Normalizing self-supervised learning for provably reliable Change Point Detection

1st Alexandra Bazarova
Applied AI Center
Skoltech
Moscow, Russia
A.Bazarova@skoltech.ru

2nd Evgenia Romanenkova
Applied AI Center
Skoltech
Moscow, Russia
shulgina@phystech.edu

3rd Alexey Zaytsev
Applied AI Center
Skoltech
Moscow, Russia
A.Zaytsev@skoltech.ru

Abstract—Change point detection (CPD) methods aim to identify abrupt shifts in the distribution of input data streams. Accurate estimators for this task are crucial across various real-world scenarios. Yet, traditional unsupervised CPD techniques face significant limitations, often relying on strong assumptions or suffering from low expressive power due to inherent model simplicity. In contrast, representation learning methods overcome these drawbacks by offering flexibility and the ability to capture the full complexity of the data without imposing restrictive assumptions. However, these approaches are still emerging in the CPD field and lack robust theoretical foundations to ensure their reliability. Our work addresses this gap by integrating the expressive power of representation learning with the groundedness of traditional CPD techniques. We adopt spectral normalization (SN) for deep representation learning in CPD tasks and prove that the embeddings after SN are highly informative for CPD. Our method significantly outperforms current state-of-the-art methods during the comprehensive evaluation via three standard CPD datasets.

Index Terms—change point detection, self-supervised learning, time series analysis

I. INTRODUCTION

Change point detection (CPD) is designed to identify unexpected changes in data streams as quickly and accurately as possible. Such a problem statement has been vital for decades for various real-world problems, spanning detailed theoretical analysis and the development of diverse applied methods. The examples include industrial quality control [1], medicine [2], video surveillance [3], finance [4], climate change [5], human activity recognition [6], and more. Most of these problems fall under the unsupervised setting due to the frequent unavailability of annotated data [3], [7]. This common scenario necessitates robust techniques that can operate effectively without labeled examples.

Traditionally, CPD problems in the unsupervised setting are solved via different theoretically-justified parametric and non-parametric approaches [8], [9]. Although these methods have a solid theoretical base, they suffer from strong assumptions about the underlying distribution and type of changes [10], [11]. For example, they look for the change in the mean or the variance in the underlying distribution [12], [13].

Such limitations can be handled via deep learning methods: with no assumptions for input data distribution, a well-designed neural network captures the key aspects of time series structure, achieving better detection [14], [15]. Moreover, they provide the only way to work with complex high-dimensional streams [3]. For the unsupervised setting, self-supervised learning (SSL) methods provide a consistent approach for, e.g., anomaly detection, among other problems. These methods have been widely used and led to advances for specific problems and modalities [7], [16]. However, only a few papers on this topic have been published in the context of change point detection [15], [17], [18]. The problem is the complexity of the representation space in an unsupervised regime and no explanation of their behavior even for anomaly detection [19], [20]. Thus, there is no theoretically rigorous understanding of how to construct the most suitable embeddings for successful CPD.

We take the best of both worlds: our method improves state-of-the-art representation learning models and aligns them with classic theory from the CPD area to produce reasonable embedding space. Within our approach, the SN forces the representation learning model to keep distributional shifts in embedding space, leading to superior empirical performance. The results are supported by theoretical justification and various experiments for different datasets and models.

Our main contributions are:

- A proof that the usage of Spectral Normalization (SN) for neural networks ensures that changes in the raw data are preserved in the representation space, maintaining the test power for standard change point detection methods.
- A framework based on self-supervised representation learning augmented with SN for change point detection. As base SSL methods, we consider two models: contrastive TS2Vec [7] and non-contrastive TS-BYOL — our adaptation of BYOL [16] for time series data. The pipeline is presented in Figure 1.
- Empirical evidence on the effectiveness of the proposed framework for three standard CPD datasets.

II. RELATED WORKS

We focus on the development of CPD methods in the unsupervised setting. It is of particular interest since it is often challenging to obtain annotated data due to the complexity and high cost of the annotation process. For example, in the oil&gas domain, well-log data segmentation [21] requires
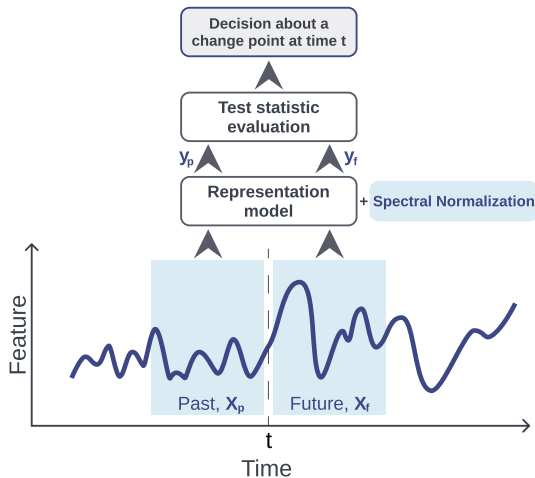
Fig. 1. The structure of the proposed change point detection procedure based on a model obtained via self-supervised learning with spectral normalization.

experts' involvement for annotation. Another example is the medical imaging domain [22], where the annotation process is very time-consuming and sometimes even includes the consultation of several experts. Thus, all methods that do not need labels from the experts are valuable from a practical point of view.

A substantial part of the unsupervised CPD methods compares a predefined statistic of time series subintervals with some threshold: if the value exceeds the threshold — the model alarms a change point. There are different ways to define such statistics, including likelihood ratio as in CUSUM methods [23], [24], the observability matrices-based dissimilarity as in subspaces analysis approaches [25], [26], Information entropy [27], or simple distance in kernel spaces [28]. However, all methods described above have significant drawbacks: they all either strongly rely on assumptions about the distribution of the time series data or the state space model, only take into consideration a few characteristics (like temporal shape) of the data, not being able to reflect the complexity of its structure, or are limited by the simplicity of the structure of the selected statistics.

Deep learning-based methods for the CPD task are particularly interesting because they do not require strong assumptions about the input data, are more flexible in terms of architecture than kernel- or graph-based methods, and allow to work with various data types, including images and video [3]. One example of such a method is KL-CPD [14], which extends kernel-based approaches by considering deep kernel learning; the authors of [14] have shown that this method outperforms multiple classical approaches. However, KL-CPD is still bounded by the shape of the kernel and does not fit the unsupervised paradigm. Given the chosen setting, we turn to self-supervised representation learning methods. These methods have shown outstanding performance [7], [16], being very versatile and thus suitable for various tasks. In the context of CPD, they are not very widely researched yet. One of the

key works in this field is TS-CP2 [15], which first applied contrastive learning to the considered task. Briefly, contrastive learning aims to learn such an embedding space in which similar sample pairs stay close to each other while dissimilar ones are far apart. The authors of [15] have shown that this approach is beneficial for change point detection; however, their method, as well as other existing ones in this area, relies on the intuitive explanation of its behavior, lacking a rigorous theoretical foundation.

Our work takes the next step in developing these methods. We propose applying the spectral normalization (SN) method, which was initially used to stabilize GANs [29], to the representation learning models. The SN method has already been used as a mechanism for ensuring the invertibility of neural networks [30]; we, in turn, extend this idea to obtain specific statistical properties of the learned representations. This technique allows us to align the expressive power of deep self-supervised learning with the existing mathematical methods, providing theoretical justification for the reasonableness of the learned representations from the change point detection point of view.

III. METHODOLOGY

Let us begin with the formal statement of change point detection problem [31]. As input, we have a sequence of independent D-dimensional observations $\mathbf{x}_1, \mathbf{x}_2, \ldots$ that possibly has a change in distribution at an unknown moment ν. So, observations $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_\nu$ come from one distribution, while $\mathbf{x}_{\nu+1}, \mathbf{x}_{\nu+2}, \ldots$ — from another distribution. Without loss of generality we can assume that $\mathbf{x}_i$ belongs to the D-dimensional sphere $\mathbb{S}^{D-1}$. The moment ν of the abrupt change is called the *change point*, with $\nu = \infty$ denoting the absence of a change. The goal is to detect the change point as precisely as possible, avoiding false alarms and long delays.

Following the state-of-the-art CPD methods [15], we propose to solve this problem in the representations space instead of the raw data space. The goal of the transition into the representation space is to capture the underlying structure of the time series, discard measurement artifacts and noise, and reduce data dimensionality. The paper [32] shows that InfoNCE loss facilitates this by maximizing mutual information between the representations.

While getting rid of the noise, we need to make sure that nothing important is lost — we hope that the representations, if not more informative than the raw data itself, are at least *as informative as the raw data* in the context of the change point detection task.

To ensure this, we propose applying the spectral normalization (SN) technique [29] on the weight matrices $\{\mathbf{W}_l\}_{l=1}^L$ for an L-layered representation learning network. At every training step (before the optimization step), the SN method evaluates the spectral norm of the matrix $\hat{\lambda} \approx \|\mathbf{W}_l\|_2$ via the power iteration method [33] and then enforces it to be less than a predefined hyperparameter c:

$$\mathbf{W}_l = \begin{cases} c \cdot \frac{\mathbf{W}_l}{\hat{\lambda}}, & \text{if } c < \hat{\lambda}, \\ \mathbf{W}_l, & \text{otherwise.} \end{cases} \quad (1)$$

This technique should be beneficial for representation learning in the context of CPD, as it enhances the test power-preserving property in neural networks. Below, we (1) describe the theoretical justification for this argument and (2) provide experimental evidence suggesting that SN benefits the change point detection quality.

A. Spectral Normalization preserves the CPD quality

Consider the following problem setting. Let $\mathbf{x}_1, \ldots, \mathbf{x}_t$ be a sequence of independent observations. Before the CP, observations come from the distribution p_∞, after — from p_0. Both distributions and the change point ν are considered unknown. The change point detection problem is a problem of testing a hypothesis H_0 versus an alternative H_1:

$$H_0 : \nu = t, \qquad \text{i.e. } \mathbf{x}_i \sim p_\infty \ \forall 1 \leq i \leq t, \quad (2)$$
$$H_1 : \nu \in \{1, \ldots, t-1\}, \quad \text{i.e. } \mathbf{x}_{\leq \nu} \sim p_\infty, \ \mathbf{x}_{>\nu} \sim p_0.$$

We will consider the two most common approaches in this setting: the two-sample tests built on kernel-based statistics and the likelihood ratio-based tests. Below, it is shown that spectral normalization ensures the transition into the space of representations:

1) in the case of two-sample kernel-based tests — does not change the type II error rate of convergence to zero (for <u>no</u> assumptions about the data distribution at all);

2) in the case of likelihood ratio-based CPD tests — does not decrease the power of such tests (for reasonable assumptions about the data distribution).

But first, let us outline the main property that the spectral normalization equips the neural networks with.

1) Bi-Lipschitz neural networks: The paper [19] showed that the SN technique ensures the bi-Lipschitz property for the neural networks G of the form:

$$G(\mathbf{X}) = h \circ g(\mathbf{X}), \quad (3)$$

where $g(\mathbf{X}) = \mathbf{A}\mathbf{X} + \mathbf{B}$, and $h(\mathbf{X})$ is a composition of L residual blocks:

$$h(\mathbf{X}) = h_L \circ \cdots \circ h_1(\mathbf{X}), \quad (4)$$
$$h_l(\mathbf{X}) = \mathbf{X} + g_l(\mathbf{X}), \ l = 1, \ldots, L,$$

where $g_l(\mathbf{X}) = \sigma(\mathbf{W}_l\mathbf{X} + \mathbf{B})$. Here, we can consider the case of consecutive mappings of $\mathbf{x}_i$ s.t. $\mathbf{y}_i = G(\mathbf{X}_{i-w:i})$ and $\mathbf{Y} = \{\mathbf{y}_i\}_{i=1}^t$ or the case when $\mathbf{Y} = G(\mathbf{X})$ is a single vector; i.e., the mapping G may (depending on the specific CPD procedure) have an image in either $\mathbb{R}^d$ or $\mathbb{R}^{t \times d}$. Inside $g_l(\mathbf{X})$, σ is a sigmoid function, but other non-linearity functions are possible. Neural networks of such form are common among the modern methods [34], [35].

More precisely, the authors of [19] proved the following lemma:

Lemma III.1. *[19] Consider a hidden mapping $h : \mathcal{X} \to \mathcal{H}$ of a form* (4). *If for $0 < \alpha \leq 1$ all g_l's are α-Lipschitz, i.e., $\|g_l(\mathbf{X}) - g_l(\mathbf{X}')\|_{\mathcal{H}} \leq \alpha\|\mathbf{X} - \mathbf{X}'\|_{\mathcal{X}} \ \forall(\mathbf{X}, \mathbf{X}') \in \mathcal{X}^2$. Then:*

$$L_1\|\mathbf{X} - \mathbf{X}'\|_{\mathcal{X}} \leq \|h(\mathbf{X}) - h(\mathbf{X}')\|_{\mathcal{H}} \leq L_2\|\mathbf{X} - \mathbf{X}'\|_{\mathcal{X}}, \quad (5)$$

where $L_1 = (1 - \alpha)^L$, $L_2 = (1 + \alpha)^L$.

Let $\mathcal{X} = \mathcal{H} = \mathbb{R}^n$, and $\|.\|_2$ denote the Euclidean distance. Since the SN technique enforces $\|\mathbf{W}_l\|_2 \leq c$, and the Lipschitz constant of the residual block (4) is bounded by $\|\mathbf{W}_l\|_2$:

$$\|g_l(\mathbf{X}) - g_l(\mathbf{X}')\|_2 \leq \|\mathbf{W}_l\mathbf{X} - \mathbf{W}_l\mathbf{X}'\|_2 \leq \|\mathbf{W}_l\|_2\|\mathbf{X} - \mathbf{X}'\|_2.$$

With $c \leq 1$, the spectral normalized hidden mapping h ensures that the Euclidean norm is preserved; the same is true for the aforementioned mapping $G(\mathbf{X}) = h \circ g(\mathbf{X})$. Note that in finite dimensional vector spaces, all norms are equivalent; hence, the preservation of the Euclidean distance implies the preservation of any vector norm distance. We will say that the mappings for which lemma III.1 holds are *distance-preserving*.

2) Preservation of kernel distances: In change point detection, the test statistic $S(\mathbf{X})$ is often a function of kernel distances between the observations $k(\mathbf{x}, \mathbf{x}')$ [28], [36]. Most popular kernels, such as RBF $k(\mathbf{x}, \mathbf{x}') = \exp\left(\frac{-\|\mathbf{x} - \mathbf{x}'\|_2^2}{2\sigma^2}\right)$, rely on the vector norm of the difference $\|\mathbf{x} - \mathbf{x}'\|$. In the previous section, we proved the bi-Lipschitz property of the SN networks, i.e., the preservation of such distances. Then, the following statement holds.

Lemma III.2. *Consider two observations $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{S}^{D-1}$ and their representations $\mathbf{y}_i, \mathbf{y}_j$. Consider the RBF kernel $k(\mathbf{x}, \mathbf{x}')$. SN networks preserve RBF kernel distance, i.e. $\exists \hat{C}, \tilde{C} > 0$:*

$$\hat{C}k(\mathbf{x}_i, \mathbf{x}_j) \leq k(\mathbf{y}_i, \mathbf{y}_j) \leq \tilde{C}k(\mathbf{x}_i, \mathbf{x}_j).$$

We provide the proof in the Appendix A-C. Note that this statement holds for both $G : \mathcal{X} \to \mathbb{R}^d$ and $G : \mathcal{X} \to \mathbb{R}^{t \times d}$.

3) Kernel-based test: In our experiments, we consider the MMD (*maximum mean discrepancy*) [36] as a kernel-based statistic. This is a nonparametric probabilistic distance commonly used in two-sample tests. Given a kernel k of the RKHS $\mathcal{H}_k$, the MMD distance between two distributions P and Q is defined as

$$\text{MMD}_k^2(P, Q) = \|\mu_P - \mu_Q\|_{\mathcal{H}^k}^2 = \quad (6)$$
$$= \mathbb{E}_P[k(\zeta, \zeta')] - 2\mathbb{E}_{P,Q}[k(\zeta, \xi)] + \mathbb{E}_Q[k(\xi, \xi')],$$

where $\mu_P = \mathbb{E}_{\zeta \sim P}[k(\zeta, .)]$, $\mu_Q = \mathbb{E}_{\xi \sim Q}[k(\xi, .)]$ are the kernel mean embeddings of distributions P and Q, accordingly. We use the biased empirical estimate of the MMD distance: given observations $Z = (\zeta_1, \ldots, \zeta_m) \sim P$, $\Xi = (\xi_1, \ldots \xi_m) \sim Q$

$$\text{MMD}_b^2(Z, \Xi) = \frac{1}{m^2}\sum_i\sum_j k(\zeta_i, \zeta_j) - \quad (7)$$
$$- \frac{2}{m^2}\sum_i\sum_j k(\zeta_i, \xi_j) + \frac{1}{m^2}\sum_i\sum_j k(\xi_i, \xi_j).$$

For kernels that are preserved under distance-preserving mappings, the following theorem holds.

Theorem 1. *Consider two sequences of observations $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_n]$ and $\hat{\mathbf{X}} = [\hat{\mathbf{x}}_1, \ldots, \hat{\mathbf{x}}_n]$, $\mathbf{x}_i, \hat{\mathbf{x}}_i \in \mathbb{S}^{D-1}$. Denote by $\mathbf{Y}, \hat{\mathbf{Y}}$ their images under the distance-preserving mapping G. Consider a bounded kernel $0 \leq k(\mathbf{x}, \mathbf{x}') \leq K$ that is preserved*

under G and the corresponding sample MMD_b^k statistic. The MMD_b^k-based two-sample test of level α has the same type II error rate $O(n^{-\frac{1}{2}})$ of convergence to zero in the space of embeddings $\mathcal{Y}$ as in the space of raw observations $\mathcal{X}$.

For proof, see Appendix A-D. Note that the MMD statistic based on RBF kernels is a particular case of Theorem 1; we use this statistic in our experiments.

To sum up, spectral normalization ensures that SN mappings preserve kernel distances, which, in turn, leads to the preservation of MMD_b-based two-sample test properties in the latent space.

4) Likelihood ratio tests: Suppose that p_0, p_∞ in the setting (2) are known. Consider the likelihood ratio of an interval $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_t] \in \mathbb{R}^{t \times D}$ in the setting (2) as its test statistic:

$$S(\mathbf{X}) = \frac{p_0^{[1:t]}(\mathbf{x}_1, \ldots, \mathbf{x}_t)}{p_\infty^{[1:t]}(\mathbf{x}_1, \ldots, \mathbf{x}_t)}. \tag{8}$$

The following proposition gives sufficient conditions for the neural network $G : \mathbb{R}^{t \times D} \to \mathbb{R}^{t \times d}$ of the form (3) to preserve likelihood ratio for elliptical distributions.

Proposition III.1. *Consider a sequence of independent variables $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_t]$ distributed according to an elliptical distribution with a possible change in mean at an unknown location ν. Consider $G(\mathbf{X})$ of the form (3), with $h(\mathbf{X})$ being an invertible function. Denote by $\mathbf{Y} = [\mathbf{y}_1, \ldots, \mathbf{y}_t]$ the transformation of the original sequence, i.e. $\mathbf{Y} = G(\mathbf{X})$.*

Let p_0 and p_∞ be the joint PDFs tested for the original sequence $\mathbf{X}$, $\hat{p}_0$ and $\hat{p}_\infty$ be the corresponding joint PDFs for the transformed sequence $\mathbf{Y}$. Then

$$\frac{p_0(\mathbf{X})}{p_\infty(\mathbf{X})} = \frac{\hat{p}_0(\mathbf{Y})}{\hat{p}_\infty(\mathbf{Y})}. \tag{9}$$

The proof is presented in the Appendix.

Therefore, for the neural network $G(\mathbf{X})$ to preserve likelihood ratio, it is sufficient for its component $h(\mathbf{X})$ to be invertible. Thankfully, the SN technique provides this property. As mentioned, the SN enforces the Lipschitz constant of g_l to be not greater than c for any predefined $c > 0$. The authors of [30] show that for $c < 1$, such restraint ensures the invertibility of the residual block h_l (4). Hence, applying SN to all layers $h_1, \ldots, h_L$ enforces the invertibility of $h(X)$.

This likelihood ratio preservation property entails the following theorem.

Theorem 2. *Consider a statistical CPD test based on the likelihood ratio. The proposed strategy ensures that for any elliptical distribution, the power of such test of level α preserves, i.e.,*

$$\mathbb{P}_{H_1}(\text{reject } H_0 \,|\, \text{raw data space}) = \tag{10}$$
$$= \mathbb{P}_{H_1}(\text{reject } H_0 \,|\, \text{representation space}).$$

The proof is given in Appendix A-B.

Summary. The SN technique ensures the test power-preserving property in neural networks. This property is valuable since it encourages the results of any considered statistical

test of considered form not to deviate too much when carried out in the latent space instead of the observation space. Test power preservation is in some way similar to information preservation but through the lens of the change point detection task.

B. General pipeline

Now that we have outlined the theoretical properties of the proposed approach let us describe the general CPD pipeline depicted in Figure 1 that we follow. The pipeline consists of three stages: (1) representation construction, (2) test statistic calculation, and (3) change point detection.

In the first stage, following the standard sliding window technique, the raw time series data is cropped into a sequence of overlapping time intervals $\mathbf{X}$ of length $2w$; for each interval, we consider separately its first $\mathbf{X}_{\text{past}} \in \mathbb{R}^{w \times D}$ and second part $\mathbf{X}_{\text{future}} \in \mathbb{R}^{w \times D}$. Then, each $\mathbf{X}_{\text{past}}$ and $\mathbf{X}_{\text{future}}$ is transformed into its representation. Denote the embeddings by $(\mathbf{Y}_{\text{past}}, \mathbf{Y}_{\text{future}})$.

In the second stage, for each pair $(\mathbf{Y}_{\text{past}}, \mathbf{Y}_{\text{future}})$, the test statistic is calculated. We consider two test statistics: the cosine distance and the MMD-score.

In the last stage, the change point detection, we estimate the change points from the obtained statistic values. We report a change at every pair of intervals where the statistic is greater than some empirically chosen threshold δ: an example is shown in Figure 2.
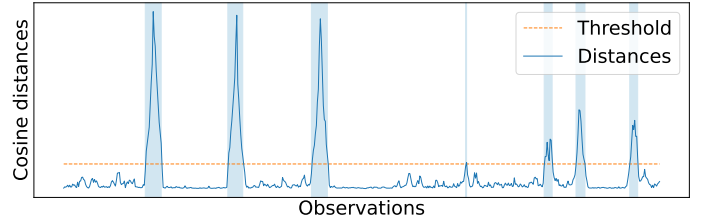

Fig. 2. Cosine distances between the subsequent subintervals of observations. The value above a predefined threshold indicates the presence of a change-point. The color filled area denotes that the corresponding intervals contain a change point.

C. Models

As base models, we consider two self-supervised models: TS2Vec [7] and our adaptation of BYOL [16] for time series data. During the CPD procedure, the embeddings of TS2Vec are obtained in two modes: either an interval is mapped to its embedding, i.e., the model performs the mapping $G : \mathbb{R}^{w \times D} \to \mathbb{R}^d$ — and the corresponding test statistic is the cosine distance — or an interval is mapped into a sequence of embeddings, i.e., $G : \mathbb{R}^{w \times D} \to \mathbb{R}^{w \times d}$ — and the MMD-score is used as the statistic. For BYOL, only the mapping to a single embedding is performed to calculate then the cosine distance.

1) TS2Vec: [7] leverages hierarchical contrastive loss for learning informative time series representations. The structure of the model is depicted in Figure 3.

In the TS2Vec encoder learning pipeline, the input time series is cropped into two overlapping subseries, creating augmented views of the original data. These samples are then encoded into timestamp-wise sequences of embeddings. The resulting embeddings are contrasted in two dimensions: temporal and instance-wise. The loss for the t-th timestamp of the i-th sample in a batch includes two corresponding terms:

$$\mathcal{L}_{\text{inst}}^{(i,t)} = -\log \frac{\exp(\mathbf{h}_{i,t} \cdot \mathbf{h}'_{i,t})}{S_i^{\text{inst}}},$$

$$S_i^{\text{inst}} = \sum_{j=1}^{B} \big(\exp(\mathbf{h}_{i,t} \cdot \mathbf{h}'_{j,t}) + \mathbb{1}(i \neq j) \exp(\mathbf{h}_{i,t} \cdot \mathbf{h}_{j,t}) \big);$$

$$\mathcal{L}_{\text{temp}}^{(i,t)} = -\log \frac{\exp(\mathbf{h}_{i,t} \cdot \mathbf{h}'_{i,t})}{S_i^{\text{temp}}},$$

$$S_i^{\text{temp}} = \sum_{t' \in \Omega} \big(\exp(\mathbf{h}_{i,t} \cdot \mathbf{h}'_{i,t'}) + \mathbb{1}(t \neq t') \exp(\mathbf{h}_{i,t} \cdot \mathbf{h}_{i,t'}) \big).$$

Here, $\mathbf{h}_{i,t}$ and $\mathbf{h}'_{i,t}$ refer to the projections of representations of the same timestamp but from two different augmentations of the original sample; S_i^{inst} is the sum over the batch; Ω refers to the overlapping part of the two subseries.

The contrasting is performed hierarchically: after each iteration of loss calculation, the max pooling operation is executed along the time axis until the dimension of the time axis reduces to one. The final loss comprises of multiple terms corresponding to different levels of representations' granularity.

2) TS-BYOL: We adapt the self-supervised BYOL [16] model, originally designed for images, as an alternative source of representations. The model (Figure 4) consists of two asymmetrical parts: the "target" and "online" networks. The online network has an additional prediction head and is trained via backpropagation, while the target network's weights are updated slowly using exponential moving average (EMA):

$$\xi := \beta\xi + (1 - \beta)\theta, \tag{11}$$

where ξ and θ denote the weights of the target and the online network, accordingly, and the hyperparameter β is usually close to 1. In our experiments, we set $\beta = 0.996$.

Online and target networks build embeddings $\mathbf{h}', \mathbf{h}''$ from augmented views $\mathbf{X}', \mathbf{X}''$ of the same sample $\mathbf{X}$. The online network minimizes the L_2-norm difference between normalized embeddings $\bar{\mathbf{h}} = \frac{\mathbf{h}}{\|\mathbf{h}\|_2}$ to match the target network's embeddings:

$$\mathcal{L}(\mathbf{h}', \mathbf{h}'') = \|\bar{\mathbf{h}}' - \bar{\mathbf{h}}''\|_2^2 = 2 - 2\frac{\langle \mathbf{h}', \mathbf{h}'' \rangle}{\|\mathbf{h}'\|_2 \|\mathbf{h}''\|_2}. \tag{12}$$

For the augmentations, we use only the random cropping of length w from an interval of length $2w$, as a similar procedure provides optimal performance in [37]. The architecture consists of a sequence of 1-D convolutions, followed by ReLU nonlinearities and dropouts, as a backbone; both projection and prediction heads are two-layer MLPs.

During CPD, we use representations $\mathbf{y}$ produced by the backbone neural network, not by the projection heads, and compute cosine distances between the representations of intervals of neighbor windows.

IV. RESULTS

A. Compared methods

- **TS-CP2** [15] is a self-supervised model intended for change point detection. It was the first method to leverage contrastive learning to obtain informative time series data representations.
- **KL-CPD** [14] is a principled CPD framework that employs deep kernel learning for two-sample hypothesis testing, being another option for NN-based CPD.
- **ESPRESSO** [27] is a hybrid approach that exploits both statistical and temporal shape properties in the CPD process. This method does not involve deep learning while showing decent performance.
- A self-supervised **TS2Vec** [7] employs hierarchical contrastive loss for representation learning, showing state-of-the-art results in time series classification, forecasting and anomaly detection. **SN-TS2Vec** is our modification of the TS2Vec model: we performed spectral normalization for each of its convolutional layers.
- An SSL approach **BYOL** [16] provides an additional way to obtain a representation learning model via a non-
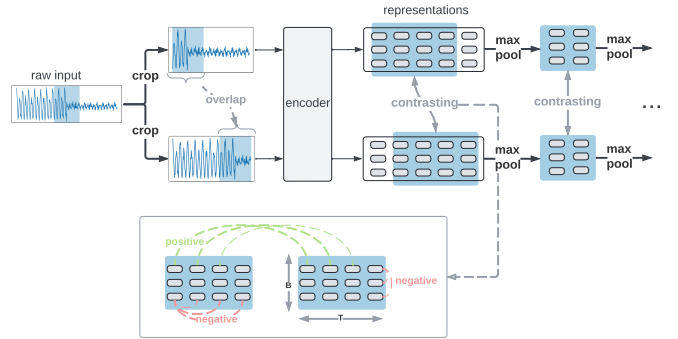

Fig. 3. TS2Vec architecture. The upper part of the figure is the hierarchical contrasting procedure; the lower part is the selection of positive and negative pairs.
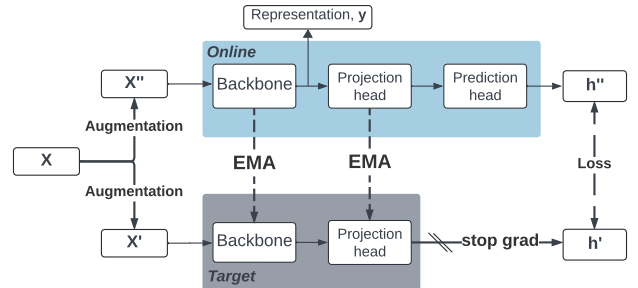

Fig. 4. BYOL's architecture. The upper part is the training *Online* network, and the lower part is the *Target* network with disabled backpropagation.

contrastive loss function. We employ **SN-BYOL** in some experiments to compare SN-equipped version to a vanilla one.

Suffixes *cos* and *MMD* denote the type of test statistic that is used during the change point detection procedure: either the cosine distance or the MMD score, accordingly.

B. Datasets

We provide the analysis of our method on the datasets below. These datasets are considered, following TS-CP2 [15], since we aimed to compare this SOTA model to our method. All datasets are publicly available.

- **Yahoo** [38]. It is a widely used anomaly detection benchmark consisting of time series that contain metrics of the various Yahoo services with manually labeled anomalies. Following [15], we used the fourth benchmark since it includes annotations of change points.
- **HASC** [39]. This dataset contains human activity data collected by three-axis accelerometers. We used the same subset as [14] and [15]. Change points in this time series indicate alternations in the type of activities (stay, walk, jog, skip, stair up, and stair down).
- **USC-HAD** [40]. The dataset also contains human activity data monitored by wearable sensors. The observed data represents basic daily life activities such as walking, sitting, sleeping, etc. We followed the pipeline from [15] and combined 30 random activities from six participants, using only the data from the accelerometer.

TABLE I
DATASETS

Dataset	Domain	#Timestamps [1]	#Sequences	#CP [2]
Yahoo	$\mathbb{R}^+$	164K	100	208
USC-HAD	$\mathbb{R}^3$	97K	6	30
HASC	$\mathbb{R}^3$	39K	1	65

C. F_1-score for change point detection

Since we aim to compare our model to the current SOTA in this field, TS-CP2 [15], we used their equivalent of F1-score for the change point detection task as the primary metric of detection quality This procedure's inputs are a sequence of binary ground truth labels (for each interval, whether it contains a change point or not) and a sequence of predicted labels, and the output is a value between 0 and 1. The metric suggests that in order to correctly detect a change point, it is sufficient for the detection algorithm to indicate the presence of a change point in at least one of the consecutive intervals containing it. For extra details, see the TS-CP2 paper [15].

D. Main results

Table II compares SN-TS2Vec to three other state-of-the-art CPD methods. For both Yahoo!A4Benchmark and USC-HAD datasets, the SN-TS2Vec approach with the cosine distance

[1] Total number of timestamps.
[2] Total number of change points.

used as a test statistic achieves two of the three best results. For the HASC dataset, all TS2Vec variants show results close to KL-CPD and ESPRESSO, outperforming the TS-CP2 approach. So, SN-TS2Vec either dominates the existing methods or performs almost as well as they do.

As for the TS-BYOL model, even though it does not achieve SOTA performance, it shows solid results on the Yahoo!A4Benchmark dataset. Comparison of the SN-BYOL to its vanilla version confirms that the SN technique enhances model properties in the context of CPD task.

TABLE II
F_1 MEASURE FOR DIFFERENT DETECTION MARGINS FOR THE PROPOSED SN-TS2VEC APPROACH VS EXISTING CPD METHODS.

Dataset	Model	Detection margin		
		24	50	75
Yahoo	TS-CP2	0.64	**0.81**	0.843
	KL-CPD	0.579	0.576	0.544
	ESPRESSO	0.224	0.340	0.4442
	TS-BYOL	0.5	0.706	0.768
	SN-BYOL	0.694	0.766	0.789
	TS2Vec	0.688	<u>0.788</u>	0.816
	SN-TS2Vec, cos	**0.726**	0.775	**0.852**
	SN-TS2Vec, MMD	<u>0.692</u>	**0.81**	<u>0.851</u>
		100	200	400
USC-HAD	TS-CP2	0.824	0.857	0.833
	KL-CPD	0.743	0.718	0.632
	ESPRESSO	0.633	0.833	0.833
	TS-BYOL	0.5	0.796	0.933
	SN-BYOL	0.5	0.636	0.722
	TS2Vec	<u>0.873</u>	**0.97**	<u>0.952</u>
	SN-TS2Vec, cos	0.736	<u>0.909</u>	**1**
	SN-TS2Vec, MMD	**0.909**	0.809	**1**
		60	100	200
HASC	TS-CP2	0.4	0.438	0.632
	KL-CPD	**0.479**	**0.473**	0.467
	ESPRESSO	0.288	0.423	**0.693**
	TS-BYOL	0.316	0.398	0.26
	SN-BYOL	0.403	0.416	0.418
	TS2Vec	<u>0.476</u>	<u>0.467</u>	0.444
	SN-TS2Vec, cos	<u>0.476</u>	0.306	<u>0.663</u>
	SN-TS2Vec, MMD	<u>0.476</u>	<u>0.467</u>	0.444

1) TS2Vec with and without SN: The results of the comparison of the model modifications are presented in Table III. The experiments were conducted for three different window sizes for each dataset; the hidden representations size was set to 16 for the Yahoo and USC-HAD datasets and to 32 for the HASC dataset. On the Yahoo!A4Benchmark dataset, the SN-TS2Vec strongly outperforms its vanilla version, with the cosine distance version being slightly dominant over the MMD-score one. The results on USC-HAD also show the superiority of the SN version: two out of three best results for each dataset belong to it. On the HASC dataset, all models demonstrate comparable results.

We also provide the extended version of this comparison in Figure 5, considering multiple code sizes (i.e., hidden dimensions). The results on the Yahoo!A4Benchmark confirm that the SN technique enhances the properties of the representations; for two other datasets, the SN improves the results for approximately half of the code sizes. Note that these two datasets contain significantly fewer change points than

Yahoo!A4Benchmark. Therefore, the results on the latter are more representative.
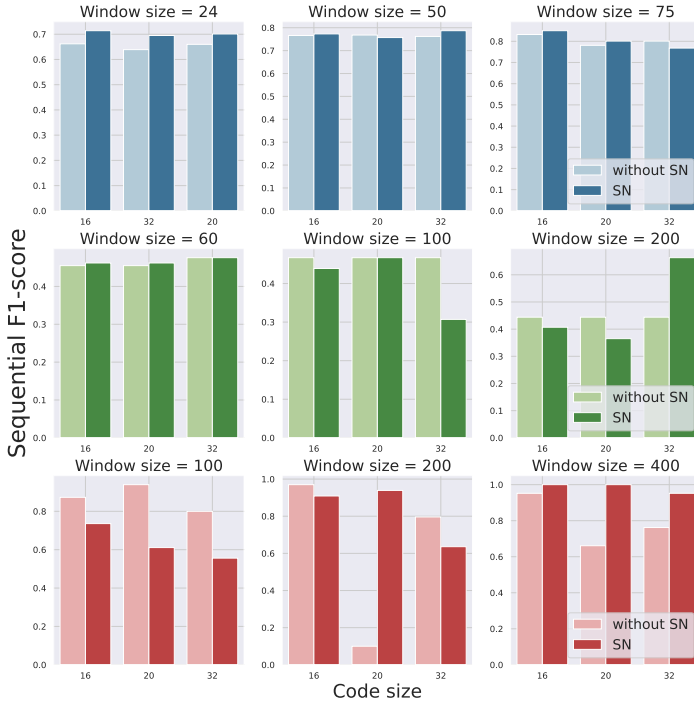

Fig. 5. Comparison of TS2Vec performance with and without spectral normalization on all datasets. The cosine distance was used as the test statistic. The upper row — Yahoo!A4Benchmark, middle — HASC, lower — USC-HAD.

TABLE III
COMPARISON OF VANILLA TS2VEC VS SN-TS2VEC PERFORMANCE FOR TWO TEST STATISTICS: COSINE DISTANCE AND MMD-SCORE.

Model	Detection margin		
	Yahoo		
	24	50	75
TS2Vec, cos	0.688±0.035	0.788±0.021	0.816±0.019
TS2Vec, MMD	0.624±0.048	0.764±0.03	0.796±0.033
SN-TS2Vec, cos	**0.726±0.021**	0.775±0.018	**0.852±0.001**
SN-TS2Vec, MMD	0.692±0.115	**0.81±0.034**	0.851±0.037
	USC-HAD		
	100	200	400
TS2Vec, cos	0.873±0.063	**0.97±0.052**	0.952±0.083
TS2Vec, MMD	0.724±0.132	0.909±0	**1±0**
SN-TS2Vec, cos	0.736±0.212	0.909±0	**1±0**
SN-TS2Vec, MMD	**0.909±0**	0.809±0	**1±0**
	HASC		
	60	100	200
TS2Vec, cos	**0.476±0**	**0.467±0**	0.444±0
TS2Vec, MMD	**0.476±0**	**0.467±0**	0.444±0
SN-TS2Vec, cos	**0.476±0**	0.306±0.142	**0.663±0.089**
SN-TS2Vec, MMD	**0.476±0**	**0.467±0**	0.444±0

2) TS-BYOL with and without spectral normalization: To enrich our comparison, we conducted additional experiments with an alternative self-supervised approach, TS-BYOL, on all three datasets. The basic TS-BYOL was compared to the TS-BYOL with SN. The results can be seen in Figure 7. It is clear that TS-BYOL with SN either outperforms or performs

on par with its vanilla version on two of three datasets — Yahoo!A4Benchmark and HASC. The results obtained on the USC-HAD dataset are different; however, it should be noted that the total number of change points in the test set of USC-HAD is the smallest among all datasets, and the standard deviation of the obtained metrics is relatively high (up to std = 0.4). Therefore, the results on USC-HAD are not as representative as those of the other two datasets.

E. The dynamics of the representations

We compared the dynamics of the representations for Vanilla TS2Vec and SN-TS2Vec. The experiment was performed as follows:

1) For each change point, we sampled a subsequence of length 300 right before the CP; let us call it $\mathbf{X}$.
2) We cloned each $\mathbf{X}$ and replaced the last half with the subsequence from right after the CP, denoting the sample obtained by $\hat{\mathbf{X}}$. Thus, for each CP, we have $\mathbf{X}$ and $\hat{\mathbf{X}}$: their first halves are identical, while their last halves belong to different distributions.
3) After that, we transformed $\mathbf{X}$ and $\hat{\mathbf{X}}$ into a sequence of representations via the sliding window procedure.
4) Finally, we calculated cosine similarities between the corresponding representations in $\mathbf{X}$ and $\hat{\mathbf{X}}$, obtaining a sequence of similarities for each CP. Supposedly, they should begin to diverge right after a CP appears.

We averaged the obtained similarity arrays over all the change points. The results are presented in Figure 6. For two out of three considered datasets, the embeddings obtained with SN-TS2Vec diverge faster, achieving greater dissimilarity values. This visualization confirms that SN-TS2Vec is more suitable for CPD tasks than its vanilla version.
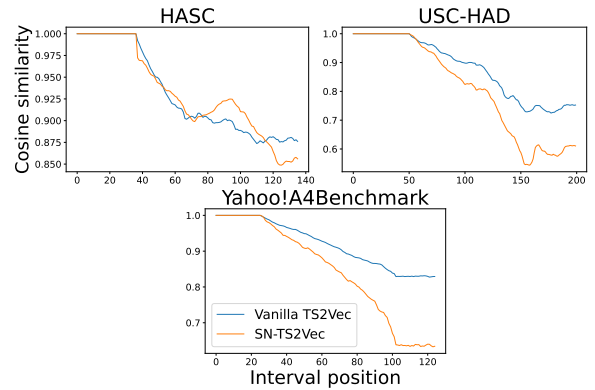

Fig. 6. The dynamics of the representations. Blue color corresponds to vanilla TS2Vec, and orange — to its SN version. Sliding window sizes: 60 for HASC, 75 for Yahoo!A4Benchmark, 100 for USC-HAD. Each value at the position i corresponds to the similarity between the representations of $\mathbf{X}[i : i + w]$ and $\hat{\mathbf{X}}[i : i + w]$.
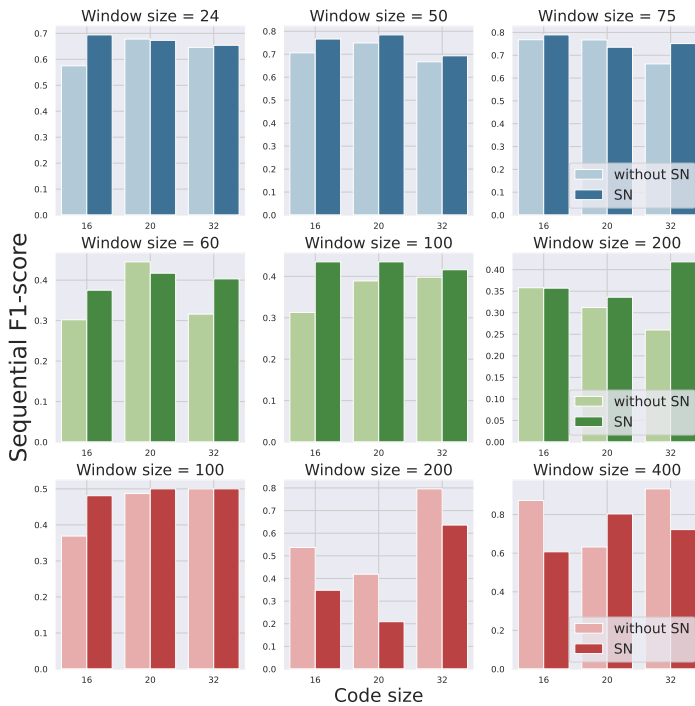
Fig. 7. Comparison of TS-BYOL performance with and without spectral normalization on all datasets. The upper row — Yahoo!A4Benchmark, middle — HASC, lower — USC-HAD.

VI. CONCLUSION

Our findings provide a step forward in the CPD field, bridging the gap between the expressive ability of representation learning models and the theoretical grounds of CPD problems.

We showed how to design self-supervised learning models with respect to CPD task specificity. The main contribution is the usage of Spectral Normalization, which, as we theoretically and empirically demonstrate, preserves test power for various two-sample tests in latent space. Thus, the resulting representations appear to be more robust for further change detection.

Through a series of experiments, we confirm the effectiveness of the self-supervised approach, showing that the modern TS2Vec model outperforms current state-of-the-art approaches for CPD. Moreover, Spectral Normalization enhances the performance of TS2Vec, leading to a 5% improvement in the target F1 metrics.

REFERENCES

[1] S. Basu and M. Meckesheimer, "Automatic outlier detection for time series: an application to sensor data," *Knowledge and Information Systems*, vol. 11, pp. 137–154, 2007.

[2] S. Liu, A. Wright, and M. Hauskrecht, "Change-point detection method for clinical decision support system rule monitoring," *Artificial intelligence in medicine*, vol. 91, pp. 49–56, 2018.

[3] E. Romanenkova, A. Stepikin, M. Morozov, and A. Zaytsev, "Indid: Instant disorder detection via a principled neural network," in *Proceedings of the 30th ACM International Conference on Multimedia*, pp. 3152–3162, 2022.

[4] S. Banerjee and K. Guhathakurta, "Change-point analysis in financial networks," *Stat*, vol. 9, no. 1, p. e269, 2020.

[5] N. Itoh and J. Kurths, "Change-point detection of climate time series by nonparametric method," in *Proceedings of the world congress on engineering and computer science*, vol. 1, pp. 445–448, 2010.

[6] S. Aminikhanghahi and D. J. Cook, "Using change point detection to automate daily activity segmentation," in *2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, pp. 262–267, IEEE, 2017.

[7] Z. Yue, Y. Wang, J. Duan, T. Yang, C. Huang, Y. Tong, and B. Xu, "Ts2vec: Towards universal representation of time series," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, pp. 8980–8987, 2022.

[8] S. Aminikhanghahi and D. J. Cook, "A survey of methods for time series change point detection," *Knowl. Inf. Syst.*, vol. 51, pp. 339–367, May 2017.

[9] C. Truong, L. Oudre, and N. Vayatis, "Selective review of offline change point detection methods," *Signal Processing*, vol. 167, p. 107299, 2020.

[10] F. Enikeeva and Z. Harchaoui, "High-dimensional change-point detection under sparse alternatives," *The Annals of Statistics*, vol. 47, no. 4, pp. 2051–2079, 2019.

[11] F. Gustafsson, "The marginalized likelihood ratio test for detecting abrupt changes," *IEEE Transactions on automatic control*, vol. 41, no. 1, pp. 66–78, 1996.

[12] S. Gharghabi, C.-C. M. Yeh, Y. Ding, W. Ding, P. Hibbing, S. LaMunion, A. Kaplan, S. E. Crouter, and E. Keogh, "Domain agnostic online semantic segmentation for multi-dimensional time series," *Data mining and knowledge discovery*, vol. 33, pp. 96–130, 2019.

[13] S. Liu, M. Yamada, N. Collier, and M. Sugiyama, "Change-point detection in time-series data by relative density-ratio estimation," *Neural Networks*, vol. 43, pp. 72–83, 2013.

[14] W.-C. Chang, C.-L. Li, Y. Yang, and B. Póczos, "Kernel change-point detection with auxiliary deep generative models," in *International Conference on Learning Representations*, (Vancouver, Canada), ICLR, 2018.

[15] S. Deldari, D. V. Smith, H. Xue, and F. D. Salim, "Time series change point detection with self-supervised contrastive predictive coding," in *Proceedings of The Web Conference 2021*, WWW '21, Association for Computing Machinery, 2021.

[16] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. Guo, M. Gheshlaghi Azar, *et al.*, "Bootstrap your own latent-a new approach to self-supervised learning," *Advances in neural information processing systems*, vol. 33, pp. 21271–21284, 2020.

[17] S. Chatterjee, "Changepoint detection using self supervised variational autoencoders," 2021.

[18] X. Bao, L. Chen, J. Zhong, D. Wu, and Y. Zheng, "A self-supervised contrastive change point detection method for industrial time series," *Engineering Applications of Artificial Intelligence*, vol. 133, p. 108217, 2024.

[19] J. Liu, Z. Lin, S. Padhy, D. Tran, T. Bedrax Weiss, and B. Lakshminarayanan, "Simple and principled uncertainty estimation with deterministic deep learning via distance awareness," *NeurIPS*, vol. 33, pp. 7498–7512, 2020.

[20] A. Vazhentsev, G. Kuzmin, A. Shelmanov, A. Tsvigun, E. Tsymbalov, K. Fedyanin, M. Panov, A. Panchenko, G. Gusev, M. Burtsev, *et al.*, "Uncertainty estimation of transformer predictions for misclassification detection," in *ACL*, pp. 8237–8252, 2022.

[21] E. Romanenkova, A. Rogulina, A. Shakirov, N. Stulov, A. Zaytsev, L. Ismailova, D. Kovalev, K. Katterbauer, and A. AlShehri, "Similarity learning for wells based on logging data," *Journal of Petroleum Science and Engineering*, vol. 215, p. 110690, 2022.

[22] R. Malladi, G. P. Kalamangalam, and B. Aazhang, "Online bayesian change point detection algorithms for segmentation of epileptic activity," in *2013 Asilomar Conference on Signals, Systems and Computers*, IEEE, Nov. 2013.

[23] E. S. Page, "Continuous inspection schemes," *Biometrika*, vol. 41, no. 1/2, pp. 100–115, 1954.

[24] X. Shao and X. Zhang, "Testing for change points in time series," *Journal of the American Statistical Association*, vol. 105, p. 1228–1240, Sept. 2010.

[25] Y. Kawahara, T. Yairi, and K. Machida, "Change-point detection in time-series data based on subspace identification," in *IEEE ICDM*, IEEE, Oct. 2007.

[26] V. Moskvina and A. Zhigljavsky, "An algorithm based on singular spectrum analysis for change-point detection," *Communications in Statistics - Simulation and Computation*, vol. 32, p. 319–352, Jan. 2003.

[27] S. Deldari, D. V. Smith, A. Sadri, and F. Salim, "Espresso: Entropy and shape aware time-series segmentation for processing heterogeneous sensor data," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 4, p. 1–24, Sept. 2020.

[28] Z. Harchaoui, F. Vallet, A. Lung-Yut-Fong, and O. Cappe, "A regularized kernel-based approach to unsupervised audio segmentation," in *IEEE ICASSP*, IEEE, Apr. 2009.

[29] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida, "Spectral normalization for generative adversarial networks," in *International Conference on Learning Representations*, 2018.

[30] J. Behrmann, W. Grathwohl, R. T. Chen, D. Duvenaud, and J.-H. Jacobsen, "Invertible residual networks," in *International conference on machine learning*, pp. 573–582, PMLR, 2019.

[31] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential analysis: Hypothesis testing and changepoint detection*. CRC press, 2014.

[32] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.

[33] R. V. Mises and H. Pollaczek-Geiringer, "Praktische verfahren der gleichungsauflösung .," *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik*, vol. 9, p. 152–164, Jan. 1929.

[34] V. Ashish, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, p. I, 2017.

[35] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE CVPR*, pp. 770–778, 2016.

[36] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel two-sample test," *Journal of Machine Learning Research*, vol. 13, no. 25, pp. 723–773, 2012.

[37] A. E. Marusov and A. Zaytsev, "Noncontrastive representation learning for intervals from well logs," *IEEE Geoscience and Remote Sensing Letters*, vol. 20, pp. 1–5, 2023.

[38] "Yahoo research webscope dataset,,"

[39] N. Kawaguchi, N. Ogawa, Y. Iwasaki, K. Kaji, T. Terada, K. Murao, S. Inoue, Y. Kawahara, Y. Sumi, and N. Nishio, "Hasc challenge: gathering large scale human activity corpus for the real-world activity understandings," in *Proceedings of the 2nd augmented human international conference*, pp. 1–5, 2011.

[40] M. Zhang and A. A. Sawchuk, "USC-HAD: a daily activity dataset for ubiquitous activity recognition using wearable sensors," in *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, (Pittsburgh, Pennsylvania), pp. 1036–1043, ACM, 2012.

[41] A. K. Gupta and D. K. Nagar, *Matrix variate distributions*. Chapman and Hall/CRC, 2018.

[42] K. Lee, K. Lee, H. Lee, and J. Shin, "A simple unified framework for detecting out-of-distribution samples and adversarial attacks," *NeurIPS*, vol. 31, 2018.

APPENDIX A
PROOFS FROM THE SECTION III

A. Proof of proposition III.1

The proof is provided for the case of Gaussian random variables, but it is clear that the similar calculations may be performed for any distribution from the elliptical family.

$\mathbf{X}$ can be viewed as a sample from the matrix normal distribution [41], and the joint PDFs p_0 and p_∞ have the densities of the following form (with the means $\mathbf{M}_0$ and $\mathbf{M}_\infty$, accordingly):

$$p(\mathbf{X}|\mathbf{M}, \mathbf{U}, \mathbf{V}) = \qquad (13)$$
$$= C \cdot \exp\left(-\frac{1}{2}\text{tr}\left[\mathbf{V}^{-1}(\mathbf{X}-\mathbf{M})^T\mathbf{U}^{-1}(\mathbf{X}-\mathbf{M})\right]\right).$$

First, consider $g(\mathbf{X}) = \mathbf{AX} + \mathbf{b}$. It is a linear transformation, so the distribution of $g(\mathbf{X})$ can be written out explicitly:

$$p(\mathbf{AX}|\mathbf{AM}, \mathbf{AUA}^T, \mathbf{V}) = \qquad (14)$$
$$= C \cdot \exp\left(-\frac{1}{2}\text{tr}\left[\mathbf{V}^{-1}(\mathbf{X}-\mathbf{M})^T\mathbf{A}^T\mathbf{A}^{-T}\mathbf{U}^{-1}\mathbf{A}^{-1}\mathbf{A}\cdot\right.\right.$$
$$\left.\left.\cdot(\mathbf{X}-\mathbf{M})\right]\right) = \tilde{C} \cdot p(\mathbf{X}|\mathbf{M}, \mathbf{U}, \mathbf{V}).$$

Here we let $b = 0$ to simplify the calculations. This implies that the linear transformation preserves the likelihood-ratio of the intervals. Now let us consider the invertible component $h(X)$. We know that for invertible transformations the PDF of the transformed distribution is calculated as

$$\hat{p}(\mathbf{y}) = p(h^{-1}(\mathbf{y}))|J_{h^{-1}}(\mathbf{y})|, \qquad (15)$$

where $J_{h^{-1}}$ is the Jacobian matrix of the inverse transformation h^{-1}. Then the likelihood-ratio after applying h becomes

$$\frac{\hat{p}_0(\mathbf{Y})}{\hat{p}_\infty(\mathbf{Y})} = \frac{p_0(\mathbf{X})|J_{h^{-1}}(\mathbf{Y})|}{p_\infty(\mathbf{X})|J_{h^{-1}}(\mathbf{Y})|} = \frac{p_0(\mathbf{X})}{p_\infty(\mathbf{X})}, \qquad (16)$$

i.e. it remains the same.

B. Proof of lemma 2

As was shown in Proposition III.1, spectral normalized neural networks of a form (3) preserve the likelihood ratio, i.e., $\Lambda(\mathbf{X}) = \hat{\Lambda}(\mathbf{Y})$, where $\Lambda(\mathbf{X})$ and $\hat{\Lambda}(\mathbf{Y})$ denote the likelihood ratio in the raw data space and the representation space, accordingly. Therefore, it is true that for any function of the likelihood ratio $S(\Lambda)$ and a predefined threshold h

$$S(\Lambda(\mathbf{X})) > h \Longleftrightarrow S(\hat{\Lambda}(\mathbf{Y})) > h. \qquad (17)$$

The rejection of the null hypothesis is usually equivalent to the statistic value exceeding the threshold; hence, due to (17), the rejection of the null hypothesis in the raw data space coincides with its rejection in the representation space.

C. Proof of the RBF kernel distance preservation

The proof is given for $G : \mathbb{R}^{w \times D} \to \mathbb{R}^d$; the same calculations for a two-dimensional codomain of G can also easily be conducted. Given:

$$L_1\|[\mathbf{x}_{i-w}, \ldots, \mathbf{x}_i] - [\mathbf{x}_{j-w}, \ldots, \mathbf{x}_j]\|_2 \le \|\mathbf{y}_i - \mathbf{y}_j\|_2 \le$$
$$\le L_2\|[\mathbf{x}_{i-w}, \ldots, \mathbf{x}_i] - [\mathbf{x}_{j-w}, \ldots, \mathbf{x}_j]\|_2 \qquad (18)$$

Therefore, taking into account that the observations lie on the unit sphere and the L_2-norm properties,

$$L_1^2\|\mathbf{x}_i - \mathbf{x}_j\|_2^2 \le \|\mathbf{y}_i - \mathbf{y}_j\|_2^2 \le L_2^2(C + \|\mathbf{x}_i - \mathbf{x}_j\|_2^2), \quad (19)$$

where $C > 0$ (a constant). Consequently,

$$\exp\left(-\frac{\|\mathbf{y}_i - \mathbf{y}_j\|_2^2}{2\sigma^2}\right) \le \exp\left(-\frac{L_1^2\|\mathbf{x}_i - \mathbf{x}_j\|_2^2}{2\sigma^2}\right) \le$$
$$\le \tilde{C}\exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|_2^2}{2\sigma^2}\right). \qquad (20)$$

where $\tilde{C}$ is some positive constant depending only on L_1 and σ^2. The last inequality uses the fact that the observations belong to the $\mathbb{S}^{D-1}$. Similarly,

$$\exp\left(-\frac{\|\mathbf{y}_i - \mathbf{y}_j\|_2^2}{2\sigma^2}\right) \geq \hat{C}\exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|_2^2}{2\sigma^2}\right). \quad (21)$$

Thus, the RBF kernel distance is indeed preserved.

D. Proof of Theorem 1

We suppose that the kernel k is preserved under G, i.e. $\exists \bar{C} > 0, \tilde{C} > 0$ such that

$$\hat{C} \cdot k(\mathbf{x}_i, \mathbf{x}_j) \leq k(\mathbf{y}_i, \mathbf{y}_j) \leq \tilde{C} \cdot k(\mathbf{x}_i, \mathbf{x}_j), \quad (22)$$

where $\mathbf{y}_i = G([\mathbf{x}_{i-w}, \ldots, \mathbf{x}_i])$.

Denote $g_w(\mathbf{x}_i) = G([\mathbf{x}_{i-w}, \ldots, \mathbf{x}_i])$. The kernel $k(\mathbf{y}_i, \mathbf{y}_j)$ in the space of the embeddings is equivalent to the kernel $\tilde{k}(\mathbf{x}_i, \mathbf{x}_j) = k(g_w(\mathbf{x}_i), g_w(\mathbf{x}_j))$ in the space of observations. Therefore, if $k(\mathbf{x}_i, \mathbf{x}_j) \leq K \ \forall \mathbf{x}_i, \mathbf{x}_j \in \mathcal{X}$, from (22) immediately follows that the upper bound $\tilde{k}(\mathbf{x}_i, \mathbf{x}_j) \leq \tilde{C}K \ \forall \mathbf{x}_i, \mathbf{x}_j \in \mathcal{X}$ holds.

Consider an MMD_b-based hypothesis test of level α with the kernel k s.t. $0 \leq k(x, y) \leq K$. It is shown in [36] that this test has the acceptance region

$$\mathrm{MMD}_b(Z, \Xi) < \sqrt{\frac{2K}{n}}\left(1 + \sqrt{2\log\alpha^{-1}}\right),$$

and is consistent:

$$\mathbb{P}\left\{|\mathrm{MMD}_b(Z, \Xi) - \mathrm{MMD}_k(P, Q)| > 4\left(\frac{K}{n}\right)^{\frac{1}{2}} + \varepsilon\right\} \leq \quad (23)$$

$$\leq 2\exp\left(\frac{-\varepsilon^2 n}{4K}\right).$$

The equation (23) shows that the type II error converges to zero at rate $O(n^{-1/2})$; note that the change of constant K to another constant $\tilde{C}K$ does not affect the convergence rate. Therefore, for the MMD_b-based hypothesis test of level α, the type II error convergence rate in the space of embeddings is the same as in the space of observations.

APPENDIX B
ADDITIONAL EXPERIMENTS

An additional experiment investigates how the SN affects embedding properties from an uncertainty estimation perspective. We conducted the experiment as follows, adopting pipeline for a classification problem [42]:

- Fit a multivariate Gaussian to the embeddings of train/val samples (windows tested for a presence of CP).
- Calculate the Mahalanobis distance to this Gaussian for each test sample.
- Discard the top 5% of test samples with the highest distances, removing the least "confident" samples, and measure the F_1-score for the remaining data. This process is repeated iteratively, reducing the dataset to 5% of its original size.

Figure 8 shows the F_1-score rejection curves for a fixed seed, comparing SN-TS2Vec with its vanilla version on the HASC dataset. The SN version consistently outperforms the vanilla one, confirming its superiority in CPD tasks.
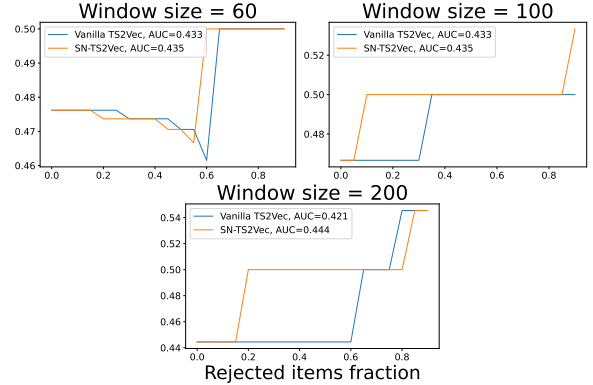

Fig. 8. F1-score rejection curves on HASC dataset. Blue color corresponds to vanilla TS2Vec, and orange — to its SN version. Higher AUCs correspond to better quality.

APPENDIX C
REPRODUCTION DETAILS

A. Train/val/test split

TABLE IV
TRAIN, VAL, AND TEST SET SIZES

Dataset	Train	Val	Test
Yahoo!A4Benchmark	0.4	0.2	0.4
USC-HAD	0.6	0.2	0.2
HASC	0.6	0.2	0.2

Each dataset was split into three parts, according to the Table C-A. The size of the test set was the same as in the TS-CP2 paper [15] for the sake of the comparison.

B. Hyperparameter choice

1) TS-BYOL: A four-layer 1D convolutional backbone with ReLU and dropout feeds two-layer MLP projection and prediction heads. Trained for 10 epochs (validation every 5), the best validation checkpoint was used for testing; average over 3 runs.

2) SN-TS2Vec: TS2Vec model parameters: $depth = 8$, $hidden\_dims = 128$, $output\_dims = code\_size$. Training, vanilla: 40 epochs (Yahoo!), 10 epochs (HASC), 1 epoch (USC-HAD). SN version: 40 epochs (Yahoo!), 5 epoch (HASC), 1 epoch (USC-HAD). Validation per epoch; the best validation checkpoint was used for testing. Averaged over 3 runs.

C. Technical details

All experiments were carried out on the NVidia L40.