

Capturing Temporal Components for Time Series Classification

Venkata Ragavendra Vavilthota¹, Ranjith Ramanathan², and Sathyanarayanan N. Aakur³[0000–0003–1062–8929]

¹ Department of Computer Science, Oklahoma State University, Stillwater OK 74078, USA

² Department of Animal and Food Sciences, Oklahoma State University, Stillwater OK 74078, USA

³ Department of Computer Science and Software Engineering, Auburn University, Auburn, 36849
san0028@auburn.edu

Abstract. Analyzing sequential data is crucial in many domains, particularly due to the abundance of data collected from the Internet of Things paradigm. Time series classification, the task of categorizing sequential data, has gained prominence, with machine learning approaches demonstrating remarkable performance on public benchmark datasets. However, progress has primarily been in designing architectures for learning representations from raw data at fixed (or ideal) time scales, which can fail to generalize to longer sequences. This work introduces a *compositional representation learning* approach trained on statistically coherent components extracted from sequential data. Based on a multi-scale change space, an unsupervised approach is proposed to segment the sequential data into chunks with similar statistical properties. A sequence-based encoder model is trained in a multi-task setting to learn compositional representations from these temporal components for time series classification. We demonstrate its effectiveness through extensive experiments on publicly available time series classification benchmarks. Evaluating the coherence of segmented components shows its competitive performance on the unsupervised segmentation task.

Keywords: Time-series classification · Temporal Compositionality · Time Series Segmentation.

1 Introduction

Time series data is ubiquitous in many domains, such as healthcare [30] and robotics [39]. Given the widespread presence of sensors and smart devices, abundant sequential (time series) data across different domains has been collected, giving rise to several important tasks in time series analysis, such as classification, segmentation, and anomaly detection. Time series classification is one task that has received significant attention in recent years. The goal is to learn robust features from sequential data to classify them into their respective categories. Machine learning approaches [31, 29], particularly deep learning approaches [36, 44],

have shown tremendous progress in learning models for time series data classification and have resulted in interesting applications such as sleep state segmentation [30] and pandemic modeling [9], to name a view.

The sequential nature of the time series data offers several challenges for classification. First, the sequence length can vary across samples within categories, which requires learning representations robust to such intra-class variations. Second, understanding the ideal time scale for extracting meaningful patterns is challenging, primarily caused by measurement errors and phase/amplitude changes across samples. Finally, long-duration sequences can have dependencies that span different time scales and pose a significant challenge to representation learning approaches. While driving tremendous progress, learning from raw signals relies heavily on representation learning mechanisms to capture intricate, compositional properties for tackling these challenges. Explicitly capturing the underlying temporal structure of signals in the representation can help alleviate this dependency and lead to more robust performance on downstream tasks. Such representations have shown tremendous potential in scene recognition tasks [28] by considering objects as atomic components that combine to compose the overall scene. However, time series may not have such clear distinctions for recognizing boundaries between components, requiring a novel paradigm for defining and detecting temporal components in sequential data.

In this work, we propose to capture the different atomic components that combine to form these signals in a compositional representation. We consider a time series data point, or signal, to be a sequence of data points ordered by some condition and can be segmented into chunks that share semantic or statistical properties. These chunks, or sub-series, are called *components* of the overall signal. Rather than learning representations over the raw sequential data, representations from this sequence of *components* can result in a compositional feature that can span longer durations with reduced computational complexity. The overall approach is illustrated in Figure 1. We first establish a multi-scale change space (Section 3.1) to segment (or tokenize) the signal into components at different temporal scales. Then, we learn compositional representations (Section 3.2) from these segments in a multi-task learning setting. Extensive evaluation (Section 4) on publicly available benchmark datasets shows that the approach performs competitively with state-of-the-art approaches and scales well to longer duration time series data. These components are remarkably similar to natural segments found in time series data (Section 5), and the approach can naturally be extended to unsupervised time series segmentation. Without bells and whistles, the approach performs competitively to state-of-the-art techniques designed explicitly for segmentation and outperforms other non-learning-based methods.

The **contributions** of our approach are four-fold: (i) we are, to the best of our knowledge, to introduce a multi-scale change space for time series data to segment them into statistically atomic components, (ii) we introduce the notion of compositional feature learning from temporally segmented components in time series data rather than modeling the raw data points, (iii) we show




Fig. 1. Overall architecture of the proposed approach is illustrated here. First, we introduce a multi-scale state change detection model to segment sequential data into components and then use a sequence-based encoder to learn compositional representations for time series classification.

that the temporal components detected by the algorithm are highly correlated with natural boundaries in time series data by evaluating it on the time series segmentation task, achieving state-of-the-art performance compared with other non-learning-based approaches, and (iv) we establish a competitive baseline that provides competitive performance with the state-of-the-art approaches on benchmark datasets for both time series classification and segmentation with limited training needs and without explicit handcrafting.

We structure the paper as follows. We review the relevant literature and techniques used in this work in Section 2, followed by an overview and detailed explanation of the proposed approach in Section 3. We present and analyze the quantitative results in Section 4 and demonstrate how it can be expanded to tackle other time series analysis tasks such as unsupervised segmentation in Section 5. Finally, in Section 6, we discuss its limitations and future directions.

2 Related Work

Time series classification has been tackled through three major types of approaches. *Classical* approaches, such as those based on handcrafted feature learning [31, 29, 27, 19], have attempted to learn discriminative features from modeling the time series at different scales through techniques such as shapelet transforms [27, 19], distance-based transforms [29, 7], and bag-of-symbols [31, 35], to

name a few. However, their computational complexity increases almost exponentially as the duration of the time series increases, and hence, they appear to hit a wall of scalability. *Deep learning*-based approaches, using architectures such as convolutional neural networks (CNNs) [24] and transformers [41], have opened a wave of large models pre-trained on significant amounts of data [22]. Deep learning models have focused on modeling the data at the ideal time scale [36,8] for capturing robust representations using different backbones such as CNNs [44,48,47,34,15,42], recurrent neural networks (RNNs) [37,38], and transformers [46,14]. *Ensemble-based* approaches [36,27,35,12], i.e., using multiple predictions from the different aspects of the same time series data, have made significant strides in establishing the state-of-the-art performance on several benchmark datasets [3,10,2]. Our work, however, offers a novel framework to capture multi-scale representations by detecting temporal components at different time scales and integrating them in a unified representation without the need for ensembling and additional overhead in the form of annotations.

Approaches to **time series segmentation** have primarily focused on detecting boundaries in sequential data through heuristic-based, domain-specific approaches. Broadly categorized into three categories [40], the time series is segmented by comparing the features of consecutive fixed-size windows using their likelihood of belonging to the same segment [21], assessing homogeneity using kernels [18] or mapping them into graph-based representation for extracting sub-graphs (segments) through heuristics such as pairwise similarity [5]. Search-based approaches [33,1,13] and learning-based approaches [32,17,11] have offered a way forward to domain-agnostic segmentation by learning sequence-level representations and segmenting them based on similarity measures. The former assigns costs to plausible boundaries and finds optimal segments by minimizing these costs, while the latter focuses on learning boundaries through pre-text tasks as self-supervision. Many of these approaches require the number of segments to be pre-defined, with learning-based approaches such as ClaSP [32] being a notable exception. Our approach built on BIC-based tokenization (Section 3.1) belongs to the search-based approach category and performs *domain-agnostic* segmentation by comparing statistical similarity measures without training.

3 Proposed Framework

In this section, we outline the proposed framework to extract temporal components from sequential data and learn a robust, compositional representation in a multi-task setting. We first outline the problem formulation to provide an overview of the approach and then introduce the multi-scale change space used to discover temporal components in signals. Finally, we introduce the representation learning mechanism used to combine these temporal components into a robust representation.

Problem Formulation. We address the task of classifying univariate time series data by decomposing the signal into its constituent parts. We aim to characterize and build a rich signal representation by detecting parts (sub-series)

that compose the overall signal. Inspired by theories of compositional event understanding [45], we consider these parts atomic, i.e., each sub-series cannot be broken down into smaller components. To this end, we consider a multi-scale approach to identify these components at different time scales to account for the unique challenges inherent in time series data, such as variations that are introduced during data collection [3,10] (i.e., sampling rate and record length) and unavoidable intra-class variations (such as amplitude offset and warping). Following prior works on state-spaces [26,25], we define a signal-dependent change scale-space that captures the multi-scale structure of the signal based on its temporal change points. The overall architecture is illustrated in Figure 1. First, we identify the temporal change points in the signal using a statistics-based multi-scale organization (Section 3.1), which allows us to break the signal down into its components. Second, we learn compositional relationships from these signal components using a bidirectional sequence learning model (Section 3.2) trained in a multi-task setting. Combined, these two steps help identify atomic components in time series signals and help capture their temporal structure in a purely bottom-up fashion without auxiliary data.

3.1 Discovering Temporal Components of Signals

The first step in our approach is to discover temporal sub-components that compose time series signals. These sub-components are temporal chunks whose statistics (mean, variance, etc.) are consistent within the sub-series yet vary significantly with neighboring chunks. Hence, detecting the change in statistics at multiple time scales allows us to discover these temporal components in univariate signals. We use the premise from statistics-based speaker-turn detection approaches [6,25] to define a function $TSCS$ (Time Series Change Space) to capture the temporal change space in time series data ($X_{0,N} = \{x_1, x_2, x_3, \dots, x_N\}$). It is a two-dimensional function over time (t) and temporal scale (δ) that characterizes the varying statistics between two sub-series $t-\delta$ and $t+\delta$ to detect a possible temporal change point (time series component) at time t , given a temporal scale δ . We cast this formulation as a hypothesis-testing problem. The null hypothesis is that two consecutive chunks are different and thus require two different models to represent them individually. The alternative hypothesis is that they are very similar and belong to a single, longer chunk one model can represent. We evaluate each hypothesis by fitting a single Gaussian model [6] for the chunks from each hypothesis. Hence, the difference in the Bayesian Information Criterion (BIC) between the two models at time t provides a measure of their separability based on their statistics. Formally, we define the state space ($TSCS(t, \delta)$) as a function of BIC given by

$$TSCS(t, \delta) = \frac{\delta}{2}(\log|\sigma_{X_{t-\delta,t}}| + \log|\sigma_{X_{t,t+\delta}}|) - \delta(\log|\Sigma_{X_{t-\delta,t+\delta}}|) + \delta P \quad (1)$$

where $\log|\sigma_{X_{t-\delta,t}}|$ and $\log|\sigma_{X_{t,t+\delta}}|$ refer to the BIC of the single Gaussian representation for the subseries from time $t-\delta$ to t and from t to $t+\delta$, respectively;

$\log|\Sigma_{X_{t-\delta:t+\delta}}|$ refers to the BIC of a multivariate jointly considering both sub-series; and P is a penalty term to account for the size of the subseries considered and is typically set to $\log(T)$ where T is the length of the subseries considered. Higher values of TSCS indicate that the two sub-series are separate components, i.e., a change in statistics is likely and indicates the presence of a change point.

Given this change space, we can build a multi-scale representation by varying the time scale δ over a range and summing up the resulting BIC curves. Formally, this can be defined as

$$MS-TSCS(t) = \sum_{\delta \in \Delta} TSCS(t, \delta) \quad (2)$$

where Δ is the set of all time scales for detecting time series components. In practice, we consider Δ to range from 10 time steps to 500 time steps. We then pass the curve from $MS-TSCS(t)$ through a low pass filter to extract peaks that provide possible time steps to segment the time series. We select peaks with high *saliency*, i.e., if it is more than two standard deviations from its neighbors. This is a common approach in statistics-based outlier detection literature [6, 25] and provides a good measure of temporal saliency for this problem. Given the temporal change locations, the *ideal* number of segments per dataset is computed as the average number of components across classes in the training set. We find that considering too few (or smaller) values in Δ will result in fewer segments and poor representations. Note that not all time series will have such components that are statistically separable. We use a uniform sampling approach to split the series into 15 equal segments in these cases. Empirically, segmenting chunks into more than 50 segments is not ideal and could degrade the performance, particularly on smaller datasets.

3.2 Capturing Signal Compositionality

The second step in our approach is to learn robust representations from the multi-scale components extracted using the MS-TSCS function defined in Section 3.1. Given the ideal number of segments K , the input sequence is *tokenized* $X_N = \{x_1, x_2, x_3, \dots, x_N\}$ into its constituent segments $\tilde{X}_K = \{\tilde{X}_1, \tilde{X}_2, \tilde{X}_3, \dots, \tilde{X}_k\}$. For capturing *compositional* representations, we then use a masked auto-encoding loss function [4] to train the encoding model (with parameters Θ). The masked auto-encoding loss randomly masks $M < k$ components and forces the encoder to independently predict the masked components by conditioning on the context provided by the unmasked components. Given the tokenized time series data $\tilde{X}_K = \{\tilde{X}_1, \tilde{X}_2, \tilde{X}_3, \dots, \tilde{X}_k\}$ and masked components $M = \{m_1, m_2, \dots, m_{|M|}\}$, the masked auto-encoding loss is

$$\mathcal{L}_{mae} = - \sum_{X_i \in \mathcal{C}} \log \prod_{m \in M} p(\tilde{X}_m | \tilde{X}_{K \setminus M}) \quad (3)$$

where $p(\tilde{X}_m | \tilde{X}_{K \setminus M})$ is the probability of predicting the randomly masked components in set $\{M\}$. This probability is computed as the mean squared error

over the masked component’s values. We use a bidirectional LSTM [20] as our encoder and ensure that the mask is bidirectional, i.e., the context for predicting the masked component is present on both sides of the mask. This masking procedure has successfully been used to train text-based [22] and image-based [43] encoders. We extend the formulation to univariate time series data. The hidden states of the forward and backward LSTM cells, h_t^f and h_t^b , respectively, are concatenated and used as the feature representation for time series classification, optimized by the cross-entropy loss ($\mathcal{L}_{CE}$). Hence, the overall objective function is given by

$$\mathcal{L}_{tot} = \lambda_1 \mathcal{L}_{mae} + \lambda_2 \mathcal{L}_{CE} \quad (4)$$

where λ_1 and λ_2 are tunable parameters that trade-off between the two losses. The values of λ_1 and λ_2 are varied according to a pre-set schedule to balance the representation learning capabilities from the self-supervised masked auto-encoding loss ($\mathcal{L}_{mae}$) and the discriminative, class-specific properties imbued by the supervised cross-entropy loss ($\mathcal{L}_{ce}$).

Implementation Details. We use a bidirectional LSTM model with a hidden size of 160 neurons, followed by a dense layer with 320 neurons, as our encoder architecture. The ReLU activation is used for all layers. All segmented components are padded as necessary to be equal in length. We use 5% of the training data for validation. We use the same pre-processing as previous work [44]. λ_1 and λ_2 are varied as follows: for the first 100 epochs, $\lambda_1 = 1$ and $\lambda_2 = 0$, then $\lambda_1 = 2$ and $\lambda_2 = 1$. The network is trained for 250 epochs or until convergence, i.e., the loss does not improve on the validation set. All experiments were conducted on a workstation server with a 32-core AMD ThreadRipper CPU, 128 GB RAM, and an NVIDIA RTX 3060.

4 Experimental Evaluation

In this section, we present the results from the experimental evaluation of the proposed approach. We begin with a discussion on the experimental setup, followed by the quantitative results, and conclude with a qualitative discussion on the representations learned by the approach.

4.1 Experimental Setup

Data. We evaluate the proposed approach on 85 datasets collated in the UCR time series archive [3]. It consists of univariate time series datasets collected from different sensors and domains such as health care, speech reorganization, and spectrum analysis, to name a few. The archive provides a comprehensive benchmark for evaluating time series classification models [36,44,15] across diverse datasets with varying characteristics. The number of classes in each dataset ranges from 2 to 6, the number of time steps per sample varies from 24 to 2709, and the number of training samples per dataset from 16 to 8926. Additionally, we evaluate the approach on 15 datasets with the longest timesteps from the UCR-85 [3] and the UCR-128 [10] datasets to evaluate its ability to capture

Table 1. Performance evaluation of the proposed approach with state-of-the-art approaches on 85 datasets from the UCR time series archive [3,10]

Approach	Ensemble?	Backbone	Accuracy
TST [46]	✗	Transformer	64.901
MDCNN [48]	✗	CNN	68.551
TWIESN [37]	✗	RNN	68.636
TS-Encoder [34]	✗	CNN	71.909
Time-CNN [47]	✗	CNN	72.284
DTW [7]	✗	Distance	74.040
TS-TCC [14]	✗	CNN-Transformer	77.764
TNC [38]	✗	Bi-RNN	77.896
PF [29]	✗	Distance	80.419
T-Loss [15]	✗	Dilated CNN	80.482
BOSS [31]	✗	Bag of Symbols	81.019
FCN [42]	✗	CNN	81.634
ResNet [42]	✗	CNN	82.201
ST [19]	✗	Shapelets	82.236
TS2Vec [44]	✗	Dilated CNN	82.934
Ours	✗	Bi-RNN	83.309
TS-CHIEF [35]	✓	Bag of Symbols	84.641
HIVE-COTE [27]	✓	Multiple	84.714
OS-CNN [36]	✓	CNN	84.774
ROCKET [12]	✓	CNN	85.077

robust representations from time series with longer duration. We use the official train and test splits on all datasets for a fair comparison with prior works. Average accuracy across all datasets is used to quantify the performance on the UCR time series archive. Code and performance for baselines are obtained from publicly available implementations of prior works [36,44].

Baselines. We compare against state-of-the-art univariate time series classification models, which use different representation learning backbones and propose robust learning methods to account for high intra-class variation common in time series data. Chiefly, we compare against models with CNN backbones [44,48,47,34,15,42], transformer backbones [46,14], RNN backbones [37,38], and other hand-crafted features such as shapelet transforms [19], distance-based metrics [29,7], and bag-of-symbols [31]. We also compare against ensembles [36,27,35,12], which explicitly capture representations at multiple time scales, which can require additional overhead for training.

4.2 Quantitative Evaluation

We present the performance of the approach on the UCR-85 archive in Table 1. We outperform other approaches on the benchmark while offering competitive performance to those designed to work in an ensemble. Interestingly, most state-of-the-art techniques are based on CNNs, with much effort spent finding optimal

Table 2. Performance on 15 longest sequence time series data from the UCR Archives [3], compared against state-of-the-art models with different backbones.

Backbone→ Dataset ↓	CNN		Transformer		Bi-RNN	
	TS2Vec [44]	OS-CNN [36]	TS-TCC [14]	TST [46]	TNC [38]	Ours
Rock	70.00	55.00	60.00	68.00	58.00	70.00
HandOutlines	92.20	92.95	72.40	73.50	93.00	94.05
HouseTwenty	91.60	94.87	79.00	81.50	78.20	<u>92.44</u>
InlineSkate	<u>41.50</u>	42.92	34.70	28.70	37.80	41.09
EthanolLevel	46.80	73.08	48.60	26.00	42.40	87.00
SemgHandSubjectCh2	95.10	71.84	75.30	48.40	77.10	<u>91.56</u>
SemgHandGenderCh2	96.30	85.61	83.70	72.50	88.20	<u>89.33</u>
SemgHandMovementCh2	86.00	56.62	61.30	42.00	59.30	<u>78.22</u>
EOGHorizontalSignal	53.90	63.97	40.10	37.30	44.20	<u>57.73</u>
EOGVerticalSignal	50.30	47.76	37.60	29.80	39.20	51.10
Haptics	52.60	<u>51.01</u>	39.60	35.70	47.40	50.32
Mallat	91.40	<u>96.38</u>	92.20	71.30	87.10	97.10
MixedShapesRegularTrain	91.70	96.09	85.50	87.90	91.10	<u>93.69</u>
MixedShapesSmallTrain	86.10	91.79	73.50	82.80	81.30	<u>87.96</u>
StarLightCurves	96.90	<u>97.51</u>	96.70	94.90	96.80	97.78
Average	<u>76.16</u>	74.49	65.35	58.69	68.07	78.62

receptive field sizes for learning robust features at multiple timescales. Sequence-based approaches, such as those based on Transformers and RNNs, have struggled in this benchmark, mostly due to the limited training examples in many datasets. We, however, significantly outperform other sequence-based approaches and provide improvements of almost 5.5% in absolute accuracy points over the closest RNN-based approach (TNC [38]). It also provides the best performance (out of non-ensemble approaches) on 17 datasets (also called *wins* in prior literature [36]) out of the 85 benchmark datasets. Additionally, it has an average rank of 5.35, performing competitively with other non-ensemble approaches. Ensemble models outperform all non-ensemble models by explicitly modeling sequential data by representing the sequential data at different time scales. However, they introduce additional overhead for handcrafting and fine-tuning multiple models.

Performance on longer sequence data. While the overall UCR-85 archive performance is excellent, we also examine the ability of the proposed approach to capture long-range dependencies when presented with time series data of longer durations. We select a subset of the UCR-128 archive, which contains additional datasets of longer duration. Specifically, we select 15 datasets with more than 1000 timesteps per sample without incomplete data. Table 2 presents a summary of the results. As can be seen, we provide competitive performance with top-performing baselines with different backbone architectures. We have an average accuracy of 78.62%, an average rank of 1.72, and provide “wins” in 6 out of the 15 long sequence datasets. It significantly improves over transformer-based (TS-TCC and TST) and RNN-based (TNC) baselines, which are trained to specifically model longer sequences through specialized training procedures such as contrastive learning. These results indicate the approach can capture robust representations from long sequences without complex ensemble processing.

Table 3. Ablation studies on the UCR-85 archive [3] to assess the impact of each component on the overall performance.

Backbone	MS-TSCS	$\mathcal{L}_{mae}$	$\mathcal{L}_{CE}$	Accuracy
Bi-LSTM	✓	✓	✓	83.31
Bi-LSTM	✗	✓	✓	73.68
Bi-LSTM	✓	✗	✓	75.31
Bi-LSTM	✓	✓	✗	74.28
Bi-LSTM	✗	✗	✓	68.33
Bi-RNN	✓	✓	✓	81.54
Uni-LSTM	✓	✓	✓	79.55

Constrained Hardware Requirements Our approach is designed to be simple and lightweight for use in settings with constrained training requirements, such as time and space budgets (i.e., limited training time, constrained hardware requirements, and limiting the number of parameters). Our model achieves competitive performance with 440k parameters and converges training on all datasets in 4 hours (on average over ten runs). For comparison, the current non-ensemble state-of-the-art approaches, TS2Vec (637k parameters) and ResNet (479k parameters), have more parameters and take longer to converge on a constrained hardware setup (32-core AMD ThreadRipper and NVIDIA RTX 3060). Similarly, on average, the BIC-based tokenization process (Section 1) takes 500 ms for a sequence of 1000 data points, running in a single-threaded CPU-only application while having significantly less overhead for storing the components compared with other approaches.

Ablation Studies. We systematically examine the impact of each module and summarize the results in Table 3. Specifically, we assess the effects of the multi-scale component discovery module (Section 3.1) and the choice of encoder model (Section 3.2). Removing the component discovery model and using a fixed number of components for all datasets (set to 25, the median number of components across datasets) significantly hurts the performance. We also evaluate the strength of the learned representations by using a kNN instead of end-to-end training by removing $\mathcal{L}_{CE}$ from Equation 4. While the loss in performance is expected, it does perform decently, indicating that the unsupervised loss function helps learn robust features. Removing $\mathcal{L}_{mae}$ results in significantly worse performance. Using bidirectional LSTMs instead of unidirectional LSTMs helps capture context and provides a more robust performance across all 85 datasets in the UCR archive.

5 Extension to Unsupervised Time Series Segmentation

In addition to evaluating the performance of our approach on time series classification, we assess the quality of the components obtained through the BIC-based segmentation (Section 3.1) by evaluating it on the time series segmentation task [16]. The goal of time series segmentation is to identify natural segments

Table 4. Evaluation of the BIC-based tokenization approach on the **time series segmentation** task [16].

Approach	Learning Phase?	Pre-Defined Window?	Mean Covering
BinSeg	✗	✓	52.4 ± 30.6
PELT	✗	✓	50.4 ± 30.0
Window	✗	✓	53.8 ± 12.9
BOCD	✗	✓	55.5 ± 14.4
ESPRESSO	✗	✓	58.0 ± 15.8
Ours	✗	✗	72.7 ± 12.5
FLOSS	✓	✓	79.0 ± 17.2
ClaSP	✓	✗	79.8 ± 20.4
Ours	✗	✓	78.3 ± 12.9

caused by *change points* in sequential data where there are sudden changes in statistical properties of the time series due to changes in events captured by the data. For example, these changes could point to transitions between actions performed by a subject. The UTSA benchmark [16] introduces a set of 32 datasets derived from the UCR archive [3] and provides human-annotated segments of datasets across 16 different use cases from biological, mechanical, and synthetic processes. Each use case in the benchmark contains, on average, 2 to 3 segments derived from real, semi-synthetic, and artificial changes and provides a considerable challenge for unsupervised time series segmentation.

We use the components discovered using the multi-scale change space model as segments and assess the quality of the segmentations on the UTSA benchmark. We compare against a variety of baselines such as BinSeg [33], PELT [23], Window [40], BOCD [1], ESPRESSO [11], FLOSS [17], and ClaSP [32], which represent the commonly used state-of-the-art unsupervised segmentation approaches. We use the mean covering with standard deviation as a metric to quantify the performance of the approaches. Based on the Jaccard index, the covering score provides a weighted overlap between the ground truth and the predicted segments. Higher values indicate better alignment between the predicted and the ground truth segments. We report results from the implementations from ClaSP [32] for a fair comparison and consistent experimental setup.

Table 4 summarizes the results. We significantly outperform other non-learning-based approaches that require a pre-defined period size (temporal window) corresponding to the ideal time scale at which the change points can be detected reliably. This value is often domain-dependent and requires extensive handcrafting (of architecture or features) to capture, especially in time series classification and segmentation. Our approach can automatically search for this using the multi-scale change space and considers change points at different temporal granularities. When given this optimal window, we establish the change space at this time scale and perform segmentation. As can be seen, we perform competitively with learning-based approaches and further widen the gap with the non-learning-




Fig. 2. Qualitative visualization of (a) a successful segmentation and (b) unsuccessful segmentation on the GreatBarbet2 and SuddenCardianDeath1 datasets, respectively. The first row shows ground truth segments, and the second shows predicted segments.

based approaches. Interestingly, we perform exceptionally well without the optimal time scale, indicating that the multi-scale change space captures the change points at time scales approaching the ideal scale. Some example segmentations are shown in Figure 2, where it can be seen that our approach can segment signals into their components without training and supervision. Although it over-segments in some instances, the segments are statistically meaningful, are captured at multiple time scales, and do not always correspond to the ground change points extracted at a single time scale. For example, in Figure 2(b), we see that over-segmentation occurs during periods of intense changes and captures fine-grained change points but has excellent coverage during stable regions on either side of this rapidly changing segment. Note that our approach detects the temporal components in a time-scale and class-agnostic manner and does not have access to the ideal time scale at which the ground truth is annotated. Despite this over-segmentation, it allows us to capture robust features for classification.

6 Discussion and Future Work

In this work, we presented a novel multi-scale change-space approach to discover temporal components in univariate time series data and provide an intuitive way to tokenize time series data using statistical measures. Given these components, we learn compositional representations using sequence-based encoders by training the model as a masked, denoising auto-encoder. Evaluation on 85 publicly available datasets on the benchmark UCR-85 archive demonstrates its effectiveness in learning robust representations. Additional experiments on segmentation benchmarks demonstrate that the detected components are highly correlated with naturally occurring segments found in time series data. We aim to extend this formulation to capture part-whole hierarchies for learning *hier-*

archical compositional representations from multi-modal and multi-variate time series data with longer temporal durations.

Acknowledgements. This work was supported by the U.S. National Science Foundation Grant IIS 2348689 and IIS 2348690 and U.S. Department of Agriculture Grant 2023-69014-39716-1030191.

References

1. Adams, R.P., MacKay, D.J.: Bayesian online changepoint detection. arXiv preprint arXiv:0710.3742 (2007)
2. Bagnall, A., Dau, H.A., Lines, J., Flynn, M., Large, J., Bostrom, A., Southam, P., Keogh, E.: The uea multivariate time series classification archive, 2018. arXiv preprint arXiv:1811.00075 (2018)
3. 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 and Knowledge Discovery* **31**, 606–660 (2017)
4. Bao, H., Dong, L., Wei, F., Wang, W., Yang, N., Liu, X., Wang, Y., Gao, J., Piao, S., Zhou, M., et al.: Unilmv2: Pseudo-masked language models for unified language model pre-training. In: *International Conference on Machine Learning*. pp. 642–652. PMLR (2020)
5. Chen, H., Chu, L.: Graph-based change-point analysis. *Annual Review of Statistics and Its Application* **10**, 475–499 (2023)
6. Chen, S., Gopalakrishnan, P., et al.: Speaker, environment and channel change detection and clustering via the bayesian information criterion. In: *DARPA Broadcast News Transcription and Understanding Workshop*. vol. 8, pp. 127–132. Citeseer (1998)
7. Chen, Y., Hu, B., Keogh, E., Batista, G.E.: Dtw-d: time series semi-supervised learning from a single example. In: *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. pp. 383–391 (2013)
8. Cui, Z., Chen, W., Chen, Y.: Multi-scale convolutional neural networks for time series classification. arXiv preprint arXiv:1603.06995 (2016)
9. Dash, S., Chakraborty, C., Giri, S.K., Pani, S.K.: Intelligent computing on time-series data analysis and prediction of covid-19 pandemics. *Pattern Recognition Letters* **151**, 69–75 (2021)
10. Dau, H.A., Bagnall, A., Kamgar, K., Yeh, C.C.M., Zhu, Y., Gharghabi, S., Ratanamahatana, C.A., Keogh, E.: The ucr time series archive. *IEEE/CAA Journal of Automatica Sinica* **6**(6), 1293–1305 (2019)
11. Deldari, S., Smith, D.V., Sadri, A., Salim, F.: Espresso: Entropy and shape aware time-series segmentation for processing heterogeneous sensor data. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* **4**(3), 1–24 (2020)
12. Dempster, A., Petitjean, F., Webb, G.I.: Rocket: exceptionally fast and accurate time series classification using random convolutional kernels. *Data Mining and Knowledge Discovery* **34**(5), 1454–1495 (2020)
13. Draayer, E., Cao, H., Hao, Y.: Reevaluating the change point detection problem with segment-based bayesian online detection. In: *ACM International Conference on Information & Knowledge Management*. pp. 2989–2993 (2021)
14. Eldele, E., Ragab, M., Chen, Z., Wu, M., Kwok, C.K., Li, X., Guan, C.: Time-series representation learning via temporal and contextual contrasting. arXiv preprint arXiv:2106.14112 (2021)

15. Franceschi, J.Y., Dieuleveut, A., Jaggi, M.: Unsupervised scalable representation learning for multivariate time series. *Advances in Neural Information Processing Systems* **32** (2019)
16. Gharghabi, S., Ding, Y., Yeh, C.C.M., Kamgar, K., Ulanova, L., Keogh, E.: Matrix profile viii: domain agnostic online semantic segmentation at superhuman performance levels. In: *IEEE International Conference on Data Mining (ICDM)*. pp. 117–126. IEEE (2017)
17. Gharghabi, S., Yeh, C.C.M., Ding, Y., Ding, W., Hibbing, P., LaMunion, S., Kaplan, A., Crouter, S.E., Keogh, E.: Domain agnostic online semantic segmentation for multi-dimensional time series. *Data Mining and Knowledge Discovery* **33**, 96–130 (2019)
18. Harchaoui, Z., Vallet, F., Lung-Yut-Fong, A., Cappé, O.: A regularized kernel-based approach to unsupervised audio segmentation. In: *IEEE International Conference on Acoustics, Speech and Signal Processing*. pp. 1665–1668. IEEE (2009)
19. Hills, J., Lines, J., Baranauskas, E., Mapp, J., Bagnall, A.: Classification of time series by shapelet transformation. *Data mining and Knowledge Discovery* **28**, 851–881 (2014)
20. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Computation* **9**(8), 1735–1780 (1997)
21. Kawahara, Y., Sugiyama, M.: Sequential change-point detection based on direct density-ratio estimation. *Statistical Analysis and Data Mining: The ASA Data Science Journal* **5**(2), 114–127 (2012)
22. Kenton, J.D.M.W.C., Toutanova, L.K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: *NAACL-HLT*. pp. 4171–4186 (2019)
23. Killick, R., Fearnhead, P., Eckley, I.A.: Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association* **107**(500), 1590–1598 (2012)
24. Kiranyaz, S., Avci, O., Abdeljaber, O., Ince, T., Gabbouj, M., Inman, D.J.: 1d convolutional neural networks and applications: A survey. *Mechanical Systems and Signal Processing* **151**, 107398 (2021)
25. Krishnan, R., Sarkar, S.: Detecting group turn patterns in conversations using audio-video change scale-space. In: *International Conference on Pattern Recognition*. pp. 137–140. IEEE (2010)
26. Laptev, I., Lindeberg, T.: A multi-scale feature likelihood map for direct evaluation of object hypotheses. In: *International Conference on Scale-Space Theories in Computer Vision*. pp. 98–110. Springer (2001)
27. Lines, J., Taylor, S., Bagnall, A.: Time series classification with hive-cote: The hierarchical vote collective of transformation-based ensembles. *ACM Transactions on Knowledge Discovery from Data* **12**(5), 1–35 (2018)
28. Locatello, F., Weissenborn, D., Unterthiner, T., Mahendran, A., Heigold, G., Uszkoreit, J., Dosovitskiy, A., Kipf, T.: Object-centric learning with slot attention. *Advances in Neural Information Processing Systems* **33**, 11525–11538 (2020)
29. Lucas, B., Shifaz, A., Pelletier, C., O’Neill, L., Zaidi, N., Goethals, B., Petitjean, F., Webb, G.I.: Proximity forest: an effective and scalable distance-based classifier for time series. *Data Mining and Knowledge Discovery* **33**(3), 607–635 (2019)
30. Ramnath, V.L., Katkoori, S.: A smart iot system for continuous sleep state monitoring. In: *2020 IEEE 63rd International Midwest Symposium on Circuits and Systems (MWSCAS)*. pp. 241–244. IEEE (2020)
31. Schäfer, P.: The boss is concerned with time series classification in the presence of noise. *Data Mining and Knowledge Discovery* **29**, 1505–1530 (2015)

32. Schäfer, P., Ermshaus, A., Leser, U.: Clasp-time series segmentation. In: ACM International Conference on Information & Knowledge Management. pp. 1578–1587 (2021)
33. Sen, A., Srivastava, M.S.: On tests for detecting change in mean. *The Annals of statistics* pp. 98–108 (1975)
34. Serra, J., Pascual, S., Karatzoglou, A.: Towards a universal neural network encoder for time series. In: International Conference of the Catalan Association for Artificial Intelligence. pp. 120–129 (2018)
35. Shifaz, A., Pelletier, C., Petitjean, F., Webb, G.I.: Ts-chief: a scalable and accurate forest algorithm for time series classification. *Data Mining and Knowledge Discovery* **34**(3), 742–775 (2020)
36. Tang, W., Long, G., Liu, L., Zhou, T., Blumenstein, M., Jiang, J.: Omni-scale cnns: a simple and effective kernel size configuration for time series classification. In: International Conference on Learning Representations (2021)
37. Tanisaro, P., Heidemann, G.: Time series classification using time warping invariant echo state networks. In: IEEE International Conference on Machine Learning and Applications. pp. 831–836. IEEE (2016)
38. Tonekaboni, S., Eytan, D., Goldenberg, A.: Unsupervised representation learning for time series with temporal neighborhood coding. In: International Conference on Learning Representations (2020)
39. Trehan, S., Aakur, S.N.: Towards active vision for action localization with reactive control and predictive learning. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 783–792 (2022)
40. Truong, C., Oudre, L., Vayatis, N.: Selective review of offline change point detection methods. *Signal Processing* **167**, 107299 (2020)
41. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *Advances in Neural Information Processing Systems* **30** (2017)
42. Wang, Z., Yan, W., Oates, T.: Time series classification from scratch with deep neural networks: A strong baseline. In: International Joint Conference on Neural Networks. pp. 1578–1585. IEEE (2017)
43. Xie, Z., Zhang, Z., Cao, Y., Lin, Y., Wei, Y., Dai, Q., Hu, H.: On data scaling in masked image modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10365–10374 (2023)
44. Yue, Z., Wang, Y., Duan, J., Yang, T., Huang, C., Tong, Y., Xu, B.: Ts2vec: Towards universal representation of time series. In: AAAI Conference on Artificial Intelligence. vol. 36, pp. 8980–8987 (2022)
45. Zacks, J.M., Tversky, B.: Event structure in perception and conception. *Psychological bulletin* **127**(1), 3 (2001)
46. Zerveas, G., Jayaraman, S., Patel, D., Bhamidipaty, A., Eickhoff, C.: A transformer-based framework for multivariate time series representation learning. In: ACM SIGKDD Conference on Knowledge Discovery & Data Mining. pp. 2114–2124 (2021)
47. Zhao, B., Lu, H., Chen, S., Liu, J., Wu, D.: Convolutional neural networks for time series classification. *Journal of Systems Engineering and Electronics* **28**(1), 162–169 (2017)
48. Zheng, Y., Liu, Q., Chen, E., Ge, Y., Zhao, J.L.: Exploiting multi-channels deep convolutional neural networks for multivariate time series classification. *Frontiers of Computer Science* **10**, 96–112 (2016)