

# Semi-Supervised Sequence Classification through Change Point Detection

Nauman Ahad, Mark A. Davenport

School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA  
 {nahad3, mdav}@gatech.edu

## Abstract

Sequential sensor data is generated in a wide variety of real-world applications. A fundamental machine learning challenge involves learning effective classifiers for such sequential data. While deep learning has led to impressive performance gains in recent years within domains such as speech, this has relied on the availability of large datasets of sequences with high-quality labels. In many applications, however, the associated class labels are often extremely limited, with precise labelling/segmentation being too expensive to perform in a high volume. However, large amounts of unlabeled data may still be available. In this paper we propose a novel framework for semi-supervised learning in such contexts. In an unsupervised manner, change point detection methods can be used to identify instances where classes change within a sequence. We show that change points provide examples of similar/dissimilar pairs of sequences which, when coupled with class labels, can be used in a semi-supervised classification setting. Pairs from labels and change points are used by a neural network to learn improved representations for classification. We provide extensive synthetic simulations and show that the learned representations are better than those learned through an autoencoder and obtain improved results on simulations and human activity recognition datasets.

## Introduction

As devices ranging from smart watches to smart toasters are equipped with ever more sensors, machine learning problems involving sequential data are becoming increasingly ubiquitous. Sleep tracking, activity recognition and characterization, and machine health monitoring are just a few applications where machine learning can be applied to sequential data. In recent years, deep networks have been widely used for such tasks as these networks are able to automatically learn suitable representations, helping them achieve state-of-the-art performance (Wang et al. 2019). However, such methods typically require large, accurately labeled training datasets in order to obtain these results. Unfortunately, especially in the context of sequential data, it is often the case that despite the availability of huge amounts of unlabeled data, labeled data is often scarce and expensive to obtain.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In such settings, *semi-supervised* techniques can provide significant advantages over traditional supervised techniques. Over the past decade, there have been great advances in semi-supervised learning methods. Impressive classification performance – particularly in the fields of computer vision – has been achieved by using large amounts of unlabeled data on top of limited labeled data. However, despite these advances, there has been comparatively much less work on semi-supervised classification of sequential data.

A key intuition that most semi-supervised learning methods share is that the data should (in the right representation) exhibit some kind of clustering, where different classes correspond to different clusters. In the context of sequential data, the equivalent assumption is that data segments within a sequence corresponding to different classes should map to distinct clusters. In the context of sequential data, the challenge is that exploiting this clustering would require the sequence to be appropriately segmented, but segment boundaries are generally unknown *a priori*. If the start/end points of each segment were actually known, it would be much easier to apply traditional semi-supervised learning methods.

In this paper, we show that standard (unsupervised) *change point detection* algorithms provide a natural and useful approach to segmenting an unlabeled sequence so that it can be more easily exploited in a semi-supervised context. Specifically, change point-detection algorithms aim to identify instances in a sequence where the data distribution changes (indicating an underlying class change). We show that the resulting change points can be leveraged to learn improved representations for semi-supervised learning.

We propose a novel framework for semi-supervised sequential classification using change point detection. We first apply unsupervised change point detection to the unlabeled data. We assume that segments between two change points belong to the same distribution and should be classified similarly, whereas adjacent segments which are on opposite sides of a change point belong to different distributions and should be classified differently. These similar/dissimilar pairs, derived from change points, can then be combined with similar/dissimilar pairs derived from labeled data. We use these combined similar/dissimilar constraints to train a neural network that preserves similarity/dissimilarity. The learned representation can then be fed into a multilayer feedforward network trained via existing semi-supervised techniques.

We show that this approach leads to improved results compared to sequential auto-encoders in a semi-supervised setting. We show that even if the final classifier is trained using standard supervised techniques that ignore the unlabeled data, the learned representations (which utilize both label and unlabeled data pairs) result in competitive performance, indicating the value of incorporating change points to learning improved representations. The proposed method is completely agnostic with respect to the change point detection procedure to be used – any detection procedure can be used as long as it does well in detecting changes.

Our main contribution is to show that pairwise information generated via change points helps neural networks achieve improved classification results in settings with limited labeled data. This, to the best of our knowledge, is the first work to recognize the utility of change points within the context of semi-supervised sequence classification. The proposed method should not be considered a substitute for existing semi-supervised methods, but should be taken as a complementary procedure that produces representations which are better suited for existing semi-supervised methods.

## Related Work

The fundamental idea of semi-supervised learning is that unlabeled data contains useful information that can be leveraged to more efficiently learn from a small subset of labeled data. For example, in the context of classification, an intuitive justification for why this might be possible might involve an implicit expectation that instances belonging to different classes will map to different clusters. More concretely, most semi-supervised approaches make assumptions on the data such as: that instances corresponding to different classes lie on different submanifolds, that class boundaries are smooth, or that class boundaries pass through regions of low data density (Van Engelen and Hoos 2020).

Perhaps the simplest semi-supervised learning method is to use transductive methods to learn a classifier on the unlabeled data and then assign “pseudo labels” to some or all of the unlabeled data, which can be used together with the labeled data to retrain the classifier. Transductive SVMs and graphical label propagation are examples of such methods (Joachims 1999; Zhu, Ghahramani, and Lafferty 2003). See (Zhu 2005) for a survey of such methods. However, such *self-training* semi-supervised methods struggle when the initial model trained from limited labels is poor.

A more common approach to semi-supervised learning is to employ methods that try to learn class boundaries that are smooth or pass through areas of low data density (Oliver et al. 2018). Entropy regularization can be used to encourage class boundaries to pass through low density regions (Grandvalet and Bengio 2005). Consistency-based methods such as denoising autoencoders, ladder networks (Rasmus et al. 2015) and the  $\pi$  method (Laine and Aila 2016) attempt to learn smooth class boundaries by augmenting the data. Specifically, unlabeled instances can be perturbed by adding noise, and while both the original and perturbed instances are unlabeled, we can ask that they both be assigned the same class. This approach is particularly effective in computer vision tasks, where rather than using only noise pertur-

bations, we can exploit class-preserving augmentations such as rotation, mirroring, and other transformations (Berthelot et al. 2019). By enforcing the classifier to produce the same labels for original and transformed images, decision boundaries are encouraged to be smooth, leading to good generalization.

Unfortunately, due to a lack of natural segmentation and the difficulty of defining class-preserving transformations, there has been comparatively little work on semi-supervised classification of sequences. Most prior work (e.g., (Dai and Le 2015; Rasmus et al. 2015) ) use sequential autoencoders (or their variants) as a consistency-based method to learn representations that lead to improved classification performance. Such autoencoders have been exploited successfully in the context of semi-supervised classification for human activity recognition (Zeng et al. 2017). However, while such consistency-based approaches do encourage smooth class boundaries, they do not necessarily promote the kind of clustering behavior that we need in cases where there are extremely few labels available.

An alternative approach that more explicitly separates different classes involves learning representations that directly incorporate pairwise similarity information about different instances. One example of this approach is *metric learning* – as an early example, (Xing et al. 2003) showed that improved classification could be achieved by learning a Mahalanobis distance using pairwise constraints based on class membership. The learned metric leads to a representation in which different classes map to different clusters. A similar approach learns a more general non-linear metric to encourage the formation of clusters while adhering to the provided pairwise constraints (Baghshah and Shouraki 2009). Neural networks such as Siamese (Koch, Zemel, and Salakhutdinov 2015) and Triplet networks also learn representations from available similar/dissimilar pairs. In (Hsu and Kira 2016) it was shown that such similar/dissimilar pairs (obtained from labeled data) can be used for clustering data where each cluster belongs to a different class in the dataset.

Our approach is similar in spirit to that of (Hsu and Kira 2016). While this prior work used pairwise similarity constraints derived from labeled images to learn clustered representations, our goal is to apply this idea in the semi-supervised context. At the core of our approach is the observation that pairwise similarity constraints on sequential data can be derived through unsupervised methods. Specifically, change point detection can be used to identify points within a sequence corresponding to distribution shifts, which can then be used to obtain pairwise similarity constraints. When the availability of labeled data is limited, this can be a valuable source of additional information.

## Proposed Method

### Change Point Detection

Given a sequence  $X : x_1, \dots, x_N$  of  $N$  vectors  $x_i \in \mathbb{R}^D$ , the first step in our procedure is to detect *all* change points within  $X$  in an unsupervised way. Note that this is a different problem than *quickest* change detection, where only a *single* change point is to be detected in the fastest possible manner.

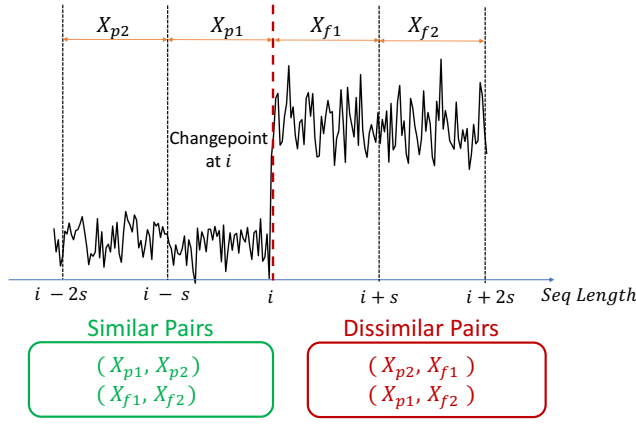


Figure 1: Using change points to generate similar and dissimilar pairs of size  $s$ .

To detect a change at a point  $i$  in the sequence, two consecutive length- $w$  windows ( $X_p^i$  and  $X_f^i$ ) are first formed:

$$X_p^i = x_{i-1}, x_{i-2} \dots x_{i-w} \quad X_f^i = x_i, x_{i+1} \dots x_{i+w}.$$

A change statistic,  $m_i$ , is then computed via some function that quantifies the difference between the distributions generating  $X_p^i$  and  $X_f^i$ . If  $m_i$  is greater than a specified constant  $\tau$ , a change point is detected at the point  $i$ .

As one example, many change point detection procedures assume a parametric form on the distributions generating  $X_p^i$  and  $X_f^i$ . In this case, the distribution parameters ( $\hat{\theta}_p^i$  and  $\hat{\theta}_f^i$ ) can be estimated from  $X_p^i$  and  $X_f^i$  via, e.g., maximum likelihood estimation. Given these parameter estimates, a symmetrical KL-divergence can be used to quantify the difference between the distributions (Liu et al. 2013):

$$m_i = \text{KL}(\hat{\theta}_p^i, \hat{\theta}_f^i) + \text{KL}(\hat{\theta}_f^i, \hat{\theta}_p^i). \quad (1)$$

More commonly in practice, the underlying distributions generating the sequence are unknown. In this case, non-parametric techniques can be used to estimate the difference between the distributions of  $X_p^i$  and  $X_f^i$ . One such approach uses the *maximum mean discrepancy* (MMD) as a change statistic (Gretton et al. 2012). The MMD has been used to identify change points in (Li et al. 2015) and (Chang et al. 2019). The MMD statistic is given below, where  $K_{a,-b}^i := k(x_{i+a}, x_{i-b})$  represents a kernel-based measure of the similarity between  $x_{i+a}$  and  $x_{i-b}$ :

$$\begin{aligned} m_i &= \text{MMD}(X_f^i, X_p^i) \\ &= \frac{1}{\binom{w}{2}} \sum_{\substack{a,b=1 \\ a \neq b}}^w 2(K_{a,b}^i + K_{-a,-b}^i) + \frac{1}{w^2} \sum_{a,b=1}^w 2K_{a,-b}^i. \end{aligned}$$

Throughout this paper, MMD with a radial basis function kernel is used to detect change points unless otherwise specified. However, we again emphasize that any change point detection method can be used as long as it performs well in identifying changes points.

The labeled data can be used to set the change point detection threshold  $\tau$  and the window size  $w$  to balance between false and missed change points. While we simply fix these parameters in advance using labeled data, these could also be considered as tuning parameters whose values can be set based on performance on a hold-out validation dataset.

## Pairwise Constraints via Change Point Detection

Equipped with the detected change points, similar and dissimilar pairs of sub-sequences can be obtained in an unsupervised manner as shown in Figure 1. The idea is to form four consecutive non-overlapping sub-sequences. The first two sub-sequences ( $X_{p1}, X_{p2}$ ) both occur before the change point. Since the change point detection algorithm did not determine that there was a change point in the combined segment of ( $X_{p1}, X_{p2}$ ), we assume these two segments are generated by the same distribution and should be classified similarly. Similarly, the last two sub-sequences ( $X_{f1}, X_{f2}$ ) both occur after the change point and are also taken as a similar pair. In contrast, the segments on opposite sides of the change point have been identified as having different underlying distributions. In order to lead to a balanced distribution of similar/dissimilar pairs, we only use the constraints that ( $X_{f1}, X_{p2}$ ) and ( $X_{f2}, X_{p1}$ ) should be classified differently. Each of the subsequences above is chosen to be of a fixed length  $s$  (determined by the spacing between change points).

The proposed method assumes that data distributions belonging to different classes do not change. This assumption is also made by most semi-supervised learning methods.

## Clustered Representation via Pairwise Constraints

Using the approach described above, we can obtain similarity constraints from the unlabeled data. We can also obtain such constraints from labeled data via the assumption that sub-sequences corresponding to the same (different) class labels are similar (dissimilar) respectively. We can represent these as a set  $\mathcal{P}_S$  consisting of sub-sequence pairs ( $X_1, X_2$ ) that are similar and a set  $\mathcal{P}_D$  of dissimilar pairs. For compactness, we use the notation  $P = (X_1, X_2)$  to refer to a sub-sequence pair belonging to  $\mathcal{P}_S$  or  $\mathcal{P}_D$ .

These sub-sequences are then fed into a 1D temporal convolutional neural network (Bai, Kolter, and Koltun 2018), as illustrated in Figure 2. The neural network consists of 6 convolutional layers (or 3 temporal blocks as defined by (Bai, Kolter, and Koltun 2018)) followed by 1 linear layer. We use a RELU activation function after every convolutional layer. We choose this architecture because the dilated filter structure leads to improved performance at classifying time series while being less computationally expensive than recurrent networks such as RNNs and LSTMs, although our framework could also easily accommodate either of these alternate network architectures.

Each instance  $x_i$ , in the input sub-sequence  $X$ , is passed through the neural network where the final linear layer transforms the output from the last convolutional layer into  $\mathbb{R}^C$ , where  $C$  is the number of classes. A softmax function is then applied to obtain the empirical distribution  $f_\theta(x_i)$  for each instance  $x_i$ . For a length- $N$  sequence  $X$ , we define the mean

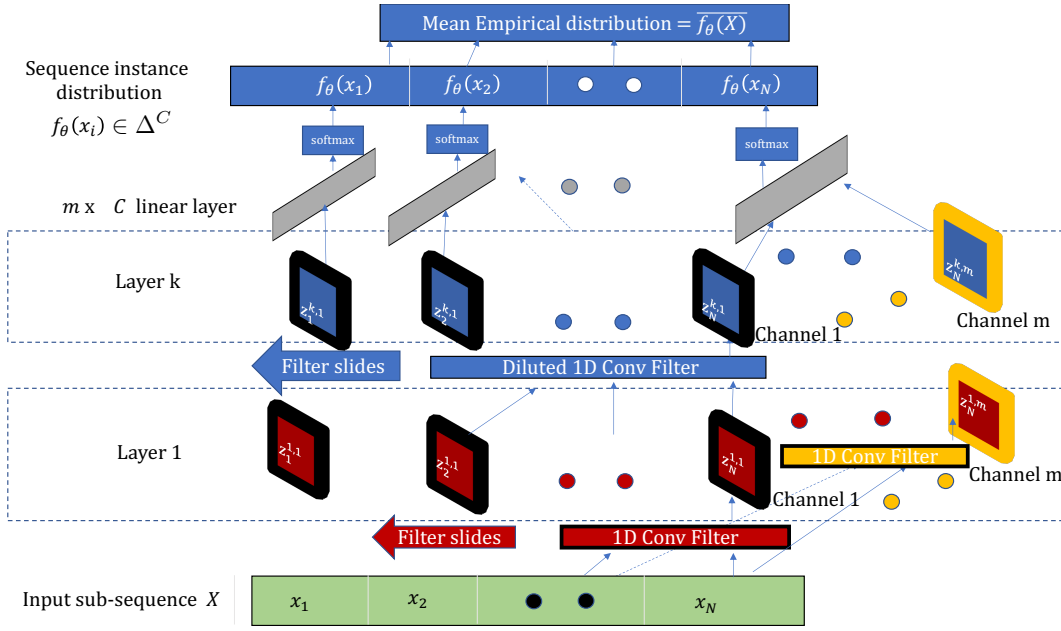


Figure 2: Neural network diagram ( $f_\theta$ ) for learning representations.

empirical distribution as:

$$\overline{f_\theta(X)} = \frac{1}{N} \sum_{i=1}^N f_\theta(x_i).$$

We then compute the KL divergence between the mean empirical distributions for each sub-sequence within a pair  $P = (X_1, X_2)$ . Our loss function is constructed applying a hinge loss (with margin parameter  $\rho$ ) to this KL divergence:

$$h_\theta(P) = \begin{cases} \text{KL}(\overline{f_\theta(X_1)}, \overline{f_\theta(X_2)}) & P \in \mathcal{X}_S, \\ \rho - \text{KL}(\overline{f_\theta(X_1)}, \overline{f_\theta(X_2)}) & P \in \mathcal{X}_D. \end{cases}$$

The network is then trained according to the loss function:

$$\mathcal{L}_R(\theta) = \frac{1}{|\mathcal{P}_L|} \sum_{P \in \mathcal{P}_L} h_\theta(P) + \frac{\lambda_R}{|\mathcal{P}_U|} \sum_{P \in \mathcal{P}_U} h_\theta(P).$$

Here,  $\mathcal{P}_L$  and  $\mathcal{P}_U$  denote the sets of sub-sequence pairs in  $\mathcal{P}_S \cup \mathcal{P}_D$  formed from the labeled and unlabeled data, respectively, and  $\lambda_r$  is a tuning parameter which controls the influence of the unsupervised part of the loss function.

### Training a Classifier

Once trained, the network  $f_\theta$  is fixed. The mean empirical distribution for an input sub-sequence  $X$ ,  $\overline{f_\theta(X)}$ , can then be used as a representation of  $X$  that can serve as input to classifier network  $f_\psi$ . We use a 2-layer feedforward neural network followed by a softmax function to obtain a distribution over the different classes. Labeled as well as unlabeled sub-sequences (which correspond to the generated pairs from change points) are passed through this classification network. Since the learned representations encourage unlabeled data points to cluster around provided labeled data points, known semi-supervised methods can be also used to

incorporate unlabeled data while training  $f_\psi$ . We use entropy regularization (Grandvalet and Bengio 2005) to exploit the unlabeled data by encouraging the classifier boundary to pass through low density regions.

The training data is comprised of two sets:  $\mathcal{X}_L$  and  $\mathcal{X}_U$ . Each element of  $\mathcal{X}_L$  consists of a pair  $(X, Y)$ , where  $X$  denotes a sequence  $x_1, \dots, x_N$  of vectors in  $\mathbb{R}^D$  and  $Y$  denotes a one-hot encoding of the class label for  $X$  (and is hence in  $\mathbb{R}^C$  where  $C$  is the number of classes). Each element of  $\mathcal{X}_U$  consists of a sub-sequence  $X$  identified by the change point detection step (i.e., the individual sub-sequences in the set  $\mathcal{P}_U$ ). The loss function that we use to train  $f_\psi$  is given by:

$$\mathcal{L}_C(\psi) = \frac{1}{|\mathcal{X}_L|} \sum_{(X,Y) \in \mathcal{X}_L} \mathcal{L}_{CE}(X, Y) + \frac{\lambda_C}{|\mathcal{X}_U|} \sum_{X \in \mathcal{X}_U} \mathcal{L}_{NE}(X).$$

Here,  $\lambda_C$  is a tuning parameter,  $\mathcal{L}_{CE}$  is the cross entropy loss, and  $\mathcal{L}_{NE}$  is the negative entropy loss:

$$\mathcal{L}_{CE}(X, Y) = - \sum_{c=1}^C Y_c \log f_\psi(\overline{f_\theta(X)})_c$$

$$\mathcal{L}_{NE}(X) = - \sum_{c=1}^C f_\psi(\overline{f_\theta(X)})_c \log f_\psi(\overline{f_\theta(X)})_c.$$

Above,  $f_\psi$  represents the output of the feedforward classification network which ends with a softmax distribution over  $C$  classes. The input to  $f_\psi$  is the mean empirical representations learned by network  $f_\theta$  for input sequence  $X$ . The negative entropy loss encourages the network  $f_\psi$  to produce low entropy empirical class distributions for unlabeled data. This encourages unlabeled data to be mapped to a distribution that concentrates on a single class, pushing the classifier boundary to  $f_\psi$  towards low-density regions.

---

**Algorithm 1** SSL via change point detection

---

**Inputs:** Unlabeled sequence  $X$ , labeled sequences  $\{X_l, Y_l\}$ , CP detection parameters  $\tau, w$ ,

**Output:** Trained networks:  $f_\theta, f_\psi$

**Init:** Add similar/dissimilar pairs from  $\{X_l\}$  to  $\mathcal{P}_S, \mathcal{P}_D$

**for**  $i = 1$  to  $\text{length}(X)$  **do**

  Form windows:  $X_p^i, X_f^i$

$m_i = \text{MMD}(X_p^i, X_f^i)$

**if**  $m_i > \tau$  **then**

    Form two segments before CP:  $X_{p1}^i, X_{p2}^i$

    Form two segments after CP:  $X_{f1}^i, X_{f2}^i$

    Add pairs  $(X_{p1}^i, X_{p2}^i)$  and  $(X_{f1}^i, X_{f2}^i)$  to  $\mathcal{P}_S$

    Add pairs  $(X_{f1}^i, X_{p2}^i)$  and  $(X_{f2}^i, X_{p1}^i)$  to  $\mathcal{P}_D$

**for**  $j = 1$  to  $\text{num\_epochs}$  **do**

  Train network  $f_\theta$  by optimizing loss  $\mathcal{L}_R$

**for**  $j = 1$  to  $\text{num\_epochs}$  **do**

  Train network  $f_\psi$  by optimizing loss  $\mathcal{L}_C$

---

A summary of our overall approach to semi-supervised learning via change point detection is given in Algorithm 1.

## Experiments

### Baselines

All of the following baselines use the same representation network  $f_\theta$  and classification network  $f_\psi$  architectures.

**Supervised** In the supervised setting, only the labeled sequence is passed through both the representation  $f_\theta$  and classifier networks  $f_\psi$ . We train the two networks in an end-to-end manner by minimizing:

$$\mathcal{L}_S(\theta, \psi) = \frac{1}{|\mathcal{X}_L|} \sum_{(X, Y) \in \mathcal{X}_L} \mathcal{L}_{\text{CE}}(X, Y).$$

**Denosing Autoencoder** A denosing autoencoder (Dai and Le 2015) or its variants such as the ladder network (where the reconstruction error for intermediate layers is also minimized) (Zeng et al. 2017) are often employed for semi-supervised learning with sequential data. Since it has been previously shown that the performance gap between these approaches is marginal (Zeng et al. 2017) – which we have observed as well – we focus only on the autoencoder as a baseline. In this approach, for every  $X \in \mathcal{X}_U$ , we also consider a perturbed version  $\hat{X}$  produced by adding noise to  $X$ . Both  $X$  and  $\hat{X}$  are passed through an encoder network  $f_\theta$  to obtain embeddings which are used by a decoder network  $f'_\theta$  to reconstruct the unlabeled data. A reconstruction loss of the form  $\mathcal{C}(X) = \|X - f'_\theta(f_\theta(\hat{X}))\|^2$  is incorporated into the loss function to exploit the unlabeled data. The labeled data is first passed through the encoder network  $f_\theta$  to obtain embeddings, which are then fed into a classifier network  $f_\psi$ . We train the two networks in an end-to-end manner by minimizing:

$$\mathcal{L}_{\text{AE}}(\theta, \psi) = \frac{1}{|\mathcal{X}_L|} \sum_{(X, Y) \in \mathcal{X}_L} \mathcal{L}_{\text{CE}}(X, Y) + \frac{\lambda_C}{|\mathcal{X}_U|} \sum_{X \in \mathcal{X}_U} \mathcal{C}(X).$$

---

Method	10 labels	30 labels
Supervised	0.90 ± 0.02	0.98 ± 0.01
Autoencoder	0.87 ± 0.03	0.99 ± 0.01
SSL-CP	0.99 ± 0.01	0.99 ± 0.01
SSL-CP (ER)	0.99 ± 0.01	0.99 ± 0.01

---

Table 1: Classifier performance for mean, variance change

### Synthetic Experiments

In all of the results below, we use the mean F1 score (unweighted) as an evaluation metric. In all synthetic simulations, we split the data in a 70/30 ratio where we use the larger split for training and the smaller split as a test dataset. We further split the training data in a ratio of 10/60/30. We use the smallest of these splits to obtain labeled data, the largest as unlabeled data for the semi-supervised setting, and the last split for validation. We use a small sub-sequence (comprising of 20 segments) in the unlabeled split to tune the parameters for change point detection. In our results, SSL-CP denotes our approach to semi-supervised learning via change points, but without the inclusion of the negative entropy term in the loss function. SSL-CP (ER) denotes our approach when including this entropy regularization term.

**Changing Mean and Variance** This example consists of data generated by a univariate normal distribution that switches its parameters  $(\mu, \sigma^2)$  every 500 samples. We use 1500 such random switches to produce a sequence of data with five classes, correspond to the parameter settings  $\{(2, 0.1), (4, 0.1), (4, 0.7), (10, 0.1), (0, 0.1)\}$ . We use the symmetrical KL divergence from (1) to detect change points in the unlabeled data. This is a simple change point detection problem where we detect all change points correctly. We use small sub-sequences of length 20 as labeled and unlabeled data. and we show the resulting performance in Table 1. This is a relatively simple sequence classification problem as it requires merely learning that the mean and variance determine class membership. Both the supervised and autoencoder baselines do reasonably well. However, classes 2 and 3 have the same mean but different variance, and both baselines struggle compared to SSL-CP in separating these classes when only 10 labels are provided.

**Mackey-Glass Equation** The Mackey-Glass equation (Glass and Mackey 2010) is a non linear time delay differential equation defined as

$$\frac{d(x(t))}{dt} = -0.1x(t) + \frac{\beta x(t)(t - \tau)}{1 + x(t - \tau)^{10}}.$$

In a manner similar to (Kohlmorgen and Lemm 2002), we generate a sequence by randomly switching between parameters  $(\beta, \tau) \in \{(0.2, 8), (0.18, 16), (0.2, 22), (0.22, 30)\}$  every 1400 samples. We define class membership according to the parameter settings of each segment. We generated 2000

Code: [https://github.com/nahad3/semi\\_sup\\_cp](https://github.com/nahad3/semi_sup_cp)




Figure 3: T-SNE visualizations for the representations learned by the representation network ( $f_\theta$ ) on the Mackey-Glass example when 5 labels are provided from each class. Figure 3(a) shows representations learned by an autoencoder using both labeled and unlabeled data. It can be seen in 3(b) that different classes overlap in this representation. Figure 3(c) show the representations learned by SSL-CP, which are clustered and non-overlapping. This leads to improved classification when limited labels are provided. True labels for these representations are shown in Figure 3(d).




Figure 4: Example switching Mackey-Glass sequence.

Model	20 labels	30 labels	60 labels
Supervised	$0.55 \pm 0.07$	$0.86 \pm 0.04$	$0.95 \pm 0.02$
Autoencoder	$0.73 \pm 0.04$	$0.90 \pm 0.02$	$0.98 \pm 0.01$
SSL-CP	$0.96 \pm 0.02$	$0.98 \pm 0.01$	$0.99 \pm 0.01$
SSL-CP (ER)	$0.99 \pm 0.02$	$0.99 \pm 0.01$	$0.99 \pm 0.01$

Table 2: Mackey-Glass: Classifier performance for different number of labeled examples

such segments and added  $\mathcal{N}(0, 0.1)$  noise to the entire sequence. A small sub-sequence is shown in Figure 4. We obtained pairs of sequences of size 100 using change points detected on the unlabeled dataset, where almost all true change points were detected correctly. There were about 4000 such pairs. We obtained 8100 non-overlapping windows of size 100 from the unlabeled-split for use by the autoencoder. Labeled data is also formed using non-overlapping windows of size 100 were used as labels. Table 2 shows results for different numbers of provided labels. We see that SSL-CP approach significantly outperforms the baselines. The representations learned by the autoencoder and SSL-CP are visualized in Figure 3, which illustrates that the autoencoder does not perform as well because it fails to learn representations that exhibit sufficient clustering. The influence of varying the number of provided pairs is shown in Table 3. We note that entropy regularization enhances the performance of SSL-CP when amount of unlabeled data is large.

Model	600 Pairs	1800 Pairs	4000 Pairs
SSL-CP	$0.87 \pm 0.2$	$0.94 \pm 0.1$	$0.96 \pm 0.1$
SSL-CP (ER)	$0.87 \pm 0.2$	$0.95 \pm 0.1$	$0.99 \pm 0.2$

Table 3: Mackey-Glass: Classifier performance for different amounts of unlabeled data

## Real World Datasets

**HCI: Gesture Recognition** The HCI gesture recognition dataset consists of a user performing 5 different gestures using the right arm (Forster, Roggen, and Troster 2009). Data is obtained from 8 IMUs placed on the arm. The gestures recorded included drawing triangle up, circle, infinity, square, and triangle down. We also consider the null case (where the user is not performing an activity) as a class. We use the free-hand subset from this dataset as it presents a relatively challenging classification problem when compared with the more controlled subset. Rather than using consecutive non-overlapping windows (as the resulting sub-sequences are too small to contain a single class, since the duration of the null class can be very small), the sequential data is first divided into 100 segments using the labels. 30 segments are left as test data.

This dataset presents a challenge to the SSL-CP approach in that most classes never appear adjacent to each other in the data set as they are always separated by a period in the null class. To obtain similarity constraints involving class pairs that do not include the null class, we generate a sequence by repeating a randomly sampled segment and concatenating it with another repeated randomly sampled segment. Change detection is then applied on this concatenated sequence to provide similar and dissimilar pairs. 600 of such similar dissimilar pairs were obtained.

When all labels within the dataset are provided, the mean F1 score for the supervised approach is 0.88. Such a score can actually sometimes be achieved by the supervised classifier even when only 1 label from each class is provided. However in this setting, the results can vary dramatically depending on exactly which instances are labeled. We obtained classification results across 30 trials, with a different ran-



Supervised	Autoencoder	SSL-CP
0.63	0.68	0.72

Table 4: HCI: Mean classifier performance when using one label per class

Supervised	Autoencoder	SSL-CP
11%	26%	63%

Table 5: HCI: Percentage of trials in which each method performs best when using one label per class

dom choice of which instance in each class were labeled. We show the average results in Table 4. In Table 5 we show the percentage of trials in which each method performed best.

**WISDM: Activity Recognition** The WISDM activity recognition dataset (Kwapisz, Weiss, and Moore 2010) consists of 36 users performing 6 activities which include running, walking, ascending stairs, descending stairs, sitting, and standing. Data is collected through an accelerometer mounted on the participant’s chest which provides 3 dimensional data sampled at 20Hz. For our experiments, we retained data from users 33, 34, and 35 as test set. We split the data from the rest of the users in a 70/30 ratio, using the large split for training and the small split for validation. We used a small sub-sequence (consisting of about 20 change points) to tune the change detection parameters. Once tuned, we obtained change points on the entire training set to obtain pairs of size 50. We obtained a total of about 4000 such pairs. We used about 7000 non-overlapping windows of size 50 as unlabeled data for the autoencoder. We used non-overlapping windows of size 50 as labeled data. In all experiments, we used a balanced number of labels from each class.

Table 6 shows results when 48 labels (6 from each class). When pairs from all detected change points within the training set (4000 in number) are used, the performance of SSL-CP is slightly worse than that of the autoencoder. This is because many false change points are detected (up to about 40% false change points) for a small number of users, leading to erroneous similarity constraints. After the removal of 10 such users, the number of falsely detected change points is reduced (to below 10% across all users) and about 1600 pairs are obtained. The performance of SSL-CP for this case (filtered users) is notably better than the autoencoder. The performance further improves when all true change points are provided. In such a case, the number of unlabeled pairs are larger leading to improved performance of entropy regularization as well. Figure 5 shows the relationship between classification performance and the number of labels available. In this experiment, only pairs derived from change points on the filtered users are used.

Method	F1 score
Supervised	$0.45 \pm 0.04$
Autoencoder	$0.54 \pm 0.02$
SSL-CP (All users)	$0.53 \pm 0.03$
SSL-CP (Filtered users)	$0.65 \pm 0.02$
SSL-CP (True CPs, all users)	$0.66 \pm 0.01$
SSL-CP-ER (Filtered users)	$0.65 \pm 0.01$
SSL-CP-ER (True CPs, all users)	$0.69 \pm 0.01$

Table 6: WISDM: Classifier performance with 48 labels




Figure 5: Performance on WISDM as the the number of provided labels increases. (Filtered users)

## Discussion and Conclusion

As highlighted by the performance on the WISDM dataset, the performance of our proposed method depends critically on the successful detection of change points. The detection of too many false change points can lead to corrupt similarity/dissimilarity constraints, that can potentially deteriorate performance. The other main limitation of the SSL-CP approach is that obtaining a rich set of similarity/dissimilarity constraints across all possible combinations of classes requires that these classes appear adjacent in the data. However, as we observed in the HCI dataset, the generation of additional sequences can provide a synthetic solution to this problem that is effective in practice.

Despite these limitations, SSL-CP consistently outperformed our baselines on both synthetic and real-world datasets. This clearly shows the potential utility of incorporating information from change points in semi-supervised learning. Moreover, the results on the WISDM dataset clearly illustrate the potential improvement that could be realized by more robust change point detection procedures. Historically, change point detection has been mostly restricted to detecting anomalies or segmenting data. We hope that this work will encourage the community to recognize the utility of change point detection in semi-supervised learning and to devote more attention to developing improved non-parametric change point detection procedures.

## Ethics Statement

Many biomedical and health assistive technologies use sensor data for fitness tracking, rehabilitative exercise monitoring etc. Machine learning models within such applications require labelled data for training which can be difficult to acquire. For example, it can be difficult to recruit wheel chair users to obtain labels for training a machine learning model that tracks rehabilitation exercises. The proposed method could help train these models without requiring labels from many participants. This is more convenient for the participants as they would need to provide lesser data. This would also be more convenient for the developers of such models as they will need to recruit less participants.

## References

- Baghshah, M. S.; and Shouraki, S. B. 2009. Semi-supervised metric learning using pairwise constraints. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI)*.
- Bai, S.; Kolter, Z.; and Koltun, V. 2018. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv:1803.01271*.
- Berthelot, D.; Carlini, N.; Goodfellow, I.; Papernot, N.; Oliver, A.; and Raffel, C. A. 2019. Mixmatch: A holistic approach to semi-supervised learning. In *Proc. Adv. in Neural Inf. Proc. Sys. (NeurIPS)*.
- Chang, W.-C.; Li, C.-L.; Yang, Y.; and Póczos, B. 2019. Kernel Change-point Detection with Auxiliary Deep Generative Models. In *Proc. Int. Conf. on Learning Representations (ICLR)*.
- Dai, A.; and Le, Q. 2015. Semi-supervised sequence learning. In *Proc. Adv. in Neural Inf. Proc. Sys. (NeurIPS)*.
- Forster, K.; Roggen, D.; and Troster, G. 2009. Unsupervised classifier self-calibration through repeated context occurrences: Is there robustness against sensor displacement to gain? In *2009 IEEE Int. Symp. on Wearable Computers (ISWC)*.
- Glass, L.; and Mackey, M. 2010. Mackey-Glass equation. *Scholarpedia* 5(3): 6908.
- Grandvalet, Y.; and Bengio, Y. 2005. Semi-supervised learning by entropy minimization. In *Proc. Adv. in Neural Inf. Proc. Sys. (NeurIPS)*.
- Gretton, A.; Borgwardt, K.; Rasch, M.; Schölkopf, B.; and Smola, A. 2012. A kernel two-sample test. *Journal of Machine Learning Research* 13: 723–773.
- Hsu, Y.-C.; and Kira, Z. 2016. Neural network-based clustering using pairwise constraints. In *Proc. Int. Conf. on Learning Representations Workshop Track*.
- Joachims, T. 1999. Transductive inference for text classification using support vector machines. In *Proc. Int. Conf. on Machine Learning (ICML)*.
- Koch, G.; Zemel, R.; and Salakhutdinov, R. 2015. Siamese neural networks for one-shot image recognition. In *Proc. Int. Conf. on Machine Learning Workshop in DeepLearning*.
- Kohlmorgen, J.; and Lemm, S. 2002. A dynamic hmm for on-line segmentation of sequential data. In *Proc. Adv. in Neural Inf. Proc. Sys. (NeurIPS)*.
- Kwapisz, J. R.; Weiss, G. M.; and Moore, S. A. 2010. Cell phone-based biometric identification. In *2010 Fourth IEEE International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, 1–7. IEEE.
- Laine, S.; and Aila, T. 2016. Temporal ensembling for semi-supervised learning. *arXiv:1610.02242*.
- Li, S.; Xie, Y.; Dai, H.; and Song, L. 2015. M-statistic for kernel change-point detection. In *Proc. Adv. in Neural Inf. Proc. Sys. (NeurIPS)*.
- Liu, S.; Yamada, M.; Collier, N.; and Sugiyama, M. 2013. Change-point detection in time-series data by relative density-ratio estimation. *Neural Networks* 43: 72–83.
- Oliver, A.; Odena, A.; Raffel, C. A.; Cubuk, E. D.; and Goodfellow, I. 2018. Realistic evaluation of deep semi-supervised learning algorithms. In *Proc. Adv. in Neural Inf. Proc. Sys. (NeurIPS)*.
- Rasmus, A.; Berglund, M.; Honkala, M.; Valpola, H.; and Raiko, T. 2015. Semi-supervised learning with ladder networks. In *Proc. Adv. in Neural Inf. Proc. Sys. (NeurIPS)*.
- Van Engelen, J.; and Hoos, H. 2020. A survey on semi-supervised learning. *Machine Learning* 109(2): 373–440.
- Wang, J.; Chen, Y.; Hao, S.; Peng, X.; and Hu, L. 2019. Deep learning for sensor-based activity recognition: A survey. *Pattern Recognition Letters* 119: 3–11.
- Xing, E.; Jordan, M.; Russell, S.; and Ng, A. 2003. Distance metric learning with application to clustering with side-information. In *Proc. Adv. in Neural Inf. Proc. Sys. (NeurIPS)*.
- Zeng, M.; Yu, T.; Wang, X.; Nguyen, L. T.; Mengshoel, O. J.; and Lane, I. 2017. Semi-supervised convolutional neural networks for human activity recognition. In *2017 IEEE Int. Conf. on Big Data (BigData)*.
- Zhu, X. 2005. Semi-supervised learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences.
- Zhu, X.; Ghahramani, Z.; and Lafferty, J. D. 2003. Semi-supervised learning using gaussian fields and harmonic functions. In *Proc. Int. Conf. on Machine Learning (ICML)*.