



HAL
open science

Layerwise Early Stopping for Test Time Adaptation

Sabyasachi Sahoo, Mostafa Elaraby, Jonas Ngnawe, Yann Pequignot, Frédéric Precioso, Christian Gagné

► **To cite this version:**

Sabyasachi Sahoo, Mostafa Elaraby, Jonas Ngnawe, Yann Pequignot, Frédéric Precioso, et al.. Layerwise Early Stopping for Test Time Adaptation. 2024. hal-04533467

HAL Id: hal-04533467

<https://hal.science/hal-04533467v1>

Preprint submitted on 4 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Layerwise Early Stopping for Test Time Adaptation

Sabyasachi Sahoo^{1,2}, Mostafa ElAraby^{2,3}, Jonas Ngnawe^{1,2}, Yann Pequignot¹,
Frédéric Precioso⁴, and Christian Gagné^{1,2,5}

¹ IID, Université Laval

² Mila

³ Université de Montréal

⁴ Université Côte d’Azur, CNRS, INRIA, I3S, Maasai

⁵ Canada CIFAR AI Chair

sabyasachi.sahoo.1@ulaval.ca

Abstract. Test Time Adaptation (TTA) addresses the problem of distribution shift by enabling pretrained models to learn new features on an unseen domain at test time. However, it poses a significant challenge to maintain a balance between learning new features and retaining useful pretrained features. In this paper, we propose Layerwise EARly STOPping (LEAST) for TTA to address this problem. The key idea is to stop adapting individual layers during TTA if the features being learned do not appear beneficial for the new domain. For that purpose, we propose using a novel gradient-based metric to measure the relevance of the current learnt features to the new domain without the need for supervised labels. More specifically, we propose to use this metric to determine dynamically when to stop updating each layer during TTA. This enables a more balanced adaptation, restricted to layers benefiting from it, and only for a certain number of steps. Such an approach also has the added effect of limiting the forgetting of pretrained features useful for dealing with new domains. Through extensive experiments, we demonstrate that Layerwise Early Stopping improves the performance of existing TTA approaches across multiple datasets, domain shifts, model architectures, and TTA losses.

Keywords: test time adaptation · parameter efficient adaptation

1 Introduction

Distribution shifts [10, 31] are frequently encountered when a deep learning model is deployed in the real world. Test Time Adaptation (TTA) [13, 17, 25] has emerged as a potent strategy for adapting pre-trained models to new domains during testing. Current TTA strategies rely on techniques like layer selection [12, 15, 28, 29], sample selection [21, 34] or regularizers [6, 19, 21] to perform adaptation while limiting potential performance degradation that may come with it. Yet, designing these techniques is challenging [5, 35], and their effectiveness may not cover all shifts or TTA losses [35].




Fig. 1: Layerwise EARly STopping or LEAST framework leverages on a mask per layer and sample for a layerwise control of test time adaptation. As a result, it proceeds by early stopping the adaptation of individual layers at different points in time during testing. However, unlike classical early stopping, LEAST can resume adaptation on later samples for certain layers if the gradients are beneficial for the new domain.

In this paper, we propose Layerwise EARly STopping (LEAST) to halt individual layer adaptation during the TTA process to prevent overfitting to new domain samples. We demonstrate that LEAST outperforms both full model early stopping and selective layer adaptation without stopping [15], supporting our hypothesis that layers adapt independently, at varying rates. LEAST also serves as an effective method for online layer selection, balancing adaptation to the new domain and preventing overfitting.

In machine learning, early stopping typically requires a labeled validation set [9, 24]. However, we lack access to such external datasets or supervision in TTA [17, 25]. To circumvent this, we propose a novel cosine distance criterion based on gradient updates at each new sample for early stopping without a validation set. Gradient update variance increases as we approach the optimization minima [1, 8, 20]. Leveraging this, we measure the cosine distance between the gradient update on a given layer for the current sample input and the average gradient updates for all samples seen so far at the same layer. Rather than maintaining a running average of the gradients, we opt for tracking the total displacement of the current parameter from its initial state, as it offers a more comprehensive measure.

Our extensive experiments on Domainbed [10] demonstrate that LEAST consistently surpasses *all layers* and *ERM* (no adaptation) baselines across various datasets, neural network backbones, and losses. It also excels in layer selection in TTA, outperforming existing baselines. We further explore LEAST on a subset of critical shifts, termed as “Tiny Domainbed”. This approach proves superior to either *layer selection-only* or *early stopping-only* methods. Interestingly, the proposed cosine distance criterion in LEAST approximates the empirical upper bound of layerwise early stopping effectively, even without a validation set. Our

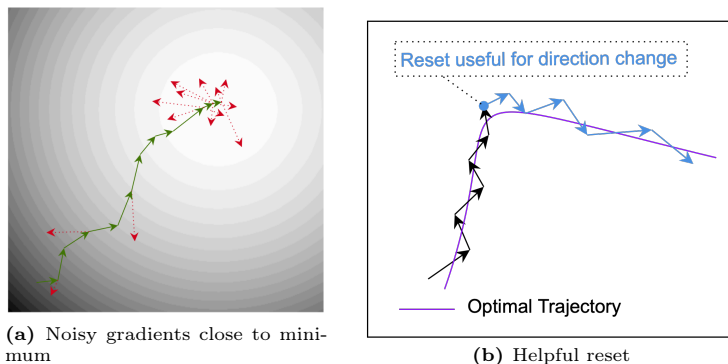


Fig. 2: Intuition for proposed approaches: **(a)** As the model reaches closer to minima, the individual sample gradients start to be misaligned with gradients of previous samples [1, 8, 20]. We leverage this misalignment to perform early stopping without a validation set. **(b)** While effective in moving in the direction of most aligned gradients, the introduced criterion based on angular deviation could prevent adaptation when a direction change is needed, even if the following updates (or gradients) are aligned. A reset of the past horizon (i.e. gradients of previous samples) considered in the alignment condition can help resolve such situations.

approach balances adapting to the new domain, avoiding source forgetting, and managing the computational cost-performance trade-off. Ablation studies reveal that LEAST’s performance is robust to hyperparameter choices.

The contributions of our paper are summarized as follows:

1. We introduce LEAST, a method designed to prevent the overfitting of individual layers to a new domain during TTA.
2. We propose a novel cosine distance-based criterion that can identify the early stopping point from gradient updates without needing a validation set.
3. Through extensive experiments across different backbones, datasets, and TTA losses, we show that LEAST outperforms standard *ERM* (no adaptation), *all layers* baselines, and other layer selection baselines (i.e., AutoRGN and AutoSNR [15]) for TTA.

2 Methodology

We propose to stop adaptation of individual layers for a given sample to prevent over-adaptation. We first introduce our proposed Layerwise EARly STopping (LEAST) framework for Test Time Adaptation (TTA) in Sec. 2.1. Next, we describe our proposed cosine distance criterion for layerwise early stopping without a validation set in Sec. 2.2. Finally, we are describing our reset window strategy to improve the proposed criterion in Sec. 2.3.




Fig. 3: Visualization of our criterion based on angular deviation. Our proposed approach applies an update to a layer only if it is aligned, otherwise, no adaption is made (early stop). In the figure, updates drawn in red are discarded, while updates drawn in green are actually applied, adding up to \mathbf{TD}_i . The update under scrutiny u_i is drawn in cyan, and its sum with \mathbf{TD}_i is drawn in blue. The decision to apply the update u_i or not is made based on the angle α_i .

2.1 Layerwise Masking in TTA

Let $f_{\theta_{\text{src}}}$ denote the model parameterized by parameters θ_{src} trained beforehand on the source domain \mathcal{D}_{src} . Let us also assume that target domain samples $\{x_i\}_{i=1}^n$ are coming in an online fashion at test time. For some sample x_i at test time, TTA adapts the model to obtain θ_i before performing inference [17, 25]. We set $\theta_0 = \theta_{\text{src}}$ and at each step θ_i is obtained by updating θ_{i-1} using the following equation:

$$\theta_i = \theta_{i-1} + \mathbf{u}_i, \quad (1)$$

where \mathbf{u}_i is a parameter update specific to the TTA algorithm. Typically, if SGD optimizer is used with learning rate η , this update takes the form $\mathbf{u}_i = -\eta \nabla \mathcal{L}(x_i; \theta_{i-1})$, where \mathcal{L} is the unsupervised loss specific to the TTA method. For notational simplicity, we adopt this general framework of online updates in this section, but it generalizes in a straightforward manner to other setups commonly used in TTA approaches such as multi-step TTA with mini-batches. In this section, we consider single-step TTA performed online on a single input sample using an SGD optimizer for notation simplicity. However, it can be easily extended to other standard experimental setups of TTA, such as multi-step TTA with a mini-batch of samples and using any optimizer like Adam.

Throughout, we assume the deep learning model is written as a certain composition of functions, which we simply refer to as layers, though any granularity would do. This allows us to write the model at step i as $f_{\theta_i} = f_{\theta_{i,L}} \circ \dots \circ f_{\theta_{i,1}}$, where $\theta_{i,l}$ denote the parameters of layer l at step i . The update equation at step i can be written for each layer as:

$$\theta_{i,l} = \theta_{i-1,l} + \mathbf{u}_{i,l}. \quad (2)$$

To perform layerwise masking, we modify this update equation by introducing a mask:

$$\theta_{i,l} = \theta_{i-1,l} + m_{i,l} \mathbf{u}_{i,l}, \quad (3)$$

where $m_{i,l} \in \{0, 1\}$ is the value of the binary mask applied to the update $\mathbf{u}_{i,l}$.

2.2 Early stopping criterion

Most gradient updates for the model across the samples initially point in the same direction [1, 8, 20], especially when the model is very far from its minima. However, as the model reaches closer to the minima, the gradients across the samples get noisy. Fig. 2 illustrates this intuition. Our objective is to identify for which new sample and for which layer the model can learn the most because it is still far from its minima. We thus exploit the results above on early stopping to design our criterion. At the start of the i^{th} step, the total displacement of parameters of layer l is given by:

$$\mathbf{TD}_{i-1,l} = \sum_{j=1}^{i-1} m_{j,l} \mathbf{u}_{j,l} = \boldsymbol{\theta}_{i-1,l} - \boldsymbol{\theta}_{0,l}. \quad (4)$$

Our proposed criterion relies on the angular deviation of the update $\mathbf{u}_{i,l}$ from the direction of the total displacement that would result from making this update:

$$\cos(\alpha_{i,l}) = \frac{\mathbf{u}_{i,l} \cdot (\mathbf{u}_{i,l} + \mathbf{TD}_{i-1,l})}{\|\mathbf{u}_{i,l}\|_2 \|\mathbf{u}_{i,l} + \mathbf{TD}_{i-1,l}\|_2}. \quad (5)$$

This angle can be interpreted as the deviation of the update under consideration from anticipated average update, since the direction of the anticipated total displacement $\mathbf{u}_{i,l} + \mathbf{TD}_{i-1,l}$ is the same as the anticipated average update. This is can be visualized in Fig. 3.

We then define our layerwise early stopping criterion to make updates only when they align well with the average anticipated update:

$$m_{i,l} = \begin{cases} 1 & \text{if } \cos(\alpha_{i,l}) > \lambda \\ 0 & \text{otherwise} \end{cases}, \quad (6)$$

where λ is the mask threshold. An advantage of this cosine metric over metrics based on norms resides in the fact that it takes value in $[-1, 1]$, which naturally allows us to compare the alignment of updates for layers with different sizes of parameters. As a result, we can intuitively use a single $\lambda > 0$ for the mask thresholding of all layers, to prevent the adaptation of updates that are misaligned with the updates applied in the past. Indeed, a λ close to 1 will only allow adaptation of updates quite aligned with past updates (i.e. $\alpha_{i,l}$ close to 0), while a λ close to 0 or worse, $\lambda < 0$, would be less restrictive.

2.3 Early Stopping in Reset Window

While the cosine distance can stop adaptation for layers if the gradients for some samples are too noisy, our criterion may fail, especially when the gradient update trajectory needs to change direction after a certain point. Fig. 2 (b) illustrates a typical simple example. If the gradient updates meet an inflection point in the loss landscape, cosine distance will prevent further adaptation, and the model

will remain stuck at this point even if the gradient update is informative. Resets can help in such cases. We propose to use a window for resetting the initial parameter point, which we will call *Anchor point*, for the computation of the total displacement of a layer, which gives us:

$$\mathbf{TD}_{i,l} = \boldsymbol{\theta}_{i,l} - \boldsymbol{\theta}_{r,l} \quad (7)$$

where $\boldsymbol{\theta}_{r,l}$ is the parameter at last reset step $r = \lfloor \frac{i-1}{s} \rfloor$, and s is the size of the reset window. We can see that the anchor point changes only when the reset window changes. This is illustrated in Fig. 1.

3 Experimental Results

This section compares our proposed approaches with existing baselines on Domainbed [10], a popular benchmark with large distribution shifts. We will show results across various shifts, backbones, and TTA losses.

3.1 Setup

This benchmark evaluates performance on four datasets: PACS [16], VLCS [7], Terra Incognita [3], and Office-Home [26]. We consider two backbones: ResNet-18 and ResNet-50, with batch normalization layers and two adaptation losses: pseudo-labeling (PL) [14] and SHOT [18]. The pretrained models are trained using default hyperparameters described in Domainbed [10]. To evaluate TTA approaches, we follow the evaluation protocol described in T3A [12], including dataset splits.

Baselines We compare the TTA performance obtained by adapting all layers vs. the layers proposed by our approach. We also compare against AutoRGN and AutoSNR [15], popular baselines in fine-tuning setup to identify optimal layers to adapt.

Model Selection Recent works [35] have shown that hyperparameter tuning significantly impacts results while evaluating different TTA losses. Therefore, we perform model selection as described in Zhao *et al.* [35] for all the benchmarks, where we report the performance of an approach for the best hyperparameter set found by sweeping over all the test samples for a given seed. We report mean and standard deviation across three seeds. We choose the best hyperparameter for TTA losses. We use a single hyperparameter set for our proposed approach.

Implementational details of LEAST Results are reported for LEAST with *window size* of 20 and *mask threshold* of 0.75 with single-layer granularity. A discussion and study of the range of hyperparameters and design choices for LEAST are done in Sec. 5. As we will see in this section, LEAST is not overly sensitive to hyperparameters, so we decided to use the above hyperparameter set. We also scale the masked updates for an initial few samples in the reset window to reduce the impact of incorrect layers selected for the initial few samples.

Table 1: Accuracy (in %) of various test time adaptation strategies on Domainbed. Each value is averaged over three repetitions, with the standard deviation between parentheses. Hyperparameters are tuned for each TTA loss [12, 35]. The best method for a given TTA loss and backbone according to accuracy is in bold, while the second-best for the average accuracy (column Mean) is underlined.

	TTA Method	PACS	VLCS	TerraIncogn.	OfficeHome	Mean	
ResNet-18	ERM -	80.99 (± 0.9)	75.14 (± 1.2)	40.80 (± 0.2)	62.18 (± 0.4)	64.78	
	PL	All layers	81.79 (± 0.7)	65.69 (± 1.5)	35.40 (± 9.7)	60.20 (± 1.4)	60.77
		AutoRGN	82.82 (± 0.6)	72.63 (± 1.3)	38.18 (± 6.1)	62.38 (± 0.2)	<u>64.00</u>
		AutoSNR	80.58 (± 1.2)	65.72 (± 1.8)	35.01 (± 10.4)	59.82 (± 0.9)	60.28
		LEAST	83.56 (± 0.6)	75.48 (± 1.2)	44.19 (± 1.1)	62.67 (± 0.2)	66.47
	SHOT	All layers	83.48 (± 0.3)	66.23 (± 2.8)	33.81 (± 1.3)	63.03 (± 0.4)	61.64
		AutoRGN	84.10 (± 0.5)	69.78 (± 1.3)	37.37 (± 0.7)	63.09 (± 0.2)	<u>63.59</u>
		AutoSNR	83.43 (± 0.3)	66.26 (± 2.7)	33.75 (± 1.2)	63.02 (± 0.4)	61.62
		LEAST	83.92 (± 0.8)	76.23 (± 1.1)	42.13 (± 1.4)	63.32 (± 0.3)	66.40
	ResNet-50	ERM -	82.84 (± 0.5)	75.83 (± 0.9)	46.14 (± 2.3)	66.93 (± 0.3)	67.93
PL		All layers	82.36 (± 2.8)	69.22 (± 1.4)	42.28 (± 3.2)	61.54 (± 3.3)	63.85
		AutoRGN	85.03 (± 1.9)	75.35 (± 1.4)	48.44 (± 2.4)	66.93 (± 0.3)	<u>68.94</u>
		AutoSNR	83.41 (± 3.4)	70.14 (± 4.6)	44.08 (± 3.4)	61.95 (± 3.0)	64.90
		LEAST	84.87 (± 0.8)	76.88 (± 1.6)	50.10 (± 2.5)	67.34 (± 0.3)	69.80
SHOT		All layers	85.15 (± 1.1)	64.25 (± 1.1)	35.33 (± 3.1)	67.37 (± 0.3)	63.03
		AutoRGN	86.34 (± 1.1)	70.2 (± 0.9)	40.59 (± 1.3)	68.10 (± 0.4)	<u>66.31</u>
		AutoSNR	85.51 (± 0.5)	64.26 (± 1.3)	34.97 (± 3.2)	67.33 (± 0.2)	63.02
		LEAST	86.13 (± 0.8)	76.48 (± 1.0)	45.94 (± 1.6)	68.13 (± 0.3)	69.17

3.2 Results

We summarize the key takeaways based on the results from Tab. 1 as follows:

- **LEAST consistently outperforms standard baselines like ERM and All layers TTA.** Tab. 1 shows that LEAST outperforms ERM (no adaptation) and all layers TTA adaptation baselines across all losses, backbones, and datasets. LEAST outperforms ERM baseline by 2% overall. Against all layers TTA, LEAST outperforms by more than 5% overall.
- **LEAST outperforms layer selection approaches.** In Tab. 1, we see that existing layer selection baselines like AutoRGN or AutoSNR can improve performance compared to all layers TTA in most setups, especially AutoRGN. However, they fail to improve against no adaptation baselines for some datasets like VLCS or TerraIncognita or some TTA losses like SHOT. On the other hand, when compared with existing baselines, LEAST consistently demonstrates equivalent or superior performance across all datasets and TTA losses, achieving an overall improvement of about 2%.
- **Most approaches improve better on large shifts than small shifts.** Domainbed has two types of datasets: small shifts (VLCS, TerraIncognita)

Table 2: Effect of early stopping and layer selection approaches on TTA accuracy, source forgetting, and execution time on Tiny Domainbed. Each value is averaged across 12 setups of six shifts and two TTA losses. They all are executed for one random seed with block granularity using a ResNet-18 backbone. Results in bold denotes best according to the metric, while underlined results are second-best.

Blocks adapted	Early stopping	TTA acc. (%) \uparrow	Forgetting (%) \downarrow	Time (sec.) \downarrow
All layers	None	42.30	49.50	198.0
	E.S. with valid. set	54.86	27.47	321.5
Single random	None	53.88	34.34	198.3
	E.S. with valid. set	55.61	<u>20.61</u>	301.2
	E.S. with cosine dis.	<u>55.11</u>	20.03	<u>220.7</u>

and large shifts (PACS, OfficeHome). LEAST improves over large shift datasets, similar to AutoRGN and AutoSNR, which also comfortably beat the ERM baseline. On small shift datasets, existing baselines struggle to outperform the no-adaptation baseline. Still, LEAST can prevent degradation by preventing over-adaptation, thereby enhancing performance over the ERM baseline and safeguard against further degradation.

4 Empirical Analysis of LEAST

In this section, we aim to study the performance of layerwise early stopping on Tiny Domainbed, a small setup we made, spanning six shifts and two TTA losses on the ResNet18 backbone with one seed and block granularity. The six shifts span the most critical shifts of the PACS, VLCS, and TerraIncognita datasets – an explanation for the choice of these shifts is given in the supplementary material. We choose the best hyperparameters for each block.

Concerning the **layerwise early stopping using a validation set**, we observe the following in the results:

- **Early stopping (with a validation set) alone helps.** In Tab. 2, we see that early stopping with a validation set (54.86%) improves over no early stopping (42.3%).
- **Layer selection alone helps.** In Tab. 2, we see that randomly choosing any layer for adaptation (53.88%) improves over adapting with all layers.
- **Layerwise early stopping, which combines the two, works the best.** In Tab. 2, we see that early stopping with random layer selection (55.61%) improves over the other two ideas when taken individually.
- **Layerwise early stopping prevents over-adaptation and can reduce source forgetting.** In Tab. 2, we see that source forgetting is minimal for layerwise early stopping (20.61%), followed by either layer selection alone (34.4%) or early stopping alone (27.47%). The standard TTA approach of training all layers without early stopping suffers from significant forgetting (49.5%), which, as a result, is reflected in degraded TTA accuracy.

Table 3: Accuracy (in %) of different partitioning and selection models of LEAST. *Block* means a single block of many layers is updated at each iteration (blocks are predefined), *single layer* corresponds to the best layer selected for the update. In contrast, *multiple layers* corresponds to k individually best layers selected for the update. Selection is based on cosine distance. The experimental setup corresponds to the best configurations reported in Tab. 1.

Partitioning	ResNet-18		ResNet-50	
	PL	SHOT	PL	SHOT
Block	66.19	65.80	70.21	68.37
Single layer	66.31	66.78	71.07	70.13
Multiple layers	63.13	65.36	69.11	68.29

Now, considering the **proposed layerwise early stopping (without a validation set)**, results show us that:

- **Cosine distance criterion well approximates the layerwise early stopping with a validation set.** In Tab. 2, we see that even without a validation set cosine distance, the proposed criterion (at 55.11%) approximates well the performance from Layerwise early stopping with a validation set (55.61%). In some sense, layerwise, early stopping with a validation set acts as an empirical upper-bound performance if we have access to a validation set with labels for early stopping over-adaptation.
- **Cosine distance can be too conservative compared to layerwise early stopping with a validation set.** Cosine distance works by trying to early stop the adaptation by preventing going too close to minima. In Tab. 2, we see that it may have stopped a little too early compared to layerwise early stopping, the empirical upper bound. As a result, it gets slightly better at forgetting but worse at TTA accuracy.
- **Cosine distance effectively balances computational time and performance.** In Tab. 2, we see that the Cosine distance is faster than layerwise early stopping with a validation set. Any early stopping with a validation set (all layers or layerwise) needs to forward pass through a validation set. It takes longer to perform this computation (~ 300 s) than for cosine distance to perform two backward passes instead of one on each sample (~ 220 s). But at the same time, the layerwise early stopping with a validation set has a higher TTA performance. Similarly, although Cosine distance is slower than methods without early stopping (~ 200 s), this slight increase in computational cost comes with increased performance for cosine distance.

5 Ablation Study of LEAST

This section aims to study the impact of different design choices (like granularity) and hyperparameters (like window size) on the performance of LEAST. This will help justify the choice of hyperparameters and design choices made in Tab. 1.

Table 4: Accuracy (in %) over two main parameters of LEAST: reset window size and mask threshold. Results reported for each backbone are averaged over four Domainbed datasets for a single seed, using single-layer selection. The parameters underlined are those used for results presented in Tab. 1 and 3.

Parameter	Value	ResNet-18		ResNet-50	
		PL	SHOT	PL	SHOT
Window size	5	66.37	66.61	70.87	70.01
	<u>20</u>	66.31	66.78	71.07	70.13
	<u><i>n</i></u>	65.71	66.59	70.76	70.42
Mask threshold	0.5	66.31	66.78	71.07	70.13
	<u>0.75</u>	66.31	66.78	71.07	70.13
	0.99	66.24	66.96	71.07	70.13

We make the following observations when ablating for design choices made in the implementation of LEAST:

- **Layer granularity performs better than Block granularity.** In Tab. 3, LEAST performs better with layer than block granularity. At layer granularity, LEAST has better fine-grained control over choosing the layers to adapt. This highlights that going even more fine-grained can improve performance.
- **Adaptation with the best single layer is much better than with the best multiple layers** In Tab. 3, we see that adapting with the single best layer is much better than adapting with all the layers identified by the cosine distance. This shows that while cosine distance can correctly identify the single best layer to train, it may still struggle to determine the exact set of all best layers. As a result, adapting with the best layer identified using cosine distance improves performance. However, the layers identified for each sample are often not the same.

We also ablate for hyperparameters of LEAST and make the following observations:

- **There exists an optimal reset-window size.** In Tab. 4, we see that a reset window size of 20 works across the backbones and the TTA losses. While there might be an optimal window size for each shift, a window size 20 generally works.
- **Not too sensitive to the choice of mask threshold.** In Tab. 4, we see that LEAST is not very sensitive to the choice of the mask threshold. We see that a threshold of 0.75 seems to work everywhere.

6 Analysis of LEAST

We are now providing some arguing on why the proposed layerwise early stopping framework is more flexible than existing masked adaptation strategies (Sec. 6.1).




Fig. 4: Different masking strategies: Layer selection approaches adapt a selected set of layers for all the samples. Early stopping approaches halt the adaptation of all layers after a certain number of samples. Layerwise Early Stopping is more flexible and can dynamically halt the adaptation of individual layers at different points in time.

We also provide explanations on the fact that cosine distance criterion depends on the *alignment* of current update with respect to total displacement over previous samples and the *magnitude* of the current update (Sec. 6.2).

6.1 Masked Adaptation

The mask formulation given in Eq. (3), with $m_{i,l} \in \{0, 1\}$ being the value of the binary mask applied to the gradient update $\mathbf{u}_{i,l}$ at layer l for the input sample x_i , is quite general and allows to recover different cases for TTA:

- **No adaptation or ERM:** $m_{i,l}$ is 0 for all layers and samples.
- **Standard TTA:** $m_{i,l}$ is 1 for all layers and samples.
- **Layer selection in TTA:** $m_{:,l} = 1$ if layer l is trainable, else 0. In other words, for every sample (or row), we train the same set of layers (or columns) that are part of a trainable layer subset.
- **Early Stopping in TTA:** $m_{i,:} = 1$ if sample x_i trainable, else 0. Like above, we train all the layers for samples (rows) that are part of a trainable sample subset. This can also be viewed as performing a form of sample selection in TTA.
- **Layerwise early stopping in TTA:** It represents the general case we introduce in this paper; decisions to adapt a given layer at a given step are taken individually. An update at step i for a layer l can be deemed beneficial ($m_{i,l} = 1$) or not ($m_{i,l} = 0$).

We can also visualize the above three formulations in Fig. 4. We can see that LEAST is more flexible than the existing layer selection and sample selection strategies in TTA for performing masked adaptation.

6.2 Cosine distance

In this section, we will analyze the proposed cosine distance criterion to understand the role of the current sample’s update with respect to the total displacement of the parameters. Expanding the numerator in the cosine distance criterion in Eq. (5) and simplifying, we get




Fig. 5: Effect of magnitude of u on Cosine distance criterion: The left figure visualizes the trajectory, whereas the middle and the right figure visualize the cosine distance in a 2D layer of update vector ($\mathbf{u}_{i,l}$ or \mathbf{u}) while keeping the previous total displacement ($\mathbf{TD}_{i-1,l}$ or \mathbf{TD}) fixed. **Left:** While Cosine distance primarily depends on the *alignment* between the update and total displacement, it also depends on the *magnitude* of the update. A larger *magnitude* ($u_1 > u_0$) results in a smaller angle ($\alpha_1 < \alpha_0$). **Middle:** When $u \ll \mathbf{TD}$, Cosine distance significantly depends on the *alignment* of u w.r.t. \mathbf{TD} and is the scenario 1 described in Sec. 6.2. **Right:** When $u \gg \mathbf{TD}$, Cosine distance can also depend on the *magnitude* of u and is the scenario 3.

$$\cos(\alpha_{i,l}) = \frac{\mathbf{u}_{i,l} \cdot \mathbf{TD}_{i-1,l}}{\|\mathbf{u}_{i,l}\|_2 \|\mathbf{u}_{i,l} + \mathbf{TD}_{i-1,l}\|_2} + \frac{\|\mathbf{u}_{i,l}\|_2}{\|\mathbf{u}_{i,l} + \mathbf{TD}_{i-1,l}\|_2} \quad (8)$$

We can see how cosine distance depends on the *alignment* of the current sample's update ($\mathbf{u}_{i,l}$) with respect to the total displacement over the previously seen samples ($\mathbf{TD}_{i-1,l}$) and the *magnitude* of the current sample's update ($\|\mathbf{u}_{i,l}\|_2$). The dependence of cosine distance on the *alignment* and *magnitude* of the current sample changes based on the magnitude or norm of the current sample's update ($\|\mathbf{u}_{i,l}\|_2$) and the norm of total displacement before the update (over all the previously seen samples) ($\|\mathbf{TD}_{i-1,l}\|_2$) as follows:

- **Scenario 1:** Current sample update is much smaller than total displacement before update, i.e. $\|\mathbf{u}_{i,l}\|_2 \ll \|\mathbf{TD}_{i-1,l}\|_2$. In this scenario, Cosine distance significantly depends on the *alignment* of the current sample and is less impacted by its *magnitude*. This scenario is more likely for most of the samples during TTA. This can also be viewed similarly to performing an exploit strategy in a given update direction, i.e., after having seen a certain number of samples, we allow adaptation of a layer if the gradients for new updates are well aligned with the previously seen updates.
- **Scenario 2:** Current sample update is comparable in norm with total displacement before update, i.e. $\|\mathbf{u}_{i,l}\|_2 \approx \|\mathbf{TD}_{i-1,l}\|_2$. In this scenario, Cosine distance depends on the current sample's *alignment* and *magnitude*, and both contend with each other. This scenario can sometimes happen for a few initial samples. As a result, we can view it as the cosine distance trying to carefully balance exploration and exploitation of update directions for the layer since we are still at the initial phase of the layer adaptation.
- **Scenario 3:** Current sample update is much larger than total displacement before update, i.e. $\|\mathbf{u}_{i,l}\|_2 \gg \|\mathbf{TD}_{i-1,l}\|_2$. In this scenario, Cosine distance

significantly depends on the *magnitude* of the current sample and is less impacted by its *alignment* with the total displacement. This is usually quite unlikely to happen during TTA. This can occur for the very first samples in rare cases or if the gradients on a layer for a particular sample are much larger than usual. Similarly, we can view it as performing a form of pure exploration strategy over update directions.

The three scenarios described above can also occur in Fig. 5. Therefore, we can see that while the proposed cosine distance criterion usually depends most on the *alignment* of the current sample’s update with the total displacement of previously seen samples, it also weakly depends on the *magnitude* of the current sample’s update.

7 Related Work

Parameter efficient fine-tuning Selectively fine-tuning a subset of parameters instead of all the parameters has seen a growing interest in recent literature on fine-tuning or supervised adaptation. Lee *et al.* [15] demonstrated that selectively training certain blocks or layers of the model can enhance performance compared to training all layers. Panigrahi *et al.* [22] highlighted the benefits of parameter selection post-comprehensive layer fine-tuning, such as improved calibration and out-of-distribution generalization. Wortsman *et al.* [30] propose to interpolate between the source pre-trained and target finetuned model. Even in the context of LLMs, low-rank approximation methods like LORA [11] have surged in popularity. Similar to these approaches, this paper proposes a novel parameter-efficient strategy in a more challenging setup of TTA. Specifically, we propose layer-wise early stopping without a validation set to balance the learning of new features and retaining the useful pretrained features on the new domain.

Parameter selection in test time adaptation Lee *et al.* [15] presented preliminary results indicating that block or layer selection could benefit TTA approaches. Park *et al.* [23] revealed that a layerwise learning rate could enhance TTA performance. Vianna *et al.* [27] suggested that optimal parameters could be selected based on the Wasserstein distance for superior TTA performance. Wang *et al.* [29] proposed to randomly restore a fraction of parameters to the pretrained model during adaptation. However, the impact of these studies on layers or parameter selection is somewhat limited to individual TTA losses, rendering them non-exhaustive and potentially inapplicable to other TTA losses proposed in the past or future. On the other hand, we propose a layerwise early stopping approach for TTA, which is more flexible than layer selection-only strategies, and we demonstrate it improves performance across different TTA losses.

Regularization in test time adaptation Various loss-based regularization approaches have been proposed to address the problem of overfitting in TTA. Niu *et al.* [21] paper proposed Fisher regularization to constrain parameter updates during adaptation. Chen *et al.* [6] proposed a nearest-neighbor voting-based regularize to improve the pseudolabelling loss. Liu *et al.* [19] propose to regularize the adaptation by minimizing the features on the new domain with respect to

the features on the source domain. Zhang *et al.* [33] paper proposed to reduce the variance of predictions across different augmentations of the input sample. While these existing works propose various strategies to regularize the loss function, we take a parameter-centric approach to regularize test time adaptation in this paper, which we argue can be more effective for efficient TTA. Moreover, these loss-based regularization approaches can be used in conjunction with our proposed approaches to improve TTA performance further.

Unsupervised Early stopping Recent works have proposed various approaches or criteria for early stopping without a validation set. Mahsereci *et al.* [20] proposed an evidence-based criterion based on the variance of gradients [1] to perform early stopping. Forouzesh *et al.* [8] proposed to use gradient disparity across samples as the criterion. Yuan *et al.* [32] proposed to perform early stopping by tracking the model’s predictions on the samples. Most of these existing works have demonstrated the effectiveness of their proposed approach often in the context of multi-epoch and supervised learning [4] or noisy learning [2]. Building upon these approaches, in this paper, we propose a novel cosine distance-based criterion to perform layerwise early stopping without a validation set in the context of test time adaptation, which is distinct in being an online and unsupervised context. We also highlight that performing early stopping without a validation set is crucial in the context of TTA, in order to prevent degradation coming from over-adaptation.

8 Conclusion

In this paper, we introduce Layerwise EARly STopping (LEAST), a novel approach designed specifically for Test Time Adaptation (TTA). To prevent overfitting during TTA, LEAST employs a novel criterion for masking the adaptation of individual layers for a given sample at test time. This is done when the updates of the current sample do not align with those of previously seen samples. The alignment is measured using a novel cosine distance-based criterion, which only requires storing a single copy of the initial model’s parameters. Consequently, there is no requirement for access to any supervised labels or validation set, a common necessity in classical early stopping methods. This makes LEAST a promising solution for efficient and effective TTA.

We have conducted comprehensive experiments on Domainbed, demonstrating that LEAST consistently outperforms all baselines and has no adaptation baselines across various datasets, neural network backbones, and losses. Furthermore, LEAST excels as an intuitive layer selection strategy in TTA, surpassing existing baselines. In the Tiny-Domainbed benchmark, we illustrate that the proposed Cosine distance-based criterion closely approximates the performance of layerwise early stopping using a validation set. Importantly, it strikes a harmonious balance between adapting to the new domain, preventing source forgetting, and managing the computational cost-performance trade-off. The results of our study underscore the efficacy of our proposed methods, paving the way for fully unlocking the potential of TTA approaches.

Acknowledgements

This work is supported by the DEEL Project CRDPJ 537462-18 funded by the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Consortium for Research and Innovation in Aerospace in Québec (CRIAQ), together with its industrial partners Thales Canada inc, Bell Textron Canada Limited, CAE inc and Bombardier inc.⁶

References

1. Agarwal, C., D’souza, D., Hooker, S.: Estimating example difficulty using variance of gradients. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10368–10378 (2022) [2](#), [3](#), [5](#), [14](#)
2. Bai, Y., Yang, E., Han, B., Yang, Y., Li, J., Mao, Y., Niu, G., Liu, T.: Understanding and improving early stopping for learning with noisy labels. *Advances in Neural Information Processing Systems* **34**, 24392–24403 (2021) [14](#)
3. Beery, S., Van Horn, G., Perona, P.: Recognition in terra incognita. In: Proceedings of the European conference on computer vision (ECCV). pp. 456–473 (2018) [6](#)
4. Bonet, D., Ortega, A., Ruiz-Hidalgo, J., Shekkizhar, S.: Channel-wise early stopping without a validation set via nnk polytope interpolation. In: 2021 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC). pp. 351–358. IEEE (2021) [14](#)
5. Burns, C., Steinhardt, J.: Limitations of post-hoc feature alignment for robustness. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2525–2533 (2021) [1](#)
6. Chen, D., Wang, D., Darrell, T., Ebrahimi, S.: Contrastive test-time adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 295–305 (2022) [1](#), [13](#)
7. Fang, C., Xu, Y., Rockmore, D.N.: Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1657–1664 (2013) [6](#)
8. Forouzesh, M., Thiran, P.: Disparity between batches as a signal for early stopping. In: Machine Learning and Knowledge Discovery in Databases. Research Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part II 21. pp. 217–232. Springer (2021) [2](#), [3](#), [5](#), [14](#)
9. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press (2016), <http://www.deeplearningbook.org> [2](#)
10. Gulrajani, I., Lopez-Paz, D.: In search of lost domain generalization. In: International Conference on Learning Representations (2021) [1](#), [2](#), [6](#)
11. Hu, E.J., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: Lora: Low-rank adaptation of large language models. In: International Conference on Learning Representations (2021) [13](#)
12. Iwasawa, Y., Matsuo, Y.: Test-time classifier adjustment module for model-agnostic domain generalization. *Advances in Neural Information Processing Systems* **34**, 2427–2440 (2021) [1](#), [6](#), [7](#)
13. Jang, M., Chung, S.Y., Chung, H.W.: Test-time adaptation via self-training with nearest neighbor information. In: The Twelfth International Conference on Learning Representations (2023) [1](#)

⁶ <https://deel.quebec>

14. Lee, D.H., et al.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In: Workshop on challenges in representation learning, ICML. vol. 3, p. 896. Atlanta (2013) [6](#)
15. Lee, Y., Chen, A.S., Tajwar, F., Kumar, A., Yao, H., Liang, P., Finn, C.: Surgical fine-tuning improves adaptation to distribution shifts. In: The Eleventh International Conference on Learning Representations (2023) [1](#), [2](#), [3](#), [6](#), [13](#)
16. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Deeper, broader and artier domain generalization. In: Proceedings of the IEEE international conference on computer vision. pp. 5542–5550 (2017) [6](#)
17. Liang, J., He, R., Tan, T.: A comprehensive survey on test-time adaptation under distribution shifts. arXiv preprint arXiv:2303.15361 (2023) [1](#), [2](#), [4](#)
18. Liang, J., Hu, D., Feng, J.: Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In: International conference on machine learning. pp. 6028–6039. PMLR (2020) [6](#)
19. Liu, Y., Kothari, P., Van Delft, B., Bellot-Gurlet, B., Mordan, T., Alahi, A.: Ttt++: When does self-supervised test-time training fail or thrive? Advances in Neural Information Processing Systems **34**, 21808–21820 (2021) [1](#), [13](#)
20. Mahsereci, M., Balles, L., Lassner, C., Hennig, P.: Early stopping without a validation set. arXiv preprint arXiv:1703.09580 (2017) [2](#), [3](#), [5](#), [14](#)
21. Niu, S., Wu, J., Zhang, Y., Chen, Y., Zheng, S., Zhao, P., Tan, M.: Efficient test-time model adaptation without forgetting. In: International conference on machine learning. pp. 16888–16905. PMLR (2022) [1](#), [13](#)
22. Panigrahi, A., Saunshi, N., Zhao, H., Arora, S.: Task-specific skill localization in fine-tuned language models. In: International conference on machine learning. PMLR (2023) [13](#)
23. Park, J., Kim, J., Kwon, H., Yoon, I., Sohn, K.: Layer-wise auto-weighting for non-stationary test-time adaptation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1414–1423 (2024) [13](#)
24. Shalev-Shwartz, S., Ben-David, S.: Understanding machine learning: From theory to algorithms. Cambridge university press (2014) [2](#)
25. Sun, Y., Wang, X., Liu, Z., Miller, J., Efros, A., Hardt, M.: Test-time training with self-supervision for generalization under distribution shifts. In: International conference on machine learning. pp. 9229–9248. PMLR (2020) [1](#), [2](#), [4](#)
26. Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep hashing network for unsupervised domain adaptation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5018–5027 (2017) [6](#)
27. Vianna, P., Chaudhary, M.S., Tang, A., Cloutier, G., Wolf, G., Eickenberg, M., Belilovsky, E.: Channel selection for test-time adaptation under distribution shift. NeurIPS 2023 Workshop on Distribution Shifts: New Frontiers with Foundation Models (2023) [13](#)
28. Wang, D., Shelhamer, E., Liu, S., Olshausen, B., Darrell, T.: Tent: Fully test-time adaptation by entropy minimization. arXiv preprint arXiv:2006.10726 (2020) [1](#)
29. Wang, Q., Fink, O., Van Gool, L., Dai, D.: Continual test-time domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7201–7211 (2022) [1](#), [13](#)
30. Wortsman, M., Ilharco, G., Gadre, S.Y., Roelofs, R., Gontijo-Lopes, R., Morcos, A.S., Namkoong, H., Farhadi, A., Carmon, Y., Kornblith, S., et al.: Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In: International Conference on Machine Learning. pp. 23965–23998. PMLR (2022) [13](#)

31. Yang, J., Wang, P., Zou, D., Zhou, Z., Ding, K., Peng, W., Wang, H., Chen, G., Li, B., Sun, Y., et al.: Openood: Benchmarking generalized out-of-distribution detection. *Advances in Neural Information Processing Systems* **35**, 32598–32611 (2022) [1](#)
32. Yuan, S., Feng, L., Liu, T.: Early stopping against label noise without validation data. In: *The Twelfth International Conference on Learning Representations* (2024) [14](#)
33. Zhang, M., Levine, S., Finn, C.: Memo: Test time robustness via adaptation and augmentation. *Advances in Neural Information Processing Systems* **35**, 38629–38642 (2022) [14](#)
34. Zhao, B., Chen, C., Xia, S.T.: Delta: Degradation-free fully test-time adaptation. In: *The Eleventh International Conference on Learning Representations* (2022) [1](#)
35. Zhao, H., Liu, Y., Alahi, A., Lin, T.: On pitfalls of test-time adaptation. In: *International conference on machine learning*. PMLR (2023) [1](#), [6](#), [7](#)