

Towards Open-Set Test-Time Adaptation Utilizing the Wisdom of Crowds in Entropy Minimization

Jungsoo Lee^{1,2*} Debasmit Das¹ Jaegul Choo² Sungha Choi^{1†}

¹Qualcomm AI Research[‡] ²KAIST

¹{jungsool, debadas, sunghac}@qti.qualcomm.com ²{bebeto, jchoo}@kaist.ac.kr

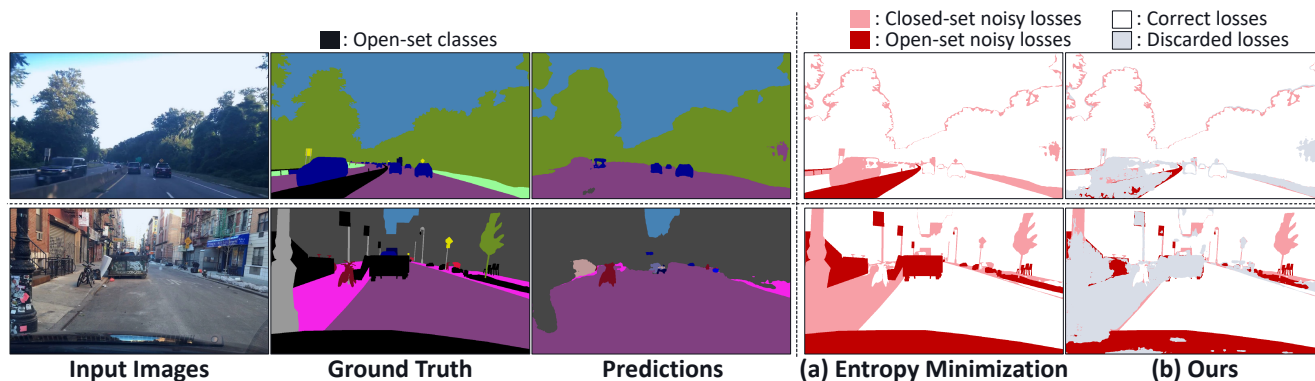


Figure 1: Discarding noisy losses in test-time adaptation (TTA). Black pixels in the ground truth indicate the open-set classes in the test set (*i.e.*, BDD-100K [60]), which were not included in the train set (*i.e.*, Cityscapes [10]), such as the guardrails in the first row or the garbage truck in the second row. (a) Since TTA models generally use their predictions for the target outputs, they are prone to utilizing noisy losses from 1) wrong predictions (pink pixels ■) and 2) open-set classes (red pixels ■). Performing TTA for a long term in such an environment degrades the performance of TTA models significantly. (b) Our method effectively filters out such noisy losses, preventing performance degradation and alarming unexpected obstacles, which is crucial for safety-critical applications such as autonomous driving (see Supplementary).

Abstract

Test-time adaptation (TTA) methods, which generally rely on the model’s predictions (*e.g.*, entropy minimization) to adapt the source pretrained model to the unlabeled target domain, suffer from noisy signals originating from 1) incorrect or 2) open-set predictions. Long-term stable adaptation is hampered by such noisy signals, so training models without such error accumulation is crucial for practical TTA. To address these issues, including open-set TTA, we propose a simple yet effective sample selection method inspired by the following crucial empirical finding. While entropy minimization compels the model to increase the probability of its predicted label (*i.e.*, confidence values), we found that noisy samples rather show decreased confidence values. To be more specific, entropy minimization attempts to raise the confidence values of an individual sample’s prediction, but individual confidence values may rise or fall due to the influence of signals from numerous other pre-

dictions (*i.e.*, wisdom of crowds). Due to this fact, noisy signals misaligned with such ‘wisdom of crowds’, generally found in the correct signals, fail to raise the individual confidence values of wrong samples, despite attempts to increase them. Based on such findings, we filter out the samples whose confidence values are lower in the adapted model than in the original model, as they are likely to be noisy. Our method is widely applicable to existing TTA methods and improves their long-term adaptation performance in both image classification (*e.g.*, 49.4% reduced error rates with TENT) and semantic segmentation (*e.g.*, 11.7% gain in mIoU with TENT).

1. Introduction

Despite the recent advancements of deep learning, models still show a significant performance degradation when confronted with large domain shifts (*e.g.*, changes of cities with different landscapes during autonomous driving) [8, 33, 39, 27, 11]. Among various studies, test-time adaptation (TTA) is at the center of attention due to its practicality in not requiring 1) the source data during the adaptation stage and 2) ground truth labels of the target domain [57].

*Work done during an internship at Qualcomm AI Research.

† Corresponding author. ‡ Qualcomm AI Research is an initiative of Qualcomm Technologies, Inc.




Figure 2: Utilizing confidence difference for selecting correct samples. Pseudo-labeling samples (*i.e.*, selecting correct samples) by using a fixed threshold does not guarantee a reasonable level of pseudo-labeling performance, which is demonstrated by the significantly low precision values. On the other hand, we maintain a reasonable level of both precision and recall by using the confidence difference between θ_o and θ_a , improving the test-time adaptation performance overall.

TTA models widely utilize a self-training strategy (*e.g.*, entropy minimization), which uses the model’s prediction as the target of the loss function [57, 9, 58, 46, 14, 13, 62, 47, 65, 24]. Since TTA models rely on their own predictions during the adaptation, they are inevitably prone to utilizing noisy signals. In this paper, noisy signals indicate supervisions that originated from 1) incorrect or 2) open-set predictions. Fig. 1 shows that performing adaptation with such noisy signals significantly degrades the TTA performance. Specifically, the pink pixels indicate the mispredicted pixels (*e.g.*, predicting sidewalks as roads in the second row), and the red ones are the predictions on open-set classes that were not included in the train set (*e.g.*, predicting guardrails and the garbage truck as roads in the first and second rows, respectively). Such an example clearly demonstrates that TTA in real-world applications needs to address such open-set classes since mispredicting guardrails as roads may cause serious accidents during autonomous driving. However, as shown in Fig. 3, previous studies focused on TTA with covariate shifts (*i.e.*, domain shifts) only and did not address TTA that also includes semantic shifts (*i.e.*, including open-set classes). Regarding its significance and practicality, adaptation with unknown classes included (*i.e.*, open-set TTA) should be also addressed.

Fig. 2 shows our empirical analysis that discloses an important finding to address such an issue. While entropy minimization enforces the model to increase the probability value of its predicted label (*i.e.*, confidence values), we found that it often fails to increase them on the wrong sam-




Figure 3: Description of open-set TTA. While previous studies assume covariate shifts (*i.e.*, Cityscapes to BDD-100K), they fail to address the semantic shifts (*i.e.*, guardrails only shown in BDD-100K). This paper addresses both closed-set and open-set test-time adaptation.

ples. While previous studies [58, 46] resorted to finding an adequate confidence value or loss value to prevent error accumulation, the process of determining it is cumbersome, and utilizing such a *static* threshold shows limited performance. We train TENT [57] with different thresholds for the analysis: (a) without thresholding, (b) selecting samples with confidence value higher or equal to 0.9¹, (c) selecting samples with loss values smaller than the entropy threshold proposed in EATA [46], and (d) selecting samples that achieve higher confidence value with the adaptation model θ_a compared to that with the original model θ_o . As shown, using the confidence difference between θ_o and θ_a for selecting correct samples outperforms utilizing the static thresholds. While b) and c) show significantly high recall values (*i.e.*, selecting actual correct samples well), it rather indicates that they simply select most of the samples and fail to filter out noisy samples considering the substantially low precision values (*i.e.*, low ratio of correct samples among the selected ones).

The intuition behind using the confidence difference is as follows. Although entropy minimization enforces the model to increase the confidence value on the predicted label of an individual sample, the individual confidence value may rise or fall, influenced by the signals that originated from numerous other predictions (*i.e.*, wisdom of crowds). To be more specific, the noisy signals that do not align with such ‘wisdom of crowds’, commonly found in the correct signals, fail to raise the individual confidence scores of wrong samples, even with the supervision from entropy minimization to increase them. By using such an observation, we select samples that achieve higher confidence value using θ_a compared to that using θ_o . Since we reflect the knowledge state of the model on each individual sample, our selection is implicitly a dynamic thresholding strategy, which outperforms the previously-used static strategies. Our *simple yet effective* sample selection method is widely applicable to existing TTA methods and improves their performances on both image classification and semantic segmentation.

¹We used the best confidence value p after grid search of $p \in \{0.5, 0.8, 0.9, 0.95, 0.99\}$.




Figure 4: As adaptation proceeds, the number of samples with decreased confidence values increases (purple graph). Additionally, among those samples, the ratio of wrongly predicted samples also increases (green graph). t_i indicates the i^{th} round during the long-term adaptation.

Our contributions are summarized as follows:

- We propose a novel sample selection method that filters out noisy samples using the confidence difference between θ_a and θ_o based on the finding that noisy samples, both closed-set wrong samples, and open-set samples, generally show decreased confidence values on the originally predicted label.
- This is the first paper to address open-set test-time adaptation, adapting to a target domain including test samples of unknown classes, which has not been explored in existing TTA studies despite its importance and practicality.
- Our proposed method can be applied to various test-time adaptation methods and improves their performances on both image classification using CIFAR-10/100-C and TinyImageNet-C (e.g., 49.38% reduced error rates with TENT in open-set TTA), and semantic segmentation (e.g., 11.69% gain in mIoU with TENT) using real-world datasets including BDD-100K and Mapillary.

2. Wisdom of Crowds in Entropy Minimization

2.1. Problem Setup

During the test-time adaptation, models adapt to a target domain with N number of test samples in the test set D_T , $\{x_i\}_{i=1}^N \in D_T$, without target labels provided. Given a pretrained model θ_o , we update θ_o to adapt to a novel target domain, where the adapted model is then defined as θ_a . For a test sample x , we define $\tilde{y} = f(x; \theta_o) \in \mathbb{R}^C$ and $\hat{y} = f(x; \theta_a) \in \mathbb{R}^C$ as the softmax outputs of the original model θ_o and the adapted model θ_a , respectively, where C denotes the number of classes. With the predicted class $c_o = \operatorname{argmax}_c f(x; \theta_o)$ of the original model, we define the probability value on the predicted label as confidence value \tilde{y}^{c_o} . Similarly, the confidence value of the adapted model θ_a on the label c_o , predicted by the original model, is defined as \hat{y}^{c_o} . The main objective of test-time adaptation is to correctly predict $c_a = \operatorname{argmax}_c f(x; \theta_a)$ using the adapted model, especially under large data distribution shifts.




Figure 5: Utilizing the confidence difference distinguishes between the correct samples (blue) and the wrong samples (red) better (AUROC of 89.42) than using the confidence values (AUROC of 58.29). We used the same model (i.e., TENT [57] adapted for 50 rounds) for the visualization.

2.2. Motivation

Decreased confidence values While entropy minimization enforces the model to increase the confidence value of its originally predicted label, we empirically found that wrong samples mostly show decreased values (i.e., $\hat{y}^{c_o} < \tilde{y}^{c_o}$). For the experiment, we perform test-time adaptation using TENT [57] for 50 rounds using CIFAR-10-C to simulate a long-term adaptation. One round includes continuously changing 15 corruption types, so we repeat it 50 times without resetting the model. With t_i indicating the i^{th} round, Fig. 4 (purple graph) shows that the number of samples achieving $\hat{y}^{c_o} < \tilde{y}^{c_o}$, showing decreased confidence values, among N number of test samples increases as adaptation proceeds even with the entropy minimization that enforces the model to increase its confidence value on the originally predicted label. In fact, the green graph in Fig. 4 shows that the ratio of wrong samples among the samples with decreased confidence values also increases as adaptation proceeds. The main reason for such an observation is due to the ‘wisdom of crowds’, the signals learned from numerous other samples influencing the confidence level of individual samples. Specifically, although the individual signal from each sample compels the model to increase the confidence value of its own predicted label, this effect may be canceled out if other dominant signals show different patterns.

Wisdom of crowds from correct samples We empirically found that models generally learn the wisdom of crowds from the correct samples. Fig. 5 demonstrates such a point with the histogram of 1) confidence values and 2) confidence difference, $\hat{y}^{c_o} - \tilde{y}^{c_o}$, using TENT [57] adapted for 50 rounds. We observe that a substantial number of the samples achieving $\hat{y}^{c_o} - \tilde{y}^{c_o} \geq 0$ are correct samples (blue). To be more specific, utilizing the confidence difference for distinguishing correct samples from wrong samples (red) achieves an AUROC of 89.42, which outperforms utilizing the confidence value of the adaptation model, achieving an AUROC of 58.29.

Such an observation discloses two findings. First, since samples achieving $\hat{y}^{c_o} \geq \tilde{y}^{c_o}$ are mostly correct ones, the




Figure 6: Overall procedure of our sample selection. We forward the mini-batch of n test images, $\{x_i\}_{i=1}^n$, to the original model θ_o and the adaptation model θ_a . Then, we compare the probability values \hat{y}^{co} and \tilde{y}^{co} and select the samples achieving $\hat{y}^{co} \geq \tilde{y}^{co}$. Finally, we apply the entropy minimization only to the selected samples.

dominant signals necessary for increasing the confidence values (*i.e.*, wisdom of crowds) are originated from the correct samples. Second, $\hat{y}^{co} - \tilde{y}^{co}$ is an adequate metric to distinguish between correct and wrong samples.

Misaligned wrong signals We further empirically analyze why signals from wrong samples fail to increase the confidence values of the original model. The main reason is that signals originated from wrong samples misalign with the ‘wisdom of crowds’ obtained from the correct samples. For the analysis, we compute the cosine similarity of gradients between two samples with the same predicted label in Fig. 7. For a given predicted label i (column), we compute $s^{j,i}$, the cosine similarity of gradients obtained between samples of ground truth label j (row) and those of predicted label i as,

$$s^{j,i} = \frac{1}{M_1 M_2} \sum_{k=1}^{M_1} \sum_{l=1}^{M_2} \frac{g_k^{j,i} \cdot g_l^{i,i}}{\|g_k^{j,i}\| \|g_l^{i,i}\|}, l \neq k \text{ if } j = i, \quad (1)$$

where $g_k^{j,i}$ indicates the gradient vector of k^{th} sample among M_1 number of samples with the ground truth label j and the predicted label i , $g_l^{i,i}$ indicates the gradient vector of l^{th} sample among M_2 number of samples with the ground truth label i and the predicted label i (*i.e.*, correct samples), and $i, j \in C$. In other words, given a certain predicted label, we compare the gradients of the correct samples and those of the samples with the same predicted label either correct

or wrong. Thus, the diagonal elements are the results obtained by comparing the gradients between correct samples and the off-diagonal elements are obtained by comparing the gradients between correct samples and the wrong samples with the same predicted label. We add the description of how the cosine similarity of each pair is computed on the right side of Fig. 7.

Given a certain column in Fig. 7, all entropy minimization losses enforce the model to increase the probability value of the same predicted label. However, we found that the signals (*i.e.*, gradients) may differ depending on the actual ground truth labels. Specifically, the correct samples show high cosine similarity of gradients (diagonal elements, *e.g.*, $s_{2,2}$) compared to the ones with wrong samples (off-diagonal elements, *e.g.*, $s_{0,2}$). Since Fig. 5 shows that the correct signals dominate the wisdom of crowds required for increasing the confidence value of the originally predicted label, signals that are different from these dominant signals can be suppressed and do not raise confidence values.

We want to clarify that the wisdom of crowds does not guarantee a model to utilize the correct signals only. Even with the wisdom of crowds, the model supervises itself with wrong predictions if the noisy losses are not filtered out. Such self-training with wrong knowledge significantly deteriorates the TTA performance of models, especially during the long-term adaptation [25]. In fact, such an issue has been widely studied in fields beyond TTA, known as the *confirmation bias* [63, 2, 56, 35, 34]. To address such an issue in TTA, we propose a sample selection method to filter out noisy samples by using the wisdom of crowds.

2.3. Proposed Method

As shown in Fig. 6, we propose a *simple yet effective* sample selection method using the confidence difference between \tilde{y}^{co} and \hat{y}^{co} . Our sample selection criterion is formulated as

$$\Phi(\hat{y}^{co}, \tilde{y}^{co}) = \mathbb{1}(\hat{y}^{co} \geq \tilde{y}^{co}), \quad (2)$$

where $\Phi(\cdot)$ is our sample selection criterion and $\mathbb{1}(\cdot)$ is the indicator function. Our total objective function using entropy minimization is formulated as

$$\mathcal{L}^{\text{main}}(x; \theta_a) = \Phi(\hat{y}^{co}, \tilde{y}^{co}) \cdot H(\hat{y}_i) - \lambda_{\text{max}} H(\bar{y}). \quad (3)$$

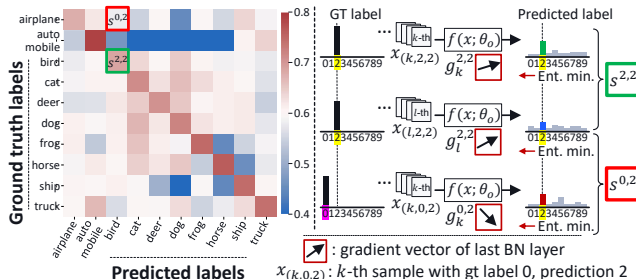


Figure 7: Cosine similarity of gradients between samples with the same predicted label. We observe that wrong signals (*i.e.*, off-diagonal elements) misalign with the correct signals (*i.e.*, diagonal elements) that dominate the wisdom of crowds.

Method	CIFAR-10-C		CIFAR-100-C		TinyImageNet-C		Average	
	Closed	Open	Closed	Open	Closed	Open	Closed	Open
Source [61]	18.27	18.27	46.75	46.75	76.71	76.71	47.24	47.24
BN Adapt [43]	14.49	15.73	39.26	42.67	61.90	63.00	38.55	40.47
GCE [64]	43.76	87.94	44.45	88.69	97.25	99.00	61.82	91.88
Conjugate [14]	49.57	92.25	98.97	98.79	99.38	99.46	82.64	96.83
ENT	87.06	89.26	56.35	98.76	99.43	99.50	80.95	95.84
+ Ours	17.33 (-69.73)	23.98 (-65.28)	37.69 (-18.66)	40.48 (-58.28)	58.93 (-40.50)	64.01 (-35.49)	37.98 (-42.97)	42.82 (-53.02)
TENT [57]	45.84	85.22	42.34	85.22	98.10	99.16	62.09	89.87
+ Ours	14.10 (-31.74)	15.77 (-69.45)	38.62 (-3.72)	42.57 (-42.65)	60.87 (-37.23)	63.13 (-36.03)	37.86 (-24.23)	40.49 (-49.38)
EATA [46]	29.78	82.05	49.31	98.75	59.82	63.47	46.30	81.42
+ Ours	14.07 (-15.71)	15.65 (-66.40)	38.44 (-10.87)	42.47 (-56.28)	59.80 (-0.02)	62.08 (-1.39)	37.44 (-8.86)	40.07 (-41.35)
SWR [9]	10.21	90.55	35.78	73.05	62.39	76.13	36.13	79.91
+ Ours	10.12 (-0.09)	72.58 (-17.97)	35.64 (-0.14)	45.68 (-27.37)	55.15 (-7.24)	61.91 (-14.22)	33.64 (-2.49)	60.06 (-19.85)

Table 1: Error rates of image classification after 50 rounds of adaptation (*i.e.*, long-term test-time adaptation). We note the performance gain by reduced error rates.

Method	CIFAR-10-C		CIFAR-100-C		TinyImageNet-C		Average	
	Closed	Open	Closed	Open	Closed	Open	Closed	Open
Source [61]	18.27	18.27	46.75	46.75	76.71	76.71	47.24	47.24
BN Adapt [43]	14.49	15.73	39.26	42.67	61.90	63.00	38.55	40.47
GCE [64]	12.81	25.70	35.83	45.78	62.84	71.41	37.16	47.63
Conjugate [14]	12.84	24.96	36.67	81.19	82.83	92.66	44.11	66.27
ENT	16.30	47.54	38.74	58.16	79.69	91.74	44.91	65.81
+ Ours	13.41 (-2.89)	16.93 (-30.61)	37.55 (-1.19)	42.60 (-15.56)	63.89 (-15.80)	69.01 (-22.73)	38.28 (-6.63)	42.85 (-22.96)
TENT [57]	12.56	27.80	36.04	45.26	68.53	80.93	39.04	51.33
+ Ours	12.39 (-0.17)	14.94 (-12.86)	36.18 (+0.14)	39.62 (-5.64)	59.90 (-8.63)	63.31 (-17.62)	36.16 (-2.88)	39.29 (-12.04)
EATA [46]	12.39	25.52	36.39	54.22	59.02	61.72	35.93	47.15
+ Ours	12.35 (-0.04)	14.92 (-10.60)	36.25 (-0.14)	39.58 (-14.64)	59.30 (+0.28)	62.11 (+0.39)	35.97 (+0.04)	38.87 (-8.28)
SWR [9]	10.76	29.32	34.21	44.79	60.34	65.18	35.10	46.43
+ Ours	10.74 (-0.02)	27.52 (-1.80)	34.23 (+0.02)	41.52 (-3.27)	58.50 (-1.84)	62.94 (-2.24)	34.49 (-0.61)	44.00 (-2.43)

Table 2: Error rates of image classification after 1 round of adaptation (*i.e.*, short-term test-time adaptation). We note the performance gain by reduced error rates.

$H(p) = \sum_{k=1}^C p^k \log p^k$, $\bar{y} = \frac{1}{N} \sum_{k=1}^C \hat{y}_i$, and λ_{max} is the scalar value for balancing the two loss values. Note that $H(\bar{y})$ has been widely used in previous studies [9, 39, 29, 38, 3, 5] to prevent the model from making imbalanced predictions towards a certain class.

Recent studies require the pre-deployment stage that obtains the necessary information needed for each method by using the samples from the source data before the adaptation phase [9, 46, 39]. However, we want to emphasize that our method does not require such a pre-deployment stage as well as those samples from the source distribution. Due to such an advantage, our method can be easily applied to existing TTA methods without additional preparations. Through extensive experiments, we demonstrate the wide applicability of our method to existing TTA methods.

3. Experiments

3.1. Experimental Setup

Datasets For the image classification task, we use the widely used corruption benchmark datasets: CIFAR-10/100-C and TinyImageNet-C. We apply 15 different types of corruptions (e.g., gaussian noise) to CIFAR-10/100 [30] and TinyImageNet [31]. Pretrained models are trained on

the clean train set and adapted to the corrupted test set. For the open-set setting, we use SVHN [44] for CIFAR-10/100-C, and ImageNet-O [20] for TinyImageNet-C, where we apply the same corruption type as the original test sets. We term the datasets as SVHN-C and ImageNet-O-C, respectively. We apply the identical corruption type in order to construct open-set samples that are drawn from the same domain shift but with unknown classes. For the semantic segmentation task under continually changing domains, we use a model pretrained on GTAV [50], and evaluate it with Cityscapes [10], BDD-100K [60], and Mapillary [45]. For semantic segmentation with a fixed target domain with multiple rounds, we use the Cityscapes for the source distribution and BDD-100K [60], GTAV [50], Mapillary [45], and SYNTHIA [51] for the target distributions. Note that the semantic segmentation task inherently includes open-set classes in the test set (e.g., traffic cones in BDD100K not shown during training with Cityscapes).

Evaluation settings Following the recent TTA studies, we evaluate TTA models under continuously changing domains without resetting the model after each domain [58, 39, 46]. For the closed-set and open-set continual long-term TTA in the image classification, we perform adaptation for 50 rounds to simulate a long-term TTA with continu-

Time	$t \longrightarrow$			Average
Method	Cityscapes	BDD-100K	Mapillary	
Source [6]	34.74	16.15	36.97	29.29
BN Adapt [43]	40.77	25.21	39.10	35.03
TTN [39]	46.28	28.07	45.46	39.94
TENT [57]	46.73	29.59	35.69	37.34
+ Ours	46.76 (+0.03)	30.55 (+0.96)	43.42 (+7.73)	40.24 (+2.90)
SWR [9]	46.17	10.70	1.28	19.38
+ Ours	46.65 (+0.48)	32.28 (+21.58)	45.09 (+43.81)	41.34 (+21.96)

(a) Average mIoU after the adaptation of each domain.


(b) mIoU changes during adaptation².

Table 3: Semantic segmentation performance (mIoU) on continuously-changing target domains with 1 round of adaptation. We evaluate with DeepLabV3Plus-ResNet-50 [6] pretrained on GTAV dataset.

Method	BDD-100K		Mapillary		GTAV		SYNTHIA		Average	
	Round 1	Round 10	Round 1	Round 10	Round 1	Round 10	Round 1	Round 10	Round 1	Round 10
Source [6]	43.50	43.50	54.37	54.37	44.55	44.55	22.78	22.78	41.30	41.30
BN Adapt [43]	43.60	43.60	47.66	47.66	43.22	43.22	25.72	25.72	40.05	40.05
TTN [39]	48.43	48.43	57.28	57.28	46.71	46.71	26.41	26.41	44.71	44.71
TENT [57]	48.90	47.57	57.94	53.36	48.14	17.91	26.88	13.36	45.47	33.05
+ Ours	48.90	48.88 (+1.31)	57.94	56.49 (+3.13)	48.28 (+0.14)	47.98 (+30.07)	26.90 (+0.02)	25.62 (+12.26)	45.51 (+0.04)	44.74 (+11.69)
SWR [9]	49.39	49.68	59.33	59.70	47.82	48.13	28.40	1.18	46.24	39.67
+ Ours	49.88 (+0.49)	50.57 (+0.89)	58.79 (-0.54)	58.89 (-0.81)	49.17 (+1.35)	49.27 (+1.14)	27.75 (-0.65)	27.82 (+26.64)	46.40 (+0.16)	46.64 (+6.97)

Table 4: Semantic segmentation performance (mIoU) on a fixed target domain with 10 rounds of adaptation. We use DeepLabV3Plus-ResNet-50 [6] pretrained on Cityscapes dataset.

ously changing domains. We report both TTA performances after 1 round (*i.e.*, short-term TTA) and 50 rounds (*i.e.*, long-term TTA). Note that we evaluate predictions made during online model adaptation, not after visiting the entire test set, strictly following the established TTA settings. For the open-set TTA, we construct the mini-batch that includes an equal number of closed-set samples (e.g., CIFAR-10-C, shot noise) and open-set samples (e.g., SVHN-C, shot noise). Although included in the mini-batch, we exclude open-set samples from the evaluation and only evaluate models with closed-set samples. To the best of our knowledge, our work is the first paper to conduct experiments with the open-set TTA. We report the error rates and mean intersection of union (mIoU) for image classification and semantic segmentation, respectively.

Baselines We mainly compare our method with previous methods addressing noisy labels [64] or improving pseudo-labeling performances in TTA [14, 46]. Note that ENT denotes updating all parameters while TENT [57] only updates affine parameters of the batch normalization layers, both utilizing the entropy minimization loss function. Gray-shaded digits indicate the performance gain by applying our method to each baseline model, and bold digits indicate the better performance between the two methods.

Implementation details For the image classification, we use the learning rate of $1e-3$ and $1e-4$ for models updating only affine parameters (TENT [57], EATA [46], GCE [64], Conjugate [14]) and all parameters (ENT, SWR [9]), respectively. We use the batch size of 200 and Adam optimizer [28] for all experiments. For experiments conducting small batch sizes in Table 8, we use the learning rate of $1e-4$

and update models after 200 steps, following TTN [39]. For the semantic segmentation, we use the learning rate of $1e-6$ and batch size of 2 following TTN. Regarding using TTN in semantic segmentation, we update the test batch statistics in an online manner to further improve the segmentation performance for all experiments. Further details on our experimental setup are included in our supplementary.

3.2. Results

Image classification As shown in Table 1, existing TTA models show a large performance degradation during the long-term adaptation. This is mainly due to the confirmation bias, caused by the unsupervised losses that inevitably include noisy losses. We significantly improve the long-term performance of the existing four different TTA models in both closed-set and open-set TTA. For example, we improve the error rate of TENT [57] by an average of 24.23% and 49.38% in the closed-set and open-set settings, respectively. Note that we do not use prior knowledge of whether the target distribution includes open-set samples or not. Additionally, Table 2 shows that our method also generally improves the short-term TTA performances.

While previous studies focused on improving the performance of closed-set TTA until now, our results show that they suffer from a large performance drop when adapted with open-set classes included. We believe that this is a practical setting since we can not guarantee that samples from the target distributions are always drawn from the

²Note that the performance variation of the source model in Cityscapes is due to the order of data samples (*e.g.*, challenging ones in the later stage), not due to the adaptation.

Method	CIFAR-10 / SVHN-C		CIFAR-100 / SVHN-C	
	AUROC \uparrow	FPR@TPR95 \downarrow	AUROC \uparrow	FPR@TPR95 \downarrow
MSP [19]	51.87	92.39	60.69	87.96
Max Logit [18]	54.68	90.31	64.88	85.45
Energy [40]	54.68	90.30	64.87	85.46
Ours	88.24	40.34	83.76	64.86

(a) Negative samples including closed-set wrong samples.

Method	CIFAR-10 / SVHN-C		CIFAR-100 / SVHN-C	
	AUROC \uparrow	FPR@TPR95 \downarrow	AUROC \uparrow	FPR@TPR95 \downarrow
MSP [19]	50.83	93.64	56.14	90.34
Max Logit [18]	56.25	90.65	62.76	87.35
Energy [40]	56.26	90.63	62.79	87.27
Ours	83.50	54.46	82.17	73.16

(b) Negative samples excluding closed-set wrong samples.

Table 5: Utilizing the confidence difference for thresholding in open-set test time adaptation. We use TENT [57] adapted to each target domain including open-set classes (SVHN-C) for 50 rounds.

Method	CIFAR-10-C	CIFAR-100-C	CIFAR-10/100-C	
	Error Rate (%)	Error Rate (%)	Memory (MB)	Time (ms)
ENT	88.16	77.56	1147	22.98
SWR [9]	50.38	54.42	1155	47.97
TENT [57]	65.53	63.78	556	18.38
EATA [46]	55.92	74.03	559	37.04
TENT [57] + Ours	14.94	40.60	565	26.62

Table 6: Comparisons on error rates (%), memory (MB), and time (ms). For the time, we report the average time after 5000 trials on NVIDIA RTX A5000.

classes learned during the training stage. Such results indicate that improving the TTA performance with open-set classes is yet to be explored in the future.

Semantic segmentation Table 3 shows the semantic segmentation performance with continuously changing domains. We evaluated a model pretrained on GTAV [50] with real-domain datasets (Cityscapes [10], BDD-100K [60], and Mapillary [45]) in order to simulate the situation where real-world target datasets are not available with only synthetic datasets provided. We observe that the performance gain by applying our method increases as the adaptation proceeds. For example, SWR [9] (Table 3b - red) suffers from a large performance drop with the last target domain, Mapillary (1.28 mIoU), while ours (Table 3b - blue) shows a stable level of performance (45.09 mIoU). Regarding Table 3b, we evaluate models after certain steps and show the average mIoU up to then. While the model without adaptation (*i.e.*, source) does not suffer from the error accumulation, it fails to bring performance gain. On the other hand, our method not only brings performance gain but also circumvents error accumulation by filtering the noisy losses.

Table 4 also reports the semantic segmentation performance with a fixed target domain over multiple rounds of adaptation. We observe that applying our method improves the performance of TENT [57] and SWR [9] by an average of 11.69 mIoU and 6.97 mIoU, respectively, after 10 rounds. As aforementioned, performing test-time adaptation in semantic segmentation needs to address not only the wrong predictions but also the inherently included open-set classes in the target distribution. Our method again improves TTA performance by effectively discarding such noisy pixels. We believe such a filtering mechanism is especially important in safety-critical applications in two aspects. First, it prevents the performance drop caused by learning with noisy losses. Second, when confronted with unknown ob-

jects, we could alarm a device immediately, which could be the starting point for it to take a different action (e.g., autonomous vehicles swerving directions to avoid running into wild animals unexpectedly shown on roads) [23].

4. Further Analysis

4.1. Utilizing Confidence Difference as Thresholds

We show that the confidence difference is an adequate metric to differentiate between correct samples and noisy samples, given that a pretrained model is adapting to a novel domain. For the evaluation, we train TENT [57] and compare utilizing confidence difference as the thresholding metric with existing prediction-based out-of-distribution (OoD) methods [19, 18, 40]. By setting the correct samples as the positive samples, we analyze two different negative samples: negative samples 1) including closed-set wrong samples and 2) excluding closed-set wrong samples. The former case shows how well a given metric differentiates between correct samples and noisy samples, including both closed-set and open-set samples. The latter case evaluates how well a given metric distinguishes between the correct samples and open-set samples only. Table 5 shows that using confidence difference outperforms the existing OoD metrics in both cases. In addition to the superior performance, another advantage of using the confidence difference is that we can filter the noisy samples immediately, while the existing OoD metrics need the entire test samples in order to choose the threshold with the best AUROC score. Such a result indicates that confidence difference can also be widely used to distinguish out-of-distribution samples in future studies with adapted models.

4.2. Comparisons on Resource Costs

Along with the TTA performances, Table 6 compares the memory usage and the time consumption of the baseline models and our method applied to TENT [57]. For the TTA performance, we average the long-term adaptation performance of closed-set and open-set TTA for each dataset. For memory usage, we use the official code of TinyTL [4] to calculate both the model parameters and the intermediate activation size, following the previous studies [22, 59, 55]. The time indicates the amount of time consumed for the forward process and the backpropagation. Since we utilize the outputs of θ_o , our method accompanies an additional for-

Method	ResNet50 [17]		WDR28 [61]	
	Closed	Open	Closed	Open
Source	48.80	48.80	43.52	43.52
BN Adapt [43]	16.01	16.89	20.43	23.61
TENT [57]	61.69	83.62	56.00	77.72
+ Ours	15.28 (-46.41)	16.99 (-66.63)	20.16 (-35.84)	23.70 (-54.02)
SWR [9]	16.19	88.53	17.94	90.15
+ Ours	16.08 (-0.11)	71.83 (-16.70)	15.35 (-2.59)	83.76 (-6.39)

(a) Long-term adaptation

Method	ResNet50 [17]		WDR28 [61]	
	Closed	Open	Closed	Open
Source	48.80	48.80	43.52	43.52
BN Adapt [43]	16.01	16.89	20.43	23.61
TENT [57]	14.03	22.76	18.23	32.74
+ Ours	13.82 (-0.21)	16.36 (-6.40)	18.32 (+0.09)	23.40 (-9.34)
SWR [9]	13.81	45.58	16.62	83.08
+ Ours	13.80 (-0.01)	43.35 (-2.23)	15.73 (-0.89)	75.89 (-7.19)

(b) Short-term adaptation

Table 7: Error rates of image classification on CIFAR-10-C using diverse architectures.

Method	Learning rate				Std.↓
	0.005	0.001	0.0005	0.0001	
Source	76.71	76.71	76.71	76.71	0
TENT [57]	99.51	89.91	75.02	63.83	15.79
+ Ours	64.14	60.04	59.59	58.76	2.40

(a) Robustness to learning rates

Method	Batch size				Std.↓
	64	32	16	8	
Source	76.71	76.71	76.71	76.71	0
TENT [57]	67.54	72.62	81.21	94.75	11.90
+ Ours	60.32	62.14	65.88	73.83	5.99

(b) Robustness to batch sizes

Table 8: Error rates of image classification on TinyImageNet-C with diverse learning rates and batch sizes. Std. is the abbreviation of the standard deviation.

ward process. However, as shown, such an additional forward process is negligible compared to the state-of-the-art models. For example, our method applied to TENT brings a significant performance gain with only half the memory and time compared to SWR [9]. Further details on resource costs, along with the results on semantic segmentation, are included in our supplementary.

4.3. Applicability on Various Models

Since our method focuses on improving the pseudo-labeling quality of entropy minimization, it does not rely on model architectures. Table 7 shows that applying our method consistently outperforms the baseline models with ResNet50 [17] and WideResNet28 [61] that were used in previous TTA studies [9, 39]. Such results demonstrate that our method is widely applicable to various architectures.

4.4. Robustness to Hyper-parameters

In real-world applications, we may not know an adequate learning rate before encountering test samples or may not use an optimal batch size due to memory constraints. In such a case, we need an approach with a stable performance regardless of such hyper-parameters. Table 8 shows that our method is more robust to such hyper-parameters compared to TENT [57], which is highly dependent on them. Such results demonstrate the scalability of our method when we do not know the optimal hyper-parameters.

5. Related Work

Test-Time Adaptation The main differences between TTA studies and other studies addressing domain shifts such as domain generalization [66, 37, 26, 15, 42, 52] or unsupervised domain adaptation (UDA) [49, 7, 38, 41, 36] is that TTA studies do not utilize 1) the source data during the adaptation stage and 2) ground truth labels on the target distribution [57, 58, 46, 14, 9, 39]. Recent studies [58, 39, 12] show that TTA models suffer from a large performance

degradation with continually changing domains and a long-term adaptation. To tackle such a challenging problem, this paper mainly evaluates long-term adaptation with continually changing domains.

Noisy Signals in Test-Time Adaptation As aforementioned, one of the key challenges in TTA is that the model is prone to utilizing wrong predictions. Preventing the model from learning with noisy supervision has been studied widely beyond TTA [64, 49, 21, 16, 1, 32, 48]. However, the main difference between TTA and these studies is that TTA studies assume that we cannot revisit the sample after performing adaptation with it. Such an assumption limits from proposing methods that require knowledge of the full data distributions [53, 1] or consistency of predictions for a given sample [49, 54]. Without such knowledge, we use the difference of confidence scores between θ_o and θ_a by using the wisdom of crowds to improve pseudo labeling.

6. Conclusion

This paper proposed a *simple yet effective* data sample selection that is widely applicable to existing various test-time adaptation methods. Based on the observation that signals from wrong samples fail to increase the confidence values of the predicted labels even with entropy minimization, we only select the samples that achieve higher confidence values with the adaptation model compared to those with the original model. This is mainly due to the wisdom of crowds, the dominant signals generally found in the correct samples influencing signals of other samples. Our method improved TTA performance on the existing TTA methods on both image classification and semantic segmentation. Additionally, we proposed a novel evaluation setting, an open-set TTA, which was overlooked until now even with its importance and practicality. We hope our work inspires future researchers to conduct more practical TTA research that improves both closed-set and open-set TTA.

Acknowledgement We would like to thank Kyuwoong Hwang, Sungrack Yun, Simyung Chang, Hyunsin Park, Janghoon Cho, Juntae Lee, Hyoungwoo Park, Seokeon Choi, Seunghan Yang, and Sunghyun Park of the Qualcomm AI Research team for their valuable discussions.

References

- [1] Unsupervised label noise modeling and loss correction. In *ICML*, 2019. 8
- [2] Eric Arazo, Diego Ortego, Paul Albert, Noel E O'Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *IJCNN*, 2020. 4
- [3] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Armand Joulin, Nicolas Ballas, and Michael G. Rabbat. Semi-supervised learning of visual features by non-parametrically predicting view assignments with support samples. In *ICCV*, 2021. 5
- [4] Han Cai, Chuang Gan, Ligeng Zhu, and Song Han. Tinytl: Reduce memory, not parameters for efficient on-device learning. *NeurIPS*, 2020. 7
- [5] Dian Chen, Dequan Wang, Trevor Darrell, and Sayna Ebrahimi. Contrastive test-time adaptation. In *CVPR*, 2022. 5
- [6] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 6
- [7] Minghao Chen, Hongyang Xue, and Deng Cai. Domain adaptation for semantic segmentation with maximum squares loss. In *ICCV*, 2019. 8
- [8] Sungha Choi, Sanghun Jung, Huiwon Yun, Joanne T Kim, Seungryong Kim, and Jaegul Choo. Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening. In *CVPR*, 2021. 1
- [9] Sungha Choi, Seunghan Yang, Seokeon Choi, and Sungrack Yun. Improving test-time adaptation via shift-agnostic weight regularization and nearest source prototypes. *ECCV*, 2022. 2, 5, 6, 7, 8
- [10] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016. 1, 5, 7
- [11] Debasmit Das, Shubhankar Borse, Hyojin Park, Kambiz Azarian, Hong Cai, Risheek Garrepalli, and Fatih Porikli. Transadapt: A transformative framework for online test time adaptive semantic segmentation. In *ICASSP*, 2023. 1
- [12] Yulu Gan, Yan Bai, Yihang Lou, Xianzheng Ma, Renrui Zhang, Nian Shi, and Lin Luo. Decorate the newcomers: Visual domain prompt for continual test time adaptation, 2022. 8
- [13] Taesik Gong, Jongheon Jeong, Taewon Kim, Yewon Kim, Jinwoo Shin, and Sung-Ju Lee. Robust continual test-time adaptation: Instance-aware bn and prediction-balanced memory. *NeurIPS*, 2023. 2
- [14] Sachin Goyal, Mingjie Sun, Aditi Raghunathan, and Zico Kolter. Test-time adaptation via conjugate pseudo-labels, 2022. 2, 5, 6, 8
- [15] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. *ICLR*, 2021. 8
- [16] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, 2018. 8
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 8
- [18] Dan Hendrycks, Steven Basart, Mantas Mazeika, Mohamadreza Mostajabi, Jacob Steinhardt, and Dawn Xiaodong Song. Scaling out-of-distribution detection for real-world settings. In *ICML*, 2022. 7
- [19] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *ICLR*, 2017. 7
- [20] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. *CVPR*, 2021. 5
- [21] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *ICML*, 2018. 8
- [22] Ziyu Jiang, Xuxi Chen, Xueqin Huang, Xianzhi Du, Denny Zhou, and Zhangyang Wang. Back razor: Memory-efficient transfer learning by self-sparsified backpropagation. In *NeurIPS*, 2022. 7
- [23] Sanghun Jung, Jungsoo Lee, Daehoon Gwak, Sungha Choi, and Jaegul Choo. Standardized max logits: A simple yet effective approach for identifying unexpected road obstacles in urban-scene segmentation. In *ICCV*, 2021. 7
- [24] Sanghun Jung, Jungsoo Lee, Nanhee Kim, Amirreza Shaban, Byron Boots, and Jaegul Choo. Cafa: Class-aware feature alignment for test-time adaptation, 2022. 2
- [25] Tommie Kerssies, Joaquin Vanschoren, and Mert Kılıçkaya. Evaluating continual test-time adaptation for contextual and semantic domain shifts. *arXiv preprint arXiv:2208.08767*, 2022. 4
- [26] Daehee Kim, Youngjun Yoo, Seunghyun Park, Jinkyu Kim, and Jaekoo Lee. Selfreg: Self-supervised contrastive regularization for domain generalization. In *ICCV*, 2021. 8
- [27] Jin Kim, Jiyoung Lee, Jungin Park, Dongbo Min, and Kwanghoon Sohn. Pin the memory: Learning to generalize semantic segmentation. In *CVPR*, 2022. 1
- [28] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *ICLR*, 2015. 6
- [29] Andreas Krause, Pietro Perona, and Ryan Gomes. Discriminative clustering by regularized information maximization. In *NeurIPS*, 2010. 5
- [30] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. 5

- [31] Ya Le and Xuan S. Yang. Tiny imagenet visual recognition challenge. 2015. 5
- [32] Jungsoo Lee, Jeonghoon Park, Daeyoung Kim, Juyoung Lee, Edward Choi, and Jaegul Choo. Revisiting the importance of amplifying bias for debiasing. *AAAI*, 2023. 8
- [33] Suhyeon Lee, Hongje Seong, Seongwon Lee, and Euntai Kim. Wildnet: Learning domain generalized semantic segmentation from the wild. In *CVPR*, 2022. 1
- [34] Jichang Li, Guanbin Li, Feng Liu, and Yizhou Yu. Neighborhood collective estimation for noisy label identification and correction. *ECCV*, 2022. 4
- [35] Junnan Li, Richard Socher, and Steven C.H. Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. In *ICLR*, 2020. 4
- [36] Rui Li, Qianfen Jiao, Wenming Cao, Hau-San Wong, and Si Wu. Model adaptation: Unsupervised domain adaptation without source data. In *CVPR*, 2020. 8
- [37] Xiaotong Li, Yongxing Dai, Yixiao Ge, Jun Liu, Ying Shan, and Ling-Yu Duan. Uncertainty modeling for out-of-distribution generalization. *ICLR*, 2022. 8
- [38] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? Source hypothesis transfer for unsupervised domain adaptation. In *ICML*, 2020. 5, 8
- [39] Hyesu Lim, Byeongeun Kim, Jaegul Choo, and Sungha Choi. TTN: A domain-shift aware batch normalization in test-time adaptation. In *ICLR*, 2023. 1, 5, 6, 8
- [40] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. *NeurIPS*, 2020. 7
- [41] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Unsupervised domain adaptation with residual transfer networks. *NeurIPS*, 2016. 8
- [42] Lucas Mansilla, Rodrigo Echeveste, Diego H. Milone, and Enzo Ferrante. Domain generalization via gradient surgery. *ICCV*, 2021. 8
- [43] Zachary Nado, Shreyas Padhy, D Sculley, Alexander D’Amour, Balaji Lakshminarayanan, and Jasper Snoek. Evaluating prediction-time batch normalization for robustness under covariate shift. *arXiv preprint arXiv:2006.10963*, 2020. 5, 6, 8
- [44] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. 2011. 5
- [45] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Bulo, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *ICCV*, 2017. 5, 7
- [46] Shuaicheng Niu, Jiayang Wu, Yifan Zhang, Yafo Chen, Shijian Zheng, Peilin Zhao, and Mingkui Tan. Efficient test-time model adaptation without forgetting. *ICML*, 2022. 2, 5, 6, 7, 8
- [47] Shuaicheng Niu, Jiayang Wu, Yifan Zhang, Zhiquan Wen, Yafo Chen, Peilin Zhao, and Mingkui Tan. Towards stable test-time adaptation in dynamic wild world. In *ICLR*, 2023. 2
- [48] Junwoo Park, Jungsoo Lee, Youngin Cho, Woncheol Shin, Dongmin Kim, Jaegul Choo, and Edward Choi. Deep imbalanced time-series forecasting via local discrepancy density, 2023. 8
- [49] Viraj Prabhu, Shivam Khare, Deeksha Kartik, and Judy Hoffman. Sentry: Selective entropy optimization via committee consistency for unsupervised domain adaptation. In *ICCV*, 2021. 8
- [50] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *ECCV*, 2016. 5, 7
- [51] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *CVPR*, 2016. 5
- [52] Yuge Shi, Jeffrey Seely, Philip H. S. Torr, N. Siddharth, Awni Hannun, Nicolas Usunier, and Gabriel Synnaeve. Gradient matching for domain generalization. *ICLR*, 2022. 8
- [53] Hwanjun Song, Minseok Kim, and Jae-Gil Lee. Selfie: Refurbishing unclean samples for robust deep learning. In *ICML*, 2019. 8
- [54] Hwanjun Song, Minseok Kim, Dongmin Park, Yooju Shin, and Jae-Gil Lee. Robust learning by self-transition for handling noisy labels. In *KDD*, 2021. 8
- [55] Junha Song, Jungsoo Lee, In So Kweon, and Sungha Choi. Ecotta: Memory-efficient continual test-time adaptation via self-distilled regularization. *CVPR*, 2023. 7
- [56] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *NeurIPS*, 2017. 4
- [57] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. *ICLR*, 2021. 1, 2, 3, 5, 6, 7, 8
- [58] Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual test-time domain adaptation. In *CVPR*, 2022. 2, 5, 8
- [59] Li Yang, Adnan Siraj Rakin, and Deliang Fan. Rep-net: Efficient on-device learning via feature reprogramming. In *CVPR*, 2022. 7
- [60] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *CVPR*, 2020. 1, 5, 7
- [61] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *BMVC*, 2016. 5, 8
- [62] Marvin Zhang, Sergey Levine, and Chelsea Finn. Memo: Test time robustness via adaptation and augmentation. *arXiv preprint arXiv:2110.09506*, 2021. 2
- [63] Zixing Zhang, Fabien Ringeval, Bin Dong, Eduardo Coutinho, Erik Marchi, and Björn Schüller. Enhanced semi-supervised learning for multimodal emotion recognition. In *ICASSP*, 2016. 4
- [64] Zhilu Zhang and Mert R. Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels. In *NeurIPS*, 2018. 5, 6, 8
- [65] Bowen Zhao, Chen Chen, and Shu-Tao Xia. Delta: degradation-free fully test-time adaptation. *ICLR*, 2023. 2
- [66] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. *ICLR*, 2021. 8