Robust Continual Test-time Adaptation: Instance-aware BN and Prediction-balanced Memory

Anonymous Author(s)
Affiliation
Address
email

Abstract

Test-time adaptation (TTA) is an emerging paradigm that addresses distributional shifts between training and testing phases without additional data acquisition or labeling cost; only unlabeled test data streams are used for continual model adaptation. Previous TTA schemes assume that the test samples are independent and identically distributed (i.i.d.), even though they are often temporally correlated (non-i.i.d.) in application scenarios, e.g., autonomous driving. We discover that most existing TTA methods fail dramatically under such scenarios. Motivated by this, we present a new test-time adaptation scheme that is robust against non-i.i.d. test data streams. Our novelty is mainly two-fold: (a) Instance-Aware Batch Normalization (IABN) that corrects normalization for out-of-distribution samples, and (b) Prediction-balanced Reservoir Sampling (PBRS) that simulates i.i.d. data stream from non-i.i.d. stream in a class-balanced manner. Our evaluation with various datasets, including real-world non-i.i.d. streams, demonstrates that the proposed robust TTA not only outperforms state-of-the-art TTA algorithms in the non-i.i.d. setting, but also achieves comparable performance to those algorithms under the i.i.d. assumption.

1 Introduction

While deep neural networks (DNNs) have been successful in several applications, their performance degrades under distributional shifts between the training data and test data [31]. This distributional shift hinders DNNs from being widely deployed in many risk-sensitive applications, such as autonomous driving, medical imaging, and mobile health care, where new types of test data unseen during training could result in undesirable disasters. For instance, Tesla Autopilot has caused 12 "deaths" until recently [2]. To address this problem, *test-time adaptation* (TTA) aims to adapt DNNs to the target/unseen domain with only unlabeled test data streams, without any additional data acquisition or labeling cost. Recent studies reported that TTA is a promising, practical direction to mitigate distributional shifts [28, 32, 40, 4, 43].

Prior TTA studies typically assume (implicitly or explicitly) that a target test sample $\mathbf{x}_t$ at time t and the corresponding ground-truth label y_t (unknown to the learner) are independent and identically distributed (i.i.d.) following a target domain, i.e., $(\mathbf{x}_t, y_t)$ is drawn independently from a time-invariant distribution $P_{\mathcal{T}}(\mathbf{x}, y)$. However, the distribution of online test samples often changes across the time axis, i.e., $(\mathbf{x}_t, y_t) \sim P_{\mathcal{T}}(\mathbf{x}, y \,|\, t)$ in many applications; for instance, AI-powered self-driving car's object encounter will be dominated by cars while driving on the highway, but less dominated by them on downtown where other classes such as pedestrians and bikes are visible. In human activity recognition, some activities last for a short term (e.g., a fall down), whereas certain activities last longer (e.g., a sleep). Figure 1 illustrates that some data distributions in the real world, such as autonomous driving and human activity recognition, are often temporally correlated. Considering that most

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

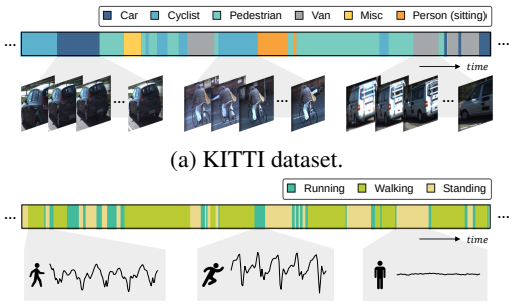

(a) KITTI dataset.

(b) HARTH dataset.

Figure 1: Illustration of test sample distributions along the time axis from real-world datasets: (a) autonomous driving (KITTI [8]) and (b) human activity recognition (HARTH [24]). They are temporally correlated.

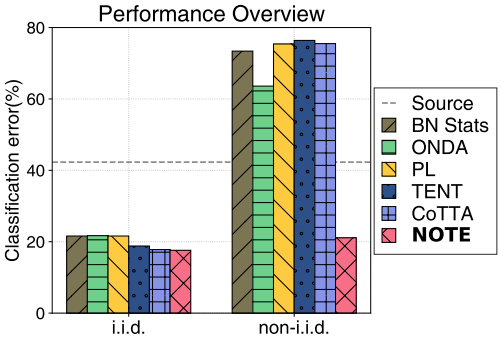

Figure 2: Average classification error (%) of existing TTA methods and our method (NOTE) on CIFAR10-C [12]. The error rates significantly increase under the non-i.i.d. setting compared with the i.i.d. setting. Lower is better.

existing TTA algorithms simply use an incoming batch of test samples for adaptation [28, 32, 40, 43], the model might be biased towards these imbalanced samples under the temporally correlated test streams. Figure 2 compares the performance of the state-of-the-art TTA algorithms under the i.i.d. and non-i.i.d.[1] conditions. While the TTA methods perform well under the i.i.d. assumption, their errors increase under the non-i.i.d. case. Adapting to temporally correlated test data results in overfitting to temporal distributions, which in turn harms the generalization of the model.

Motivated by this, we present a NOn-i.i.d. TEst-time adaptation scheme, *NOTE*, that consists of two components: (a) Instance-Aware Batch Normalization (IABN) and (b) Prediction-Balanced Reservoir Sampling (PBRS). First, we propose a novel normalization layer, IABN, that eliminates the dependence on temporally correlated data for adaptation while being robust to distribution shifts. IABN detects out-of-distribution instances *sample by sample* and corrects via instance-aware normalization. The key idea of IABN is synthesizing Batch Normalization (BN) [15] with Instance Normalization (IN) [36] in a unique way; it calculates how different the learned knowledge (BN) is from the current observation (IN) and corrects the normalization by the deviation between IN and BN. Second, we present PBRS that resolves the problem of overfitting to non-i.i.d. samples by mimicking i.i.d. samples from non-i.i.d. streams. By utilizing predicted labels of the model, PBRS aims for both time-uniform sampling and class-uniform sampling from the non-i.i.d. streams, and stores the 'simulated' i.i.d. samples in memory. With the i.i.d.-like batch in the memory, PBRS enables the model to adapt to the target domain without being biased to temporal distributions.

We evaluate NOTE with state-of-the-art TTA baselines [28, 32, 21, 26, 4, 43] on multiple datasets, including common TTA benchmarks (CIFAR10-C and CIFAR100-C [12]) and real-world non-i.i.d. datasets (KITTI [8], HARTH [24], and ExtraSensory [37]). Our results suggest that NOTE not only significantly outperforms the baselines under non-i.i.d. test data, e.g., it achieves a 21.1% error rate on CIFAR10-C which is on average 15.1% lower than the state-of-the-art method [4], but also shows comparable performance even under the i.i.d. assumption, e.g., 17.6% error on CIFAR10-C where the best baseline [43] achieves 17.8% error. Our ablative study demonstrates the individual effectiveness of IABN and PBRS and further highlights their synergy when jointly used.

Finally, we summarize the key characteristics of NOTE. First, NOTE is a batch-free inference algorithm (requiring a single instance for inference), different from the state-of-the-art TTA algorithms [28, 32, 40, 4, 43] where a batch of test data is necessary for inference to estimate normalization statistics (mean and variance). Second, while some recent TTA methods leverage augmentations to improve performance at the cost of additional forwarding passes [34, 43], NOTE requires only a single forwarding pass. NOTE updates only the normalization statistics and affine parameters in IABN, which is, e.g., approximately 0.02% of the total trainable parameters in ResNet18 [11]. Third, NOTE's additional memory overhead is negligible. It merely stores predicted labels of test data to run PBRS. These characteristics make NOTE easy to apply to any existing AI system and particularly, are beneficial in latency-sensitive tasks such as autonomous driving and human health monitoring.

[1] We use the terms *temporally correlated* and *non-i.i.d.* interchangeably in the context of test-time adaptation.

2 Background

2.1 Problem setting: test-time adaptation with non-i.i.d. streams

Test-time adaptation. Let $\mathcal{D}_\mathcal{S} = \{\mathcal{X}^\mathcal{S}, \mathcal{Y}\}$ be the data from the source domain and $\mathcal{D}_\mathcal{T} = \{\mathcal{X}^\mathcal{T}, \mathcal{Y}\}$ be the data from the target domain to adapt to. Each data instance and the corresponding ground-truth label pair $(\mathbf{x}_i, y_i) \in \mathcal{X}^\mathcal{S} \times \mathcal{Y}$ in the source domain follows a probability distribution $P_\mathcal{S}(\mathbf{x}, y)$. Similarly, each target test sample and the corresponding label at test time t, $(\mathbf{x}_t, y_t) \in \mathcal{X}^\mathcal{T} \times \mathcal{Y}$, follows a probability distribution $P_\mathcal{T}(\mathbf{x}, y)$ where y_t is unknown for the learner. The standard covariate shift assumption in domain adaptation is defined as $P_\mathcal{S}(\mathbf{x}) \neq P_\mathcal{T}(\mathbf{x})$ and $P_\mathcal{S}(y|\mathbf{x}) = P_\mathcal{T}(y|\mathbf{x})$ [31]. Unlike traditional domain adaptation that uses $\mathcal{D}_\mathcal{S}$ and $\mathcal{X}^\mathcal{T}$ collected beforehand for adaptation, test-time adaptation (TTA) continually adapts a pre-trained model $f_\theta(\cdot)$ from $\mathcal{D}_\mathcal{S}$ by utilizing only $\mathbf{x}_t$ obtained at test time t.

TTA on non-i.i.d. streams. Note that previous TTA mechanisms typically assume that each target sample $(\mathbf{x}_t, y_t) \in \mathcal{X}^\mathcal{T} \times \mathcal{Y}$ is independent and identically distributed (i.i.d.) following a time-invariant distribution $P_\mathcal{T}(\mathbf{x}, y)$. However, the data obtained at test time is non-i.i.d. in many scenarios. By non-i.i.d., we refer to distribution changes over time, i.e., $(\mathbf{x}_t, y_t) \sim P_\mathcal{T}(\mathbf{x}, y \mid t)$, which is a practical setting in many real world applications [45].

2.2 Batch normalization in TTA

Batch Normalization (BN) [15] is a widely-used training technique in deep neural networks as it reduces the internal covariant shift problem. Let $\mathbf{f} \in \mathbb{R}^{B \times C \times L}$ denote a batch of feature maps in general, where B, C, and L denote the batch size, the number of channels, and the size of each feature map, respectively. Given the statistics of the feature maps for normalization, say mean $\boldsymbol{\mu}$ and variance $\boldsymbol{\sigma}^2$, BN is *channel-wise*, i.e., $\boldsymbol{\mu}, \boldsymbol{\sigma}^2 \in \mathbb{R}^C$ and computes:

$$\mathrm{BN}(\mathbf{f}_{:,c,:}; \boldsymbol{\mu}_c, \boldsymbol{\sigma}_c^2) := \gamma \cdot \frac{\mathbf{f}_{:,c,:} - \boldsymbol{\mu}_c}{\sqrt{\boldsymbol{\sigma}_c^2 + \epsilon}} + \beta, \tag{1}$$

where γ and β are the affine parameters followed by the normalization, and $\epsilon > 0$ is a small constant for numerical stability.

Although a conventional way of computing BN in test-time is to set $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}^2$ as those estimated from *training* (or source) data, say $\bar{\boldsymbol{\mu}}$ and $\bar{\boldsymbol{\sigma}}^2$, the state-of-the-art TTA methods based on adapting BN layers [28, 32, 40, 43] instead use the statistics computed directly from the recent test batch to de-bias distributional shifts at test-time, i.e.:

$$\hat{\boldsymbol{\mu}}_c := \frac{1}{BL} \sum_{b,l} \mathbf{f}_{b,c,l}, \text{ and } \hat{\boldsymbol{\sigma}}_c^2 := \frac{1}{BL} \sum_{b,l} (\mathbf{f}_{b,c,l} - \hat{\boldsymbol{\mu}}_c)^2. \tag{2}$$

This practice is simple yet effective under distributional shifts and is thus adopted in many recent TTA studies [28, 32, 40, 43]. Based on the test batch statistics, they often further adapt the affine parameters via entropy minimization of the model outputs [40] or update the entire parameters with self-training [43].

3 Method

In the same vein as previous work [28, 32, 40], we focus on adapting BN layers in the given model to perform TTA, and this includes essentially two approaches: (a) re-calibrating (or adapting) channel-wise statistics for normalization (instead of using those learned from training), and (b) adapting the affine parameters (namely, γ and β) after the normalization with respect to a certain objective based on test samples, e.g., the entropy minimization of model outputs [40].

Under scenarios where test data are temporally correlated, however, naïvely adapting to the incoming batch of test samples [28, 32, 40, 43] could be problematic for both approaches: the batch is now more likely to (a) remove instance-wise variations that are actually useful to predict y, i.e., the "contents" rather than "styles" through normalization, and (b) include a bias in $p(y)$ rather than uniform, which can negatively affect the test-time adaptation objective such as entropy minimization.

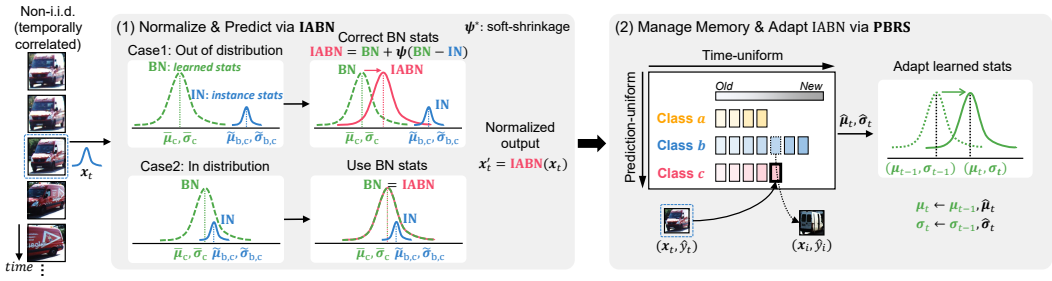

Figure 3: An overview of the proposed methodology: Instance-Aware Batch Normalization (IABN) and Prediction-Balanced Reservoir Sampling (PBRS). IABN detects out-of-distribution samples and in turn corrects the normalization for inference. PBRS manages data in a time- and prediction-uniform manner from non-i.i.d. data streams and gradually adapts IABN layers with the balanced data afterwards.

We propose two approaches to tackle each of the failure modes of adapting BN under temporal correlation. Our method consists of two components: (a) Instance-Aware Batch Normalization (IABN) (Section §3.1) to overcome the limitation of BN under distribution shift and (b) Prediction-Balanced Reservoir Sampling (PBRS) (Section §3.2) to combat with the temporal correlation of test batches. Figure 3 illustrates the overall workflow of NOTE with IABN and PBRS.

3.1 Instance-Aware Batch Normalization

As described in Section §2.2, recent TTA algorithms rely solely on the test batch to re-calculate BN statistics. We argue that this common practice does not successfully capture the feature statistics to normalize the feature map $\mathbf{f} \in \mathbb{R}^{B \times C \times L}$ under temporal correlation in the test batch B. In principle, standardizing a given feature map $\mathbf{f}_{:,c,:}$ by the statistics $\hat{\boldsymbol{\mu}}_c, \hat{\boldsymbol{\sigma}}_c^2$ computed across B and L is posited on premise that averaging information across B can marginalize out uninformative instance-wise variations for predicting y. Under temporal correlation in B, however, this assumption is no longer valid, and averaging across B may not successfully de-correlate useful information in $\mathbf{f}_{:,c,:}$ from $\boldsymbol{\mu}_c$ and $\boldsymbol{\sigma}_c^2$.

In an attempt to bypass such an "over-whitening" effect of using $\hat{\boldsymbol{\mu}}_c$ and $\hat{\boldsymbol{\sigma}}_c^2$ in test-time under temporal correlation, we propose correcting normalization statistics on a *per-sample basis*: specifically, instead of completely switching from the original statistics of $(\bar{\boldsymbol{\mu}}, \bar{\boldsymbol{\sigma}}^2)$ into $(\hat{\boldsymbol{\mu}}_c, \hat{\boldsymbol{\sigma}}_c^2)$, our proposed *Instance-Aware Batch Normalization* (IABN) considers the *instance-wise* statistics $\tilde{\boldsymbol{\mu}}, \tilde{\boldsymbol{\sigma}}^2 \in \mathbb{R}^{B,C}$ of $\mathbf{f}$ similarly to Instance Normalization (IN) [36], namely:

$$\tilde{\boldsymbol{\mu}}_{b,c} := \frac{1}{L} \sum_l \mathbf{f}_{b,c,l} \quad \text{and} \quad \tilde{\boldsymbol{\sigma}}_{b,c}^2 := \frac{1}{L} \sum_l (\mathbf{f}_{b,c,l} - \tilde{\boldsymbol{\mu}}_{b,c})^2. \tag{3}$$

We assume that $\tilde{\boldsymbol{\mu}}_{b,c}$ and $\tilde{\boldsymbol{\sigma}}_{b,c}^2$ follow the *sampling distribution* of a sample size L in $\mathcal{N}(\bar{\boldsymbol{\mu}}, \bar{\boldsymbol{\sigma}}^2)$ as the population. Then the corresponding variances for the sample mean $\tilde{\boldsymbol{\mu}}_{b,c}$ and the sample variance $\tilde{\boldsymbol{\sigma}}_{b,c}^2$ can be calculated as:

$$s_{\tilde{\boldsymbol{\mu}},c}^2 := \frac{\bar{\boldsymbol{\sigma}}_c^2}{L} \quad \text{and} \quad s_{\tilde{\boldsymbol{\sigma}}^2,c}^2 := \frac{2\bar{\boldsymbol{\sigma}}_c^4}{L-1}. \tag{4}$$

IABN corrects $(\bar{\boldsymbol{\mu}}, \bar{\boldsymbol{\sigma}}^2)$ only in cases when $\tilde{\boldsymbol{\mu}}_{b,c}$ (and $\tilde{\boldsymbol{\sigma}}_{b,c}^2$) significantly differ from $\bar{\boldsymbol{\mu}}_c$ (and $\bar{\boldsymbol{\sigma}}_c^2$). Specifically, we propose to use the following statistics for TTA:

$$\boldsymbol{\mu}_{b,c}^{\texttt{IABN}} := \bar{\boldsymbol{\mu}}_c + \psi(\tilde{\boldsymbol{\mu}}_{b,c} - \bar{\boldsymbol{\mu}}_c; \alpha s_{\tilde{\boldsymbol{\mu}},c}), \text{ and } (\boldsymbol{\sigma}_{b,c}^{\texttt{IABN}})^2 := \bar{\boldsymbol{\sigma}}_c^2 + \psi(\tilde{\boldsymbol{\sigma}}_{b,c}^2 - \bar{\boldsymbol{\sigma}}_c^2; \alpha s_{\tilde{\boldsymbol{\sigma}}^2,c}),$$

$$\text{where } \psi(x;\lambda) = \begin{cases} x - \lambda, & \text{if } x > \lambda \\ x + \lambda, & \text{if } x < -\lambda \\ 0, & \text{otherwise} \end{cases} \text{ is the soft-shrinkage function.} \tag{5}$$

Algorithm 1 Prediction-Balanced Reservoir Sampling

Input: target stream $\mathbf{x}_t \sim P_{\mathcal{T}}(\mathbf{x}|t)$,
for $t \in \{1...K\}$ **do**
  if memory is not full **then**
    Add $(\mathbf{x}_t, \hat{y}_t)$
  else
    $M \leftarrow$ majority class(es) in memory
    if $\hat{y}_t \notin M$ **then** ▷ Prediction-Balanced
      Randomly discard one instance $(\mathbf{x}_i, \hat{y}_i)$ in memory where $\hat{y}_i \in M$
      Add $(\mathbf{x}_t, \hat{y}_t)$
    else ▷ Reservoir Sampling
      $m_c \leftarrow$ number of instances in memory where $\hat{y}_i = \hat{y}_t$
      $n_c \leftarrow$ number of instances encountered so far where $\hat{y}_i = \hat{y}_t$
      Sample $p \sim$Uniform(0,1)
      if $p < m_c/n_c$ **then**
        Discard a random instance $(\mathbf{x}_i, \hat{y}_i)$ in memory where $\hat{y}_i = \hat{y}_t$
        Add $(\mathbf{x}_t, \hat{y}_t)$
      else
        Discard $(\mathbf{x}_t, \hat{y}_t)$

$\alpha \geq 0$ is the hyperparameter of IABN that determines the confidence level of the BN statistics. A high value of α relies more on the learned statistics (BN), while a low value of α is in favor of the current statistics measured from the instance. Finally, the output of IABN can be described as:

$$\text{IABN}(\mathbf{f}_{b,c,:}; \bar{\boldsymbol{\mu}}_c, \bar{\boldsymbol{\sigma}}_c^2; \tilde{\boldsymbol{\mu}}_{b,c}, \tilde{\boldsymbol{\sigma}}_{b,c}^2) := \gamma \cdot \frac{\mathbf{f}_{b,c,:} - \boldsymbol{\mu}_{b,c}^{\texttt{IABN}}}{\sqrt{(\boldsymbol{\sigma}_{b,c}^{\texttt{IABN}})^2 + \epsilon}} + \beta. \tag{6}$$

Observe that IABN becomes IN and BN when $\alpha = 0$ and $\alpha = \infty$, respectively. If one chooses too small $\alpha \geq 0$, IABN may remove useful features, e.g., styles, of input (as with IN), which can degrade the overall classification (or regression) performance [29]. Hence, it is important to choose an appropriate α. Nevertheless, we found that a good choice of α is not too sensitive across tested scenarios, where we chose $\alpha = 4$ for all experiments. This way, IABN can be robust to distributional shifts without the risk of eliminating crucial information to predict y.

3.2 Adaptation via Prediction-Balanced Reservoir Sampling

Temporally correlated distributions lead to an undesirable bias in $p(y)$, and thus adaptation with a batch of consecutive test samples negatively impacts the adaptation objective, such as entropy minimization [40]. To combat this imbalance, we propose Prediction-Balanced Reservoir Sampling (PBRS) that mimics i.i.d. samples from temporally correlated streams with the assistance of a small (e.g., a mini-batch size) memory. PBRS combines *time-uniform* sampling and *prediction-uniform* sampling to simulate i.i.d. samples from the non-i.i.d. streams. For time-uniform sampling, we adopt reservoir sampling (RS) [39], a proven random sampling algorithm to collect time-uniform data in a single pass on a stream without the prior knowledge of the total length of data. For prediction-uniform sampling, we first use the predicted labels to compute the majority class(es) in the memory. We then replace a random instance of the majority class(es) with a new sample. We detail the algorithm of PBRS as a pseudo-code in Algorithm 1. We found that these two heuristics can effectively balance samples among both time and class axes, which mitigates the bias in temporally correlated data.

With the stored samples in the memory, we update the normalization statistics and affine parameters in the IABN layers. Note that IABN assumes $\tilde{\boldsymbol{\mu}}_{b,c}$ and $\tilde{\boldsymbol{\sigma}}_{b,c}^2$ follow the sampling distribution of $\mathcal{N}(\bar{\boldsymbol{\mu}}, \bar{\boldsymbol{\sigma}}^2)$ and corrects the normalization if $\tilde{\boldsymbol{\mu}}_{b,c}$ and $\tilde{\boldsymbol{\sigma}}_{b,c}^2$ are out of distribution. While IABN is resilient to distributional shifts to a certain extent, the assumption might not hold under severe distributional shifts. Therefore, we aim to find better estimates of $\bar{\boldsymbol{\mu}}, \bar{\boldsymbol{\sigma}}^2$ in IABN under distributional shifts via PBRS. Specifically, we update the normalization statistics, namely the means $\boldsymbol{\mu}$ and variances $\boldsymbol{\sigma}^2$, via exponential moving average: (a) $\boldsymbol{\mu}_t = (1-m)\boldsymbol{\mu}_{t-1} + m\frac{N}{N-1}\hat{\boldsymbol{\mu}}_t$ and (b) $\boldsymbol{\sigma}_t^2 = (1-m)\boldsymbol{\sigma}_{t-1}^2 + m\frac{N}{N-1}\hat{\boldsymbol{\sigma}}_t^2$ where m is a momentum and N is the size of the memory. We further

optimize the affine parameters, scaling factor γ and bias term β via a single backward pass with entropy minimization, similar to a previous study [40]. These parameters account for only around 0.02% of the total trainable parameters in ResNet18 [11]. The IABN layers are adapted with the N samples in the memory every N test samples. We set the memory size N as 64 following the common batch size of existing TTA methods [32, 4, 40] to ensure a fair memory constraint.

3.3 Inference

NOTE infers each sample via a single forward pass with IABN layers. Note that NOTE requires only a single instance for inference, different from the state-of-the-art TTA methods [28, 32, 40, 4, 43] that require batches for every inference. Moreover, NOTE requires only one forwarding pass for inference, while multiple forward passes are required in other TTA methods that utilize augmentations [34, 43]. The batch-free single-forward inference of NOTE is beneficial in latency-sensitive tasks such as autonomous driving and human health monitoring. After inference, NOTE determines whether to store the sample and predicted label to the memory via PBRS.

4 Experiments

We implemented NOTE and the baselines via the PyTorch framework [30]. We ran all experiments with three random seeds and specified the standard deviations. Additional experimental details are in the supplementary material.

Baselines. We consider the following baselines including the state-of-the-art test-time adaptation algorithms: **Source** evaluates the model trained from the source data directly on the target data without adaptation. Test-time normalization (**BN stats**) [28, 32] updates the BN statistics from a batch of test data. Online Domain Adaptation (**ONDA**) [26] adapts batch normalization statistics to target domains via a batch of target data with exponential moving average. Pseudo-Label (**PL**) [21] optimizes the trainable parameters in BN layers via hard pseudo labels. We update the BN layers only in PL, as done in previous studies [40, 43]. Test entropy minimization (**TENT**) [40] updates the BN parameters via entropy minimization. Laplacian Adjusted Maximum-likelihood Estimation (**LAME**) [4] takes a more conservative approach; it modifies the classifier's output probability and not the internal parameters of the model itself. By doing so, it prevents the model parameters from over-adapting to the test batch. Continual test-time adaptation (**CoTTA**) [43] reduces the error accumulation by using weight-averaged and augmentation-averaged predictions. It avoids catastrophic forgetting by stochastically restoring a part of the neurons to the source pre-trained weights.

Adaptation and hyperparameters. We assume the model pre-trained with source data is available for TTA. In NOTE, we replaced BN with IABN during training. We set the test batch size as 64 and the adaptation epoch as one for adaptation, which is the most common setting among the baselines [32, 4, 40]. Similarly, we set the memory size N as 64 and adapt the model every 64 samples in NOTE to ensure a fair memory constraint. We conduct online adaptation and evaluation, where the model is continually updated. For the baselines, we adopt the best values for the hyperparameters reported in their papers or the official codes. We followed the guideline to tune the hyperparameters when such a guideline was available [43]. We use fixed values for the hyperparameters of NOTE, soft-shrinkage width $\alpha = 4$ and exponential moving average momentum $m = 0.01$, and update the affine parameters via the Adam optimizer [17] with a learning rate $l = 0.0001$ for all experiments. We detailed hyperparameter information of the baselines in the supplementary material.

4.1 Robustness under corruptions

Datasets. CIFAR10-C/CIFAR100-C [12] are common TTA benchmarks for evaluating the robustness to corruptions [28, 32, 40, 43]. Both CIFAR10/CIFAR100 [18] have 50,000/10,000 training/test data. CIFAR10/CIFAR100 have 10/100 classes respectively. CIFAR10-C/CIFAR100-C apply 15 types of corruption to CIFAR10/CIFAR100 test data. Similar to previous studies [28, 32, 40, 43], we use the most severe corruption level of 5. We use ResNet18 [11] as the backbone network and pre-trained it on the clean training data. Following prior studies [22, 14, 42, 41], we adopt Dirichlet distribution to generate synthetic non-i.i.d. test streams from the originally i.i.d. CIFAR10/100 data. The details are provided in the supplementary material. We vary the Dirichlet concentration parameter

Table 1: Average classification error (%) and their corresponding standard deviations on CIFAR10-C/100-C under temporally correlated (non-i.i.d.) and uniformly distributed (i.i.d.) test data stream. **Bold** fonts indicate the lowest classification errors, while Red fonts show performance degradation after adaptation. Values encompassed by parentheses refers to NOTE used directly with test batches (without using PBRS). Averaged over 3 runs.

	Temporally correlated test stream			**Uniformly distributed test stream**		
Method	CIFAR10-C	CIFAR100-C	Avg	CIFAR10-C	CIFAR100-C	Avg
Source	42.3 ± 1.1	66.6 ± 0.1	54.4	42.3 ± 1.1	66.6 ± 0.1	54.4
BN Stats [28]	73.4 ± 1.3	65.0 ± 0.3	69.2	21.6 ± 0.4	46.6 ± 0.2	34.1
ONDA [26]	63.6 ± 1.0	49.6 ± 0.3	56.6	21.7 ± 0.4	46.5 ± 0.1	34.1
PL [21]	75.4 ± 1.8	66.4 ± 0.4	70.9	21.6 ± 0.2	43.1 ± 0.3	32.3
TENT [40]	76.4 ± 2.7	66.9 ± 0.6	71.7	18.8 ± 0.2	**40.3 ± 0.2**	29.6
LAME [4]	36.2 ± 1.3	63.3 ± 0.3	49.7	44.1 ± 0.5	68.8 ± 0.1	56.4
CoTTA [43]	75.5 ± 0.7	64.2 ± 0.2	69.8	17.8 ± 0.3	44.3 ± 0.2	31.1
NOTE	**21.1 ± 0.6**	**47.0 ± 0.1**	**34.0**	20.1 ± 0.5 (**17.6 ± 0.3**)	46.4 ± 0.0 (41.0 ± 0.2)	33.2 (**29.3**)

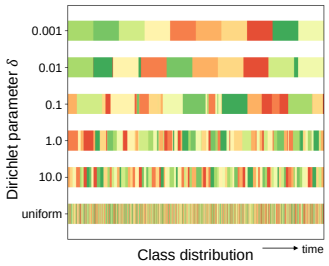

Figure 4: Illustration of synthetic non-i.i.d. streams sampled from a Dirichlet distribution varying δ on CIFAR10-C. *uniform* denotes an i.i.d. condition. The lower the δ, the more temporally correlated the distribution.

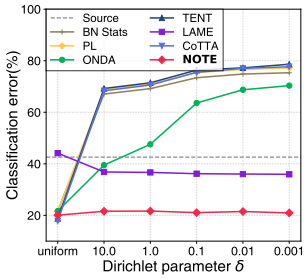

(a) Effect of Dirichlet concentration parameter δ.

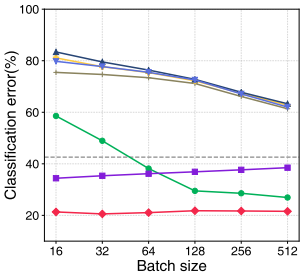

(b) Effect of batch size.

Figure 5: Average classification error (%) under the non-i.i.d. setting with CIFAR10-C dataset. We vary (a) the Dirichlet concentration parameter δ to investigate the impact of the degree of temporal correlation and (b) batch size to understand the behaviors of the TTA methods. Averaged over 3 runs. Lower is better.

δ to simulate diverse streams and visualize the resulting data in Figure 4. We use $\delta = 0.1$ as the default value unless specified.

Overall result. Tables 1 shows the result under the temporally correlated (non-i.i.d.) data and the uniform (i.i.d.) data, respectively. We observe significant performance degradation in the baselines under the temporally correlated setting. For BN Stats, PL, TENT, and CoTTA, this degradation is particularly due to the dependence on the test batch for the re-calculation of the BN statistics. Updating the batch statistics from test data via exponential moving average (ONDA) also suffers from the temporally correlated data. This indicates relying on the test batch for re-calculating the BN statistics indeed cancels out meaningful instance-wise variations under temporal correlation. Interestingly, LAME works better in the non-i.i.d. setting than in the i.i.d. setting, which is consistent with previous reports [4]. The primary reason is, as stated by the authors, it "discourages deviations from the predictions of the pre-trained model," and thus it "does not noticeably help in i.i.d and class-balanced scenarios." NOTE achieves on average 15.7% improvement over the best baseline (LAME) under the non-i.i.d. setting. With the i.i.d. assumption, NOTE still achieves comparable performance to the baselines. When we know the target samples are i.i.d., we can simply use the test batch without using PBRS. We report the result of the variant version of NOTE in the parentheses, which achieves on average 3.9% improvement further when the i.i.d. assumption is known.

Effect of the degree of temporal correlation. We also investigate the effect of the degree of temporal correlation for TTA algorithms. Figure 5a shows the result. The lower δ is, the severer

Table 2: Average classification error (%) and their corresponding standard deviations on three real test data streams: KITTI, HARTH, and ExtraSensory. **Bold** fonts indicate the lowest classification errors, while Red fonts show performance degradation after adaptation. Averaged over 3 runs.

	Real test stream			
Method	KITTI	HARTH	ExtraSensory	Avg
Source	12.3 ± 2.3	62.6 ± 8.5	50.2 ± 2.2	41.7
BN Stats [28]	35.4 ± 0.5	68.6 ± 1.1	56.0 ± 0.9	53.4
ONDA [26]	26.3 ± 0.5	69.3 ± 1.1	48.2 ± 1.5	47.9
PL [21]	39.0 ± 0.3	64.8 ± 0.6	56.0 ± 0.9	53.3
TENT [40]	39.6 ± 0.2	64.1 ± 0.7	56.0 ± 0.8	53.2
LAME [4]	11.3 ± 2.9	61.0 ± 10.0	50.7 ± 2.7	41.0
CoTTA [43]	35.4 ± 0.6	68.7 ± 1.1	56.0 ± 0.9	53.4
NOTE	**10.9 ± 3.6**	**51.0 ± 5.6**	**45.4 ± 2.6**	**35.8**

temporal correlation becomes. The error rates of most of the baselines deteriorate as δ decreases, which shows that the existing TTA baselines are susceptible to temporally correlated data. NOTE shows consistent performance among all δ values, indicating its robustness under temporal correlation.

Effect of batch size. While we experiment with a widely-used value of 64 as the batch size (or memory size in NOTE), one might be curious about the impact of batch size under temporally correlated streams. Figure 5b shows the result with six different batch sizes. As shown, NOTE is not much affected by the batch size, while most of the baselines recovers performance degradation as the batch size increases. This is because a higher batch size has a better chance of adaptation with balanced samples under temporally correlated streams. Increasing the batch size, however, mitigates temporal correlation at the expense of inference latency and adaptation speed.

4.2 Real-distributions with domain shift

Datasets. We evaluate NOTE under three real-world distribution datasets: object detection in autonomous driving (KITTI [8]), human activity recognition (HARTH [24]), and user behavioral context recognition (ExtraSensory [37]). KITTI is a well-known autonomous driving dataset that provides consecutive frames that contains natural temporal correlation in driving contexts. We adapt from KITTI to KITTI-Rain [10] - a dataset that converted KITTI images to rainy images. This contains 7,481/7,800 train/test samples with nine classes. We use ResNet50 [11] pre-trained on ImageNet [7] as the backbone network. HARTH was collected from 22 users in free-living environments for seven days. Each user was equipped with two three-axial Axivity AX3 accelerometers for recording human activities. We use 15 users collectively as the source domain and the remaining seven users as each target domain, which entails natural domain shifts from source users to target users as different physical conditions make domain shifts across users. HARTH contains 82,544/39,377 train/test samples with 12 classes. We report the averaged error over all target domains. We use four one-dimensional convolutional layers followed by one fully-connected layer as the backbone network for HARTH. The Extrasensory dataset collected users' own smartphone sensory data (motion sensors, audio, etc.) in the wild for seven days, aiming to capture people's authentic behaviors in their regular activities. We use 16 users as the source domain and seven users as target domains. ExtraSensory includes 17,777/4,862 train/test data with five classes. For ExtraSensory, we use two one-dimensional convolutional layers followed by one fully-connected layer as the backbone network. For both HARTH and ExtraSensory models, a single BN layer follows each convolutional layer. Additional details are in the supplementary material.

Result. Table 2 shows the result for the real-world datasets. The overall trend is similar to the temporal correlation experiments with CIFAR10-C/CIFAR100-C datasets, which indicates that the real-world datasets are indeed temporally correlated. NOTE consistently reduces errors after adaptation under real-world distributions. We believe this demonstrates NOTE is a promising method to be utilized in various real-world ML applications with distributional shifts.

Table 3: Average classification error (%) and corresponding standard deviations of varying ablation settings on CIFAR10-C/100-C under temporally correlated (non-i.i.d.) and uniformly distributed (i.i.d.) test data stream. **Bold** type indicates those of lowest classification error. Averaged over 3 runs.

	Temporally correlated test stream			**Uniformly distributed test stream**		
Method	CIFAR10-C	CIFAR100-C	Avg	CIFAR10-C	CIFAR100-C	Avg
Source	42.3 ± 1.1	66.6 ± 0.1	54.4	42.3 ± 1.1	66.6 ± 0.1	54.4
IABN	24.6 ± 0.6	54.5 ± 0.1	39.5	24.6 ± 0.6	54.5 ± 0.1	39.5
PBRS	27.5 ± 1.0	51.7 ± 0.2	39.6	25.8 ± 0.2	51.3 ± 0.1	38.5
IABN+RS	**20.5 ± 1.5**	48.2 ± 0.2	34.3	20.7 ± 0.6	48.3 ± 0.3	34.5
IABN+PBRS	21.1 ± 0.6	**47.0 ± 0.1**	**34.0**	**20.1 ± 0.5**	**46.4 ± 0.0**	**33.2**

4.3 Ablation study

We conduct an ablative study to further investigate the effectiveness of the individual components. Table 3 shows the result under both i.i.d. and non-i.i.d. settings. Using IABN alone significantly reduces error rates over Source, demonstrating the effectiveness of correcting normalization for out-of-distribution samples. Using PBRS with BN shows comparable improvement with the IABN-only result. Note that there is only a marginal gap (around 1%) between the non-i.i.d. and i.i.d. results in PBRS. This indicates that PBRS could effectively simulate i.i.d. samples from non-i.i.d. streams. The joint use of IABN and PBRS outperforms using either of them, meaning that PBRS provides IABN with better estimates for the normalizing operation. In addition, PBRS is better than Reservoir Sampling (RS) that has been a strong baseline in continual learning [16, 5]. This shows storing prediction-balanced sampling in addition to time-uniform sampling leads to better adaptation in TTA. We also investigated the joint use of IN and PBRS with the combination of IABN and PBRS on CIFAR100-C, and the result shows that IABN+PBRS (47.0%) achieves lower error rate than IN+PBRS (52.5%) on CIFAR100-C under temporal correlation.

5 Discussion and conclusion

This paper highlights that real-world distributions often change across the time axis, and existing test-time adaptation algorithms mostly suffer from the non-i.i.d. test data streams. To address this problem, we present a NOn-i.i.d. TEst-time adaptation algorithm, NOTE. Our experiments evaluated robustness under corruptions and domain adaptation on real-world distributions. The results demonstrate that NOTE not only outperforms the baselines under the non-i.i.d./real distribution settings, but it also shows comparable performance under the i.i.d. assumption. We believe that the insights and findings from this study are a meaningful step toward the practical impact of the test-time adaptation paradigm.

Limitations. NOTE and most state-of-the-art TTA algorithms [28, 21, 26, 32, 40, 43] assume that the backbone networks are equipped with BN (or IABN) layers. While BN is a widely-used component in deep learning, several architectures, such as LSTMs [13] and Transformers [38], do not embed BN layers. A recent study uncovered that BN is advantageous in Vision Transformers [44], showing potential room to apply our idea to architectures without BN layers. However, more in-depth studies are necessary to identify the actual applicability of BN (or IABN) to those architectures. While LAME [4] is applicable to models without BN, its limitation is the performance drop in i.i.d. scenarios, as shown in both its paper and our evaluation. While NOTE shows its effectiveness in both non-i.i.d and i.i.d. scenarios, a remaining challenge is to design an algorithm that generalizes to any architecture. We believe the findings and contributions of our work could give valuable insights to future endeavors on this end.

Potential negative societal impacts. As TTA relies on unlabeled test samples and changes the model accordingly, the model is exposed to potential data-driven biases after adaptation, such as fairness issues [3] and adversarial attacks [35]. In some sense, the utility of TTA comes at the expense of exposure to threats. This vulnerability is another crucial problem that both ML researchers and practitioners need to take into consideration. In addition, TTA entails additional computations for adaptation with test data, which may have negative impacts on environments, e.g., increasing electricity consumption and carbon emissions [33]. Nevertheless, we believe NOTE would not exacerbate this issue as it is computationally efficient as mentioned in Section §1.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section §5.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section §5.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See the supplementary material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section §4 and the supplementary material.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] We ran the entire experiments with three different random seems (0,1,2) and reported the standard deviations.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See the supplementary material.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

References

[1] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. *Advances in neural information processing systems*, 32, 2019.

[2] Elon Bachman and I Capulet. Digital record of tesla crashes resulting in death, May 2022.

[3] Solon Barocas, Moritz Hardt, and Arvind Narayanan. Fairness in machine learning. *Nips tutorial*, 1:2, 2017.

[4] Malik Boudiaf, Romain Mueller, Ismail Ben Ayed, and Luca Bertinetto. Parameter-free online test-time adaptation. In *Proceedings of Conference on Computer Vision and Pattern Recognition*, 2022.

[5] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc'Aurelio Ranzato. Continual learning with tiny episodic memories. *ICML Workshop: Multi-Task and Lifelong Reinforcement Learning*, 2019.

[6] Aristotelis Chrysakis and Marie-Francine Moens. Online continual learning from imbalanced data. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1952–1961. PMLR, 13–18 Jul 2020.

[7] 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.

[8] A Geiger, P Lenz, C Stiller, and R Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.

[9] Clément Godard, Oisin Mac Aodha, and Gabriel J. Brostow. Unsupervised monocular depth estimation with left-right consistency, 2016.

[10] Shirsendu Sukanta Halder, Jean-François Lalonde, and Raoul de Charette. Physics-based rendering for improving robustness to rain, 2019.

[11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, Los Alamitos, CA, USA, jun 2016. IEEE Computer Society.

[12] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*, 2019.

[13] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[14] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*, 2019.

[15] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 448–456, Lille, France, 07–09 Jul 2015. PMLR.

[16] David Isele and Akansel Cosgun. Selective experience replay for lifelong learning. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence*, AAAI'18/IAAI'18/EAAI'18. AAAI Press, 2018.

[17] Diederick P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.

[18] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. *Master's thesis, Department of Computer Science, University of Toronto*, 2009.

[19] Nicholas D. Lane, Emiliano Miluzzo, Hong Lu, Daniel Peebles, Tanzeem Choudhury, and Andrew T. Campbell. A survey of mobile phone sensing. *IEEE Communications Magazine*, 48(9):140–150, 2010.

[20] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[21] 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*, page 896, 2013.

[22] Qinbin Li, Yiqun Diao, Quan Chen, and Bingsheng He. Federated learning on non-iid data silos: An experimental study. *arXiv preprint arXiv:2102.02079*, 2021.

[23] Yuejiang Liu, Parth Kothari, Bastien van Delft, Baptiste Bellot-Gurlet, Taylor Mordan, and Alexandre Alahi. Ttt++: When does self-supervised test-time training fail or thrive? *Advances in Neural Information Processing Systems*, 34, 2021.

[24] Aleksej Logacjov, Kerstin Bach, Atle Kongsvold, Hilde Bremseth Bårdstu, and Paul Jarle Mork. Harth: A human activity recognition dataset for machine learning. *Sensors*, 21(23), 2021.

[25] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. 2017.

[26] Massimiliano Mancini, Hakan Karaoguz, Elisa Ricci, Patric Jensfelt, and Barbara Caputo. Kitting in the wild through online domain adaptation. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1103–1109. IEEE, 2018.

[27] Norman Mu and Justin Gilmer. Mnist-c: A robustness benchmark for computer vision. *arXiv preprint arXiv:1906.02337*, 2019.

[28] Zachary Nado, Shreyas Padhy, D Sculley, Alexander D'Amour, Balaji Lakshminarayanan, and Jasper Snoek. Evaluating prediction-time batch normalization for robustness under covariate shift. *arXiv preprint arXiv:2006.10963*, 2020.

[29] Hyeonseob Nam and Hyo-Eun Kim. Batch-instance normalization for adaptively style-invariant neural networks. *Advances in Neural Information Processing Systems*, 31, 2018.

[30] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.

[31] Joaquin Quiñonero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D Lawrence. *Dataset shift in machine learning*. Mit Press, 2008.

[32] Steffen Schneider, Evgenia Rusak, Luisa Eck, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. Improving robustness against common corruptions by covariate shift adaptation. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 11539–11551. Curran Associates, Inc., 2020.

[33] Roy Schwartz, Jesse Dodge, Noah A. Smith, and Oren Etzioni. Green ai. *Commun. ACM*, 63(12):54–63, nov 2020.

[34] 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.

[35] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.

[36] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.

[37] Yonatan Vaizman, Katherine Ellis, and Gert Lanckriet. Recognizing detailed human context in the wild from smartphones and smartwatches. *IEEE Pervasive Computing*, 16(4):62–74, 2017.

[38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[39] Jeffrey S. Vitter. Random sampling with a reservoir. *ACM Trans. Math. Softw.*, 11(1):37–57, mar 1985.

[40] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *International Conference on Learning Representations*, 2021.

[41] Hongyi Wang, Mikhail Yurochkin, Yuekai Sun, Dimitris Papailiopoulos, and Yasaman Khazaeni. Federated learning with matched averaging. *arXiv preprint arXiv:2002.06440*, 2020.

[42] Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H Vincent Poor. Tackling the objective inconsistency problem in heterogeneous federated optimization. *Advances in neural information processing systems*, 33:7611–7623, 2020.

[43] Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual test-time domain adaptation. In *Proceedings of Conference on Computer Vision and Pattern Recognition*, 2022.

[44] Zhuliang Yao, Yue Cao, Yutong Lin, Ze Liu, Zheng Zhang, and Han Hu. Leveraging batch normalization for vision transformers. In *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pages 413–422, 2021.

[45] Hangyu Zhu, Jinjin Xu, Shiqing Liu, and Yaochu Jin. Federated learning on non-iid data: A survey. *Neurocomputing*, 465:371–390, 2021.