

Weakly Supervised Temporal Anomaly Segmentation with Dynamic Time Warping

Dongha Lee^{1*}, Sehun Yu², Hyunjun Ju², Hwanjo Yu^{2†}

¹University of Illinois at Urbana-Champaign (UIUC), Urbana, IL, United States

²Pohang University of Science and Technology (POSTECH), Pohang, South Korea

donghal@illinois.edu, {hunul2, hyunjunju, hwanjoyu}@postech.ac.kr

Abstract

Most recent studies on detecting and localizing temporal anomalies have mainly employed deep neural networks to learn the normal patterns of temporal data in an unsupervised manner. Unlike them, the goal of our work is to fully utilize instance-level (or weak) anomaly labels, which only indicate whether any anomalous events occurred or not in each instance of temporal data. In this paper, we present WETAS, a novel framework that effectively identifies anomalous temporal segments (i.e., consecutive time points) in an input instance. WETAS learns discriminative features from the instance-level labels so that it infers the sequential order of normal and anomalous segments within each instance, which can be used as a rough segmentation mask. Based on the dynamic time warping (DTW) alignment between the input instance and its segmentation mask, WETAS obtains the result of temporal segmentation, and simultaneously, it further enhances itself by using the mask as additional supervision. Our experiments show that WETAS considerably outperforms other baselines in terms of the localization of temporal anomalies, and also it provides more informative results than point-level detection methods.

1. Introduction

Anomaly detection, which refers to the task of identifying anomalous (or unusual) patterns in data, has been extensively researched in a wide range of domains, such as fraud detection [1], network intrusion detection [18], and medical diagnosis [38]. In particular, detecting the anomaly from temporal data (e.g., multivariate time-series and videos) has gained much attention in many real-world applications, for finding out anomalous events that resulted in unexpected changes of a temporal pattern or context.

Recently, several studies on anomaly detection started to

*This work was done when the author was at POSTECH.

†Corresponding author.

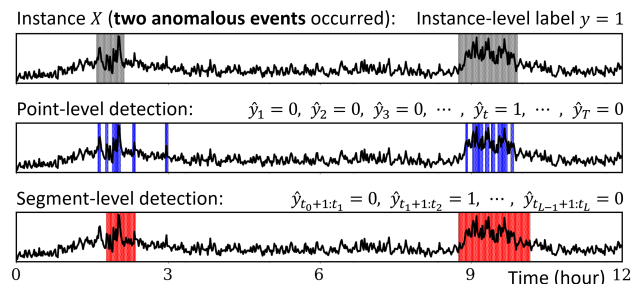


Figure 1: Two different strategies for localizing temporal anomalies. Given an input instance of length T , the point-level and segment-level anomaly detection produce the series of binary labels of length T and length L , respectively.

localize and segment the anomalies within an input instance based on deep neural networks [4, 5, 14], unlike conventional methods which simply classify each input instance as positive (i.e., anomalous) or negative (i.e., normal). In this work, we aim to precisely localize the anomalies in temporal data by detecting anomalous temporal segments, defined as the group of consecutive time points relevant to anomalous events. Note that labeling every anomalous time point is neither practical nor precise, similarly to other segmentation problems [7, 21, 31]. In this sense, the main challenge of temporal anomaly segmentation is to distinguish anomalous time points from normal time points without using point-level anomaly labels for model training.

The most dominant approach to the localization of the temporal anomalies is point-level anomaly detection based on the reconstruction of input instances [3, 12, 15, 25, 26, 30, 35, 40]. To learn the point-level (or local) anomaly score in an unsupervised manner, they focus on modeling the *normal* patterns (i.e., temporal context of each time point, usually given in a form of its past inputs [6]) from *normal* data, while considering all unlabeled data as normal. Specifically, they mainly employ variational autoencoder (VAE) to learn normal latent vectors of temporal inputs, or train sequence models to predict the next input by using RNNs or CNNs.

Methods	Anomaly Prediction	Instance Label
Park et al. [30]	point-level	✗
Su et al. [35]	point-level	✗
Xu et al. [40]	point-level	✗
Sultani et al. [36]	(static) segment-level	✓
WETAS (ours)	(dynamic) segment-level	✓

Table 1: Comparison of anomaly detection methods that are able to localize the temporal anomalies in each instance.

Then, they compute the anomaly score for each time point based on the magnitude of error between an actual input and the reconstructed (or predicted) one.

However, such unsupervised learning methods show limited performance due to the absence of information about the target anomalies. They are not good at discovering specific patterns caused by an anomalous event, especially in the case that such patterns reside in training data. As a solution, we argue that instance-level labels are often easy to acquire by simply indicating the occurrence of any anomalous events, whereas acquiring point-level labels is hard in practice. For example, given a time-series instance collected for a fixed period of time, a human annotator can easily figure out whether any anomalous events occurred or not during the period, then generate a binary label for the instance.

In this paper, we propose a novel deep learning framework, named as WETAS, which leverages WEak supervision for TEmporal Anomaly Segmentation. WETAS learns from labeled “instances” of temporal data in order to detect anomalous “segments” within the instance. The segment-level anomaly detection is more realistic than the point-level detection (Figure 1), because anomalous events usually result in variable-length anomalous segments in temporal data. The point-level detection cannot tell whether two nearby points detected as the anomaly come from a single anomalous event or not, and thus the continuity of detected points should be investigated for further interpretation of the results. Our problem setting is similar to the multiple-instance learning (MIL) [36] which learns from bag-level labels,¹ but differs in that MIL makes predictions for *static* segments of the same length. Table 1 summarizes recent anomaly detection methods for temporal data.

WETAS fully exploits the instance-level (or weak) labels for training its model and also for inferring the sequential anomaly labels (i.e., pseudo-labeling) that can be used as rough segmentation masks. Based on dynamic time warping (DTW) alignment [34] between an input instance and its sequential pseudo-label, WETAS effectively finds variable-length anomalous segments. To be specific, WETAS optimizes the model to accurately classify an input instance as its instance-level label, and simultaneously to best align the instance with its sequential pseudo-label based on the DTW.

¹The term “multiple-instance” used in MIL refers to same-length temporal segments that compose a single bag. Thus, “instance” and “bag” in MIL respectively correspond to “segment” and “instance” in our approach.

As the training progresses, the model generates more accurate pseudo-labels, and this eventually improves the model itself by the guidance for better alignment between inputs and their pseudo-labels.

Our extensive experiments on real-world datasets, including multivariate time-series and surveillance videos, demonstrate that WETAS successfully learns the normal and anomalous patterns under the weak supervision. Among various types of baselines, WETAS achieves the best performance in detecting anomalous points. Furthermore, the qualitative comparison of the detection results shows that point-level detection methods identify non-consecutive anomalous points even for a single event, whereas our framework obtains the smoothed detection results which specify the starting and end points of each event.

2. Related Work

2.1. Anomaly Detection for Temporal Data

Most recent studies on anomaly detection for temporal data (e.g., multivariate time-series) utilize deep neural networks to learn the temporal normality (or regularity) from training data [3, 15, 26]. In particular, VAE-based models [12, 25, 30, 35, 40] have gained much attention because of their capability to capture the normal patterns in an unsupervised manner. Their high-level idea is to compute the point-level anomaly scores by measuring how well a temporal context for each time point (i.e., a sliding window-based temporal input) can be reconstructed using the VAE. They are able to localize the temporal anomalies within an input instance to some extent by the help of this point-level detection (or alert) approach. However, they do not utilize information about the anomaly (e.g., anomaly labels) for model training, which makes it difficult to accurately identify the anomalies. In general, as the anomalies are not limited to simple outliers or extreme values, modeling the *abnormality* as well is helpful to learn useful features for discrimination between normal and anomalous time points [11, 28].

Several studies make use of anomaly labels by adopting U-Net architectures [39, 43] which are known to be effective for spatial segmentation [24, 32]. However, their models need to be trained with full supervision, which means that the model training is guided by anomaly labels for every time point. It makes the detector impractical because labeling every point in each instance or obtaining such point-level labels is infeasible or costs too much in practice.

2.2. Weakly Supervised Temporal Segmentation

To address the limitation of fully supervised temporal segmentation, weakly supervised approaches have been actively researched for video action segmentation (or detection) tasks [7, 31, 41, 42]. They aim to learn weakly annotated data whose labels are not given for every time

point (frame) in an instance (video). To this end, dynamic program-based approaches are employed to search for the best results of segmentation over the time axis under the weak supervision, including the Viterbi algorithm [31] and dynamic time warping [7, 41, 42]. Nevertheless, they cannot be directly applied to anomaly segmentation, because they require instance-level *sequential* labels being used as rough segmentation masks, or they impose strong constraints on the occurrence of target events (actions).²

Without the help of sequential labels and constraints, the multiple-instance learning (MIL) approach [36] showed promising results in detecting anomalous events by leveraging weakly labeled videos. It considers each video as a bag after dividing the video into a fixed number of same-length segments, then trains an anomaly detection model by using both positive (i.e., anomalous) and negative (i.e., normal) bags. The trained model is able to predict the label of each video segment as well as a bag, thus the results can be utilized for anomaly segmentation. However, due to its statically-segmented inputs, there exists a trade-off between accurate anomaly detection and precise segmentation with respect to the number of segments in a single bag. The more (and shorter) segments make it harder to capture the long-term contexts within the segments, whereas the fewer (and longer) segments generate coarse-grained segmentation results which could much differ from the actual observation.

3. Temporal Anomaly Segmentation

3.1. Problem Formulation

The goal of temporal anomaly segmentation is to specify variable-length anomalous segments (i.e., the starting and end points) within an input instance of temporal data. Formally, given a D -dimensional input instance $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T] \in \mathbb{R}^{D \times T}$ of temporal length T , we aim to produce segment-level anomaly predictions $\hat{\mathbf{y}} = [\hat{y}_{t_0+1:t_1}, \hat{y}_{t_1+1:t_2}, \dots, \hat{y}_{t_{L-1}+1:t_L}] \in \{0, 1\}^L$ where t_l denotes the end point of the l -th segment (i.e., $t_0 = 0$ and $t_L = T$). This task can be treated similarly to the point-level anomaly detection from the perspective of its final output $\hat{\mathbf{y}}$, but differs in that it guarantees a one-to-one mapping between identified segments and anomalous events.

In this work, we focus on *weakly supervised learning*, in order to bypass the challenge of obtaining densely labeled temporal data for model training. Specifically, a training set consists of N instances³ with their instance-level binary labels, $\{(\mathbf{X}^{(1)}, y^{(1)}), (\mathbf{X}^{(2)}, y^{(2)}), \dots, (\mathbf{X}^{(N)}, y^{(N)})\}$. In this context, the weak supervision means that each label simply indicates whether any anomalous events are observed or not

²They assume that each action occurs following a specific transition diagram [41] or occurs at most once in a single instance [42].

³Each instance is collected for the fixed period of time T , or obtained by splitting data streams into fixed-length temporal data.

in the instance. Only with the weakly annotated dataset, we train the model for temporal anomaly segmentation, where temporal data are given without any labels at test time.

3.2. Dynamic Time Warping (DTW) Alignment for Temporal Anomaly Segmentation

We first present how to obtain the segmentation result of an input instance by using a rough segmentation mask. To this end, we define an additional type of anomaly labels which can serve as the segmentation mask, referred to as sequential anomaly label.⁴ This label indicates the existence (and order) of normal and anomalous events for each input instance; $\mathbf{z} = [z_1, z_2, \dots, z_L] \in \{0, 1\}^L$ where L is the length of the sequential label. For example, $\mathbf{z} = [0, 1, 0, 1, 0]$ in case of $L = 5$ means that two anomalous events are observed within the instance. Note that the sequential label \mathbf{z} itself cannot be the final output of our framework, because it does not contain the information about the starting and end points of each segment.

For temporal anomaly segmentation, we utilize dynamic time warping (DTW) [34] which outputs the optimal alignment between a target instance and the sequential anomaly label. The goal of DTW is to find the optimal alignment (i.e., point-to-point matching) between \mathbf{z} and \mathbf{X} , which minimizes their total alignment cost with the time consistency among the aligned point pairs. To be precise, the total alignment cost is defined by the inner product of the cost matrix $\Delta(\mathbf{z}, \mathbf{X}) \in \mathbb{R}^{L \times T}$ and the binary alignment matrix $A \in \{0, 1\}^{L \times T}$. Each entry of the cost matrix, $[\Delta(\mathbf{z}, \mathbf{X})]_{lt} := \delta(z_l, \mathbf{x}_t)$, encodes the cost for aligning \mathbf{x}_t with z_l (i.e., the penalty for labeling \mathbf{x}_t as z_l), and similarly, the entry of A indicates the alignment between \mathbf{z} and \mathbf{X} ; that is, $A_{lt} = 1$ if \mathbf{x}_t is aligned with z_l and $A_{lt} = 0$ otherwise. The optimal alignment matrix A^* is obtained by

$$A^* = \operatorname{argmin}_{A \in \mathcal{A}} \langle A, \Delta(\mathbf{z}, \mathbf{X}) \rangle, \quad (1)$$

where \mathcal{A} is the set of possible binary alignment matrices.

In order to align the two series (\mathbf{z} and \mathbf{X}) in a temporal order, each alignment matrix is enforced to represent a single path on a $L \times T$ matrix that connects the upper-left $(1, 1)$ -th entry to the lower-right (L, T) -th entry using \rightarrow, \searrow moves. Based on the obtained alignment A^* , the starting and end points of the l -th temporal segment and its anomaly label are determined by $\hat{y}_{t_{l-1}+1:t_l} = z_l$ such that $A_{lt}^* = 1$ for all $t_{l-1} < t \leq t_l$. For example, in Figure 2 (Left), the red arrows represent non-zero entries in the optimal alignment matrix, and we can easily identify the segments whose segment-label is 1 (i.e., $z_l = 1$) as the anomaly. Unlike the conventional DTW, \downarrow moves are not available in our warping paths because each time point should be aligned with only a single label.

⁴The details about synthesizing this label are discussed in Section 3.4.

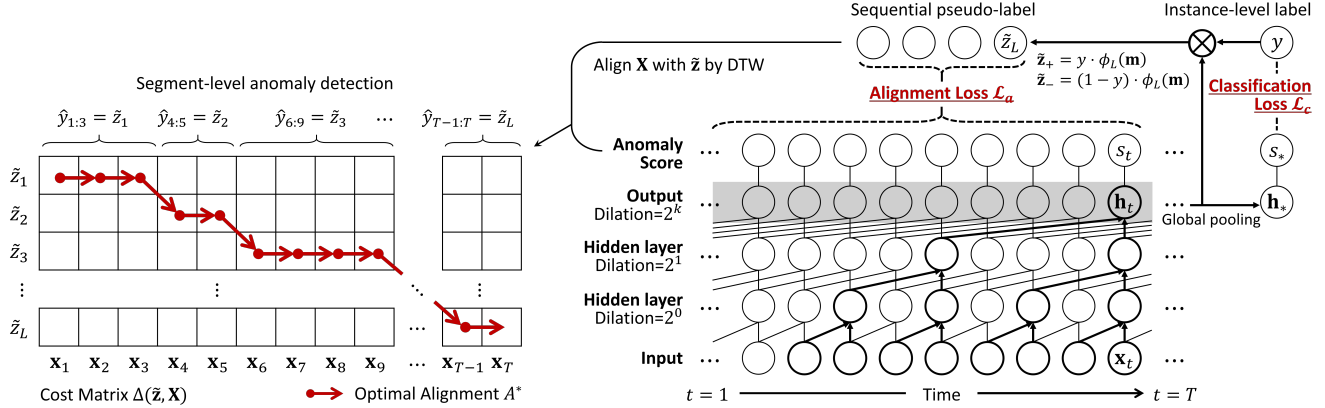


Figure 2: The overall framework of WETAS, optimized by both the classification loss \mathcal{L}_c and alignment loss \mathcal{L}_a . The red arrows in the cost matrix $\Delta(\tilde{\mathbf{z}}, \mathbf{X})$ are the desired alignment between a temporal input (column) and its sequential label (row) with time consistencies. Based on the optimal alignment, WETAS can produce the segment-level anomaly predictions.

The two key challenges are raised here as follows. 1) We need to carefully model the cost matrix based on the anomaly score so that it can effectively capture the *normality* and *abnormality* at each point, while considering the long-term dependency. In addition, 2) we need to obtain sequential labels used for the DTW alignment by distilling the information from given instance-level labels.

3.3. Anomaly Score Modeling for Cost Matrix

To model the cost function $\delta(z_l, \mathbf{x}_t)$ that is required to define the cost matrix $\Delta(\mathbf{z}, \mathbf{X})$, we parameterize the anomaly score at time t (i.e., the probability that \mathbf{x}_t comes from an anomalous event) by using deep neural networks. Among various types of autoregressive networks designed for temporal data, we employ the basic architecture of dilated CNN (DiCNN) introduced by WaveNet [27]. As illustrated in Figure 2 (Right), DiCNN is basically stacks of dilated *causal* convolutions, which is the convolution operation not considering future inputs when computing the output at each point. Specifically, the n -th layer has a convolution with filter size k and dilation rate k^{n-1} , thus the output vector from layer n at time t (denoted by \mathbf{h}_t) depends on its previous k^n points in the end; i.e., $[\mathbf{x}_{t-k^n+1}, \dots, \mathbf{x}_t]$ becomes its receptive field.

Using the inner product of the output vector at each point $\mathbf{h}_t \in \mathbb{R}^d$ and the anomaly weight vector $\mathbf{w} \in \mathbb{R}^d$, whose parameters are trainable, we compute the series of local anomaly scores $\mathbf{s} \in \mathbb{R}^T$ by

$$\mathbf{s} = [s_1, s_2, \dots, s_T], \text{ where } s_t = \sigma(\mathbf{w}^\top \mathbf{h}_t). \quad (2)$$

We apply the sigmoid function σ to make the score range in $[0, 1]$. Finally, we define the cost $\delta(z_l, \mathbf{x}_t)$ by the negative log-posterior probability using the anomaly score s_t :

$$\begin{aligned} \delta(z_l, \mathbf{x}_t) &= -\log p(z_l | \mathbf{x}_t) \\ &= -\{z_l \cdot \log s_t + (1 - z_l) \cdot \log(1 - s_t)\}. \end{aligned} \quad (3)$$

The cost becomes large when the computed score s_t is far from a target label z_l . In this sense, DTW can identify the anomalous segments by finding the optimal point-to-point matching between the series of anomaly scores and its sequential label, which minimizes the total alignment cost.

3.4. Learning with Weak Supervision

Our proposed framework, termed as WETAS, optimizes the DiCNN model by leveraging only weak supervision for the anomaly segmentation. To fully utilize the given instance-level labels, two different types of losses are considered: 1) the classification loss for correctly classifying an input instance as its instance-level anomaly label, and 2) the alignment loss for matching the input instance well with the sequential label, which is synthesized by the model by distilling the instance-level label.

Learning from Instance-level Label. Similar to the *local* (or point-level) anomaly score in Equation (2), we define the *global* (or instance-level) anomaly score for its binary classification. The global anomaly score $s_* \in \mathbb{R}$ is computed by using the anomaly weight $\mathbf{w} \in \mathbb{R}^d$ and the global output vector $\mathbf{h}_* \in \mathbb{R}^d$, obtained by global max (or average) pooling on all the output vectors along the time axis,

$$\begin{aligned} \mathbf{h}_* &= \text{global-pooling}(\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_T) \\ s_* &= \sigma(\mathbf{w}^\top \mathbf{h}_*). \end{aligned} \quad (4)$$

To optimize our model based on the instance-level labels, we define the classification loss by the binary cross entropy between y and s_* . Thereby, we can predict the anomaly label for an input instance based on its global anomaly score:

$$\mathcal{L}_c = -\sum_{i=1}^N \left\{ y^{(i)} \cdot \log s_*^{(i)} + (1 - y^{(i)}) \cdot \log(1 - s_*^{(i)}) \right\}. \quad (5)$$

The classification loss optimizes the model to differentiate

between normal patterns and anomalous patterns that are commonly observed in weakly labeled temporal data.

Learning from Sequential Pseudo-Label. To begin with, we propose the technique to infer the sequential anomaly label (i.e., pseudo-label generation) of an input instance, then utilize the generated sequential pseudo-label to further improve our DiCNN model so that it can compute more accurate local anomaly scores. The pseudo-labeling technique is motivated by temporal class activation map (CAM) [20, 37] which helps to analyze temporal regions that most influence the instance-level classification. We obtain the anomaly activation map $\mathbf{m} = [m_1, m_2, \dots, m_T] \in \mathbb{R}^T$ by multiplying each output vector \mathbf{h}_t with the weight vector \mathbf{w} that is used for the instance-level anomaly classification. That is, m_t becomes proportional to how strongly the time point t contributes to classifying the input instance as the anomaly. For the condition $0 \leq m_t \leq 1$, their values are min-max normalized along the time axis within each instance,

$$m_t = \frac{\mathbf{w}^\top \mathbf{h}_t - \min_{t'} \mathbf{w}^\top \mathbf{h}_{t'}}{\max_{t'} \mathbf{w}^\top \mathbf{h}_{t'} - \min_{t'} \mathbf{w}^\top \mathbf{h}_{t'}}. \quad (6)$$

Then, we introduce a pseudo-labeling function $\phi_L : \mathbb{R}^T \rightarrow \{0, 1\}^L$ which partitions a whole anomaly activation map into L segments (or disjoint intervals) of the same length $\lceil T/L \rceil$ and generates a binary label for each segment, i.e., $\tilde{\mathbf{z}} = \phi_L(\mathbf{m}) \in \{0, 1\}^L$. The pseudo-label \tilde{z}_l is determined by whether the maximum activation value in the l -th segment is larger than the anomaly threshold τ or not:

$$\tilde{z}_l = \mathbb{I} \left[\left(\max_{(l-1) \cdot \lceil T/L \rceil \leq t \leq l \cdot \lceil T/L \rceil} m_t \right) \geq \tau \right]. \quad (7)$$

This sequential pseudo-label $\tilde{\mathbf{z}}$ offers information on how many anomalous segments do exist and also their relative locations. Note that it is used as the rough segmentation mask for the DTW alignment, explained in Section 3.2.

The alignment loss is designed to reduce the DTW discrepancy between an input \mathbf{X} and its sequential pseudo-label $\tilde{\mathbf{z}}$ for their better alignment. However, the gradient of DTW is not very well defined (i.e., not differentiable) with respect to the cost matrix, due to its discontinuous hard-min operation taking only the minimum value. Thus, we instead adopt the soft-DTW [8] that provides the continuous relaxation of DTW by combining it with global alignment kernels [9]. The soft-DTW utilizes the soft-min operation,

$$\text{DTW}_\gamma(\tilde{\mathbf{z}}, \mathbf{X}) = \min_\gamma \{ \langle A, \Delta(\tilde{\mathbf{z}}, \mathbf{X}) \rangle, \forall A \in \mathcal{A} \} \quad (8)$$

where \min_γ with a smoothing parameter γ is defined by

$$\min_\gamma \{ a_1, a_2, \dots, a_n \} = \begin{cases} \min_{i \leq n} a_i, & \gamma = 0 \\ -\gamma \log \sum_{i=1}^n e^{-a_i/\gamma}, & \gamma > 0. \end{cases}$$

The soft-DTW distance in Equation (8) can be obtained by solving a dynamic program based on Bellman's recursions.

Please refer to [8, 20] for its forward and backward recursions to compute $\text{DTW}_\gamma(\tilde{\mathbf{z}}, \mathbf{X})$ and $\nabla_{\Delta(\tilde{\mathbf{z}}, \mathbf{X})} \text{DTW}_\gamma(\tilde{\mathbf{z}}, \mathbf{X})$.

We remark that only minimizing $\text{DTW}_\gamma(\tilde{\mathbf{z}}, \mathbf{X})$ causes the problem of *degenerated alignment* [7] that assigns each label to a single time point. To obtain more precise boundaries of segmentation results, we use the discriminative modeling under the supervision of the obtained pseudo-label. In detail, we optimize the model so that the alignment with a positive pseudo-label $\tilde{\mathbf{z}}_+$ costs less than that with a negative pseudo-label $\tilde{\mathbf{z}}_-$ by a margin β as follows.

$$\mathcal{L}_a = \sum_{i=1}^N \left[\frac{1}{T} \text{DTW}_\gamma(\tilde{\mathbf{z}}_+^{(i)}, \mathbf{X}^{(i)}) - \frac{1}{T} \text{DTW}_\gamma(\tilde{\mathbf{z}}_-^{(i)}, \mathbf{X}^{(i)}) + \beta \right]_+, \quad (9)$$

where $[z]_+ = \max(z, 0)$. The positive and negative pseudo-labels are respectively obtained by $\tilde{\mathbf{z}}_+ = y \cdot \phi_L(\mathbf{m})$ and $\tilde{\mathbf{z}}_- = (1 - y) \cdot \phi_L(\mathbf{m})$, thereby the instance-level label y plays the role of a binary mask. In other words, we regard the sequential pseudo-label corrupted by the wrong instance-level label (i.e., $1 - y$) as negative. This loss is helpful to produce more accurate local anomaly scores by making them align better with its sequential pseudo-label. Consequently, the model is capable of exploiting richer supervision than only using the instance-level labels.

3.5. Optimization and Inference

The final loss is described by the sum of the two losses, $\mathcal{L} = \mathcal{L}_c + \mathcal{L}_a$, where we can control the importance of the alignment loss by adjusting the margin size β . Note that the anomaly weight vector \mathbf{w} and the network parameters in the DiCNN model are effectively optimized by minimizing the final loss, whereas L , τ , and β are the hyperparameters of WETAS. Figure 2 shows the overall framework of WETAS.

The segmentation result for a test instance is obtained by the DTW alignment with its sequential pseudo-label $\tilde{\mathbf{z}} = \hat{y}_* \cdot \phi_L(\mathbf{m})$, where $\hat{y}_* = \mathbb{I}[s_* \geq \tau_*]$. Since the instance-level label is not given for the test input, we instead impose the predicted label \hat{y}_* as a binary mask; that is, we filter out normal instances based on the global anomaly score s_* . Once the model is trained, τ_* is automatically determined by its optimal value that achieves the best instance-level classification performance on the validation set.

4. Experiments

4.1. Experimental Settings

Dataset. For our experiments, we use four real-world temporal datasets collected from a range of tasks for detecting anomalous events, including multivariate time-series (MTS) and surveillance videos (Table 2). We split the set of all the instances by 5:2:3 ratio into a training set, a validation set, and a test set. Note that only the instance-level labels are given for the training and validation set.

- **Electromyography Dataset⁵ (EMG)** [23]: The 8-channel myographic signals recorded by bracelets worn on subjects’ forearms. Among several types of gestures, ulnar deviation is considered as anomalous events. Each time-series instance contains the signals for 5 seconds, which are downsampled to 500 points.
- **Gasoil Plant Heating Loop Dataset⁶ (GHL)** [10]: The control sequences of a gasoil plant heating loop, which suffered cyber-attacks. As done in [39], we crop 10 different instances of length 50,000 from each time-series, then downsample each of them to 1,000 points.
- **Server Machine Dataset⁷ (SMD)** [35]: The 38-metric multivariate time-series from server machines over 5 weeks, collected from an internet company. We split them every 720 points (i.e., 12 hours) to build the set of time-series instances.
- **Subway Exit Dataset (Subway)** [2]: The surveillance video for a subway exit gate, where each anomalous event corresponds to the passenger walking towards a wrong direction. We extract the visual features of each frame by using the pre-trained ResNet-34 [13], and make a single video instance include 450 frames.

Baselines. We compare the performance of WETAS with that of anomaly detection methods for temporal data, which are based on three different approaches.

- **Unsupervised learning:** VAE-based methods that compute the point-level anomaly scores from the reconstruction of temporal inputs — Donut [40], LSTM-VAE [30], LSTM-NDT [15], and OmniAnomaly [35].
- **Semi-supervised learning:** The variants of the unsupervised methods, whose models are trained by using only normal instances in order to make them utilize given instance-level labels⁸ — Donut⁺⁺, LSTM-VAE⁺⁺, LSTM-NDT⁺⁺, and OmniAnomaly⁺⁺.
- **Weakly supervised learning:** The multiple-instance learning method that can produce the anomaly prediction for each fixed-length segment — DeepMIL [36].

For fair comparisons, DeepMIL employs the same model architecture with WETAS (i.e., DiCNN). Moreover, we consider different numbers of the segments in a single instance, denoted by DeepMIL-4, 8, 16. In case of the video dataset (i.e., Subway), the additional comparison with other frame-level video anomaly detection methods [16, 17, 22, 29] is presented in the supplementary material.

Evaluation Metrics. For quantitative evaluation of the anomaly detection (and segmentation) results, we measure

⁵<http://archive.ics.uci.edu/ml/datasets/EMG+data+for+gestures>

⁶https://kas.pr/ics-research/dataset_ghl_1

⁷<https://github.com/smallcowbaby/OmniAnomaly>

⁸These methods are categorized as semi-supervised learning [6] in that they only utilize densely-labeled (normal) points from normal instances.

Table 2: Statistics of real-world temporal datasets.

Dataset	EMG	GHL	SMD	Subway
Data Type	MTS	MTS	MTS	Video
#Variables	8	19	38	1024
#Points (Train)	211,892	240,000	354,200	32,051
#Points (Valid)	84,734	96,000	141,670	13,050
#Points (Test)	127,199	144,000	212,550	19,800
Anomaly Ratio (%)	5.97	0.49	4.16	2.60
Instance Length	500	1,000	720	450

the F1-score (denoted by F1) and the intersection over union (denoted by IoU) by using the point-level ground truth: $F1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$ where $\text{Precision} = \frac{TP}{TP+FP}$ and $\text{Recall} = \frac{TP}{TP+FN}$, and $\text{IoU} = \frac{TP}{TP+FP+FN}$. Several previous work [30, 36, 40] mainly reported the area under the ROC curve (AUROC) as their evaluation metric, but it is known to produce misleading results especially for severely imbalanced classification with few samples of the minority class [33], such as anomaly detection tasks.

To compute the precision and recall, the weakly supervised methods (i.e., DeepMIL and WETAS) can find the optimal anomaly threshold (applied to instance-level or segment-level scores) by utilizing the instance-level labels from the validation set. However, in cases of the unsupervised and semi-supervised methods, they require point-level anomaly labels to tune the anomaly threshold (applied to point-level scores), thus we report their F1 and IoU by using the anomaly threshold that yields the best F1 among all possible thresholds (denoted by F1-best and IoU-best, respectively). This can be interpreted as a measure of the discrimination power between normal and anomalous points. For some of the baselines (i.e., LSTM-NDT and OmniAnomaly) that have their own techniques to automatically determine the anomaly threshold, we additionally report their F1 and IoU by using the selected threshold.

Implementation Details. We implement our WETAS and all the baselines using PyTorch, and train them with the Adam optimizer [19]. For the unsupervised methods, we tune their hyperparameters in the ranges suggested by the previous work [35] that considered the same baselines. In case of VAE-based methods, we set the size of a temporal context for each point (i.e., sliding window) to 128 (for MTS) and 16 (for video). For DiCNN-based methods, we stack 7 (for MTS) and 4 (for video) layers of dilated convolutions with filter size 2 to keep the size of its receptive field ($=2^7, 2^4$) the same with the others’. The dimensionality d of hidden (and output) vectors in DiCNN and the smoothing factor γ of soft-DTW are set to 128 and 0.01, respectively. We provide the in-depth sensitivity analysis on the hyperparameters (i.e., L, τ , and β) in our supplementary material.

4.2. Experimental Results

Performances on Anomaly Detection. We first measure

Table 3: Performances of WETAS and other baselines in detecting anomalous points from real-world temporal data. All the results support statistically significant improvement of WETAS over the best baseline ($p \leq 0.05$ from the paired t -test).

Methods	EMG				GHL			
	F1	IoU	F1-best	IoU-best	F1	IoU	F1-best	IoU-best
Donut [40]	-	-	0.1748(0.002)	0.0958(0.001)	-	-	0.0363(0.013)	0.0185(0.007)
LSTM-VAE [30]	-	-	0.1728(0.000)	0.0946(0.000)	-	-	0.0746(0.015)	0.0388(0.008)
LSTM-NDT [15]	0.1317(0.016)	0.0705(0.009)	0.1880(0.010)	0.1199(0.001)	0.0640(0.030)	0.0332(0.016)	0.1025(0.023)	0.0541(0.013)
OmniAnomaly [35]	0.1574(0.003)	0.0854(0.002)	0.1793(0.001)	0.0985(0.001)	0.0611(0.033)	0.0318(0.018)	0.0743(0.026)	0.0387(0.014)
Donut++ [40]	-	-	0.1784(0.001)	0.0980(0.001)	-	-	0.0850(0.036)	0.0447(0.019)
LSTM-VAE++ [30]	-	-	0.1745(0.000)	0.0956(0.000)	-	-	0.0828(0.023)	0.0433(0.013)
LSTM-NDT++ [15]	0.1344(0.005)	0.0720(0.003)	0.1916(0.011)	0.1201(0.002)	0.0984(0.091)	0.0538(0.053)	0.1043(0.106)	0.0578(0.064)
OmniAnomaly++ [35]	0.1536(0.002)	0.0832(0.001)	0.1807(0.002)	0.0993(0.001)	0.1209(0.101)	0.0668(0.058)	0.2096(0.143)	0.1231(0.094)
DeepMIL-4 [36]	0.4699(0.022)	0.3073(0.019)	-	-	0.0690(0.035)	0.0359(0.098)	-	-
DeepMIL-8 [36]	0.4317(0.029)	0.2755(0.023)	-	-	0.0571(0.023)	0.0323(0.015)	-	-
DeepMIL-16 [36]	0.3182(0.056)	0.1902(0.039)	-	-	0.1497(0.040)	0.0813(0.023)	-	-
WETAS (ours)	0.5803(0.068)	0.4118(0.064)	-	-	0.2295(0.028)	0.1298(0.018)	-	-

Methods	SMD				Subway			
	F1	IoU	F1-best	IoU-best	F1	IoU	F1-best	IoU-best
Donut [40]	-	-	0.3206(0.011)	0.1909(0.008)	-	-	0.5080(0.018)	0.3406(0.016)
LSTM-VAE [30]	-	-	0.2671(0.018)	0.1542(0.012)	-	-	0.5329(0.024)	0.3635(0.022)
LSTM-NDT [15]	0.1145(0.018)	0.0608(0.010)	0.3588(0.019)	0.2187(0.014)	0.0000(0.000)	0.0000(0.000)	0.5658(0.005)	0.3945(0.005)
OmniAnomaly [35]	0.1176(0.002)	0.0625(0.001)	0.1223(0.006)	0.0651(0.004)	0.5367(0.043)	0.3704(0.038)	0.6065(0.020)	0.4355(0.020)
Donut++ [40]	-	-	0.2875(0.067)	0.1693(0.047)	-	-	0.5100(0.026)	0.3449(0.021)
LSTM-VAE++ [30]	-	-	0.2477(0.015)	0.1414(0.010)	-	-	0.5452(0.016)	0.3749(0.016)
LSTM-NDT++ [15]	0.1211(0.010)	0.0645(0.006)	0.3819(0.020)	0.2361(0.015)	0.0000(0.000)	0.0000(0.000)	0.5723(0.004)	0.4009(0.004)
OmniAnomaly++ [35]	0.1435(0.081)	0.0790(0.049)	0.1750(0.077)	0.0974(0.047)	0.5479(0.028)	0.3790(0.026)	0.6198(0.023)	0.4494(0.024)
DeepMIL-4 [36]	0.3561(0.052)	0.2176(0.038)	-	-	0.5138(0.081)	0.3738(0.073)	-	-
DeepMIL-8 [36]	0.3450(0.032)	0.2088(0.023)	-	-	0.6471(0.066)	0.4885(0.064)	-	-
DeepMIL-16 [36]	0.3568(0.016)	0.2173(0.012)	-	-	0.6102(0.077)	0.4391(0.072)	-	-
WETAS (ours)	0.4358(0.046)	0.2795(0.037)	-	-	0.7414(0.023)	0.5907(0.028)	-	-

Table 4: F1-scores of WETAS that ablates each component, Dataset: SMD. C and R denote the ‘‘corruption’’ and ‘‘random sampling’’ for negative pseudo-labels, respectively.

Model Arch.	\mathcal{L}_c	\mathcal{L}_a	DTW Seg.	Global Pool.	Neg. Label	F1
LSTM	✓	✓	✓	MAX	C	0.3685
	✓	-	-	MAX	C	0.1225
	✓	✓	-	MAX	C	0.1265
	✓	-	✓	MAX	C	0.3384
DiCNN	✓	✓	✓	AVG	C	0.3046
	✓	✓	✓	MAX	R	0.4272
	✓	✓	✓	MAX	C	0.4358

the detection performance of WETAS and the other baselines. In this experiment, we do not consider the point-adjust approach [35, 40] for our evaluation strategy: if any point in a ground truth anomalous segment is detected as the anomaly, all points in the segment are considered to be correctly detected as the anomalies.⁹ We repeat to train each model five times using different random seeds, and report the averaged results with their standard deviations.

In Table 3, the unsupervised and semi-supervised methods show considerably worse performance than the weakly supervised methods in terms of both F1 and IoU. Even considering F1-best and IoU-best, their performance is at most comparable to that of the weakly supervised methods. In

other words, though they additionally use the point-level labels for finding the best anomaly threshold that clearly distinguishes anomalous points from normal ones, they cannot achieve the detection performance as high as the weakly supervised methods due to the lack of supervision.

Compared to the unsupervised methods that use all training instances regardless of their instance-level labels, the semi-supervised methods that selectively use them (i.e., only normal instances) sometimes perform better and sometimes worse. That is, learning the normality of temporal data from anomalous instances can degrade the detection performance due to anomalous patterns that reside in the instances; but in some cases, the increasing number of training data points is rather helpful to improve the generalization power despite the existence of anomalous points. This strongly indicates that how to utilize the instance-level labels does matter and largely affects the final performance.

Our WETAS achieves the best performance for all the datasets, and notably, it significantly beats DeepMIL whose performance highly depends on the number of segments. For the EMG dataset whose instance length is only 500, DeepMIL-16 is not able to effectively detect anomalies because each input segment is too short to accurately compute the segment-level anomaly score; for the GHL dataset, DeepMIL-4 suffers from coarse-grained segmentation, which leads to the limited F1 and IoU. On the contrary, WETAS successfully finds anomalous segments by using dynamic alignment with the sequential pseudo-label.

⁹This approach skews the results by excessively increasing the true positives, i.e., the F1 and IoU are overestimated.

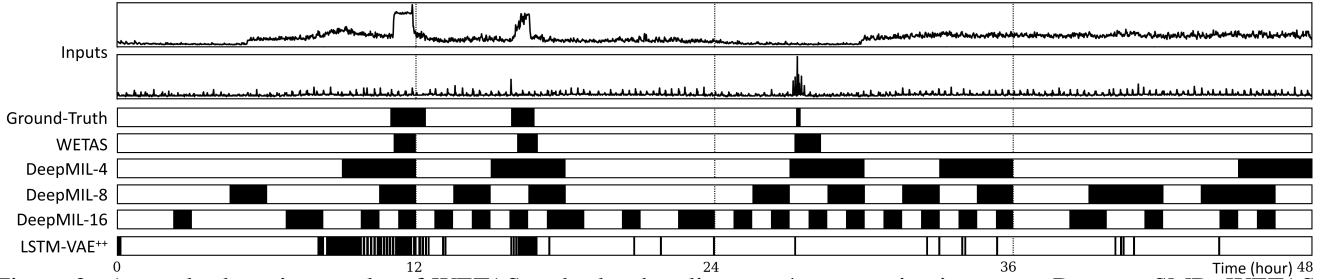


Figure 3: Anomaly detection results of WETAS and other baselines on 4 consecutive instances, Dataset: SMD. WETAS correctly detects anomalous points in terms of the number of ground truth anomalous segments and their boundaries.

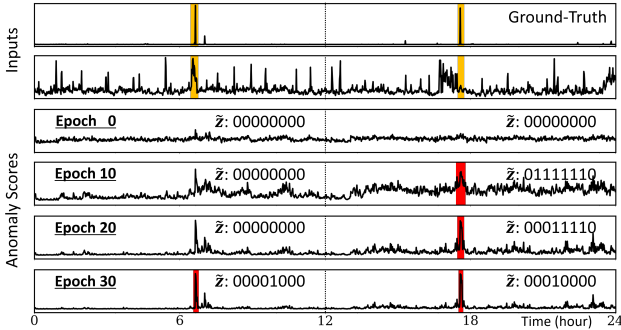


Figure 4: Test instances (upper) and the local anomaly scores computed by WETAS every 10 epoch (lower), Dataset: SMD. The ground truth and identified anomalous segments are colored in yellow and red, respectively.

Ablation Study. To further investigate the effectiveness of WETAS, we measure the F1-score of several variants that ablate each of the following components: 1) DiCNN (vs. LSTM) for capturing the temporal dependency, 2) DTW-based alignment loss, 3) DTW-based segmentation, 4) global max (vs. average) pooling for the classification loss, and 5) corruption by the instance-level label (vs. random sampling) for the negative pseudo-label \tilde{z}_- in Equation (9). In case of WETAS without the DTW-based segmentation, we measure both the F1¹⁰ and F1-best using the series of local (or point-level) anomaly scores.

Table 4 summarizes the results of our ablation study on the SMD dataset. It is worth noting that the DTW-based segmentation considerably improves the performance compared to the case simply using the obtained local anomaly scores. This is because the alignment with the pseudo-label filters out normal instances based on the global anomaly score, and also considers the continuity of anomalous points while limiting the number of anomalous segments. Nevertheless, their F1-best scores (for Rows 2 and 3) are respectively 0.4472 and 0.4590, which means that WETAS outperforms other point-level detection methods even without the DTW-based segmentation. In addition, it is obvious that DiCNN is more effective to capture the long-term dependency

¹⁰For this measure, we set the anomaly threshold (applied to our point-level anomaly scores) to τ that is used for pseudo-labeling in Equation (7).

of temporal data than LSTM. The global max pooling and the corrupted pseudo-labels also turn out to help the DiCNN model to learn useful features from the instance-level labels, compared to their alternatives.

Qualitative Analyses. We also qualitatively compare the anomaly detection results for test instances. Figures 3 and 4 present the results on consecutive instances for the SMD dataset. We plot only the two series of input variables that are directly related to observed anomalous events.

Compared to the baselines, WETAS more accurately finds anomalous segments in terms of their number as well as the boundaries. To be specific, DeepMIL-4 with fewer segments correctly identifies the segments that include anomalous events, but the boundaries are far from the ground truth due to their fixed length and location. In contrast, DeepMIL-16 with more segments makes unreliable predictions, because each input segment is not long enough that the normality and abnormality of the segment are effectively captured. The point-level detection method, LSTM-VAE⁺⁺, finds out non-consecutive anomalous points, which results in a large number of separate segments. It is obvious that such discontinuous prediction makes it difficult to investigate the occurrence of anomalous events.

5. Conclusion

This paper proposes a weakly supervised learning framework for temporal anomaly segmentation, which effectively finds anomalous segments by leveraging instance-level anomaly labels for model training. For each input instance of temporal data, the learning objective of WETAS includes its accurate classification as the instance-level label and also strong alignment with the sequential pseudo-label. In the end, WETAS utilizes the DTW alignment between an input instance and its pseudo-label to obtain the temporal segmentation result. Our empirical evaluation shows that WETAS outperforms all baselines in detecting temporal anomalies as well as specifying precise segments that correspond to anomalous events.

Acknowledgement. This work was supported by the NRF grant (No. 2020R1A2B5B03097210) and the IITP grant (No. 2018-0-00584, 2019-0-01906) funded by the MSIT.

References

- [1] Aisha Abdallah, Mohd Aizaini Maarof, and Anazida Zainal. Fraud detection system: A survey. *Journal of Network and Computer Applications*, 68:90–113, 2016. [1](#)
- [2] Amit Adam, Ehud Rivlin, Ilan Shimshoni, and Daviv Reinitz. Robust real-time unusual event detection using multiple fixed-location monitors. *TPAMI*, 30(3):555–560, 2008. [6](#)
- [3] Subutai Ahmad, Alexander Lavin, Scott Purdy, and Zuha Agha. Unsupervised real-time anomaly detection for streaming data. *Neurocomputing*, 262:134–147, 2017. [1](#), [2](#)
- [4] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtec ad—a comprehensive real-world dataset for unsupervised anomaly detection. In *CVPR*, pages 9592–9600, 2019. [1](#)
- [5] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings. In *CVPR*, pages 4183–4192, 2020. [1](#)
- [6] Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM computing surveys*, 41(3):1–58, 2009. [1](#), [6](#)
- [7] Chien-Yi Chang, De-An Huang, Yanan Sui, Li Fei-Fei, and Juan Carlos Niebles. D3tw: Discriminative differentiable dynamic time warping for weakly supervised action alignment and segmentation. In *CVPR*, pages 3546–3555, 2019. [1](#), [2](#), [3](#), [5](#)
- [8] Marco Cuturi and Mathieu Blondel. Soft-dtw: a differentiable loss function for time-series. In *ICML*, pages 894–903, 2017. [5](#)
- [9] Marco Cuturi, Jean-Philippe Vert, Oystein Birkenes, and Tomoko Matsui. A kernel for time series based on global alignments. In *ICASSP*, volume 2, pages II–413, 2007. [5](#)
- [10] Pavel Filonov, Andrey Lavrentyev, and Artem Vorontsov. Multivariate industrial time series with cyber-attack simulation: Fault detection using an lstm-based predictive data model. *arXiv preprint arXiv:1612.06676*, 2016. [6](#)
- [11] Nico Görnitz, Marius Kloft, Konrad Rieck, and Ulf Brefeld. Toward supervised anomaly detection. *JAIR*, 46:235–262, 2013. [2](#)
- [12] Mahmudul Hasan, Jonghyun Choi, Jan Neumann, Amit K Roy-Chowdhury, and Larry S Davis. Learning temporal regularity in video sequences. In *CVPR*, pages 733–742, 2016. [1](#), [2](#)
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [6](#)
- [14] Dan Hendrycks, Steven Basart, Mantas Mazeika, Mohammadreza Mostajabi, Jacob Steinhardt, and Dawn Song. A benchmark for anomaly segmentation. *arXiv preprint arXiv:1911.11132*, 2019. [1](#)
- [15] Kyle Hundman, Valentino Constantinou, Christopher Laporte, Ian Colwell, and Tom Soderstrom. Detecting spacecraft anomalies using lstms and nonparametric dynamic thresholding. In *KDD*, pages 387–395, 2018. [1](#), [2](#), [6](#), [7](#)
- [16] Radu Tudor Ionescu, Sorina Smeureanu, Bogdan Alexe, and Marius Popescu. Unmasking the abnormal events in video. In *ICCV*, pages 2895–2903, 2017. [6](#)
- [17] Radu Tudor Ionescu, Sorina Smeureanu, Marius Popescu, and Bogdan Alexe. Detecting abnormal events in video using narrowed normality clusters. In *WACV*, pages 1951–1960, 2019. [6](#)
- [18] Hyunjun Ju, Dongha Lee, Junyoung Hwang, Junghyun Namkung, and Hwanjo Yu. Pumad: Pu metric learning for anomaly detection. *Information Sciences*, 523:167–183, 2020. [1](#)
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [20] Dongha Lee, Seonghyeon Lee, and Hwanjo Yu. Learnable dynamic temporal pooling for time series classification. In *AAAI*, volume 35, pages 8288–8296, 2021. [5](#)
- [21] Jungbeom Lee, Eunji Kim, Sungmin Lee, Jangho Lee, and Sungroh Yoon. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In *CVPR*, pages 5267–5276, 2019. [1](#)
- [22] Yusha Liu, Chun-Liang Li, and Barnabás Póczos. Classifier two sample test for video anomaly detections. In *BMVC*, page 71, 2018. [6](#)
- [23] Sergey Lobov, Nadia Krilova, Innokentiy Kastalskiy, Victor Kazantsev, and Valeri A Makarov. Latent factors limiting the performance of semg-interfaces. *Sensors*, 18(4):1122, 2018. [6](#)
- [24] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015. [2](#)
- [25] Weixin Luo, Wen Liu, and Shenghua Gao. Remembering history with convolutional lstm for anomaly detection. In *ICME*, pages 439–444, 2017. [1](#), [2](#)
- [26] Mohsin Munir, Shoaib Ahmed Siddiqui, Andreas Dengel, and Sheraz Ahmed. Deepant: A deep learning approach for unsupervised anomaly detection in time series. *IEEE Access*, 7:1991–2005, 2018. [1](#), [2](#)
- [27] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016. [4](#)
- [28] Guansong Pang, Chunhua Shen, and Anton van den Hengel. Deep anomaly detection with deviation networks. In *KDD*, pages 353–362, 2019. [2](#)
- [29] Guansong Pang, Cheng Yan, Chunhua Shen, Anton van den Hengel, and Xiao Bai. Self-trained deep ordinal regression for end-to-end video anomaly detection. In *CVPR*, pages 12173–12182, 2020. [6](#)
- [30] Daehyung Park, Yuuna Hoshi, and Charles C Kemp. A multimodal anomaly detector for robot-assisted feeding using an lstm-based variational autoencoder. *IEEE Robotics and Automation Letters*, 3(3):1544–1551, 2018. [1](#), [2](#), [6](#), [7](#)
- [31] Alexander Richard, Hilde Kuehne, Ahsan Iqbal, and Juergen Gall. Neuralnetwork-viterbi: A framework for weakly supervised video learning. In *CVPR*, pages 7386–7395, 2018. [1](#), [2](#), [3](#)

- [32] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015. [2](#)
- [33] Takaya Saito and Marc Rehmsmeier. The precision-recall plot is more informative than the roc plot when evaluating binary classifiers on imbalanced datasets. *PloS one*, 10(3):e0118432, 2015. [6](#)
- [34] Hiroaki Sakoe and Seibi Chiba. Dynamic programming algorithm optimization for spoken word recognition. *IEEE transactions on acoustics, speech, and signal processing*, 26(1):43–49, 1978. [2](#), [3](#)
- [35] Ya Su, Youjian Zhao, Chenhao Niu, Rong Liu, Wei Sun, and Dan Pei. Robust anomaly detection for multivariate time series through stochastic recurrent neural network. In *KDD*, pages 2828–2837, 2019. [1](#), [2](#), [6](#), [7](#)
- [36] Waqas Sultani, Chen Chen, and Mubarak Shah. Real-world anomaly detection in surveillance videos. In *CVPR*, pages 6479–6488, 2018. [2](#), [3](#), [6](#), [7](#)
- [37] Zhiguang Wang, Weizhong Yan, and Tim Oates. Time series classification from scratch with deep neural networks: A strong baseline. In *IJCNN*, pages 1578–1585, 2017. [5](#)
- [38] Qi Wei, Yin hao Ren, Rui Hou, Bibo Shi, Joseph Y Lo, and Lawrence Carin. Anomaly detection for medical images based on a one-class classification. In *Medical Imaging 2018: Computer-Aided Diagnosis*, volume 10575, page 105751M, 2018. [1](#)
- [39] Tailai Wen and Roy Keyes. Time series anomaly detection using convolutional neural networks and transfer learning. *arXiv preprint arXiv:1905.13628*, 2019. [2](#), [6](#)
- [40] Haowen Xu, Wenxiao Chen, Nengwen Zhao, Zeyan Li, Jiahao Bu, Zhihan Li, Ying Liu, Youjian Zhao, Dan Pei, Yang Feng, et al. Unsupervised anomaly detection via variational auto-encoder for seasonal kpis in web applications. In *WWW*, pages 187–196, 2018. [1](#), [2](#), [6](#), [7](#)
- [41] Tan Yu, Zhou Ren, Yuncheng Li, Enxu Yan, Ning Xu, and Junsong Yuan. Temporal structure mining for weakly supervised action detection. In *ICCV*, pages 5522–5531, 2019. [2](#), [3](#)
- [42] Zehuan Yuan, Jonathan C Stroud, Tong Lu, and Jia Deng. Temporal action localization by structured maximal sums. In *CVPR*, pages 3684–3692, 2017. [2](#), [3](#)
- [43] Yong Zhang, Yu Zhang, Zhao Zhang, Jie Bao, and Yunpeng Song. Human activity recognition based on time series analysis using u-net. *arXiv preprint arXiv:1809.08113*, 2018. [2](#)