

Gradient Alignment Improves Test-Time Adaptation for Medical Image Segmentation

Ziyang Chen¹, Yiwen Ye¹, Yongsheng Pan¹ †, Yong Xia^{1,2,3} †

¹ National Engineering Laboratory for Integrated Aero-Space-Ground-Ocean Big Data Application Technology, School of Computer Science and Engineering, Northwestern Polytechnical University, Xi'an, China

² Research & Development Institute of Northwestern Polytechnical University in Shenzhen, Shenzhen, China

³ Ningbo Institute of Northwestern Polytechnical University, Ningbo, China
{zychen, ywye}@mail.nwpu.edu.cn, {yspan, yxia}@nwpu.edu.cn

Abstract

Although recent years have witnessed significant advancements in medical image segmentation, the pervasive issue of domain shift among medical images from diverse centres hinders the effective deployment of pre-trained models. Many Test-time Adaptation (TTA) methods have been proposed to address this issue by fine-tuning pre-trained models with test data during inference. These methods, however, often suffer from less-satisfactory optimization due to suboptimal optimization direction (dictated by the gradient) and fixed step-size (predicated on the learning rate). In this paper, we propose the **Gradient alignment-based Test-time adaptation (GraTa)** method to improve both the gradient direction and learning rate in the optimization procedure. Unlike conventional TTA methods, which primarily optimize the pseudo gradient derived from a self-supervised objective, our method incorporates an auxiliary gradient with the pseudo one to facilitate gradient alignment. Such gradient alignment enables the model to excavate the similarities between different gradients and correct the gradient direction to approximate the empirical gradient related to the current segmentation task. Additionally, we design a dynamic learning rate based on the cosine similarity between the pseudo and auxiliary gradients, thereby empowering the adaptive fine-tuning of pre-trained models on diverse test data. Extensive experiments establish the effectiveness of the proposed gradient alignment and dynamic learning rate and substantiate the superiority of our GraTa method over other state-of-the-art TTA methods on a benchmark medical image segmentation task.

Code — <https://github.com/Chen-Ziyang/GraTa>

Introduction

Medical image segmentation assumes a pivotal role in computer-aided diagnosis, offering precise delineation of specific anatomical structures. Over the past years, considerable research efforts (Tajbakhsh et al. 2020; Salpea, Tzouveli, and Kollias 2022; Liu et al. 2020; Wang et al. 2022; Tang et al. 2023; Xu et al. 2024) have been devoted to medical image segmentation utilizing deep-learning techniques, resulting in significant progress. However, domain shift, primarily caused by variations in scanners, imaging protocols,

and operators (Gibson et al. 2018; Ghafoorian et al. 2017), poses challenges for these models pre-trained on the labeled dataset (source domain) to generalize across different medical centers (target domain).

Test-time adaptation (TTA) has emerged as a prospective paradigm to mitigate domain shift with minimal data demand, relying only on test data during inference. Several studies on TTA attempt to alleviate the domain shift by modifying the statistics stored in batch normalization (BN) layers (Wang et al. 2023b; Nado et al. 2020; Mirza et al. 2022; Zhang et al. 2023; Park et al. 2024). A typical example is DUA (Mirza et al. 2022), which incrementally adjusts the statistics within BN layers from the source to the target domain to enhance feature representations. Although these methods improve adaptation, their performance potential is limited since the model’s parameters remain fixed. The mainstream TTA methods focus on designing self-supervised objectives (*e.g.*, entropy minimization, consistency constraint, and rotation prediction) to fine-tune the pre-trained models to boost their generalization performance on the target domain (Wang et al. 2021; Sun et al. 2020; Zhang et al. 2023; Niu et al. 2023; Nguyen et al. 2023; Sinha et al. 2023; Yang et al. 2022; Bateson, Lombaert, and Ben Ayed 2022; Wen et al. 2024). For instance, TENT (Wang et al. 2021) introduces an entropy minimization objective to TTA to optimize the affine-transformation parameters of pre-trained models, and TTT (Sun et al. 2020) updates the model parameters by predicting the rotation angle of test data in a self-supervised manner.

However, all these methods, based on gradient descent, overlook two critical elements in the optimization procedure: the direction and the step-size. The i th optimization procedure of the pre-trained model with parameters θ can be formulated as $\theta_{i+1} \leftarrow \theta_i - \eta \nabla \mathcal{L}_{pse}(\theta_i)$, where η is the learning rate, and $\nabla \mathcal{L}_{pse}$ denotes the gradient produced by the self-supervised objective function, called pseudo gradient. Obviously, the optimization direction is determined by $\nabla \mathcal{L}_{pse}(\theta_i)$, and the optimization step-size depends on η .

In Figure 1(a), we displayed two gradients: an empirical gradient, customized for the specific task (segmentation in this study) using provided labels and not encountered in the TTA setup; and a pseudo gradient, primarily employed for fine-tuning pre-trained models in existing TTA methods. Ideally, the pseudo gradient should exhibit a direction




Figure 1: Illustration of our motivation. (a) We display a pseudo gradient $\nabla\mathcal{L}_{pse}(\theta)$, which is primarily utilized for optimization but may diverge far from the empirical gradient $\nabla\mathcal{L}_{emp}(\theta)$, tailored to the specific task (segmentation in this study). Existing methods typically optimize $\nabla\mathcal{L}_{pse}(\theta)$ in a straightforward manner. (b) Our GraTa introduces an auxiliary gradient $\nabla\mathcal{L}_{aux}(\theta)$ to minimize the angle between $\nabla\mathcal{L}_{pse}(\theta)$ and $\nabla\mathcal{L}_{aux}(\theta)$, resulting in a novel auxiliary objective, *i.e.*, gradient alignment. (c) Although achieving complete alignment is challenging due to the different objectives of these two gradients, the model can learn to align their task-relevant components ∇_{\checkmark} through this auxiliary objective, approximating $\nabla\mathcal{L}_{emp}(\theta)$ and facilitating effective fine-tuning. \angle denotes the angle. ‘Obj’: Abbreviation of ‘Objective’.

similar to the empirical gradient. Unfortunately, due to the lack of reliable supervision, their directions may differ significantly, posing challenges to model optimization. Current TTA methods simply optimize the pseudo gradient, failing to address this issue. Moreover, concerning the learning rate, almost all these methods employ a fixed value, thereby restricting the pre-trained models from being adaptively fine-tuned on diverse test data.

In this paper, we focus on optimizing both the optimization direction and the learning rate. We propose a novel TTA method, namely **Gradient alignment-based Test-time adaptation (GraTa)**. GraTa aims to improve the gradient direction through gradient alignment and enables pre-trained models to adapt to test data using a dynamic learning rate. As shown in Figure 1(b), we incorporate a new auxiliary objective, *i.e.*, gradient alignment, by introducing an auxiliary gradient to the pseudo one. Specifically, we employ the consistency loss to produce the pseudo gradient, while leveraging the gradient derived from the entropy loss as the auxiliary one. The entropy loss is computed on the original test data. The consistency loss is conducted on the weak and strong augmentation variants of the test data. Through alignment, the model is capable of excavating the similarities between distinct gradients, especially the components relevant to the current segmentation task, and approximating the empirical gradient (see Figure 1(c)). Furthermore, we also present a dynamic learning rate, which is inversely proportional to the angle between these two gradients, to adaptively fine-tune the pre-trained models. A larger angle implies greater conflict in the optimization of the two gradients, thus requiring a smaller learning rate, and vice versa. Extensive experiments demonstrate that our GraTa achieves a superior performance over other state-of-the-art TTA methods.

Our contributions are three-fold: (1) we rethink the optimization procedure in existing TTA methods and propose GraTa to improve both the optimization direction and step-size; (2) we align the pseudo and auxiliary gradients, aimed at distinct objectives, to reduce the divergence of their task-

specific components, thereby improving the optimization direction; and (3) we design a variable learning rate considering the angle between distinct gradients, aiding in the dynamical determination of the optimization step-size for adaptive fine-tuning.

Related Work

Test-time Adaptation

Test-time Adaptation (TTA) addresses domain shift by adapting the pre-trained source model to the distribution of the target domain using test data during inference (Liang, He, and Tan 2024). TTA can be employed in a source-free and online manner, allowing the model to adapt without accessing the source data. The mainstream TTA methods are based on fine-tuning the model by constructing self-supervised auxiliary tasks to guide the model in training on the test data, enhancing its performance and adaptability (Wang et al. 2021; Sun et al. 2020; Zhang, Levine, and Finn 2022; Iwasawa and Matsuo 2021; Nguyen et al. 2023; Niu et al. 2023; Yang et al. 2022; Zhang et al. 2023; Wen et al. 2024). (Wang et al. 2021) proposed a test-time entropy minimization scheme to decrease the entropy of model predictions by fine-tuning the affine parameters within batch normalization layers. (Zhang et al. 2023) extended the vanilla entropy minimization loss to a generalized one to better utilize the information in the test data. (Niu et al. 2023) filtered and removed partial noise with large uncertainty in the supervised information and employed sharpness-aware minimization to optimize the parameters of the model toward a flat minimum. These methods demonstrate effective adaptation when the supervised information is reliable but may lead to unexpected performance degradation when the supervised information is unreliable. In this paper, our purpose is to improve the optimization direction and step-size by gradient alignment and dynamic learning rate, respectively. After aligning the gradients, we can obtain more reliable supervised information, thereby facilitating more effective training of the source model.

Gradient Alignment

In deep learning, the objective function serves as a measure to quantify the discrepancy between a model’s prediction and the expected output. Objective functions with similar gradients generally indicate similar objectives and contribute to robust training, which is particularly important in multi-task learning. A typical approach in multi-task learning is PCGrad (Yu et al. 2020), which is developed to mitigate gradient conflicts between task gradients via replacing the gradient by its projection onto the normal plane of another gradient. In contrast, our optimization objective does not involve resolving conflicts between multiple gradients, as it does not directly optimize multiple gradients simultaneously. Instead, the goal of our GraTa is to improve the direction of pseudo gradient, thereby enhancing the model’s fine-tuning process. Recently, many methods seek to align the gradients of multiple objective functions to improve training (Gao et al. 2021; Dandi, Barba, and Jaggi 2022; Joo et al. 2023; Le and Woo 2024; Zhao et al. 2023, 2021). (Gao et al. 2021) devised a gradient discriminator to align the gradients from the source and target domains to improve the similarity between two distributions for better adversarial domain adaptation. (Dandi, Barba, and Jaggi 2022) proposed an algorithm for federated learning that induces the implicit regularization to promote the alignment of gradients across different clients. (Joo et al. 2023) presented to combine both L2 and cosine distance-based criteria as regularization terms, leveraging their strengths in aligning the local gradient to produce more robust interpretations. Different from these methods, we implicitly perform gradient alignment to capture the similarity between gradients of two distinct objective functions to improve the gradient direction for test-time adaptation.

Dynamic Learning Rate

The learning rate is a critical hyperparameter in deep learning, as it determines the step-size to update the model. An excessively large learning rate may hinder the model’s convergence, while an excessively small learning rate may make the model trapped in local optima. Extensive research has been conducted to dynamically adjust the learning rate to enhance model training (Yang et al. 2022; Xiao and Zhang 2021; Wan et al. 2021; Xiong et al. 2022; Yu, Chen, and Cheng 1995). (Yang et al. 2022) presented a dynamic learning rate by constructing a memory bank and calculating the cosine similarity scores between the feature of the current test image and previous features from the memory bank. (Xiao and Zhang 2021) devised the dynamic weighted learning to dynamically weights the learning losses of alignment and discriminability to avoid excessive alignment learning and excessive discriminant learning. (Xiong et al. 2022) proposed a graph-network-based scheduler to learn a specific scheduling mechanism without restrictions and control the learning rate via reinforcement learning. In this paper, our focus is on enhancing the optimization direction by aligning two distinct gradients. The cosine similarity between these two gradients can be employed as a natural metric to measure the degree of alignment, allowing our GraTa to adjust the learning rate dynamically.

Method

Problem Definition

Let the labeled source domain dataset and unlabeled testing target dataset be $\mathcal{D}^s = \{\mathcal{X}_i^s, \mathcal{Y}_i^s\}_{i=1}^{N^s}$ and $\mathcal{D}^t = \{\mathcal{X}_i^t\}_{i=1}^{N^t}$, respectively, where $\mathcal{X}_i^* \in \mathbb{R}^{H \times W \times C}$ is the i -th image with C channels and size of $H \times W$, and \mathcal{Y}_i^* is its label. Our goal is to fine-tune the model $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ pre-trained on \mathcal{D}^s using \mathcal{X}_i^t to adapt to \mathcal{D}^t . An overview of our GraTa is illustrated in Figure 2. Now we delve into the details.

The Objective of GraTa

In this study, we introduce the entropy loss \mathcal{L}_{ent} as the auxiliary objective function to align with the consistency loss \mathcal{L}_{con} . Our GraTa aims at minimizing not only \mathcal{L}_{con} but also the angle between gradients derived from \mathcal{L}_{con} and \mathcal{L}_{ent} for gradient alignment. \mathcal{L}_{ent} is calculated on the prediction P_i of original test data \mathcal{X}_i^t as follows:

$$\mathcal{L}_{ent}(\theta; \mathcal{X}_i^t) = -P_i \log P_i, P_i = \text{Sigmoid}(f_\theta(\mathcal{X}_i^t)). \quad (1)$$

\mathcal{L}_{con} is conducted on the set of weak augmentation variants $\mathcal{S}_i = \{\hat{\mathcal{X}}_{ij}^t | \hat{\mathcal{X}}_{ij}^t = \text{Aug}_w^j(\mathcal{X}_i^t)\}_{j=1}^6$ and the strong augmentation variant $\tilde{\mathcal{X}}_i^t = \text{Aug}_s(\mathcal{X}_i^t)$ as follows:

$$\begin{aligned} \mathcal{L}_{con}(\theta; \mathcal{X}_i^t) &= l_{ce}(\hat{P}_i, \tilde{P}_i), \tilde{P}_i = \text{Sigmoid}(f_\theta(\tilde{\mathcal{X}}_i^t)), \\ \hat{P}_i &= \frac{1}{j} \sum_{j=1}^6 \text{Sigmoid}(f_\theta(\hat{\mathcal{X}}_{ij}^t)), \end{aligned} \quad (2)$$

where l_{ce} denotes the cross-entropy loss. The weak augmentation strategy Aug_w consists of six augmentation techniques: identity mapping, horizontal flipping, vertical flipping, rotation 90° , 180° , and 270° . We sequentially apply each Aug_w^j to \mathcal{X}_i^t to obtain \mathcal{S} . The strong augmentation strategy Aug_s incorporates brightness adjustment, contrast adjustment, Gamma transformation, Gaussian noise, and Gaussian blur. \mathcal{X}_i^t is transformed to $\tilde{\mathcal{X}}_i^t$ utilizing all of these techniques.

Then the overall objective can be formulated as

$$\min_{\theta} \mathcal{L}_{con}(\theta; \mathcal{X}_i^t) + \angle(\nabla \mathcal{L}_{con}(\theta; \mathcal{X}_i^t), \nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t)), \quad (3)$$

where \angle denotes the angle. However, the angle is non-differentiable and cannot be directly minimized, and a common solution is to maximize the inner product of these two gradients (Wang et al. 2023a), *i.e.*, $\nabla \mathcal{L}_{con}(\theta; \mathcal{X}_i^t) \cdot \nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t)$. Unfortunately, calculating the inner product explicitly will introduce the calculation of the Hessian matrix during back-propagation, which is computationally complex and unstable. Inspired by (Li et al. 2018), we rewrite the objective to implicitly maximize the inner product as follows:

$$\min_{\theta} \mathcal{L}_{con}(\theta'; \mathcal{X}_i^t), \quad (4)$$

where $\theta' = \theta - \nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t)$.

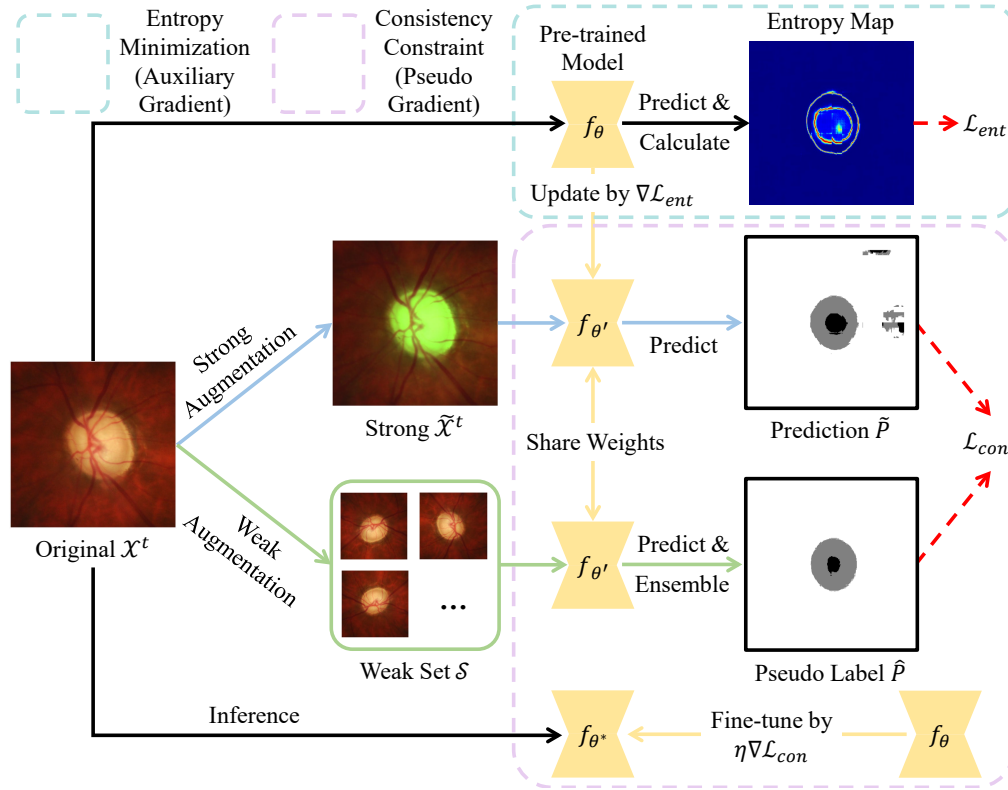


Figure 2: Overview of our GraTa. For each test image, we first calculate the entropy loss on its prediction to update the pre-trained model f_θ by $\theta' = \theta - \nabla \mathcal{L}_{ent}$. Then we perform weak and strong augmentation on the original test image to obtain a strong augmentation variant and a set of weak augmentation variants and calculate the consistency loss on their predictions produced by $f_{\theta'}$. Finally, the consistency loss is utilized to fine-tune the model by $\theta^* \leftarrow \theta - \eta \nabla \mathcal{L}_{con}$, and the test image is fed into f_{θ^*} for inference.

Proof. For Eq. 4, we perform the first-order Taylor expansion around θ as follows:

$$\begin{aligned}
& \min_{\theta} \mathcal{L}_{con}(\theta - \nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t); \mathcal{X}_i^t) \\
&= \min_{\theta} \mathcal{L}_{con}(\theta; \mathcal{X}_i^t) - \nabla \mathcal{L}_{con}(\theta; \mathcal{X}_i^t) \cdot \nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t) + O \\
&\approx \min_{\theta} \mathcal{L}_{con}(\theta; \mathcal{X}_i^t) - \nabla \mathcal{L}_{con}(\theta; \mathcal{X}_i^t) \cdot \nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t), \quad (5)
\end{aligned}$$

where O denotes the remainder term, which can be omitted. Then, the objective in Eq. 4 can be rewritten as:

$$\min_{\theta} \mathcal{L}_{con}(\theta; \mathcal{X}_i^t) - \nabla \mathcal{L}_{con}(\theta; \mathcal{X}_i^t) \cdot \nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t). \quad (6)$$

Eq. (6) reveals that our objective is to minimize $\mathcal{L}_{con}(\theta; \mathcal{X}_i^t)$ and maximize the inner product of these two gradients. As $\nabla \mathcal{L}_{con}(\theta; \mathcal{X}_i^t)$ and $\nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t)$ become more closely aligned, their inner product increases, reaching the maximum when they are completely aligned.

Dynamic Learning Rate

We design a dynamic learning rate η for adaptive fine-tuning. η is determined based on the cosine similarity

Algorithm 1: The Algorithm of GraTa.

- Initialize :** Pre-trained source model f_{θ_0} and scaling factor β .
- Input:** For each time step i (starting from 0), current test image \mathcal{X}_i^t .
- ▷ **Fine-tune the model:**
- 1: Calculate $\mathcal{L}_{ent}(\theta_i; \mathcal{X}_i^t)$ by Eq. (1) and perform back-propagation
 - 2: Update the model by $\theta'_i = \theta_i - \nabla \mathcal{L}_{ent}(\theta_i; \mathcal{X}_i^t)$
 - 3: Calculate $\mathcal{L}_{con}(\theta'_i; \mathcal{X}_i^t)$ by Eq. (2) and perform back-propagation
 - 4: Obtain η utilizing β through Eq. (5)
 - 5: Fine-tune the model by $\theta_{i+1} \leftarrow \theta_i - \eta \nabla \mathcal{L}_{con}(\theta'_i; \mathcal{X}_i^t)$
- ▷ **Inference:**
- 6: Forward $\mathbb{P}_i = f_{\theta_{i+1}}(\mathcal{X}_i^t)$
- Output:** Adapted prediction \mathbb{P}_i
-

tween $\nabla \mathcal{L}_{con}(\theta'; \mathcal{X}_i^t)$ and $\nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t)$, formulated as:

$$\eta = \beta * \text{Cos}(\frac{\nabla \mathcal{L}_{con}(\theta'; \mathcal{X}_i^t) \cdot \nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t)}{\|\nabla \mathcal{L}_{con}(\theta'; \mathcal{X}_i^t)\| \|\nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t)\|}), \quad (7)$$

where β is a scaling factor, and $\text{Cos}(x) = \frac{1}{4}(x + 1)^2$ is a custom increasing activation function to map the cosine sim-

ilarity to $[0, 1]$. As the cosine similarity increases, the optimization directions of $\nabla \mathcal{L}_{con}(\theta'; \mathcal{X}_i^t)$ and $\nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t)$ become more aligned, thus assigning a larger learning rate, and conversely. The overall process of our GraTa is summarized in Algorithm 1.

Experiments and Results

Datasets and Evaluation Metrics

We evaluate our proposed GraTa and other state-of-the-art TTA methods on the joint optic disc (OD) and cup (OC) segmentation task, which comprises five public datasets collected from different medical centres, denoted as domain A (RIM-ONE-r3 (Fumero et al. 2011)), B (REFUGE (Orlando et al. 2020)), C (ORIGA (Zhang et al. 2010)), D (REFUGE-Validation/Test (Orlando et al. 2020)), and E (Drishti-GS (Sivaswamy et al. 2014)). These datasets consist of 159, 400, 650, 800, and 101 images, respectively. For each image, we cropped a region of interest (ROI) centered at the OD with a size of 800×800 , and each ROI is further resized to 512×512 and normalized by min-max normalization following (Hu, Liao, and Xia 2022). We utilize the Dice score metric (DSC) for evaluation.

Implementation Details

We trained a ResUNet-34 (He et al. 2016) backbone following (Hu, Liao, and Xia 2022) as the baseline individually on each domain (source domain) and subsequently tested it on each remaining domain (target domain), computing the mean metrics to evaluate all the methods across diverse scenarios (4×5 in total). For a fair comparison, we conducted single-iteration adaptation for each batch of test data using a batch size of 1 across all experiments following (Yang et al. 2022). To deploy our GraTa, we utilized the Adam (Kingma and Ba 2014) optimizer to train the affine-transformation parameters within pre-trained models, and the BN statistics were recollected from test data following (Wang et al. 2021). The scaling factor β is set to 0.0001 empirically.

Results

We compare our GraTa with the ‘No Adapt’ baseline (testing without adaptation), and eight competing TTA methods, including three methods based on refining the BN statistics (DUA (Mirza et al. 2022), DIGA (Wang et al. 2023b), and MedBN (Park et al. 2024)), and five methods based on fine-tuning the pre-trained models (TENT (Wang et al. 2021), DLTTA (Yang et al. 2022), DomainAdaptor (Zhang et al. 2023), SAR (Niu et al. 2023), and DeTTA (Wen et al. 2024)). Specifically, we re-implemented all competing methods using the same baseline as our GraTa and integrated DLTTA with the entropy loss proposed in TENT (Yang et al. 2022). The result of each domain is the average metric calculated by treating the current domain as the source domain and employing the pre-trained model on remaining domains. **Comparing to other TTA methods.** The results of our GraTa, the ‘No Adapt’ baseline, and eight competing TTA methods are detailed in Table 1. The BN-based methods exhibit stable but suboptimal enhancements across all domains due to the absence of model updating. TENT exhibits

limited performance, while other variants (*i.e.*, DLTTA, DomainAdaptor, and SAR) enhance the performance significantly. It can be observed that our GraTa achieves the best performance across most scenarios and improves the baseline by 5.87%, which demonstrates the effectiveness and superiority of our GraTa. Specifically, our GraTa boosts the performance in Domain D by 12.77%, while other methods exhibit marginal improvement, which means our approach extends the upper limit of TTA in challenging scenarios.

Ablation study. To evaluate the contributions of our optimization objective and dynamic learning rate, we conducted a series of ablation experiments, as shown in Table 2. It shows that (1) using $\mathcal{L}_{con}(\theta; \mathcal{X}_i^t)$ only obtains limited performance similar to TENT; (2) introducing $\mathcal{L}_{con}(\theta - \nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t); \mathcal{X}_i^t)$ and η enhances the performance by 2.14% and 2.81%, respectively; (3) the best performance is achieved when both are utilized jointly (*i.e.*, our GraTa).

Qualitative analysis. We selected the model trained on Domain D (source) as a case study and visualized four test images from four scenarios, denoted as ‘‘D \rightarrow target domain,’’ along with their corresponding segmentation results produced by our GraTa method, eight competing methods, and the ground truth, as shown in Figure 3. The regions of OD and OC are highlighted in green and blue, respectively. For comparative purposes, we drew the bounding boxes of the OD and OC on the ground truth and overlaid them on each segmentation result. Compared to other competing methods, our GraTa demonstrates reduced over-segmentation and under-segmentation across all samples, indicating the effectiveness of the improved optimization procedure.

Discussions

Exchange the Location of \mathcal{L}_{con} and \mathcal{L}_{ent} . It is particularly critical to determine which objective function is the optimal one to produce the pseudo gradient. We exchanged the location of \mathcal{L}_{con} and \mathcal{L}_{ent} in our objective and repeated the experiments for evaluation. The results are listed in Table 3. Compared to TENT which optimizes $\mathcal{L}_{ent}(\theta; \mathcal{X}_i^t)$ only, the variant of our GraTa using $\mathcal{L}_{ent}(\theta - \nabla \mathcal{L}_{con}(\theta; \mathcal{X}_i^t); \mathcal{X}_i^t)$ as the optimization objective and η as the learning rate also boosts the performance. Our GraTa achieves the best performance, which demonstrates that \mathcal{L}_{con} is more suitable to produce the pseudo gradient due to its superior supervision.

Obtain η using other activation functions. To ascertain the optimal activation function for mapping the cosine similarity to obtain η , we conducted experiments using various activation functions and listed the results in Table 4. It reveals that (1) our designed function achieves superior performance to others by amplifying the distinctions among learning rates generated by diverse cosine similarity scores; (2) ReLU exhibits the worst average performance which underscores the necessity of assigning appropriate learning rates for fine-tuning even if distinct gradients diverge largely.

Alignment between the empirical and pseudo gradients. To demonstrate that GraTa can align the empirical gradient and pseudo gradient more closely, we repeated the experiments on all scenarios and visualized the cosine similarity between these two gradients, where the empirical gradient is produced by the cross-entropy loss in a supervised manner,

Methods		Domain A	Domain B	Domain C	Domain D	Domain E	Average
		<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i> ↑
No Adapt		67.96	76.71	74.71	54.07	68.88	68.47
BN-based	DUA (Mirza et al. 2022)	73.10	77.10	75.34	58.62	72.99	71.43
	DIGA (Wang et al. 2023b)	76.57	77.10	73.64	62.04	71.54	72.18
	MedBN (Park et al. 2024)	75.14	76.77	74.46	59.91	72.00	71.65
Fine-tune-based	TENT (Wang et al. 2021)	74.06	78.55	74.56	51.45	69.66	69.65
	DLTTA (Yang et al. 2022)	75.19	78.48	76.37	57.50	71.35	71.78
	DomainAdaptor (Zhang et al. 2023)	75.88	77.43	75.66	58.64	72.38	72.00
	SAR (Niu et al. 2023)	75.26	78.07	75.07	61.24	73.69	72.67
	DeTTA (Wen et al. 2024)	75.15	78.17	75.27	61.26	73.61	72.69
GraTa (Ours)		76.61	78.81	76.52	66.84	72.94	74.34

Table 1: Performance of our GraTa, ‘No Adapt’ baseline, and eight TTA methods. The best results are highlighted in **bold**.

Objective	Learning Rate	Domain A	Domain B	Domain C	Domain D	Domain E	Average
		<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i> ↑
-	-	67.96	76.71	74.71	54.07	68.88	68.47
$\mathcal{L}_{con}(\theta; \mathcal{X}_i^t)$	β	78.35	78.50	69.24	52.13	68.74	69.39
$\mathcal{L}_{con}(\theta - \nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t); \mathcal{X}_i^t)$	β	77.97	78.77	73.90	56.53	70.46	71.53
$\mathcal{L}_{con}(\theta - \nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t); \mathcal{X}_i^t)$	η	76.61	78.81	76.52	66.84	72.94	74.34

Table 2: Performance of using various objectives and learning rates for optimization. The best results are highlighted in **bold**.

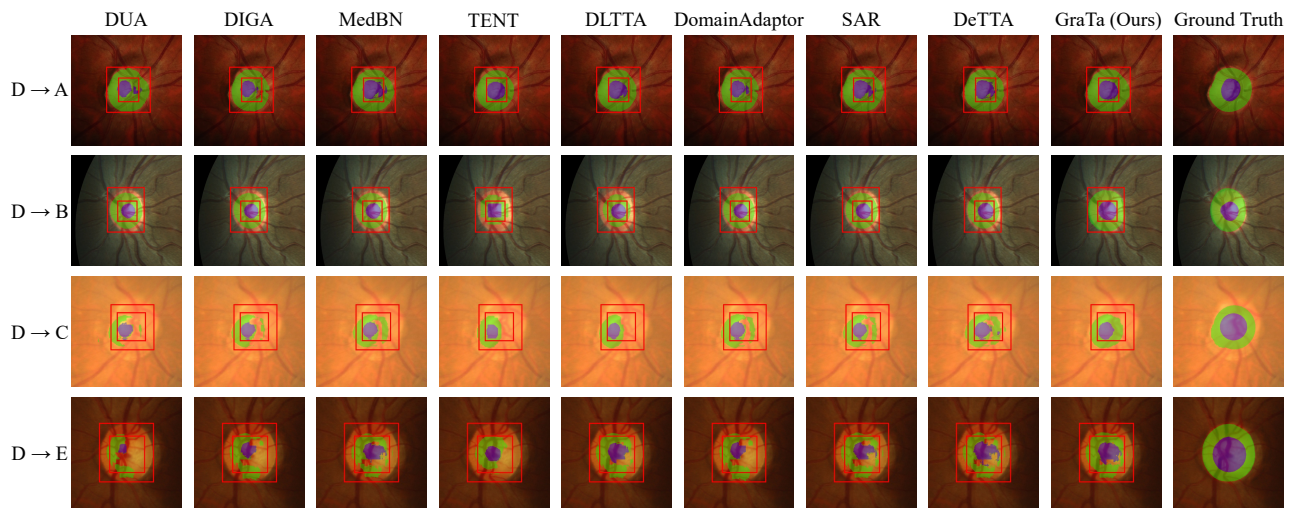


Figure 3: Qualitative results of our GraTa and eight TTA methods using Domain D as the source domain and remaining domains as target domains. We displayed the results of two simple samples (rows 1-2) and two hard samples (rows 3-4). The ground-truth bounding boxes of OD and OC were overlaid on each segmentation result to highlight potential over- or under-segmentation. Best viewed in color.

Optimization Objective	Learning Rate	Domain A	Domain B	Domain C	Domain D	Domain E	Average
		<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i> ↑
$\mathcal{L}_{ent}(\theta; \mathcal{X}_i^t)$	β	74.06	78.55	74.56	51.45	69.66	69.65
$\mathcal{L}_{ent}(\theta - \nabla \mathcal{L}_{con}(\theta; \mathcal{X}_i^t); \mathcal{X}_i^t)$	η	74.69	78.46	75.60	56.75	71.86	71.47
$\mathcal{L}_{con}(\theta - \nabla \mathcal{L}_{ent}(\theta; \mathcal{X}_i^t); \mathcal{X}_i^t)$	η	76.61	78.81	76.52	66.84	72.94	74.34

Table 3: Performance of TENT, our GraTa, and its variant. The best results are highlighted in **bold**.

as shown in Figure 4. We reported the average results using each domain as the source domain and remaining ones as target domains. The greater the cosine similarity, the closer the pseudo gradient aligns with the empirical gradient. The results reveal the effectiveness of our GraTa in improving gradient direction and also indicate that even small alignment

improvements can lead to significant performance gains.

More combinations of objective functions. For further discussion, we introduced four additional loss functions, *i.e.*, the reconstruction loss, the rotation prediction loss, the super-resolution loss, and the denoising loss. We utilized the cross-entropy loss to calculate the rotation prediction loss

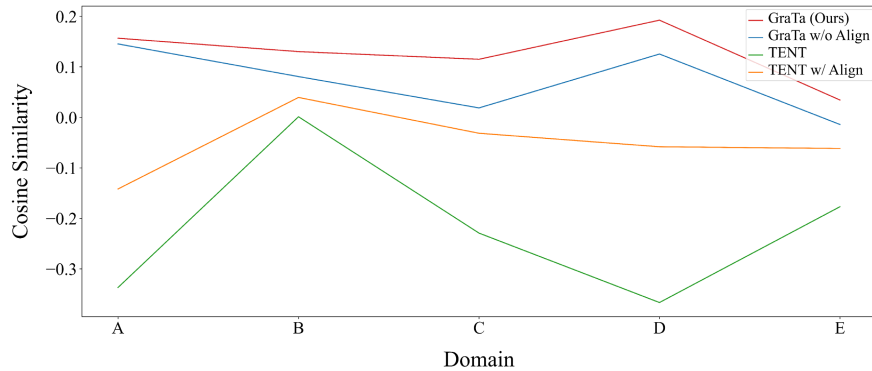


Figure 4: Cosine similarity between the pseudo gradient and empirical gradient w/ or w/o our proposed gradient alignment. The objectives of ‘‘TENT’’ and ‘‘TENT w/ Align’’ are $\mathcal{L}_{ent}(\theta; \mathcal{X}_i^t)$ and $\mathcal{L}_{ent}(\theta - \nabla \mathcal{L}_{con}(\theta; \mathcal{X}_i^t); \mathcal{X}_i^t)$, respectively.

Function	Do. A	Do. B	Do. n C	Do. D	Do. E	Average
	<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i> \uparrow
Linear: $\frac{1}{2}(x+1)$	77.15	79.11	76.07	63.79	72.42	73.71
Sigmoid: $\frac{1}{1+e^{-x}}$	77.31	79.04	76.01	64.69	72.10	73.83
ReLU: $\max(0, x)$	75.92	78.29	75.16	62.14	73.07	72.92
Softplus: $\ln(1+e^x)$	77.60	78.81	75.44	63.04	71.68	73.32
Ours: $\frac{1}{4}(x+1)^2$	76.61	78.81	76.52	66.84	72.94	74.34

Table 4: Performance of using different activation functions to map the cosine similarity. e denotes the natural base. The best results are highlighted in **bold**.

and the mean-square-error loss to calculate others. To construct the super-resolution loss, we downsampled the input image by a factor of four. For the denoising loss, we added the Gaussian noise sampled from a standard normal distribution to the image. We evaluated all possible combinations of them to determine the optimal one and displayed the results in Table 5. It reveals that (1) using \mathcal{L}_{con} as the pseudo objective function achieves superior performance over others; (2) using \mathcal{L}_{rec} , \mathcal{L}_{rot} , \mathcal{L}_{sup} , or \mathcal{L}_{den} as the pseudo objective function results in similar but suboptimal performance since they all require an additional branch to construct the loss, making it challenging to provide appropriate gradients during inference; (3) using \mathcal{L}_{rec} , \mathcal{L}_{rot} , \mathcal{L}_{sup} , or \mathcal{L}_{den} as the auxiliary objective function can enhance the performance of \mathcal{L}_{con} and \mathcal{L}_{ent} , demonstrating the effectiveness of GraTa.

Conclusion

In this paper, we proposed the GraTa, a TTA method to focus on optimizing the optimization direction and step-size during inference by gradient alignment and utilizing a dynamic learning rate, respectively. Specifically, we introduced an auxiliary gradient to the pseudo one to perform alignment to excavate the similarities between distinct gradients and further approximate the empirical gradient. The dynamic learning rate is constructed based on the cosine similarity between the pseudo and auxiliary gradients, encouraging pre-trained models to enhance learning as these gradients become closely aligned. Extensive experiment results on diverse scenarios demonstrated the superiority of our GraTa and the effectiveness of each special design. In our future

Aux	Pse	Do. A	Do. B	Do. C	Do. D	Do. E	Average	
		<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i> \uparrow	
\mathcal{L}_{con}	\mathcal{L}_{con}	76.61	78.81	76.52	66.84	72.94	74.34	
		76.76	78.94	74.74	66.43	72.19	73.81	
		76.58	79.02	74.39	67.90	72.53	74.08	
		\mathcal{L}_{sup}	76.68	78.99	74.32	67.82	72.31	74.02
		\mathcal{L}_{den}	76.71	78.94	75.11	65.18	72.26	73.64
		\mathcal{L}_{rec}	74.69	78.46	75.60	56.75	71.86	71.47
\mathcal{L}_{ent}	\mathcal{L}_{ent}	75.27	78.54	76.36	60.29	70.97	72.28	
		\mathcal{L}_{rot}	75.19	78.54	76.34	56.29	70.78	71.43
		\mathcal{L}_{sup}	75.20	78.53	76.53	57.41	70.79	71.69
		\mathcal{L}_{den}	75.37	78.52	76.67	59.38	70.72	72.13
		\mathcal{L}_{con}	75.18	77.85	75.13	62.33	73.71	72.84
		\mathcal{L}_{ent}	74.47	77.74	74.71	59.64	73.49	72.01
\mathcal{L}_{rec}	\mathcal{L}_{rec}	\mathcal{L}_{rot}	75.74	78.04	75.12	60.13	73.91	72.59
		\mathcal{L}_{sup}	74.58	77.51	74.91	62.99	73.67	72.73
		\mathcal{L}_{den}	75.44	77.82	75.66	60.12	73.73	72.55
		\mathcal{L}_{con}	75.31	78.01	75.02	61.66	73.81	72.76
		\mathcal{L}_{ent}	75.37	78.17	75.07	62.14	73.75	72.90
		\mathcal{L}_{rec}	75.19	78.10	74.39	61.57	73.78	72.60
\mathcal{L}_{rot}	\mathcal{L}_{rot}	\mathcal{L}_{sup}	75.31	78.08	75.95	61.24	73.73	72.86
		\mathcal{L}_{den}	75.27	78.21	74.84	61.56	73.65	72.70
		\mathcal{L}_{con}	75.36	78.17	75.27	61.49	73.77	72.81
		\mathcal{L}_{ent}	75.51	77.86	74.67	59.23	73.47	72.15
		\mathcal{L}_{rec}	75.26	78.09	74.53	59.78	73.66	72.27
		\mathcal{L}_{rot}	75.17	77.87	74.92	58.48	73.75	72.04
\mathcal{L}_{sup}	\mathcal{L}_{sup}	\mathcal{L}_{den}	74.94	77.91	74.83	60.65	73.63	72.39
		\mathcal{L}_{con}	75.34	78.02	74.47	61.74	73.71	72.66
		\mathcal{L}_{ent}	75.20	77.44	74.77	59.22	73.38	72.00
		\mathcal{L}_{rec}	75.40	78.09	74.51	60.71	73.59	72.46
		\mathcal{L}_{rot}	75.32	77.71	76.49	55.30	73.44	71.65
		\mathcal{L}_{sup}	75.48	77.87	75.08	59.82	73.68	72.38

Table 5: Performance of various optimization objective combinations, where \mathcal{L}_{con} , \mathcal{L}_{ent} , \mathcal{L}_{rec} , \mathcal{L}_{rot} , \mathcal{L}_{sup} , \mathcal{L}_{den} denote the consistency loss, the entropy loss, the reconstruction loss, the rotation prediction loss, the super-resolution loss, and the denoising loss, respectively. The best results are highlighted in **bold**. ‘Aux’: Abbreviation of ‘Auxiliary’. ‘Pse’: Abbreviation of ‘Pseudo’.

work, we intend to explore the alignment of other auxiliary gradients to further enhance performance, as well as investigate other deployment scenarios, such as continual test-time adaptation (Chen et al. 2024). We hope that our efforts can promote research on gradient optimization for TTA.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (Nos. 62171377, 92470101, 6240012686), in part by grants from National Key R&D Program of China (2022YFC2009903/2022YFC2009900), in part by the Ningbo Clinical Research Center for Medical Imaging under Grant 2021L003 (Open Project 2022LYK-FZD06), in part by Shenzhen Science and Technology Program under Grants JCYJ20220530161616036, in part by the Fundamental Research Funds for the Central Universities (No. D5000230376), and in part by the Innovation Foundation for Doctor Dissertation of Northwestern Polytechnical University under Grant CX2024016.

References

- Bateson, M.; Lombaert, H.; and Ben Ayed, I. 2022. Test-time adaptation with shape moments for image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 736–745. Springer.
- Chen, Z.; Pan, Y.; Ye, Y.; Lu, M.; and Xia, Y. 2024. Each test image deserves a specific prompt: Continual test-time adaptation for 2d medical image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11184–11193.
- Dandi, Y.; Barba, L.; and Jaggi, M. 2022. Implicit gradient alignment in distributed and federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 6454–6462.
- Fumero, F.; Alayón, S.; Sanchez, J. L.; Sigut, J.; and Gonzalez-Hernandez, M. 2011. RIM-ONE: An open retinal image database for optic nerve evaluation. In *2011 24th International Symposium on Computer-based medical Systems (CBMS)*, 1–6. IEEE.
- Gao, Z.; Zhang, S.; Huang, K.; Wang, Q.; and Zhong, C. 2021. Gradient distribution alignment certificates better adversarial domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8937–8946.
- Ghafoorian, M.; Mehrtash, A.; Kapur, T.; Karssemeijer, N.; Marchiori, E.; Pesteie, M.; Guttmann, C. R.; de Leeuw, F.-E.; Tempny, C. M.; Van Ginneken, B.; et al. 2017. Transfer learning for domain adaptation in MRI: Application in brain lesion segmentation. In *Medical Image Computing and Computer Assisted Intervention- MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11-13, 2017, Proceedings, Part III 20*, 516–524. Springer.
- Gibson, E.; Hu, Y.; Ghavami, N.; Ahmed, H. U.; Moore, C.; Emberton, M.; Huisman, H. J.; and Barratt, D. C. 2018. Inter-site variability in prostate segmentation accuracy using deep learning. In *Medical Image Computing and Computer Assisted Intervention- MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part IV 11*, 506–514. Springer.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Hu, S.; Liao, Z.; and Xia, Y. 2022. Prosfa: Prompt learning based source-free domain adaptation for medical image segmentation. *arXiv preprint arXiv:2211.11514*.
- Iwasawa, Y.; and Matsuo, Y. 2021. Test-time classifier adjustment module for model-agnostic domain generalization. *Advances in Neural Information Processing Systems*, 34: 2427–2440.
- Joo, S.; Jeong, S.; Heo, J.; Weller, A.; and Moon, T. 2023. Towards more robust interpretation via local gradient alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 8168–8176.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Le, B. M.; and Woo, S. S. 2024. Gradient alignment for cross-domain face anti-spoofing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 188–199.
- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. 2018. Learning to generalize: Meta-learning for domain generalization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Liang, J.; He, R.; and Tan, T. 2024. A comprehensive survey on test-time adaptation under distribution shifts. *International Journal of Computer Vision*, 1–34.
- Liu, L.; Cheng, J.; Quan, Q.; Wu, F.-X.; Wang, Y.-P.; and Wang, J. 2020. A survey on U-shaped networks in medical image segmentations. *Neurocomputing*, 409: 244–258.
- 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/CVF Conference on Computer Vision and Pattern Recognition*, 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*.
- Nguyen, A. T.; Nguyen-Tang, T.; Lim, S.-N.; and Torr, P. H. 2023. TIPI: Test Time Adaptation with Transformation Invariance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24162–24171.
- 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 *International Conference on Learning Representations*.
- Orlando, J. I.; Fu, H.; Breda, J. B.; Van Keer, K.; Bathula, D. R.; Diaz-Pinto, A.; Fang, R.; Heng, P.-A.; Kim, J.; Lee, J.; et al. 2020. REFUGE challenge: A unified framework for evaluating automated methods for glaucoma assessment from fundus photographs. *Medical Image Analysis*, 59: 101570.
- Park, H.; Hwang, J.; Mun, S.; Park, S.; and Ok, J. 2024. MedBN: Robust Test-Time Adaptation against Malicious Test Samples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5997–6007.

- Salpea, N.; Tzouveli, P.; and Kollias, D. 2022. Medical image segmentation: A review of modern architectures. In *European Conference on Computer Vision*, 691–708. Springer.
- Sinha, S.; Gehler, P.; Locatello, F.; and Schiele, B. 2023. TeST: Test-time Self-Training under Distribution Shift. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2759–2769.
- Sivaswamy, J.; Krishnadas, S.; Joshi, G. D.; Jain, M.; and Tabish, A. U. S. 2014. Drishti-GS: Retinal image dataset for optic nerve head (onh) segmentation. In *2014 IEEE 11th International Symposium on Biomedical Imaging (ISBI)*, 53–56. IEEE.
- 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 *International Conference on Machine Learning*, 9229–9248. PMLR.
- Tajbakhsh, N.; Jeyaseelan, L.; Li, Q.; Chiang, J. N.; Wu, Z.; and Ding, X. 2020. Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation. *Medical Image Analysis*, 63: 101693.
- Tang, F.; Xu, Z.; Huang, Q.; Wang, J.; Hou, X.; Su, J.; and Liu, J. 2023. DuAT: Dual-aggregation transformer network for medical image segmentation. In *PRCV*.
- Wan, R.; Zhu, Z.; Zhang, X.; and Sun, J. 2021. Spherical motion dynamics: Learning dynamics of normalized neural network using sgd and weight decay. *Advances in Neural Information Processing Systems*, 34: 6380–6391.
- Wang, D.; Shelhamer, E.; Liu, S.; Olshausen, B.; and Darrell, T. 2021. Tent: Fully test-time adaptation by entropy minimization. In *International Conference on Learning Representations*.
- Wang, J.; Huang, Q.; Tang, F.; Meng, J.; Su, J.; and Song, S. 2022. Stepwise feature fusion: Local guides global. In *MICCAI*. Springer.
- Wang, P.; Zhang, Z.; Lei, Z.; and Zhang, L. 2023a. Sharpness-aware gradient matching for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3769–3778.
- Wang, W.; Zhong, Z.; Wang, W.; Chen, X.; Ling, C.; Wang, B.; and Sebe, N. 2023b. Dynamically Instance-Guided Adaptation: A Backward-Free Approach for Test-Time Domain Adaptive Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24090–24099.
- Wen, R.; Yuan, H.; Ni, D.; Xiao, W.; and Wu, Y. 2024. From Denoising Training to Test-Time Adaptation: Enhancing Domain Generalization for Medical Image Segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 464–474.
- Xiao, N.; and Zhang, L. 2021. Dynamic weighted learning for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15242–15251.
- Xiong, Y.; Lan, L.-C.; Chen, X.; Wang, R.; and Hsieh, C.-J. 2022. Learning to schedule learning rate with graph neural networks. In *International conference on learning representation (ICLR)*.
- Xu, Z.; Tang, F.; Chen, Z.; Zhou, Z.; Wu, W.; Yang, Y.; Liang, Y.; Jiang, J.; Cai, X.; and Su, J. 2024. Polyp-Mamba: Polyp Segmentation with Visual Mamba. In *MICCAI*. Springer.
- Yang, H.; Chen, C.; Jiang, M.; Liu, Q.; Cao, J.; Heng, P. A.; and Dou, Q. 2022. DLTTA: Dynamic Learning Rate for Test-time Adaptation on Cross-domain Medical Images. *IEEE Transactions on Medical Imaging*, 41(12): 3575–3586.
- Yu, T.; Kumar, S.; Gupta, A.; Levine, S.; Hausman, K.; and Finn, C. 2020. Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems*, 33: 5824–5836.
- Yu, X.-H.; Chen, G.-A.; and Cheng, S.-X. 1995. Dynamic learning rate optimization of the backpropagation algorithm. *IEEE Transactions on Neural Networks*, 6(3): 669–677.
- Zhang, J.; Qi, L.; Shi, Y.; and Gao, Y. 2023. Domainadaptor: A novel approach to test-time adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 18971–18981.
- Zhang, M.; Levine, S.; and Finn, C. 2022. Memo: Test time robustness via adaptation and augmentation. *Advances in neural information processing systems*, 35: 38629–38642.
- Zhang, Z.; Yin, F. S.; Liu, J.; Wong, W. K.; Tan, N. M.; Lee, B. H.; Cheng, J.; and Wong, T. Y. 2010. ORIGA-light: An online retinal fundus image database for glaucoma analysis and research. In *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*, 3065–3068. IEEE.
- Zhao, B.; Chen, C.; Wang, Q.-W.; He, A.; and Xia, S.-T. 2023. Combating unknown bias with effective bias-conflicting scoring and gradient alignment. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 3561–3569.
- Zhao, F.; Wang, W.; Liao, S.; and Shao, L. 2021. Learning anchored unsigned distance functions with gradient direction alignment for single-view garment reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12674–12683.