{X}^{\text{train}}$ and $\mathbf{y}^{\text{train}}$, the classifier can also access some unlabeled time-series $\mathcal{X}^{\text{unlabeled}}$, which does not overlap with $\mathcal{X}^{\text{test}}$.

- **Transductive semi-supervised setting**: In addition to $\mathcal{X}^{\text{train}}$, $\mathbf{y}^{\text{train}}$ and $\mathcal{X}^{\text{unlabeled}}$, the classifier is exposed to testing time-series $\mathcal{X}^{\text{test}}$. Note that $\mathbf{y}^{\text{test}}$ is not accessible in training.

The above settings differ in how much unlabeled time-series data the classifier can access. With more unlabeled data, the classifier could better learn the underlying data distributions and often achieve better accuracy.

**2.2 Time-Series Similarity Measure** Given two time-series $\mathbf{X}_1$ and $\mathbf{X}_2$, we aim at providing a distance $d(\mathbf{X}_1, \mathbf{X}_2)$, such that similar time-series tend to have smaller $d(\mathbf{X}_1, \mathbf{X}_2)$. Dynamic Time Warping (DTW) [23] is one of the most popular ones. The key idea is to calculate the optimal match between two time-series such that the sum of matched series has the smallest values. The troughs and peaks of the same pattern can be matched even if they are not perfectly synced up. DTW is a standard tool with many efficient implementations, such as the UCR Suite [26]. In this work, we adopt DTW as the similarity measure; one can also use other similarity measures under our framework.

**2.3 Deep Neural Networks for TSC** Numerous deep learning models have been developed for time-series classification [30, 15, 16, 10]. In this work, we mainly focus on Residual Network (ResNet) since it is shown to have superior performance on the majority of UCR Time Series Classification Archive [9]. The network consists of multiple residual blocks. Each block consists of three 1D convolutional layers followed by batch normalization and a ReLU activation function, with shortcuts to enable a direct flow of the gradient:

$$(2.1) \qquad \mathbf{H}_1 = \text{ReLU}(\text{BatchNorm}(\text{Conv1d}(\mathbf{X}))),$$

$$(2.2) \qquad \mathbf{H}_2 = \text{ReLU}(\text{BatchNorm}(\text{Conv1d}(\mathbf{H}_1))),$$

$$(2.3) \qquad \mathbf{H}_3 = \text{ReLU}(\text{BatchNorm}(\text{Conv1d}(\mathbf{H}_2))),$$

$$(2.4) \qquad \hat{\mathbf{H}} = \text{ReLU}(\mathbf{H}_3 + \mathbf{X}),$$

where $\hat{\mathbf{H}}$ is the output of residual block, Conv1d($\cdot$) denotes 1D convolutional, BatchNorm($\cdot$) is batch normalization, and ReLU($\cdot$) is ReLU activation function.

**2.4 Graph Neural Networks** Graph neural networks (GNNs) have achieved great success in modeling node dependencies in graph [8, 17, 37]. To capture the node dependency, Graph Convolution Network (GCN) [17] performs joint learning of feature extraction and aggregation of neighboring nodes. Given

a graph $G = (\widetilde{\mathbf{A}}, \widetilde{\mathbf{X}})$, where $\widetilde{\mathbf{A}} \in \mathbb{R}^{N \times N}$ is the normalized adjacency matrix, $\widetilde{\mathbf{X}} \in \mathbb{R}^{N \times M}$ is the attribute information matrix, $N$ is the number of nodes, and $M$ is the feature dimension, a GCN layer performs feature aggregation of neighboring nodes with

$$(2.5) \qquad \hat{\mathbf{Z}} = \widetilde{\mathbf{A}}\widetilde{\mathbf{X}}\mathbf{W},$$

where $\mathbf{W}$ denotes trainable parameters, and $\hat{\mathbf{Z}}$ denotes the output of the GCN layer. Similar to ResNet, we can stack multiple GCN layers with activation functions. While there are hundreds of GNNs, we adopt the basic GCN to make our contribution focused; one can also use other GNNs under our framework.

## 3 Methodology

Figure 2 shows an overview of **Sim**ilarity-Aware **T**ime-**S**eries **C**lassification (SimTSC), which consists of three modules: (1) a graph construction module that connects the time-series based on a similarity measure (e.g., DTW), (2) a backbone that performs feature extraction with deep neural networks (e.g., ResNet), and (3) a GNN module that aggregates the features of neighboring time-series (e.g., GCN). The graph construction is unsupervised so that it can flexibly adapt to all the three settings defined in Section 2.1.

**3.1 Graph Construction with Similarity** This subsection describes how we construct a graph based on a similarity measure of time-series, e.g., DTW.

Let $\mathcal{X} = [\mathbf{X}_1, \mathbf{X}_2, ..., \mathbf{X}_N]$ denote all the accessible time-series. In the supervised setting, $\mathcal{X}$ is simply all the training data. In the semi-supervised settings defined in Section 2.1, the $\mathcal{X}$ consists of both labeled and unlabeled time-series. Based on the learned similarity measure $d(\cdot, \cdot)$, we can obtain a similarity matrix for $\mathcal{X}$ as

(3.6)
$$\mathbf{D} = \begin{bmatrix} d(\mathbf{X}_1, \mathbf{X}_1) & d(\mathbf{X}_1, \mathbf{X}_2) & \cdots & d(\mathbf{X}_1, \mathbf{X}_N) \\ d(\mathbf{X}_2, \mathbf{X}_1) & d(\mathbf{X}_2, \mathbf{X}_2) & \cdots & d(\mathbf{X}_2, \mathbf{X}_N) \\ \vdots & \vdots & \ddots & \vdots \\ d(\mathbf{X}_N, \mathbf{X}_1) & d(\mathbf{X}_N, \mathbf{X}_2) & \cdots & d(\mathbf{X}_N, \mathbf{X}_N) \end{bmatrix}.$$

Given the similarity matrix $\mathbf{D}$, we construct the graph as follows. First, we introduce a scaling hyper-parameter $\alpha \in [0, \infty)$ to control the importance of top neighbors. Specifically, let $\mathbf{D}_{ij}$ denote the $(i, j)^{th}$ entry of $\mathbf{D}$. The adjacency matrix $\mathbf{A}$ is obtained by

$$(3.7) \qquad \mathbf{A}_{ij} = \frac{1}{e^{\alpha \mathbf{D}_{ij}}}, \forall i, j,$$

where $\mathbf{A}_{ij}$ denotes the $(i, j)^{th}$ entry of $\mathbf{A}$. A larger $\alpha$ will give more importance to the top neighbors. When $\alpha = 0$, each node will equally aggregate the features of all
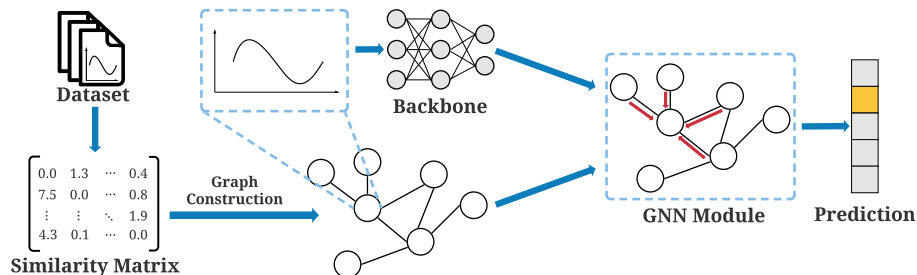
Figure 2: An overview of SimTSC framework. The graph is constructed based on the pair-wise similarities (e.g., DTW distances) of the time-series. Each time-series is processed by a backbone (e.g., ResNet) for feature extraction. The GNN module will aggregate the features and produce the final representations for classification.

the nodes, and all the nodes will have indistinguishable features. When $\alpha \to \infty$, our framework reduces to the backbone itself since $\mathbf{A}$ reduces to a diagonal matrix.

Second, to filter out irrelevant neighbors, we sample the top-$K$ neighbors for each node. Specifically, for each row $\mathbf{a}_i$ in $\mathbf{A}$, we only keep $K$ entries with the largest weights and zero out the others, which leads to a sparse matrix. Finally, we normalize the adjacency matrix with $\widetilde{\mathbf{A}}_{ij} = \frac{\mathbf{A}_{ij}}{\sum_{j'} \mathbf{A}_{ij'}}, \forall i, j$.

### 3.2 Joint Learning of Backbone and Graph Convolution Layers
This subsection introduces how we optimize the backbone and the GNN module.

The backbone takes all the accessible time-series $\mathcal{X}$ as inputs and produces an attribute information matrix $\widetilde{\mathbf{X}} \in \mathbb{R}^{N \times M}$, where $N$ denotes the number of all the accessible time-series, and $M$ denotes the feature dimension. Here, the output of the backbone is assumed to be flattened to be 1-dimensional. Given $\widetilde{\mathbf{X}}$ and $\widetilde{\mathbf{A}}$, the GNN aggregates the node features and produces the final representation followed by a softmax layer for classification.

Let $\widetilde{\mathbf{Z}} \in \mathbb{R}^{N \times C}$ denote the final output of the model, where $C$ is the number classes, and $\widetilde{\mathbf{Z}}^{\text{train}} \in \mathbb{R}^{N^{\text{train}} \times C}$ denotes the rows that have labels. The objective is to minimize the cross-entropy over the labeled time-series:

$$(3.8) \qquad \mathcal{L} = - \sum_{i=1}^{N^{\text{train}}} \sum_{j=1}^{C} \mathbf{Y}_{ij}^{\text{train}} \log \widetilde{\mathbf{Z}}_{ij}^{\text{train}},$$

where $\mathbf{Y}_{ij}^{\text{train}}$ denotes the $j^{th}$ element of one-hot encoded label of the $i^{th}$ labeled time-series. The weights of the backbone and the GNN module can be jointly optimized using gradient descent.

### 3.3 Batch Training with Negative Sampling
This subsection presents how we handle large datasets with batch training and how we use negative sampling to improve efficiency.

A naive training strategy is to put all the time-series into the GPU memory in the first place and then train the network with backpropagation. However, this strategy has the following limitations. First, it cannot scale to large datasets with many time-series or very long time-series. Second, each time-series can only interact with a fixed number of neighbors since the constructed graph is static. However, the top-$K$ dropping strategy may improperly drop some important connections, which leads to sub-optimal performance.

To overcome the above limitations, we propose a batch training strategy with negative sampling as follows. Given a batch size $B$, we sample $B/2$ labeled and $B/2$ unlabeled time-series, respectively, as a batch for training. The unlabeled time-series are "negatively sampled" in that we can usually access far more unlabeled data in real-world applications. Then, we construct a graph with this batch of time-series and update the model accordingly. The above training procedure is summarized in Algorithm 1. Here, $\mathcal{X}^{\text{unlabeled}}$ will also cover testing data in the transductive setting[1].

This design has three nice properties. First, it can scale to large datasets since we only need to put a batch of data into GPU memory in each update step. Second, each node will interact with more neighbors for aggregation since a new graph will be constructed in every randomly sampled batch. Third, compared with random sampling, negative sampling ensures that half of the data in each batch is labeled so that we can have sufficient learning signals to update the weights. In particular, if very few labels are available, random sampling may result in very few or even no labeled data in a sampled batch, which will reduce training efficiency.

At testing time, we similarly sample $B/2$ testing and $B/2$ non-testing time-series in each batch to construct the graph. In this way, we can aggregate some of the non-testing representations to enhance the representations of the testing samples.

---

[1]In the supervised setting, we only sample labeled time-series.

---

**Algorithm 1** Batch training of SimTSC with negative sampling

---

1: **Input:** Labeled time-series $\mathcal{X}^{\text{train}}$, unlabeled time-series $\mathcal{X}^{\text{unlabeled}}$, similarity measure $d(\cdot, \cdot)$, batch size $B$
2: $\mathcal{X} \leftarrow \{\mathcal{X}^{\text{train}}, \mathcal{X}^{\text{unlabeled}}\}$
3: Precompute the similarity matrix $\mathbf{D}$ for $\mathcal{X}$ based on Eq. 3.6
4: **for** each epoch **do**
5:     **for** each batch $\mathcal{X}^{\text{train}}_{\text{batch}}$ of size $B/2$ in $\mathcal{X}^{\text{train}}$ **do**
6:         Sample a batch $\mathcal{X}^{\text{unlabeled}}_{\text{batch}}$ of size $B/2$ from $\mathcal{X}^{\text{unlabeled}}$
7:         $\mathcal{X}_{\text{batch}} \leftarrow \{\mathcal{X}^{\text{train}}_{\text{batch}}, \mathcal{X}^{\text{unlabeled}}_{\text{batch}}\}$
8:         Obtain submatrix $\mathbf{D}_{\text{batch}}$ from $\mathbf{D}$ with sampled indices
9:         Construct adjacency matrix $\mathbf{A}$ from $\mathbf{D}_{\text{batch}}$ based on Eq. 3.7
10:        Compute normalized matrix $\widetilde{\mathbf{A}}$ with $\widetilde{\mathbf{A}}_{ij} = \frac{\mathbf{A}_{ij}}{\sum_{j'} \mathbf{A}_{ij'}}, \forall i, j$
11:         Obtain model output $\widetilde{\mathbf{Z}}$ with $\mathcal{X}_{\text{batch}}$ and $\widetilde{\mathbf{A}}$
12:         Update the model weights based on Eq. 3.8
13:     **end for**
14: **end for**

---

## 4 Experiments

We evaluate SimTSC across various settings to answer the following research questions: **RQ1**: How does SimTSC compare with existing deep learning and similarity-based methods (Section 4.2)? **RQ2:** Can SimTSC also enhance other neural architectures, such as MLP and FCN (Section 4.3)? **RQ3:** How will the number of graph convolutional layers impact the performance (Section 4.4)? **RQ4:** How will graph structure impact the performance of SimTSC (Section 4.5)? **RQ5:** Is the proposed negative sampling strategy effective (Section 4.6)? **RQ6:** Can SimTSC also be applied to multivariate time-series classification tasks (Section 4.7)? **RQ7:** How does SimTSC learn the representations with similarity information (Section 4.8)?

**4.1 Experimental Setup** As suggested in [7], we evaluate the performance on the full 128 UCR datasets[2]. We merge the original training and testing data to create new splits for all the datasets to simulate the three settings defined in Section 2.1. First, we randomly split 20% of the data as the hold-out set for testing purposes, denoted as $\mathcal{X}^{\text{test}}$. Second, we vary the number of training labels per class from the set $\{1, 5, 10, 15, 20, 25, 30, 35, 40, 45, 50\}$ to test different levels of supervision and sample a subset from the remain-

---

ing 80% of data as the training data to create few-shot settings, denoted as $\mathcal{X}^{\text{train}}$. Third, the time-series out of the above two splits will serve as another split of some unlabeled data, denoted as $\mathcal{X}^{\text{unlabeled}}$. We further consider four multivariate datasets: Character Trajectories [3], ECG [25], KickVsPunch [5], and NetFlow [5].

**Metric.** For UCR datasets, we follow previous work [7, 9] and rank the algorithms on each dataset according to the mean accuracy and report the average ranks across the 128 datasets. We perform Wilcoxon signed-rank test with a significance level of 0.05. For the multivariate datasets, we report accuracies.

**Baselines.** For deep learning models, we consider four **supervised** architectures, including MLP, Fully Convolutional Network (FCN) [30], ResNet [11], and InceptionTime [10], and a state-of-the-art **semi-supervised** framework TapNet [33]. For **similarity-based** methods, we include DTW with a 1 Nearest-Neighbor (1-NN) classifier [27]. We train SimTSC in supervised, inductive semi-supervised, and transductive semi-supervised settings, denoted as SimTSC-S, SimTSC-I, and SimTSC-T, respectively.

**Implementation Details.** We use DTW as the similarity measure and ResNet as the backbone. The hyperparamters are set based on the accuracy on the training data, with the scaling factor $\alpha$ as 0.3, the number of neighbors $K$ as 3, one GCN layer, the batch size as 128, and the number of epochs to be 500, across all the experiments. For a fair comparison, the backbones used in SimTSC are exactly the same as the baselines. We use the authors' implementations of InceptionTime[3] and TapNet[4] with the default hyperparameters. We run five times and report the average accuracy. More details of the neural architectures, hardware, and the dataset statistics are provided in supplementary materials.

**4.2 Performance Comparison on Benchmarks** To study **RQ1**, we report the average ranks of SimTSC on the UCR datasets in Table 2, separately present TapNet in Table 3 because TapNet does not support batch training and suffers from memory explosion on 45 datasets. We make the following observations.

**First**, SimTSC outperforms the baselines with very few labels. Given 1, 5, or 10 labels per class, SimTSC trained in all the settings achieve better ranks than ResNet significantly. Given more labels, i.e, 15, 20, 25, 30 , or 35, SimTSC also beats ResNet consistently. An interesting observation is that SimTSC-S performs well even though it only uses very few time-series to construct the graph. A possible explanation is that the graph may serve as a regularizer for ResNet to help

---

Table 2: Average ranks ($\downarrow$) of SimTSC and baselines on the full 128 UCR datasets with different numbers of training labels per class. †, ▲, and ▽ to denote the cases where SimTSC-S, SimTSC-I, and SimTSC-T are significantly better than the other algorithms w.r.t. the Wilcoxon signed rank test ($p < 0.05$), respectively.

| Algorithm \ Labels | 1 | 5 | 10 | 15 | 20 | 25 | 30 | 35 | 40 | 45 | 50 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| DTW | 3.776 | 4.163 | 4.465▲▽ | 4.738▲▽ | 4.824†▲▽ | 5.048†▲▽ | 4.965▲▽ | 5.160†▲▽ | 5.309†▲▽ | 5.199†▲▽ | 5.211†▲▽ |
| MLP | 5.504†▲▽ | 5.496†▲▽ | 5.438†▲▽ | 5.309†▲▽ | 5.316†▲▽ | 5.256†▲▽ | 5.367†▲▽ | 5.477†▲▽ | 5.195†▲▽ | 5.402†▲▽ | 5.348†▲▽ |
| FCN | 4.630▲▽ | 4.310 | 4.383▲▽ | 4.508▲▽ | 4.723†▲▽ | 4.803†▲▽ | 4.699†▲▽ | 4.910†▲▽ | 4.883†▲▽ | 4.852†▲▽ |
| ResNet | 4.846†▲▽ | 4.857†▲▽ | 4.617†▲▽ | 4.047 | 4.449▲▽ | 4.039 | 4.102 | 4.090 | 4.086 | 3.840 | 3.895 |
| InceptionTime | 5.484†▲▽ | 5.302†▲▽ | 5.438†▲▽ | 5.434†▲▽ | 5.215†▲▽ | 5.145†▲▽ | 5.168†▲▽ | 4.914†▲▽ | 4.941†▲▽ | 5.066†▲▽ | 5.039†▲▽ |
| SimTSC-S | 4.224▲ | 4.278▲▽ | 4.074 | 4.277▲▽ | 4.141▲▽ | 4.044▽ | 4.148▲ | 3.988 | 3.887 | 3.918 | 4.047 |
| SimTSC-I | **3.724** | 3.817 | 3.793 | **3.836** | 3.746 | 4.031▲▽ | **3.762** | 3.734 | **3.852** | 3.867 | **3.797** |
| SimTSC-T | 3.811 | **3.778** | **3.781** | 3.852 | **3.586** | **3.632** | 3.789 | **3.727** | 3.957 | **3.824** | 3.812 |

Table 3: Average ranks of TapNet and SimTSC on 83 datasets, on which TapNet does not suffer from memory explosion. ▽ suggests SimTSC is significantly better.
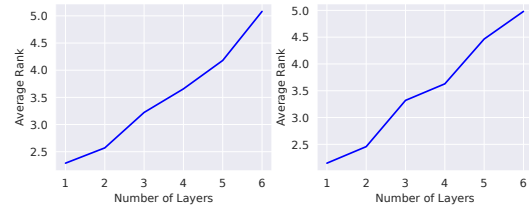
| Algorithm \ Labels | 5 | 10 | 15 | 20 | 25 |
|---|---|---|---|---|---|
| TapNet | 1.645▽ | 1.524 | 1.548 | 1.536 | 1.530 |
| SimTSC | **1.355** | **1.476** | **1.452** | **1.464** | **1.470** |

Table 4: Average ranks of using MLP and FCN as backbones. †, ▲, and ▽ suggest SimTSC-S, SimTSC-I, and SimTSC-T are significantly better, respectively.
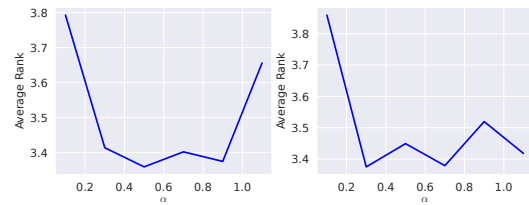
| Algorithm \ Labels | 10 | 20 | 30 | 40 |
|---|---|---|---|---|
| MLP | 3.137†▲▽ | 2.906†▲▽ | 2.879†▲▽ | 2.898†▲▽ |
| SimTSC-S with MLP backbone | 2.738▲▽ | 2.633▲▽ | 2.618▲▽ | 2.578▲▽ |
| SimTSC-I with MLP backbone | 2.078 | 2.258 | 2.414▽ | **2.258** |
| SimTSC-T with MLP backbone | **2.047** | **2.203** | **2.090** | 2.266 |
| FCN | 2.672 | 2.695▽ | 2.641 | 2.563 |
| SimTSC-S with FCN backbone | 2.523 | 2.492 | 2.504 | **2.438** |
| SimTSC-I with FCN backbone | 2.500▽ | 2.590▽ | 2.539 | 2.547 |
| SimTSC-T with FCN backbone | **2.309** | **2.223** | **2.316** | 2.453 |



(a) 10 labels per class   (b) 20 labels per class
Figure 3: Impact of the number of GCN layers.



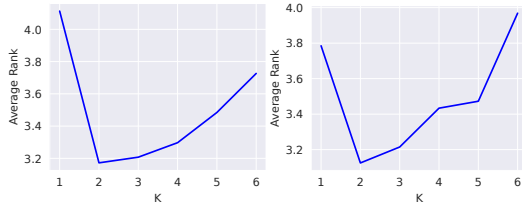(a) 10 labels per class   (b) 20 labels per class
Figure 4: Impact of $\alpha$.

alleviate the overfitting issue. **Second**, similarity information is less effective when we have sufficient labels. Given 40, 45, or 50 labels, ResNet and SimTSC achieve similar ranks. Nevertheless, the results suggest that SimTSC can still deliver competitive performance with enough labels, which shows the flexibility of SimTSC. **Third**, SimTSC tends to perform better with more unlabeled data. With very few exceptions, SimTSC-T > SimTSCI > SimTSC-S. This is because we can construct a better graph to capture the underlying data distributions with more unlabeled time-series. **Fourth**, SimTSC consistently outperforms TapNet across all different numbers of labels when both evaluated under transductive semi-supervised setting. The result again verifies the effectiveness of SimTSC.

**4.3   Results on Other Neural Architectures** To investigate **RQ2**, we show the results of applying SimTSC on MLP and FCN on the UCR datasets in Table 4. SimTSC can also enhance these two neural architectures. In particular, we observe significant performance gains when applying SimTSC on MLP.

**4.4   Analysis of Graph Convolutional Layers** To answer **RQ3**, we report the ranks using more GCN layers. We focus on the transductive setting with 10 and 20 labels per class (Figure 3). SimTSC achieves the best performance with only one GCN layer. We speculate that this is because the graph is constructed based on the pair-wise similarity, which can fully capture the relationship between each pair of time-series. As such, stacking more GCN layers will not introduce more information but instead makes the model more susceptible to over-smoothing [35, 36]. Nevertheless, stacking more GCN layers could help in larger time-series datasets by computing a submatrix and leverage multiple GCN layers to capture the multi-hop connections, which is deferred as our future work.

**4.5   Analysis of Graph Structure** For **RQ4**, we vary the hyperparameters of graph construction. In Figure 4, we vary $\alpha$ from the set $\{0.1, 0.3, 0.5, 0.7, 0.9, 1.1\}$ with $K$ fixed to be 3. We observe a significant performance drop when $\alpha = 0.1$. This is because when $\alpha \to 0$, each node will aggregate more information from the neighbors, and the resulting representation will become

(a) 10 labels per class    (b) 20 labels per class

Figure 5: Impact of K.



(a) ResNet (83% accuracy) (b) SimTSC (100% accuracy)

Figure 6: Learned representations of ResNet and SimTSC on Coffee with 56 time-series, two classes marked in blue and green, respectively, and only one time-series labeled in each class (circled in red).

Table 5: Comparison of negative and random sampling on the full UCR datasets (top). Comparison with the variant without batch training on 121 UCR datasets (bottom), on which w/o batch training does not suffer from GPU memory explosion. ▲ suggests negative sampling is significantly better than the other algorithms.

| Algorithm          Labels | 10 | 20 | 30 | 40 |
|---|---|---|---|---|
| random sampling | 1.637▲ | 1.609▲ | 1.547 | 1.582▲ |
| negative sampling | **1.363** | **1.391** | **1.453** | **1.418** |
| random sampling | 2.152▲ | 1.979▲ | 1.868 | 1.926▲ |
| w/o batch training | 2.068▲ | 2.277▲ | 2.331▲ | 2.322▲ |
| negative sampling | **1.780** | **1.744** | **1.802** | **1.752** |

less distinguishable. There is also a performance drop when $\alpha = 1.1$. A larger $\alpha$ will reduce the similarity information in the model, which suggests that incorporating similarity information is indeed helpful. In Figure 5, we vary $K$ from the set $\{1, 2, 3, 4, 5, 6\}$ with $\alpha$ fixed to be 0.3. The best performance is achieved with $K = 2$. When $K = 1$, our framework reduces to backbone itself since there will be only self-connections in the graph. When $K$ becomes larger, the performance also drops since aggregating more information from the neighbors makes the representations less distinguishable. Overall, we find that incorporating an appropriate amount of similarity information lead to the best performance.

**4.6 Ablation Study** For **RQ5**, we consider two ablations to study the effectiveness of negative sampling. First, we consider a variant that uses random sampling for batch training. Second, we consider a variant that puts all the time-series into GPU memory without batch training. Unfortunately, this variant will explode the GPU memory for 7 of the datasets so that only 121 datasets are reported. We perform a grid-search of $K$ and find that $K = 10$ achieves the best performance.

Table 5 summarizes the results. First, we observe that negative sampling outperforms random sampling significantly. Second, the variant without batch training does not perform well. A possible explanation is that the graph is pre-constructed so that each time-series can only interact with a limited number of connected neighbors, which may lead to sub-optimal performance.
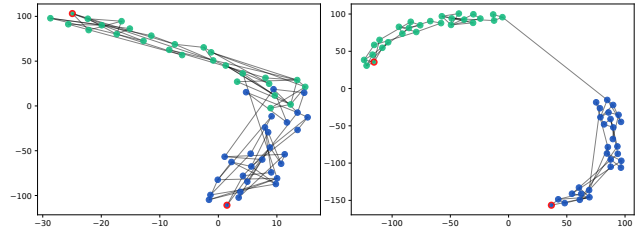
**4.7 Results on Multivariate Datasets** For **RQ6**, we evaluate SimTSC on multivariate time-series classification tasks in Table 6. First, SimTSC outperforms ResNet and DTW with very few labels. Given 5, 10, or 15 labels, at least one of the SimTSC variants achieves the best performance. Second, when more labels are given, the similarity information helps on some datasets but worsens some others' performance. SimTSC delivers poor performance on ECG in the transductive setting given 50 labels, while we observe a consistent improvement on CharacterTrajectories and NetFlow.

**4.8 Visualization of Learned Representations** To answer **RQ7**, we conduct a case study on the Coffee dataset from UCR Archive. Figure 6 visualizes the learned representations of SimTSC and ResNet as well as the constructed graphs. We observe that there is an overlap between the two classes' representations learned by ResNet. Thus, ResNet cannot distinguish those overlapped time-series and only gives 83% accuracy. Whereas, the representations learned by SimTSC form clear clusters so that SimTSC achieves 100% accuracy. A possible reason is that the overlapped time-series tend to be close to those in the same class in terms of DTW, and thus their representations are corrected by aggregation. SimTSC achieves better accuracy by jointly performing feature extraction and aggregation.

## 5 Related Work

**Deep Learning for TSC.** Prior deep learning models for TSC can be mainly grouped into (1) learning representations of time-series in an unsupervised manner and then applying a classifier to the learned representations [22], and (2) training a classifier in an end-to-end fashion [11, 4, 15]. However, the existing models mainly focus on feature extraction but cannot capture similarity information. Although [32] has introduced GNNs to capture time-series dependencies, they focus on time-series forecasting and can only model the dependencies among multivariate time-series. Whereas,

Table 6: Classification accuracy on multivariate time-series datasets across different numbers of training labels per class. For KickvsPunch, we only report the performance up to 15 labels since there are less than 15 labels per class. TapNet is not reported on Character Trajectories and Netflow due to memory explosion.

| Dataset | Algorithm (Labels) | 5 | 10 | 15 | 20 | 25 | 30 | 35 | 40 | 45 | 50 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Character Trajectories | DTW | .847±.014 | .881±.005 | .895±.009 | .900±.014 | .908±.014 | .907±.013 | .906±.010 | .906±.007 | .909±.010 | .913±.008 |
| | ResNet | .834±.024 | .898±.017 | .920±.010 | .937±.010 | .939±.008 | .941±.009 | .949±.008 | .956±.011 | .958±.007 | .958±.007 |
| | InceptionTime | .883±.010 | .939±.007 | .947±.006 | **.968±.006** | .964±.003 | .974±.005 | .979±.003 | .978±.005 | .979±.001 | **.986±.003** |
| | TapNet | - | - | - | - | - | - | - | - | - | - |
| | SimTSC-S | .894±.020 | .939±.009 | .949±.007 | .947±.017 | .964±.011 | .975±.003 | .977±.011 | .975±.004 | .981±.007 | .982±.005 |
| | SimTSC-I | **.914±.012** | .944±.009 | .951±.015 | .953±.012 | **.969±.011** | **.978±.006** | **.981±.007** | .979±.005 | .977±.008 | .980±.003 |
| | SimTSC-T | .903±.014 | **.946±.005** | **.957±.011** | .964±.009 | .967±.012 | .973±.009 | .976±.009 | **.981±.006** | **.983±.008** | **.986±.004** |
| ECG | DTW | .605±.124 | .670±.086 | .740±.112 | .755±.103 | .805±.043 | .825±.050 | .805±.053 | .800±.057 | .805±.053 | .800±.057 |
| | ResNet | .745±.048 | .795±.037 | .805±.058 | .800±.079 | **.860±.030** | **.855±.048** | .850±.052 | **.855±.029** | .830±.037 | **.870±.029** |
| | InceptionTime | .750±.045 | .805±.033 | .785±.020 | .800±.037 | .820±.037 | .830±.043 | .825±.016 | .850±.027 | **.855±.0010** | .850±.016 |
| | TapNet | .770±.043 | .780±.012 | .755±.025 | .795±.048 | .810±.037 | .795±.029 | .785±.025 | .815±.037 | .830±.019 | .845±.024 |
| | SimTSC-S | .795±.043 | .810±.020 | **.855±.040** | .840±.051 | .830±.056 | .840±.020 | **.860±.041** | .825±.047 | .830±.071 | .860±.025 |
| | SimTSC-I | .790±.062 | .765±.072 | .830±.070 | .730±.159 | .740±.087 | .800±.091 | .830±.048 | .750±.052 | .790±.108 | .735±.108 |
| | SimTSC-T | **.810±.041** | **.815±.046** | .770±.108 | .815±.115 | .730±.118 | .745±.075 | .745±.099 | .780±.051 | .775±.071 | .710±.101 |
| KickvsPunch | DTW | .433±.082 | .433±.082 | .433±.082 | - | - | - | - | - | - | - |
| | ResNet | .667±.183 | **.833±.149** | .833±.183 | - | - | - | - | - | - | - |
| | InceptionTime | .667±.000 | .533±.125 | .567±.226 | - | - | - | - | - | - | - |
| | TapNet | .700±.125 | .767±.082 | .733±.013 | - | - | - | - | - | - | - |
| | SimTSC-S | **.733±.200** | .767±.133 | **.867±.125** | - | - | - | - | - | - | - |
| | SimTSC-I | .700±.125 | **.833±.105** | .800±.125 | - | - | - | - | - | - | - |
| | SimTSC-T | .600±.133 | .767±.133 | .767±.082 | - | - | - | - | - | - | - |
| NetFlow | DTW | .611±.016 | .559±.128 | .607±.132 | .595±.118 | .546±.103 | .568±.125 | .523±.154 | .481±.203 | .503±.217 | .504±.214 |
| | ResNet | .613±.074 | .714±.063 | .749±.022 | .763±.038 | .739±.058 | .767±.050 | .769±.054 | .767±.049 | .787±.026 | .797±.039 |
| | InceptionTime | .418±.052 | .456±.046 | .484±.058 | .618±.049 | .642±.036 | .657±.024 | .678±.014 | .675±.036 | .681±.018 | .681±.015 |
| | TapNet | - | - | - | - | - | - | - | - | - | - |
| | SimTSC-S | .519±.108 | .720±.071 | .705±.055 | .709±.089 | .738±.082 | .786±.036 | .765±.091 | .790±.045 | .784±.063 | .799±.047 |
| | SimTSC-I | .766±.043 | .788±.036 | .689±.139 | **.776±.042** | .731±.084 | .755±.104 | **.834±.037** | .798±.066 | .810±.065 | .839±.035 |
| | SimTSC-T | **.769±.052** | **.805±.035** | **.785±.101** | .766±.095 | **.745±.092** | **.825±.029** | .801±.065 | **.827±.059** | **.847±.023** | **.852±.028** |

we adopt GNNs to jointly perform feature extraction and model the time-series similarities for TSC.

**Similarity-Based TSC Methods.** Similarity-based methods first obtain time-series similarities and then use a $k$-NN classifier for classification [24, 31]. Recently, some similarity learning approaches have been proposed to learn the similarities [2, 12]. However, they often rely on a separate procedure for classification. In contrast, we connect time-series similarity with deep models under a unified framework.

**Leveraging Unlabeled Data in Time-Series.** Prior work has explored semi-supervised learning [29], domain adaptation [20], and anomaly detection [21, 34, 18, 19] on the unlabeled data, most of which do not target deep models. Recently, [14] proposes a deep TSC method with auxiliary forecasting tasks. [33] augments networks with task-adaptive projection. However, they do not support batch training. Unlike the previous work, our framework brings benefits to not only semi-supervised setting but also the supervised setting.

## 6 Conclusions and Future Work

This work explores connecting the research efforts of time-series similarity measuring and deep learning for TSC. To jointly model feature extraction and similarity information, we formulate TSC as a node classification problem in graphs and introduce GNNs on the top of a backbone to enable end-to-end training. We instantiate our framework with ResNet and DTW with extensive experiments on the full 128 UCR datasets and several multivariate datasets. Experimental results suggest that incorporating similarity information can improve deep models significantly. In the future, we will investigate differentiable DTW [6] for graph construction.

## Acknowledgements

## References

[1] A. ABANDA, U. MORI, AND J. A. LOZANO, *A review on distance based time series classification*, Data Mining and Knowledge Discovery, 33 (2019), pp. 378–412.

[2] A. ABID AND J. Y. ZOU, *Learning a warping distance from unlabeled time series using sequence autoencoders*, in NeurIPS, 2018.

[3] A. BAGNALL, H. A. DAU, J. LINES, M. FLYNN, J. LARGE, A. BOSTROM, P. SOUTHAM, AND E. KEOGH, *The uea multivariate time series classification archive, 2018*, arXiv preprint arXiv:1811.00075, (2018).

[4] A. Bagnall, J. Lines, A. Bostrom, J. Large, and E. Keogh, *The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances*, Data mining and knowledge discovery, 31 (2017), pp. 606–660.

[5] M. G. Baydogan, *Multivariate time series classification datasets*, 2019.

[6] M. Cuturi and M. Blondel, *Soft-dtw: a differentiable loss function for time-series*, in ICML, 2017.

[7] H. A. Dau, A. Bagnall, K. Kamgar, C.-C. M. Yeh, Y. Zhu, S. Gharghabi, C. A. Ratanamahatana, and E. Keogh, *The ucr time series archive*, IEEE/CAA Journal of Automatica Sinica, 6 (2019).

[8] M. Defferrard, X. Bresson, and P. Vandergheynst, *Convolutional neural networks on graphs with fast localized spectral filtering*, in NeurIPS, 2016.

[9] H. I. Fawaz, G. Forestier, J. Weber, L. Idoumghar, and P.-A. Muller, *Deep learning for time series classification: a review*, Data Mining and Knowledge Discovery, 33 (2019), pp. 917–963.

[10] H. I. Fawaz, B. Lucas, G. Forestier, C. Pelletier, D. F. Schmidt, J. Weber, G. I. Webb, L. Idoumghar, P.-A. Muller, and F. Petitjean, *Inceptiontime: Finding alexnet for time series classification*, Data Mining and Knowledge Discovery, (2020).

[11] Y. Geng and X. Luo, *Cost-sensitive convolution based neural networks for imbalanced time-series classification*, arXiv preprint arXiv:1801.04396, (2018).

[12] J. Grabocka and L. Schmidt-Thieme, *Neuralwarp: Time-series similarity with warping networks*, arXiv preprint arXiv:1812.08306, (2018).

[13] W. L. Hamilton, R. Ying, and J. Leskovec, *Inductive representation learning on large graphs*, in NeurIPS, 2017.

[14] S. Jawed, J. Grabocka, and L. Schmidt-Thieme, *Self-supervised learning for semi-supervised time series classification*, in PAKDD, 2020.

[15] F. Karim, S. Majumdar, H. Darabi, and S. Chen, *Lstm fully convolutional networks for time series classification*, IEEE access, 6 (2017), pp. 1662–1669.

[16] F. Karim, S. Majumdar, H. Darabi, and S. Harford, *Multivariate lstm-fcns for time series classification*, Neural Networks, 116 (2019), pp. 237–245.

[17] T. N. Kipf and M. Welling, *Semi-supervised classification with graph convolutional networks*, in ICLR, 2017.

[18] K.-H. Lai, D. Zha, G. Wang, J. Xu, Y. Zhao, D. Kumar, Y. Chen, P. Zumkhawaka, M. Wan, D. Martinez, et al., *Tods: An automated time series outlier detection system*, in AAAI, 2021.

[19] K.-H. Lai, D. Zha, J. Xu, Y. Zhao, G. Wang, and X. Hu, *Revisiting time series outlier detection: Definitions and benchmarks*, in NeurIPS, 2021.

[20] Y. Li, Z. Chen, D. Zha, M. Du, D. Zhang, H. Chen, and X. Hu, *Learning disentangled representations for time series*, arXiv preprint arXiv:2105.08179, (2021).

[21] Y. Li, D. Zha, P. Venugopal, N. Zou, and X. Hu, *Pyodds: An end-to-end outlier detection system with automated machine learning*, in Companion Proceedings of the Web Conference 2020, 2020.

[22] R. Mittelman, *Time-series modeling with undecimated fully convolutional neural networks*, arXiv preprint arXiv:1508.00317, (2015).

[23] H. Ney and S. Ortmanns, *Dynamic programming search for continuous speech recognition*, IEEE Signal Processing Magazine, 16 (1999), pp. 64–83.

[24] M.-I. Nicolae, É. Gaussier, A. Habrard, and M. Sebban, *Similarity learning for time series classification*, arXiv preprint arXiv:1610.04783, (2016).

[25] R. T. Olszewski, *Generalized feature extraction for structural pattern recognition in time-series data*, tech. rep., CARNEGIE-MELLON UNIV PITTSBURGH PA SCHOOL OF COMPUTER SCIENCE, 2001.

[26] T. Rakthanmanon, B. Campana, A. Mueen, G. Batista, B. Westover, Q. Zhu, J. Zakaria, and E. Keogh, *Searching and mining trillions of time series subsequences under dynamic time warping*, in KDD, 2012.

[27] J. Serra and J. L. Arcos, *An empirical evaluation of similarity measures for time series classification*, Knowledge-Based Systems, 67 (2014), pp. 305–314.

[28] N. Strodthoff and C. Strodthoff, *Detecting and interpreting myocardial infarction using fully convolutional neural networks*, Physiological measurement, 40 (2019), p. 015001.

[29] H. Wang, Q. Zhang, J. Wu, S. Pan, and Y. Chen, *Time series feature learning with labeled and unlabeled data*, Pattern Recognition, 89 (2019), pp. 55–66.

[30] Z. Wang, W. Yan, and T. Oates, *Time series classification from scratch with deep neural networks: A strong baseline*, in IJCNN, 2017.

[31] L. Wu, I. E.-H. Yen, J. Yi, F. Xu, Q. Lei, and M. Witbrock, *Random warping series: A random features method for time-series embedding*, in AISTATS, 2018.

[32] Z. Wu, S. Pan, G. Long, J. Jiang, X. Chang, and C. Zhang, *Connecting the dots: Multivariate time series forecasting with graph neural networks*, in KDD, 2020.

[33] X. Zhang, Y. Gao, J. Lin, and C.-T. Lu, *Tapnet: Multivariate time series classification with attentional prototypical network*, in AAAI, 2020.

[34] Y. Zhao, Z. Nasrullah, and Z. Li, *Pyod: A python toolbox for scalable outlier detection*, Journal of Machine Learning Research, 20 (2019), pp. 1–7.

[35] K. Zhou, X. Huang, Y. Li, D. Zha, R. Chen, and X. Hu, *Towards deeper graph neural networks with differentiable group normalization*, in NeurIPS, 2020.

[36] K. Zhou, X. Huang, D. Zha, R. Chen, L. Li, S.-H. Choi, and X. Hu, *Dirichlet energy constrained learning for deep graph neural networks*, in NeurIPS, 2021.

[37] K. Zhou, Q. Song, X. Huang, D. Zha, N. Zou, and X. Hu, *Multi-channel graph neural networks*, in IJCAI, 2020.

Table 7: Statistics of the multivariate time series datasets.

| Dataset \ Attribute | Data Size | Dimension | # Classes | Length |
|---|---|---|---|---|
| Character Trajectories | 2858 | 3 | 20 | 109-205 |
| ECG | 200 | 2 | 2 | 39-152 |
| KickvsPunch | 26 | 62 | 2 | 274-841 |
| NetFlow | 1337 | 2 | 4 | 50-997 |

## A Experimental Details

**A.1 Datasets** We use both univariate datasets and multivariate datasets in our experiments. We provide detailed descriptions below.

**Univariate Time-Series Datasets.** The experiments are conducted on the full 128 datasets from the UCR Time Series Classification Archive[5] [7]. These datasets are collected from various domains, including Image, Spectro, Sensor, Simulated, Device, Motion, ECG, Traffic, EOG, HRM, Trajectory, and Hemodynamics. These datasets also have diverse data sizes, numbers of classes, and lengths. For example, ElectricDevices, one of the largest datasets in the archive, has 16,637 time-series in total, while BeetleFly only has 40 time-series. Similarly, ShapesAll has 60 classes, while many other datasets have only 2 or 3 classes. The time-series length can also be up to 2,709 and can also be as short as 60. Moreover, 11 of the datasets have varying lengths for different time-series. Thus, the UCR datasets provide a rigorous test for time-series classification.

**Multivariate Time-Series Datasets.** We focus on four multivariate datasets from different domains with varying characteristics.

- **Character Trajectories:** is a handwriting dataset captured at 200Hz by a WACOM tablet. The three dimensions are the positions of the x-axis, the y-axis, and the pen tip force. The data has been numerically differentiated, and Gaussian smoothed, with a sigma value of 2.

- **ECG:** traces the electrical activity recorded during heartbeats. The two classes are the normal heartbeat and the myocardial infarction.

- **KickVsPunch:** is a motion capture dataset collected by CMU-MOCAP. Each dimension is a motion marker. The two classes are the two actions, i.e., the kick and the punch.

- **NetFlow:** is the traffic flow of websites.

We summarize the statistics of the above datasets in Table 4.

---

[5]https://www.cs.ucr.edu/~eamonn/time_series_data_2018/

**A.2 Data Preprocessing** For both univariate and multivariate time-series data, we store them as 3-dimensional Numpy arrays, where the first dimension is the number of time-series, the second dimension is the number of time-series (for univariate time-series, there is only one time-series), and the third dimension is the length of the time-series. We impute missing values with zeros for the datasets with varying lengths to make the time-series have the same length. All the time-series are z-normalized before feeding into the models.

**A.3 Data Splitting** The original splits in UCR datasets are very diverse. For example, ElectricDevices has 8,926 training samples, while ECGFiveDays only has 23 training samples. This makes it difficult to understand how the models will behave under different amounts of supervision. Specifically, when an algorithm performs well, it is hard to tell how it performs when we have very few labels and how it will perform when we have enough labels. Motivated by few-shot learning, we simulate the settings where different numbers of labels per class are given. Specifically, we merge the original training and testing data to create new splits as follows. First, we merge the Numpy arrays of training and testing data, where the training data is put before the testing data. Second, we randomly shuffle all the indices. Third, we separate out the last 20% of the shuffled indices as the hold-out testing time-series. Fourth, given the target number of labels per class, we iterate over the first 80% of the shuffled indices sequentially until we find enough number of labeled data for each class. It is possible that we can not find enough labeled data for some classes even after iterating all the indices. In this case, we simply use as many labels as we can. For example, for a class A and a target number 10, if the number of class A data in the first 80% of the shuffled indices is only eight, we simply use eight training data for class A. The remaining data in the first 80% of the shuffled indices will serve as the third split (unlabeled data).

The above three splits are used to simulate the supervised, inductive semi-supervised, and transductive semi-supervised settings defined in Section 2.1. The performance will be evaluated on the hold-out testing set. The above splitting procedure is applied to all the univariate and multivariate time-series data. Since data splitting may significantly affect the performance, particularly when we have very few labels, we run each experiment 5 times on different splits.

**A.4 Dynamic Time Warping** Dynamic Time Warping (DTW) is s standard algorithm for measuring the similarity between two time-series. The main

idea of DTW is to calculate the optimal match between two time-series such that the sum of matched series has the smallest values. In this work, the DTW is computed based on a Python wrapper of The UCR Suite[6]. This suite provides a highly efficient C++ implementation of DTW via dynamic programming. For all the datasets, we set the size of the warping window to be 100. If either of the two time-series is shorter than 100, we use the shortest time-series length.

For multivariate datasets, we compute independent DTW. Specifically, we first compute the DTW for each pair of univariate time-series and then sum them up to represent the distance between two multivariate time-series. We have uploaded the pre-computed DTW to Google Drive for reproducibility.

Since most of the datasets are small, it will not take much time to compute the full similarity matrix. For most of the datasets, the computation of DTW can be finished in minutes. For some larger datasets, it takes at most a day using one CPU core.

**A.5 Neural Architecture of Backbone** We use PyTorch to implement all the neural networks. For most of the experiments, we use ResNet as the backbone. The network consists of three residual blocks. Each residual block consists of three 1-D convolution layers. The kernel sizes of the three convolution layers are 7, 5, and 3, respectively. After each convolution layer, we use a 1-D batch normalization layer to stabilize training, followed by a ReLU activation function. The number of channels is set to be 64 for all the convolution layer. We find that using more channels will lead to unsatisfactory performance with very little training data due to overfitting issue. A skip connection is added in each block to enable direct flow to alleviate the gradient vanishing issue. The three blocks are stacked sequentially to perform feature extraction. To reduce the feature dimension, we add a global average pooling layer to the last residual block's output. For our SimTSC, we directly use the global average pooling layer's output as the extracted features. For the ResNet baseline, these features will be further processed by a fully-connected layer with a softmax activation for classification purposes. The weights of the networks are initialized with the default initializers in PyTorch. Figure 7 summarizes the neural architecture of ResNet.
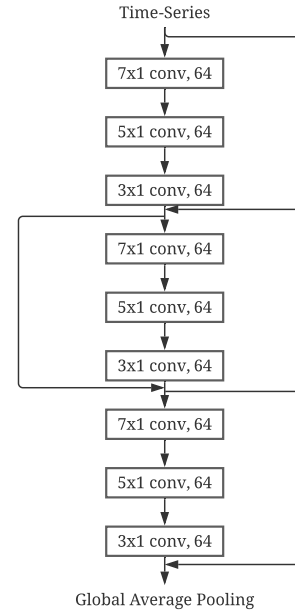


Figure 7: Neural architecture of ResNet. Each layer is followed by a batch normalization layer and a ReLU activation.

For the MLP baseline, we use four fully-connected layers with a hidden size of 500. We use a ReLU activation function after each fully-connected layer followed by a dropout layer to avoid overfitting. The dropout rates are set to be 0.1, 0.2, 0.2, 0.3 for the four layers. We flatten the time-series into one dimension so that it can be directly used in MLP. Different from convolution layers, MLP can not capture temporal information, which leads to unsatisfactory performance. For the FCN baseline, we use three 1-D convolution layers. Each convolution layer is followed by a batch normalization layer. The kernel sizes are set to be 7, 5, and 3. Similar to ResNet, the number of channels is set to 64. For both MLP and PyTorch, we use the default initializers in PyTorch to initialize the weights.

Note that, for all the experiments, we use exactly the same architecture for the backbone of SimTSC and the baseline. The only difference of SimTSC is adding a GCN layer on the top of the backbone. Thus, the comparison is fair. The performance gain is solely attributed to modeling the similarity information.

**A.6 Graph Convolution Layers** The graph convolution layers are implemented based on the standard implementation of GCN[7]. Specifically, each GCN layer takes node features and an adjacency matrix as the input, where the adjacency matrix is a sparse tensor. Then it aggregates the neighbors' features by performing matrix multiplication of the input features, weights, and the adjacency matrix. Finally, a bias term is added

---

[6]https://www.cs.ucr.edu/~eamonn/UCRsuite.html

[7]https://github.com/tkipf/pygcn

to the obtained features. For multiple GCN layers, we add a dropout layer after each GCN layer to avoid over-fitting.

The adjacency matrix used in the graph convolution layers is constructed as follows. We maintain a full pre-computed similarity matrix in the memory. In each update step, we sample a batch of indices for training. We then use the sampled index to obtain a submatrix of the full similarity matrix. This submatrix will only contain the indices in this batch. Further, we rank each row's values in ascending order and only keep the top-$K$ similar neighbors for each row. We finally use the top-$K$ neighbors to construct the graph, which is represented by a sparse matrix. The above constructing procedure is efficient since we only need to take care of a batch of indices instead of all the indices.

**A.7 Hyperparameter Settings** We summarize the hyperparmeters of graph, optimizer, and how we train SimTSC and all the baselines as follows.

- **Graph Construction:** We set the scaling factor $\alpha = 0.3$ and number neighbors for each node $K = 3$.

- **Graph Convolution:** We use one GCN layer for most of the experiments. For multiple GCN layers, the feature dimension is set to be 64, and the dropout rate is set to be 0.5.

- **Optimizer:** We use Adam optimizer. The learning rate is set to be 0.0001. The $\epsilon$ is set to be $10^{-8}$.

- **Training Procedure:** For SimTSC and all the baselines, we use the model that achieves the best performance on the training data for evaluation. Specifically, we calculate the accuracy based on the training data after each training epoch and store the model's weights with the highest accuracy. Then the stored weights will be reloaded for evaluation purposes. For all the models, we train 500 epochs. While validating on the training data may lead to overfitting, we find in practice that it works better than separating a validation set from the training data. This is because a separated validation set will be too small to perform a meaningful evaluation.

**A.8 Hardware and Software Descriptions** We conduct all the experiments on a server with two AMD EPYC 7282 16-Core processors, four GeForce RTX 3090 GPUs, and 252 GB memory. We use Ubuntu 18.04.5 LTS system and PyTorch 1.7.0.