

Test-time Adaptation in the Dynamic World with Compound Domain Knowledge Management

Junha Song¹, Kwanyong Park¹, InKyu Shin¹, Sanghyun Woo¹, Chaoning Zhang², and In So Kweon¹

Abstract—Prior to the deployment of robotic systems, pre-training the deep-recognition models on all potential visual cases is infeasible in practice. Hence, test-time adaptation (TTA) allows the model to adapt itself to novel environments and improve its performance during test time (i.e., lifelong adaptation). Several works for TTA have shown promising adaptation performances in continuously changing environments. However, our investigation reveals that existing methods are vulnerable to dynamic distributional changes and often lead to overfitting of TTA models. To address this problem, this paper first presents a robust TTA framework with compound domain knowledge management. Our framework helps the TTA model to harvest the knowledge of multiple representative domains (i.e., compound domain) and conduct the TTA based on the compound domain knowledge. In addition, to prevent overfitting of the TTA model, we devise novel regularization which modulates the adaptation rates using domain-similarity between the source and the current target domain. With the synergy of the proposed framework and regularization, we achieve consistent performance improvements in diverse TTA scenarios, especially on dynamic domain shifts. We demonstrate the generality of proposals via extensive experiments including image classification on ImageNet-C and semantic segmentation on GTA5, C-driving, and Cityscapes datasets.

I. INTRODUCTION

One of the main challenges in deep neural networks is the decrease in performance during test time caused by domain shift between the pre-training domain (source) and the test domain (target) [1]. This issue should be resolved especially for safety-critical applications like robots and autonomous vehicles, as it could pose a significant risk to application users. For example, in autonomous driving scenarios, failures of the recognition model on novel driving conditions may cause accidents and harm occupants.

Several research fields attempt to address this problem, such as domain generalization (DG) [2] and unsupervised domain adaptation (UDA) [3]. In particular, UDA performs joint training using both the source and target domain during the *pre-training* phase, applying techniques such as adversarial [1] or pseudo-labeling [4] learning. However, it may be challenging to utilize the UDA techniques in the following realistic cases: 1) where the target data is unavailable, or 2) when the model is deployed on unseen domains which are different from the target domain used in pre-training. Motivated by this issue, test-time adaptation (TTA) has been

¹J. Song, K. Park, I. Shin, S. Woo, and I. S. Kweon are with the School of Electrical Engineering, KAIST, Daejeon 34141, Republic of Korea. {sb020518, pkyong7, dlsrbgg33, shwoo93, and iskweon77}@kaist.ac.kr

²C. Zhang is with the School of Artificial Intelligence, Kyung Hee, Yongin-si, 17104, Republic of Korea. chaoningzhang1990@gmail.com

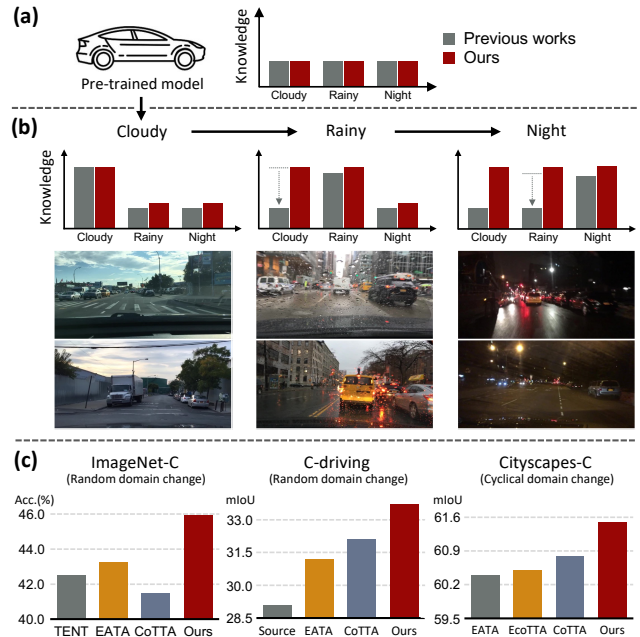


Fig. 1. We present a conceptual illustration of our paper (a,b) and the experiment results (c). (a) **Before model deployment**, limited data may result in insufficient knowledge of some domains. (b) **During Test-time adaptation (TTA)**, the model can adapt to unfamiliar domains itself. Previous works only focus on adapting to the current domain while ignoring previous adaptations. On the other hand, our approach learns multiple knowledge from the compound domain, resulting in robust performance in diverse TTA scenarios, as shown in the **experimental results** (c).

presented [5], assuming that the networks are pre-trained with only the source domain and the adaptation to the target domain is conducted during the *testing* phase.

TTA is a promising research direction for lifelong adaptation in robotic systems, which enables robots to quickly adapt to a new environment and leads to a more robust performance over time. Due to its usefulness, TTA has garnered a great deal of attention and enjoyed a rapid performance boost by improved unsupervised loss [6], [7], [8] or small-batch adaptability [9], [10]. In particular, CoTTA [11] and EcoTTA [12] presented a promising potential for lifelong adaptation by facilitating the prevention of catastrophic forgetting [13] and error accumulation [12]. Despite these TTA breakthroughs, we have identified a notable issue with previous works. As depicted in Figure 1(c), they demonstrate decreased performance in situations where dynamic domain changes occur, such as when encountering tunnels, sun glare, or rain while driving. We speculate that it is because previous works barely consider managing the knowledge of multiple sub-target domains (e.g., weather or location changes), only focusing on the adaptation to the current domain while

discarding the knowledge acquired from previous domains, as illustrated in Figure 1(b).

Therefore, this paper proposes a simple yet effective approach to address this issue. We first introduce a robust TTA framework that empowers the TTA model to learn and manage multiple knowledge from the compound domain. This framework utilizes our proposed continual domain-matching algorithm that efficiently discriminates the current image’s domain type (*i.e.*, pseudo-domain label). Specifically, we continually match incoming target images with our domain prototypes and then estimate the domain type among the multiple sub-target domains. Moreover, we introduce a novel regularization term to prevent overfitting of the TTA model in the compound domain. We observe that the quality of unsupervised signals is significantly varying depending on which domain the image belongs. For example, the unsupervised signals generated by night images are often noisy, compared to the counterpart of the daytime images. Thus, we modulate the impact of the unsupervised signals on the model by leveraging the domain-similarity between the source and the sub-target domain.

Our paper presents the following contributions:

- We propose a new TTA framework that helps the TTA model to manage compound domain knowledge, which effectively copes with dynamic domain shifts.
- We present a domain similarity-based loss regularization to adjust different adaptation rates for various domain types, improving TTA performance in the wild.
- We evaluate our approach in diverse scenarios that can occur in the dynamic world. Our experiments consist of 1) image classification task with ImageNet-C dataset and 2) semantic segmentation task with GTA5, C-driving, and Cityscapes datasets.

II. RELATED WORK

A. Domain Adaptation

To mitigate the performance drop caused by the difference between the train and test domains, numerous methods for unsupervised domain adaptation (UDA) [3], [1] have been proposed. Most UDA methods alleviate the domain shift by utilizing adversarial training [4] and self-training [14]. While the UDA methods have demonstrated their efficacy, their experimental settings had a practical limitation; the target domain was assumed to come from a single distribution (*e.g.*, only sunny) even though the target domain can be diverged into sub-target domains (*e.g.*, cloudy, rainy, and snowy) due to time and weather variations. For this reason, the setup called compound domain adaptation (CDA) emerged [15], [16], taking multiple sub-target domains (*i.e.*, compound domain) into account. For example, DHA-CDA [16] learns invariant representations under the compound domain by adversarially pre-training the networks with both source and multiple target datasets. Despite the development of DA techniques, the DA field has overlooked an important fact that the pre-trained networks can undergo the target domain which was not considered during the pre-training phase.

For instance, self-driving cars typically operate in pleasant daytime conditions, but they may also encounter unfamiliar target domains such as snow, fog, sun glare, and tunnel.

B. Test-time Adaptation

Pre-training the model with all potential target domains is difficult and expensive. This issue has led to the development of a new setting called as test-time adaptation (TTA), in which the adaptation is performed using the unlabeled target sample encountered by the model after deployment, regardless of the source domain on which the model was pre-trained. TENT [5], a seminal work in TTA, proposes updating only small parameters (*i.e.*, batch normalization layers) by minimizing the entropy of model predictions during test time. Following TENT, several TTA works improve TTA performance by designing sophisticated unsupervised loss [6], [17], [18], [19] or improving small-batch scalability [10], [20]. Moreover, since adaptation with unsupervised loss can cause error accumulation [12] and catastrophic forgetting [13], CoTTA [11] and EcoTTA [12] perform long-term TTA under a continually changing target domain (*i.e.*, continual TTA) with weight restoration or self-distilled regularization, respectively. Nevertheless, previous studies only consider domains that change in a continuous or monotonic fashion, and as a result, they tend to underperform when faced with more complex types of domain shifts, such as random or cyclical changes. Given this issue, this paper proposes a novel solution that involves managing compound domain knowledge and regulating the model’s adaptation rate for each domain separately.

III. APPROACH

A. Prerequisite

TTA models conduct the inference and adaptation simultaneously during test time to acquire more accurate prediction results in the target domain. However, performing TTA for a long period of time inevitably causes target domain overfitting [11], [12], which means the model collapses and become to infer inaccurate predictions. This issue occurs because TTA models utilize an unsupervised loss with unlabeled target images, *i.e.*, an unreliable backpropagation signal compared with supervised loss. Obviously, this issue can be avoided by adaptation early-stopping [21], but this approach is not suitable for TTA because TTA is conducted under the scenario that it is unknown when the unfamiliar target domain will emerge.

Our work, therefore, considers long-term TTA as a basic TTA setting and leverages the adaptation loss and weight regularization proposed in EATA [22] as a baseline to prevent overfitting. In detail about EATA, they only update the parameters Θ of batch normalization (BN) layers and freeze the rest in the networks during the adaptation process. To generate the backward propagation from i -th data x_i in a mini-batch of n target images, entropy minimization with sample filtering \mathcal{L}^{adapt} is computed, which is formulated as:

$$\mathcal{L}^{adapt} = \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{H(\hat{y}_i) < H_0\}} \cdot H(\hat{y}_i), \quad (1)$$

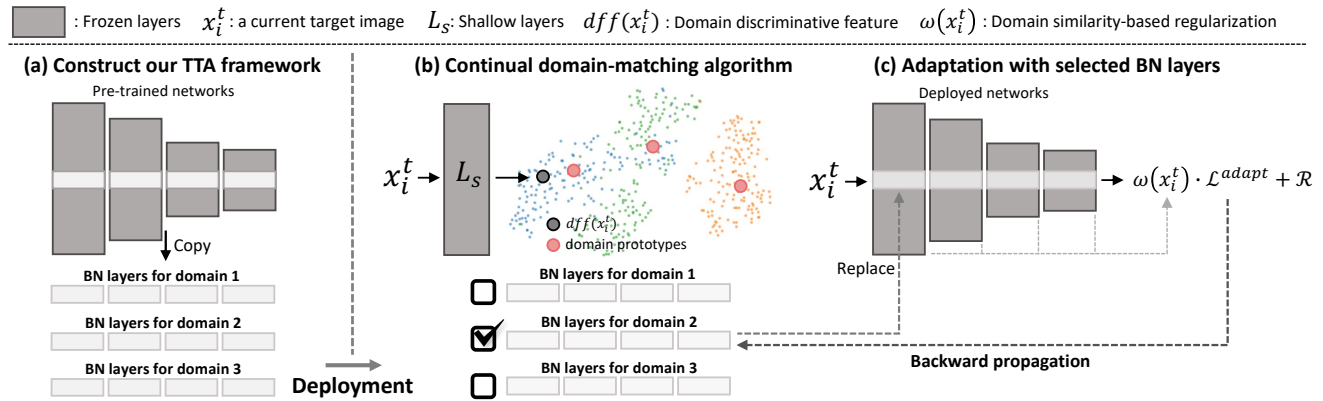


Fig. 2. **Overview of our approach.** (a) Before deployment, we construct a TTA framework for compound domain knowledge, cloning the batch normalization modules in the pre-trained networks. (b) In order to select one of the modules, we determine the domain type of the current target image by our proposed continual domain-matching algorithm, helping the TTA model to manage the adaptation knowledge of the compound domain. (c) We perform test-time adaptation with the selected modules, where our domain similarity-based regularization $\omega(x_i^t)$ is also applied to adjust different adaptation rates for the multiple sub-target domains. \mathcal{L}^{adapt} and \mathcal{R} are described in Section III-A.

where $H(y) = -\sum_C p(y) \log p(y)$ is a formula of entropy minimization and $\mathbb{I}_{\{\cdot\}}$ is an indicator function. \hat{y}_i means the logits output of x_i , and H_0 , $p(\cdot)$, and C refer to a pre-defined threshold, the softmax function, and the number of task classes, respectively. Moreover, they apply the following weight regularization \mathcal{R} to prevent model parameters from changing too much and thus avoid overfitting,

$$\mathcal{R}(\tilde{\theta}, \tilde{\theta}^o) = \sum_{\theta_i \in \tilde{\theta}} w(\theta_i) (\theta_i - \theta_i^o)^2. \quad (2)$$

$\tilde{\theta}$ are adapted parameters of the networks, $\tilde{\theta}^o$ are the corresponding frozen parameters of the pre-trained networks which are extra-stored, and $w(\theta_i)$ denotes the importance of θ_i which is calculated by the diagonal fisher information matrix [23]. Note that our approach introduced in the next subsections, which is orthogonal to the approach proposed in EATA and can also be applied to existing continual TTA works, such as CoTTA [11] and EcoTTA [12].

B. TTA Framework for compound domain knowledge

In the dynamic world, TTA models can face various environmental changes, including sudden domain shifts such as driving through tunnels and experiencing sun glare. Therefore, it is crucial for TTA models to accommodate compound domain knowledge and conduct the TTA without latency. In this subsection, we (i) present our TTA framework for managing compound domain knowledge, (ii) domain distinctive features, and (iii) continual domain-matching algorithm to estimate the domain type of a current target image. Figure 2 provides the overview of our approach.

Compound domain TTA framework. Previous works for adapting to multiple domains, such as DSNB [24] and DSON [25], indicate that composing domain-specific modules helps the model to effectively learn multiple domain knowledge. Inspired by such a finding, we also construct K -specific modules consisting of BN layers to manage domain knowledge acquired from multiple sub-target domains. To be more specific, as shown in Figure 2(a), we initialize the modules as cloning BN layers of the pre-trained model before deployment. During TTA, each module becomes specialized for each sub-target domain, as updating each module

matched with the estimated domain type of an incoming image, as depicted in Figure 2(b,c).

Domain distinctive features. To interpret the domain information implied in an input image, we need to define domain-distinctive features ddf , a prerequisite for our domain-matching algorithm (detailed in the next paragraph). We utilize style features proposed in the field of style transfer [26]. Specifically, convolutional feature statistics from small parts of the shallow layers L_s (e.g., conv1 and layer1 in the case of ResNet) are used as the ddf , which can indicate the type of domain (i.e., style) of the current image [27]. Consequently, our ddf vectors are defined as:

$$ddf(x_i) = \{\mu(\phi_0(x_i), \sigma(\phi_0(x_i), \mu(\phi_1(x_i), \sigma(\phi_1(x_i))))\}, \quad (3)$$

where μ and σ denotes the feature statistics and ϕ_0 and ϕ_1 are the final output of the conv1 and layer1, respectively.

Continual domain-matching algorithm. With the ddf defined, we need a fast way to determine the pseudo-domain label to make K -specific modules optimized to each sub-target domain. As proposed in [27], kmeans clustering [28] can be one of the potential solutions, but we note that such iterative clustering algorithms require long running time when we deal with enormous vectors (i.e., all ddf vectors of incoming target samples). Instead of the clustering algorithm, we focus on an effective matching algorithm and thus propose a continual domain-matching algorithm. Our algorithm, visualized in Figure 3 and defined as the following, is continuously executed as the target image incomes:

- 1) Compute the distances between the domain prototypes $\{d_k\}_{k=1}^K$ and the ddf of the current image x_i .
- 2) Assign the index l of the nearest domain prototype d_l as the pseudo domain label for the x_i .
- 3) Update the selected domain prototype by moving average technique formulated as $d_l \leftarrow \eta \cdot d_l + (1-\eta) \cdot ddf(x_i)$, where η represents momentum.
- 4) Conduct test-time adaptation with the l -th module, as shown in Figure 2(c).

Specifically, the distance function is used as the Bhattacharya function [29] which is useful for comparing two statistics




Fig. 3. **Continual domain-matching algorithm** Our fast matching algorithm allows the model to effectively predict the pseudo domain label of target images. Through this process, each domain prototype becomes representative of each type of sub-domains.

(p, q) and formulated as:

$$D_B(p, q) = \frac{1}{4} \ln \left(\frac{1}{4} \left(\frac{\sigma_p^2}{\sigma_q^2} + \frac{\sigma_q^2}{\sigma_p^2} + 2 \right) \right) + \frac{1}{4} \left(\frac{(\mu_p - \mu_q)^2}{\sigma_p^2 + \sigma_q^2} \right). \quad (4)$$

Through our proposed algorithm, each domain prototype is derived to represent each sub-target domain. In addition, the domain prototypes are initialized before model deployment by utilizing kmeans clustering with the *ddf*s of the source dataset. We highlight that access to the source data prior to deployment is possible in cases where the source dataset is publicly accessible or if the owner of a pre-trained model attempts to adapt the model to a target domain, as stated in previous TTA works [30], [12], [22].

C. Domain Similarity-based Loss Regularization.

TTA models adapt to new domains using *unsupervised* loss (e.g., entropy minimization [5]) since *unlabeled* target images are typically collected in edge devices. Unfortunately, the unsupervised loss is less reliable than the supervised loss, occasionally containing false backpropagation signals when the calculated loss comes from wrong predictions. Existing TTA works, such as EATA [22] and EcoTTA [12], demonstrate that this issue leads to error accumulation and model failure (i.e., overfitting).

We observe that the pace of error accumulation varies according to the sub-target domain since images gathered in the compound domain make different qualities of unsupervised loss. For example, night images typically provide less accurate predictions than daytime images and thus the night-specific module becomes quickly collapsed (i.e., overfitted). Therefore, we attempt to alleviate error accumulation by regularizing the adaptation pace, especially when an unsupervised loss is likely to be more unreliable since the sub-target domain is significantly distinct from the pre-training domain. Specifically, we compute a loss regularization term ω leveraging the similarity between target statistics calculated from a single image and source statistics \mathcal{S} in BN layers (i.e., running estimates [31]), which is defined as:

$$\omega(x_i) = \mathbb{E}_{l \sim L} [\text{sim} \{ (\mu(\phi_l(x_i)), \sigma(\phi_l(x_i))), \mathcal{S}_l \}^\gamma], \quad (5)$$

where we use cosine similarity, L denotes last BN layers in layer1~layer4 in the case of ResNet backbone [31]. The scale parameter γ controls the strength of the regularization. Overall, the total loss \mathcal{L}^{total} for test-time adaptation to the target domain is formulated as:

$$\mathcal{L}^{total} = \omega(x_i) \cdot \mathcal{L}^{adapt} + \mathcal{R}. \quad (6)$$

It should be noted that ω and \mathcal{R} respectively help to alleviate catastrophic forgetting and error accumulation so as to facilitate long-term adaptation.

IV. EXPERIMENTAL RESULTS

In this section, we present and verify the effectiveness of our approach on various benchmarks, including image classification and semantic segmentation. Additional empirical studies that provide a deeper understanding of our approach are referenced in [32].

A. Classification Experiments

Experimental setup. We use ImageNet and ImageNet-C dataset [33] as the source (i.e., pre-training) and the target (i.e., adaptation) domain, respectively. Following the previous TTA studies [22], [11], we conduct experiments in a continual TTA setup, where we continually adapt the deployed model to each corruption type sequentially without resetting the model. Specifically, by using the official code of CoTTA [11], evaluations are executed on the continual TTA setup with ten diverse corruption-type sequences, which provide a more comprehensive evaluation of TTA methods. We emphasize that diverse sequences include both the continuous flow as well as *random shifts* in corruption types.

Implementation Details. We use the ResNet-50 pre-trained with AugMix [34], which is officially provided in RobustBench [33]. After deployment, we utilize the SGD optimizer, the learning rate of 5e-3, the batch size of 32, and the same hyperparameters for adaptation loss \mathcal{L}^{adapt} and weight regularization \mathcal{R} as EATA [22] (such as H_0 of $0.4 \times \ln C$, where C is the number of task classes). We set the number of specific modules K , the scale parameter γ , the momentum η , to 3, 1.5, and 0.9, respectively.

Comparison with TTA methods. We report the average error rate (%) of ten diverse sequences in Table I and we detail two notable sequences among the ten sequences in Table II. The first sequence means a scenario where corruptions belonging to the same main category *continuously* appear, whereas the second sequence depicts a situation where all corruption types are *randomly* distributed. The results demonstrate that our technique achieves robust TTA performance even with random domain change, compared to competing TTA methods. To be more specific, our approach shows similar performance regardless of continuous and random changing. In contrast, EATA, CoTTA, and EcoTTA show significantly degraded performance in the random domain change. We suggest that separating and managing

TABLE I

COMPARISON (μ) ON IMAGENET-C. WE REPORT THE AVERAGE ERROR RATE (%) OVER 10 DIVERSE CORRUPTION SEQUENCES. THE VALUES IN PARENTHESES REPRESENT STANDARD DEVIATIONS FOR TEN SEQUENCES. THE WALL-CLOCK TIME REFERS TO THE TOTAL TIME TAKEN TO ADAPT TO ONE SEQUENCE COMPOSED OF 15 CORRUPTIONS IN IMAGENET-C.

ImageNet-C	Source	TENT [5]	TTT+ [30]	EATA [22]	CoTTA [11]	EcoTTA [12]	Ours
Avg. error ↓	74.5 (0)	56.1 (0.6)	58.2 (0.8)	54.9 (2.3)	54.6 (3.9)	54.4 (2.7)	53.5 (1.6)
Wall time ↓	4m 35s	7m 18s	8m 12s	7m 23s	22m 37s	7m 51s	7m 52s

TABLE II

COMPARISON ON TWO NOTABLE SEQUENCES. WE CONDUCT EXPERIMENTS ON CONTINUAL TTA SETUP. AVG. ERR MEANS THE AVERAGE ERROR RATE (%) OF ALL 15 CORRUPTIONS. TWO BELOW SEQUENCES (IN SHORT, SEQ.) REPRESENT SITUATIONS IN WHICH 15 CORRUPTIONS OF THE FOUR MAIN CATEGORIES, *i.e.*, NOISE, BLUR, WEATHER, AND DIGITAL, OCCUR IN A CONTINUOUS AND RANDOM MANNER, RESPECTIVELY.

Time	$t \rightarrow$															
Continuous seq.	Gaus. N	Shot. N	Impu. N	Defo. B	Glas. B	Moti. B	Zoom B	Snow W	Fros. W	Fog. W	Brig. W	Cont. D	Elas. D	Pixe. D	Jpeg D	Avg. err ↓
Source	94.3	88.5	94.6	74.3	82.9	71.0	64.7	76.8	72.2	76.1	49.2	84.3	71.4	65.2	52.2	74.5
TENT [5]	61.6	58.9	56.8	59.5	61.6	56.0	53.9	56.9	59.6	52.4	45.2	66.3	52.9	53.9	55.9	56.8
TTT++ [30]	63.8	60.4	58.1	60.7	64.1	55.6	51.8	57.7	59.2	51.0	49.5	65.9	55.1	58.7	57.2	57.9
EATA [22]	61.1	57.1	53.2	55.6	59.6	54.2	51.6	55.9	56.7	51.1	45.4	52.9	49.1	49.2	49.8	53.5
CoTTA [11]	61.9	58.3	54.6	57.3	59.8	55.2	53.6	56.2	56.6	51.3	49.6	53.3	50.6	51.3	49.3	54.6
EcoTTA [12]	61.1	57.3	52.8	56.4	59.2	54.6	51.4	56.8	56.8	50.8	45.6	53.6	48.9	49.1	49.4	53.5
Ours	61.2	56.4	52.6	56.0	59.4	53.3	52.4	54.9	56.8	50.5	45.1	53.2	48.6	49.0	50.2	53.3

Random seq.	Cont. D	Gaus. N	Defo. B	Zoom B	Fog. W	Glas. B	Jpeg D	Fros. W	Pixe. D	Elas. D	Shot. N	Impu. N	Snow W	Moti. B	Brig. W	Avg. err ↓
Source	84.3	94.3	74.3	64.7	76.1	82.9	52.2	72.2	65.2	71.4	88.5	94.6	76.8	71.0	49.2	74.5
TENT [5]	56.4	60.6	62.6	52.2	57.2	62.4	50.8	60.8	54.8	51.0	62.5	59.1	61.9	59.8	54.0	57.5
TTT++ [30]	60.2	62.9	63.6	52.0	58.5	64.8	55.8	62.6	56.2	51.3	64.7	61.7	60.4	61.2	58.3	59.6
EATA [22]	54.9	66.4	62.2	57.0	58.9	62.4	52.4	57.9	52.1	50.3	57.6	57.2	57.7	55.8	48.5	56.8
CoTTA [11]	59.4	68.1	64.6	63.3	61.6	62.8	53.7	57.8	51.6	53.4	57.2	57.0	58.6	57.5	51.3	58.5
EcoTTA [12]	53.8	64.7	63.6	57.3	57.1	62.5	53.0	57.1	51.8	50.3	56.6	56.4	57.1	56.4	49.1	56.4
Ours	53.9	62.0	55.8	53.1	52.2	61.0	51.2	56.8	49.7	48.1	56.1	54.0	55.7	54.0	47.6	54.1

the compound domain knowledge enables our approach to adapt widely to all corruption types and perform well even in randomly changing domains.

In addition, we measure the wall-clock time for each TTA method to adapt to all 15 corruptions as presented in Table I. Our approach requires only a negligible overhead compared to CoTTA. Furthermore, we evaluate the efficiency of our domain-matching algorithm against the kmeans clustering. To adapt to only one corruption (*i.e.*, 5000 samples), our TTA approach takes 32s (156fps), whereas it takes 28m 16s (3fps) when we simply replace our domain-matching algorithm with the kmeans clustering. This suggests that the iterative clustering algorithm may not be suitable for TTA in practice.

B. Segmentation Experiments

Experimental setup. We conduct our experiments with {GTA5, C-driving} and {Cityscapes, Cityscapes with corruption} where the former and latter dataset represents the source and the target domain, respectively. The GTA5 [35] dataset contains synthetic images from the video game, Grand Theft Auto 5, whereas the C-driving dataset [15] includes the real-world samples which are separated by weather attributes (*i.e.*, cloudy, rainy, and snowy) or time-of-day attributes (*i.e.*, daytime, twilight, and night). The Cityscapes with corruption [6] (Cityscapes-C) are created by applying the 4 types of corruption selected from each of the main corruption categories to the validation set of Cityscapes [36]; we use motion blur, speckle noise (close to sensor noise), brightness (close to sun glare), and contrast (close to tunnel), which are illustrated in Figure 4. Following the experimental setup of CoTTA [11], we repeat the same sequence group (*e.g.*, weather or time-of-day attributes of



Fig. 4. Examples of Cityscapes with corruption. We evaluate our approach under the scenario where the four types of corruption cyclically come across. The experiment results can be found in Table IV.

C-driving dataset and four types of corrupted Cityscapes) 10 times so that we mimic the real-life scenario where driving environments might be revisited. It should be noted that this setup also shows the adaptation performance in the long term. Additionally, when the target is the C-driving dataset, we utilize 400 unlabeled images from the train set of each sub-attribute for the adaptation, as suggested in CoTTA.

Implementation Details. The source models for GTA5 and Cityscapes datasets are respectively based on DeepLabV3 and DeepLabV3+ with ResNet-50. The GTA5 model is pre-trained according to the official code and specifications of CBST [14], while the Cityscapes model is used as the publicly-available one from RobustNet [37]. For adaptation to the target domain, we use the batch size of 2, the SGD optimizer, and the learning rate of $1e-5$. The image size of the C-driving and Cityscapes datasets is set to 1280×720 and 1440×720 , respectively. As mentioned in Section III-A, our method can be utilized with existit TTA works; therefore, for semantic segmentation tasks, we evaluate our approach using the adaptation loss of CoTTA [11] which is cross entropy loss using the pseudo label with multi-scaling input (scales=[0.75, 1.0, 1.25, 1.5]) and flip. The rest of the implementation details are the same as for image classification.

Comparison with TTA methods. In the experiment of the GTA5 to C-driving dataset, we report the average mIoU (*i.e.*, mean intersection over union) score using six domain sequences made by *randomly* mixing the order of three sub-attributes. As shown in Table III, our approach achieves

TABLE III

COMPARISON ON C-DRIVING DATASET. THE AVERAGE mIoU SCORE FOR SIX DOMAIN SEQUENCES, CREATED BY RANDOMLY MIXING THREE SUB-ATTRIBUTES, IS REPORTED FOR GTA5 TO C-DRIVING EXPERIMENT.

Random seq.	Time-of-day				Weather			
	daytime	night	twilight	Mean	cloudy	rainy	snowy	Mean
Source	29.6	11.2	26.1	22.3	32.8	27.0	27.5	29.1
TENT [5]	18.9	10.5	17.4	15.6	23.9	17.4	22.0	21.1
EATA [22]	30.8	15.1	29.4	25.1	35.5	29.1	29.0	31.2
CoTTA [11]	32.1	15.4	30.2	25.9	36.4	29.6	30.2	32.1
Ours w/o \mathcal{R}	21.3	12.1	18.4	17.7	27.2	18.4	24.6	23.4
Ours w/o $\omega(x_i)$	33.3	15.6	32.3	27.0	37.1	30.8	31.8	33.2
Ours	33.6	16.5	32.7	27.6	37.5	31.2	32.3	33.7

TABLE IV

COMPARISON ON THE CORRUPTED CITYSCAPES DATASET. WE SIMULATE THE CYCLICAL DRIVING CONDITIONS ENCOUNTERED IN THE REAL WORLD. THE RESULTS SHOW IF THE TTA METHOD IS ABLE TO ADAPT OVER A LONG PERIOD OF TIME TO MULTIPLE SUB-TARGET DOMAINS.

Round Domain	1				4				7				10				mIoU \uparrow Mean
	Moti. B	Spec. N	Brig. W	Cont. D	Moti. B	Spec. N	Brig. W	Cont. D	Moti. B	Spec. N	Brig. W	Cont. D	Moti. B	Spec. N	Brig. W	Cont. D	
Source	31.3	11.5	60.9	27.3	31.3	11.5	60.9	27.3	31.3	11.5	60.9	27.3	31.3	11.5	60.9	27.3	32.8
TENT [5]	57.9	53.2	69.1	59.8	54.4	49.5	63.9	54.1	44.1	39.9	52.0	42.9	37.7	34.2	44.0	36.3	49.6
EATA [22]	58.1	53.5	69.6	60.7	57.5	53.5	69.6	60.7	57.5	53.5	69.6	60.7	57.5	53.5	69.6	60.7	60.4
EcoTTA [12]	58.1	53.5	69.8	60.7	57.8	53.5	69.8	60.7	57.8	53.5	69.8	60.7	57.8	53.5	69.8	60.7	60.5
CoTTA [11]	57.9	53.6	70.8	61.2	57.9	53.4	70.8	61.0	57.8	53.4	70.8	61.0	57.8	53.4	70.8	61.2	60.8
Ours	58.0	54.5	71.6	62.0	57.9	54.3	71.8	62.2	57.9	54.3	71.8	62.2	57.9	54.3	71.8	62.2	61.5

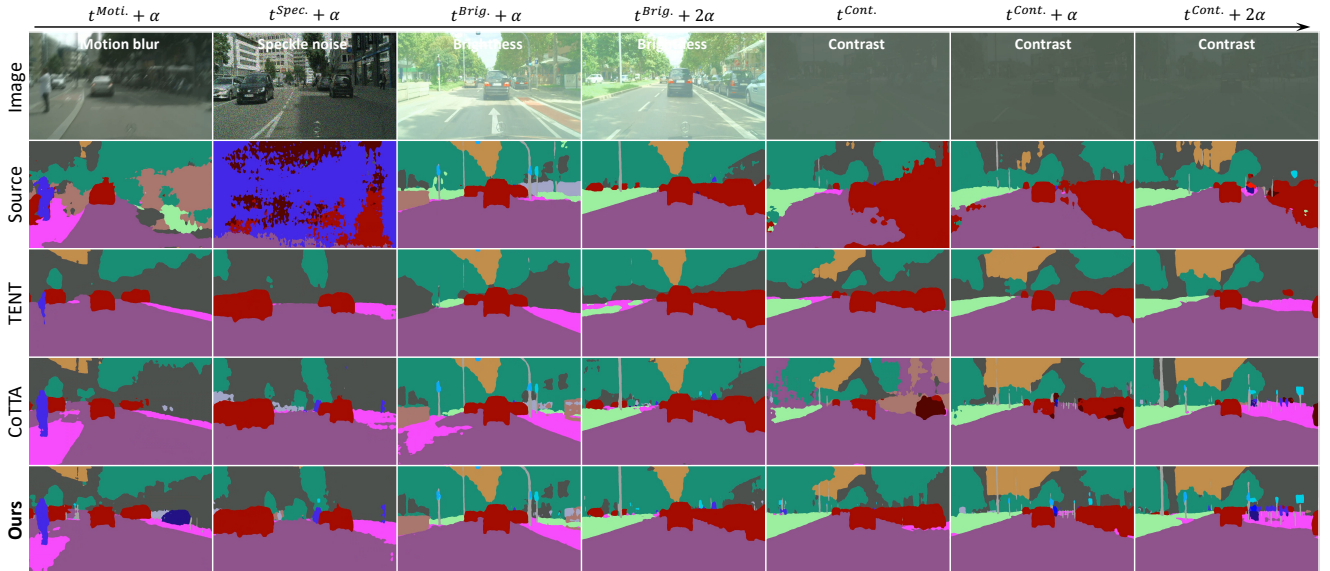


Fig. 5. **Qualitative comparison with TTA methods.** We visualize the prediction results in the fourth round. t^X denotes the moment samples corrupted by X type come in and α represents a short period of time (e.g., one second). The results of Source and TENT [5] reflect the problem of domain shift and overfitting to the target domain, respectively. In comparison to CoTTA [11], our approach shows robust performance even in dynamic domain change.

the best performance. We also describe the adaptation performance of our approach without our loss regularization term $\omega(x_t)$. The results indicate that our loss regularization contributes to improved performance, especially in domains that are challenging to generate a high-quality unsupervised loss, such as night, rainy, and snowy conditions.

We report the mIoU score of four types of corrupted Cityscapes across 10 *cyclical* rounds, as described in Table IV. Also, we provide qualitative comparisons as shown in Figure 5, where the prediction results in the *fourth* round are visualized using Cityscape video demo [36]. Our approach outperforms existing TTA approaches by acquiring comprehensive knowledge from the compound domain. TENT [5] suffers from a gradual decline in performance because it does not attempt to alleviate the overfitting problem in long-term adaptation. Even though CoTTA [11] shows consistent performance over time, their performance is inferior to ours. Figure 5 shows the problem that CoTTA fails to retain the acquired knowledge from cyclical sub-target domains. Specifically, although CoTTA adapted to the contrast domain (close to tunnel and underground) in the previous rounds, once the contrast domain re-encounters (i.e., $t^{cont.}$), the CoTTA model predicts unreliable results since the model may be strongly optimized for the last domain, brightness (close to daytime and sun glare). On the other hand, our approach makes stable predictions even in abruptly changing domains.

We believe that the adaptation knowledge across multiple sub-target domains is effectively stored and managed by our TTA framework for compound domain knowledge, allowing our method to outperform its competitors.

V. CONCLUSION

In this paper, we propose a simple yet effective TTA approach for dynamic domain changes since previous works suffer from performance degradation in this scenario which is likely to occur in the dynamic world. We first propose a TTA framework for managing compound domain knowledge via a continual domain-matching algorithm. Our framework helps the TTA model manage multiple domain knowledge by estimating the domain type of an input image among the compound domain. Moreover, in order to prevent overfitting, we modulate the adaptation rates utilizing domain-similarity between the current sub-target and the source domain. With extensive experiments, we verify the improved TTA performance of our approach in diverse dynamic TTA scenarios. In this regard, we hope that our efforts will contribute to developing a reliable technique for test-time adaptation and thus facilitating lifelong adaptation in robotics.

ACKNOWLEDGMENT

This work was supported in part by the Institute of Information and Communications Technology Planning and Evaluation (IITP) grant funded by the Korea Government (MSIT) (AI Innovation Hub) under Grant 2021-0-02068.

REFERENCES

- [1] Y.-H. Tsai, W.-C. Hung, S. Schuler, K. Sohn, M.-H. Yang, and M. Chandraker, "Learning to adapt structured output space for semantic segmentation," in *CVPR*, 2018.
- [2] I. Gulrajani and D. Lopez-Paz, "In search of lost domain generalization," in *ICLR*, 2021.
- [3] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," in *ICML*, 2015.
- [4] I. Shin, S. Woo, F. Pan, and I. Kweon, "Two-phase pseudo label densification for self-training based domain adaptation," in *ECCV*, 2020.
- [5] D. Wang, E. Shelhamer, S. Liu, B. Olshausen, and T. Darrell, "Tent: Fully test-time adaptation by entropy minimization," in *ICLR*, 2021.
- [6] S. Choi, S. Yang, S. Choi, and S. Yun, "Improving test-time adaptation via shift-agnostic weight regularization and nearest source prototypes," in *ECCV*, 2022.
- [7] D. Chen, D. Wang, T. Darrell, and S. Ebrahimi, "Contrastive test-time adaptation," in *CVPR*, 2022.
- [8] I. Shin, Y.-H. Tsai, B. Zhuang, S. Schuler, B. Liu, S. Garg, I. S. Kweon, and K.-J. Yoon, "Mm-tta: Multi-modal test-time adaptation for 3d semantic segmentation," in *CVPR*, 2022.
- [9] S. Niu, J. Wu, Y. Zhang, Z. Wen, Y. Chen, P. Zhao, and M. Tan, "Towards stable test-time adaptation in dynamic wild world," in *ICLR*, 2023.
- [10] H. Lim, B. Kim, J. Choo, and S. Choi, "TTN: A domain-shift aware batch normalization in test-time adaptation," in *ICLR*, 2023.
- [11] Q. Wang, O. Fink, L. Van Gool, and D. Dai, "Continual test-time domain adaptation," in *CVPR*, 2022.
- [12] J. Song, J. Lee, I. S. Kweon, and S. Choi, "Ecotta: Memory-efficient continual test-time adaptation via self-distilled regularization," in *CVPR*, 2023.
- [13] T. Kersiesis, J. Vanschoren, and M. Kılıçkaya, "Evaluating continual test-time adaptation for contextual and semantic domain shifts," *arXiv preprint arXiv:2208.08767*, 2022.
- [14] Y. Zou, Z. Yu, B. Kumar, and J. Wang, "Unsupervised domain adaptation for semantic segmentation via class-balanced self-training," in *ECCV*, 2018.
- [15] Z. Liu, Z. Miao, X. Pan, X. Zhan, D. Lin, S. X. Yu, and B. Gong, "Open compound domain adaptation," in *CVPR*, 2020.
- [16] K. Park, S. Woo, I. Shin, and I. S. Kweon, "Discover, hallucinate, and adapt: Open compound domain adaptation for semantic segmentation," in *NeurIPS*, 2020.
- [17] M. Döbler, R. A. Marsden, and B. Yang, "Robust mean teacher for continual and gradual test-time adaptation," in *CVPR*, 2023.
- [18] J. Lee, D. Das, J. Choo, and S. Choi, "Towards open-set test-time adaptation utilizing the wisdom of crowds in entropy minimization," in *ICCV*, 2023.
- [19] S. Park, S. Yang, J. Choo, and S. Yun, "Label shift adapter for test-time adaptation under covariate and label shifts," in *ICCV*, 2023.
- [20] F. You, J. Li, and Z. Zhao, "Test-time batch statistics calibration for covariate shift," *arXiv preprint arXiv:2110.04065*, 2021.
- [21] H. Jabbar and R. Z. Khan, "Methods to avoid over-fitting and under-fitting in supervised machine learning (comparative study)," *Computer Science, Communication and Instrumentation Devices*, 2015.
- [22] S. Niu, J. Wu, Y. Zhang, Y. Chen, S. Zheng, P. Zhao, and M. Tan, "Efficient test-time model adaptation without forgetting," in *ICML*, 2022.
- [23] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska et al., "Overcoming catastrophic forgetting in neural networks," *Proceedings of the national academy of sciences*, 2017.
- [24] W.-G. Chang, T. You, S. Seo, S. Kwak, and B. Han, "Domain-specific batch normalization for unsupervised domain adaptation," in *CVPR*, 2019.
- [25] S. Seo, Y. Suh, D. Kim, G. Kim, J. Han, and B. Han, "Learning to optimize domain specific normalization for domain generalization," in *ECCV*, 2020.
- [26] L. A. Gatys, A. S. Ecker, and M. Bethge, "Image style transfer using convolutional neural networks," in *CVPR*, 2016.
- [27] T. Matsuura and T. Harada, "Domain generalization using a mixture of multiple latent domains," in *AAAI*, 2020.
- [28] J. A. Hartigan and M. A. Wong, "Algorithm as 136: A k-means clustering algorithm," *Journal of the royal statistical society. series c (applied statistics)*, 1979.
- [29] T. Kailath, "The divergence and bhattacharyya distance measures in signal selection," *IEEE transactions on communication technology*, 1967.
- [30] Y. Liu, P. Kothari, B. van Delft, B. Bellot-Gurlet, T. Mordan, and A. Alahi, "Ttt++: When does self-supervised test-time training fail or thrive?" in *NeurIPS*, 2021.
- [31] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, and e. a. Lin, "Pytorch: An imperative style, high-performance deep learning library," in *NeurIPS*, 2019. [Online]. Available: <https://github.com/pytorch/vision/blob/main/torchvision/models/resnet.py>
- [32] J. Song, K. Park, I. Shin, S. Woo, C. Zhang, and I. S. Kweon, "Test-time adaptation in the dynamic world with compound domain knowledge management," *IEEE Robotics and Automation Letters*, pp. 7583–7590, 2023.
- [33] F. Croce, M. Andriushchenko, V. Sehwag, E. Debenedetti, N. Flammarion, M. Chiang, P. Mittal, and M. Hein, "Robustbench: a standardized adversarial robustness benchmark," in *NeurIPS Datasets and Benchmarks Track*, 2021. [Online]. Available: <https://openreview.net/forum?id=SSKZPJc7B>
- [34] D. Hendrycks, N. Mu, E. D. Cubuk, B. Zoph, J. Gilmer, and B. Lakshminarayanan, "AugMix: A simple data processing method to improve robustness and uncertainty," in *ICLR*, 2020.
- [35] S. R. Richter, V. Vineet, S. Roth, and V. Koltun, "Playing for data: Ground truth from computer games," in *ECCV*, 2016.
- [36] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *CVPR*, 2016.
- [37] S. Choi, S. Jung, H. Yun, J. T. Kim, S. Kim, and J. Choo, "Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening," in *CVPR*, 2021.