

Learning to Generate Gradients for Test-Time Adaptation via Test-Time Training Layers

Qi Deng^{*1}, Shuaicheng Niu^{*2}, Ronghao Zhang¹, Yaofu Chen¹,
Runhao Zeng^{3†}, Jian Chen^{1†}, Xiping Hu³

¹South China University of Technology,

²Nanyang Technological University,

³Artificial Intelligence Research Institute, Shenzhen MSU-BIT University

{dengqi.kei; niushuaicheng; zhangronghao16; runhaozeng.cs}@gmail.com; ellachen@scut.edu.cn;

Abstract

Test-time adaptation (TTA) aims to fine-tune a trained model online using unlabeled testing data to adapt to new environments or out-of-distribution data, demonstrating broad application potential in real-world scenarios. However, in this optimization process, unsupervised learning objectives like entropy minimization frequently encounter noisy learning signals. These signals produce unreliable gradients, which hinder the model’s ability to converge to an optimal solution quickly and introduce significant instability into the optimization process. In this paper, we seek to resolve these issues from the perspective of optimizer design. Unlike prior TTA using manually designed optimizers like SGD, we employ a learning-to-optimize approach to automatically learn an optimizer, called Meta Gradient Generator (MGG). Specifically, we aim for MGG to effectively utilize historical gradient information during the online optimization process to optimize the current model. To this end, in MGG, we design a lightweight and efficient sequence modeling layer – gradient memory layer. It exploits a self-supervised reconstruction loss to compress historical gradient information into network parameters, thereby enabling better memorization ability over a long-term adaptation process. We only need a small number of unlabeled samples to pre-train MGG, and then the trained MGG can be deployed to process unseen samples. Promising results on ImageNet-C/R/Sketch/A indicate that our method surpasses current state-of-the-art methods with fewer updates, less data, and significantly shorter adaptation times. Compared with a previous SOTA SAR, we achieve 7.4% accuracy improvement and $4.2\times$ faster adaptation speed on ImageNet-C.

Code — <https://github.com/keikeiqi/MGTTA>

1 Introduction

Since the emergence of test-time adaptation (TTA) (Sun et al. 2020; Wang et al. 2021), it has made significant progress (Liu et al. 2021; Zhang, Levine, and Finn 2022; Iwasawa and Matsuo 2021; Mirza et al. 2022; Boudiaf et al. 2022a; Gandelsman et al. 2022; Niu et al. 2024) and demonstrated broad potential across various scenarios to enhance

^{*}Equal Contribution

[†]Corresponding Author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

model performance on out-of-distribution (OOD) data or novel environments, also known as distribution shifts. By adapting to each test data immediately after inference in an unsupervised manner, TTA offers minimal overhead, distinguishing it as a practical choice for real-world applications.

Although the online unsupervised setting enhances TTA’s practicality, it also introduces certain challenges for TTA. This is because the model performance often degrades significantly on OOD data, leading to unsupervised objectives like entropy minimization (Wang et al. 2021) and self-learning (Goyal et al. 2022) encountering noisy supervision, which in turn produces unreliable gradients. Such unreliable gradients may hinder TTA from converging to an optimal solution quickly and also introduce instability into the learning process. This issue is especially critical in more complex test settings, such as label distribution shifts and mixed domain shifts, as highlighted by Niu et al. (2023).

To address this issue, various TTA methods (Niu et al. 2022, 2023; Yuan, Xie, and Li 2023; Lee et al. 2024) have been developed. For example, both EATA (Niu et al. 2022) and DeYO (Lee et al. 2024) devise different sample filtering strategies to select partial samples for TTA. However, the sample filtering strategy sometimes might be threshold-sensitive, making it difficult to set a reasonable threshold when test data are unknown. Moreover, filtering samples may result in the loss of valuable information, leading to insufficient learning, particularly when only a few test samples are available. In addition to sample filtering, EATA (Niu et al. 2022) also exploits weight regularization and SAR (Niu et al. 2023) devises a sharpness-aware minimization strategy to stabilize the online TTA process, etc.

Unlike prior methods, we seek to resolve the above issue from a new perspective of optimizer design. Instead of using manually designed optimizers like SGD and Adam, we cast the design of optimization algorithms as a learning problem, *i.e.*, learning to optimize (L2O) (Andrychowicz et al. 2016). Benefiting from the strong power of end-to-end learning, L2O has been extensively validated in supervised settings, even with noisy input gradients, demonstrating its ability to enhance model performance and accelerate convergence. Inspired by this, we devise an automatically learned optimizer for TTA, as shown in Figure 1. Considering that in online TTA, all unreliable gradients are immediately used for model updates after inference, their cumulative effect

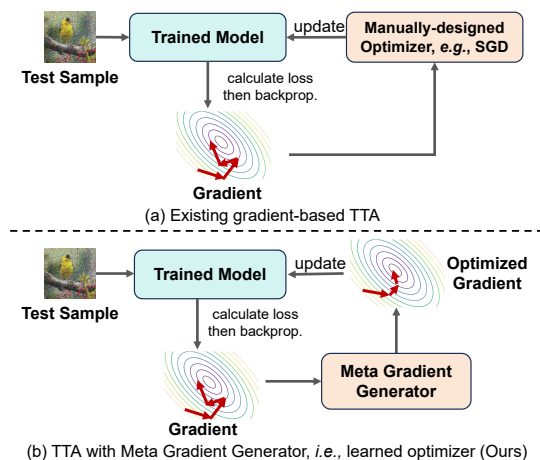


Figure 1: Method Differences. We devise an automatically learned meta gradient generator to optimize the original gradients produced by a TTA loss to be more reliable.

can amplify the impact of these unreliable gradients. However, noisy gradients are often short-term fluctuations, while the model optimization/gradients typically exhibit regularity over a longer time scale. In this paper, we suggest that the issue of unreliable gradients indeed can be alleviated a lot if we have access to all test sample gradients before adaptation, as leveraging the patterns in the historical gradient path would allow for a collective analysis to determine a more reliable gradient descent direction. This shares a similar motivation with SGDM or Adam that exploit historical gradients to improve the optimization. Thus, we propose memorizing long-term historical gradients during online adaptation to help and guide L2O in generating more reliable gradients.

Based on the above motivation, we devise a **Meta Gradient Generator (MGG)-guided Test-Time Adaptation** method, termed (**MGTTA**), in which the MGG is built upon L2O to replace manually designed optimizers used in TTA. MGG first memorizes the input original gradients and then outputs the refined gradients for TTA. For memorization, we devise a lightweight and efficient sequence modelling layer, termed the **Gradient Memory Layer (GML)**. GML is inspired by a most recent advanced architecture in the large language model community, termed **test-time training** layer (Sun et al. 2024). GML leverages a self-supervised reconstruction loss to encode historical gradient information into the model parameters, thereby enhancing GML’s capacity to retain all historical gradients throughout a long-term online adaptation process. For gradient optimization/correction, we utilize a feature discrepancy loss and a prediction entropy loss as our TTA objectives, guiding the MGG in automatically refining the original gradients to enhance their reliability through a learning-to-optimize manner. The training of MGG only requires a small number of unlabeled OOD samples (*e.g.*, 128), and then the trained MGG can be deployed to unseen samples for TTA. Extensive results indicate that our method surpasses existing SOTAs with fewer updates, fewer data, and significantly shorter adaptation times.

Main Novelty and Contributions. 1) We devise a novel

Meta Gradient Generator (MGG), which is automatically learned in a learning-to-optimize manner, to replace manually designed optimizers for TTA. This generator takes the original unreliable gradients as input and produces optimized gradients, thereby alleviating the impact of noisy learning signals encountered during TTA. 2) We introduce a lightweight yet efficient sequential modeling network, namely Gradient Memory Layer, which can memorize the historical gradient information during a long-term online TTA process by encoding this information into model parameters via a reconstruction loss. Then, we use the memorized gradients to guide the optimization/correction of gradients for the current model adaptation. 3) Extensive experiments demonstrate that using a small number (*e.g.*, 128) of unlabeled samples from the ImageNet-C validation set is sufficient to train an effective MGG. With this pre-trained MGG, our method outperforms existing methods on various unseen datasets, including ImageNet-C/R/Sketch/A. The fast convergence observed in experiments makes our method practical for real-world scenarios especially when the computational resource budget is limited, as in Table 2.

2 Related Work

Test-Time Adaptation (TTA) aims to adapt a trained model to new environments or OOD data online using unlabeled test data and it has made considerable progress (Nado et al. 2020; Khurana et al. 2021; Boudiaf et al. 2022b; Zeng et al. 2023; Chen et al. 2024a,b; Wen et al. 2023). According to whether relying on gradient computation, existing TTA methods can be categorized into gradient-free methods, such as LAME (Boudiaf et al. 2022a) and T3A (Iwasawa and Matsuo 2021), and gradient-based methods, such as TENT (Wang et al. 2021), EATA (Tan et al. 2024), etc. By directly updating the model parameters using some unsupervised loss, the later one often achieves much better performance. However, in this unsupervised learning, the objectives may generate unreliable gradients due to the interference of noisy signals, leading to unstable adaptation or sub-optimal performance. To address this, several methods have been proposed. For example, CoTTA (Wang et al. 2022) mitigates error accumulation by utilizing weight-averaged and augmentation-averaged pseudo-labels. SAR (Niu et al. 2023) introduces a sharpness-aware entropy minimization, while both SAR and DeYO (Lee et al. 2024) devise sample filtering strategies to exclude certain samples from adaptation. However, these methods may sometimes be threshold-sensitive or require substantial additional computational resources. Unlike existing methods, in this paper, we explore a new perspective for resolving the unreliable gradients issue, *i.e.*, exploiting a learning-to-optimize framework to automatically learn a gradient optimizer for TTA, thereby enhancing the quality of gradients used for TTA.

Learning to Optimize (L2O) aims to automatically learn optimizers (Li and Malik 2016), improving on traditional methods like Bayesian optimization, random search, and gradient-based approaches. Andrychowicz et al. (2016) treated optimization as a learning problem, and subsequent methods extended it by introducing techniques such as task-independent optimization (Li and Malik 2017), hierarchical

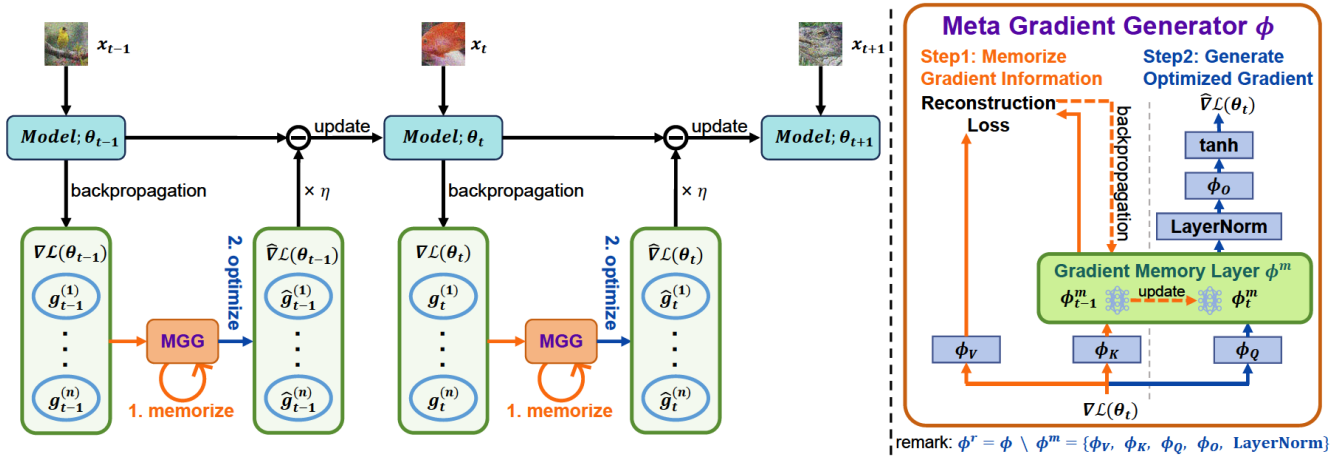


Figure 2: An overall illustration of MGTTA, in which we design a two-step meta gradient generator (MGG) to generate optimized gradients for TTA. Given a trained model $f(\cdot; \theta)$, for each batch of test samples, we first calculate predictions and obtain gradients by backpropagation. Then, in Step 1 MGG first memorizes gradients and then in Step 2 MGG generates optimized gradients based on the historical gradient information. Finally, the model parameters θ are updated using the optimized gradients. Here, the learnable parameters within θ only involve norm layers and the rest are kept frozen during adaptation.

RNN architectures (Wichrowska et al. 2017), and random scaling to speed up training (Lv, Jiang, and Li 2017). To address gradient truncation, solutions like dynamic weighting (Metz et al. 2019) and progressive unroll length (Chen et al. 2020) were proposed. HALO (Li et al. 2020) enhanced generalization through a novel regularizer, while PES (Vicol, Metz, and Sohl-Dickstein 2021) removed truncation bias using accumulated correction terms. SL2O (Chen et al. 2022) focused on learning within a restricted subspace. The above methods show promise in speeding up model convergence and highlight the potential of L2O approaches. However, they rely on ground truth for training, which is often unavailable in real-world scenarios. In this paper, we explore the concept of L2O to optimize gradients during the TTA process and design a self-supervised training method that enables efficient L2O without the need for labeled data.

3 Proposed Method

3.1 Problem Statement

Consider a set of source training images $\{x_n\}_{n=1}^N$, where each image x_n is drawn from the distribution $P(x)$. A model $f(\cdot; \theta)$, parameterized by θ , is trained using the labeled dataset $\{(x_n, y_n)\}_{n=1}^N$. Ideally, $f(\cdot; \theta)$ performs effectively on test samples originating from the same distribution as the training data, denoted by $x \sim P(x)$. However, in practice, this condition is frequently unmet. The test data may often be out of distribution, possibly corrupted, represented by $x \sim U(x)$, where $U(x) \neq P(x)$. Such discrepancies can significantly impair the model’s performance.

Test-time adaptation (TTA) attempts to mitigate this issue by adapting the model using only unlabeled test samples. During the TTA phase, the model $f(\cdot; \theta)$ often employs an unsupervised loss $\mathcal{L}(\theta)$ to fine-tune its parameters, thus facilitating adaptation to out-of-distribution (OOD) data or new environments. Traditional TTA methods typically use

manually designed optimizers like stochastic gradient descent (SGD) to update the parameters at the step $t+1$ by

$$\theta_{t+1} = \theta_t - \eta \cdot \nabla \mathcal{L}(\theta_t), \quad (1)$$

where η denotes the learning rate. However, such unsupervised learning objectives may produce unreliable gradients due to noise interference, resulting in unstable optimization and challenges in swiftly converging to an optimal solution.

3.2 General Scheme

To address the above issues, we propose a learning-to-learn framework to develop a novel optimizer tailored for TTA scenarios. We learn to generate gradients by developing a neural network-based optimizer termed the **meta gradient generator** $f_{\text{MGG}}(\cdot; \phi)$, where ϕ is learnable parameters. This optimizer refines the current gradient $\nabla \mathcal{L}(\theta_t)$ to yield more precise and stable gradients and updating parameters by

$$\theta_{t+1} = \theta_t - \eta \cdot f_{\text{MGG}}(\nabla \mathcal{L}(\theta_t); \phi). \quad (2)$$

Here, the key challenge is how can we design MGG to make it optimize the $\nabla \mathcal{L}(\theta_t)$ to be more reliable. In this paper, we posit that unreliable gradients primarily arise due to the online nature of TTA. Since TTA processes each test sample only once for immediate adaptation post-inference, any unreliable gradients shall be incorporated into the model adaptation process, thereby degrading performance. Ideally, if we could access all test samples’ gradients before adaptation, we could analyze them collectively and determine a more reliable direction for gradient descent, thus mitigating the impact of unreliable gradients. However, this approach conflicts with the online nature of TTA, potentially converting it into an offline process. Instead, another straightforward solution is storing gradients from previous test samples during online learning, which, however, will result in much higher memory consumption. To address this, we propose to use

MGG to memorize historical gradients during the online adaptation process, and then leverage this historical information to help MGG generate more reliable gradients.

To be specific, we propose a compact yet effective sequential modeling scheme, termed as **gradient memory layer** $f_{\text{GML}}(\cdot; \phi^m)$, which serves as a pivotal component within our MGG ($\phi^m \in \phi$). Inspired by Test-Time Training layers (Sun et al. 2024), our design idea for $f_{\text{GML}}(\cdot)$ is “parameter as memory”—compresses the continuous gradient update information into the neural network parameters. Without loss of generality, consider the t -th step of TTA, our method operates in two steps: **1) memorize**, input the gradient $\nabla \mathcal{L}(\theta_t)$, directly computed from the loss, into f_{GML} , and updating the parameters ϕ^m of $f_{\text{GML}}(\cdot; \phi^m)$ via self-supervised learning to encode the current information. **2) optimize**, use updated ϕ^m to optimize the input gradient, ultimately applying Eqn. (2) for TTA. The schematic depiction of our approach is shown in Figure 2. We term our method as **Meta Gradient Generator-guided Test-Time Adaptation (MGTTA)**. Below we first introduce the details of MGG and then the pre-training and TTA pipelines of our MGTTA.

3.3 Meta Gradient Generator (MGG)

Unlike prior gradient-based TTA methods that utilize manually designed optimizers, *e.g.*, SGD, we seek to update the model at test time using an auto-learned optimizer, termed the meta gradient generator (MGG). MGG takes the gradients computed from the unsupervised loss as input and outputs optimized gradients for more stable parameter updates.

Gradient Memory Layer TTA typically operates online and processes sequential samples, in which each sample (with its gradients) is discarded immediately after adaptation. This process is highly sensitive to unreliable gradients since in this way all unreliable signals will be accumulated and finally degrade model performance a lot. Therefore, we envision the MGG being capable of memorizing historical gradient information. In this sense, MGG shall have the potential to understand the previous and current gradients together, and thus determine a more reliable gradient descent direction, alleviating the impacts of unreliable signals.

To achieve the above goal, one can employ Long Short-Term Memory (LSTM) networks (Hochreiter and Schmidhuber 1997) to store the history of time series in a hidden state vector. However, this method is often constrained by the expressive power of a single vector when managing long sequences. Inspired by the Test-Time Training layer (Sun et al. 2024) that most recently emerged in the large language model community, we introduce a gradient memory layer $f_{\text{GML}}(\cdot; \phi^m)$, to compress the sequence information within network parameters, which has much stronger expression power as model parameters have larger capacity than the hidden states of LSTM. *i.e.*, higher dimension. We design a two-step learning approach to implement MGG.

Step 1: Memorize Gradient Information The central concept in incorporating gradient information compression into f_{GML} is to utilize a reconstruction loss as the self-supervision to update ϕ^m . This approach is similar to how

language models often employ reconstruction or mask prediction loss to compress knowledge from the training corpus into neural network parameters through gradient descent. During test time, GML is continuously updated by learning from test data, allowing it to “memorize” and leverage historical information in future TTA steps. Specifically, suppose the input gradient at time t is $g_t = \nabla \mathcal{L}(\theta_t)$, the reconstruction loss can be expressed as

$$\mathcal{L}_{\text{GML}}(g_t; \phi^m) = \| f_{\text{GML}}(\phi_K g_t; \phi^m) - \phi_V g_t \|^2, \quad (3)$$

where ϕ_K and ϕ_V are two learnable projection matrices that upscale g_t , similar but different from Test-Time Training (Sun et al. 2024) in which the projection matrices are low-rank. We then update ϕ^m by

$$\phi^m \leftarrow \phi^m - \eta_{\text{GML}} \cdot \nabla \mathcal{L}_{\text{GML}}(g_t; \phi^m), \quad (4)$$

where η_{GML} is the learning rate. Here, since the learning rate is crucial in gradient descent, we use an adaptive learning rate that is determined by a learnable vector ϕ_{lr} , denoted by

$$\eta_{\text{GML}} = \sigma(\phi_{lr} \cdot g_t), \quad (5)$$

where σ is the sigmoid function. Up to this time point, ϕ^m contains gradient information from before and at time step t , achieving the memorization of current gradient information in the network parameters.

Step 2: Generate Optimized Gradient With the updated memory that encapsulates both current and historical gradient information, we can commence the optimization of gradients. The gradient g_t in need of optimization is first processed through a projection layer ϕ_Q (akin to ϕ_K and ϕ_V), subsequently fed into the updated function $f_{\text{GML}}(\cdot; \phi^m)$, followed by a layer normalization operation and another projection layer ϕ_O . The process of obtaining the optimized gradient can be represented as

$$\hat{g}_t = f_{\text{MGG}}(g_t; \phi) = \tanh(\phi_O \cdot \text{LN}(f_{\text{GML}}(\phi_Q g_t; \phi^m))), \quad (6)$$

where LN denotes a LayerNorm layer. Then, this gradient \hat{g}_t can be applied to the TTA of our target model (Eqn. (2)).

3.4 Learning MGG before Test-Time Adaptation

Considering that MGG divides the gradient optimization process into two steps, we accordingly design a multi-step iterative update strategy for the MGG parameters. Without loss of generality, let ϕ be the parameters of MGG, by excluding those of $f_{\text{GML}}(\cdot; \phi^m)$, we denote the **remaining parameters** as $\phi^r = \phi \setminus \phi^m$. In other words, all these parameters ϕ_Q, ϕ_K, ϕ_V and ϕ_O belong to ϕ^r (as shown in the right part of Figure (2)). **First**, input the data x_t into the model $f(\cdot; \theta)$ to obtain predictions \hat{y} (*e.g.*, classification results). Then, use the loss function $\mathcal{L}_{\text{TTA}}(\theta)$ to perform backpropagation to obtain gradients *w.r.t.* ϕ^r and θ . **Second**, update ϕ^r using the corresponding gradients. **Third**, fix the parameters ϕ^r and use Eqn. (3) to update the f_{GML} parameters ϕ^m . **Lastly**, with the parameters ϕ^m and ϕ^r fixed, utilize MGG to obtain the optimized gradient and employ Eqn. (1) to update the parameters θ of the model requiring TTA.

For the TTA loss, we select FOA (Niu et al. 2024) as our test-time learning objective as FOA loss is one of the most

recent advanced objectives and has shown promising performance in the existing TTA literature. Beyond entropy minimization (Wang et al. 2021), FOA introduces a feature statistics discrepancy loss. To be specific, FOA first collects a small set of unlabelled in-distribution samples (e.g., 64 samples) to calculate the source feature mean $\{\mu_i^S\}_{i=1}^N$ and variances $\{\sigma_i^S\}_{i=1}^N$ of each layer, where i denotes layer index. During TTA, FOA also calculates the corresponding statistics $\{\mu_i(\mathcal{X}_t)\}_{i=1}^N$ and $\{\sigma_i(\mathcal{X}_t)\}_{i=1}^N$ of testing samples and aligns them with pre-calculated source statistics. Formally, given a batch of test samples \mathcal{X}_t , the loss is defined by

$$\mathcal{L}_{TTA}(f(\mathcal{X}_t; \theta)) = \sum_{x \in \mathcal{X}} \sum_{c \in \mathcal{C}} -\hat{y}_c \log \hat{y}_c + \lambda \sum_{i=1}^N (\|\mu_i(\mathcal{X}_t) - \mu_i^S\|^2 + \|\sigma_i(\mathcal{X}_t) - \sigma_i^S\|^2), \quad (7)$$

where C denotes the number of categories, \hat{y}_c is the prediction for category c , and λ is a trade-off parameter. Note that the entire training process of MGG only requires a few number of unlabeled test samples. In our main experiments, we random sample 128 images without labels from the held-out validation set of ImageNet-C as the training set of MGG, and then test the trained MGG on all ImageNet-C testing datasets and other ImageNet variants. Experiments in Table 6 demonstrate that 128 images are sufficient for our method to achieve excellent performance.

3.5 Test-Time Adaptation with MGG

With a trained MGG, we can begin to employ it for TTA. TTA is an online optimization process with limited available resources. To save computational and storage overhead while efficiently adapting the model, we update only the parameters of the model’s normalization layers (e.g., 0.044% of parameters in ViT-Base (Dosovitskiy et al. 2021)).

Parameter-wise Memorizing To enhance the capability of MGG in modeling the temporal variations of gradients, we treat each parameter independently rather than as a whole. Specifically, within TTA, each parameter θ^i within the model parameters $\theta \in \mathbb{R}^n$ evolves over time, forming a sequence $\{\theta_0^i, \theta_1^i, \dots, \theta_t^i\}$. For each parameter θ^i , we employ an independent ϕ^m to model the corresponding sequence of gradient information, while the other parameters of the MGG, i.e., ϕ^r , are shared across all parameters of θ . This approach allows the GML to concentrate on the temporal changes of each individual parameter rather than the interrelationships among them. Although maintaining historical information for each parameter incurs certain memory costs, the GML is a compact neural network with a small number of parameters, resulting in minimal additional memory overhead during TTA. Taking ViT-base as an example ($n = 38,400$), in our implementation, ϕ^m for each θ^i is a linear layer with input and output dimensions of 8, making the total parameters of all GML summing up to only 2.76 M.

Deploying Pre-trained MGG for TTA After a small set of held-out samples are used for training, the MGG can be directly employed for TTA. Initially, input the data x into

Algorithm 1: The pre-training/TTA pipeline of MGTTA.

// Pre-training&TTA share the same pipeline but differ from:
For **pre-training**, we use \mathcal{D}_{val} and random initialized $\phi^r \in \phi$.
For **TTA**, we use \mathcal{D}_{test} and ϕ^r is inherited from pre-training.

Require: Trained model $f(\cdot; \theta)$, MGG model $f_{MGG}(\cdot; \phi)$, samples $\mathcal{D} = \{x_j\}_{j=1}^M$, hyper-parameters T and η .

- 1: Random initialize GML’s parameters as ϕ_1^m
- 2: Calculate predictions on batch \mathcal{X}_1 from \mathcal{D} via $f(\cdot; \theta)$
- 3: Calculate the loss $\mathcal{L}_{TTA}(f(\mathcal{X}_1; \theta))$ in Eqn. (7)
- 4: Calculate the gradient $g_1 = \nabla \mathcal{L}_{TTA}(\theta)$
- 5: **for** t in $[2, 3, \dots, T]$ **do**
- 6: // **Step1: Memorize Gradient Information**
- 7: Calculate the loss $\mathcal{L}_{GML}(g_{t-1}; \phi_{t-1}^m)$ in Eqn. (3)
- 8: Obtain ϕ_t^m with $\nabla \mathcal{L}_{GML}(\phi_{t-1}^m)$ via Eqn. (4)
- 9: // **Step2: Generate Optimized Gradient for TTA**
- 10: Optimize the gradient $\hat{g}_{t-1} = f_{MGG}(g_{t-1}; \phi)$
- 11: Update θ with optimized gradient \hat{g}_{t-1} via Eqn. (2)
- 12: Calculate predictions on batch \mathcal{X}_t from \mathcal{D} via $f(\cdot; \theta)$
- 13: Calculate the loss $\mathcal{L}_{TTA}(f(\mathcal{X}_t; \theta))$ in Eqn. (7)
- 14: Calculate gradients $\nabla \mathcal{L}_{TTA}(\phi^r)$ and $g_t = \nabla \mathcal{L}_{TTA}(\theta)$
- 15: Update ϕ^r via $\nabla \mathcal{L}_{TTA}(\phi^r)$
- 16: **end for**

Ensure: The trained MGG $f_{MGG}(\cdot; \phi)$ for **pre-training or**
The predictions of all samples in \mathcal{D} for **TTA**.

the model $f(\cdot; \theta)$ and compute the initial gradient $\mathcal{L}_{TTA}(\theta)$. Subsequently, input g into the MGG to perform one memorization and one optimization step, then use the optimized gradient to update the model $f(\cdot; \theta)$. Note that during TTA, one can choose whether to update ϕ^r . However, our preliminary studies indicate that this does not clearly improve performance. Therefore, for higher efficiency, we keep ϕ^r frozen during the whole TTA phase. We summarize the overall pseudo-code of our method in algorithm 1.

4 Experiments

4.1 Datasets, Models and Compared Methods

We conduct experiments on four benchmark datasets, including 1) **ImageNet-C** (Hendrycks and Dietterich 2019) contains corrupted images in 15 types of 4 main categories and each type has 5 severity levels. **In ours experiments, all results are evaluated on the severity level 5.** 2) **ImageNet-R** (Hendrycks et al. 2021a) contains various artistic renditions of 200 ImageNet classes. 3) **ImageNet-Sketch** (Wang et al. 2019) includes sketch-style images representing 1,000 ImageNet classes. 4) **ImageNet-A** (Hendrycks et al. 2021b) consists of natural adversarial examples. We use ViT-Base (Dosovitskiy et al. 2021) as the source model, which is trained on ImageNet-1K and adapted to the above datasets.

Baselines: LAME (Boudiaf et al. 2022a), T3A (Iwasawa and Matsuo 2021), TENT (Wang et al. 2021), CoTTA (Wang et al. 2022), SAR (Niu et al. 2023), FOA (Niu et al. 2024), EATA (Niu et al. 2022), DeYO (Lee et al. 2024).

4.2 Implementation Details

For pre-training MGG, we randomly select 128 unlabeled samples from the ImageNet-C validation set. The learning

Method	Noise			Blur				Weather				Digital				Avg.
	Gauss.	Shot	Impul.	Defoc.	Glass	Motion	Zoom	Snow	Frost	Fog	Brit.	Contr.	Elastic	Pixel	JPEG	
NoAdapt	56.8	56.8	57.5	46.9	35.6	53.1	44.8	62.2	62.5	65.7	77.7	32.6	46.0	67.0	67.6	55.5
TENT	60.3	61.6	61.8	59.2	56.5	63.5	59.2	54.3	64.5	2.3	79.1	67.4	61.5	72.5	70.6	59.6
CoTTA	63.6	63.8	64.1	55.5	51.1	63.6	55.5	70.0	69.4	71.5	78.5	9.7	64.5	73.4	71.2	61.7
SAR	59.2	60.5	60.7	57.5	55.6	61.8	57.6	65.9	63.5	69.1	78.7	45.7	62.4	71.9	70.3	62.7
FOA	61.5	63.2	63.3	59.3	56.7	61.4	57.7	69.4	69.6	73.4	81.1	67.7	62.7	73.9	73.0	66.3
EATA	61.2	62.3	62.7	60.0	59.2	64.7	61.7	69.0	66.6	71.8	79.7	66.8	65.0	74.2	72.3	66.5
DeYO	62.4	64.0	63.9	61.0	60.7	66.4	62.9	70.9	69.6	73.7	80.5	67.2	69.9	75.7	73.7	68.2
MGTTA (ours)	64.5	66.5	66.3	63.8	65.0	70.1	69.7	74.5	72.8	77.0	81.3	71.0	75.0	77.7	75.1	71.3

Table 1: Comparisons with state-of-the-art methods on ImageNet-C *w.r.t.* accuracy(%).

Method	Time Budget for Adaptation (seconds)								# Data Budget for Adaptation (number of batch)							
	2.0	5.0	10.0	20.0	30.0	50.0	90.0	∞	10	20	35	50	75	100	200	782
NoAdapt	55.5	55.5	55.5	55.5	55.5	55.5	55.5	55.5	55.5	55.5	55.5	55.5	55.5	55.5	55.5	55.5
TENT	56.8	55.4	56.3	57.8	58.7	59.4	59.8	59.6	56.7	55.7	55.4	56.0	56.7	57.4	58.8	59.6
CoTTA	45.8	37.5	37.1	36.3	37.7	38.3	39.5	61.7	36.6	36.8	38.4	37.3	38.1	39.6	42.6	61.7
SAR	55.9	56.6	56.4	55.8	56.4	61.3	58.9	62.7	56.9	57.0	52.1	55.8	56.6	57.0	62.1	62.7
FOA	54.3	55.3	56.1	56.7	57.1	57.7	59.1	66.3	56.9	57.3	58.3	59.5	61.1	62.1	64.2	66.3
EATA	57.0	58.8	61.5	63.9	64.8	65.8	66.4	66.5	56.7	58.1	59.6	60.6	62.1	63.1	65.0	66.5
DeYO	57.0	59.2	62.4	64.8	66.0	67.1	68.0	68.2	57.1	59.1	60.9	61.9	63.8	64.9	66.9	68.2
MGTTA (ours)	64.1	68.7	70.4	70.9	71.1	71.3	71.3	71.3	63.2	66.4	69.3	70.1	70.6	70.8	71.2	71.3

Table 2: Comparisons under limited adaptation budgets on ImageNet-C *w.r.t.* acc(%). Total #batches is 782, with batch size 64.

Method	NoAdapt	TENT	SAR	FOA	EATA	DeYO	MGTTA
R	59.5	63.9	63.3	63.8	63.3	66.1	70.2
Sketch	44.9	49.1	48.7	49.9	50.9	52.2	53.3
A	0.1	52.9	52.5	51.5	53.4	54.1	56.7
Avg.	34.8	55.3	54.8	55.1	55.9	57.5	60.1

Table 3: Comparisons with state-of-the-arts on ImageNet-R, ImageNet-Sketch and ImageNet-A *w.r.t.* accuracy(%).

rate is set to $1e-4$ for θ and $1e-2$ for ϕ^r . We update θ and ϕ^r for $T=2,000$ iterations with a batch size of 2. The GML hidden size is set to 8. **During TTA**, the batch size is 64, and ϕ^r is fixed, the learning rate for θ is set to $1e-3$.

4.3 Comparisons with State-of-the-arts

TTA Results on ImageNet-C. We report the results of 15 different corruptions on the ImageNet-C dataset (severity level 5) in Table 1. The experimental results indicate that our method significantly outperforms other approaches across all corruptions. Compared to T3A, TENT, and CoTTA, our method achieves an average performance improvement of approximately 10% on ImageNet-C. When compared to the previous best-performing method, DeYO, our method still shows an improvement of 3.1% (68.2% vs. 71.3%), demonstrating the effectiveness of our approach.

Performance under Limited Adaptation Budget on ImageNet-C. In practical applications, the number of samples available for adaptation and the time allotted are often limited. We conduct two experiments: **1) limited update time:** set a maximum update time t_{max} , beyond which the model is frozen and subsequent data are only used for inference. Results in Table 2 indicate that our method significantly enhances performance with only limited time used for adaptive updates, achieving a 68.7% accuracy in just 5 seconds, surpassing other methods over their entire update

Variants of our method	C	R	Sketch	A
Ours	71.3	70.2	53.3	56.7
Ours <i>replace GML with LSTM</i>	71.3	69.4	50.3	56.8
Ours <i>remove MGG</i>	70.0	67.2	51.9	55.1

Table 4: Effect of MGG and GML on ImageNet-C/R/S/A.

duration. **2) limited available samples:** perform TTA using only the first u_{max} batches of samples. From Table 2, our method, adapting with only 35 out of 782 batches (4.5%), achieves a 69.3% accuracy, already significantly surpassing other methods. In contrast, methods like SAR, EATA, and DeYO, which involve filtering samples and do not utilize the entire data within a batch, show lower efficiency. MGTTA does not filter samples but instead optimizes the gradient using MGG, offering superior effectiveness and efficiency.

TTA Results on ImageNet-R/Sketch/A. We further explore the generalization ability of the MGG trained on ImageNet-C validation set to datasets with totally different distribution shift types, including ImageNet-R/Sketch/A datasets. The results, shown in Table 3, indicate that our method outperforms others on these three challenging ImageNet variant datasets by a large margin. Compared to the previous best-performing method, our method achieves an average performance improvement of 2.6% (57.5% vs. 60.1%). This highlights the strong generalization capability of MGG, making it effective on various OOD datasets.

4.4 Ablation Studies

Effectiveness of MGG and GML. We conduct experiments by removing MGG or replacing GML with LSTM. From Table 4, without MGG’s gradient optimization, accuracy significantly decreases on ImageNet-C/R/Sketch/A, indicating that MGG generates more reliable gradients. On ImageNet-C/A, both our GML and LSTM achieve bet-

	NoAdapt	LAME	T3A	TENT	CoTTA	SAR	FOA	EATA	DeYO	Ours ($u_{max}=20$)	Ours ($u_{max}=50$)	Ours ($u_{max}=100$)	Ours
Acc.(%)	55.5	54.1	56.9	59.6	61.7	62.7	66.3	66.5	68.2	66.4	70.1	70.8	71.3
Runtime(s)	54.3	54.8	125.5	122.6	619.3	242.7	1636.7	127.4	172.4	55.8	58.5	63.2	125.5

Table 5: Wall-clock runtime for processing 50,000 images of ImageNet-C on a RTX 4090 GPU, and Acc. averaged over 15 corruptions. u_{max} : adapt the model with only the first u_{max} batches. If u_{max} is not specified, all samples are used for adaptation.

Dataset	$N=64$	$N=128$	$N=256$	$N=1000$
ImageNet-C	71.3	71.3	71.2	71.3
ImageNet-R	70.0	70.2	70.1	70.1

Table 6: Effect on numbers of training samples for MGG *w.r.t.* average accuracy over 15 corruptions on ImageNet-C.

	NoAdapt	TENT	CoTTA	SAR	FOA	EATA	DeYO	Ours
Acc.	53.8	52.1	46.3	60.7	61.7	64.6	64.0	65.9

Table 7: Comparisons *w.r.t.* accuracy on ImageNet-C under mixed domain shifts, *i.e.*, the mixture of 15 corruptions.

Batch size	TENT	SAR	FOA	EATA	DeYO	Ours
2	57.4	64.3	66.5	67.0	63.6	70.3
4	59.2	63.8	66.6	67.2	65.5	70.9

Table 8: Comparisons with small batch size on ImageNet-C *w.r.t.* average accuracy(%) over 15 corruptions.

ter performance than without MGG, demonstrating that our idea of enhancing TTA performance through MGG-optimized gradient is general and effective. Notably, when using LSTM, the performance on ImageNet-R/Sketch is lower than GML (50.3% vs. 53.3%), possibly because MGG was trained on ImageNet-C val data, which visually differs significantly from ImageNet-R/Sketch. Our GML exhibits better generalization performance compared to LSTM.

Convergence Comparison. We perform TTA with a fixed number of iterations, followed by accuracy evaluations on remaining samples without additional adaptation. Figure 3 shows that MGTTA achieves the best convergence and final accuracy, with a steep rise in accuracy within the initial 0-50 updates, suggesting the effectiveness of MGG-optimized gradients. Here, MGTTA without MGG also achieves better performance than baselines, which mainly benefits from the advanced loss in Eqn. (7). Even though, our method with MGG still achieves much better convergence than that without MGG, showing the superiority of MGG.

Effect of the Number of Samples for Pre-training MGG. Considering the challenge of data acquisition in practical applications, we explore the feasibility of training MGG with as minimal data as possible. We train MGG with N images randomly sampled from ImageNet-C validation set. From Table 6, using only 128 samples suffices to achieve optimal accuracies of 71.3%/70.2% on ImageNet-C/-R, suggesting our method is able to learn optimizing gradients efficiently.

Results under Wild Scenarios. We consider two wild scenarios on ImageNet-C following (Niu et al. 2023): **1)** small TTA batch size and **2)** mixed distribution shift: mix

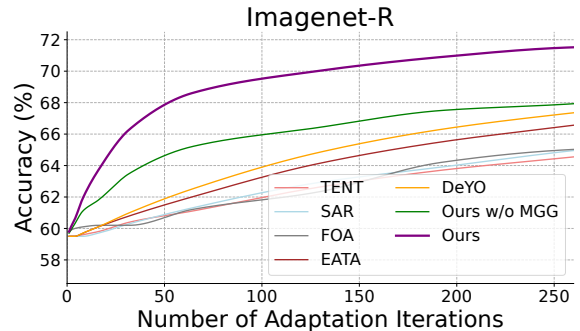


Figure 3: Convergence speed comparisons on ImageNet-R.

and randomly shuffle the 15 corruption types. Results in Tables 7 and 8 show that our method consistently outperforms other methods. Reducing the batch size has a minimal impact on the performance of our method (compared to other methods), indicating that our approach exhibits superior and more stable performance under challenging conditions.

Runtime and Memory Comparison. We report the wall-clock time of MGTTA and its variants in Table 5, in which the MGTTA variants adapt the model using only the first u_{max} out of 782 batches (with batch size 64), and then the model is frozen for the subsequent inference. From the results, MGTTA ($u_{max} = 50$) surpasses all baselines (adapt on all 782 batches) by adapting the model on only 50 batches, suggesting our effectiveness. Moreover, MGTTA requires 5,193 MB of GPU memory for TTA, whereas without MGG (*i.e.*, directly updating model parameters via back-propagated gradients) requires 5,165 MB (TENT). Thus, our method consumes only an additional 28 MB of GPU memory, demonstrating that MGG is lightweight and efficient.

5 Conclusion

In this paper, we have proposed a Meta Gradient Generator-guided TTA (MGTTA) method, to alleviate the issues of unreliable gradients encountered during the unsupervised online TTA process. Specifically, we developed a novel neural network-based optimizer to replace manually tuned optimizers. This optimizer has been capable of optimizing gradients, thereby enhancing their reliability and stability. Recognizing that historical gradient information can guide the generation of current gradients, we designed a lightweight and efficient sequence modeling layer, termed gradient memory layer, to store historical information and provide a better direction for gradient optimization. We trained this optimizer using a learning-to-optimize approach with a TTA loss of feature discrepancy and prediction entropy. The superior performance and fast convergence observed in extensive experiments facilitate our application in the real world.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (NSFC) (Grant Nos. 62202311, 62376099), Excellent Science and Technology Creative Talent Training Program of Shenzhen Municipality (Grant No. RCBS20221008093224017), Guangdong Basic and Applied Basic Research Foundation (Grant No. 2023A1515011512), Key Scientific Research Project of the Department of Education of Guangdong Province (Grant No. 2024ZDZX3012), Natural Science Foundation of Guangdong Province (Grant No. 2024A1515010989), and Ministry of Education, Singapore, under its Academic Research Fund Tier 1.

References

- Andrychowicz, M.; Denil, M.; Gomez, S.; Hoffman, M. W.; Pfau, D.; Schaul, T.; Shillingford, B.; and De Freitas, N. 2016. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems*, 3981–3989.
- Boudiaf, M.; Mueller, R.; Ben Ayed, I.; and Bertinetto, L. 2022a. Parameter-free online test-time adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 8344–8353.
- Boudiaf, M.; Mueller, R.; Ben Ayed, I.; and Bertinetto, L. 2022b. Parameter-free online test-time adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 8344–8353.
- Chen, G.; Niu, S.; Chen, D.; Zhang, S.; Li, C.; Li, Y.; and Tan, M. 2024a. Cross-Device Collaborative Test-Time Adaptation. In *NeurIPS*.
- Chen, T.; Zhang, W.; Jingyang, Z.; Chang, S.; Liu, S.; Amini, L.; and Wang, Z. 2020. Training stronger baselines for learning to optimize. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS*.
- Chen, X.; Chen, T.; Cheng, Y.; Chen, W.; Awadallah, A.; and Wang, Z. 2022. Scalable learning to optimize: A learned optimizer can train big models. In *Computer Vision - ECCV 2022 - 17th European Conference*, 389–405. Springer.
- Chen, Y.; Niu, S.; Wang, Y.; Xu, S.; Song, H.; and Tan, M. 2024b. Towards Robust and Efficient Cloud-Edge Elastic Model Adaptation via Selective Entropy Distillation. In *ICLR*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2021. An image is worth 16x16 words: Transformers for image recognition at scale.
- Gandelsman, Y.; Sun, Y.; Chen, X.; and Efros, A. 2022. Test-time training with masked autoencoders. In *Advances in Neural Information Processing Systems (NeurIPS)*, 29374–29385.
- Goyal, S.; Sun, M.; Raghunathan, A.; and Kolter, J. Z. 2022. Test time adaptation via conjugate pseudo-labels. In *Advances in Neural Information Processing Systems*, volume 35, 6204–6218.
- Hendrycks, D.; Basart, S.; Mu, N.; Kadavath, S.; Wang, F.; Dorundo, E.; Desai, R.; Zhu, T.; Parajuli, S.; Guo, M.; et al. 2021a. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8340–8349.
- Hendrycks, D.; and Dietterich, T. 2019. Benchmarking neural network robustness to common corruptions and perturbations. In *7th International Conference on Learning Representations, ICLR*.
- Hendrycks, D.; Zhao, K.; Basart, S.; Steinhardt, J.; and Song, D. 2021b. Natural adversarial examples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15262–15271.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Iwasawa, Y.; and Matsuo, Y. 2021. Test-time classifier adjustment module for model-agnostic domain generalization. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, 2427–2440.
- Khurana, A.; Paul, S.; Rai, P.; Biswas, S.; and Aggarwal, G. 2021. Sita: Single image test-time adaptation. *arXiv preprint arXiv:2112.02355*.
- Lee, J.; Jung, D.; Lee, S.; Park, J.; Shin, J.; Hwang, U.; and Yoon, S. 2024. Entropy is not enough for test-time adaptation: From the perspective of disentangled factors. In *The Twelfth International Conference on Learning Representations, ICLR*.
- Li, C.; Chen, T.; You, H.; Wang, Z.; and Lin, Y. 2020. Halo: Hardware-aware learning to optimize. In *European Conference on Computer Vision*, 500–518.
- Li, K.; and Malik, J. 2016. Learning to optimize. *arXiv preprint arXiv:1606.01885*.
- Li, K.; and Malik, J. 2017. Learning to optimize neural nets. *arXiv preprint arXiv:1703.00441*.
- Liu, Y.; Kothari, P.; Van Delft, B.; Bellot-Gurlet, B.; Mordan, T.; and Alahi, A. 2021. TTT++: When does self-supervised test-time training fail or thrive? In *Advances in Neural Information Processing Systems (NeurIPS)*, 21808–21820.
- Lv, K.; Jiang, S.; and Li, J. 2017. Learning gradient descent: Better generalization and longer horizons. In *Proceedings of the 34th International Conference on Machine Learning, ICML, 2247–2255*.
- Metz, L.; Maheswaranathan, N.; Nixon, J.; Freeman, D.; and Sohl-Dickstein, J. 2019. Understanding and correcting pathologies in the training of learned optimizers. In *Proceedings of the 36th International Conference on Machine Learning, ICML, 4556–4565*.
- Mirza, M. J.; Micorek, J.; Possegger, H.; and Bischof, H. 2022. The norm must go on: dynamic unsupervised domain adaptation by normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 14765–14775.

- Nado, Z.; Padhy, S.; Sculley, D.; D’Amour, A.; Lakshminarayanan, B.; and Snoek, J. 2020. Evaluating prediction-time batch normalization for robustness under covariate shift. *arXiv preprint arXiv:2006.10963*.
- Niu, S.; Miao, C.; Chen, G.; Wu, P.; and Zhao, P. 2024. Test-Time Model Adaptation with Only Forward Passes. In *The International Conference on Machine Learning*.
- Niu, S.; Wu, J.; Zhang, Y.; Chen, Y.; Zheng, S.; Zhao, P.; and Tan, M. 2022. Efficient test-time model adaptation without forgetting. In *Proceedings of the International Conference on Machine Learning (ICML)*, 16888–16905.
- Niu, S.; Wu, J.; Zhang, Y.; Wen, Z.; Chen, Y.; Zhao, P.; and Tan, M. 2023. Towards stable test-time adaptation in dynamic wild world. In *The Eleventh International Conference on Learning Representations, ICLR*.
- Sun, Y.; Li, X.; Dalal, K.; Xu, J.; Vikram, A.; Zhang, G.; Dubois, Y.; Chen, X.; Wang, X.; Koyejo, S.; et al. 2024. Learning to (learn at test time): Rnns with expressive hidden states. *arXiv preprint arXiv:2407.04620*.
- Sun, Y.; Wang, X.; Liu, Z.; Miller, J.; Efros, A.; and Hardt, M. 2020. Test-time training with self-supervision for generalization under distribution shifts. In *Proceedings of the International Conference on Machine Learning (ICML)*, 9229–9248.
- Tan, M.; Chen, G.; Wu, J.; Zhang, Y.; Chen, Y.; Zhao, P.; and Niu, S. 2024. Uncertainty-Calibrated Test-Time Model Adaptation without Forgetting. *arXiv preprint arXiv:2403.11491*.
- Vicol, P.; Metz, L.; and Sohl-Dickstein, J. 2021. Unbiased gradient estimation in unrolled computation graphs with persistent evolution strategies. In *Proceedings of the 38th International Conference on Machine Learning, ICML*, 10553–10563. PMLR.
- Wang, D.; Shelhamer, E.; Liu, S.; Olshausen, B.; and Darrell, T. 2021. Tent: Fully test-time adaptation by entropy minimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Wang, H.; Ge, S.; Lipton, Z.; and Xing, E. P. 2019. Learning robust global representations by penalizing local predictive power. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems*, 10506–10518.
- Wang, Q.; Fink, O.; Van Gool, L.; and Dai, D. 2022. Continual test-time domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7201–7211.
- Wen, Z.; Niu, S.; Li, G.; Wu, Q.; Tan, M.; and Wu, Q. 2023. Test-time model adaptation for visual question answering with debiased self-supervisions. *IEEE Transactions on Multimedia*.
- Wichrowska, O.; Maheswaranathan, N.; Hoffman, M. W.; Colmenarejo, S. G.; Denil, M.; Freitas, N.; and Sohl-Dickstein, J. 2017. Learned optimizers that scale and generalize. In *Proceedings of the 34th International Conference on Machine Learning, ICML*, 3751–3760.
- Yuan, L.; Xie, B.; and Li, S. 2023. Robust test-time adaptation in dynamic scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15922–15932.
- Zeng, R.; Deng, Q.; Xu, H.; Niu, S.; and Chen, J. 2023. Exploring Motion Cues for Video Test-Time Adaptation. In *Proceedings of the 31st ACM International Conference on Multimedia*, 1840–1850.
- Zhang, M.; Levine, S.; and Finn, C. 2022. Memo: Test time robustness via adaptation and augmentation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 38629–38642.