

Continual Test-time Domain Adaptation via Dynamic Sample Selection

Yanshuo Wang^{1,2}, Jie Hong^{1,2}, Ali Cheraghian^{1,2}, Shafin Rahman³
David Ahmed-Aristizabal², Lars Petersson², Mehrtash Harandi^{2,4}

¹Australian National University, ²Data61-CSIRO, Australia

³North South University, Bangladesh, ⁴Monash University, Australia

{Yanshuo.Wang, Ali.Cheraghian, david.ahmedtaristizabal, Lars.Petersson}@data61.csiro.au,
shafin.rahman@northsouth.edu, jie.hong@anu.edu.au, mehrtash.harandi@monash.edu

Abstract

The objective of Continual Test-time Domain Adaptation (CTDA) is to gradually adapt a pre-trained model to a sequence of target domains without accessing the source data. This paper proposes a Dynamic Sample Selection (DSS) method for CTDA. DSS consists of dynamic thresholding, positive learning, and negative learning processes. Traditionally, models learn from unlabeled unknown environment data and equally rely on all samples' pseudo-labels to update their parameters through self-training. However, noisy predictions exist in these pseudo-labels, so all samples are not equally trustworthy. Therefore, in our method, a dynamic thresholding module is first designed to select suspected low-quality from high-quality samples. The selected low-quality samples are more likely to be wrongly predicted. Therefore, we apply joint positive and negative learning on both high- and low-quality samples to reduce the risk of using wrong information. We conduct extensive experiments that demonstrate the effectiveness of our proposed method for CTDA in the image domain, outperforming the state-of-the-art results. Furthermore, our approach is also evaluated in the 3D point cloud domain, showcasing its versatility and potential for broader applicability.

1. Introduction

Consider an intelligent agent moving around non-stationary environments where the input domain gradually changes over time. As an illustration, a self-driving car could transition from daylight to darkness and then enter a snowy scenario. Another example is an automated robot processing images captured from multiple continual domains like blurry and bright domains. Therefore, the intelligent agent requires *Continual Test-Time domain Adaptation* (CTDA) to adapt to gradually changing environments. Because of its real-life applicability, researchers have recently

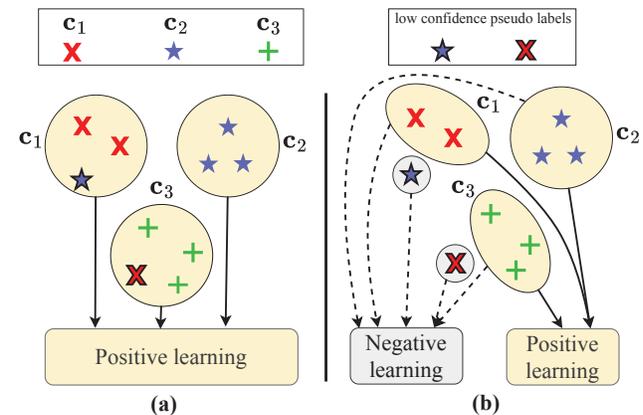


Figure 1. Illustration of our learning strategy: (a) Pseudo labeling described is accomplished by the mean-teacher model, where a teacher model yields a pseudo label for a given sample. Some of the unlabeled samples probably assign incorrect pseudo-labels. Hence, treating all samples equally leads to error accumulation. (b) Our proposed method introduces Dynamic Sample Selection (DSS), to monitor the pseudo label generated by the mean teacher model and it selects low-quality pseudo labels via our proposed dynamic thresholding process. Ultimately, we use positive learning for samples with high-quality pseudo-label and negative learning for all samples.

started exploring this new area [6, 27]. In contrast to traditional domain adaptation tasks, CTDA imposes some constraints: 1) There should be multiple continual test domains instead of one, and the model should gradually adapt to continual domains using unlabeled test data. 2) The source data with which the model is trained cannot be used during the adaptation process. Instead, it is only allowed to leverage the pre-trained model developed from the source data.

Error accumulation is a well-known challenge in CTDA. In this task, the model adapts continual domains during test time on unlabeled test samples. The model normally relies on samples' pseudo-labels for the adaptation. However,

not all pseudo-labels are reliable. As the erroneous pseudo-labels arise from the continual domains, the network tends to absorb inaccurate feedback, leading to an incorrect adjustment of its parameters over time. This accumulation of errors from the previous domains would have an adverse impact on the model’s performance in the subsequent domains. To address this, the existing SOTA method [27] generates augmented samples of a given input sample and calculates an augmentation-averaged prediction score from a teacher model. The assumption is that pseudo-labels calculated from the augmentation-averaged prediction score become more accurate over time.

To address the aforementioned problem, we introduce a Dynamic Sample Selection (DSS) method, which incorporates jointly dynamic thresholding, positive, and negative learning. More specifically, to mitigate the issue of error accumulation, using dynamic thresholding, we design a class-wise confidence metric to evaluate the prediction quality of the pseudo-label for each sample in a batch. Dynamic thresholding monitors the confidence and identifies high- and low-quality samples. As shown in Figure 1, our proposed method treats the unlabeled samples differently based on confidence. We employ only the samples with higher confidence for positive learning, thus reducing the amount of noise introduced. At the same time, we utilize all examples for negative learning. The negative learning aims to eliminate the impact of low-confidence class prediction. A prediction that is potentially considered wrong is replaced by a negative complementary label. Traditionally, CTDA methods are tested on only 2D image datasets. However, in addition to challenging 2D image datasets (CIFAR-10-C [10], CIFAR-100-C [10], and ImageNet-C [10]), we experimented with 3D point cloud data ScanObjectNN-C [22]), demonstrating the robustness of our proposed method. Our method consistently outperforms existing methods by a decent margin.

In summary, contributions of this work are three-fold:

- We develop a novel method, Dynamic Sample Selection (DSS) for the CTDA task. DSS selects high- and low-quality samples for training, and it effectively reduces the negative impact of error accumulation of CTDA.
- We apply joint positive and negative learning on noisy pseudo-labels to reduce the risk of propagating misleading information during the CTDA task.
- Our study is pioneering in investigating the CTDA problem in the domain of 3D point cloud objects. We have benchmarked the performance of existing 2D approaches on 3D data and identified new challenges associated with 3D point cloud objects.

2. Related work

Test-time Domain Adaptation. Compared with traditional unsupervised domain adaptation, Test-Time domain Adaptation (TTA) adapts the model trained from the source domain to the novel target domain without access to the original source data during inference time. One popular approach to reducing the domain gap in the absence of source data is to fine-tune the source model by adopting an unsupervised loss function based on the target distribution. Test entropy minimization (TENT) [26] updates the trainable batch normalization parameters from a pre-trained model at test time by minimizing the entropy of the model prediction. Source hypothesis transfer (SHOT) [17] employs an entropy minimization and diversity regularizer with label smoothing techniques to train a general feature extractor from a pre-trained source model. Alternatively, Test-Time Training (TTT) [23] introduced a self-supervised rotation prediction auxiliary task to update model parameters for novel target samples. Test-Time Training with Masked Autoencoders [8] similarly adopts the idea of masked autoencoders as an auxiliary task to train the model for each test sample. Recently AdaContrast [3] first leveraged contrastive learning with online pseudo refinement to learn better feature representations with less noisy pseudo labels. Conjugate PL [9] propose a general way of obtaining test-time adaptation loss which is used for more robust predictions under distribution shifts.

Continual Test-time Domain Adaptation. In addressing the challenge of Continual Test-time Domain Adaptation (CTDA), various solutions have been proposed. Notably, the online version of Tent, introduced by Wang *et al.* [26], presents an applicable approach. However, a drawback of earlier methods lies in their reliance on source data during inference. A significant advancement in this field is the Continual Test-time Adaptation Approach (CoTTA) developed by Wang *et al.* [27]. CoTTA is notably the first method explicitly tailored to the demands of online continual test-time adaptation. This method employs a weighted augmentation-averaged mean teacher framework, building upon the insights from prior work such as the mean teacher predictions introduced by Tarvainen and Valpola [24]. Remarkably, the student-teacher framework proposed by CoTTA [27] serves as a foundational architecture for numerous subsequent studies. In particular, Niu *et al.* [18], adopt a similar student-teacher framework. They incorporate continuous batch normalization statistics updates to reduce computational costs, thereby refining the efficiency of the adaptation process. Another notable avenue explored by researchers is the utilization of a mean teacher setup with symmetric cross-entropy and contrastive learning, as demonstrated in the work by Dobler *et al.* [6] under the name RMT. While this approach introduces valuable insights, it remains dependent on source data to establish the

source prototypes during a warm-up stage. Consequently, RMT cannot be categorized as a truly source-free approach within the realm of continual test-time domain adaptation. Alternatively, Gan *et al.* [7] present an innovative strategy that leverages visual prompt learning in conjunction with a homeostasis-based adapting strategy.

Pseudo Labelling. Pseudo-labelling is a widely used technique in semi-supervised learning or self-learning. It uses the model’s output class probability as a label for training. FixMatch [21] is a semi-supervised method that generates pseudo labels using the model’s predictions on weakly augmented unlabeled images. Then the model is trained to match the pseudo-label with the prediction on strongly-augmented images. In contrast, UPS [20] not only produces reliable pseudo-labels with high confidence and low uncertainty, but it also incorporates negative learning to further reduce the model calibration error. In the context of unsupervised domain adaptation, the process of pseudo-labeling frequently involves generating labels for unlabeled target samples by utilizing the predicted class probability of the source model. Existing methods [6, 27] mainly rely on pseudo labels as a form of “supervision” in order to compensate for the absence of ground truth labels in the target domain. However, they did not closely investigate the quality of pseudo labels, as mislabeled samples used in self-learning accelerate error accumulation ultimately. In contrast, in this paper, we propose joint positive and negative learning with dynamic threshold modules to minimize the effect of error accumulation from mislabeled pseudo labels.

3. Dynamic Sample Selection

3.1. Problem Formulation

Given a sequence of domains $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_t, \dots, \mathcal{D}_T\}$, we define the domain at time step t as $\mathcal{D}_t = \{\mathbf{x}_{t,i}\}_{i=1}^{n_t}$, where $\mathbf{x}_{t,i}$ is the i^{th} sample, and n_t is the number of samples. The first domain, \mathcal{D}_1 , is referred to as the source domain. We assume that a deep learning model $h_{\theta_1}(\cdot)$, where θ_1 is the pre-trained parameters on the source data $\mathbf{x}_{1,i}$ in \mathcal{D}_1 . However, we discard the source data thereafter, typically because of privacy or memory concerns. More precisely, we cannot access source data \mathcal{D}_1 for future new domains. Therefore, the future domains \mathcal{D}_t with $t > 1$, are termed as the target domains. Within the CTDA framework, our objective is to improve the performance of the model $h_{\theta_1}(\mathbf{x})$ during test time for a dynamically changing target domain in an online manner. This involves feeding the model unlabeled test data from the target domain sequentially. To be more specific, at target domain \mathcal{D}_t , the unlabeled data \mathbf{x}_t is given to the model $h_{\theta_t}(\mathbf{x}_t)$, and the model needs to make the prediction and adapt itself accordingly ($\theta_t \rightarrow \theta_{t+1}$) for the next target domain, \mathcal{D}_{t+1} . In the process of test-time training, we ensure that the model is adapted to each current

sample from the target domain. Following this adaptation, the updated model is utilized to predict the class label of the respective sample.

3.2. Model Overview

This section provides a brief overview of our proposed architecture, Dynamic Sample Selection (DSS), which is illustrated in Figure 2. Similar to other methods in Continual Target Domain Adaptation (CTDA) [27], we adopt the student and teacher setup from the mean teacher framework [24] as the foundational structure for self-learning in our approach. Given the input \mathbf{x}_t from domain \mathcal{D}_t , the teacher model computes the pseudo label $\hat{\mathbf{y}}_t$ based on an augmentation module, while the student model produces the output $\bar{\mathbf{y}}_t$. The model is then trained using the loss function that enforces consistency between $\hat{\mathbf{y}}_t$ and $\bar{\mathbf{y}}_t$, thus promoting alignment between the predicted labels. It is worth noting that $\hat{\mathbf{y}}_t$ is typically used as the final prediction.

As indicated in Figure 2, we do not treat all samples in \mathcal{D}_t equally for adaptation. In contrast, we assume that not all samples have enough quality for training during test time. To achieve this objective, we employ adaptive pseudo-labelling from recent advances in semi-supervised learning [28, 32], a learning approach that involves continuous monitoring of the model’s learning progress in accordance with its learning status [1]. In other words, we utilize a dynamic threshold module to modulate the learning progress adaptively each time we encounter new domains. Therefore, we design a DSS framework on top of the mean teacher model to select samples for different training strategies. To be more specific, given input \mathbf{x}_t in \mathcal{D}_t , we leverage the student-teacher setup to generate the prediction (or pseudo) label $\hat{\mathbf{y}}_t$. Using $\hat{\mathbf{y}}_t$, we compute the adaptive class-wise threshold $\pi_{t,c}(\hat{\mathbf{y}}_t)$ for class c . Based on $\pi_{t,c}$, we group high- and low-quality samples, $\mathbf{x}_{t,high}$ and $\mathbf{x}_{t,low}$. Then, both $\mathbf{x}_{t,high}$ and $\mathbf{x}_{t,low}$ are used for negative learning. Meanwhile, the model is trained with the positive learning loss with high-quality samples $\mathbf{x}_{t,high}$ alone.

3.3. Dynamic Thresholding

In CTDA, the learning progress of the model h_{θ} varies over time since different target domains have significant domain gaps. When adopting a model to new test samples from different domains, pseudo labels inevitably suffers from label noise due to the presence of domain shifts. To minimize noise in test-time, We adopt a threshold mechanism to adjust the confidence threshold dynamically.

As shown in Figure 2, the samples \mathbf{x}_t from \mathcal{D}_t are first given to the teacher model to compute the initial predictions $\hat{\mathbf{y}}_t$. Subsequently, taking inspiration from the threshold strategy [28, 32] and in order to track the learning progress, we implement a dynamic approach that involves detecting low-confidence predictions in samples through the use of

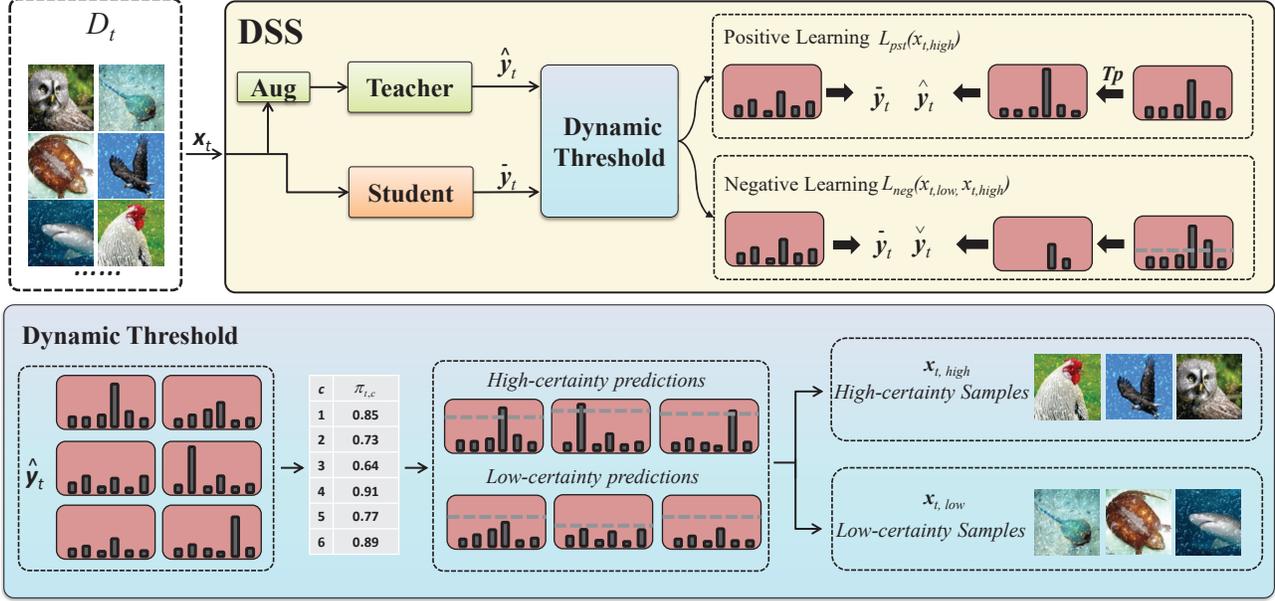


Figure 2. The proposed Dynamic Sample Selection (DSS) framework for continual test-time domain adaptation (CTDA). Unlabeled target domains are changing over time. At time step t , the model first makes the predictions for samples \mathbf{x}_t from D_t . Taking the input \mathbf{x}_t from domain D_t , the teacher model computes the pseudo label $\hat{\mathbf{y}}_t$ based on augmentation module while the student model gives the output $\bar{\mathbf{y}}_t$. The dynamic threshold module then outputs a suitable threshold to split the samples \mathbf{x}_t into two groups, high- and low-quality groups. Positive learning is exclusively applied to the high-quality group of samples, ensuring that the model focuses its learning efforts on reliable and accurate data. Conversely, in negative learning, complementary labels are generated for classes that yield low prediction scores. This approach enables the model to learn from and address the challenging samples actively, contributing to improved performance and adaptability in evolving domains.

a dynamic threshold. The common practice is to use a fixed confidence threshold to remove suspected noisy labels. However, model confidence on out-of-distribution samples could differ vastly for each domain, which may hurt the adaptation performance under a fixed threshold. Thus, we propose the adaptive threshold function to output a suitable selection threshold π_t :

$$\pi_{t,j} = \lambda \pi_{t,j-1} + (1 - \lambda) \cdot \frac{1}{N} \sum_{i=1}^N \max \hat{\mathbf{y}}_{t,j,i} \quad (1)$$

where j denotes the order of the batch. For simplicity, we drop the most j in the following context. π_t denotes the adaptive threshold value for the samples \mathbf{x}_t . $\hat{\mathbf{y}}_{t,i}$ is the predicted confidence vector for i^{th} sample. $\max \hat{\mathbf{y}}_{t,i}$ indicates the highest confidence value in $\hat{\mathbf{y}}_{t,i}$. The number of samples from each batch is denoted as N , and λ represents the exponential moving average factor.

An appropriate initial threshold could help the algorithm to converge more swiftly to an effective threshold in an online manner. It is important to note that DSS algorithm operates in an online fashion during the inference phase. Starting with a threshold value that is considerably distant from the current confidence level may result in a significant time

delay before reaching a satisfactory threshold. Therefore, we adopt the following approach to initialize the threshold at the start of each domain: we take the average of the previous model confidence at the beginning, denoted as $\pi_{t,0} = (\pi_{t-1,final} + \frac{1}{C})/2$ where $\pi_{t-1,final}$ is the computed threshold from the previous domain and C represents the number of classes. It should be mentioned that for the first domain, we set $\pi_{0,0} = \frac{1}{C}$. By incorporating the average of the previous domain's threshold, we aim to achieve an accurate estimation of the initial threshold by leveraging the knowledge acquired from the previous domains. Unlike other existing methods in semi-supervised learning [2], we use flexible threshold initialization in CTDA to provide a more appropriate initial value when facing a new domain instead of using a fixed average constant.

To reflect the discrepancy between the adaptation status among different classes, we further give the threshold in a class-wise manner to ensure that the class with average low confidence (low-certainty class) is not ignored. The class-wise manner guarantees sufficient samples for training as there are many low-certainty classes. The aim is to prevent the model from being biased toward the majority classes. We first compute the rescale ratio $\tau_{t,c}$ for class c as follows:

$$\tau_{t,c} = \frac{\delta_{t,c}}{\max \delta_t} \quad (2)$$

where δ_t is the vector where one element is the average confidence for one class. $\delta_{t,c}$ is the average confidence for class c of samples in D_t . The formulation $\delta_{t,c}$ is given as follows:

$$\delta_{t,c} = \frac{1}{N} \sum_{i=1}^N \max \hat{y}_{t,i} \cdot \mathbb{1}(\arg \max \hat{y}_{t,i} = c) \quad (3)$$

where $\max \hat{y}_{t,i}$ is the maximum prediction probability in the vector $\hat{y}_{t,i}$. Hence, the improved version of the adaptive threshold in a class-wise manner is calculated as:

$$\pi_{t,c} = \pi_t \cdot \tau_{t,c} \quad (4)$$

For D_t , after achieving $\pi_{t,c}$ for samples, we divide \mathbf{x}_t into the high- and low-quality groups, $\mathbf{x}_{t,high}$ and $\mathbf{x}_{t,low}$. $\mathbf{x}_{t,high}$ are samples where their predicted confidence is above the threshold $\pi_{t,c}$, and vice versa.

In the following algorithmic steps, the different training strategies will be used for $\mathbf{x}_{t,high}$ and $\mathbf{x}_{t,low}$. Overall, the self-adaptive threshold accurately identifies correctly- and wrongly-predicted samples as high- and low-certainty groups. Grouped samples will facilitate positive and negative learning, introduced in the following subsection.

3.4. Joint Positive and Negative Learning

Training the model in CTDA without access to data labels is challenging. To be more specific, assigning the computed pseudo label to the unlabeled data runs the risk of letting the model learn incorrect information as the pseudo label is noisy, which could lead to lower overall accuracy. To address this issue, as shown in Figure 2 we first group samples into high- and low-quality groups using dynamic threshold. Then, we train two groups of samples in two different ways (See ‘‘DSS: Positive and Negative Learning’’ of Figure 2).

More specifically, we restrict positive learning to the high-quality group $\mathbf{x}_{t,high}$, as these samples are expected to have a higher probability of being accurately classified by the network. Conversely, negative learning [13, 14] is applied to both groups (*i.e.*, $\mathbf{x}_{t,high}$ and $\mathbf{x}_{t,low}$) to mitigate the risk of providing inaccurate information to the model by generating the complementary labels for the training.

To train our proposed DSS, we employ a cross-entropy loss function. We first have the pseudo-label generated by the teacher model $\hat{y}_{k,high}$, and then we apply temperature scaling to minimize the prediction entropy by sharpening the prediction distribution [2].

$$\text{Sharpen}(\hat{y}_{k,high}, Tp) := \hat{y}_{k,high}^{Tp} / \sum_{k=1}^C \hat{y}_{k,high}^{Tp} \quad (5)$$

where Tp is the temperature hyper-parameter. Most existing methods construct pseudo labels in one-hot way [21,32]. However, in CTDA, we mildly give pseudo labels by applying the sharpening. Specifically, we propose to minimize entropy by constructing sharpened pseudo labels with Tp for selected high-certainty samples. It somewhat avoids overfitting by preserving prediction structures from the teacher model.

The loss function for positive learning is formulated as the cross entropy loss between the student and teacher predictions:

$$\mathcal{L}_{pst}(\bar{y}_{k,high}, \hat{y}_{k,high}) = - \sum_{k=1}^C \hat{y}_{k,high} \log \bar{y}_{k,high} \quad (6)$$

where $\bar{y}_{k,high}$ the prediction from the student model. This loss function will enforce alignment between the student prediction and the more accurate teacher prediction after sharpening. To adequately utilize \hat{y} , the complementary label is also created based on low class confidences. The complementary label represents ‘‘the input does not belong to particular classes.’’ This provides more reliable information regarding the classes that are most likely not to be the true label. Following this, we present the negative learning loss as follows:

$$\mathcal{L}_{neg}(\bar{y}_k, \hat{y}_k) = - \sum_{k=1}^C \check{y}_k \log (1 - \bar{y}_k) \quad (7)$$

where $\check{y}_k = \mathbb{1}(\hat{y}_k < \alpha)$ is the complementary label based on the prediction from the teacher model, \hat{y}_k . Applying the complementary label \check{y} in the negative learning pushes the prediction away from those classes with low confidence. α is the threshold vector whose element value is set to 0.05. Notably, we keep the teacher model fixed and back-propagate the student model only. The reason is that we aim to ensure a stable student-teacher setup. Finally, having the Equations 6 and 7 at hand, we build the total loss function as follows:

$$\mathcal{L}_{total} = \mathcal{L}_{neg} + \mathcal{L}_{pst} \quad (8)$$

Moreover, the teacher model with $\hat{\theta}$ is updated by the moving average of the parameters of the student model, $\bar{\theta}$. Hence, we have

$$\hat{\theta}_{t+1} = \beta \hat{\theta}_t + (1 - \beta) \bar{\theta}_{t+1} \quad (9)$$

where β is the coefficient. We set θ to 0.999 which is the same as that in [27]. The overall training process is described in Algorithm 1.

4. Experiments

Extensive experiments are conducted in this section to demonstrate the effectiveness of our DSS approach. We

Algorithm 1 The proposed DSS for CTDA

Initialization: A source model h_{θ_1} pre-trained on source domain \mathcal{D}_1 , teacher model $h_{\theta_t}^t$ and student model $h_{\theta_t}^s$ initialized from h_{θ_1} .

Input: Unlabeled test data \mathbf{x}_t for target domain \mathcal{D}_t , $t > 1$,

- 1: **for** $t = 2$ to T **do**
- 2: Augment \mathbf{x}_t and get pseudo-labels $\hat{\mathbf{y}}_t$;
- 3: Generate class-wise threshold $\pi_{t,c}$ via Eq. (4);
- 4: Divide \mathbf{x}_t into $\mathbf{x}_{t,low}$ and $\mathbf{x}_{t,high}$ by $\pi_{t,c}$;
- 5: Compute \mathcal{L}_{pst} via Eq. (6);
- 6: Compute \mathcal{L}_{neg} via Eq. (7);
- 7: Update student model via \mathcal{L}_{total} in Eq. (8);
- 8: Update teacher model via Eq. (9);

Output: Prediction $\hat{\mathbf{y}}_t$, updated student and teacher models $h_{\theta_t}^s$ and $h_{\theta_t}^t$, selection threshold $\pi_{t,c}$.

evaluate our method on three 2D image continual test-time adaptation benchmark tasks, CIFAR-10-C [11], CIFAR-100-C [11], and ImageNet-C [11], designed to assess the robustness of machine learning models to corruptions and perturbations in the input data. In addition, we test our proposed method on a 3D point cloud dataset ScanObjectNN-C derived from ScanObjectNN [25].

4.1. Datasets

CIFAR10-C. CIFAR10-C is an extension of the CIFAR-10 dataset [15], which consists of 32×32 color images from 10 classes. CIFAR10-C includes 15 different corruptions, each at five levels of severity. The corruptions are applied to test images of CIFAR-10, resulting in a total of 10,000 images.

CIFAR100-C. CIFAR100-C is an extension of the CIFAR-100 dataset [15], which consists of 32×32 color images in 100 classes. CIFAR100-C includes 15 different corruptions, each at five levels of severity. The corruptions are applied to the test images of CIFAR-100, resulting in a total of 10,000 images.

ImageNet-C. ImageNet-C is an extension of the ImageNet dataset [5], which contains over 14 million images in more than 20,000 categories. ImageNet-C includes 15 different corruptions, with 5 severity levels. The corruptions are applied to the validation images of ImageNet.

ScanObjectNN-C. ScanObjectNN [25] is a point cloud classification dataset that is collected from the real world. It contains 15 classes, with 2309 samples in the train set and 581 in the test set. To build ScanObjectNN-C, we employ the setting proposed by [22], to generate 15 corruptions in the test set of ScanObjectNN for our experiments.

4.2. Implementation

We strictly follow the setting of CTDA that no source data is accessed [27]. All models are evaluated based on the largest corruption severity level of five for all datasets in an online fashion. Model predictions are first generated before adapting to the current test stream.

Following [27], we adopt standard pre-trained WideResNet [31], ResNeXt-29 [30] and ResNet-50 [4] on CIFAR10-C, CIFAR100-C, and ImageNet-C as the source models. The sharpening factor Tp used in temperature scaling is set to 0.6, and augmentation module from [27] is used to generate the augmentation-weighted pseudo-label. For 3D experiments, we use DGCNN [29] pre-trained on the clean set as the backbone.

4.3. Main Results

2D Results. We first examine the adaptation performance on 2D tasks, and the experimental results are summarized in Table 1, 2, and 3. Directly testing the source model on target domains in sequence yields high average errors of 43.5%, 46.4%, and 77.2% on CIFAR10-C, CIFAR100-C, and ImageNet-C respectively. Applying BN Stats Adapt [16] to update the batch normalization statistics from the current test stream, the average error across all target domains is significantly reduced on all datasets. The TENT-based method [26] also helps the model to adapt to the target domain in sequence, but it may suffer severe error accumulation in the long term. As shown in Table 2, the TENT-based method [26] yields a substantially higher error rate of 60.9% in the long run on CIFAR100-C. Similarly, Conjugate PL [9] perform well in early adaptation to multiple initial domains, but soon experiences a gradual increase in error rate over time. Based on more accurate weighted augmentation-averaged predictions, CoTTA [27] has the lowest average error rate in comparison with other adaptation methods.

By leveraging joint positive and negative learning with an adaptive threshold, our proposed method DSS constantly outperforms its baseline CoTTA in all datasets. This indicates that DSS helps the model better adapt to continual target domains with less suffering from error accumulation generated from noisy pseudo labels. We notice that this phenomenon becomes more evident in more difficult datasets, where the model has low certainty on proceeding target domains. Compared with CoTTA, DSS successfully reduces the average error from 32.5% to 30.9% and 66.8% to 64.6% on CIFAR100-C and ImageNet-C respectively, which demonstrates the advantage of our DSS module.

To further investigate the effectiveness of DSS over the baseline, we also evaluate its adaptation performance over ten different sequences on ImageNet-C (See Table 4). There is a 1.4% decrease on average over 10 diverse sequences in comparison with CoTTA, indicating that our method is

Method	Gaussian	shot	impulse	defocus	glass	motion	zoom	snow	frost	fog	brightness	contrast	elastic_trans	pixelate	jpeg	Mean
Source	72.3	65.7	72.9	46.9	54.3	34.8	42.0	25.1	41.3	26.0	9.3	46.7	26.6	58.5	30.3	43.5
BN Adapt [16]	28.1	26.1	36.3	12.8	35.3	14.2	12.1	17.3	17.4	15.3	8.4	12.6	23.8	19.7	27.3	20.4
TENT-cont [26]	24.8	20.6	28.6	14.4	31.1	16.5	14.1	19.1	18.6	18.6	12.2	20.3	25.7	20.8	24.9	20.7
AdaContrast [3]	29.1	22.5	30.0	14.0	32.7	14.1	12.0	16.6	14.9	14.4	8.1	10.0	21.9	17.7	20.0	18.5
Conjugate PL [9]	27.6	25.7	35.9	12.7	34.8	14.1	12.0	17.1	17.3	15.0	8.4	12.3	23.5	19.3	26.7	20.2
CoTTA [27]	24.3	21.3	26.6	11.6	27.6	12.2	10.3	14.8	14.1	12.4	7.5	10.6	18.3	13.4	17.3	16.2
DSS (Ours)	24.1	21.3	25.4	11.7	26.9	12.2	10.5	14.5	14.1	12.5	7.8	10.8	18.0	13.1	17.3	16.0

Table 1. Classification error rate (%) on CIFAR10-C. The best numbers are in bold.

Method	Gaussian	shot	impulse	defocus	glass	motion	zoom	snow	frost	fog	brightness	contrast	elastic_trans	pixelate	jpeg	Mean
Source	73.0	68.0	39.4	29.3	54.1	30.8	28.8	39.5	45.8	50.3	29.5	55.1	37.2	74.7	41.2	46.4
BN Adapt [16]	42.1	40.7	42.7	27.6	41.9	29.7	27.9	34.9	35.0	41.5	26.5	30.3	35.7	32.9	41.2	35.4
TENT-cont [26]	37.2	35.8	41.7	37.9	51.2	48.3	48.5	58.4	63.7	71.1	70.4	82.3	88.0	88.5	90.4	60.9
AdaContrast [3]	42.3	36.8	38.6	27.7	40.1	29.1	27.5	32.9	30.7	38.2	25.9	28.3	33.9	33.3	36.2	33.4
Conjugate PL [9]	39.2	37.1	36.9	26.3	39.4	28.4	26.4	32.9	33.2	38.2	29.3	29.3	34.0	30.3	39.0	33.1
CoTTA [27]	40.1	37.7	39.7	26.9	38.0	27.9	26.4	32.8	31.8	40.3	24.7	26.9	32.5	28.3	33.5	32.5
DSS (Ours)	39.7	36.0	37.2	26.3	35.6	27.5	25.1	31.4	30.0	37.8	24.2	26.0	30.0	26.3	31.1	30.9

Table 2. Classification error rate (%) on CIFAR100-C. The best numbers are in bold.

Method	Gaussian	shot	impulse	defocus	glass	motion	zoom	snow	frost	fog	brightness	contrast	elastic_trans	pixelate	jpeg	Mean
Source	95.3	94.6	95.3	84.9	91.1	86.8	77.2	84.4	80.0	77.3	44.4	95.6	85.2	76.9	66.7	77.2
BN Adapt [16]	87.6	87.4	87.8	87.7	88.0	78.2	64.5	67.6	70.6	54.9	36.4	89.3	58.0	56.4	66.6	66.2
TENT-cont [26]	85.7	80.0	78.3	82.2	79.2	70.9	59.1	65.6	66.4	55.4	40.6	80.3	55.5	53.5	59.0	67.4
Conjugate PL [9]	85.2	79.6	77.1	82.4	79.8	70.8	59.3	65.2	66.1	54.8	39.8	79.5	55.1	52.7	58.7	67.1
CoTTA [27]	87.5	86.0	84.4	85.1	84.4	73.9	61.5	63.6	64.2	51.9	38.6	74.8	51.1	45.1	50.2	66.8
DSS (Ours)	84.6	80.4	78.7	83.9	79.8	74.9	62.9	62.8	62.9	49.7	37.4	71.0	49.5	42.9	48.2	64.6

Table 3. Classification error rate (%) on ImageNet-C. The best numbers are in bold.

more robust to the order of the target domain sequence.

3D Results. While our method is designed to address the CTDA task for 2D image recognition, we also adapt and evaluate our method for the case of 3D point cloud recognition [12, 19, 29]. To this end, we evaluate our proposed methods on ScanObject-C. As shown in Table 5, DSS is again able to reach the lowest error rate of 43.1% using DGCNN [29] as the backbone. The CoTTA method, on the other hand, performs 0.2% worse on average due to the rapid accumulation of errors during the adaptation process.

4.4. Ablation Study

Component analysis. Our method selectively matches the standard student distribution, in contrast to CoTTA which attempts to align prediction probability for all test data. Aligning distributions for all data leads to error accumulation, as demonstrated in Table 6 where CoTTA naively matching distributions without considering the noisy condition of pseudo labels consistently generates larger errors over the continual domains. In comparison, we demonstrate that using our adaptive threshold (DSS w/DT) can yield additional improvements of 0.7% and 0.8% on CIFAR100-C and Imagenet-C respectively. Furthermore, the positive learning (DSS w/DT&PL) leads to 0.8% and 1.2% improvements respectively over DSS w/DT on CIFAR100-C and Imagenet-C. By applying negative learning (DSS w/DT&PL&NL) to utilize more unlabelled low-confidence samples, we achieve further improvements of approximately 0.1% and 0.2%. Apart from these, we observe that

the overall improvement in CIFAR10-C is marginal compared with other cases, and this may be due to the fact that the model is already very confident on those related target domains with average confidence greater than 90%. In that case, the source model is able to produce low-entropy predictions even without online adaptation, and our approach has a limited effect to overcome the bias that lies in the high-confidence predictions.

Threshold. Here, we validate the effectiveness of the dynamic threshold via comparisons with the simple fixed threshold approach across all domains. The results, as presented in Table 7, first show that using a fixed threshold with $\pi = 0.2$ may help to reduce the average error rate by 0.1% on CIFAR100-C. However, using a fixed threshold still causes incorrect predictions to be utilized for self-learning, resulting in a higher error rate of 0.6% compared with the dynamic threshold. Note that the averaged model confidence on out-of-distribution data clearly varies across all target domains (See the green line of Figure 3), and an inappropriate fixed threshold could limit the generalization performance. To illustrate this point, in Table 7, we also show that fixed $\pi = 0.2$ actually increases the error rate to 67.8% from 66.8% on ImageNet-C, whereas our dynamic threshold still consistently decrease the error rate. This is mainly due to the fact that a fixed threshold may filter out excessive samples, leaving inadequate data for model adaptation. Thus, we argue that our adaptive threshold approach strikes a better balance between maximizing target adaptation performance and minimizing error accumulation.

Avg. Error (%)	Source	BN Adapt [16]	Test Aug	TENT [26]	CoTTA [27]	DSS (Ours)
ImageNet-C	82.4	72.1	71.4	66.5	63.0 (± 1.8)	61.6 (± 0.3)

Table 4. Classification error rate (%) on ImageNet-C over 10 sequences. The best numbers are in bold.

Method	uniform	gaussian	background	impulse	upsampling	rbf	rbf-inv	den-dec	dens-inc	shear	rot	cut	distort	oclsion	lidar	Mean
Source	58.7	49.6	55.4	43.5	53.9	37.5	33.2	24.1	19.6	32.9	38.4	24.8	34.8	91.7	92.6	46.0
TENT-cont [26]	53.2	41.3	59.9	46.5	49.7	37.7	36.0	31.0	26.3	37.5	40.3	34.8	39.1	90.5	90.9	47.6
BN [16]	51.3	41.1	53.0	48.7	49.6	34.8	31.7	29.4	22.2	32.9	37.2	32.4	34.8	90.2	90.4	45.3
CoTTA [27]	56.6	47.2	57.0	39.2	49.6	33.2	31.0	21.2	16.2	29.8	34.6	21.3	30.3	90.9	91.4	43.3
DSS (Ours)	56.6	46.8	57.1	39.1	49.6	32.9	31.0	21.2	16.0	29.8	34.1	21.9	30.1	89.8	90.0	43.1

Table 5. Classification error rate (%) on ScanObjectNN-C. DGCNN [29] is adopted as the backbone. The best numbers are in bold.

Avg. Error (%)	CIFAR10-C	CIFAR100-C	Imagenet-C
CoTTA [27]	16.2	32.5	66.8
DSS (w/ DT)	16.1	31.8	66.0
DSS (w/ DT&PL)	16.1	31.0	64.8
DSS (w/ DT&PL&NL)	16.0	30.9	64.6

Table 6. Component analysis. DT, PL, and NL are dynamic thresholding, positive and negative learning. DSS (w/ DT) indicates solely using the high-quality samples and training the model with \mathcal{L}_{cst} , as CoTTA does. DSS (w/ DT&PL) and (w/ DT&PL&NL) are trained with \mathcal{L}_{pst} and $\mathcal{L}_{neg} + \mathcal{L}_{pst}$, respectively.

Avg. Error (%)	CIFAR100-C	Imagenet-C
CoTTA [27]	32.5	66.8
DSS (w/ fixed π)	32.4	67.8
DSS (w/ fixed π &PL)	32.1	77.4
DSS (w/ fixed π &PL&NL)	31.9	76.8
DSS (w/ DT)	31.8	66.0
DSS (w/ DT&PL)	31.0	64.8
DSS (w/ DT&PL&NL)	30.9	64.6

Table 7. Adaptive threshold versus Fixed threshold. DT, PL, and NL are dynamic threshold, positive and negative learning.

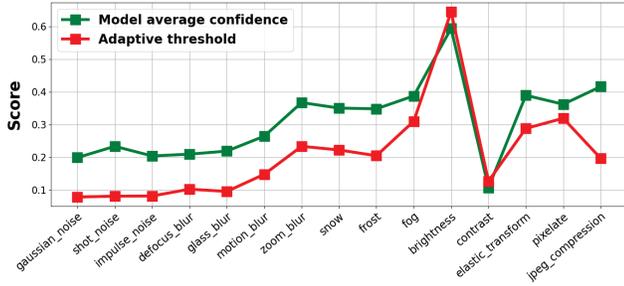


Figure 3. The average prediction confidence (green line) and adaptive threshold (red line) on ImageNet-C across domains.

Moreover, applying positive learning with a fixed threshold on ImageNet-C instead resulted in an increased average error rate of 77.4% from 66.8%, indicating that naive entropy minimization by positive learning would lead to an increase in error accumulation rate. As such, the positive learning component can only be effectively applied with an appropriate threshold that filters out most incorrect pseudo labels. Our adaptive threshold module, however, can always provide a reliable value to filter suspected noisy labels.

Sharpening. Sharpening is employed to decrease the entropy of teacher predictions to assist with model generalization. Here, we demonstrate that in Table 8 when $Tp = 0.6$, it has the lowest average error rate over all domains. Nevertheless, the average error rate increases when the Tp value falls below or exceeds 0.6. The reasoning behind this phenomenon may be attributed to two different factors. Firstly, as Tp approaches 1, the model’s ability to gain knowledge from predictions with higher entropy is limited, resulting in

Avg. Error (%)	CIFAR100-C	Imagenet-C
CoTTA [27]	32.5	66.8
DSS ($Tp = 0.2$)	32.1	74.8
DSS ($Tp = 0.4$)	31.7	66.6
DSS ($Tp = 0.6$)	30.9	64.6
DSS ($Tp = 0.8$)	31.7	64.7
DSS ($Tp = 1$)	32.0	64.9

Table 8. Temperature scaling with Tp ranging from 0.2 to 1.

lower accuracy. On the other hand, decreasing Tp below 0.6 may push the models to output overconfident predictions, potentially resulting in noise overfitting.

5. Conclusion

In this paper, we introduce a novel method, termed Dynamic Sample Selection (DSS), based on joint positive and negative learning with a dynamic threshold for the CTDA task. Traditional methods for CTDA struggle with error accumulation since the adaptation method relies on suspected noisy pseudo-labels as a part of the adaptation process. To address this, in this paper, we consistently monitor the prediction confidence in an online manner and select low- and high-quality samples for different training strategies. To this end, we use positive learning for high-quality samples and negative learning for both low- and high-quality samples. Moreover, we have shown that our proposed method can also be applied to 3D point cloud data as well as 2D images, which showcases its versatility and potential for wide applicability.

References

- [1] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48, 2009. 3
- [2] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *Advances in neural information processing systems*, 32, 2019. 4, 5
- [3] Dian Chen, Dequan Wang, Trevor Darrell, and Sayna Ebrahimi. Contrastive test-time adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2, 7
- [4] Francesco Croce, Maksym Andriushchenko, Vikash Sehwag, Edoardo Debenedetti, Nicolas Flammarion, Mung Chiang, Prateek Mittal, and Matthias Hein. Robustbench: a standardized adversarial robustness benchmark. *arXiv preprint arXiv:2010.09670*, 2020. 6
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 6
- [6] Mario Döbler, Robert A Marsden, and Bin Yang. Robust mean teacher for continual and gradual test-time adaptation. *arXiv preprint arXiv:2211.13081*, 2022. 1, 2, 3
- [7] Yulu Gan, Xianzheng Ma, Yihang Lou, Yan Bai, Renrui Zhang, Nian Shi, and Lin Luo. Decorate the newcomers: Visual domain prompt for continual test time adaptation. *arXiv preprint arXiv:2212.04145*, 2022. 3
- [8] Yossi Gandelsman, Yu Sun, Xinlei Chen, and Alexei A Efros. Test-time training with masked autoencoders. *arXiv preprint arXiv:2209.07522*, 2022. 2
- [9] Sachin Goyal, Mingjie Sun, Aditi Raghunathan, and J Zico Kolter. Test time adaptation via conjugate pseudo-labels. *Advances in Neural Information Processing Systems*, 35:6204–6218, 2022. 2, 6, 7
- [10] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2019. 2
- [11] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019. 6
- [12] Jie Hong, Shi Qiu, Weihao Li, Saeed Anwar, Mehrtash Harandi, Nick Barnes, and Lars Petersson. Pointcam: Cut-and-mix for open-set point cloud learning. *arXiv preprint arXiv:2212.02011*, 2023. 7
- [13] Youngdong Kim, Junho Yim, Juseung Yun, and Junmo Kim. Nlnl: Negative learning for noisy labels. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 101–110, 2019. 5
- [14] Y. Kim, J. Yun, H. Shon, and J. Kim. Joint negative and positive learning for noisy labels. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9437–9446, Los Alamitos, CA, USA, jun 2021. IEEE Computer Society. 5
- [15] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6
- [16] Yanghao Li, Naiyan Wang, Jianping Shi, Jiaying Liu, and Xiaodi Hou. Revisiting batch normalization for practical domain adaptation. *arXiv preprint arXiv:1603.04779*, 2016. 6, 7, 8
- [17] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *International Conference on Machine Learning*, pages 6028–6039, 2020. 2
- [18] Shuaicheng Niu, Jiayang Wu, Yifan Zhang, Yafo Chen, Shijian Zheng, Peilin Zhao, and Mingkui Tan. Efficient test-time model adaptation without forgetting. *arXiv preprint arXiv:2204.02610*, 2022. 2
- [19] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. 7
- [20] Mamshad Nayeem Rizve, Kevin Duarte, Yogesh S Rawat, and Mubarak Shah. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. *arXiv preprint arXiv:2101.06329*, 2021. 3
- [21] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020. 3, 5
- [22] Jiachen Sun, Qingzhao Zhang, Bhavya Kailkhura, Zhiding Yu, Chaowei Xiao, and Z Morley Mao. Benchmarking robustness of 3d point cloud recognition against common corruptions. *arXiv preprint arXiv:2201.12296*, 2022. 2, 6
- [23] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *International conference on machine learning*, pages 9229–9248, 2020. 2
- [24] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017. 2, 3
- [25] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Duc Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *International Conference on Computer Vision (ICCV)*, 2019. 6
- [26] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *International Conference on Learning Representations*, 2021. 2, 6, 7, 8
- [27] Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual test-time domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7201–7211, 2022. 1, 2, 3, 5, 6, 7, 8
- [28] Yidong Wang, Hao Chen, Qiang Heng, Wenxin Hou, Marios Savvides, Takahiro Shinozaki, Bhiksha Raj, Zhen Wu, and

- Jindong Wang. Freematch: Self-adaptive thresholding for semi-supervised learning. *arXiv preprint arXiv:2205.07246*, 2022. 3
- [29] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)*, 38(5):1–12, 2019. 6, 7, 8
- [30] Dong Yin, Raphael Gontijo Lopes, Jon Shlens, Ekin Dogus Cubuk, and Justin Gilmer. A fourier perspective on model robustness in computer vision. *Advances in Neural Information Processing Systems*, 32, 2019. 6
- [31] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In Edwin R. Hancock Richard C. Wilson and William A. P. Smith, editors, *Proceedings of the British Machine Vision Conference (BMVC)*, pages 87.1–87.12. BMVA Press, September 2016. 6
- [32] Bowen Zhang, Yidong Wang, Wenxin Hou, Hao Wu, Jindong Wang, Manabu Okumura, and Takahiro Shinozaki. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. *Advances in Neural Information Processing Systems*, 34:18408–18419, 2021. 3, 5