

Each Test Image Deserves A Specific Prompt: Continual Test-Time Adaptation for 2D Medical Image Segmentation

Ziyang Chen¹ Yongsheng Pan^{1†} Yiwen Ye¹ Mengkang Lu¹ Yong Xia^{1,2,3†}

¹ School of Computer Science and Engineering, Northwestern Polytechnical University, China

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

³ Ningbo Institute of Northwestern Polytechnical University, China

zychen@mail.nwpu.edu.cn, yspan@nwpu.edu.cn, ywye@mail.nwpu.edu.cn

lmk@mail.nwpu.edu.cn, yxia@nwpu.edu.cn

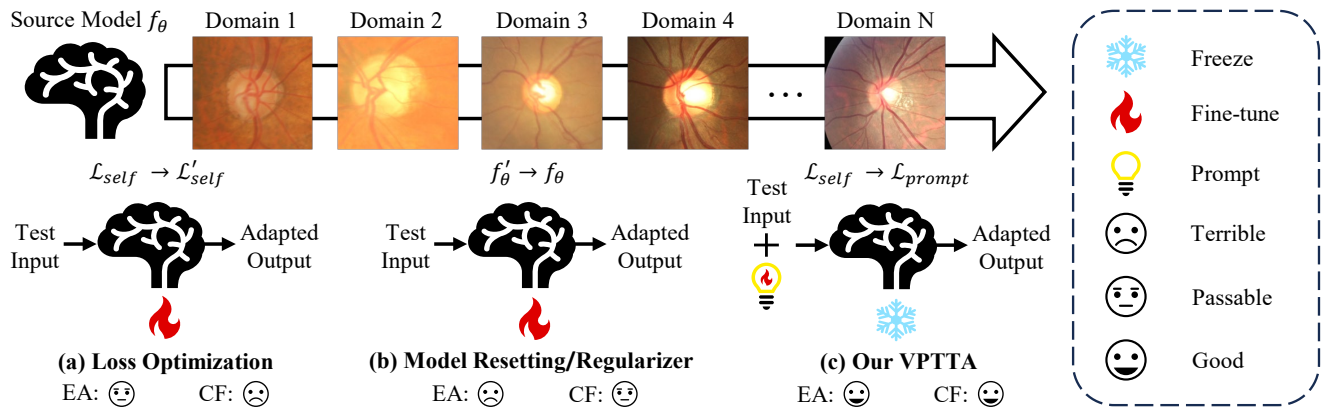


Figure 1. Comparison between our VPTTA and existing solutions under the CTTA setup. Our VPTTA avoids both error accumulation (EA) and catastrophic forgetting (CF) by freezing model parameters and achieves adaptation by training a prompt for each test image. The commonly used self-supervised loss and the optimized one are denoted by \mathcal{L}_{self} and \mathcal{L}'_{self} , respectively.

Abstract

Distribution shift widely exists in medical images acquired from different medical centres and poses a significant obstacle to deploying the pre-trained semantic segmentation model in real-world applications. Test-time adaptation has proven its effectiveness in tackling the cross-domain distribution shift during inference. However, most existing methods achieve adaptation by updating the pre-trained models, rendering them susceptible to error accumulation and catastrophic forgetting when encountering a series of distribution shifts (i.e., under the continual test-time adaptation setup). To overcome these challenges caused by updating the models, in this paper, we freeze the pre-trained model and propose the **Visual Prompt-based Test-Time Adaptation (VPTTA)** method to train a specific prompt for each test image to align the statistics in the batch normalization layers. Specifically, we present the low-frequency prompt, which is lightweight with only a few parameters

and can be effectively trained in a single iteration. To enhance prompt initialization, we equip VPTTA with a memory bank to benefit the current prompt from previous ones. Additionally, we design a warm-up mechanism, which mixes source and target statistics to construct warm-up statistics, thereby facilitating the training process. Extensive experiments demonstrate the superiority of our VPTTA over other state-of-the-art methods on two medical image segmentation benchmark tasks. The code and weights of pre-trained source models are available at <https://github.com/Chen-Ziyang/VPTTA>.

[†]Yong Xia and Yongsheng Pan are the corresponding authors. This work was supported in part by the National Natural Science Foundation of China under Grants 62171377, in part by STI2030-Major Project under Grants 2022ZD0213100, in part by Shenzhen Science and Technology Program under Grants JCYJ20220530161616036, and in part by the Ningbo Clinical Research Center for Medical Imaging under Grant 2021L003 (Open Project 2022LYKFZD06).

1. Introduction

Semantic segmentation serves a pivotal role in medical image processing to outline specific anatomical structures. However, the segmentation model pre-trained on the source dataset may be hindered from deploying in real-world applications due to the commonly existing distribution shift, which is typically caused by the variations in imaging protocols, operators, and scanners [16, 34]. Test-time adaptation (TTA) has emerged as a promising paradigm to replace domain adaptation (DA) [5, 43, 46, 53], as TTA requires only test data during the inference phase, while DA necessitates a sufficiently large and representative target dataset [45]. The mainstream TTA methods [29–31, 37, 40, 41, 47, 48] leverage the self-supervised tasks to construct auxiliary losses, such as entropy minimization [41] and self-training [37], to update the parameters of pre-trained models. Although these TTA methods successfully alleviate the domain gap and performance degradation, they are based on the assumption that the target domains exhibit static distributions. Unfortunately, the target domains are continuously changing rather than static in most real-world scenarios.

This was first presented as the continual test-time adaptation (CTTA) setup [42] that addresses TTA on a series of distribution shifts. Under this setup, pre-trained models are required to adapt to several distinct domains and therefore are more prone to suffer error accumulation and catastrophic forgetting. As mentioned above, many TTA methods rely on self-supervised tasks and may be affected by the noisy losses due to unreliable supervision, leading to error accumulation [2, 23]. Catastrophic forgetting typically occurs when the model is continually trained on new domains, resulting in unexpected performance degradation [21, 30]. Some research has been studied in addressing these two challenges based on optimizing the loss functions for better supervision [6, 31, 42, 42, 45] (see Figure 1 (a)) or prevent the model parameters from excessive alterations by model resetting or introducing regularizers [30, 31, 39, 42] (see Figure 1 (b)). Despite their efforts, loss optimization cannot ensure a robust training process to avoid catastrophic forgetting during the long-term inference, insufficient or excessive updating of the model parameters may lead to inadequate adaptation or error accumulation. Since the pre-trained model is updated by the self-supervised loss where the noise cannot be completely removed, these methods may still suffer from these two issues.

Considering the above limitations in CTTA, updating the pre-trained model during inference seems to be inappropriate. To this end, we propose the **Visual Prompt-based Test-Time Adaptation (VPTTA)** method to avoid error accumulation and catastrophic forgetting associated with updating model by freezing the model parameters and conduct adaptation by learning a specific visual prompt for each test image (see Figure 1 (c)). The prompt can adapt each test im-

age to the frozen pre-trained model to take into account both adaptation ability and knowledge retaining. We attempt to achieve this by answering the following three questions.

- How to design the prompt? Each prompt is restricted to a limited number of training iterations (sometimes just one) to satisfy the demands of online inference, requiring the prompt to be lightweight. Based on this principle, we introduce the prompt for only the low-frequency components of each sample, resulting in the low-frequency prompt. Since the low-frequency components are strongly associated with style textures and can serve as the primary source of distribution shift [46], our low-frequency prompt allows for maintaining lightweight and being able to reduce distribution shifts by modifying the low-frequency components.
- How to initialize the prompt? It is well known that model training can benefit from proper parameter initialization, resulting in not only fewer training iterations but also improved performance [3, 52]. We argue that this is also applicable to prompt training and therefore construct a memory bank to store the prompts and low-frequency components of previous test images. By comparing the similarity between the low-frequency components of the current and previous images, we select several historical prompts of the most similar images to initialize the current prompt.
- How to train the prompt? Since the statistics mismatch between the source and target domains in the batch normalization (BN) layers is a significant cause of distribution shift [33, 35, 44], we optimize the prompt with the objective of minimizing the distances between BN statistics and the statistics extracted from the features of test data for statistics alignment [19]. To further address the training difficulty posed by the large distribution shift at the beginning of the inference phase, we devise a statistics-fusion-based warm-up mechanism to assist in learning the prompt from easy to hard.

To summarise, our contributions are three-fold: (1) We present the low-frequency prompt to adjust the test image to adapt the model instead of updating the model and therefore prevent the model from error accumulation and catastrophic forgetting. (2) We propose the novel VPTTA framework to alleviate the distribution shift by training a sample-specific visual prompt, which is lightweight, well-initialized, and follows a gradual learning process. (3) The powerful framework we provide can be effectively deployed, requiring only a single iteration and a few dozen parameters for adaptation.

2. Related Work

2.1. Test-time Adaptation (TTA)

TTA aims to adapt the model pre-trained on the source domain to the test data during inference in a source-free

and online manner [40, 41]. The mainstream TTA methods are based on model updating, which constructs self-supervised auxiliary tasks to guide the model to learn on the test data [4, 13, 26, 29, 37, 40, 41, 48, 49]. Sun *et al.* [40] utilized an auxiliary task to predict the rotation to assist the model learning on the test data. Wang *et al.* [41] presented a test-time entropy minimization scheme to reduce generalization error by reducing the entropy of model predictions on test data. Zhang *et al.* [49] introduced a supplementary network to fit the test data, which can work in a plug-and-play fashion, necessitating minimal hyperparameter tuning. Some other methods are based on BN statistics, which alleviate the distribution shifts by modifying the statistics within BN layers [25, 27, 44]. Mirza *et al.* [25] constantly adapted the statistics of the BN layers to modify the feature representations of the model. These TTA methods improve the adaptation performance on the static target domain but neglect the continually changing target domains in most real-world scenarios.

Recently, the continual test-time adaptation (CTTA) [42] is proposed to tackle TTA under a series of continually changing target domains. Due to long-term inference, both error accumulation [2, 23] and catastrophic forgetting [21, 30] are more prone to occur in this setup. Many methods [6, 30, 31, 39, 42, 45, 47] optimize the self-supervised losses or prevent excessive alterations to the model parameters via model resetting or regularizer to improve this situation but are still impacted by these issues. Wang *et al.* [42] refined pseudo labels by the weight- and augmentation-averaged predictions and randomly restored a small part of model parameters. Niu *et al.* [31] removed partial noisy samples with large gradients and encouraged model weights toward a flat minimum. Yang *et al.* [45] designed a dynamic strategy to adjust the learning rate for each test image. Zhang *et al.* [47] presented to adaptively fuse source and test statistics and improve the entropy minimization loss to a generalized one. In contrast, our proposed VPPTA employs a learnable prompt for each test image to mitigate distribution shifts during inference, requiring no changes to the pre-trained model and therefore alleviating the error accumulation and catastrophic forgetting.

2.2. Prompt Learning

Prompt learning was initially employed to devise additional text instructions for input text to fine-tune the large-scale Natural Language Processing (NLP) models on downstream tasks [24]. With remarkable success in NLP, it has also been involved in computer vision and greatly benefited visual models in recent years [12, 15, 19, 51]. Jia *et al.* [15] introduced a small number of trainable parameters in the input space while keeping the model backbone frozen to fine-tune the model with high quality and efficiency. Zhou *et al.* [51] presented the conditional context optimization

to generate an input-conditional token for each image to generalize the pre-trained vision-language model to unseen classes. Hu *et al.* [19] employed the visual prompt at the pixel level and added it to the input images to explicitly minimize the domain discrepancy. Gan *et al.* [12] learned a domain-specific and domain-agnostic visual prompt for adaptation by using the exponential-moving-average strategy under the guidance of self-training. Different from them, our VPPTA presents the low-frequency prompt to address a series of changing distribution shifts with a few dozen parameters. We further consider optimizing the initialization and warm-up mechanism, enabling our visual prompts to be effectively trained within a single iteration to satisfy the demands of CTTA.

3. Methodology

3.1. Problem Definition and Method Overview

Let the labeled source domain dataset and unlabeled testing target dataset be $\mathbb{D}^s = \{X_i^s, Y_i^s\}_{i=1}^{N^s}$ and $\mathbb{D}^t = \{X_i^t\}_{i=1}^{N^t}$, respectively, where $X_i^* \in \mathbb{R}^{H \times W \times C}$ is the i -th image with C channels and size of $H \times W$, and Y_i^* is its label. Our goal is to learn a low-frequency prompt \mathcal{P}_i for each test image X_i^t to obtain the modified image \tilde{X}_i^t that can be adapted to the model $f_\theta : X \rightarrow Y$ pre-trained on \mathbb{D}^s .

To accomplish this goal, our VPPTA considers three key facets: prompt design, prompt initialization, and prompt training. For prompt design, the prompt is designed to change the low-frequency component of the amplitude of the test image. For prompt initialization, we adopt a memory bank to initialize the current prompt with previous ones. For prompt training, we employ the absolute distance to train the prompt to align the statistics between features extracted from test images and the batch normalization (BN) layers and also present a statistics-fusion-based warm-up mechanism. An overview of our VPPTA is illustrated in Figure 2. We now delve into its details.

3.2. Prompt Design: Low-frequency Prompt

Let $\mathcal{F}(\cdot)$ and $\mathcal{F}^{-1}(\cdot)$ be the Fast Fourier Transform (FFT) and Inverse Fast Fourier Transform (IFFT) operations [10], respectively. The amplitude component and phase component are denoted by $\mathcal{F}^A(\cdot)$ and $\mathcal{F}^P(\cdot)$ respectively. Note that the low-frequency component in $\mathcal{F}^A(\cdot)$ is shifted to the center. Utilizing the prompt $\mathcal{P}_i \in \mathbb{R}^{(\alpha \times H) \times (\alpha \times W) \times C}$ to obtain the adapted image \tilde{X}_i^t can be formalized as

$$\tilde{X}_i^t = \mathcal{F}^{-1}([\text{OnePad}(\mathcal{P}_i) \odot \mathcal{F}^A(X_i^t), \mathcal{F}^P(X_i^t)]), \quad (1)$$

where \odot is the element-wise multiplication, and *OnePad* means padding one around \mathcal{P}_i to expand it to the size of $H \times W$. Under the control of the small coefficient $\alpha \in (0, 1)$, we can ensure that the prompt \mathcal{P}_i remains lightweight and focuses on the low-frequency component.

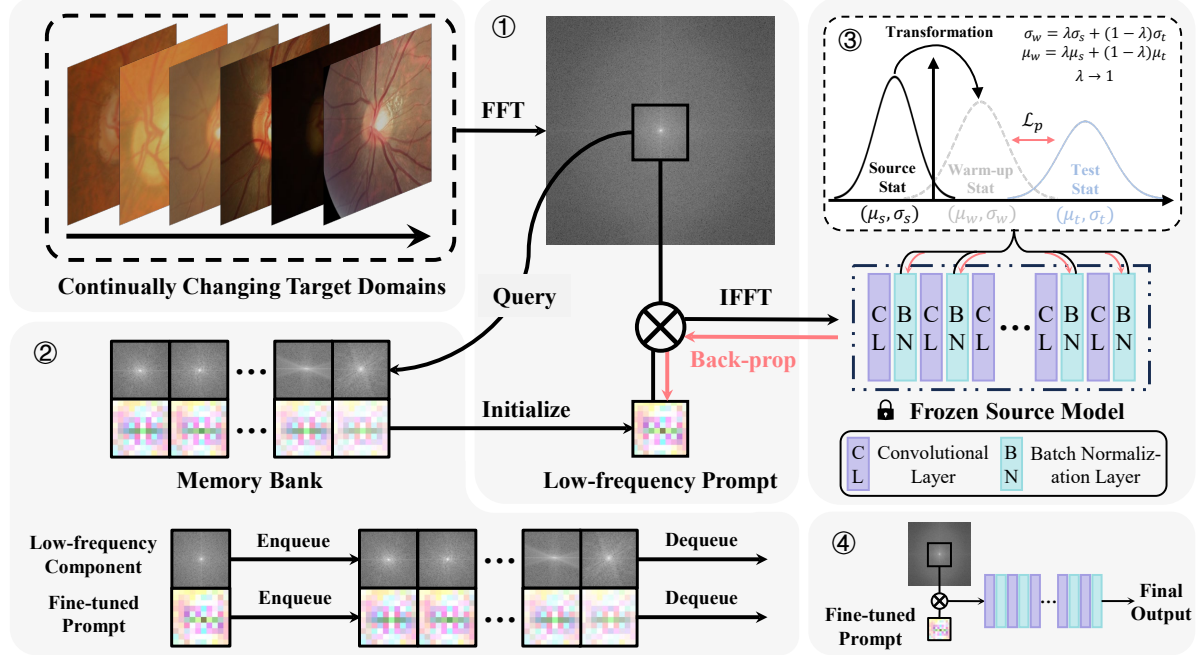


Figure 2. Overview of our VPTTA. For each test image, (1) the Fast Fourier Transform (FFT) is first applied to transform it into the frequency domain, where the low-frequency component of amplitude is used to query in a memory bank to initialize the current prompt, and then the amplitude is multiplied with the prompt at the low-frequency component and transformed back to the spatial domain using the Inverse Fast Fourier Transform (IFFT). (2) The memory bank is built on the previous low-frequency components and their corresponding prompts and is updated using the First In First Out (FIFO) strategy. (3) We convert the source statistics stored in BN layers into the warm-up statistics and calculate the absolute distance between warm-up and target statistics as the loss to fine-tune the prompt. (4) Finally, we feed the image corrected by its fine-tuned prompt to the pre-trained model to produce the output. ‘Stat’: Abbreviation of ‘Statistics’.

Simultaneously, we recognize that the low-rank decomposition is widely employed in the field of data compression [22] and model fine-tuning [18]. Consequently, we also design a low-rank prompt $\mathcal{P}_i^{LoRa} = \mathcal{B}_i @ \mathcal{A}_i$ as a variant (denoted as VPTTA-LR), where $\mathcal{B}_i \in \mathbb{R}^{H \times r \times C}$ ($r \ll W$), $\mathcal{A}_i \in \mathbb{R}^{r \times W \times C}$ ($r \ll H$), and @ denotes the matrix multiplication operation. In this way, the adapted image can be formalized as $\hat{X}_i^t = X_i^t + \mathcal{P}_i^{LoRa}$.

3.3. Prompt Initialization: Memory Bank

To achieve efficient training, it is critical to produce a suitable initialization for \mathcal{P}_i . Therefore, we present a memory-bank-based initialization strategy to benefit the current prompt from previous ones. Specifically, the memory bank M stores S pairs of keys and values $\{q_s, v_s\}_{s=1}^S$ and holds the First In First Out (FIFO) principle. As shown in Figure 2, the keys are the low-frequency components of previous test images, and the values are their corresponding prompts. Note that when using VPTTA-LR, the keys are the low-rank reconstruction matrix of previous test images. To initialize the current prompt \mathcal{P}_i , we first extract the low-frequency component of the current test image $\mathcal{F}_{low}^A(X_i^t)$ and calculate the cosine similarity between $\mathcal{F}_{low}^A(X_i^t)$ and

each key q_s as follows:

$$\text{Cos}(\mathcal{F}_{low}^A(X_i^t), q_s) = \frac{\langle \mathcal{F}_{low}^A(X_i^t), q_s \rangle}{\|\mathcal{F}_{low}^A(X_i^t)\| * \|q_s\|}, \quad (2)$$

where $\langle \cdot, \cdot \rangle$ denotes the dot product operation, and $\|\cdot\|$ denotes the Euclidean norm. Since the low-frequency component can represent the style texture of an image, we rank the similarity scores and then retrieve the values of the K most relevant keys from M to construct a support set $R_i = \{q_k, v_k\}_{k=1}^K$ for the current test image. Finally, we assign a weight w_k to v_k to initialize the current prompt by $\mathcal{P}_i = \sum_{k=1}^K w_k v_k$, where the weights $\{w_k\}$ are calculated based on the similarity scores [45] and $\sum_{k=1}^K w_k = 1$.

3.4. Prompt Training: Statistics Alignment

Inspired by [19, 27, 47], the distribution shift largely exists in distinct BN statistics calculated on the test features (i.e., μ_t and σ_t) and stored in the source model f_θ (i.e., μ_s and σ_s), where μ and σ respectively denote the mean and standard deviation. To train the prompts to align the statistics, we adopt the absolute-distance-based loss \mathcal{L}_p [19] as follows:

$$\mathcal{L}_p = \frac{1}{j} \sum |\mu_s^j - \mu_t^j| + |\sigma_s^j - \sigma_t^j|, \quad (3)$$

Algorithm 1: VPTTA algorithm.

Initialize : Frozen source model f_θ , initial prompt \mathcal{P} filled with ones, memory bank M , hyperparameters S , K and τ .

Input: For each time step i , current test image X_i^t .

- 1: \triangleright **Initialize the prompt \mathcal{P}_i :**
- 2: **if** $i < K$ **then**
- 3: $\mathcal{P}_i = \mathcal{P}$
- 4: **else**
- 5: Construct the support set R_i from M
- 6: $\mathcal{P}_i = \sum_{k=1}^K w_k v_k, v_k \in R_i$
- 7: **end if**
- 8: \triangleright **Train the prompt \mathcal{P}_i :**
- 9: Obtain \tilde{X}_i^t using \mathcal{P}_i by Eq. (1)
- 10: Forward $f_\theta(\tilde{X}_i^t)$ and calculate the warm-up statistics μ_w and σ_w by Eq. (4) to replace the μ_s and σ_s stored in f_θ
- 11: Backward and Update \mathcal{P}_i by Eq. (5)
- 12: \triangleright **Inference:**
- 13: Obtain \tilde{X}_i^t using updated \mathcal{P}_i by Eq. (1)
- 14: Forward $O_i = f_\theta(\tilde{X}_i^t)$
- 15: \triangleright **Update the memory bank M :**
- 16: $M.enqueue(q_i, v_i), q_i = \mathcal{F}^A(X_i^t), v_i = \mathcal{P}_i$
- 17: **if** $len(M) > S$ **then**
- 18: $M.dequeue(q_{first}, v_{first})$
- 19: **end if**

Output: Adapted prediction O_i

where j indicates the j -th BN layer in f_θ .

However, aligning the source and target domains within a limited number of iterations (one in this study) provides challenges due to the empty memory bank and undiminished distribution shifts at the beginning of the inference phase. Therefore, we design a warm-up mechanism to address this issue by simulating the warm-up statistics (*i.e.*, μ_w and σ_w) as follows:

$$\mu_w = \lambda\mu_t + (1 - \lambda)\mu_s, \sigma_w = \lambda\sigma_t + (1 - \lambda)\sigma_s, \quad (4)$$

where $\lambda = \frac{1}{\sqrt{i/\tau+1}}$, i indicates the index of the current test image and starts from 1, and τ is a temperature coefficient to control the rate of transition from the warm-up statistics to the source statistics. Now, the Eq. (3) can be rewritten by:

$$\mathcal{L}_p = \frac{1}{j} \sum |\mu_w^j - \mu_t^j| + |\sigma_w^j - \sigma_t^j|, \quad (5)$$

and we utilize μ_w and σ_w for normalization instead of μ_s and σ_s . As the inference progresses, the warm-up statistics will gradually shift towards the source statistics. By simulating the warm-up statistics between the source and target statistics, we can significantly reduce the training difficulty at the beginning of the inference phase, resulting in learning prompts from easy to difficult. The overall process of VPTTA is summarized in Algorithm 1.

4. Experiments

We evaluate our proposed VPTTA on two 2D segmentation benchmark tasks: the joint optic disc (OD) and cup (OC) segmentation task, and the polyp segmentation task.

4.1. Datasets and Evaluation Metrics

The OD/OC segmentation dataset comprises five public datasets collected from different medical centres, denoted as domain A (RIM-ONE-r3 [11]), B (REFUGE [32]), C (ORIGA [50]), D (REFUGE-Validation/Test [32]), and E (Drishti-GS [38]). There are 159, 400, 650, 800, and 101 images from these datasets. We cropped a region of interest (ROI) centering at OD with size of 800×800 for each image following [19], and each ROI is further resized to 512×512 and normalized by min-max normalization. The Dice score metric (DSC) is utilized for evaluation in this task.

The polyp segmentation dataset consists of four public datasets collected from different medical centres, denoted as domain A (BKAI-IGH-NeoPolyp [28]), B (CVC-ClinicDB [1]), C (ETIS-LaribPolypDB [36]), and D (Kvasir-Seg [20]), which have 1000, 612, 196, and 1000 images, respectively. We followed [9] to resize each image to 352×352 and normalize the resized image by the statistics computed on ImageNet. The DSC , enhanced-alignment metric (E_ϕ^{max}) [8], and structural similarity metric (S_α) [7] are utilized for evaluation in this task.

4.2. Implementation Details

For each task, we trained the source model on each single domain (source domain) and tested it on the left domains (target domains) to calculate the mean metrics to evaluate the methods on different scenarios. In the source-training phase, we trained a ResUNet-34 [17] backbone, following [19], as the baseline for OD/OC segmentation task and trained the PraNet [9] with a Res2Net-based [14] backbone for polyp segmentation task. Due to the complexity of PraNet’s decoder, we only calculated the prompt loss \mathcal{L}_p for each BN layer in the encoder instead of the whole network. In the test-adaptation phase, We performed one-iteration adaptation for each batch of test data with a batch size of 1 on all experiments of our VPTTA and other competing methods [45]. As our VPTTA involves an iteration to update and an inference step, distinguishing it from other methods like [41] that do not require a dedicated inference step, we applied the same configuration as VPTTA to other competing methods to ensure consistency in the experimental setup. To deploy our VPTTA, we utilized the Adam optimizer with a learning rate of 0.05 and 0.01 for the OD/OC segmentation task and polyp segmentation task, respectively. The hyperparameters α (size of prompt), S (size of memory bank), K (size of support set), and τ (temperature coefficient in warm-up) are set to 0.01, 40, 16, and 5 for both segmentation tasks.

Table 1. Performance of our VPTTA, ‘Source Only’ baseline, and six competing methods on the OD/OC segmentation task. The best and second-best results in each column are highlighted in **bold** and underline, respectively.

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> ↑
Source Only (ResUNet-34)	64.53	76.06	71.18	52.67	64.87	65.86
TENT-continual (ICLR 2021) [41]	73.07	78.66	71.94	46.81	70.20	68.13
CoTTA (CVPR 2022) [42]	75.39	75.98	69.14	53.99	70.40	68.98
DLTTA (TMI 2022) [45]	<u>75.11</u>	<u>78.85</u>	<u>73.89</u>	51.64	69.71	69.84
DUA (CVPR 2022) [25]	72.28	76.59	70.13	56.17	71.38	69.31
SAR (ICLR 2023) [31]	74.55	77.71	70.78	55.40	<u>71.72</u>	<u>70.03</u>
DomainAdaptor (CVPR 2023) [47]	74.50	76.39	71.81	56.78	70.55	70.01
VPTTA (Ours)	73.91	79.36	74.51	<u>56.51</u>	75.35	71.93

Table 2. Performance of our VPTTA, ‘Source Only’ baseline, and six competing methods on the polyp segmentation task. The best and second-best results in each column are highlighted in **bold** and underline, respectively.

Methods	Domain A			Domain B			Domain C			Domain D			Average		
	<i>DSC</i>	E_{ϕ}^{max}	S_{α}	<i>DSC</i>	E_{ϕ}^{max}	S_{α}	<i>DSC</i>	E_{ϕ}^{max}	S_{α}	<i>DSC</i>	E_{ϕ}^{max}	S_{α}	<i>DSC</i> ↑	E_{ϕ}^{max} ↑	S_{α} ↑
Source Only (PraNet)	<u>79.90</u>	<u>87.97</u>	<u>84.66</u>	66.33	78.51	76.72	73.89	84.64	81.28	82.95	90.84	88.08	75.77	85.49	82.69
TENT-continual (ICLR 2021) [41]	74.86	84.58	80.52	67.51	78.66	78.05	17.79	40.04	53.30	73.55	83.38	82.41	58.43	71.67	73.57
CoTTA (CVPR 2022) [42]	76.46	85.37	82.56	66.77	76.75	79.17	71.39	83.42	80.18	70.71	79.81	82.54	71.33	81.34	81.11
DLTTA (TMI 2022) [45]	76.27	85.23	82.41	66.58	77.00	79.24	63.72	78.23	75.56	71.20	81.32	83.47	69.44	80.45	80.17
DUA (CVPR 2022) [25]	78.93	87.37	83.96	66.84	78.52	77.51	<u>76.53</u>	<u>86.45</u>	<u>83.05</u>	<u>86.24</u>	<u>93.23</u>	<u>89.82</u>	<u>77.13</u>	<u>86.39</u>	<u>83.58</u>
SAR (ICLR 2023) [31]	76.48	85.89	81.49	66.45	77.35	78.05	71.46	83.23	79.40	70.41	80.11	81.07	71.20	81.65	80.00
DomainAdaptor (CVPR 2023) [47]	77.48	86.31	82.40	<u>70.82</u>	<u>81.76</u>	<u>80.88</u>	71.96	83.06	79.97	76.89	85.89	84.45	74.29	84.26	81.93
VPTTA (Ours)	81.00	88.91	84.91	76.87	87.31	84.08	77.58	87.48	83.64	86.39	93.47	89.87	80.46	89.29	85.62

Table 3. Results of ablation study on the OD/OC segmentation task. The best and second-best results in each column are highlighted in **bold** and underline, respectively.

Methods			Domain A	Domain B	Domain C	Domain D	Domain E	Average
Low-frequency Prompt	Memory Bank	Warm-up	<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i>	<i>DSC</i> ↑
			64.53	76.06	71.18	52.67	64.87	65.86
✓			65.99	76.58	71.75	52.86	67.30	66.90
✓	✓		70.70	79.36	72.82	<u>55.50</u>	<u>72.84</u>	<u>70.24</u>
✓		✓	74.41	<u>77.47</u>	<u>74.14</u>	54.03	67.85	69.58
✓	✓	✓	<u>73.91</u>	79.36	74.51	56.51	75.35	71.93

4.3. Experimental Results

We compare our VPTTA with the ‘Source Only’ baseline (training on the source domain and testing without adaptation), and six competing methods, including a pseudo-label-based method (CoTTA [42]), two entropy-based methods (TENT-continual [41] and SAR [31]), a method dynamically adjusting the learning rate (DLTTA [45]), a method combining entropy-loss and BN statistics fusion (DomainAdaptor [47]), and a BN statistics-based method (DUA [25]) on two segmentation tasks. Specifically, we re-implemented all the competing methods using the same baseline as VPTTA and combined DLTTA with the entropy loss proposed in TENT [45]. The result of each domain is calculated by using the current domain as the source domain and applying the pre-trained model on the left domains.

Comparison on the OD/OC segmentation task. The results of the proposed VPTTA, the ‘Source Only’ baseline, and six competing methods on the OD/OC segmentation task are detailed in Table 1. Obviously, all competing methods outperform the ‘Source Only’ baseline, indicating the

effectiveness of adaptation towards the distribution shifts. It can be found that our VPTTA achieves the best overall performance across all domains, underscoring its superior applicability and robustness. Meanwhile, our VPTTA also shows competitive performance in domains where most methods exhibit satisfactory performance (such as D) and also performs well in domains where other methods exhibit suboptimal results (such as C).

Comparison on the polyp segmentation task. We conducted the comparison experiment on the polyp segmentation task under the same experimental setup, as shown in Table 2. In contrast to the results of the OD/OC segmentation task, the methods based on model updating (*i.e.*, TENT-continual, CoTTA, DLTTA, SAR, and DomainAdaptor) exhibit significantly worse performance in this task and are even inferior to the ‘Source Only’ baseline. It can be attributed to the fact that the polyps are relatively hidden, and the models tend to segment almost nothing rather than poorly segment under the impact of distribution shifts, resulting in confident and low entropy predictions but wrong

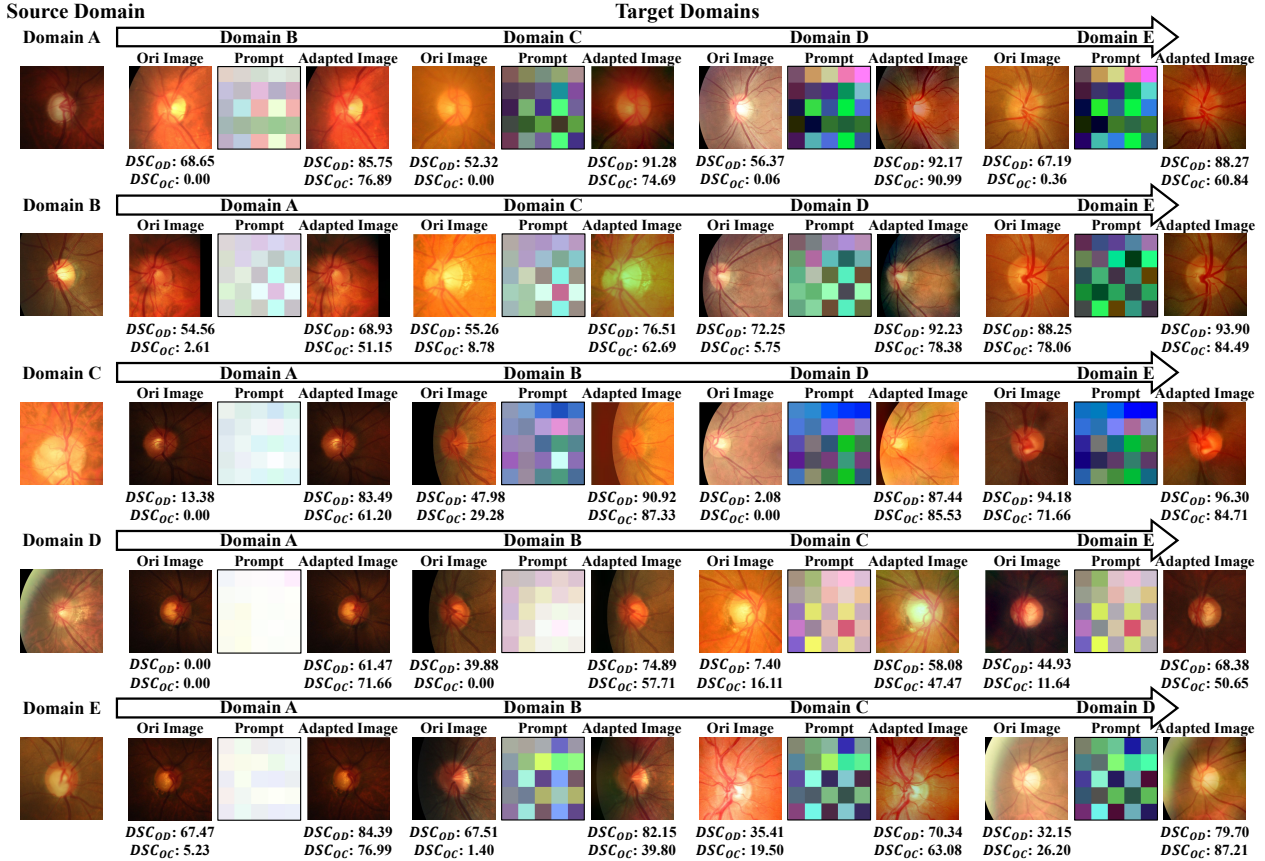


Figure 3. Visualization of the original images, estimated prompts, and adapted images on the OD/OC segmentation task. We normalize the prompts to [0, 1] for better visualization. The DSC of applying the frozen source model on the original and adapted images is displayed below each image. We also show an example of each source domain on the left side of this diagram. ‘Ori’: Abbreviation of ‘Original’.

gradients for these methods. Since DUA is backward-free, it still performs well in this task. However, DUA also fails when using domain A as the source domain and has limited performance due to the absence of training for the current test image. Our VPTTA simultaneously benefits from training on the current test image and refraining from updating the model, thus surpassing all competing methods and demonstrating advanced performance.

Visualization of prompts and adapted images. We visualized the estimated prompts and adapted images on the OD/OC segmentation task in Figure 3 and displayed the DSC values predicted on the original and adapted images below the corresponding images. It can be observed that: (1) Applying prompts to the original test images can achieve significant performance gain; (2) Images from different target domains have similar prompts, proving that involving the memory bank to initialize the current prompt using previous ones is reasonable; (3) The appearances of adapted images are close to the source domain, explicitly demonstrating the reduction of distribution shifts via our VPTTA.

Comparison in a long-term continual test-time adap-

tation. To evaluate our VPTTA in a long-term continual test-time adaptation [42], we conducted the experiments for the OD/OC segmentation task multiple rounds and reported the results in Table 4. Since DUA is a backward-free method and DomainAdaptor resets the pre-trained model before each adaptation, they are unsuitable to be evaluated in this experiment. We compared our VPTTA with TENT, CoTTA, DLTTA, and SAR, all of which involve continuous training of the pre-trained model. The performance across different rounds with the same source domain reflects the extent of catastrophic forgetting in the model, and the average performance across multiple rounds highlights the error accumulation of the model within a series of distinct domains. The results reveal that both optimizing loss functions (CoTTA, DLTTA, SAR) and resetting the model (CoTTA, SAR) are effective for mitigating error accumulation and catastrophic forgetting, while TENT with vanilla entropy-minimization loss suffers from severe performance degradation. Additionally, after multiple rounds of testing, our VPTTA maintains superior performance with minimal performance degradation, emphasizing the effectiveness of

Table 4. Performance of our VPTTA and four competing methods on the OD/OC segmentation task in a long-term continual test-time adaptation. The red numbers indicate the performance gain relative to the ‘Source Only’ baseline. The performance degradation is calculated between the overall average DSC and the average DSC of round 1. ‘Ave’: Abbreviation of ‘Average’.

Time Round	→															Average DSC ↑	Perform. Degra. ↓			
Methods	1					2					3									
	A	B	C	D	E	Ave	A	B	C	D	E	Ave	A	B	C	D	E	Ave		
Source Only	64.53	76.06	71.18	52.67	64.87	65.86	64.53	76.06	71.18	52.67	64.87	65.86	64.53	76.06	71.18	52.67	64.87	65.86	65.86	-
TENT	73.07	78.66	71.94	46.81	70.20	68.13	62.09	69.32	70.67	39.02	68.22	61.86	57.05	62.47	70.20	39.02	66.37	59.02	63.01 (-2.85)	5.12
CoTTA	75.39	75.98	69.14	53.99	70.40	68.98	74.31	75.00	67.99	51.04	68.28	67.32	73.22	74.33	66.72	50.23	67.08	66.32	67.54 (+1.68)	1.44
DLTTA	75.11	78.85	73.89	51.64	69.71	69.84	74.14	79.65	74.25	45.05	69.04	68.43	72.28	78.93	72.87	42.37	69.26	67.14	68.47 (+2.61)	1.37
SAR	74.55	77.71	70.78	55.40	71.72	70.03	74.74	78.09	71.00	52.13	69.02	69.00	74.90	78.24	71.18	50.16	68.44	68.58	69.20 (+3.34)	0.83
VPTTA (Ours)	73.91	79.36	74.51	56.51	75.35	71.93	73.57	78.84	73.61	56.91	74.80	71.55	73.12	78.45	72.63	57.11	74.04	71.07	71.51 (+5.65)	0.42

training prompts over updating the pre-trained model.

Ablation study. To evaluate the contributions of the low-frequency prompt, memory-bank-based initialization, and warm-up mechanism, we conducted ablation studies on the OD/OC segmentation task, as shown in Table 3. Note that Eq. (3)/(5) is used as the loss function when training the prompt without/with our warm-up mechanism. The results show that (1) using the low-frequency prompt only has limited performance gain due to the poorly trained prompt; (2) both the memory bank and warm-up mechanism are beneficial to training the prompt; (3) the best performance is achieved when the memory bank and warm-up mechanism are jointly used, which means they are complementary.




Figure 4. Performance of our VPTTA with various α on the OD/OC segmentation task.




Figure 5. Performance of our VPTTA with various S , K , and τ on the OD/OC segmentation task.

Analysis of hyperparameter α , S , K , and τ . We conducted the experiments for the OD/OC segmentation task to discuss the hyper-parameters used in our VPTTA and displayed the results in Figure 4 and Figure 5. It can be seen that $\alpha = 0.01$ is optimal, as an excessively small prompt makes adaptation challenging, while an overly large prompt hinders effective training. Similar trends are also reflected

Table 5. Performance of our VPTTA and VPTTA-LR on two segmentation tasks.

Methods	OD/OC		Polyp			
	DSC ↑	Param ↓	DSC ↑	E_{ϕ}^{max} ↑	S_{α} ↑	Param ↓
VPTTA-LR ($r=3$)	69.01	9,216	78.25	88.27	84.20	6,336
VPTTA	71.93	75	80.46	89.94	86.84	27

in S and K : as they increase, the DSC values improve, while the excessively large S and K may introduce unrelated samples, affecting the initialization. With the increase of τ , performance gradually approaches the upper bound and tends to remain stable, suggesting that the warm-up has reached saturation. Therefore, we finally chose $\alpha = 0.01$, $S = 40$, $K = 16$, and $\tau = 5$ as the best configuration.

Low-frequency prompt vs. low-rank prompt. To compare the effectiveness of VPTTA and VPTTA-LR, we repeated the experiments on two segmentation tasks and showed results in Table 5. It can be seen that VPTTA significantly outperforms VPTTA-LR, attributed to the fewer learnable parameters and superior representation of style texture in the frequency domain. Benefiting from the low-frequency prompt, VPTTA can be effectively trained with over a hundred times fewer parameters than VPTTA-LR.

5. Conclusion

In this paper, we proposed the VPTTA, a prompt-based method to alleviate the potential risks (*i.e.*, error accumulation, and catastrophic forgetting) by freezing the pre-trained model. Specially, we presented the low-frequency prompt that trains a lightweight prompt with only a small number of parameters in a single iteration to alleviate the distribution shifts in BN layers, constructed a memory bank to better initialize each prompt by utilizing the prompts of previous and the most similar test images, and designed the statistics-fusion-based warm-up mechanism to simulate the warm-up statistics to overcome the training difficulty at the beginning of the inference phase. The extensive experimental results on two medical image segmentation benchmark tasks with multiple domains verified the superiority of our VPTTA over other methods and the effectiveness of each special design. In this regard, we hope that our efforts can offer a different paradigm for continual test-time adaptation.

References

- [1] Jorge Bernal, F Javier Sánchez, Gloria Fernández-Esparrach, Debora Gil, Cristina Rodríguez, and Fernando Vilariño. Wm-dova maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians. *Comput. Med. Imaging and Graph.*, 43:99–111, 2015. [5](#)
- [2] Paola Cascante-Bonilla, Fuwen Tan, Yanjun Qi, and Vicente Ordonez. Curriculum labeling: Revisiting pseudo-labeling for semi-supervised learning. In *AAAI*, pages 6912–6920, 2021. [2](#), [3](#)
- [3] Oscar Chang, Lampros Flokas, and Hod Lipson. Principled weight initialization for hypernetworks. In *Int. Conf. Learn. Represent.*, 2019. [2](#)
- [4] Dian Chen, Dequan Wang, Trevor Darrell, and Sayna Ebrahimi. Contrastive test-time adaptation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 295–305, 2022. [3](#)
- [5] Lin Chen, Huaian Chen, Zhixiang Wei, Xin Jin, Xiao Tan, Yi Jin, and Enhong Chen. Reusing the task-specific classifier as a discriminator: Discriminator-free adversarial domain adaptation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7181–7190, 2022. [2](#)
- [6] Mario Döbler, Robert A Marsden, and Bin Yang. Robust mean teacher for continual and gradual test-time adaptation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7704–7714, 2023. [2](#), [3](#)
- [7] Deng-Ping Fan, Ming-Ming Cheng, Yun Liu, Tao Li, and Ali Borji. Structure-measure: A new way to evaluate foreground maps. In *Int. Conf. Comput. Vis.*, pages 4548–4557, 2017. [5](#)
- [8] Deng-Ping Fan, Cheng Gong, Yang Cao, Bo Ren, Ming-Ming Cheng, and Ali Borji. Enhanced-alignment measure for binary foreground map evaluation. In *IJCAI*, pages 698–704, 2018. [5](#)
- [9] Deng-Ping Fan, Ge-Peng Ji, Tao Zhou, Geng Chen, Huazhu Fu, Jianbing Shen, and Ling Shao. Pronet: Parallel reverse attention network for polyp segmentation. In *Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, pages 263–273. Springer, 2020. [5](#)
- [10] Matteo Frigo and Steven G Johnson. FFTW: An adaptive software architecture for the fft. In *ICASSP*, pages 1381–1384. IEEE, 1998. [3](#)
- [11] Francisco Fumero, Silvia Alayón, José L Sanchez, Jose Sigut, and M Gonzalez-Hernandez. RIM-ONE: An open retinal image database for optic nerve evaluation. In *Int. Symp. Comput.-based Med. Syst.*, pages 1–6. IEEE, 2011. [5](#)
- [12] Yulu Gan, Yan Bai, Yihang Lou, Xianzheng Ma, Renrui Zhang, Nian Shi, and Lin Luo. Decorate the newcomers: Visual domain prompt for continual test time adaptation. In *AAAI*, pages 7595–7603, 2023. [3](#)
- [13] Yossi Gandelsman, Yu Sun, Xinlei Chen, and Alexei Efros. Test-time training with masked autoencoders. *Adv. Neural Inform. Process. Syst.*, 35:29374–29385, 2022. [3](#)
- [14] Shang-Hua Gao, Ming-Ming Cheng, Kai Zhao, Xin-Yu Zhang, Ming-Hsuan Yang, and Philip Torr. Res2net: A new multi-scale backbone architecture. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(2):652–662, 2019. [5](#)
- [15] Yunhe Gao, Xingjian Shi, Yi Zhu, Hao Wang, Zhiqiang Tang, Xiong Zhou, Mu Li, and Dimitris N Metaxas. Visual prompt tuning for test-time domain adaptation. *arXiv preprint arXiv:2210.04831*, 2022. [3](#)
- [16] Mohsen Ghafoorian, Alireza Mehrtash, Tina Kapur, Nico Karssemeijer, Elena Marchiori, Mehran Pesteie, Charles RG Guttmann, Frank-Erik de Leeuw, Clare M Tempany, Bram Van Ginneken, et al. Transfer learning for domain adaptation in mri: Application in brain lesion segmentation. In *Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, pages 516–524. Springer, 2017. [2](#)
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–778, 2016. [5](#)
- [18] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *Int. Conf. Learn. Represent.*, 2022. [4](#)
- [19] Shishuai Hu, Zehui Liao, and Yong Xia. Prosfda: Prompt learning based source-free domain adaptation for medical image segmentation. *arXiv preprint arXiv:2211.11514*, 2022. [2](#), [3](#), [4](#), [5](#)
- [20] Debesh Jha, Pia H Smedsrud, Michael A Riegler, Pål Halvorsen, Thomas de Lange, Dag Johansen, and Håvard D Johansen. Kvasir-seg: A segmented polyp dataset. In *Int. Conf. Multimedia Modeling*, pages 451–462. Springer, 2020. [5](#)
- [21] Christiaan Lamers, René Vidal, Nabil Belbachir, Niki van Stein, Thomas Bäeck, and Paris Giampouras. Clustering-based domain-incremental learning. In *Int. Conf. Comput. Vis.*, pages 3384–3392, 2023. [2](#), [3](#)
- [22] Joonseok Lee, Seungyeon Kim, Guy Lebanon, and Yoram Singer. Local low-rank matrix approximation. In *Int. Conf. Mach. Learn.*, pages 82–90. PMLR, 2013. [4](#)
- [23] Jian Liang, Ran He, Zhenan Sun, and Tieniu Tan. Exploring uncertainty in pseudo-label guided unsupervised domain adaptation. *Pattern Recognition*, 96:106996, 2019. [2](#), [3](#)
- [24] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Comput. Surveys*, 55(9):1–35, 2023. [3](#)
- [25] M Jehanzeb Mirza, Jakub Micorek, Horst Possegger, and Horst Bischof. The norm must go on: Dynamic unsupervised domain adaptation by normalization. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 14765–14775, 2022. [3](#), [6](#)
- [26] Chaithanya Kumar Mummadi, Robin Huttmacher, Kilian Rambach, Evgeny Levinkov, Thomas Brox, and Jan Hendrik Metzen. Test-time adaptation to distribution shift by confidence maximization and input transformation. *arXiv preprint arXiv:2106.14999*, 2021. [3](#)
- [27] Zachary Nado, Shreyas Padhy, D Sculley, Alexander D’Amour, Balaji Lakshminarayanan, and Jasper Snoek. Evaluating prediction-time batch normalization for robustness under covariate shift. *arXiv preprint arXiv:2006.10963*, 2020. [3](#), [4](#)

- [28] Phan Ngoc Lan, Nguyen Sy An, Dao Viet Hang, Dao Van Long, Tran Quang Trung, Nguyen Thi Thuy, and Dinh Viet Sang. Neounet: Towards accurate colon polyp segmentation and neoplasm detection. In *Adv. Vis. Comput. Int. Symp.*, pages 15–28. Springer, 2021. [5](#)
- [29] A Tuan Nguyen, Thanh Nguyen-Tang, Ser-Nam Lim, and Philip HS Torr. Tipi: Test time adaptation with transformation invariance. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 24162–24171, 2023. [2, 3](#)
- [30] Shuaicheng Niu, Jiayang Wu, Yifan Zhang, Yaofu Chen, Shijian Zheng, Peilin Zhao, and Mingkui Tan. Efficient test-time model adaptation without forgetting. In *Int. Conf. Mach. Learn.*, pages 16888–16905. PMLR, 2022. [2, 3](#)
- [31] Shuaicheng Niu, Jiayang Wu, Yifan Zhang, Zhiquan Wen, Yaofu Chen, Peilin Zhao, and Mingkui Tan. Towards stable test-time adaptation in dynamic wild world. In *Int. Conf. Learn. Represent.*, 2023. [2, 3, 6](#)
- [32] José Ignacio Orlando, Huazhu Fu, João Barbosa Breda, Karel van Keer, Deepti R Bathula, Andrés Diaz-Pinto, Ruogu Fang, Pheng-Ann Heng, Jeyoung Kim, JoonHo Lee, et al. REFUGE Challenge: A unified framework for evaluating automated methods for glaucoma assessment from fundus photographs. *Med. Image Anal.*, 59:101570, 2020. [5](#)
- [33] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *Eur. Conf. Comput. Vis.*, pages 464–479, 2018. [2](#)
- [34] Swami Sankaranarayanan, Yogesh Balaji, Arpit Jain, Ser Nam Lim, and Rama Chellappa. Learning from synthetic data: Addressing domain shift for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3752–3761, 2018. [2](#)
- [35] Steffen Schneider, Evgenia Rusak, Luisa Eck, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. Improving robustness against common corruptions by covariate shift adaptation. *Adv. Neural Inform. Process. Syst.*, 33:11539–11551, 2020. [2](#)
- [36] Juan Silva, Aymeric Histace, Olivier Romain, Xavier Dray, and Bertrand Granado. Toward embedded detection of polyps in wce images for early diagnosis of colorectal cancer. *Int. J. Comput.-Assist. Rad. Sur.*, 9(2):283–293, 2014. [5](#)
- [37] Samarth Sinha, Peter Gehler, Francesco Locatello, and Bernt Schiele. Test: Test-time self-training under distribution shift. In *IEEE Wint. Conf. Appl. Comput. Vis.*, pages 2759–2769, 2023. [2, 3](#)
- [38] Jayanthi Sivaswamy, SR Krishnadas, Gopal Datt Joshi, Madhulika Jain, and A Ujjwaft Syed Tabish. Drishti-GS: Retinal image dataset for optic nerve head (onh) segmentation. In *IEEE Int. Symp. on Bio. Imaging*, pages 53–56. IEEE, 2014. [5](#)
- [39] Junha Song, Jungsoo Lee, In So Kweon, and Sungha Choi. Ecotta: Memory-efficient continual test-time adaptation via self-distilled regularization. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11920–11929, 2023. [2, 3](#)
- [40] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *Int. Conf. Mach. Learn.*, pages 9229–9248. PMLR, 2020. [2, 3](#)
- [41] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *Int. Conf. Learn. Represent.*, 2021. [2, 3, 5, 6](#)
- [42] Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual test-time domain adaptation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7201–7211, 2022. [2, 3, 6, 7](#)
- [43] Shujun Wang, Lequan Yu, Kang Li, Xin Yang, Chi-Wing Fu, and Pheng-Ann Heng. Boundary and entropy-driven adversarial learning for fundus image segmentation. In *Int. Conf. Med. Image Comput. Assist. Intervent.*, pages 102–110. Springer, 2019. [2](#)
- [44] Wei Wang, Zhun Zhong, Weijie Wang, Xi Chen, Charles Ling, Boyu Wang, and Nicu Sebe. Dynamically instance-guided adaptation: A backward-free approach for test-time domain adaptive semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 24090–24099, 2023. [2, 3](#)
- [45] Hongzheng Yang, Cheng Chen, Meirui Jiang, Quande Liu, Jianfeng Cao, Pheng Ann Heng, and Qi Dou. Dltta: Dynamic learning rate for test-time adaptation on cross-domain medical images. *IEEE Trans. Med. Imaging*, 41(12):3575–3586, 2022. [2, 3, 4, 5, 6](#)
- [46] Yanchao Yang and Stefano Soatto. FDA: Fourier domain adaptation for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4085–4095, 2020. [2](#)
- [47] Jian Zhang, Lei Qi, Yinghuan Shi, and Yang Gao. Domainadaptor: A novel approach to test-time adaptation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 18971–18981, 2023. [2, 3, 4, 6](#)
- [48] Marvin Zhang, Sergey Levine, and Chelsea Finn. Memo: Test time robustness via adaptation and augmentation. *Adv. Neural Inform. Process. Syst.*, 35:38629–38642, 2022. [2, 3](#)
- [49] Yizhe Zhang, Tao Zhou, Yuhui Tao, Shuo Wang, Ye Wu, Benyuan Liu, Pengfei Gu, Qiang Chen, and Danny Z Chen. TestFit: A plug-and-play one-pass test time method for medical image segmentation. *Medical Image Analysis*, 92:103069, 2024. [3](#)
- [50] Zhuo Zhang, Feng Shou Yin, Jiang Liu, Wing Kee Wong, Ngan Meng Tan, Beng Hai Lee, Jun Cheng, and Tien Yin Wong. ORIGA-light: An online retinal fundus image database for glaucoma analysis and research. In *Int. Conf. IEEE Eng. in Med. and Bio.*, pages 3065–3068. IEEE, 2010. [5](#)
- [51] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 16816–16825, 2022. [3](#)
- [52] Chen Zhu, Renkun Ni, Zheng Xu, Kezhi Kong, W Ronny Huang, and Tom Goldstein. Gradinit: Learning to initialize neural networks for stable and efficient training. *Adv. Neural Inform. Process. Syst.*, 34:16410–16422, 2021. [2](#)
- [53] Yukun Zuo, Hantao Yao, and Changsheng Xu. Attention-based multi-source domain adaptation. *IEEE Trans. Image Process.*, 30:3793–3803, 2021. [2](#)