



An Empirical Study of Graph-Based Approaches for Semi-supervised Time Series Classification

*Dominik Bünger, Miriam Gondos, Lucile Peroche and Martin Stoll**

Department of Mathematics, Chair of Scientific Computing, TU Chemnitz, Chemnitz, Germany

OPEN ACCESS

Edited by:

Stefan Kunis,
Osnabrück University, Germany

Reviewed by:

Tarek Emmrich,
Osnabrück University, Germany
Alex Jung,
Aalto University, Finland

*Correspondence:

Martin Stoll
martin.stoll@
mathematik.tu-chemnitz.de

Specialty section:

This article was submitted to
Mathematics of Computation and
Data Science,
a section of the journal
Frontiers in Applied Mathematics and
Statistics

Received: 28 September 2021

Accepted: 28 December 2021

Published: 20 January 2022

Citation:

Bünger D, Gondos M, Peroche L and
Stoll M (2022) An Empirical Study of
Graph-Based Approaches for
Semi-supervised Time Series
Classification.
Front. Appl. Math. Stat. 7:784855.
doi: 10.3389/fams.2021.784855

Time series data play an important role in many applications and their analysis reveals crucial information for understanding the underlying processes. Among the many time series learning tasks of great importance, we here focus on semi-supervised learning based on a graph representation of the data. Two main aspects are studied in this paper. Namely, suitable distance measures to evaluate the similarities between different time series, and the choice of learning method to make predictions based on a given number of pre-labeled data points. However, the relationship between the two aspects has never been studied systematically in the context of graph-based learning. We describe four different distance measures, including (Soft) DTW and MPDist, a distance measure based on the Matrix Profile, as well as four successful semi-supervised learning methods, including the recently introduced graph Allen–Cahn method and Graph Convolutional Neural Network method. We provide results for the novel combination of these distance measures with both the Allen–Cahn method and the GCN algorithm for binary semi-supervised learning tasks for various time-series data sets. In our findings we compare the chosen graph-based methods using all distance measures and observe that the results vary strongly with respect to the accuracy. We then observe that no clear best combination to employ in all cases is found. Our study provides a reproducible framework for future work in the direction of semi-supervised learning for time series with a focus on graph representations.

Keywords: semi-supervised learning, time series, graph Laplacian, Allen–Cahn equation, graph convolutional networks

1. INTRODUCTION

Many processes for which data are collected are time-dependent and as a result the study of time series data is a subject of great importance [1–3]. The case of time series is interesting for tasks such as anomaly detection [4], motif computation [5] or time series forecasting [6]. We refer to [7–10] for more general introductions.

We here focus on the task of classification of time series [11–16] in the context of semi-supervised learning [17, 18] where we want to label all data points¹ based on the fact that only a small portion of the data is already pre-labeled.

An example is given in **Figure 1** where we see some time series reflecting ECG (electrocardiogram) data and the classification into normal heartbeats on the one hand and myocardial infarction on the other hand. In our applications, we assume that only for some of the time series the corresponding class is known a priori. Our main contribution is to introduce a novel combination of incorporating the data into a *graph* and then incorporate this representation into several recently introduced methods for *semi-supervised learning*. For this, each time series becomes a node within a weighted undirected graph and the edge-weight is proportional to the similarity between different time series. Graph-based approaches have become a standard tool in many learning tasks (cf. [19–24] and the references mentioned therein). The matrix representation of the graph via its Laplacian [25] leads to studying the network using matrix properties. The Laplacian is *the* representation of the network that is utilized from machine learning to mathematical imaging. Recently, it has also been used network-Lasso-based learning approaches focusing on data with an inherent network structure, see e.g., [26, 27]. A very important ingredient in the construction of the Laplacian is the choice of the appropriate weight function. In many applications, the computation of the distance between time series or subsequences becomes a crucial task and this will be reflected in our choice of weight function. We consider several distance measures such as dynamic time warping DTW [28], soft DTW [29], and matrix profile [30].

We will embed these measures via the graph Laplacian into two different recently proposed semi-supervised learning frameworks. Namely, a diffuse interface approach that originates from material science [31] via the graph Allen-Cahn equation as well as a method based on graph convolutional networks [21]. Since these methods have originally been introduced outside of the field of time series learning, their relationship with time series distance measures has never been studied. Our goal is furthermore to compare these approaches with the well-known 1NN approach [11] and a simple optimization formulation solved relying on a linear system of equations. Our motivation follows that of [32, 33], where many methods for supervised learning in the context of time series were compared, namely that we aim to provide a wide-ranging overview of recent methods based on a graph representation of the data and combined with several distance measures.

We structure the paper as follows. In section 2, we introduce some basic notations and illustrate the basic notion of graph-based learning motivated with a clustering approach. In section 3, we discuss several distance measures with a focus on the well-known DTW measure as well as two recently emerged alternatives, i.e., Soft DTW and the MP distance. We use section 4 to introduce the two semi-supervised learning methods in more detail, followed by a shorter description of their well-known

competitors. section 5 will allow us to compare the methods and study the hyperparameter selection.

2. BASICS

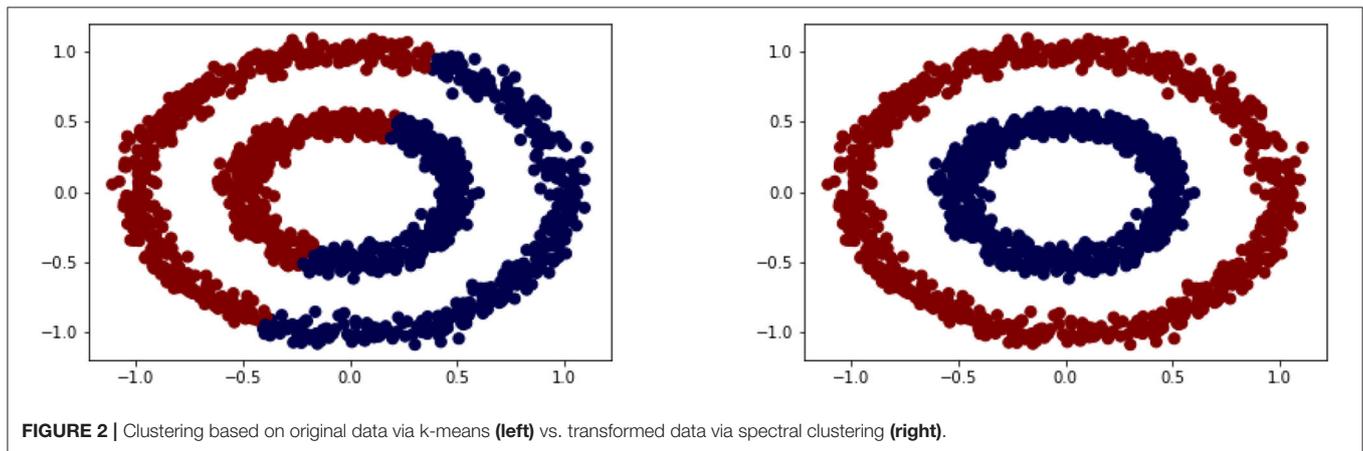
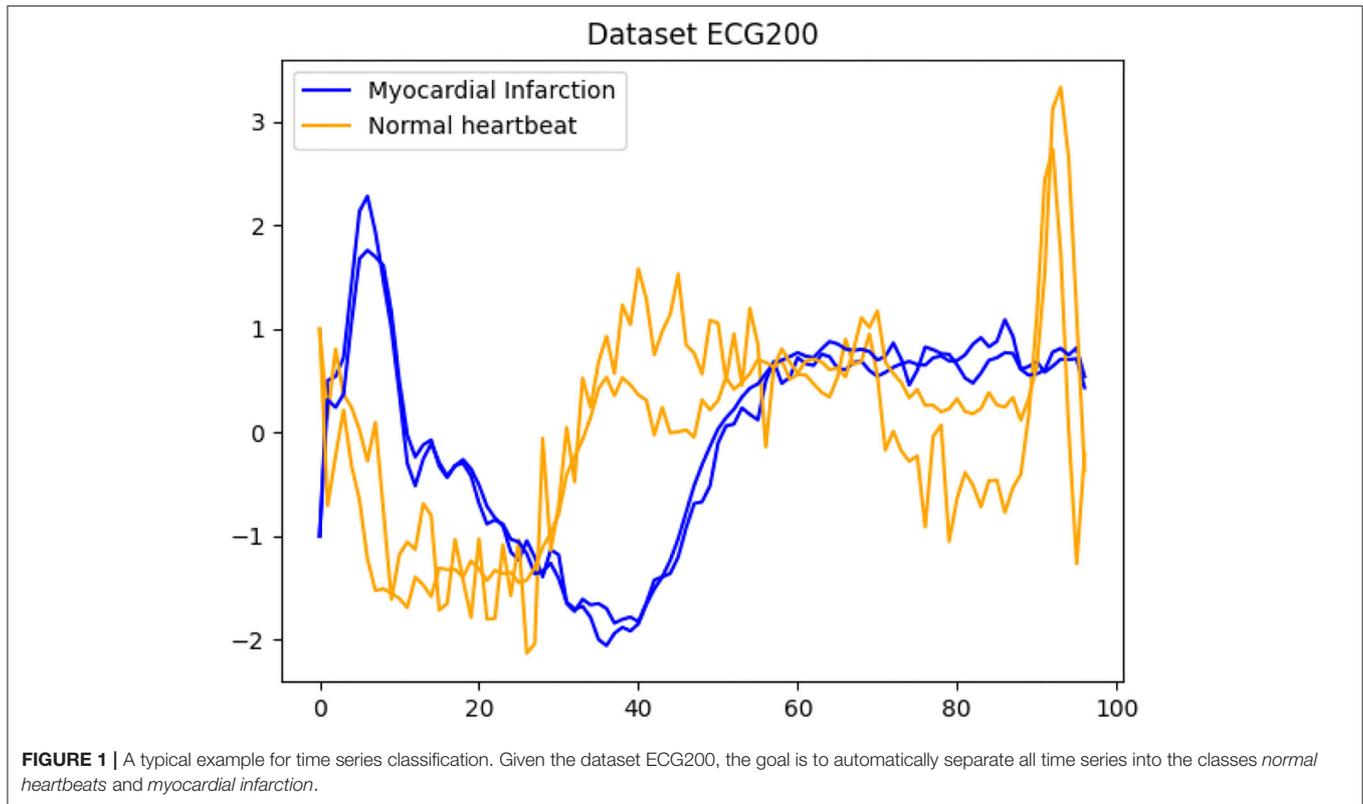
We consider discrete time series \mathbf{x}_i given as a vector of real numbers of length m_i . In general, we allow for the time series to be of different dimensionality; later we often consider all $m_i = m$. We assume that we are given n time series $\mathbf{x}_i \in \mathbb{R}^{m_i}$. The goal of a classification task is to group the n time series into a number k of different *clusters* C_j with $j = 1, \dots, k$. In this paper we focus on the task of semi-supervised learning [17] where only some of the data are already labeled but we want to classify all available data simultaneously. Nevertheless, we review some techniques for unsupervised learning first as they deliver useful terminology. As such the *k-means* algorithm is a prototype-based² clustering algorithm that divides the given data into a predefined number of k clusters [34]. The idea behind *k-means* is rather simple as the cluster centroids are repeatedly updated and the data points are assigned to the nearest centroid until the centroids and data points have converged. Often the termination condition is not handled that strictly. For example, the method can be terminated when only 1% of the points change clusters. The starting classes are often chosen at random but can also be assigned in a more systematic way by calculating the centers first and then assign the points to the nearest center. While *k-means* remains very popular it also has certain weaknesses coming from its minimization of the sum of squared errors loss function [35]. We discuss this method in some detail here to point out the main mechanism and this is based on assigning points to clusters and hence the cluster centroids based on the distance being the Euclidean norm, which would also be done when *k-means* is applied to time series. As a result the clusters might not capture the shape of the data manifold as illustrated in a simple two-dimensional example shown in **Figure 2**. In comparison, the alternative method shown, i.e., a spectral clustering technique, performs much better. We briefly discuss this method next as it forms the basis of the main techniques introduced in this paper.

2.1. Graph Laplacian and Spectral Clustering

As we illustrated in **Figure 2** the separation of the data into two-classes is rather difficult for *k-means* as the centroids are based on a 2-norm minimization. One alternative to *k-means* is based on interpreting the data points as nodes in a graph. For this, we assume that we are given data points x_1, \dots, x_n and some measure of similarity [23]. We define the weighted undirected similarity graph $G = (V, E)$ with the *vertex* or *node* set V and the edge set E . We view the data points \mathbf{x}_i as vertices, $V = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, and if two nodes $(\mathbf{x}_i, \mathbf{x}_j)$ have a positive similarity function value, they are connected by an edge with weight w_{ij} equal to that similarity. With this reformulation of the data we turn the clustering problem into a graph partitioning problem where we want to cut the graph into two or possibly more classes.

¹We here view one time-series as a data point and the feature vector for this data point is the vector with the associated data collected in a vector.

²Here the prototype of the cluster is the centroid.



This is usually done in such a way that the weight of the edges across the partition is minimal.

We collect all edge weights in the *adjacency matrix* $W = (w_{ij})_{i,j=1,\dots,n}$. The degree of a vertex x_i is defined as $d_i = \sum_{j=1}^n w_{ij}$ and the degree matrix D is the diagonal matrix holding all n node degrees. In our case we use a fully connected graph with the *Gaussian similarity function*

$$w(x_i, x_j) = \exp\left(-\frac{\text{dist}(x_i, x_j)^2}{\sigma^2}\right), \quad (1)$$

where σ is a scaling parameter and $\text{dist}(x_i, x_j)$ is a particular distance function such as the Euclidean distance $\text{dist}(x_i, x_j) :=$

$\|x_i - x_j\|^2$. Note that for similar nodes, the value of the *distance* function is smaller than it would be for dissimilar nodes while the *similarity* function is relatively large.

We now use both the degree and weight matrix to define the *graph Laplacian* as $L = D - W$. Often the *symmetrically normalized Laplacian* defined via

$$L_{\text{sym}} = D^{-\frac{1}{2}} L D^{-\frac{1}{2}} = I - D^{-\frac{1}{2}} W D^{-\frac{1}{2}} \quad (2)$$

provides better clustering information [23]. It has some very useful properties that we will exploit here. For example, given a

non-zero vector $u \in \mathbb{R}^n$ we obtain the energy term

$$u^T L_{\text{sym}} u = \frac{1}{2} \sum_{ij} w_{ij} \left(\frac{u_i}{\sqrt{d_i}} - \frac{u_j}{\sqrt{d_j}} \right)^2. \tag{3}$$

Using this it is easy to see that L_{sym} is positive semi-definite with non-negative eigenvalues $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$. The main advantage of the graph Laplacian is that based on its spectral information one can usually rely on transforming the data into a space where they are easier to separate [23, 25, 36]. As a result one typically requires the spectral information corresponding to the smallest eigenvalues of L_{sym} . The most famed eigenvector is the *Fiedler vector*, i.e., the eigenvector corresponding to the first non-zero eigenvalue, which is bound to have a sign change and as a result can be used for binary classification. The weight function (1) is also found in kernel methods [37, 38] when the radial basis kernel is applied.

2.2. Self-Tuning

In order to improve the performance of the methods based on the graph Laplacian, tuning the parameter σ is crucial. While hyperparameter tuning based on a grid search or cross validation is certainly possible we also consider a σ that adapts to the given data. For spectral clustering, such a procedure was introduced in [39]. Here we use this technique to learning with time series data. For each time series \mathbf{x}_i we assume a local scaling parameter σ_i . As a result, we have the generalized square distance as

$$\frac{\text{dist}(\mathbf{x}_i, \mathbf{x}_j)}{\sigma_i} \frac{\text{dist}(\mathbf{x}_i, \mathbf{x}_j)}{\sigma_j} = \frac{\text{dist}(\mathbf{x}_i, \mathbf{x}_j)^2}{\sigma_i \sigma_j} \tag{4}$$

and this gives the entries of the adjacency matrix W via

$$w_{i,j} = \exp \left(- \frac{\text{dist}(\mathbf{x}_i, \mathbf{x}_j)^2}{\sigma_i \sigma_j} \right). \tag{5}$$

The authors in [39] choose σ_i as the distance to the K -th nearest neighbor of \mathbf{x}_i where K is a fixed parameter, e.g., $K = 9$ is used in [31].

In section 5, we will explore several different values for K and their influence on the classification behavior.

3. DISTANCE MEASURES

We have seen from the definition of the weight matrix that the Laplacian depends on the choice of distance measure $\text{dist}(\mathbf{x}_i, \mathbf{x}_j)$. If all time series are of the same length then the easiest distance measure would be a Euclidean distance, which especially for large n is fast to compute. This makes the Euclidean distance incredibly popular but it suffers from being sensitive to small shifts in the time series. As a result we discuss several popular and efficient methods for different distance measures. Our focus is to illustrate in an empirical study how the choice of distance measure impacts the performance of graph-based learning and to provide further insights for future research (cf. [40]).

3.1. Dynamic Time Warping

We first discuss the distance measure of Dynamic Time Warping (DTW, [28]). By construction, DTW is an algorithm to find an optimal alignment between time series.

In the following, we adapt the notation of [28] to our case. Consider two time series \mathbf{x} and $\tilde{\mathbf{x}}$ of lengths m and \tilde{m} , respectively, with entries $x_i, \tilde{x}_i \in \mathbb{R}$ for $i = 1, \dots, m$ and $j = 1, \dots, \tilde{m}$. We obtain the local cost matrix $C \in \mathbb{R}^{m \times \tilde{m}}$ by assembling the local differences for each pair of elements, i.e., $C_{ij} = |x_i - \tilde{x}_j|$.

The DTW distance is defined via (m, \tilde{m}) -warping paths, which are sequences of index tuples $p = ((i_1, j_1), \dots, (i_L, j_L))$ with boundary, monotonicity, and step size conditions

$$1 = i_1 \leq i_2 \leq \dots \leq i_L = m, \quad 1 = j_1 \leq j_2 \leq \dots \leq \tilde{m}, \\ (i_{\ell+1} - i_\ell, j_{\ell+1} - j_\ell) \in \{(1, 0), (0, 1), (1, 1)\} \quad (\ell = 1, \dots, L - 1).$$

The total cost of such a path with respect to $\mathbf{x}, \tilde{\mathbf{x}}$ is defined as

$$c_p(\mathbf{x}, \tilde{\mathbf{x}}) = \sum_{\ell=1}^L |x_{i_\ell} - \tilde{x}_{j_\ell}|.$$

The DTW distance is then defined as the minimum cost of any warping path:

$$\text{DTW}(\mathbf{x}, \mathbf{y}) := \min\{c_p(\mathbf{x}, \mathbf{y}) \mid p \text{ is a } (m, \tilde{m})\text{-warping path}\}. \tag{6}$$

Both the warping and the warping path are illustrated in **Figure 3**.

Computing the optimal warping path directly quickly becomes infeasible. However, we can use dynamic programming to evaluate the accumulated cost matrix D recursively via

$$D(i, j) := |x_i - \tilde{x}_j| + \min\{D(i, j-1), D(i-1, j), D(i-1, j-1)\}. \tag{7}$$

The actual DTW distance is finally obtained as

$$\text{DTW}(\mathbf{x}, \mathbf{y}) = D(m, \tilde{m}). \tag{8}$$

The DTW method is a heavily used distance measure for capturing the sometimes subtle similarities between time series. In the literature it is typically stated that the computational cost of DTW being prohibitively large. As a result one is interested in accelerating the DTW algorithm itself. One possibility arises from imposing additional constraints (cf. [28, 41]) such as the Sakoe-Chiba Band and the Itakura parallelogram as these simplify the identification of the optimal warping path. While these are appealing concepts the authors in [42] observe that the well-known FastDTW algorithm [41] is in fact slower than DTW. For our purpose we will hence rely on DTW and in particular on the implementation of DTW provided via <https://github.com/wannesm/dtaidistance>. We observe that for this implementation of DTW indeed FastDTW is outperformed frequently.

3.2. Soft Dynamic Time Warping

Based on a slight reformulation of the above DTW scheme, we want to look at another time series distance measure, the *Soft Dynamic Time Warping* (Soft DTW). It is an extension of DTW designed allowing a differentiable loss function and it was

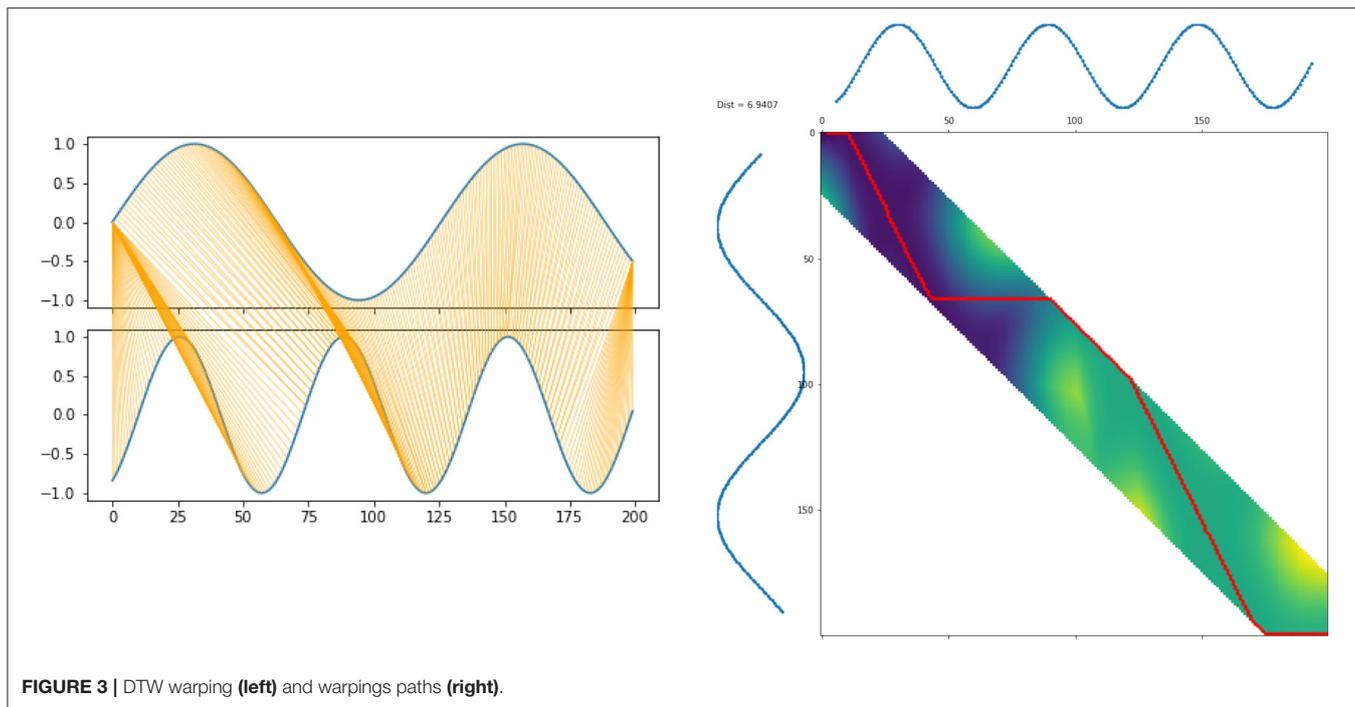


FIGURE 3 | DTW warping (left) and warping paths (right).

introduced in [29, 43]. We again start from the cost matrix C with $C(i, j) = |x_i - \tilde{x}_j|$ for time series \mathbf{x} and $\tilde{\mathbf{x}}$. Each warping path can equivalently be described by a matrix $A \in \{0, 1\}^{m \times \tilde{m}}$ with the following condition: The ones in A form a path starting in $(1, 1)$ going to (m, \tilde{m}) , only using steps downwards, to the right and diagonal downwards. A is called monotonic alignment matrix and we denote the set containing all these alignment matrices with $\mathcal{A}(m, \tilde{m})$. The Frobenius inner product $\langle A, C \rangle$ is then the sum of costs along the alignment A . Solving the following minimization problem leads us to a reformulation of the dynamic time warping introduced above as

$$DTW(C) = \min_{A \in \mathcal{A}(N, M)} \langle A, C \rangle. \tag{9}$$

With Soft DTW we involve all alignments possible in $\mathcal{A}(N, M)$ by replacing the minimization with a *soft minimum*:

$$\min_{x \in S} f(x) \approx \min_{\gamma} f(x) := -\gamma \log \sum_{x \in S} \exp\left(\frac{-f(x)}{\gamma}\right) \tag{10}$$

where S is a discrete subset of the real numbers. This function approximates the minimum of $f(x)$ and is differentiable. The parameter γ controls the tuning between smoothness and approximation of the minimum. Using the DTW-function (9) within (10) yields the expression for Soft Dynamic Time Warping written as

$$\begin{aligned} DTW_{\gamma}(\mathbf{x}, \tilde{\mathbf{x}}) &= \min_{A \in \mathcal{A}(m, n)} \langle A, C \rangle \\ &= -\gamma \log \sum_{A \in \mathcal{A}(m, n)} \exp\left(\frac{-\langle A, C \rangle}{\gamma}\right). \end{aligned} \tag{11}$$

This is now a differentiable alternative to DTW, which involves all alignments in our cost matrix.

Due to entropic bias³, Soft DTW can generate negative values, which would cause issues for our use in time series classification. We apply the following remedy to overcome this drawback:

$$Div(\mathbf{x}, \mathbf{y}) = DTW_{\gamma}(\mathbf{x}, \mathbf{y}) - \frac{1}{2} \cdot (DTW_{\gamma}(\mathbf{x}, \mathbf{x}) + DTW_{\gamma}(\mathbf{y}, \mathbf{y})). \tag{12}$$

This measure is called Soft DTW divergence [43] and will be employed in our experiments.

3.3. Matrix Profile Distance

Another alternative time series measure that has recently been introduced is the *Matrix Profile Distance* (MP distance, [30]). This measure is designed for fast computation and finding similarities between time series.

We will again introduce the concept of the matrix profile of two time series \mathbf{x} and $\tilde{\mathbf{x}}$. The matrix profile is based on the subsequences of these two time series. For a fixed window length L , the subsequence $\mathbf{x}_{i,L}$ of a time series \mathbf{x} is defined as a contiguous L -element subset of \mathbf{x} via $\mathbf{x}_{i,L} = (x_i, x_{i+1}, \dots, x_{i+L-1})$. The *all-subsequences set* A of \mathbf{x} contains all possible subsequences of \mathbf{x} with length L , $A = \{\mathbf{x}_{1,L}, \mathbf{x}_{2,L}, \dots, \mathbf{x}_{m-L+1,L}\}$, where m is again the length of \mathbf{x} .

For the matrix profile, we need the all-subsequences sets A and B of both time series \mathbf{x} and $\tilde{\mathbf{x}}$. The matrix profile \mathbf{P}_{ABBA} is the set consisting of the closest Euclidean distances from each

³This term is commonly used when the regression results shrink toward a mass at the barycenter of a target [44].

subsequence in A to any subsequence in B and vice versa:

$$P_{ABBA} = \left\{ \min_{\tilde{x}_{j,L} \in B} \|x_{i,L} - \tilde{x}_{j,L}\| \mid x_{i,L} \in A \right\} \cup \left\{ \min_{x_{i,L} \in A} \|\tilde{x}_{j,L} - x_{i,L}\| \mid \tilde{x}_{j,L} \in B \right\}$$

With the matrix profile, we can finally define the MP distance based on the idea that two time series are similar if they have many similar subsequences. We do not consider the smallest or the largest value of P_{ABBA} because then the MP distance could be too rough or too detailed. For example, if we would have two rather similar time series, but either one has a noisy spike or some missing values, then the largest value of the matrix profile could give a wrong impression about the similarity of these two time series. Instead, the distance is defined as

$$MPdist(X, Y) = k\text{-th smallest value in sorted } P_{ABBA},$$

where the parameter k is typically set to 5% of $2N$ [30].

We now illustrate the MP distance using an example as illustrated in section 3.3, where we display three time series of length $N = 100$. Our goal is to compare these time series using the MP distance. We observe that X_1 and X_2 have quite similar oscillations. The third time series X_3 does not share any obvious features with the first two sequences.

The MP distance compares the subsequences of the time series, depending on the window length L . Choosing the window length to be $L = 40$, we get the following distances:

$$\begin{aligned} MPdist(X_1, X_2) &= 0.433, \\ MPdist(X_1, X_3) &= 5.425, \\ MPdist(X_2, X_3) &= 5.404. \end{aligned}$$

As we can see, the MP distance identified the similarity between X_1 and X_2 shows that X_1, X_2 differ from X_3 . We also want to show that the MP Distance depends on the window length L . Let us look at the MP distance between the lower oscillation time series X_2 and X_3 , which is varying a lot for different values of L as indicated in **Table 1**. Choosing $L = 10$ there is not a large portion of both time series to compare with and as a result we observe a small value for the MP distance, which does not describe the dissimilarity of X_2 and X_3 in a proper way. If we look at $L = 40$, there is a larger part of the time series structure to compare the two series. If there is a special recurring pattern in the time series, the length L should be large enough to cover one recurrence. We illustrate the comparison based on different window lengths in **Figure 4**.

For the tests all data sets consist of time series with a certain length, varying for each data set. Thus we have to decide which

TABLE 1 | MP distance depending on the window length.

L	10	20	30	40
MPdist(X_2, X_3)	0.270	2,034	3,955	5,404

window length L should be chosen automatically in the classifier. An empirical study showed that choosing $L \approx N/2$ gives good classification results.

We briefly illustrate the computing times of the different distance measures when applied to time series of increasing length shown in **Figure 5**. It can be seen that DTW is faster than fastDTW. Obviously, the Euclidean distance shows the best scalability. We also observe that the computation of the SDTW is scaling worse than the competing approaches when applied to longer time series.

4. SEMI-SUPERVISED LEARNING BASED ON GRAPH LAPLACIANS

In this section, we focus mainly on two methods that have recently gained wide attention. This first method is inspired by a partial differential equation model originating from material science and the second approach is based on neural networks that incorporate the graph structure of the labeled and unlabeled data.

4.1. Semi-supervised Learning With Phase Field Methods: Allen–Cahn Model

Within the material science community phase field methods have been developed to model the phase separation of a multicomponent alloy system (cf. [45, 46]). The evolution of the phases over time is described by a partial differential equation (PDE) model, such as the Allen-Cahn [46] or Cahn-Hilliard equation [47] both non-linear reaction-diffusion equations of second and fourth order, respectively. These equations can be obtained as gradient flows of the Ginzburg–Landau energy functional

$$\mathcal{E}(u) = \int \frac{\varepsilon}{2} |\nabla u|^2 + \frac{1}{\varepsilon} \phi(u)$$

where u is the order parameter and ε a parameter reflecting the width of the interface between the pure phases. The polynomial ϕ is chosen to have minima at the pure phases, namely $u = -1$ and $u = 1$, to enforce that a minimization of the Ginzburg–Landau energy will lead to phase separation. A common choice is the well-known double-well potential $\phi(u) = \frac{1}{4}(1 - u^2)^2$. The Dirichlet energy term $|\nabla u|^2$ corresponds to minimization of the interfacial length. The minimization is then performed using a gradient flow, which leads to the Allen-Cahn equation

$$u_t = \Delta u - \frac{1}{\varepsilon} \phi'(u) \tag{13}$$

equipped with appropriate boundary and initial conditions. A modified Allen–Cahn equation was used for image inpainting, i.e., restoring damage parts in an image, where a misfit $\omega(f - u)$ term is added to Equation (13) (cf. [48, 49]). Here, ω is a penalty parameter and f is a function equal to the undamaged image parts or later training data. In [31], Bertozzi and Flenner extended this idea to the case of semi-supervised learning where the training data correspond to the undamaged image parts, i.e., the function

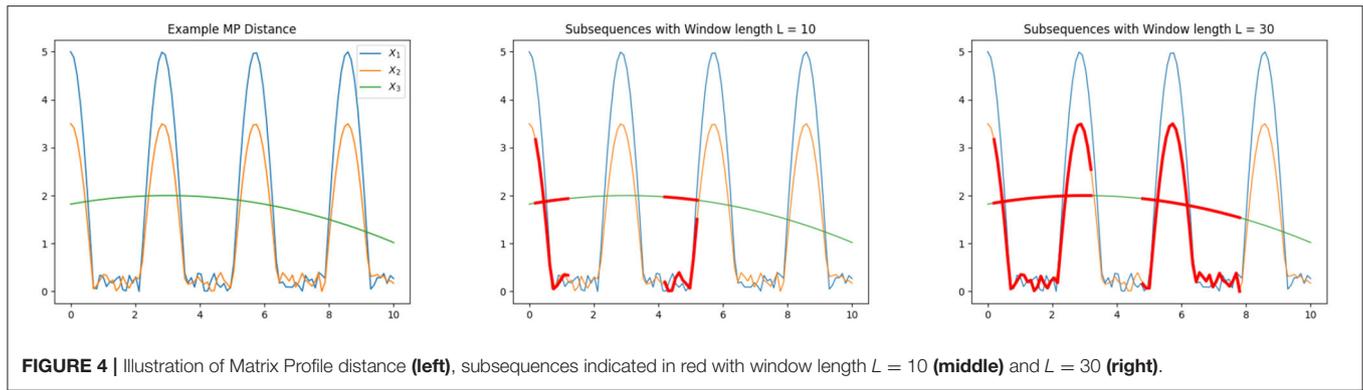


FIGURE 4 | Illustration of Matrix Profile distance (left), subsequences indicated in red with window length $L = 10$ (middle) and $L = 30$ (right).

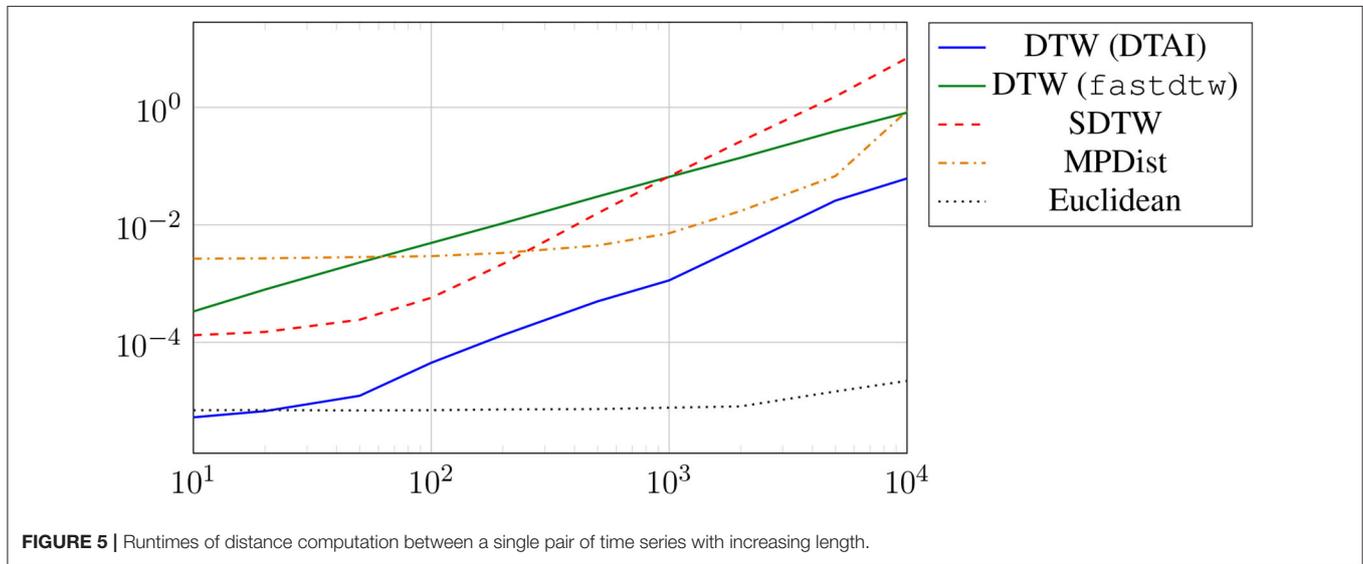


FIGURE 5 | Runtimes of distance computation between a single pair of time series with increasing length.

TABLE 2 | Default parameters used in the experiments.

Method	Parameters and default values
Allen-Cahn	$m_e = 20, \varepsilon = \frac{1}{\sqrt{n}}, c = \frac{3}{\varepsilon} + \omega, \omega = 1e10, \tau = 0.01, tol = 1e-8$
GCN	10-NN sparsification, $h = 32$, dropout $p = 0.5$, Adam optimization [62], learning rate 0.01, weight decay 0.0005, 500 epochs
Linear System	$\beta = 1, tol = 1e-5$
1NN	—

f. Their idea is to consider the modified energy of the following form

$$E(u) = \frac{\varepsilon}{2} u^T L_{\text{sym}} u + \frac{1}{4\varepsilon} \sum_{i=1}^n (u_i^2 - 1)^2 + \sum_{i=1}^n \frac{\omega_i}{2} (f_i - u_i) \quad (14)$$

where f_i holds the already assigned labels. Here, the first term in (14) reflects the RatioCut based on the graph Laplacian, the second term enforces the pure phases, and the third term corresponds to incorporating the training data. Numerically, this system is solved using a convexity splitting approach [31] where

we write

$$E(u) = E_1(u) - E_2(u)$$

with

$$E_1(u) := \frac{\varepsilon}{2} u^T L_{\text{sym}} u + \frac{c}{2} u^T u$$

and

$$E_2(u) := \frac{c}{2} u^T u - \frac{1}{4\varepsilon} + \sum_{i=1}^n (u_i^2 - 1)^2 - \sum_{i=1}^n \frac{\omega_i}{2} (f_i - u_i)$$

where the positive parameter $c \in \mathbb{R}$ ensures convexity of both energies. In order to compute the minimizer of the above energy we use a gradient scheme where

$$\frac{u^{l+1} - u^l}{\tau} = -\nabla E_1(u^{l+1}) + \nabla E_2(u^l)$$

where the indices $k, k + 1$ indicate the current and next time step, respectively. The variable τ is a hyperparameter but can be

TABLE 3 | Study of self-tuning parameters.

		$k = 7$ (%)	$k = 20$ (%)	$k = \sqrt{n}$ (%)	$k = 0.1n$ (%)	$k = 0.05n$ (%)
ECG200 (n = 200)						
MPDist	GCN	83.58	81.74	81.90	81.74	82.54
	Allen-Cahn	81.00	79.00	80.00	79.00	80.00
SDTW	GCN	91.95	91.34	90.70	91.43	90.55
	Allen-Cahn	92.00	90.00	91.00	90.00	91.00
DTW	GCN	88.92	86.76	87.43	86.76	88.97
	Allen-Cahn	82.00	82.00	83.00	82.00	82.00
SonyAIBORobotSurface1 (n = 621)						
MPDist	GCN	95.45	88.74	93.08	78.10	89.62
	Allen-Cahn	75.54	72.88	73.04	75.37	73.71
SDTW	GCN	90.32	91.46	92.48	87.34	92.85
	Allen-Cahn	93.68	85.19	82.36	81.36	82.36
DTW	GCN	97.59	97.58	97.48	96.49	97.35
	Allen-Cahn	84.03	86.85	87.69	87.19	88.19
ECGFiveDays (n = 884)						
MPDist	GCN	99.70	99.77	99.51	99.66	99.15
	Allen-Cahn	89.89	90.71	95.35	95.82	96.40
SDTW	GCN	97.30	97.11	97.31	96.49	97.06
	Allen-Cahn	82.00	86.99	85.48	86.76	87.57
DTW	GCN	97.22	97.19	97.39	97.20	97.35
	Allen-Cahn	77.35	76.31	75.72	73.17	74.68
TwoLeadECG (n = 1,162)						
MPDist	GCN	99.81	99.78	99.81	99.62	99.74
	Allen-Cahn	99.12	97.10	96.49	97.72	96.57
SDTW	GCN	92.10	90.74	90.53	89.98	90.72
	Allen-Cahn	97.19	93.24	91.04	87.27	87.71
DTW	GCN	92.94	94.04	94.98	93.97	96.49
	Allen-Cahn	93.85	92.36	92.10	94.12	93.50

Bold values indicate most accurate classification.

interpreted as a pseudo time-step. In more detail following the notation of [20], this leads to

$$\frac{u^{l+1} - u^l}{\tau} + \varepsilon L_{\text{sym}} u^{l+1} + cu^{l+1} = cu^l - \frac{1}{\varepsilon} \nabla \psi(u^l) + \nabla \phi(u^l)$$

with

$$\psi(u^l) = \sum_{i=1}^n ((u_i^l)^2 - 1)^2, \quad \phi(u^l) = \sum_{i=1}^n \frac{\omega_i}{2} (f_i - u_i^l).$$

Expanding the order parameter in a number of the small eigenvectors ϕ_i of L_{sym} via $u = \sum_{i=1}^{m_e} a_i \phi_i = \Phi_{m_e} a$ where a is a coefficient vector and $\Phi_{m_e} = [\phi_1, \dots, \phi_{m_e}]$. This lets us arrive at

$$(1 + \varepsilon \tau \lambda_j a_j^{l+1} + c \tau) a_j^{l+1} = (1 + \tau c) a_j^l - \frac{1}{\varepsilon} b_j^l + d_j^l, \quad \forall j = 1, \dots, m_e$$

using

$$b^l = \Phi_{m_e}^T \nabla \psi(\Phi_{m_e} a^l), \quad d^l = \Phi_{m_e}^T \nabla \phi(\Phi_{m_e} a^l).$$

In [50], the authors extend this to the case of multiple classes where again the spectral information of the graph Laplacian are crucial as the energy term includes $\frac{\varepsilon}{2} \text{tr}(U^T L_{\text{sym}} U)$ with $U \in \mathbb{R}^{n \times s}$, s being the number of classes for segmentation, and tr being the trace of the matrix. Details of the definition of the potential and the fidelity term incorporating the training data are found in [50]. Further extensions of this approach have been suggested in [20, 22, 51–55].

4.2. Semi-supervised Learning Based on Graph Convolutional Networks

Artificial neural networks and in particular deep neural networks have shown outstanding performance in many learning tasks [56, 57]. The incorporation of additional structural information via a graph structure has received wide attention [24] with particular success within the semi-supervised learning formulation [21].

Let $\mathbf{h}_i^{(l)}$ denote the hidden feature vector of the i -th node in the l -th layer. The feature mapping of a simple multilayer perceptron (MLP) computes the new features by multiplying with a weight matrix $\Theta^{(l)T}$ and adding a bias vector $\mathbf{b}^{(l)}$, then applying a (potentially layer-dependent) ReLU activation function σ_l in all

TABLE 4 | Varying the number of eigenpairs for the reduced Allen–Cahn equation.

Number of eigenvalues	10 (%)	20 (%)	30 (%)	150 (%)	190 (%)
Dataset ECG200					
MPDist	82.00	81.00	86.00	62.00	56.00
SDTW	78.00	92.00	92.00	68.00	66.00
DTW	78.00	82.00	87.00	69.00	54.00
Number of eigenvalues	10 (%)	20 (%)	30 (%)	500 (%)	600 (%)
SonyAIBORobotSurface1					
MPDist	85.36	75.54	73.04	51.58	51.08
SDTW	96.17	93.68	83.19	52.08	49.92
DTW	90.01	84.03	72.71	52.41	48.58
Number of eigenvalues	10 (%)	20 (%)	30 (%)	700 (%)	800 (%)
ECGFiveDays					
MPDist	87.19	89.89	85.95	50.29	51.22
SDTW	91.52	82.00	84.20	54.00	52.38
DTW	68.87	77.35	77.00	49.82	50.29

Bold values indicate most accurate classification.

layers except the last. This layer operation can be written as $\mathbf{h}_i^l = \sigma_i \left(\Theta^{(l)T} \mathbf{h}_i^{(l-1)} + \mathbf{b}^{(l)} \right)$.

In Graph Neural Networks, the features are additionally propagated along the edges of the graph. This is achieved by forming weighted sums over the local neighborhood of each node, leading to

$$\mathbf{h}_i^l = \sigma_i \left(\sum_{j \in \mathcal{N}_i \cup \{i\}} \frac{\hat{w}_{ij}}{\sqrt{\hat{d}_i \hat{d}_j}} \Theta^{(l)T} \mathbf{h}_j^{(l-1)} + \mathbf{b}^{(l)} \right). \quad (15)$$

Here, \mathcal{N}_i denotes the set of neighbors of node i , $\Theta^{(l)}$ and $\mathbf{b}^{(l)}$ the trainable parameters of layer l , the \hat{w}_{ij} denote the entries of the adjacency matrix W with added self loops, $\hat{W} = W + I$, and the \hat{d}_i denote the row sums of that matrix. By adding the self loops, it is ensured that the original features of that node are maintained in the weighted sum.

To obtain a matrix formulation, we can accumulate state matrices $X^{(l)}$ whose n rows are the feature vectors $\mathbf{h}_i^{(l)T}$ for $i = 1, \dots, n$. The propagation scheme of a simple two-layer graph convolutional network can then be written as

$$\begin{aligned} X^{(1)} &= \sigma \left(\hat{D}^{-1/2} \hat{W} \hat{D}^{-1/2} X^{(0)} \Theta^{(1)} + \mathbf{b}^{(1)} \right) \\ X^{(2)} &= \hat{D}^{-1/2} \hat{W} \hat{D}^{-1/2} X^{(1)} \Theta^{(2)} + \mathbf{b}^{(2)} \end{aligned} \quad (16)$$

where \hat{D} is the diagonal matrix holding the \hat{d}_i .

Multiplication with $\hat{D}^{-1/2} \hat{W} \hat{D}^{-1/2}$ can also be understood in a spectral sense as performing *graph convolution* with the spectral filter function $\varphi(\lambda) = 1 - \lambda$. This filter originates from truncating a Chebyshev polynomial to first order as discussed in [58]. As a result of this filter the eigenvalues λ of the graph Laplacian operator \mathcal{L} (formed in this case *after* adding the self

loops) are transformed via φ to obtain damping coefficients for the corresponding eigenvectors. This filter has been shown to lead to convolutional layers equivalent to aggregating node representations from their direct neighborhood (cf. [58] for more information).

It has been noted, e.g., in [59] that traditional graph neural networks including GCN are mostly targeted at the case of *sparse* graphs, where each node is only connected to a small number of neighbors. The fully connected graphs that we utilize in this work present challenges for GCN through their spectral properties. Most notably, these *dense* graphs typically have large eigengaps, i.e., the gap between the smallest eigenvalue $\lambda_1 = 0$ and the second eigenvalue $\lambda_2 > 0$ may be close to 1. Hence the GCN filter acts almost like a projection onto the undesirable eigenvector ϕ_1 . However, it has been observed in the same work that in some applications, GCNs applied to *sparsified* graphs yield comparable results to dedicated dense methods. Our experiments justified only using Standard GCN on a k -nearest neighbor subgraph.

4.3. Other Semi-supervised Learning Methods

In the context of graph-based semi-supervised learning a rather straightforward approach follows from minimizing the following objective

$$\min_u \frac{1}{2} \|u - f\|_2^2 + \frac{\beta}{2} u^T L_{\text{sym}} u \quad (17)$$

where f holds the values 1, -1 , and 0 according to the labeled and unlabeled data. Calculating the derivative shows that in order to obtain u , we need to solve the following *linear system* of equations

$$(I + \beta L_{\text{sym}}) u = f$$

where I is the identity matrix of the appropriate dimensionality.

Furthermore, we compare our previously introduced approaches to the well-known one-nearest neighbor (1NN) method. In the context of time series classification this method was proposed in [11]. In each iteration, we identify the indices i, j with the shortest distance between the labeled sample \mathbf{x}_i and the unlabeled sample \mathbf{x}_j . The label of \mathbf{x}_i is then copied to \mathbf{x}_j . This process is repeated until no unlabeled data remain.

In [60], the authors construct several graph Laplacians and then perform the semi-supervised learning based on a weighted sum of the Laplacian matrices.

5. NUMERICAL EXPERIMENTS

In this section, we illustrate how the algorithms discussed in this paper perform when applied to multiple time series data sets. We here focus on binary classification and use time series taken from the UCR time series classification archive ⁴ [61]. All our codes are to be found at <https://github.com/dominikalfke/TimeSeriesSSL>. The distance measure we use here are the previously introduced

⁴We focussed on all binary classification series listed in `TwoClassProblems.csv` within http://www.timeseriesclassification.com/Downloads/Archives/Univariate2018_arff.zip.

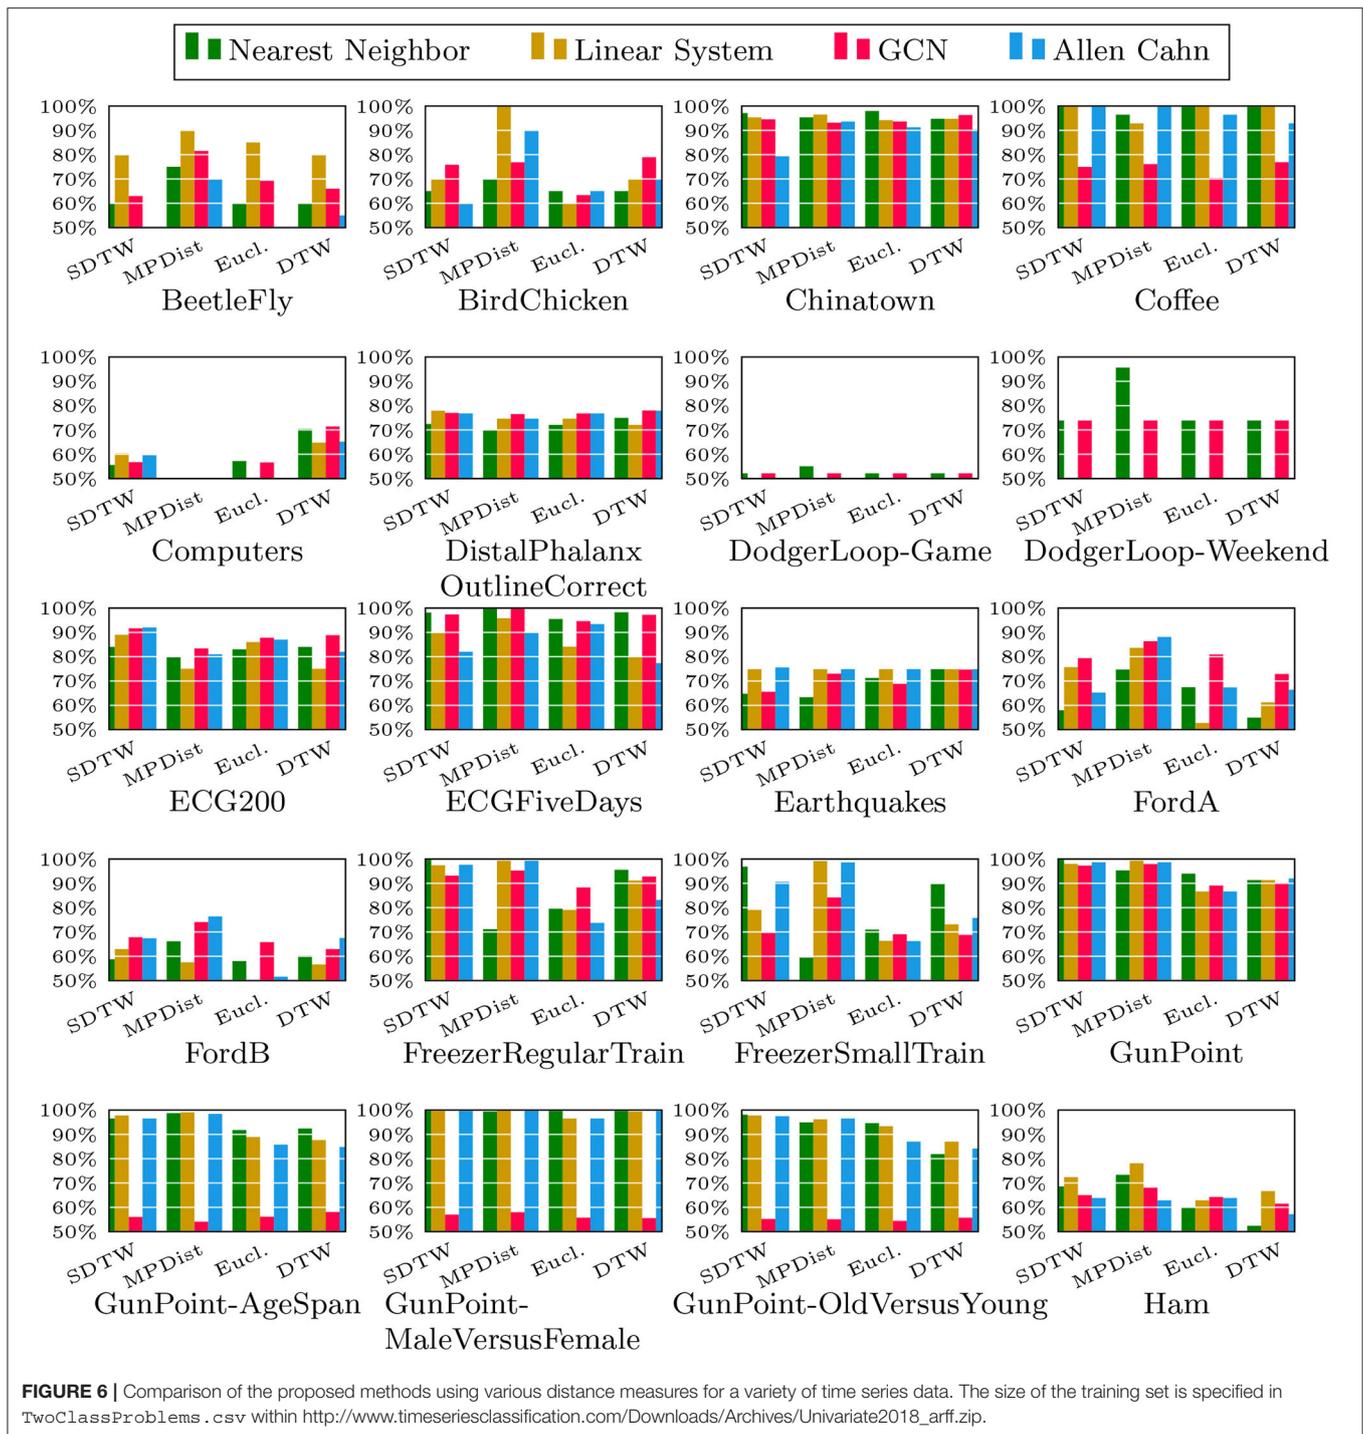


FIGURE 6 | Comparison of the proposed methods using various distance measures for a variety of time series data. The size of the training set is specified in `TwoClassProblems.csv` within http://www.timeseriesclassification.com/Downloads/Archives/Univariate2018_arff.zip.

DTW, Soft DTW divergence, MP, and Euclidean distances. For completeness, we list the default parameters for all methods in **Table 2**.

We split the presentation of the numerical results in the following way. We start by exploring the dependence of our schemes on some of the hyperparameters inherent in their derivation. We start by investigating the self-tuning parameters, namely the value of the chosen neighbor to

compute the local scaling. We then study the performance of the Allen-Cahn model depending on the number of eigenpairs used for the approximation of the graph Laplacian. For our main study, we pair up all distance measures with all learning methods and report the results on all datasets. Furthermore, we investigate how the method's performance depends on the number of available training data using random training splits.

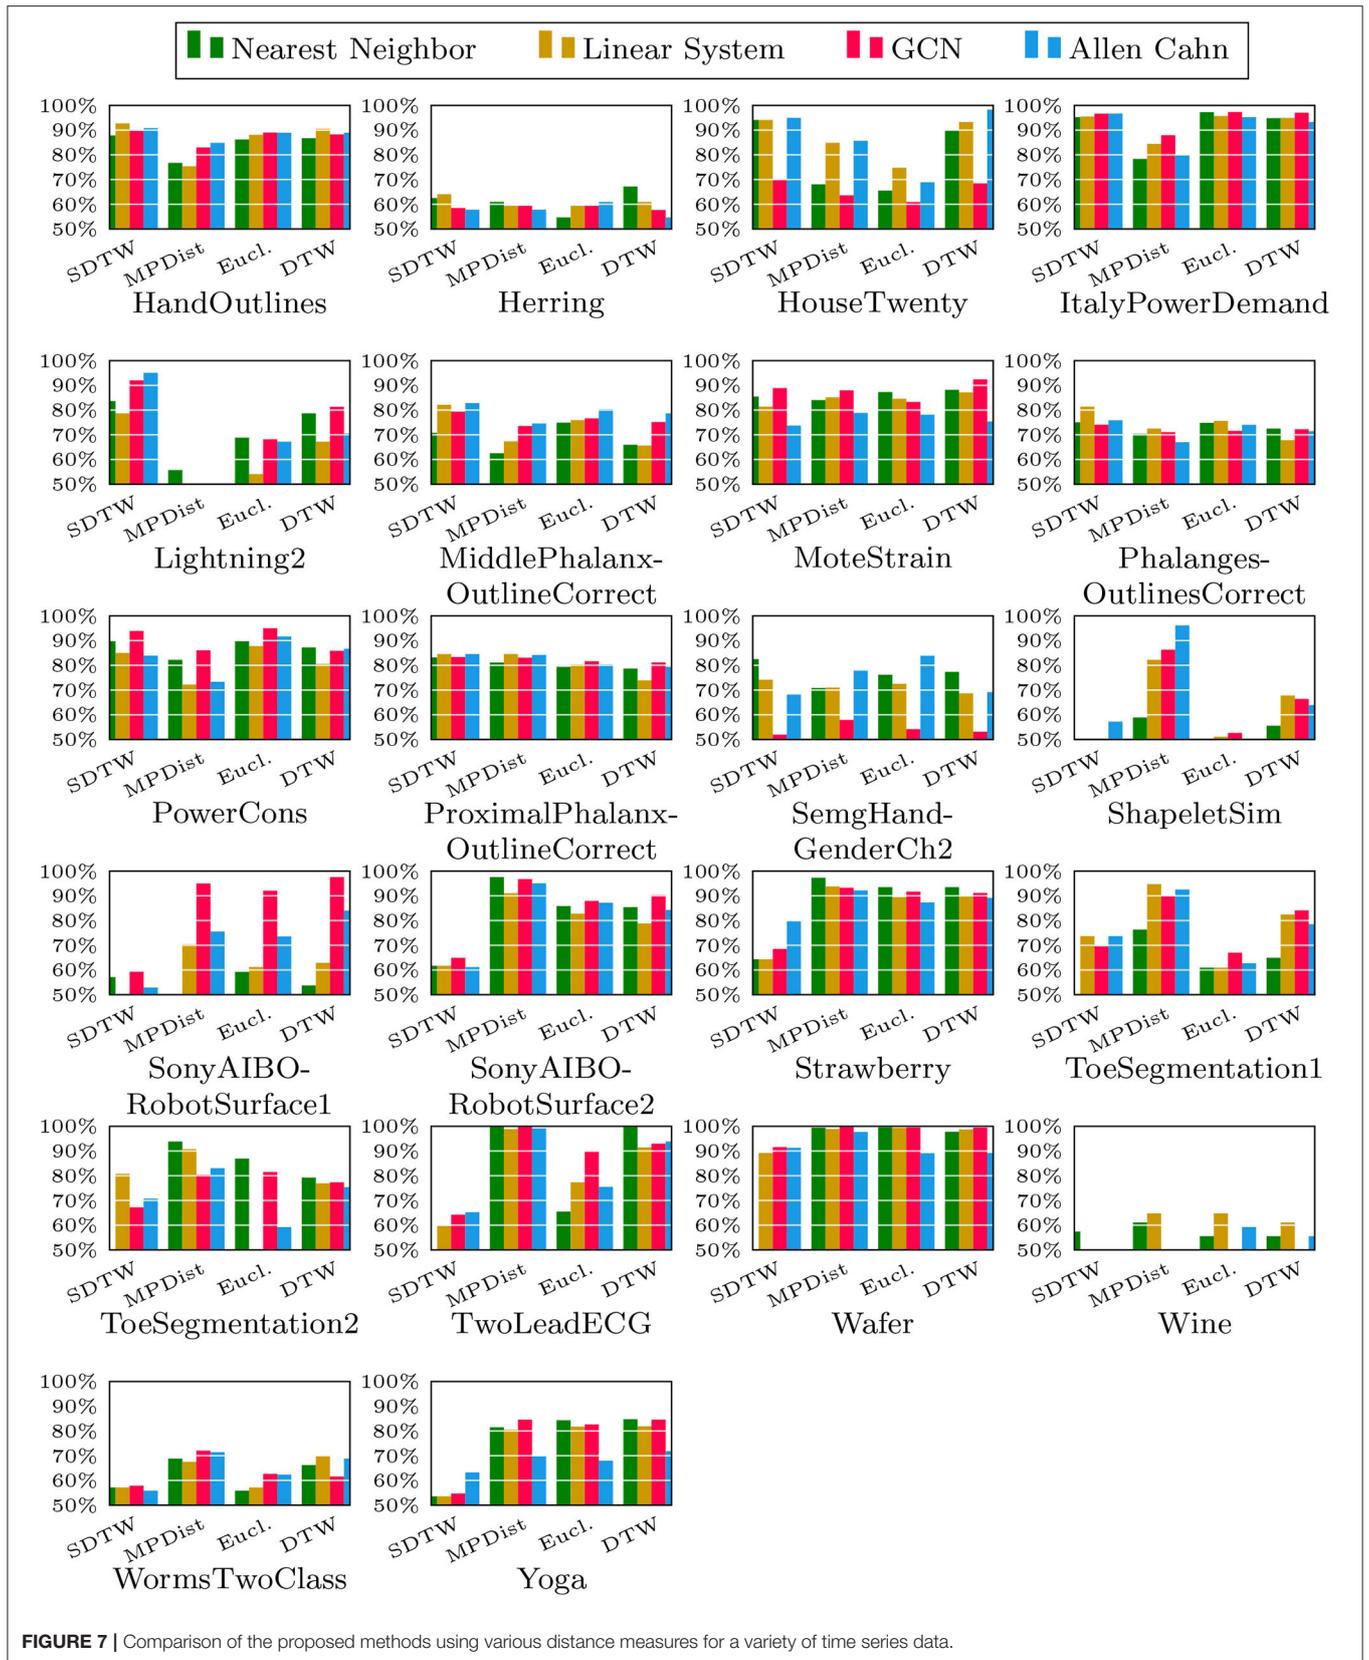


FIGURE 7 | Comparison of the proposed methods using various distance measures for a variety of time series data.

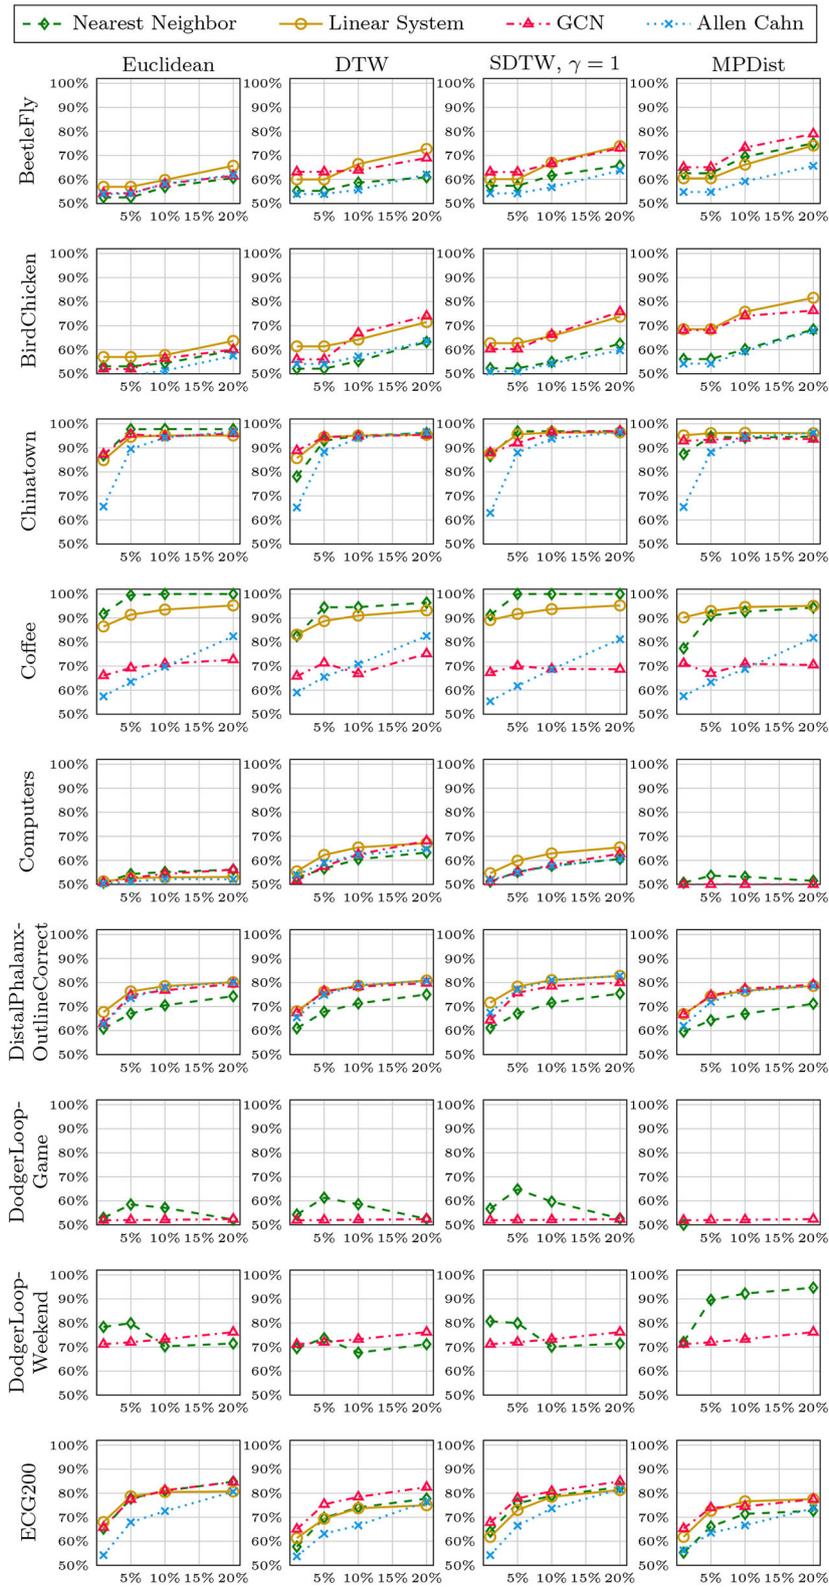


FIGURE 8 | Method accuracy comparison for random training splits of different sizes (part 1/5).

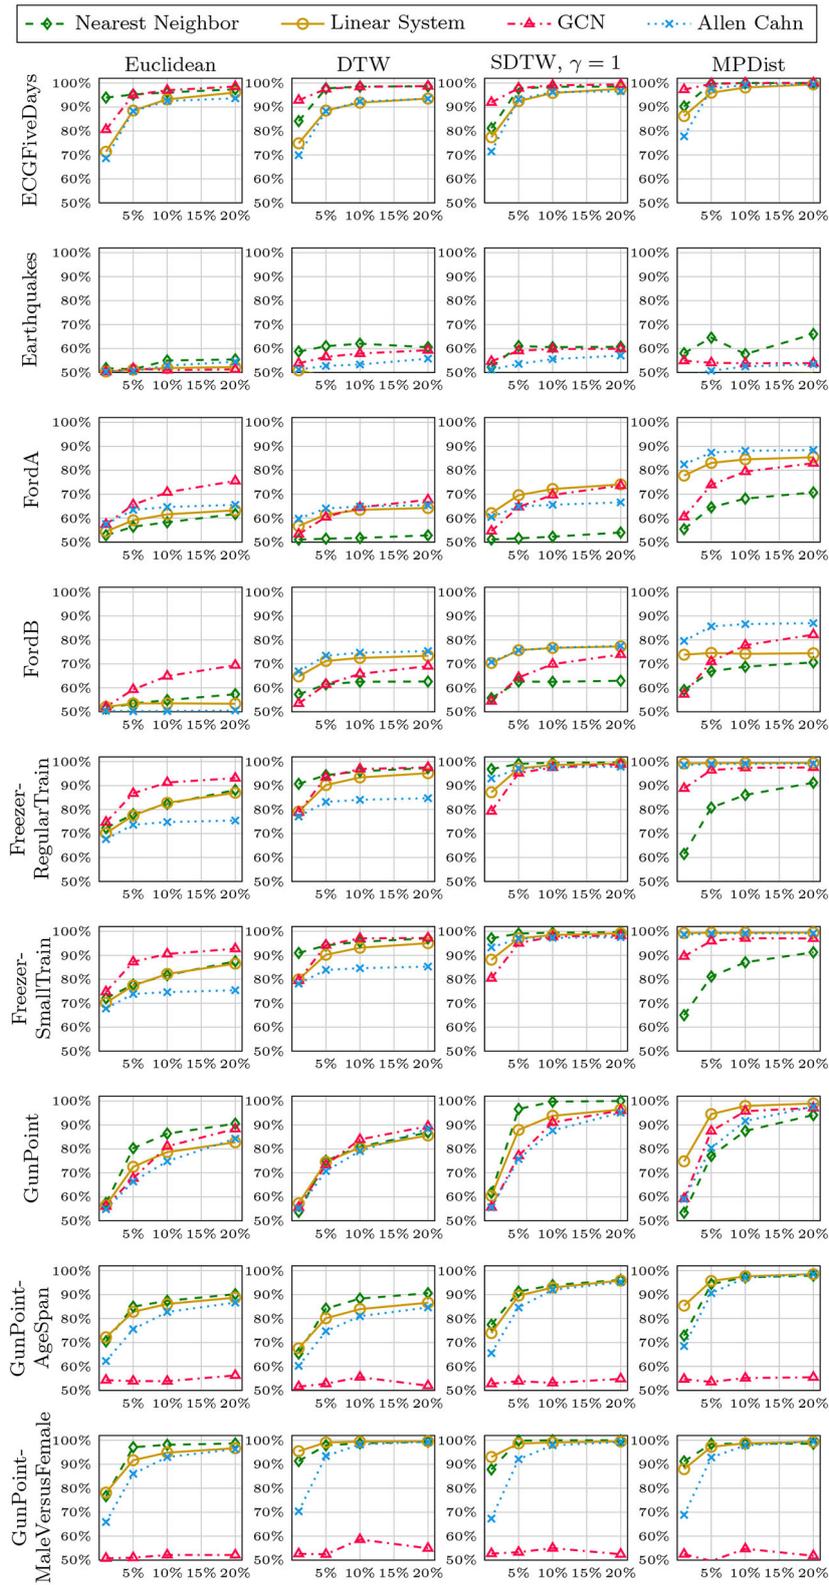


FIGURE 9 | Method accuracy comparison for random training splits of different sizes (part 2/5).

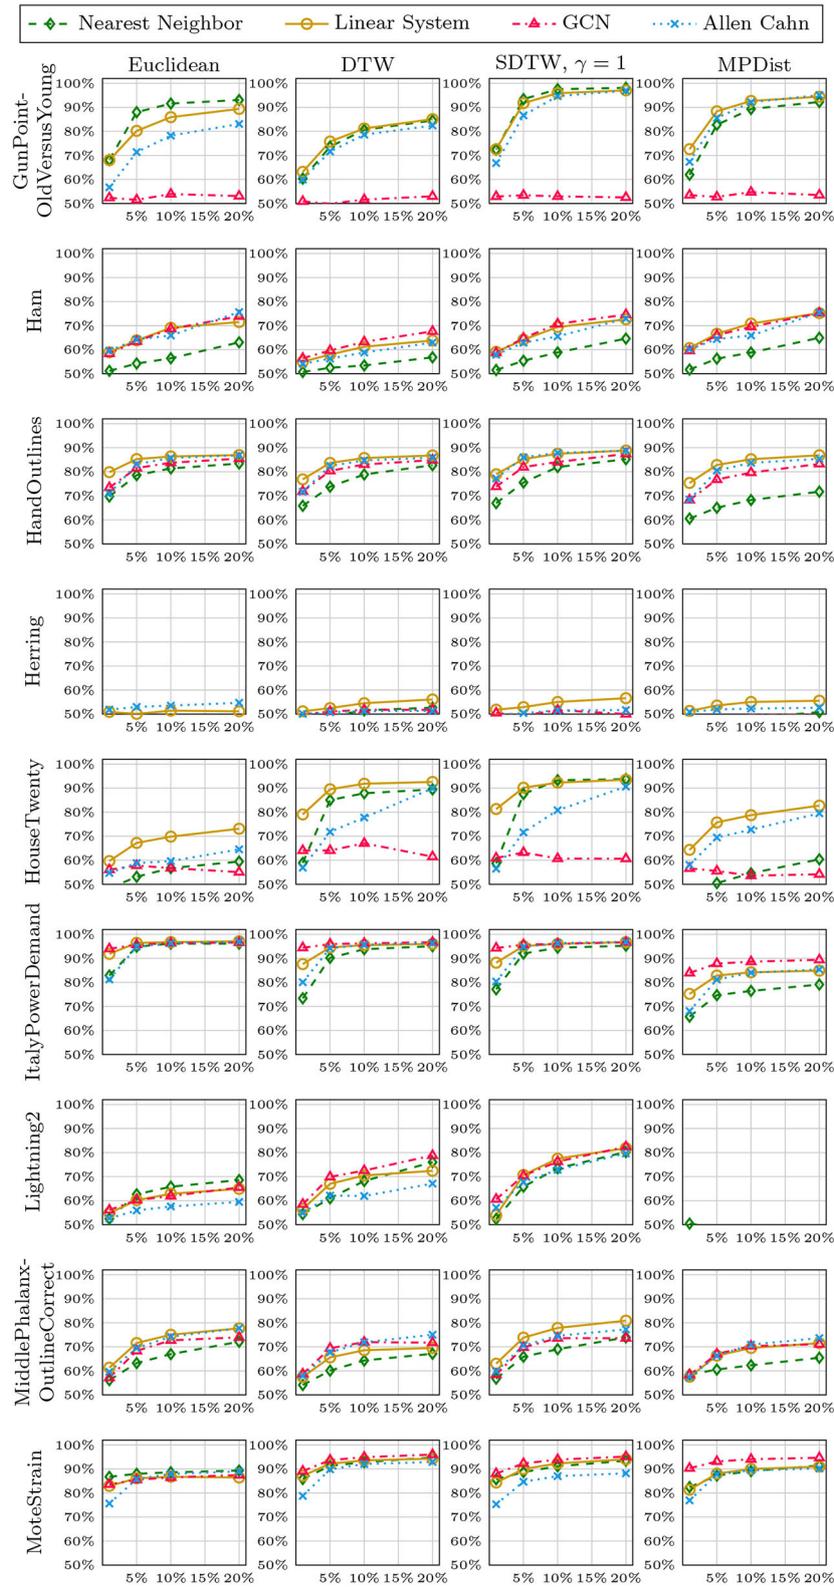


FIGURE 10 | Method accuracy comparison for random training splits of different sizes (part 3/5).

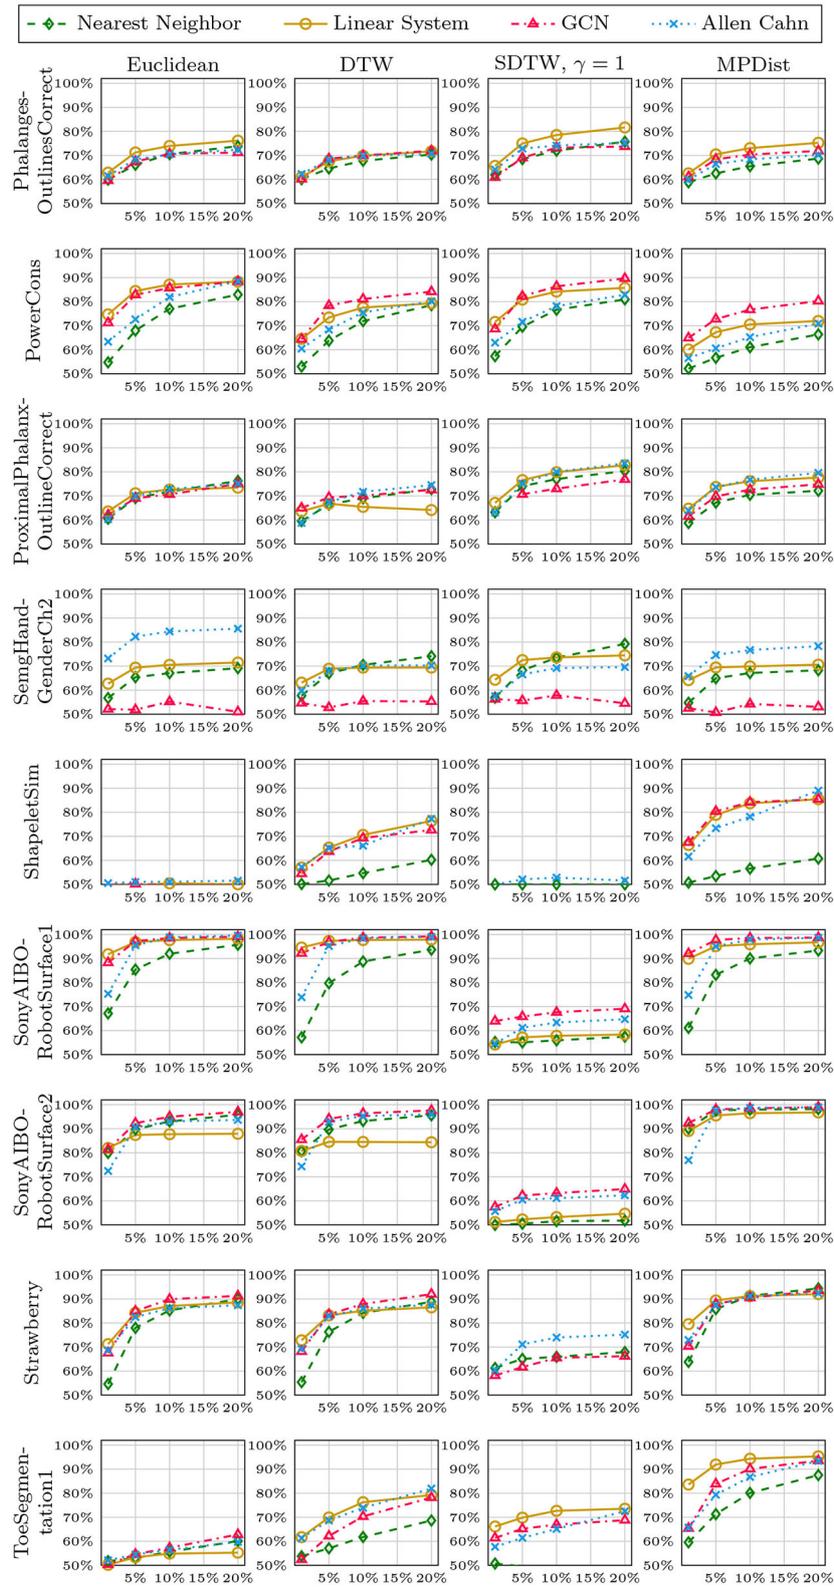


FIGURE 11 | Method accuracy comparison for random training splits of different sizes (part 4/5).

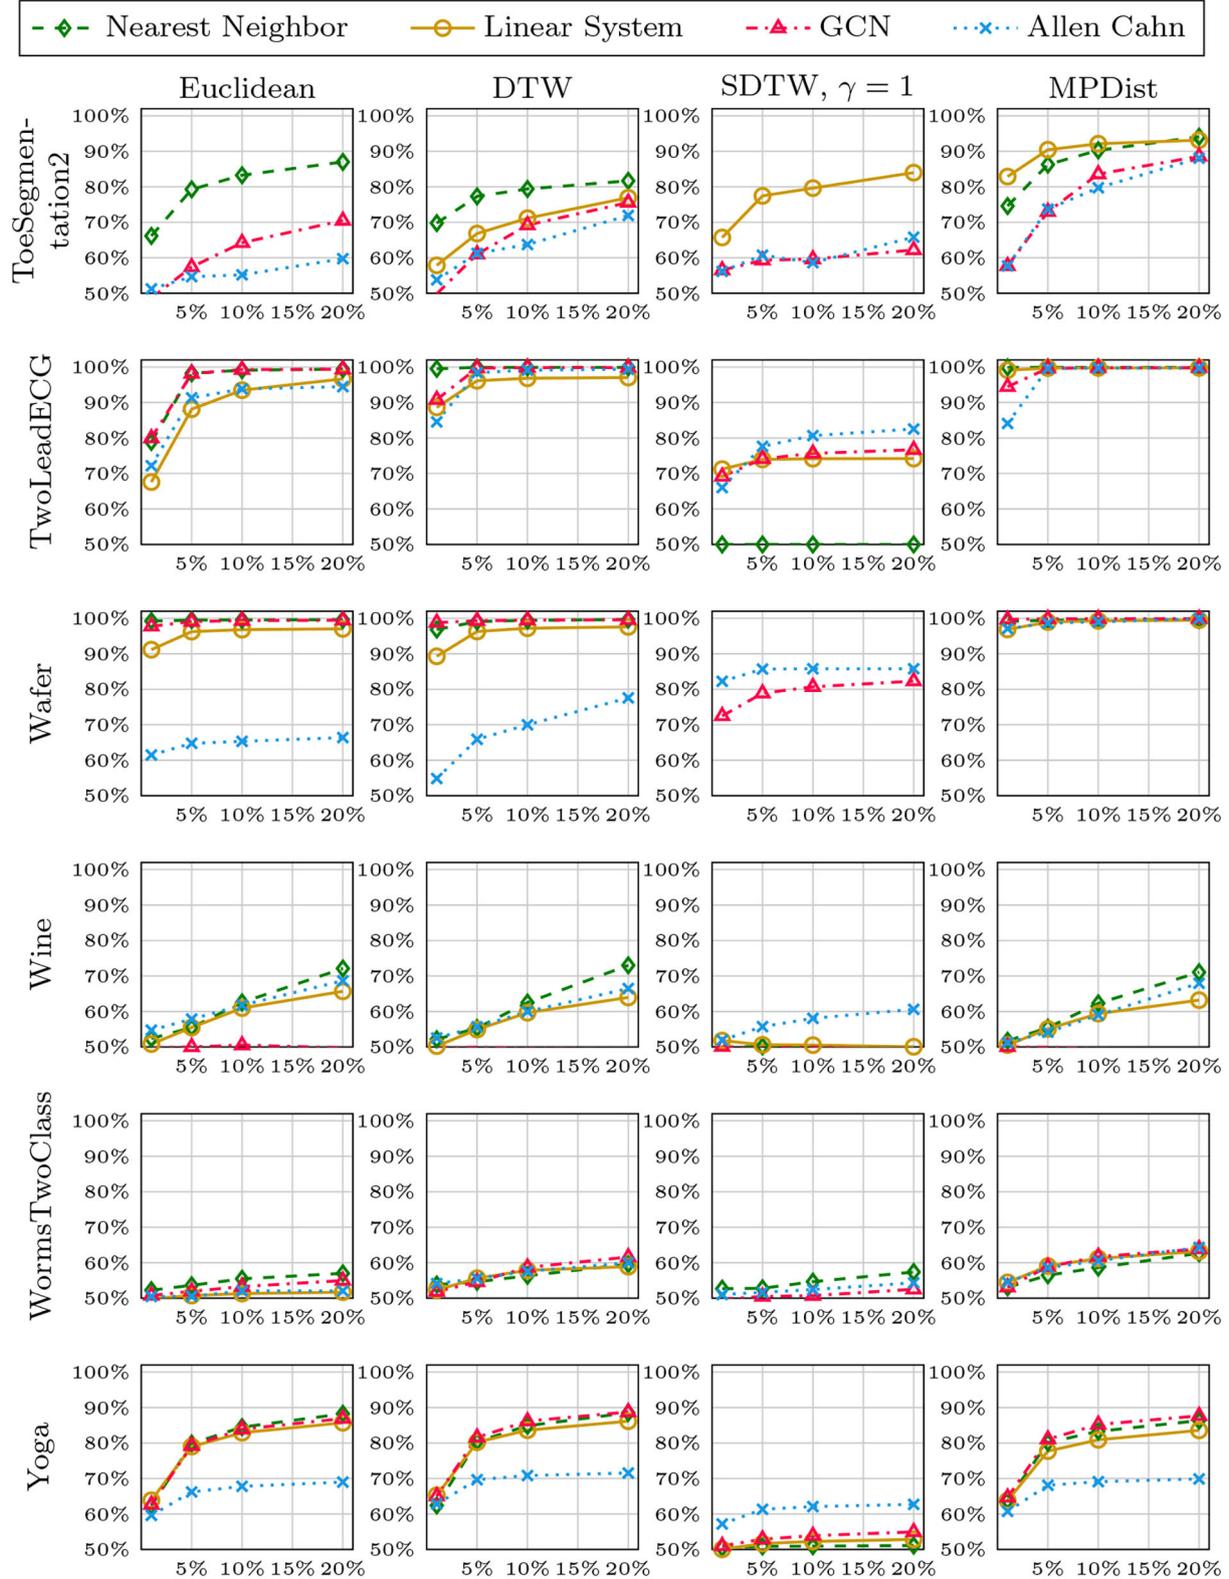


FIGURE 12 | Method accuracy comparison for random training splits of different sizes (part 5/5).

5.1. Self-Tuning Values

In section 2, we proposed the use of the self-tuning approach for the Gaussian function within the weight matrix. The crucial hyperparameter we want to explore now is the choice of neighbor k for the construction of $\sigma_i = \text{dist}(\mathbf{x}_i, \mathbf{x}_{k,i})$ with $\mathbf{x}_{k,i}$ the k -th nearest neighbor of the data point \mathbf{x}_i . We can see from **Table 3** that the small values $k = 7, 20$ perform quite well in comparison to the larger self-tuning parameters. As a result we will use these smaller values in all further computations.

5.2. Spectral Approximation

As described in section 4 the Allen–Cahn equation is projected to a lower-dimensional space using the insightful information provided by the eigenvectors to the smallest eigenvalues of the graph Laplacian. We now investigate how the number of used eigenvectors impacts the accuracy. In the following we vary the number of eigenvalues from 10 to 190 and compare the performance of the Allen–Cahn method on three different datasets. The results are shown in **Table 4** and it becomes clear that a vast number of eigenvectors does not lead to better classification accuracy. As a result we require a smaller number of eigenpair computations and also fewer computations within the Allen–Cahn scheme itself. The comparison was done for the self-tuning parameter $k = 7$.

5.3. Full Method Comparison

We now compare the Allen–Cahn approach, the GCN scheme, the linear systems based method, and the INN algorithm, each paired up with each of the distance measures introduced in section 3. Full results are listed in **Figures 6, 7**. We show the comparison for all 42 datasets.

As can be seen there are several datasets where the performance of all methods is fairly similar even when the distance measure is varied. Here, we name *Chinatown*, *Earthquakes*, *GunPoint*, *ItalyPowerDemand*, *MoteStrain*, *Wafer*. There are several examples where the methods do not seem to perform well, with GCN and INN relatively similar outperforming the Linear System and Allen–Cahn approach. Such examples are *DodgerLoopGame*, *DodgerLoopWeekend*. The GCN method clearly does not perform well with the *GunPoint* datasets where the other methods clearly perform well. It is surprising to note that the Euclidean distance, given its computational speed and simplicity, does not come out as underperforming with respect to the accuracy across the different methods. There are very few datasets where one distance clearly outperforms the other choice. We name *ShapeletSim*, *ToeSegmentation1* here. One might conjecture that the varying sizes of the training data might be a reason for the difference in performance of the models. To investigate this further we will next vary the training splits for all datasets and methods.

5.4. Varying Training Splits

In **Figures 8–12**, we vary the size of the training set from 1 to 20% of the available data. All reported numbers are averages over 100 random splits. The numbers we observe mirror the performance of the full training size. We see that the methods show reduced performance when only 1% of the training data are used but often reach an accuracy plateau when 5 to 10% of the training data are used. We observe that the size of the training set alone does not explain the different performance in the various datasets and methods applied here.

6. CONCLUSION

In this paper we took to the task of classifying time series data in a semi-supervised learning setting. For this we proposed to represent the data as a fully-connected graph where the edge weights are created based on a Gaussian similarity measure (1). The heart of this function is the difference measure between the time series, for which we used the (Soft) Dynamic Time Warping and Matrix Profile based distance measures as well as the Euclidean distance. We then investigated several learning algorithms, namely, the Allen–Cahn-based method, the Graph Convolutional Network scheme, and a linear system approach, all reliant on the graph Laplacian, as well as the Nearest Neighbor method. We then illustrated the performance of all pairs of distance measure and learning methods. In this empirical study we observed that the methods tend to show an increased performance adding more training data. Studying all binary time-series with the *timeseriesclassification.com* repository gives results that in accordance with the no free lunch theorem show no clear winner. On the positive side the methods often perform quite well and there are only a few datasets with decreased performance. The comparison of the distance measures indicates there are certain cases where they outperform their competitors but also there is no clear winner with regards to accuracy. We believe that this empirical, reproducible study will encourage further research in this direction. Additionally, it might be interesting to consider model-based representations of time-series such as ARMA [63, 64] to use within the graph representations used here.

DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author/s.

AUTHOR CONTRIBUTIONS

MG provide the initial implementation of some of the methods. MS supervised the other members, wrote parts of the manuscript, and implemented the Allen Cahn scheme. DB implemented the GCN approach, wrote parts of the manuscript, and oversaw the design of the tests. LP implemented several algorithms and wrote parts of the manuscript. All authors contributed to the article and approved the submitted version.

FUNDING

MS and LP acknowledge the funding of the BMBF grant 01|S20053A. DB was partially supported by KINTUC project [Sächsische Aufbaubank–Förderbank–(SAB) 100378180].

REFERENCES

- Fu TC. A review on time series data mining. *Eng Appl Artif Intell.* (2011) 24:164–81. doi: 10.1016/j.engappai.2010.09.007
- Bello-Orgaz G, Jung JJ, Camacho D. Social big data: recent achievements and new challenges. *Inform Fusion.* (2016) 28:45–59. doi: 10.1016/j.inffus.2015.08.005
- Chen F, Deng P, Wan J, Zhang D, Vasilakos AV, Rong X. Data mining for the internet of things: literature review and challenges. *Int J Distribut Sensor Netw.* (2015) 11:431047. doi: 10.1155/2015/431047
- Laptev N, Amizadeh S, Flint I. Generic and scalable framework for automated time-series anomaly detection. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* (2015). p. 1939–47. doi: 10.1145/2783258.2788611
- Chiu B, Keogh E, Lonardi S. Probabilistic discovery of time series motifs. In: *Proceedings of the ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* (2003). p. 493–8. doi: 10.1145/956750.956808
- De Gooijer JG, Hyndman RJ. 25 years of time series forecasting. *Int J Forecast* (2006). 22:443–73. doi: 10.1016/j.ijforecast.2006.01.001
- Wei WW. Time series analysis. In: Todd little editor. *The Oxford Handbook of Quantitative Methods in Psychology.* Vol. 2. (2006).
- Chatfield C, Xing H. *The Analysis of Time Series: An Introduction with R.* CRC Press. (2019) doi: 10.1201/9781351259446
- Fawaz HI, Forestier G, Weber J, Idoumghar L, Muller PA. Deep learning for time series classification: a review. *Data Mining Knowledge Discov.* (2019) 33:917–63. doi: 10.1007/s10618-019-00619-1
- Abanda A, Mori U, Lozano JA. A review on distance based time series classification. *Data Mining Knowledge Discov.* (2019) 33:378–412. doi: 10.1007/s10618-018-0596-4
- Wei L, Keogh E. Semi-supervised time series classification. In: *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* (2006). p. 748–53. doi: 10.1145/1150402.1150498
- Liao TW. Clustering of time series data—a survey. *Pattern Recogn.* (2005) 38:1857–74. doi: 10.1016/j.patcog.2005.01.025
- Aghabozorgi S, Shirkhorshidi AS, Wah TY. Time-series clustering—a decade review. *Inform Syst.* (2015) 53:16–38. doi: 10.1016/j.is.2015.04.007
- Shifaz A, Pelletier C, Petitjean F, Webb GI. TS-CHIEF: a scalable and accurate forest algorithm for time series classification. *Data Mining Knowledge Discov.* (2020) 34:742–75. doi: 10.1007/s10618-020-00679-8
- Dempster A, Petitjean F, Webb GI. ROCKET: exceptionally fast and accurate time series classification using random convolutional kernels. *Data Mining Knowledge Discov.* (2020) 34:1454–95. doi: 10.1007/s10618-020-00701-z
- Fawaz HI, Lucas B, Forestier G, Pelletier C, Schmidt DF, Weber J, et al. Inceptiontime: Finding alexnet for time series classification. *Data Mining Knowledge Discov.* (2020) 34:1936–62. doi: 10.1007/s10618-020-00710-y
- Zhu X, Goldberg A. *Introduction to Semi-supervised Learning.* Morgan & Claypool Publishers (2009). doi: 10.2200/S00196ED1V01Y200906AIM006
- Chapelle O, Schölkopf B, Zien A. Semi-supervised learning. *IEEE Trans Neural Netw.* (2009) 20:542. doi: 10.1109/TNN.2009.2015974
- Stoll M. A literature survey of matrix methods for data science. *GAMM-Mitt.* (2020) 43:e202000013:4. doi: 10.1002/gamm.202000013
- Mercado P, Bosch J, Stoll M. Node classification for signed social networks using diffuse interface methods. In: *ECMLPKDD.* (2019). doi: 10.1007/978-3-030-46150-8_31
- Kipf TN, Welling M. Semi-supervised classification with graph convolutional networks. *arXiv [Preprint]. arXiv:160902907* (2016).
- Bertozi AL, Luo X, Stuart AM, Zygalakis KC. Uncertainty quantification in graph-based classification of high dimensional data. *SIAM/ASA J Uncertainty Quant.* (2018) 6:568–95. doi: 10.1137/17M1134214
- von Luxburg U. A tutorial on spectral clustering. *Stat Comput.* (2007) 17:395–416. doi: 10.1007/s11222-007-9033-z
- Bruna J, Zaremba W, Szlam A, LeCun Y. Spectral networks and locally connected networks on graphs. *arXiv [Preprint]. arXiv:13126203* (2013).
- Chung FR, Graham FC. *Spectral graph Theory.* American Mathematical Soc. (1997).
- Jung A. Networked exponential families for big data over networks. *IEEE Access.* (2020) 8:202897–909. doi: 10.1109/ACCESS.2020.3033817
- Jung A, Tran N. Localized linear regression in networked data. *IEEE Signal Process Lett.* (2019) 26:1090–4. doi: 10.1109/LSP.2019.2918933
- Müller M. *Information Retrieval for Music and Motion.* vol. 2. Springer (2007). doi: 10.1007/978-3-540-74048-3
- Cuturi M, Blondel M. Soft-DTW: a differentiable loss function for time-series. In: *International Conference on Machine Learning.* PMLR (2017). p. 894–903.
- Gharghabi S, Imani S, Bagnall A, Darvishzadeh A, Keogh E. An ultra-fast time series distance measure to allow data mining in more complex real-world deployments. *Data Mining Knowledge Discov.* (2020) 34:1104–35. doi: 10.1007/s10618-020-00695-8
- Bertozi AL, Flenner A. Diffuse interface models on graphs for classification of high dimensional data. *Multiscale Model Simul.* (2012) 10:1090–118. doi: 10.1137/11083109X
- Bagnall A, Lines J, Bostrom A, Large J, Keogh E. The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *Data Mining Knowledge Discov.* (2017) 31:606–60. doi: 10.1007/s10618-016-0483-9
- Ruiz AP, Flynn M, Large J, Middlehurst M, Bagnall A. The great multivariate time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *Data Mining Knowledge Discov.* (2021) 35:401–49. doi: 10.1007/s10618-020-00727-3
- MacQueen J. Some methods for classification and analysis of multivariate observations. In: *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability.* Oakland, CA (1967). p. 281–97.
- MacKay DJ, Mac Kay DJ. *Information Theory, Inference and Learning Algorithms.* Cambridge University Press (2003).
- Belkin M, Niyogi P. Laplacian eigenmaps and spectral techniques for embedding and clustering. In: *Advances in Neural Information Processing Systems.* (2001). p. 585–91.
- Shawe-Taylor J, Cristianini N. *Kernel Methods for Pattern Analysis.* Cambridge University Press (2004). doi: 10.1017/CBO9780511809682
- Hofmann T, Schölkopf B, Smola AJ. Kernel methods in machine learning. *Ann Stat.* (2008) 36:1171–220. doi: 10.1214/00905360700000677
- Zelnik-Manor L, Perona P. Self-tuning spectral clustering. In: *Advances in Neural Information Processing Systems.* (2005). p. 1601–8.
- Keogh E, Kasetty S. On the need for time series data mining benchmarks: a survey and empirical demonstration. *Data Mining Knowledge Discov.* (2003) 7:349–71. doi: 10.1023/A:1024988512476
- Salvador S, Chan PK. Toward accurate dynamic time warping in linear time and space. *Intell Data Anal.* (2004) 11:70–80. doi: 10.3233/IDA-2007-11508
- Wu R, Keogh EJ. FastDTW is approximate and generally slower than the algorithm it approximates. *IEEE Trans Knowledge Data Eng.* (2020). doi: 10.1109/TKDE.2020.3033752
- Blondel M, Mensch A, Vert JP. Differentiable divergences between time series. *arXiv [Preprint]. arXiv:201008354* (2020).

ACKNOWLEDGMENTS

All authors would like to acknowledge the hard work and dedication by the team maintaining www.timeseriesclassification.com/. The publication of this article was funded by Chemnitz University of Technology.

44. Lin H, Hong X, Ma Z, Wei X, Qiu Y, Wang Y, et al. Direct measure matching for crowd counting. *arXiv [Preprint]*. *arXiv:210701558* (2021). doi: 10.24963/ijcai.2021/116
45. Taylor JE, Cahn JW. Linking anisotropic sharp and diffuse surface motion laws via gradient flows. *J Statist Phys.* (1994) 77:183–97. doi: 10.1007/BF02186838
46. Allen SM, Cahn JW. A microscopic theory for antiphase boundary motion and its application to antiphase domain coarsening. *Acta Metall.* (1979) 27:1085–95. doi: 10.1016/0001-6160(79)90196-2
47. Cahn JW, Hilliard JE. Free energy of a nonuniform system. I. Interfacial free energy. *J Chem Phys.* (1958) 28:258–67. doi: 10.1063/1.1744102
48. Bosch J, Kay D, Stoll M, Wathen A. Fast solvers for Cahn-Hilliard inpainting. *SIAM J Imaging Sci.* (2014) 7:67–97. doi: 10.1137/130921842
49. Bertozzi AL, Esedoglu S, Gillette A. Inpainting of binary images using the Cahn-Hilliard equation. *IEEE Trans Image Process.* (2007) 16:285–91. doi: 10.1109/TIP.2006.887728
50. Garcia-Cardona C, Merkurjev E, Bertozzi AL, Flenner A, Percus AG. Multiclass data segmentation using diffuse interface methods on graphs. *IEEE Trans Pattern Anal Mach Intell.* (2014) 36:1600–13. doi: 10.1109/TPAMI.2014.2300478
51. Bosch J, Klamt S, Stoll M. Generalizing diffuse interface methods on graphs: nonsmooth potentials and hypergraphs. *SIAM J Appl Math.* (2018) 78:1350–77. doi: 10.1137/17M1117835
52. Bergermann K, Stoll M, Volkmer T. Semi-supervised learning for multilayer graphs using diffuse interface methods and fast matrix vector products. *SIAM J Math Data Sci.* (2021). doi: 10.1137/20M1352028
53. Budd J, van Gennip Y. Graph MBO as a semi-discrete implicit Euler scheme for graph Allen-Cahn. *arXiv [Preprint]*. *arXiv:190710774* (2019). doi: 10.1137/19M1277394
54. Budd J, van Gennip Y, Latz J. Classification and image processing with a semi-discrete scheme for fidelity forced Allen-Cahn on graphs. *arXiv [Preprint]*. *arXiv:201014556* (2020). doi: 10.1002/gamm.202100004
55. Calatroni L, van Gennip Y, Schönlieb CB, Rowland HM, Flenner A. Graph clustering, variational image segmentation methods and Hough transform scale detection for object measurement in images. *J Math Imaging Vision.* (2017) 57:269–91. doi: 10.1007/s10851-016-0678-0
56. Goodfellow I, Bengio Y, Courville A, Bengio Y. *Deep Learning*. vol. 1. Cambridge: MIT Press (2016).
57. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature.* (2015) 521:436–44. doi: 10.1038/nature14539
58. Zhang S, Tong H, Xu J, Maciejewski R. Graph convolutional networks: a comprehensive review. *Comput Soc Netw.* (2019) 6:1–23. doi: 10.1186/s40649-019-0069-y
59. Alfke D, Stoll M. Pseudoinverse graph convolutional networks: fast filters tailored for large eigengaps of dense graphs and hypergraphs. *Data Mining Knowledge Discov.* (2021). doi: 10.1007/s10618-021-00752-w
60. Xu Z, Funaya K. Time series analysis with graph-based semi-supervised learning. In: *2015 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. IEEE (2015). p. 1–6. doi: 10.1109/DSAA.2015.7344902
61. Dau HA, Bagnall A, Kamgar K, Yeh CCM, Zhu Y, Gharghabi S, et al. The UCR time series archive. *IEEE/CAA J Automat Sin.* (2019) 6:1293–305. doi: 10.1109/JAS.2019.1911747
62. Kingma D, Ba JL. Adam: a method for stochastic optimization. In: *Proc Int Conf Learn Represent. ICLR'15*. (2015).
63. Brockwell PJ, Davis RA. *Time Series: Theory and Methods*. Springer Science & Business Media (2009).
64. Spiegel S. *Time series distance measures*. Ph.D. thesis. Berlin, Germany.

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2022 Bünger, Gondos, Peroche and Stoll. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.