

A Probabilistic Framework for Lifelong Test-Time Adaptation

Dhanajit Brahma

Indian Institute of Technology Kanpur

dhanajit@cse.iitk.ac.in

Piyush Rai

Indian Institute of Technology Kanpur

piyush@cse.iitk.ac.in

Abstract

Test-time adaptation (TTA) is the problem of updating a pre-trained source model at inference time given test input(s) from a different target domain. Most existing TTA approaches assume the setting in which the target domain is stationary, i.e., all the test inputs come from a single target domain. However, in many practical settings, the test input distribution might exhibit a lifelong/continual shift over time. Moreover, existing TTA approaches also lack the ability to provide reliable uncertainty estimates, which is crucial when distribution shifts occur between the source and target domain. To address these issues, we present PETAL (Probabilistic liFELong Test-time Adaptation with seLf-training prior), which solves lifelong TTA using a probabilistic approach, and naturally results in (1) a student-teacher framework, where the teacher model is an exponential moving average of the student model, and (2) regularizing the model updates at inference time using the source model as a regularizer. To prevent model drift in the lifelong/continual TTA setting, we also propose a data-driven parameter restoration technique which contributes to reducing the error accumulation and maintaining the knowledge of recent domains by restoring only the irrelevant parameters. In terms of predictive error rate as well as uncertainty based metrics such as Brier score and negative log-likelihood, our method achieves better results than the current state-of-the-art for online lifelong test-time adaptation across various benchmarks, such as CIFAR-10C, CIFAR-100C, ImageNetC, and ImageNet3DCC datasets. The source code for our approach is accessible at <https://github.com/dhanajitb/petal>.

1. Introduction

Deep learning models exhibit excellent performance in settings where the model is evaluated on data from the same distribution as the training data. However, the performance of such models degrades drastically when the distribution of the test inputs at inference time is different from the distribution of the train data (source) [11, 16, 36]. Thus, there is

a need to robustify the network to handle such scenarios. A particularly challenging setting is when we do not have any labeled target domain data to finetune the source model, and unsupervised adaptation must happen at test time when the unlabeled test inputs arrive. This problem is known as *test-time adaptation* (TTA) [28, 35, 38]. Moreover, due to the difficulty of training a single model to be robust to all potential distribution changes at test time, standard fine-tuning is infeasible, and TTA becomes necessary. Another challenge in TTA is that the source domain training data may no longer be available due to privacy/storage requirements, and we only have access to the source pre-trained model.

Current approaches addressing the problem of TTA [28, 35, 38, 41] are based on techniques like self-training based pseudo-labeling or entropy minimization in order to enhance performance under distribution shift during testing. One crucial challenge faced by existing TTA methods is that real-world machine learning systems work in non-stationary and continually changing environments. Even though the self-training based approaches perform well when test inputs are from a different domain but all still i.i.d., it has been found that the performance is unstable when target test inputs come from a continually changing environment [32]. Thus, it becomes necessary to perform test-time adaptation in a continual manner.

Such a setting is challenging because the continual adaptation of the model in the long term makes it more difficult to preserve knowledge about the source domain. Continually changing test distribution causes pseudo-labels to become noisier and miscalibrated [10] over time, leading to error accumulation [3] which is more likely to occur if early predictions are incorrect. When adapting to new test input, the model tends to forget source domain knowledge, triggering *catastrophic forgetting* [27, 31, 33]. Moreover, existing TTA methods do not account for *model/predictive uncertainty*, which can result in miscalibrated predictions.

Recently, [41] proposed CoTTA, an approach to address the continual/lifelong TTA setting using a stochastic parameter reset mechanism to prevent forgetting. Their reset mechanism however is based on randomly choosing a subset of weights to reset and is not data-driven. Moreover,

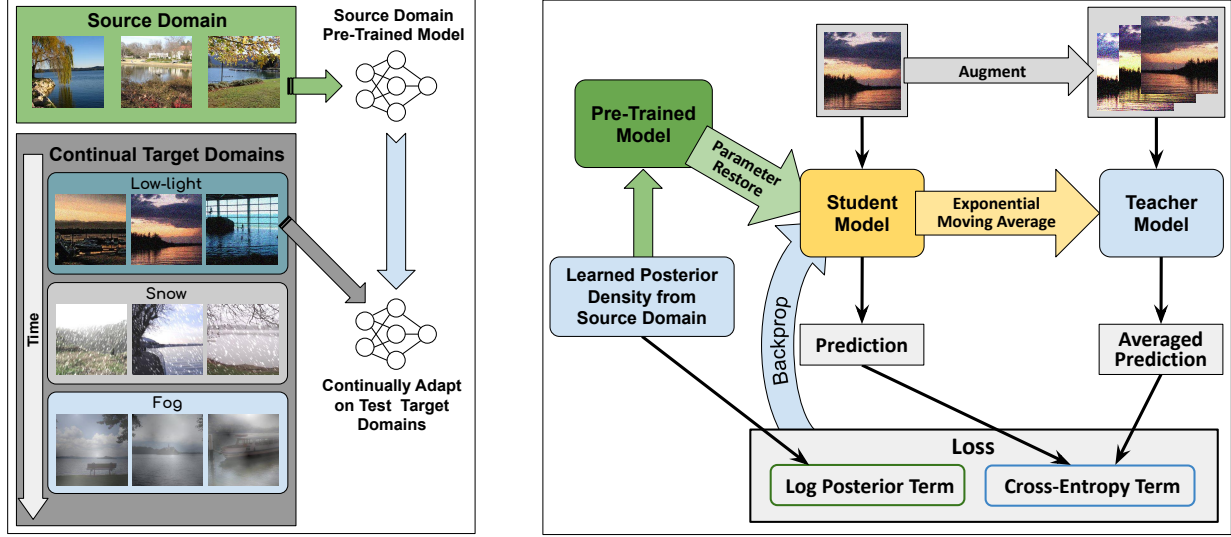


Figure 1. **Left:** Problem setup of online lifelong TTA. During adaptation on test input, the source domain data is no longer available, and only the model pre-trained on the source domain is provided. Test inputs from different domains arrive continually, and the model has no knowledge about change in the domain. **Right:** Our proposed probabilistic framework for online lifelong TTA. We obtain a source domain pre-trained model from the posterior density learned using training data from the source domain. The posterior density is used to initialize the student model. A test sample is provided as input to the student model. Using multiple augmentations of a test sample, we obtain augmentation averaged prediction from the teacher model. The loss term consists of log posterior and cross-entropy terms utilizing student and teacher model predictions. We utilize backpropagation to update student model and exponential moving average for teacher model.

their method does not take into account model/predictive uncertainty and is therefore susceptible to overconfident and miscalibrated predictions.

To improve upon these challenges of continual/lifelong TTA, we propose a principled, probabilistic framework for lifelong TTA. Our framework (shown in Fig. 1 (Right)) constructs a posterior distribution over the source model weights and a data-dependent prior which results in a self-training based cross-entropy loss, with a regularizer term in the learning objective. This regularizer arises from terms corresponding to the posterior, which incorporates knowledge of source (training) domain data.

Moreover, our framework also offers a probabilistic perspective and justification to the recently proposed CoTTA [41] approach, which arises as a special case of our probabilistic framework. In particular, only considering the data-driven prior in our approach without the regularizer term, corresponds to the student-teacher based cross-entropy loss used in CoTTA. Further, to improve upon the stochastic restore used by [41], we present a *data-driven* parameter restoration based on Fisher Information Matrix (FIM). In terms of improving accuracy and enhancing calibration during distribution shift, our approach surpasses existing approaches in various benchmarks.

Main Contributions

1. From a probabilistic perspective, we arrive at the student-teacher training framework in our proposed

Probabilistic liFelong Test-time Adaptation with seLf-training prior (PETAL) approach. Inspired from the self-training framework [19, 42], the teacher model is the exponential moving average of the student model, as depicted in Fig. 1 (Right).

2. The student-teacher cross-entropy loss with a regularizer term corresponding to posterior of source domain data naturally emerges in the probabilistic formulation.
3. We propose a data-driven parameter restoration based on Fisher Information Matrix (FIM) to handle error accumulation and catastrophic forgetting.

2. Problem Setup and Background

In this section, we define the notation used, the problem setup of lifelong/continual TTA, and the basic probabilistic framework on which our approach is based.

2.1. Problem Setup

Let \mathbf{x} denote the inputs sampled i.i.d from a generative model having parameters ψ , and $p(y|\mathbf{x}, \theta)$, having parameters θ , be the conditional distribution from which the corresponding labels are sampled. We denote the prior distributions of θ and ψ using $p(\theta)$ and $p(\psi)$, respectively.

A typical test-time adaptation setting is as follows: We have a model with parameters θ_0 trained on the source training data $\mathcal{X} = \{\mathbf{x}_n, y_n\}_{n=1}^N$. The aim is to adapt θ_0 and per-

form well on the test inputs $\mathcal{U}^d = \{\mathbf{x}_m^d\}_{m=1}^{M_d}$ from an unlabeled target domain d . In case of multiple target domains, the adaptation happens for each target domain separately: $\theta_0 \rightarrow \theta_d$.

In *lifelong/continual* test-time adaptation, unlabeled test inputs from D different target domains $\{\mathcal{U}^d : d = 1, \dots, D\}$ arrive continually, and, thus, the model can utilize only the data available for the current target domain. Note that there is no information available to the learner about the change in domain. At step t , for a test input \mathbf{x}_m , we make predictions using $p(y_m|\mathbf{x}_m, \theta_t)$ as well as adapt the parameters for future steps, i.e., $\theta_t \rightarrow \theta_{t+1}$. Note that there is a continual domain shift in the data distribution of \mathbf{x}_m . Moreover, the model evaluation is performed based on predictions obtained online. Fig. 1 (Left) depicts the online lifelong TTA problem setup.

2.2. The Underlying Probabilistic Framework

In this section, we review standard probabilistic discriminative models for supervised learning and semi-supervised learning (SSL), which incorporates unlabeled data via a *partly* data-dependent prior. Then we discuss a formulation of self-training based Bayesian SSL with partly data-driven cross-entropy prior. Further, we describe a modification to this Bayesian SSL formulation to handle the situation when the distribution of unlabeled inputs is different from the labeled inputs (covariate shift).

In Section 3, we describe how this Bayesian SSL formulation can be further extended for our problem, i.e., the lifelong/continual TTA setting where we learn a source model using only labeled data, and then the model has to predict labels of unlabeled test inputs coming from target domains with different distributions.

Bayesian Supervised Learning

The Bayesian setup for supervised learning typically assumes that we are given labeled data $\mathcal{D} = \{\mathbf{x}_n, y_n\}_{n=1}^N$, and we estimate θ using its posterior distribution

$$p(\theta|\mathcal{D}) \propto p(\theta) \prod_{n=1}^N p(y_n|\mathbf{x}_n, \theta). \quad (1)$$

Given a novel test input \mathbf{x}_{N+1} , we make predictions using posterior predictive distribution obtained by marginalizing over the posterior distribution $p(y_{N+1}|\mathbf{x}_{N+1}) = \int p(y_{N+1}|\mathbf{x}_{N+1}, \theta)p(\theta|\mathcal{D})d\theta$.

Bayesian Semi-Supervised Learning

Here, we are provided with unlabeled data along with some labeled data. We denote the unlabeled data points as $\mathcal{U} = \{\mathbf{x}_m\}_{m=1}^M$. To circumvent the inability to use unlabeled data while inferring θ , one needs to make assumptions about the dependency between distributions of inputs and labels.

To this end, [9] proposed a prior that is partly data-dependent via the inputs \mathbf{x} : $p(\theta|\psi) \propto p(\theta) \exp(-\lambda H_{\theta, \psi}(y|\mathbf{x}))$, where $p(\theta)$ is the prior in

Eq. 1, and H is the conditional entropy of the class label. Here, *partly data-dependent* prior refers to a prior defined using only the input \mathbf{x} treated as given, not the label y which is treated as a random variable.

Bayesian Semi-Supervised Learning with Self-Training

The self-training framework [19, 42] has demonstrated significant success in semi-supervised learning. Our proposed framework is also based on self-training wherein we use an exponential moving average of the parameters θ of the student model $p(y|\mathbf{x}, \theta)$ (which is initialized with the source pre-trained model parameters θ_0), and refer to the averaged model as the teacher model (θ'): $\theta'_{t+1} = \pi\theta'_t + (1-\pi)\theta_{t+1}$, where π is the smoothing factor. For brevity, we will omit the time step index t from here onwards.

In the semi-supervised setting, the teacher model can be utilized to obtain augmentation-averaged pseudo-labels y' corresponding to an unlabeled input \mathbf{x} . To prevent error accumulation, augmentation is only used when the domain difference is substantial. Defining $\hat{y}' = p(y|\mathbf{x}, \theta')$; $\tilde{y}' = \frac{1}{K} \sum_{i=1}^K p(y|\alpha_i(\mathbf{x}), \theta')$, the pseudo-label is defined as

$$y' = \begin{cases} \hat{y}', & \text{if } C(p(y|\mathbf{x}, \theta_0)) \geq \tau \\ \tilde{y}', & \text{otherwise.} \end{cases} \quad (2)$$

Here, K is number of times augmentation is applied, $\alpha_i()$ is augmentation function, $C()$ gives the confidence of the prediction, and τ is threshold for selecting confident predictions. The prediction confidence of current input using source domain pre-trained model θ_0 gives us an estimate of the domain difference between source and target domain.

Using these pseudo labels, we formulate the following partly data-driven cross-entropy prior

$$\begin{aligned} p(\theta|\psi) &\propto p(\theta) \exp(-\lambda H_{\theta, \psi}^{\text{xc}}(y', y|\mathbf{x})) \\ &= p(\theta) \exp(\lambda \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x}|\psi), y' \sim p(y|\mathbf{x}, \theta')} [\log p(y|\mathbf{x}, \theta)]). \end{aligned} \quad (3)$$

Here, $y = p(y|\mathbf{x}, \theta)$ is the prediction of the student model and H^{xc} is the conditional cross-entropy of labels conditioned on the inputs. This cross-entropy term leverages the knowledge from the teacher model as it is incorporated into the partly data-driven prior.

Bayesian Semi-Supervised Learning with Unlabeled Data Distribution Shift

The above semi-supervised learning formulation assumes that the unlabeled inputs come from the same distribution as the labeled inputs. To handle the situation when the unlabeled inputs come from a different distribution, we introduce additional generative parameters $\tilde{\psi}$ while using the same conditional model $p(y|\mathbf{x}, \theta)$ parameters θ for both distributions. Unlabeled inputs $\tilde{\mathbf{x}}$ which come from a different distribution, are assumed to be sampled from the generative model with parameters $\tilde{\psi}$.

Incorporating the additional generative parameters $\bar{\psi}$ in Eq. 3, the prior becomes

$$p(\theta|\psi, \bar{\psi}) \propto p(\theta) \exp(-\lambda H_{\theta, \psi}^{\text{xe}}(y', y|\mathbf{x})) \exp(-\bar{\lambda} H_{\theta, \bar{\psi}}^{\text{xe}}(y', y|\bar{\mathbf{x}})). \quad (4)$$

The conditional entropies present in Eq. 4 require expectations over the distributions of the (labeled and unlabeled) inputs.

Replacing $p(\mathbf{x}|\psi)$ and $p(\bar{\mathbf{x}}|\bar{\psi})$ with the empirical distributions of \mathbf{x} and $\bar{\mathbf{x}}$, we get

$$p(\theta|\psi, \bar{\psi}) \propto p(\theta) \exp\left(-\frac{\lambda}{N} \sum_{n=1}^N H^{\text{xe}}(y'_n, y_n|\mathbf{x}_n, \theta)\right) \exp\left(-\frac{\bar{\lambda}}{M} \sum_{m=1}^M H^{\text{xe}}(y'_m, y_m|\bar{\mathbf{x}}_m, \theta)\right). \quad (5)$$

Using Eq. 5 as the prior in Eq. 1 and taking logarithm on both sides (ignoring additive normalization constants), the new posterior distribution of the network parameters becomes: $\log p(\theta|\mathcal{D}, \mathcal{U}) = \log p(\theta) + \sum_{n=1}^N \log p(y_n|\mathbf{x}_n, \theta) - \frac{\lambda}{N} \sum_{n=1}^N H^{\text{xe}}(y'_n, y_n|\mathbf{x}_n) - \frac{\bar{\lambda}}{M} \sum_{m=1}^M H^{\text{xe}}(y'_m, y_m|\bar{\mathbf{x}}_m)$.

Since the labeled data is already used in the likelihood term, in the prior we ignore the cross-entropy term for the labeled data by setting $\lambda = 0$. Thus, the log-posterior density is simplified to

$$\log p(\theta|\mathcal{D}, \mathcal{U}) = \log p(\theta) + \sum_{n=1}^N \log p(y_n|\mathbf{x}_n, \theta) - \frac{\bar{\lambda}}{M} \sum_{m=1}^M H^{\text{xe}}(y'_m, y_m|\bar{\mathbf{x}}_m). \quad (6)$$

3. Probabilistic Test-Time Adaptation

Our formulation for TTA is similar to the formulation we described above for the problem of Bayesian semi-supervised learning with unlabeled data distribution shift, with a key difference. In contrast to Bayesian SSL in which labeled and unlabeled data are available during training, in *test-time adaptation*, we need to predict labels of the test inputs from a different domain, but we only have access to the source model weights without the source domain training data. In our probabilistic approach, we assume that the source model is given in form of the approximate posterior distribution $q(\theta)$ over the weights of the source model

$$q(\theta) \approx p(\theta|\mathcal{D}). \quad (7)$$

At test-time, we use $q(\theta)$ to represent the source domain knowledge. Substituting Eq. 1 for the posterior of the source domain data, we take the logarithm in Eq. 7 and simplify further to get the following

$$\log q(\theta) = \log p(\theta) + \sum_{n=1}^N \log p(y_n|\mathbf{x}_n, \theta). \quad (8)$$

For our TTA setting, substituting this approximate posterior above in Eq. 6, the log-posterior density with both labeled (source) and unlabeled (test inputs) data becomes

$$\log p(\theta|\mathcal{D}, \mathcal{U}) = \log q(\theta) - \frac{\bar{\lambda}}{M} \sum_{m=1}^M H^{\text{xe}}(y'_m, y_m|\bar{\mathbf{x}}). \quad (9)$$

Since posterior inference for deep neural networks is challenging, we leverage the Gaussian posterior approximation based on the SWAG-diagonal [26] method. It uses the SGD iterates to construct the mean and (diagonal) covariance matrix of the Gaussian posterior approximation and requires minimal changes to the training mechanism on source domain training data.

3.1. Parameter Restoration

Stochastic Restoration

In lifelong TTA, in order to reduce the error accumulation over the long term in self-training and handle catastrophic forgetting, [41] proposed stochastic restoration of weights by additionally updating the parameters. Let θ^f denote the flattened parameter θ of the student model, and D be the dimension of θ^f . After the gradient update at time step t , stochastic restore further updates the parameters:

$$\mathbf{m} \sim \text{Bernoulli}(\rho), \quad (10)$$

$$\theta_{t+1}^f = \mathbf{m} \odot \theta_0^f + (\mathbf{1} - \mathbf{m}) \odot \theta_{t+1}^f. \quad (11)$$

Here, \odot is element-wise multiplication, and ρ is stochastic restore probability. \mathbf{m} is mask to determine which parameters within θ_{t+1}^f to restore to original source weight θ_0^f .

Fisher Information Based Restoration

To improve upon stochastic restoration, we propose a data-driven parameter restoration. Fisher Information Matrix (FIM) is widely used as a metric of parameter importance for a given data [15]. Thus, we use FIM, F , of the student model parameterized by θ as a measure of the importance of the parameters. For a given time step t with unlabeled test inputs batch $\mathcal{B} = \{\bar{\mathbf{x}}_m\}_{m=t|\mathcal{B}|}^{(t+1)|\mathcal{B}|}$, we consider the following diagonal approximation of FIM: $F = \text{Diag}((\nabla_{\theta} \mathcal{L})(\nabla_{\theta} \mathcal{L})^T)$, where, $\mathcal{L} = \log q(\theta) - \frac{\bar{\lambda}}{|\mathcal{B}|} \sum_{m=t|\mathcal{B}|}^{(t+1)|\mathcal{B}|} H^{\text{xe}}(y'_m, y_m|\bar{\mathbf{x}}_m)$. Here, y'_m and y_m are the teacher and student model predictions, respectively. Note that F has the same dimension as θ . Thus, upon using FIM based restore, the parameter restoration in Eq. 10 becomes

$$\mathbf{m}_p = \begin{cases} 1, & \text{if } F_p < \gamma \\ 0, & \text{otherwise.} \end{cases}, p = 1, \dots, P. \quad (12)$$

Here, $\gamma = \text{quantile}(F, \delta)$ is the threshold value which is the δ -quantile of F . Thus, the elements in \mathbf{m} corresponding to

Algorithm 1 Proposed Approach PETAL

Input: Training dataset $\mathcal{X} = \{\mathbf{x}_n, y_n\}_{n=1}^N$;
Test domain data $\mathcal{U}^d = \{\mathbf{x}_m^d\}_{m=1}^{M_d}, d = 1, \dots, D$

Output: Prediction \mathcal{Y}^d

Training:

1: Compute $q(\theta) \approx p(\theta|\mathcal{X})$

Continual Adaptation:

2: Initialize time step $t = 0$, prediction set $\mathcal{Y}^d = \emptyset$
3: Initialize student model $\theta_0 = \max_{\theta} \log q(\theta)$
4: Initialize teacher model $\theta'_0 = \theta_0$
5: **for** $d = 1, \dots, D$ **do**
6: **for each batch** $\mathcal{B} \in \mathcal{U}^d$ **do**
7: $\theta'_{t+1}, \theta_{t+1}, \mathcal{B}_y = \text{Adapt}(\theta'_t, \theta_t, \theta_0, \mathcal{B}, q(\theta), t)$
8: $\mathcal{Y}^d = \mathcal{Y}^d \cup \mathcal{B}_y$
9: $t = t + 1$
10: **end for**
11: **end for**

Algorithm 2 Adapt

Input: $\theta'_t, \theta_t, \theta_0, \mathcal{B}, q(\theta), t$

Require: Number of augmentations K ; learning rate η ;
Threshold for confident predictions τ ;
Confidence function C ; Cross-entropy weight $\bar{\lambda}$;
Quantile of FIM δ ; Augmentation function α ;
Smoothing factor π

Output: $\theta'_{t+1}, \theta_{t+1}, \mathcal{B}_y$

1: Initialize: $\mathcal{H} = 0, \mathcal{B}_y = \emptyset$
2: **for each test input** $\bar{\mathbf{x}} \in \mathcal{B}$ **do**
3: Teacher model prediction: $\hat{y}' = p(y|\bar{\mathbf{x}}, \theta')$
4: Augmented Average: $\tilde{y}' = \frac{1}{K} \sum_{i=1}^K p(y|\alpha_i(\bar{\mathbf{x}}), \theta')$
5: Augment based on domain gap:

$$y' = \begin{cases} \hat{y}', & \text{if } C(p(y|\bar{\mathbf{x}}, \theta_0)) \geq \tau \\ \tilde{y}', & \text{otherwise.} \end{cases}$$

6: Using student model, predict: $y = p(y|\bar{\mathbf{x}}, \theta)$
7: Update: $\mathcal{H} = \mathcal{H} + H^{\text{se}}(y', y|\bar{\mathbf{x}})$
8: $\mathcal{B}_y = \mathcal{B}_y \cup \{y'\}$

9: **end for**

10: Compute: $\mathcal{L} = \left(\log q(\theta) - \frac{\bar{\lambda}}{|\mathcal{B}|} \mathcal{H} \right)$
11: Adapt student model: $\theta_{t+1} = \theta_t + \eta \nabla_{\theta} \mathcal{L}$
12: Update teacher model: $\theta'_{t+1} = \pi \theta'_t + (1 - \pi) \theta_{t+1}$
13: Compute FIM: $F = \text{Diag}((\nabla_{\theta} \mathcal{L})(\nabla_{\theta} \mathcal{L})^T)$
14: Compute mask \mathbf{m} for resetting:

$$\gamma = \text{quantile}(F, \delta)$$

$$\mathbf{m}_p = \begin{cases} 1, & \text{if } F_p < \gamma \\ 0, & \text{otherwise.} \end{cases}, \quad p = 1, \dots, P.$$

15: Reset updated student model back to source model:

$$\theta_{t+1} = \mathbf{m} \odot \theta_0 + (\mathbf{1} - \mathbf{m}) \odot \theta_{t+1}$$

FIM value less than γ would be 1, implying that the corresponding parameters would be restored to original source weight $\theta_f^{(0)}$. The algorithm for PETAL is given in Algorithm 1. Fig. 1 (Right) provides an overview of our approach PETAL.

4. Related Work

4.1. Unsupervised Domain Adaptation

The goal of unsupervised domain adaptation (UDA) [6, 24, 30, 40] is to enhance the performance of the learning model when there is a change in distribution between the training data domain and the test data domain. UDA approaches often assume that the source and (unlabeled) target domain data are accessible simultaneously. Most existing methods address UDA by ensuring that the feature distributions [6, 24, 37] or the input spaces [12, 44] of the source and the target domains are brought closer.

4.2. Test-Time Adaptation

Some works also refer to TTA as *source-free* domain adaptation. Recent works explore source-free domain [18, 20, 23] setting in which training data is unavailable and only unlabeled data is available during adaptation. Test entropy minimization (TENT) [38] starts from a source pre-trained model and updates only the batch-norm (BN) parameters by minimizing entropy in test predictions. [34] address TTA by updating the source domain BN statistics using test input statistics. Continual Test-Time Adaptation (CoTTA) [41] addresses online lifelong TTA by employing weight averaging and augmentation averaging, and random parameter restoration back to source pre-trained model parameters. [8] adapts to continually changing target domains by utilizing a normalization layer to handle the out-of-distribution examples and balanced reservoir sampling to store the simulated i.i.d. data in the memory. It would be an interesting future work to extend PETAL for the temporally correlated test stream setting proposed by [8]. EATA [29] is another related work that looks at preventing forgetting in the context of TTA; however, EATA mainly focuses on preventing the forgetting of the source task model and is not designed to handle forgetting in a lifelong TTA setting. Our work PETAL is a principled probabilistic approach for lifelong test-time adaptation that uses an approximate posterior during test-time adaptation obtained from source domain data. PETAL also offers a probabilistic perspective and justification to CoTTA which arises a special case of PETAL.

Bayesian Adaptation for Covariate Shift (BACS) [46] proposes a Bayesian perspective for TENT, which naturally gives rise to the entropy term along with a regularizer that captures knowledge from posterior density obtained from training data. However, BACS only addresses standard TTA setting, and regularized entropy minimization lacks the abil-

Time	t															
Method	Gaussian	shot	impulse	defocus	glass	motion	zoom	snow	frost	fog	brightness	contrast	elastic	pixelate	jpeg	Mean
Source	72.33	65.71	72.92	46.94	54.32	34.75	42.02	25.07	41.30	26.01	9.30	46.69	26.59	58.45	30.30	43.51
BN Adapt	28.08	26.12	36.27	12.82	35.28	14.17	12.13	17.28	17.39	15.26	8.39	12.63	23.76	19.66	27.30	20.44
Pseudo-label	26.70	22.10	32.00	13.80	32.20	15.30	12.70	17.30	17.30	16.50	10.10	13.40	22.40	18.90	25.90	19.80
TENT-online ⁺	24.80	23.52	33.04	11.93	31.83	13.71	10.77	15.90	16.19	13.67	7.86	12.05	21.98	17.29	24.18	18.58
TENT-continual	24.80	20.60	28.60	14.40	31.10	16.50	14.10	19.10	18.60	18.60	12.20	20.30	25.70	20.80	24.90	20.70
CoTTA	23.92	21.40	25.95	11.82	27.28	12.56	10.48	15.31	14.24	13.16	7.69	11.00	18.58	13.83	17.17	16.29 (0.02)
PETAL (S-Res)	23.44	21.20	25.50	11.80	27.22	12.54	10.45	15.14	14.31	12.89	7.61	10.72	18.42	13.83	17.37	16.16 (0.02)
PETAL (FIM)	23.42	21.13	25.68	11.71	27.24	12.19	10.34	14.76	13.91	12.65	7.39	10.49	18.09	13.36	16.81	15.95 (0.04)

Table 1. Experimental results for CIFAR10-to-CIFAR10C online lifelong test-time adaptation task. The numbers denote the classification error (%) obtained with the highest corruption of severity level 5. TENT-online uses domain information, denoted using +.

ity to handle error accumulation and catastrophic forgetting encountered in lifelong TTA.

4.3. Continual Learning

The objective of Continual Learning (CL) is to learn from a sequential series of tasks, enabling the model to retain previously acquired knowledge while learning a new task, preventing catastrophic forgetting [27, 31, 33]. Elastic weight consolidation (EWC) [15] is a regularization-based technique that penalizes parameter changes having a significant impact on prediction. [22] proposed learning without forgetting (LwF), which preserves knowledge of previous tasks using knowledge distillation. Gradient episodic memory (GEM) [25] maintains a limited number of samples to retrain while constraining fresh task updates from interfering with prior task knowledge. [1, 39] address the continual semi-supervised learning problem where continually arriving tasks consist of labeled and unlabeled data.

5. Experiments

We thoroughly evaluate PETAL and compare it to other state-of-the-art approaches on image classification lifelong test-time adaptation benchmark tasks: CIFAR10-to-CIFAR10C, CIFAR100-to-CIFAR100C, ImageNet-to-ImageNetC, and ImageNet-to-ImageNet3DCC.

5.1. Benchmark Datasets

[11] developed CIFAR10C, CIFAR100C, and ImageNetC datasets to serve as benchmarks for the robustness of classification models. In each dataset, there are 15 different types of corruption and five different levels of severity. These corruptions are applied to test images of original CIFAR10 and CIFAR100 [17] datasets and validation images of original ImageNet [5] dataset. Further, we experiment with ImageNet 3D Common Corruptions (ImageNet3DCC) dataset, recently proposed by [14], which utilizes the geometry of the scene in transformations, leading to more realistic corruptions. ImageNet3DCC dataset consists of 12 different types of corruptions, each with five levels of severity. Refer to Appendix for details.

In online lifelong TTA, we begin with a network trained on CIFAR10, CIFAR100, and ImageNet clean training set for the respective experiments. At the time of testing, the model gets corrupted images online.

Following CoTTA, we continually adjust the source pre-trained model to each corruption type as they sequentially arrive, as opposed to conventional TTA in which the pre-trained model is separately adapted to each corruption type. We evaluate the model using online predictions obtained immediately as the data is encountered. We follow the online lifelong test-time adaptation setting for all the experiments. For ImageNet-to-ImageNetC experiments, we evaluate using 10 different sequences of corruptions.

5.2. Model Architectures and Hyperparameters

Following TENT [38] and CoTTA [41], we adopt pre-trained WideResNet-28 [45] model for CIFAR10-to-CIFAR10C, pre-trained ResNeXt-29 [43] model for CIFAR100-to-CIFAR100C, and standard pre-trained ResNet-50 model for both ImageNet-to-ImageNetC and ImageNet-to-ImageNet3DCC experiments from RobustBench [4].

We utilize SWAG-D [26] to approximate the posterior density from the source domain training data. SWAG-D approximates the posterior using a Gaussian distribution with diagonal covariance from the SGD trajectory. Before adapting the model, we initialize it with the maximum a posteriori (MAP) of the approximate posterior that corresponds to the solution obtained by Stochastic Weight Averaging [13]. This is effectively the source domain pre-trained model.

We update all the trainable parameters in all experiments. We use $K = 32$ number of augmentations. We adopt the same augmentation confidence threshold described in [41]. For FIM based parameter restoration, we set the quantile value $\delta = 0.03$. We refer the readers to the Appendix for more details on the hyperparameters.

5.3. Baselines and Compared Approaches

To evaluate the efficacy of PETAL, we compare the PETAL with CoTTA and five other methods in online lifelong test-time adaptation. *Source* denotes the baseline pre-

Time	$t \rightarrow$															
Method	Gaussian	shot	impulse	defocus	glass	motion	zoom	snow	frost	fog	brightness	contrast	elastic	pixelate	jpeg	Mean
Source	73.00	68.01	39.37	29.32	54.11	30.81	28.76	39.49	45.81	50.30	29.53	55.10	37.23	74.69	41.25	46.45
BN Adapt	42.14	40.66	42.73	27.64	41.82	29.72	27.87	34.88	35.03	41.50	26.52	30.31	35.66	32.94	41.16	35.37
Pseudo-label	38.10	36.10	40.70	33.20	45.90	38.30	36.40	44.00	45.60	52.80	45.20	53.50	60.10	58.10	64.50	46.20
TENT-continual	37.20	35.80	41.70	37.90	51.20	48.30	48.50	58.40	63.70	71.10	70.40	82.30	88.00	88.50	90.40	60.90
CoTTA	40.09	37.67	39.77	26.91	37.82	28.04	26.26	32.93	31.72	40.48	24.72	26.98	32.33	28.08	33.46	32.48 (0.02)
PETAL (S-Res)	38.37	36.43	38.69	25.87	37.06	27.34	25.55	32.10	31.02	38.89	24.38	26.38	31.79	27.38	32.98	31.62 (0.04)
PETAL (FIM)	38.26	36.39	38.59	25.88	36.75	27.25	25.40	32.02	30.83	38.73	24.37	26.42	31.51	26.93	32.54	31.46 (0.04)

Table 2. Experimental results for CIFAR100-to-CIFAR100C online lifelong test-time adaptation task. The numbers denote the classification error rate (%) obtained with the highest corruption of severity level 5.

trained model that has not been adapted to test inputs. In *BN Adapt* [21, 34], the network parameters are kept frozen, and only Batch Normalization statistics are adapted to produce predictions for test inputs. *Pseudo-label* updates the Batch-Norm trainable parameters with hard pseudo-labels [19]. *TENT-online* denotes the TENT [38] approach in this setting, but it has access to extra information about the change in the domain and, thus, resets itself to the original pre-trained model upon encountering test inputs from the new domain and adapts afresh. But this additional knowledge is unavailable in real-life scenarios. *TENT-continual* has no extra information about domain change. *CoTTA* [41] uses weight averaging and augmentation averaging, along with randomly restoring parameters to the original pre-trained model. However, it lacks explicit uncertainty modeling and data-driven parameter restoration.

5.4. Evaluation Metrics

We evaluate our model using the error rate in predictions. To evaluate the uncertainty estimation, we use negative log-likelihood (NLL) and Brier score [2]. Both NLL and Brier are proper scoring rules [7], and they are minimized if and only if the predicted distribution becomes identical to the actual distribution. In Table 1 and 2, the number within brackets is the standard deviation over 5 runs. Refer to the Appendix for more details on evaluation metrics.

Table 3. CIFAR10-to-CIFAR10C results for gradually changing severity level before changing corruption types. The numbers are averaged over all 15 corruption types. The number after \pm is the standard deviation over 10 random corruption sequences. Our method surpasses all baselines in the depicted settings.

Method \ Metric	Source	BN Adapt	TENT	CoTTA	PETAL (FIM)
Error (%)	23.94	13.54	29.46	10.40 \pm 0.22	10.11 \pm 0.23
Brier	0.408	0.222	0.575	0.159 \pm 0.003	0.158 \pm 0.004

5.5. CIFAR10-to-CIFAR10C Results

In Table 1, we observe that directly using the pre-trained model (*Source*) leads to poor performance with an average error rate of 43.51%, suggesting the necessity of adapta-

tion. Adapting the Batch Normalization (BN) statistics improves the average error rate to 20.44%. Using hard pseudo-labels and updating only the BN parameters further improves the performance to 19.8%. TENT-online reduces the error rate to 18.58% using extra information about domain change, but access to such information is mostly unavailable in real-world scenarios. As expected, the error rate of TENT-continual increases to 20.7% without access to domain change information. Further, CoTTA improves the average error rate to 16.29%. Our proposed approach consistently outperforms other approaches for most individual corruption types, reducing the average error rate to 15.95%. Moreover, our proposed approach demonstrates no performance degradation in the long term.

To investigate the contribution of FIM based parameter restore, we show the results of our approach with stochastic restore in place of FIM based restore, denoted using S-Res. We observe that FIM based restore performs better than S-Res for most of the corruption types, highlighting the effectiveness of FIM based restore. Further, the improved performance of PETAL (S-Res) and PETAL (FIM) over various baselines, including CoTTA which is specifically designed for continual TTA, demonstrates the effectiveness of our probabilistic framework where the source model’s posterior induces a regularizer and the data-driven resetting helps make an informed selection of weights to reset/keep. Moreover, the contribution of the regularizer term (induced by the source model’s posterior) is evident from the superior performance of PETAL over CoTTA, since CoTTA is a special case of PETAL without the regularizer term.

Gradually changing corruptions: Following [41], we evaluate PETAL in the setting where the severity of corruption changes gradually before the change in corruption type. When the corruption type changes, the severity level is lowest, and thus, domain shift is gradual. In addition, distribution shifts within each corruption type are also gradual. Refer to the Appendix for details.

We report the average error and average Brier score over 10 randomly shuffled orders of corruptions. We can observe in Table 3 that our approach PETAL performs better than other approaches in terms of error and Brier score.

Table 4. CIFAR100-to-CIFAR100C results with the most severe level of corruption, 5, averaged over all corruption types. Our method surpasses all baselines in terms of NLL and Brier uncertainty estimation measures.

Method	Source	BN Adapt	TENT	CoTTA	PETAL (FIM)
NLL	2.4945	1.3937	7.3789	1.2767 (0.0006)	1.2206 (0.0011)
Brier	0.6704	0.4744	1.1015	0.4430 (0.0003)	0.4317 (0.0003)

5.6. CIFAR100-to-CIFAR100C Results

To illustrate the efficacy of the proposed approach, we conduct an evaluation on the more challenging CIFAR100-to-CIFAR100C task. In Table 2, we compare the results with Source, BN Adapt, Pseudo-label, and TENT-continual approaches. We observe that TENT-continual performs better initially, but as new domains arrive continually, the performance degrades drastically in the long term. With an average error rate of 31.46%, our proposed approach PETAL (FIM) consistently outperforms the other approaches.

To measure the ability of uncertainty estimation of our approach, we compare with the other approaches in Table 4 in terms of NLL and Brier score. We obtain both NLL and Brier score for corruption with a severity level of 5 and average over all corruption types. Our approach performs better than all other approaches in terms of average NLL and average Brier, demonstrating the ability of our approach to improve the uncertainty estimation. Moreover, the FIM based restore outperforms stochastic restore for most corruption types, illustrating the utility of data-driven resetting.

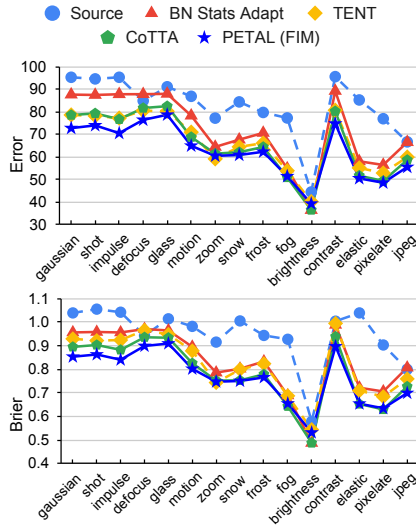


Figure 2. ImageNet-to-ImageNetC results averaged over 10 different corruption orders with level 5 corruption severity.

5.7. ImageNet-to-ImageNetC Results

In Table 5, we investigate the performance of our proposed approach for ImageNet-to-ImageNetC task with 10

different sequences of corruption. We obtain the performance scores by averaging over all corruption types and all corruption orders. In terms of average error, average NLL, and average Brier, our approach performs better than the other approaches. In Fig. 2, we compare the performance by averaging over 10 different corruption orders in terms of error rate and Brier score. For most corruption types, our approach performs better than other existing approaches.

Table 5. ImageNet-to-ImageNetC results averaged over all corruption types and over 10 diverse corruption orders (highest corruption severity level 5).

Method	Source	BN Adapt	TENT	CoTTA	PETAL (FIM)
Error (%)	82.35	72.07	66.52	63.18	62.71
NLL	5.0701	3.9956	3.6076	3.3425	3.3252
Brier	0.9459	0.8345	0.8205	0.7681	0.7663

5.8. ImageNet-to-ImageNet3DCC Results

For ImageNet-to-ImageNet3DCC dataset, we experiment with 10 different random sequences of corruptions. We provide the results in Table 6 by averaging over 10 random sequences of corruptions and 12 corruption types at severity level 5. PETAL consistently outperforms all other approaches in terms of error rate, NLL, and Brier score.

Table 6. ImageNet-to-ImageNet3DCC results averaged over all corruption types and over 10 diverse corruption orders (highest corruption severity level 5).

Method	Source	BN Adapt	TENT	CoTTA	PETAL (FIM)
Error (%)	69.21	67.32	95.93	59.91	59.61
NLL	5.0701	3.9956	3.6076	3.3425	3.3252
Brier	3.9664	3.7163	19.0408	3.2636	3.2560

6. Conclusion

We proposed a probabilistic framework for lifelong TTA using a partly data-driven prior. Addressing the problem via the probabilistic perspective naturally gives rise to the student-teacher framework along with a regularizer that captures the source domain knowledge. In lifelong TTA, we have demonstrated that our principled use of an approximate training posterior surpasses prior heuristic approaches. Our proposed approach also provides more reliable uncertainty estimates demonstrated with better NLL and Brier score. Further, we developed a Fisher information matrix based parameter restoration, which is driven by the data to improve upon existing stochastic restore. In terms of error rate, NLL and Brier score, PETAL yields state-of-the-art results across CIFAR10-to-CIFAR10C, CIFAR100-to-CIFAR100C, ImageNet-to-ImageNetC, and ImageNet-to-ImageNet3DCC benchmark tasks.

Acknowledgments: The authors acknowledge support from the Qualcomm Innovation Fellowship and travel grant support from the Research-I Foundation at IIT Kanpur.

References

- [1] Dhanajit Brahma, Vinay Kumar Verma, and Piyush Rai. Hypernetworks for continual semi-supervised learning. In *International Workshop on Continual Semi-Supervised Learning, IJCAI*, 2021. 6
- [2] Glenn W Brier et al. Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78(1):1–3, 1950. 7
- [3] Chaoqi Chen, Weiping Xie, Wenbing Huang, Yu Rong, Xinghao Ding, Yue Huang, Tingyang Xu, and Junzhou Huang. Progressive feature alignment for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 627–636, 2019. 1
- [4] Francesco Croce, Maksym Andriushchenko, Vikash Sehwag, Edoardo Debenedetti, Nicolas Flammarion, Mung Chiang, Prateek Mittal, and Matthias Hein. Robustbench: a standardized adversarial robustness benchmark. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. 6
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. IEEE, 2009. 6
- [6] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. 5
- [7] Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007. 7
- [8] Taesik Gong, Jongheon Jeong, Taewon Kim, Yewon Kim, Jinwoo Shin, and Sung-Ju Lee. Robust continual test-time adaptation: Instance-aware bn and prediction-balanced memory. *arXiv preprint arXiv:2208.05117*, 2022. 5
- [9] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. *Advances in neural information processing systems*, 17, 2004. 3
- [10] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR, 2017. 1
- [11] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2019. 1, 6
- [12] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pages 1989–1998. Pmlr, 2018. 5
- [13] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. In *34th Conference on Uncertainty in Artificial Intelligence 2018, UAI 2018*, pages 876–885. Association For Uncertainty in Artificial Intelligence (AUAI), 2018. 6
- [14] Oğuzhan Fatih Kar, Teresa Yeo, Andrei Atanov, and Amir Zamir. 3d common corruptions and data augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18963–18974, 2022. 6
- [15] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. 4, 6
- [16] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, pages 5637–5664. PMLR, 2021. 1
- [17] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. 6
- [18] Jogendra Nath Kundu, Naveen Venkat, R Venkatesh Babu, et al. Universal source-free domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4544–4553, 2020. 5
- [19] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, page 896, 2013. 2, 3, 7
- [20] Rui Li, Qianfen Jiao, Wenming Cao, Hau-San Wong, and Si Wu. Model adaptation: Unsupervised domain adaptation without source data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9641–9650, 2020. 5
- [21] Yanghao Li, Naiyan Wang, Jianping Shi, Jiaying Liu, and Xiaodi Hou. Revisiting batch normalization for practical domain adaptation. *arXiv preprint arXiv:1603.04779*, 2016. 7
- [22] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. 6
- [23] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *International Conference on Machine Learning*, pages 6028–6039. PMLR, 2020. 5
- [24] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *International conference on machine learning*, pages 97–105. PMLR, 2015. 5
- [25] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30, 2017. 6
- [26] Wesley J Maddox, Pavel Izmailov, Timur Garipov, Dmitry P Vetrov, and Andrew Gordon Wilson. A simple baseline for bayesian uncertainty in deep learning. *Advances in Neural Information Processing Systems*, 32, 2019. 4, 6
- [27] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989. 1, 6

- [28] 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. 1
- [29] Shuaicheng Niu, Jiayang Wu, Yifan Zhang, Yaofo Chen, Shijian Zheng, Peilin Zhao, and Minghui Tan. Efficient test-time model adaptation without forgetting. In *International Conference on Machine Learning*. PMLR, 2022. 5
- [30] Sinno Jialin Pan, Ivor W Tsang, James T Kwok, and Qiang Yang. Domain adaptation via transfer component analysis. *IEEE transactions on neural networks*, 22(2):199–210, 2010. 5
- [31] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019. 1, 6
- [32] Viraj Prabhu, Shivam Khare, Deeksha Kartik, and Judy Hoffman. Sentry: Selective entropy optimization via committee consistency for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8558–8567, 2021. 1
- [33] Roger Ratcliff. Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychological review*, 97(2):285, 1990. 1, 6
- [34] Steffen Schneider, Evgenia Rusak, Luisa Eck, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. Improving robustness against common corruptions by covariate shift adaptation. *Advances in Neural Information Processing Systems*, 33:11539–11551, 2020. 5, 7
- [35] 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 *International conference on machine learning*, pages 9229–9248. PMLR, 2020. 1
- [36] Rohan Taori, Achal Dave, Vaishaal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig Schmidt. Measuring robustness to natural distribution shifts in image classification. *Advances in Neural Information Processing Systems*, 33:18583–18599, 2020. 1
- [37] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7472–7481, 2018. 5
- [38] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *ICLR*, 2021. 1, 5, 6, 7
- [39] Liyuan Wang, Kuo Yang, Chongxuan Li, Lanqing Hong, Zhenguo Li, and Jun Zhu. Ordisco: Effective and efficient usage of incremental unlabeled data for semi-supervised continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5383–5392, 2021. 6
- [40] Mei Wang and Weihong Deng. Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153, 2018. 5
- [41] Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual test-time domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 1, 2, 4, 5, 6, 7
- [42] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10687–10698, 2020. 2, 3
- [43] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. 6
- [44] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4085–4095, 2020. 5
- [45] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *British Machine Vision Conference 2016*. British Machine Vision Association, 2016. 6
- [46] Aurick Zhou and Sergey Levine. Training on test data with bayesian adaptation for covariate shift. *NeurIPS*, 2021. 5

Appendix: A Probabilistic Framework for Lifelong Test-Time Adaptation

Dhanajit Brahma

Indian Institute of Technology Kanpur

dhanajit@cse.iitk.ac.in

Piyush Rai

Indian Institute of Technology Kanpur

piyush@cse.iitk.ac.in

In this Appendix, we provide further details about benchmark datasets, computational resources, training details, evaluation metrics and additional experimental analysis.

1. Benchmark Datasets

CIFAR10C and CIFAR100C datasets are corrupted versions of the standard CIFAR10 and CIFAR100 [7] datasets, respectively. ImageNetC [5] and ImageNet3DCC [6] are both corrupted versions of the standard ImageNet [4] dataset.

Both CIFAR10C and CIFAR100C datasets contain 10,000 images per corruption type, making 150,000 images in total for each dataset. For ImageNetC dataset, there are 50,000 images for each corruption type. CIFAR10C, CIFAR100C, and ImageNet-C datasets consists of 15 diverse corruption types with 4 extra corruption types for validation. Every corruption has five different levels of severity. The various corruptions, along with a brief description, are as follows: i. Gaussian noise: often occurs in low light; ii. Shot noise: electronic noise due to discrete nature of light; iii. Impulse noise: color equivalent of salt-and-pepper noise, and it may be due to bit errors; iv. Defocus blur: caused when a photograph is taken out of focus; v. Frosted Glass Blur: due to an image through a frosted glass window; vi. Motion blur: occurs when a camera is rapidly moving; vii. Zoom blur: happens when a camera quickly approaches an object; viii. Snow: visually an obscuring kind of precipitation; ix. Frost arises when ice crystals adhere to windows; x. Fog: Objects are cloaked in fog, which is rendered using the diamond-square algorithm; xi. Brightness: varying brightness with the sunshine intensity; xii. Contrast: depends on the lighting conditions and color of the photographed item; xiii. Elastic transformations: Small image areas are stretched or contracted via elastic transformations; xiv. Pixelation: happens when a low-resolution picture is upsampled; xv. JPEG: lossy image compression that results in compression artifacts.

Recently proposed by [6], the ImageNet 3D Common Corruptions (ImageNet3DCC) dataset exploits the geometry of the scene in transformations, resulting in more realistic corruptions. The ImageNet3DCC dataset has 50,000

images for each corruption type. It consists of 12 different types of corruption, each with 5 severity levels. The corruptions are as follows: i. Near focus: changing the focus area to the near portion of the scene at random; ii. Far focus: randomly change the focus to the scene’s far part; iii. Bit error: caused by an imperfect video transmission channel; iv. Color quantization: decreases the RGB image’s bit depth; v. Flash: caused by positioning a light source near the camera; vi. Fog 3D: generated by using a standard optical model for fog; vii. H.265 ABR: codec H.265 for compression with Average Bit Rate control mode; viii. H.265 CRF: codec H.265 for compression with Constant Rate Factor (CRF) control mode; ix. ISO noise: noise using a Poisson-Gaussian distribution; x. Low light: simulate low-light imaging setting by lowering the pixel intensities and adding Poisson-Gaussian distributed noise; xi. XY-motion blur: the main camera is moving along the image XY-plane; xii. Z-motion blur: the main camera is moving along the image Z-axis.

These datasets are developed to serve as benchmarks for measuring robustness of classification models.

2. Computational Resource Details

All the experiments were run locally on a GPU server with Nvidia Titan RTX GPUs with 24 GB memory, Intel(R) Xeon(R) Silver 4110 CPU with 128 GB RAM, and Ubuntu 18.04.6 LTS OS. PETAL is implemented using PyTorch version 1.10.0. More details about the other relevant libraries are provided in the source code.

For CIFAR-10 experiments, further training the pre-trained WideResNet-28 model from RobustBench benchmark [3] using the source domain training data takes approximately 20 minutes. For CIFAR-100 experiments, we further train the pre-trained ResNeXt-29 [10] model from RobustBench benchmark [3] using the source domain training data for approximately 15 minutes. For ImageNet experiments, training the pre-trained ResNet50 model from RobustBench benchmark [3] further using the source domain training data takes approximately 5 hours.

For the CIFAR10-to-CIFAR10C adaptation experiment on the highest severity level of 5 using our approach

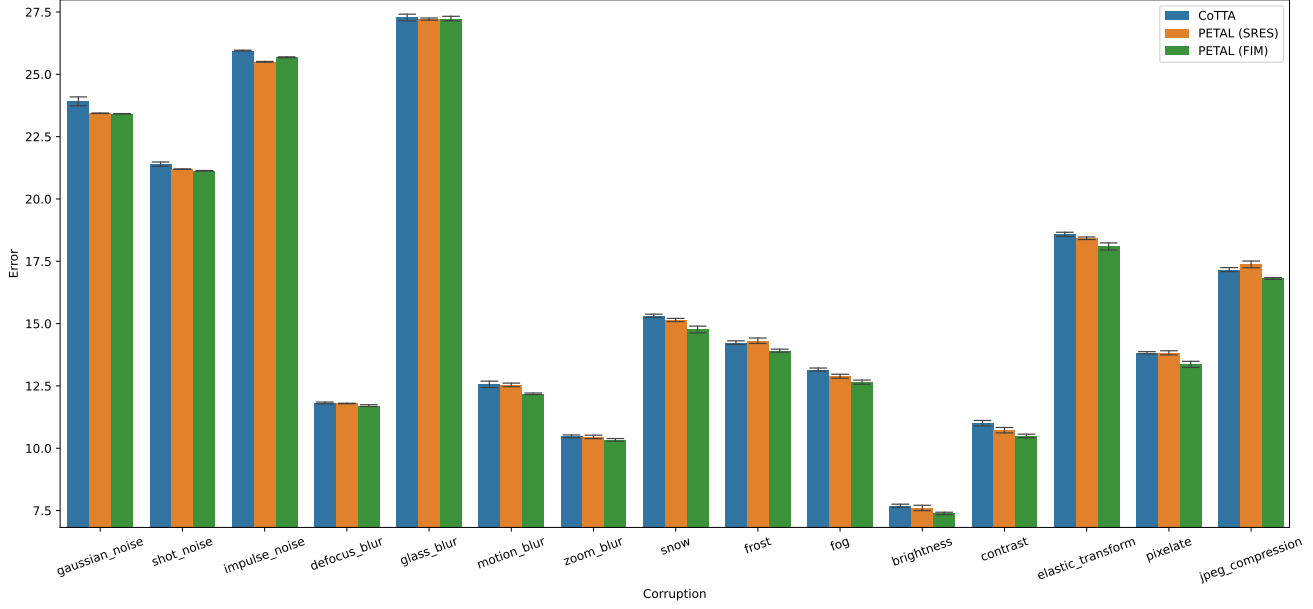


Figure 1. CIFAR10-to-CIFAR10C Results for the corruption order as depicted in the figure with the corruption of severity level 5. The error rate is averaged over 5 runs. The standard deviation is depicted in the bar plot. PETAL (FIM) performs better in most of the settings.

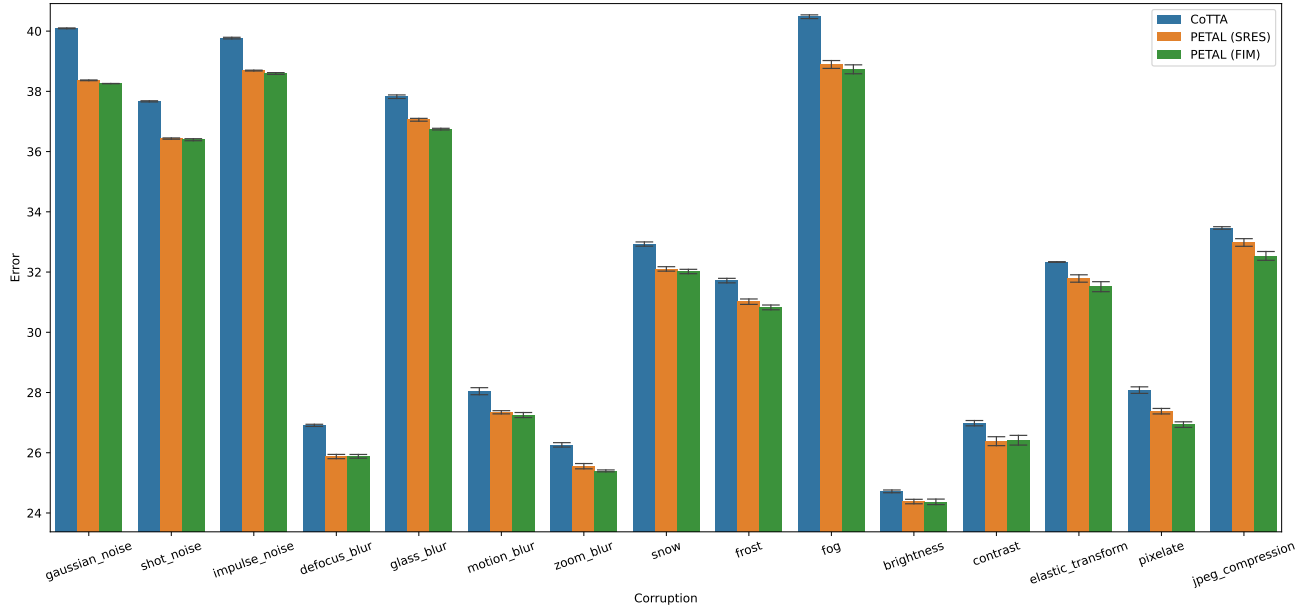


Figure 2. CIFAR100-to-CIFAR100C Results for the corruption order as depicted in the figure with the corruption of severity level 5. The error rate is averaged over 5 runs. The standard deviation is depicted in the bar plot. PETAL (FIM) performs better in most of the settings.

PETAL, the total time taken for adapting the pre-trained WideResNet-28 model on 15 corruption types takes approximately 3 hours. In CIFAR10-to-CIFAR10C gradual experiment, since the severity level changes gradually before the change in corruption type, the total number of different corruption-severity level pairs is 131. Further, the experiment is conducted for 10 random orders of corruption se-

quences and so the total time taken is about 6 days and 19 hours.

CIFAR100-to-CIFAR100C adaptation experiment using pre-trained ResNeXt-29 for the highest severity level of 5 takes approximately 1 hour and 10 minutes.

For ImageNet-to-ImageNetC adaptation experiment on the highest severity level of 5, the total time taken for

adapting the pre-trained ResNet50 model on 15 corruption types repeated for 10 diverse corruption sequences is about 2 hours and 30 minutes in total. The ImageNet-to-ImageNet3DCC adaptation experiment on the highest severity level of 5 for 12 corruptions took a total time of about 2 hours and 15 minutes.

3. Training Details

3.1. Training on Source Domain data

To obtain the approximate posterior using the source domain training data, we use SWAG-D [8] posterior. This requires further training of the pre-trained models on the source domain training data for a few more epochs.

For CIFAR-10 experiments, the pre-trained WideResNet-28 model from RobustBench benchmark [3] is used, which is already trained using CIFAR-10 training data. We further train it using SGD with momentum for 5 epochs with a learning rate of $8e-4$, collecting iterates for SWAG-D [8] once in each epoch. Similarly, for CIFAR-100 experiments, we use the pre-trained ResNeXt-29 [10] model from RobustBench benchmark [3] by training it further using SGD with momentum for 5 more epochs with a learning rate of $8e-4$. For ImageNet experiments, we further train the pre-trained standard ResNet50 model from RobustBench benchmark [3] for 2 more epochs and collect 4 iterates per epoch with a learning rate of $1e-4$ for obtaining SWAG-D approximate posterior.

3.2. Adaptation on Target Domain Test data

For the online lifelong test-time adaptation (TTA), we use the same hyperparameters mentioned in CoTTA [9] using a learning rate of 0.001 with Adam optimizer. Based on [2] and [9], we apply random augmentations that include color jitter, random affine, Gaussian blur, random horizontal flip, and Gaussian noise. For stochastic restore probability, we use the same value used by CoTTA, i.e., 0.01 for CIFAR10-to-CIFAR10C/CIFAR100-to-CIFAR100C and 0.001 for ImageNet-to-ImageNetC experiments. We tune the FIM based parameter restoration quantile value δ using the extra four validation corruptions and use 0.03 for CIFAR10-to-CIFAR10C/CIFAR100-to-CIFAR100C, and use 0.003 for ImageNet-to-ImageNetC and ImageNet-to-ImageNet3DCC.

To adapt the model, we optimize the following training objective from Equation 7:

$$\max_{\theta} \log q(\theta) - \frac{\bar{\lambda}}{M} \sum_{m=1}^M H^{xe}(y', y|\bar{\mathbf{x}}) \quad (11)$$

Here, $q(\theta)$ is the approximate posterior density obtained using SWAG-D [8]. $\bar{\lambda}$ is a hyperparameter that determines

the importance of the cross-entropy minimization term relative to the posterior term. For the implementation, putting $\alpha = 1/\bar{\lambda}$, we rewrite the training objective as follows:

$$\max_{\theta} \alpha \log q(\theta) - \frac{1}{M} \sum_{m=1}^M H^{xe}(y', y|\bar{\mathbf{x}}) \quad (12)$$

For all the experiments, we tune the α hyperparameter using the corresponding extra four validation corruptions and use one of the values from $\{1e-6, 1e-7, 1e-9, 1e-10, 5e-10, 1e-11, 5e-11, 1e-12\}$ for which the average error is the lowest. We follow CoTTA [9] for setting the hyperparameters: augmentation threshold for confident predictions τ and exponential moving average smoothing factor $\pi = 0.999$. CoTTA [9] discusses choice of τ in detail in their supplementary.

Gradually changing corruptions: In this setting, the severity of corruption changes gradually before the change in corruption type. For example, if A_s , B_s , and C_s are three different corruptions of severity level s , then there are 23 corruption-severity pairs in total. They arrive as follows:

$$\begin{aligned} & A_5 \rightarrow A_4 \rightarrow A_3 \rightarrow A_2 \rightarrow A_1 \xrightarrow[\text{corruption}]{\text{change}} B_1 \rightarrow B_2 \rightarrow B_3 \rightarrow B_4 \\ & \rightarrow B_5 \rightarrow B_4 \rightarrow B_3 \rightarrow B_2 \rightarrow B_1 \xrightarrow[\text{corruption}]{\text{change}} C_1 \rightarrow C_2 \rightarrow C_3 \rightarrow C_4 \\ & \rightarrow C_5 \rightarrow C_4 \rightarrow C_3 \rightarrow C_2 \rightarrow C_1 \end{aligned}$$

When the corruption type changes, the severity level is 1, and thus, the domain shift is gradual. In addition, distribution shifts within each corruption type are also gradual.

4. Evaluation Metrics

Let y'_n be the predicted probability vector and y_n be the true label ($y_{ni} = 1$ if i is the true class label, else $y_{ni} = 0$) for input x_n . Let the number of examples be N and the number of classes be D .

4.1. Error

The average error rate is defined as:

$$\text{Error} = \frac{1}{N} \sum_{n=1}^N \mathbb{I}(y'_n \neq y_n)$$

where $\mathbb{I}()$ is the indicator function.

4.2. Brier Score

The average Brier score [1] is defined as:

$$\text{Brier score} = \frac{1}{N} \sum_{n=1}^N \sum_{i=1}^D (y'_{ni} - y_{ni})^2$$

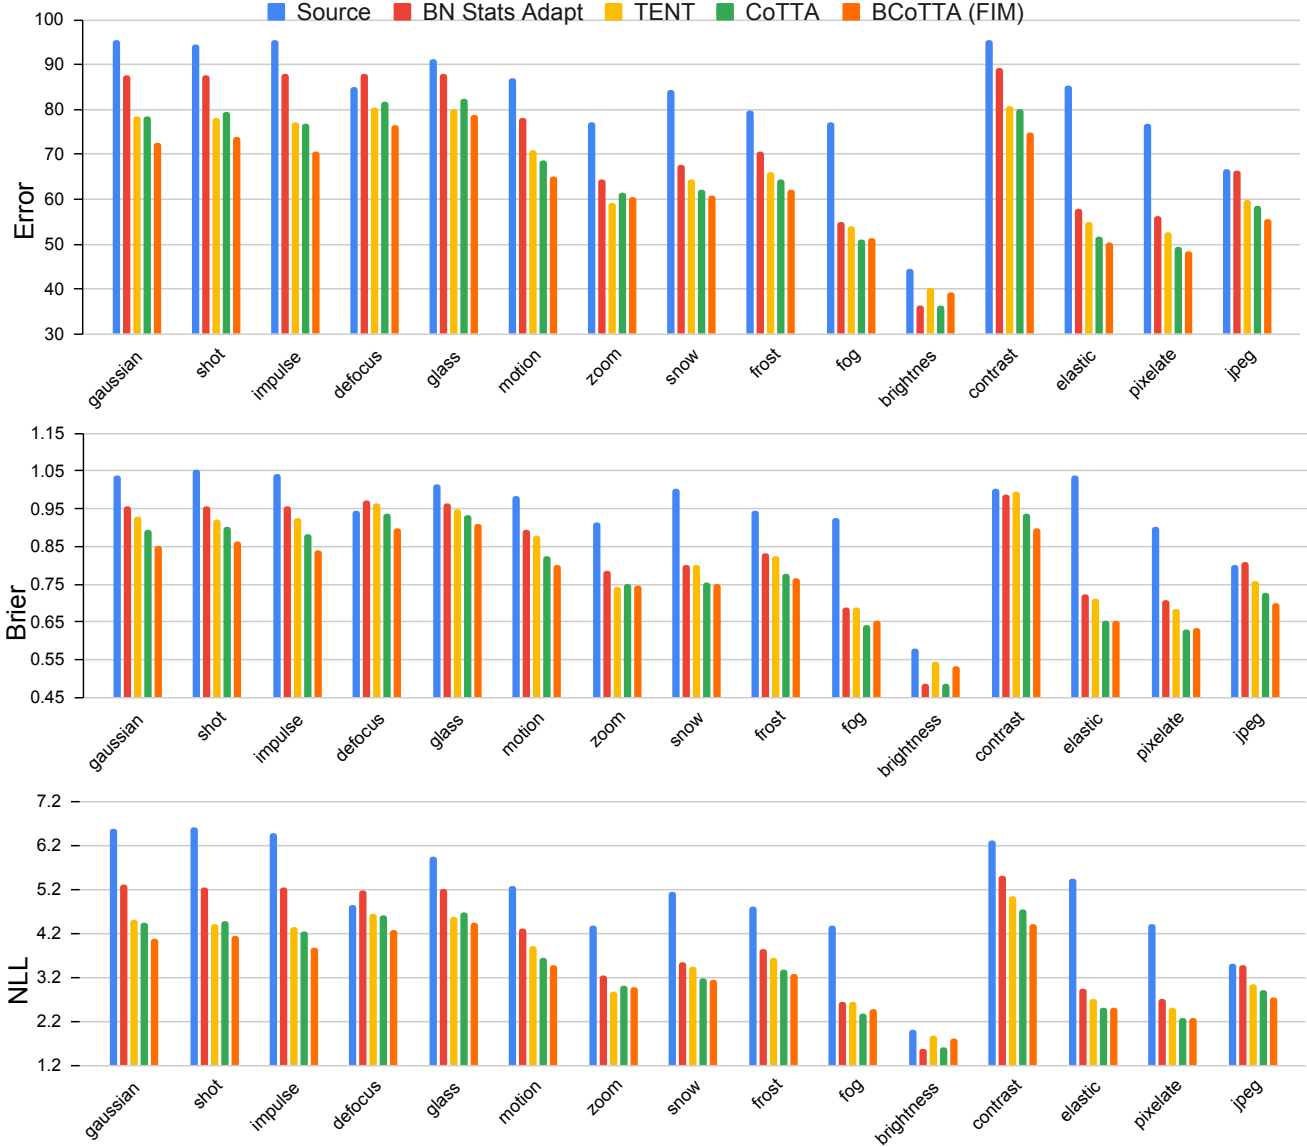


Figure 3. ImageNet-to-ImageNetC results averaged over 10 different corruption sequences with corruption of severity level 5. PETAL (FIM) outperforms other approaches for most of the corruptions.

4.3. Negative Log-Likelihood

The average negative log-likelihood (NLL) is defined as:

$$\text{NLL} = -\frac{1}{N} \sum_{n=1}^N \sum_{i=1}^D (y_{ni} \log y'_{ni})$$

5. Additional Experimental Results

5.1. CIFAR10-to-CIFAR10C Results

Figure 1 shows the results of CIFAR10-to-CIFAR10C experiments in which the corruption types arrive in the depicted order. The standard deviations are depicted in the figure. PETAL performs better in most of the settings. More-

over, PETAL (FIM) performs better than PETAL (SRES) demonstrating the effectiveness of data-driven FIM based parameter resetting.

5.2. CIFAR100-to-CIFAR100C Results

The CIFAR100-to-CIFAR100C adaptation results are shown in Figure 2 in which the corruption types arrive in the depicted order. The standard deviations are depicted in the figure. We observe that PETAL outperforms most of the approaches.

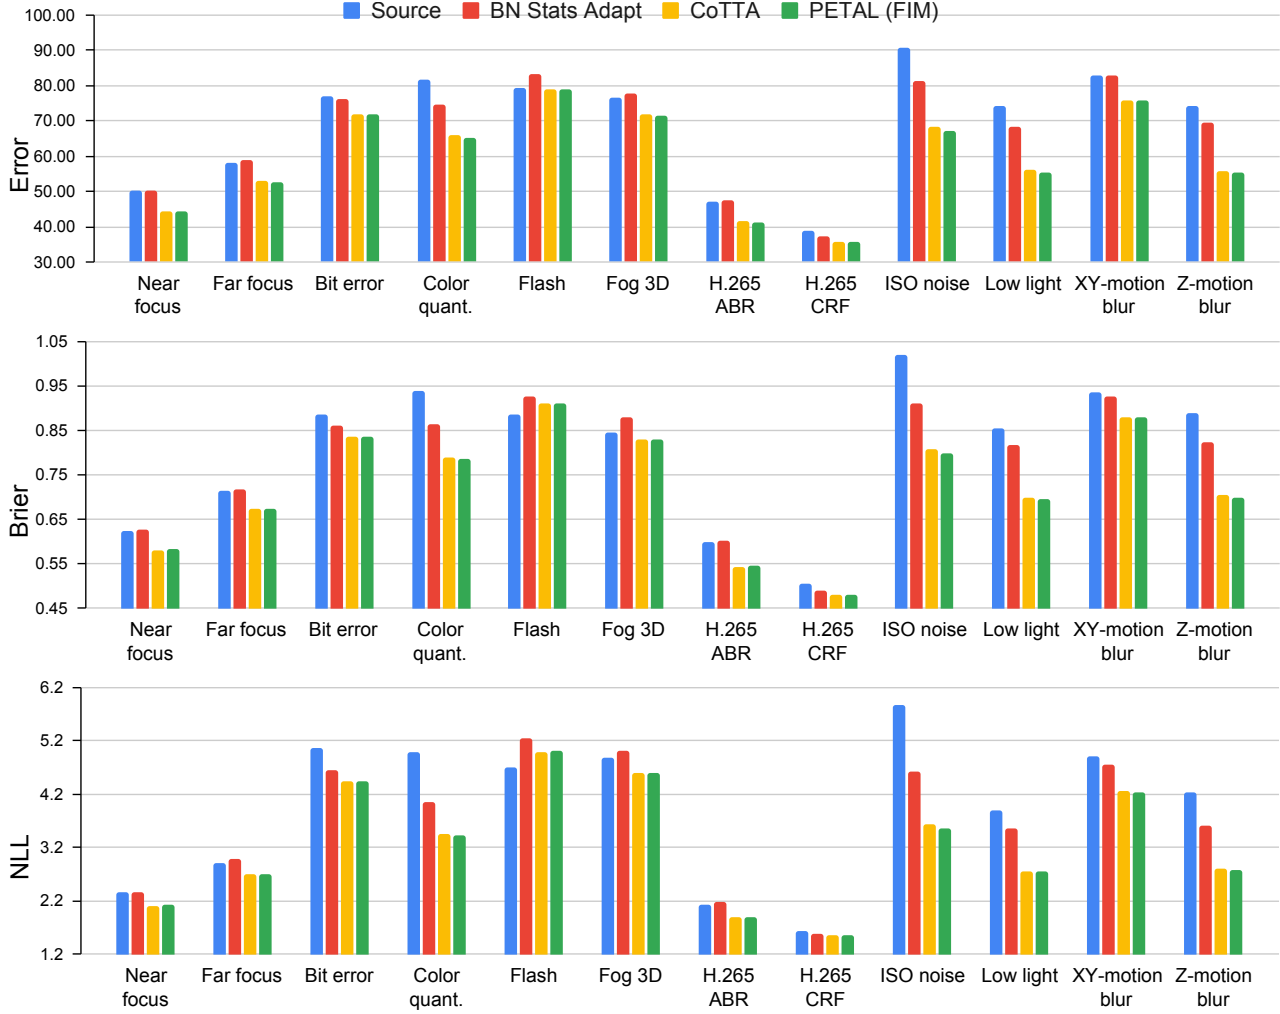


Figure 4. ImageNet-to-ImageNet3DCC results averaged over 10 different corruption sequences with corruption of severity level 5. PETAL (FIM) outperforms other approaches for most of the corruptions.

5.3. ImageNet-to-ImageNetC Results

For ImageNet-to-ImageNetC experiments, we provide the error rate, Brier score, and NLL in Figure 3. The experiments are conducted for 10 diverse corruption sequences. The numbers in Figure 3 are obtained by averaging over the 10 different orders of corruption sequences. PETAL (FIM) outperforms CoTTA and PETAL (SRES) for most of the corruptions in terms of average error rate, average Brier score, and average NLL.

5.4. ImageNet-to-ImageNet3DCC Results

We report the results for ImageNet-to-ImageNet3DCC experiments in Fig. 4. The numbers are averaged over 10 random corruption orders. Our approach PETAL performs better than other approaches for most of the corruption types.

Results demonstrate that our probabilistic approach, in

which the source model’s posterior leads to a regularizer and the data-driven reset helps make an informed choice of weights to reset/keep, significantly outperforms a number of baselines, including CoTTA, which is specifically designed for continual test time adaptation.

5.5. Experiments with Similar Resetting Rate

We compare the performance of PETAL (FIM) and PETAL (SRes) with a similar resetting rate of 0.03. For CIFAR10-to-CIFAR10C, the average error rate over 5 runs for PETAL (SRes) with 0.03 restoring rate is 16.51 ± 0.02 , whereas that of PETAL (FIM) is 15.95 ± 0.04 . This indicates that FIM enables more parameters to be restored than stochastic restore (SRes).

5.6. Effectiveness of Various Components

In Table 1 and Table 2, the better performance of PETAL (FIM) than PETAL (S-Res) highlights the contribution of FIM for improved performance over the stochastic restore. For CIFAR10-to-CIFAR10C, the average error rate for PETAL **with no parameter restore** over 5 runs is 16.37 ± 0.03 , which is worse than having a parameter restore. In addition, the contribution of the regularizer term (caused by the source model’s posterior) is clear from the better performance of PETAL over CoTTA [9], given that CoTTA is a special case of PETAL with no regularizer term.

5.7. Limitation and Future Scope

A limitation of our proposed approach is that we use an approximate posterior obtained using SWAG-diagonal, and we can experiment with better approximations for the posterior in future work. Further, we can explore the problem setting of lifelong TTA where we have a small number of examples per batch.

References

- [1] Glenn W Brier et al. Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78(1):1–3, 1950. 3
- [2] Gilad Cohen and Raja Giryes. Katana: Simple post-training robustness using test time augmentations. *arXiv preprint arXiv:2109.08191*, 2021. 3
- [3] Francesco Croce, Maksym Andriushchenko, Vikash Sehwag, Edoardo Debenedetti, Nicolas Flammarion, Mung Chiang, Prateek Mittal, and Matthias Hein. Robustbench: a standardized adversarial robustness benchmark. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. 1, 3
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. IEEE, 2009. 1
- [5] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2019. 1
- [6] Oğuzhan Fatih Kar, Teresa Yeo, Andrei Atanov, and Amir Zamir. 3d common corruptions and data augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18963–18974, 2022. 1
- [7] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. 1
- [8] Wesley J Maddox, Pavel Izmailov, Timur Garipov, Dmitry P Vetrov, and Andrew Gordon Wilson. A simple baseline for bayesian uncertainty in deep learning. *Advances in Neural Information Processing Systems*, 32, 2019. 3
- [9] Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual test-time domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 3, 6
- [10] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. 1, 3