

Neural Collapse in Test-Time Adaptation

Xiao Chen¹, Zhongjing Du², Jiazhen Huang¹, Xu Jiang²,
Li Lu³, Jingyan Jiang^{4*}, Zhi Wang^{1*}

¹ Shenzhen International Graduate School, Tsinghua University

² Peking University ³ Sichuan University ⁴ Shenzhen Technology University

chen-x25@mails.tsinghua.edu.cn, jiangjingyan@sztu.edu.cn, wangzhi@sz.tsinghua.edu.cn


Abstract

Test-Time Adaptation (TTA) enhances model robustness to out-of-distribution (OOD) data by updating the model online during inference, yet existing methods lack theoretical insights into the fundamental causes of performance degradation under domain shifts. Recently, Neural Collapse (NC) has been proposed as an emergent geometric property of deep neural networks (DNNs), providing valuable insights for TTA. In this work, we extend NC to the sample-wise level and discover a novel phenomenon termed Sample-wise Alignment Collapse (NC3+), demonstrating that a sample's feature embedding, obtained by a trained model, aligns closely with the corresponding classifier weight. Building on NC3+, we identify that the performance degradation stems from sample-wise misalignment in adaptation which exacerbates under larger distribution shifts. This indicates the necessity of realigning the feature embeddings with their corresponding classifier weights. However, the misalignment makes pseudo-labels unreliable under domain shifts. To address this challenge, we propose NCTTA, a novel feature-classifier alignment method with hybrid targets to mitigate the impact of unreliable pseudo-labels, which blends geometric proximity with predictive confidence. Extensive experiments demonstrate the effectiveness of NCTTA in enhancing robustness to domain shifts. For example, NCTTA outperforms Tent by 14.52% on ImageNet-C. Project page is publicly available at <https://github.com/Cevaaa/NCTTA>.

1. Introduction

Deep neural networks (DNNs) [33] have achieved remarkable success across various tasks but suffer significant performance degradation under out-of-distribution (OOD) data. Traditional approaches like Domain Adaptation (DA) [19, 46] aim to address this issue but often involve complex training procedures and assume access to test data,

*Co-corresponding authors



(a) Proposition: Sample-wise Alignment Collapse (NC3+). (b) Observation: Sample-wise Misalignment in Adaptation.

Figure 1. **Overview of Main Contributions.** (a) NC3+ highlights the convergence of sample feature embeddings with their corresponding classifier weights. (b) Sample feature embeddings deviate from the ground-truth classifier weights, leading to performance degradation.

which is rarely feasible in real-world. Test-Time Adaptation (TTA) offers a more practical solution by enabling models to adapt online during inference using only a mini-batch of test data, without altering the training process [5, 18]. This lightweight and efficient paradigm improves robustness to domain shifts, facilitating reliable deployment in dynamic and unpredictable environments.

Existing TTA methods can be broadly categorized into four main strategies: prototype-based methods [12, 17], consistency regularization methods [37, 40], normalization layer-based methods [8, 13, 24, 31], and entropy minimization methods [2, 15, 23, 36]. Although these approaches have shown promising results through algorithmic optimization during inference, they generally lack a theoretical exploration of the fundamental causes of model degradation when facing domain shifts.

Recently, the phenomenon known as *Neural Collapse* (NC) [1, 26, 44, 47] has emerged as a significant characterization of the geometric properties of DNNs during the Terminal Phase of Training (TPT). During the TPT, after achieving zero classification error, DNNs asymptotically drive the training loss toward zero. Unlike previous TTA methods, this work bridges the theory of NC with TTA by leveraging the well-defined geometric structures of a trained DNN formalized by NC properties. This connection offers novel theoretical insights into TTA, grounding it in the

structured behavior of DNNs. Specifically, during the TPT, NC consistently appears in the penultimate layer and its linear classifier [7], exhibiting a highly structured and intriguing behavior defined by four properties:

- **Variability Collapse (NC1):** Sample feature embeddings within the same class collapse to their corresponding class means, reducing intra-class variance.
- **Convergence to a Simplex ETF (NC2):** Class means converge to the vertices of a simplex equiangular tight frame (ETF), forming a maximally separated structure.
- **Convergence to Self-Duality (NC3):** Class means and corresponding classifier weights will converge to each other during the TPT.
- **Simplification to Nearest Class-Center (NC4):** The classification function effectively reduces to choosing the nearest class mean in the feature space.

While NC unveils elegant geometric patterns in DNNs during the TPT, its analysis predominantly focuses on a *class-wise* perspective. However, this assumption breaks down in TTA scenarios, where models are adapted using only mini-batches of unlabeled test samples. The absence of labeled data precludes the computation of *class means*, which are indispensable in NC properties, thereby rendering class-wise analyses inapplicable in TTA. This challenge motivates us to explore whether similar collapse phenomena emerge at the sample-wise level. As illustrated in Figure 1(a), our analysis extends the NC framework by identifying an additional property, NC3+, that holds at the sample-wise granularity:

- **Sample-wise Alignment Collapse (NC3+):** Each sample feature embedding and the corresponding classifier weight will converge to each other during the TPT.

Building on this observation, we further analyze the performance degradation under domain shifts during test time. As illustrated in Figure 1(b), for OOD data, the feature embedding of a test sample tends to deviate from its corresponding classifier weight and drift toward that of an incorrect one, leading to misclassification. This misalignment intensifies as the domain shift becomes more severe, highlighting a key challenge in TTA:

Misalignment between feature embeddings and classifier weights is the key challenge under domain shifts, which underscores the necessity of feature-classifier alignment.

To mitigate this effect, we propose a novel **NC-guided TTA (NCTTA)** method to promote feature-classifier alignment. The sample-wise misalignment implies the unreliability of pseudo-labels. Therefore, instead of simply aligning the feature embedding with the classifier weight of its pseudo-label, NCTTA blends geometric proximity with predictive confidence to obtain hybrid targets for realignment. This strategy substantially improves robustness under domain shifts by alleviating the pseudo-label unreliability

caused by Sample-wise Misalignment in Adaptation. Our main contributions are as follows:

- NC theory is, for the first time, extended to the sample-wise level through Sample-wise Alignment Collapse (NC3+), validated theoretically and empirically.
- Comprehensive empirical evidence is provided to show that the performance degradation of pre-trained models on OOD data originates from sample-wise misalignment.
- A novel NC-guided TTA method is proposed to effectively promote alignment between sample feature embeddings and their corresponding classifier weights.
- Extensive experiments demonstrate the effectiveness of NCTTA. For instance, it achieves 78.30% average accuracy on CIFAR-10-C under severe corruption and 66.61% on ImageNet-C, outperforming prior methods.

2. Related Work

2.1. Test-Time Adaptation

TTA addresses performance degradation on OOD data by fine-tuning a pretrained model during inference using a mini-batch in an online manner [22, 39]. Existing TTA methods can be categorized into four types: prototype-based methods (e.g., SHOT [17], T3A [12]) refine decision boundaries using class prototypes; consistency regularization methods (e.g., MEMO [40], CoTTA [37]) stabilize predictions through constraints across augmented views or over time; normalization-based methods (e.g., NOTE [8], SAR [24]) explore the role of normalization layers in TTA and carefully design their adaptation; and entropy minimization methods (e.g., Tent [36], EATA [23], DeYO [15]) enhance reliability by minimizing prediction entropy with confident samples. Despite their empirical effectiveness, these methods lack theoretical understanding of model behavior under domain shifts. To address this, we provide a principled analysis of performance degradation during TTA and propose a theory-driven method that aligns test-time feature embeddings with their corresponding classifier weights, improving robustness to OOD scenarios.

2.2. Neural Collapse

NC describes a set of geometric properties that emerge in DNNs during the TPT [26]. These properties include the collapse of within-class variability (NC1), the convergence of class means to a Simplex Equiangular Tight Frame (NC2), the alignment of classifier weights with class means (NC3), and classification via the nearest class mean (NC4). These phenomena have been consistently observed across diverse datasets and architectures under both cross-entropy and MSE losses [7], and are often interpreted as optimal structures arising from simplified learning objectives [1, 44, 47]. However, existing works on NC [6, 20, 38, 45, 48] have primarily focused on class-wise geometric structures

between class centers and classifier weights observed in the TPT, relying on class labels and full supervision to define its geometric properties, whereas TTA operates on unlabeled target data with incremental updates using mini-batches, making NC’s class-wise assumptions inapplicable. To overcome this limitation, NC is extended to a sample-wise perspective through the introduction of NC3+, as detailed in Section 3.2. Building on the insights provided by NC3+, the causes of misclassification on OOD data are analyzed in Section 3.3. To address these challenges, NCTTA is proposed, a feature-classifier alignment method, leveraging the principles of NC to enhance TTA performance under domain shifts.

3. Rethinking Neural Collapse

3.1. Preliminaries

In this section, we establish the notations used throughout the paper. Let the training and test datasets be denoted as $D_{\text{train}} = \{(x_i, y_i)\}_{i=1}^{N_{\text{train}}} \sim \mathcal{P}_{\text{train}}$ and $D_{\text{test}} = \{x_i\}_{i=1}^{N_{\text{test}}} \sim \mathcal{P}_{\text{test}}$, where N_{train} and N_{test} represent the number of samples in the training and test sets, respectively. Under the TTA setting, there exists a domain shifts such that $\mathcal{P}_{\text{train}} \neq \mathcal{P}_{\text{test}}$, and the D_{test} is unlabeled. Moreover, during adaptation, only a mini-batch of test samples is accessible at each iteration.

We denote a DNN by $f_{\theta}(\cdot)$, which consists of a feature extractor $h_{\theta-\omega}(\cdot)$ and a classifier head $g_{\omega}(\cdot)$, i.e., $f_{\theta}(\cdot) = (g_{\omega} \circ h_{\theta-\omega})(\cdot)$, where \circ denotes function composition. Here, θ represents the full set of model parameters, and $\omega = [w_1, w_2, \dots, w_K] \in \mathbb{R}^{K \times L}$ denotes the classifier weights, with K the number of classes and L the feature embedding dimension. For simplicity, we omit the bias term in the classifier.

Given an input sample x_i , its feature embedding is computed as $\mathbf{h}_i = h_{\theta-\omega}(x_i) \in \mathbb{R}^L$. The corresponding logits are given by $\mathbf{z}_i = g_{\omega}(\mathbf{h}_i) \in \mathbb{R}^K$. Applying the softmax function $\sigma(\cdot)$ to \mathbf{z}_i yields the predicted class probabilities $p_i = \sigma(\mathbf{z}_i) \in \mathbb{R}^K$, where $p_{ij} = \exp(\mathbf{z}_{ij}) / \sum_{k=1}^K \exp(\mathbf{z}_{ik})$.

3.2. Extension of Neural Collapse

This subsection analyzes the training phase of f_{θ} , focusing on the NC3. The remaining NC properties are detailed in Appendix A.1. NC3 describes the *self-duality* phenomenon, where class means from $h_{\theta-\omega}$ align with the classifier weights ω in g_{ω} , up to a rescaling factor. Formally, the *global mean* and *class mean* for each class c are defined as: $\mu_G = 1/|D_{\text{train}}| \sum_{x_i \in D_{\text{train}}} h_{\theta-\omega}(x_i)$, $\mu_c = 1/|D_{\text{train}}^c| \sum_{x_i \in D_{\text{train}}^c} h_{\theta-\omega}(x_i)$, where $D_{\text{train}}^c = \{(x_i, y_i) \in D_{\text{train}} \mid y_i = c\}$ and $|\cdot|$ denotes the cardinality of a set. Prior work [26] quantify NC3 with the metric:

$$\left\| \frac{\mathbf{M}}{\|\mathbf{M}\|_2} - \frac{\omega}{\|\omega\|_2} \right\|_2 \rightarrow 0, \quad (1)$$

where $\mathbf{M} = [\mu_0 - \mu_G, \mu_1 - \mu_G, \dots, \mu_{K-1} - \mu_G]$.

NC3 captures the convergence of class means and classifier weights, revealing an elegant alignment between feature and parameter spaces. However, it is inherently class-centric, relying on class labels to compute μ_c and the full training set to estimate μ_G . These assumptions are incompatible with TTA, where only unlabeled mini-batches are accessible during inference. To address this, we propose NC3+, a sample-wise extension of NC3 that examines the alignment between individual sample feature embeddings and classifier weights. Specifically, we define a Feature-Classifier Alignment (FCA) distance d_{ij} of a test sample x_i as the normalized Euclidean distance between its feature embedding $\mathbf{h}_i = h_{\theta-\omega}(x_i)$ and classifier weight w_j of j -th class. Formally, the FCA distance d_{ij} is calculated as:

$$d_{ij} = \left\| \frac{\mathbf{h}_i}{\|\mathbf{h}_i\|_2} - \frac{w_j}{\|w_j\|_2} \right\|_2, \quad j = 0, 1, \dots, K-1. \quad (2)$$

To validate NC3+ during the TPT, we provide a theoretical proof (Appendix A.3) demonstrating that, for a labeled sample $(x_i, y_i) \in D_{\text{train}}$, with the cross-entropy loss $\mathcal{L}_{\text{CE}}(x_i, y_i) = -\log p_{y_i}$, the Ground-truth FCA (G-FCA) distance d_{iy_i} decreases monotonically and converges to zero. In contrast, for $j \neq y_i$, d_{ij} remains largely unaffected. We formally define NC3+ as:

Proposition: Sample-wise Alignment Collapse (NC3+). *During the TPT, the G-FCA distance d_{iy_i} between a sample’s feature embedding and its corresponding classifier weight is expected to converge to zero, formulated as:*

$$\text{G-FCA: } d_{iy_i} = \left\| \frac{\mathbf{h}_i}{\|\mathbf{h}_i\|_2} - \frac{w_{y_i}}{\|w_{y_i}\|_2} \right\|_2 \rightarrow 0. \quad (3)$$

NC3+ encapsulates the phenomenon of sample-wise alignment collapse during the TPT, extending NC theory to settings where class-wise and global means are not feasible. We empirically validate it through controlled experiments. As shown in Figure 2, G-FCA distance d_{iy_i} consistently decreases and asymptotically approaches zero during training, substantiating the presence of NC3+. *This implies that well-trained DNNs generally satisfy the NC3+ property, extending the Neural Collapse phenomenon to settings where class-wise and global means are not feasible.*

3.3. Explaining Performance Degradation in TTA

NC3+ inspires us to understand performance degradation under domain shifts from the perspective of FCA distance. Prior work [1] observed that feature embeddings of In-Distribution (ID) and OOD data tend to become orthogonal as training progresses. Building on this, we analyze the relationship between G-FCA distance and Pseudo-label FCA

(P-FCA) distance to better understand the misalignment responsible for performance degradation. Figure 3 illustrates the distributions of G-FCA d_{iy_i} and P-FCA $d_{i\hat{y}_i}$ distances of test samples $x_i \in D_{\text{test}}$ that are defined as:

$$\text{G-FCA: } d_{iy_i} = \left\| \frac{\mathbf{h}_i}{\|\mathbf{h}_i\|_2} - \frac{\mathbf{w}_{y_i}}{\|\mathbf{w}_{y_i}\|_2} \right\|_2, \quad (4)$$

$$\text{P-FCA: } d_{i\hat{y}_i} = \left\| \frac{\mathbf{h}_i}{\|\mathbf{h}_i\|_2} - \frac{\mathbf{w}_{\hat{y}_i}}{\|\mathbf{w}_{\hat{y}_i}\|_2} \right\|_2. \quad (5)$$

where y_i is the ground-truth label and \hat{y}_i is the pseudo-label. Assume that the FCA distance $d_{ij}^{(\cdot)}$ follows a normal distribution, $d_{ij}^{(\cdot)} \sim N(\mu_j^{(\cdot)}, \sigma_j^{(\cdot)})$, where $\mu_j^{(\cdot)}$ and $\sigma_j^{(\cdot)}$ represent the mean and variance. Here, (\cdot) indicates whether the test sample x_i is classified correctly or incorrectly. For misclassified samples, $\mu_{y_i}^{\text{wrong}}$ is significantly larger, reflecting severe misalignment with ground-truth class weights caused by domain shifts, while $\mu_{\hat{y}_i}^{\text{wrong}}$ is smaller, indicating drift toward incorrect class weights. Figure 4 further quantifies this trend, showing that as corruption severity increases, the gap between $\mu_{y_i}^{\text{wrong}}$ and $\mu_{\hat{y}_i}^{\text{wrong}}$ widens. Figure 3 and Figure 4 reveal that feature-classifier misalignment is the primary driver of performance degradation for OOD data:

Observation: Sample-wise Misalignment in Adaptation.

The feature embedding of an OOD sample tends to deviate from its corresponding classifier weight and drift toward that of an incorrect class, leading to misclassification. Formally, the means of G-FCA distance d_{iy_i} and P-FCA distance $d_{i\hat{y}_i}$ satisfy:

$$\mu_{\hat{y}_i}^{\text{correct}} = \mu_{y_i}^{\text{correct}} < \mu_{\hat{y}_i}^{\text{wrong}} \ll \mu_{y_i}^{\text{wrong}}. \quad (6)$$

NC3+ and the observation indicate that performance degradation under domain shifts stems from the displacement of feature embeddings away from corresponding classifier weights. This misalignment increases d_{iy_i} and reduces $d_{i\hat{y}_i}$, leading to erroneous predictions.




Figure 2. **Empirical validation of NC3+.** We evaluate NC3+ on ImageNet-100 [29] using various backbones. The G-FCA distance d_{iy_i} decreases throughout training, indicating sample-wise alignment collapse. More details are provided in Appendix A.2.




Figure 3. **Histograms of G-FCA and P-FCA distances.** We present the distributions of $d_{iy_i}^{\text{correct}}$, $d_{iy_i}^{\text{wrong}}$, and $d_{i\hat{y}_i}^{\text{wrong}}$ on correctly and incorrectly classified samples from ImageNet-C datasets under severity level 5 Gaussian noise and Snow corruption. The results reveal that NC3+ is violated on OOD data, leading to Sample-wise Misalignment in Adaptation.




Figure 4. **Violin plots of G-FCA and P-FCA distances for misclassified samples.** The plots show the distributions of distances from misclassified OOD samples to both G-FCA $d_{iy_i}^{\text{wrong}}$ and P-FCA $d_{i\hat{y}_i}^{\text{wrong}}$ under increasing Gaussian noise or Snow severity on ImageNet-C. The results reveal that misalignment becomes progressively more severe with higher corruption levels.

4. Method

Building on the NC3+ property from Section 3.2 and the observation from Section 3.3, we note that under domain shifts, OOD sample feature embeddings deviate from their corresponding classifier weights, thereby increasing G-FCA distance d_{iy_i} and inducing misclassification as feature embeddings align with incorrect classifier weights. To address this issue, we propose a new method, NCTTA, that explicitly constrains d_{iy_i} to realign feature embeddings with their correct classifier weights. Since TTA operates in an unlabeled regime, directly computing G-FCA distance d_{iy_i} is infeasible. A straightforward workaround is to use the pseudo-label \hat{y}_i as a proxy for y_i and constrain P-FCA distance $d_{i\hat{y}_i}$. However, as demonstrated in Section 3.3, pseudo-labels become unreliable under severe domain shifts. To overcome the limitation, NCTTA is designed as an NC-guided contrastive alignment mechanism with hybrid targets, blending geometric proximity and predictive confidence.

Neural Collapse in Test-Time Adaptation (NCTTA).

During TTA, pretrained models are updated online using unlabeled data, which prevents direct computation of G-FCA distance d_{iy_i} . Section 3.3 shows that the offset between \mathbf{h}_i and w_{y_i} makes P-FCA distance $d_{i\hat{y}_i}$ unreliable. Therefore, simply constraining P-FCA distance $d_{i\hat{y}_i}$ is im-

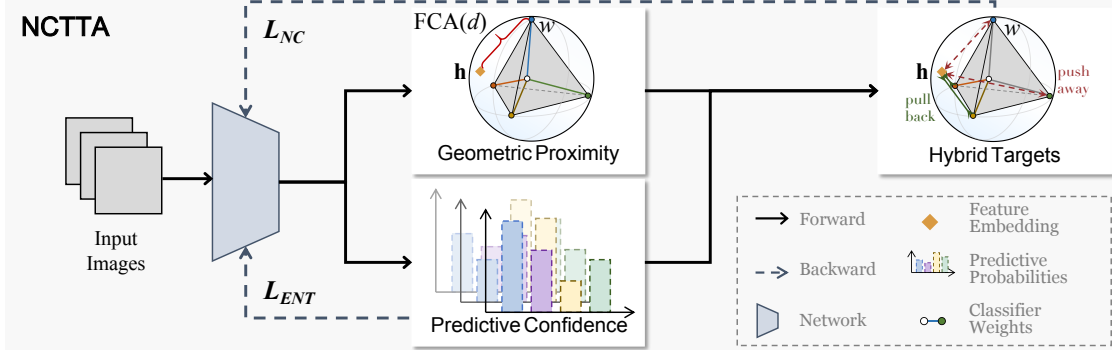


Figure 5. **Overview of our proposed NCTTA.** During test-time adaptation, NCTTA blends geometric proximity (FCA distance) and predictive confidence to form hybrid targets, pulling features toward plausible classifier weights while pushing away negatives via \mathcal{L}_{NC} .

practical. Our NCTTA replace \hat{y}_i with dual-guided hybrid targets determined by $\tilde{\mathbf{y}} \in \mathbb{R}^K$ via:

$$\tilde{\mathbf{y}}_i = (1 - \alpha) \hat{d}_i + \alpha p_i, \quad (7)$$

where \hat{d}_i is obtained by first applying Z-score normalization to d_i , then taking the negative normalized distances followed by a softmax operation. This ensures that smaller FCA distances correspond to larger probability values. The coefficient α balances geometric proximity and predictive confidence. A larger $\tilde{\mathbf{y}}_{ij}$ indicates that the feature embedding of x_i is geometrically closer to classifier weight w_j and the model is more confident in class j . Therefore, we introduce the dual-guided order $\mathcal{O}_i = \text{argsort}(\tilde{\mathbf{y}}_i)$, ranking classes in descending order of $\tilde{\mathbf{y}}_i$. Obviously, we have $\mathcal{O}_i(\text{argmax } \tilde{\mathbf{y}}_i) = 0$ and $\mathcal{O}_i(\text{argmin } \tilde{\mathbf{y}}_i) = K - 1$. Hybrid targets are obtained using:

$$\mathcal{T}_i = \{\mathcal{O}_i(0), \mathcal{O}_i(1), \dots, \mathcal{O}_i(k-1)\}, \quad (8)$$

where \mathcal{T}_i is the set of the k classes with the largest $\tilde{\mathbf{y}}_i$ values. When $\alpha = 1.0$ and $k = 1$, \mathcal{T}_i contains a single element, namely \hat{y}_i . Given \mathcal{T}_i , we then compute an NC-guided soft alignment loss:

$$\mathcal{L}_{NC}(x_i) = \ell(\{d_{ij}\}_{j \in \mathcal{T}_i}, \{d_{ij}\}_{j \notin \mathcal{T}_i}), \quad (9)$$

where ℓ enforces a monotonic preference for ‘‘closer to positives, farther from negatives.’’ Specifically, it is nonincreasing in d_{ij} for $j \in \mathcal{T}_i$ which means $\partial \ell / \partial d_{ij} \geq 0$ ($j \in \mathcal{T}_i$) according to gradient descent [28], thereby encouraging reductions in d_{ij} for positives; and it is nondecreasing in d_{ij} for $j \notin \mathcal{T}_i$ which means $\partial \ell / \partial d_{ij} \leq 0$ ($j \notin \mathcal{T}_i$). We instantiate ℓ using InfoNCE-style, L2-style, or Triplet-style objectives, as detailed in Section 5.3.

Final Objective. While \mathcal{L}_{NC} leverages hybrid targets to mitigate misalignment between \mathbf{h}_i and w_{y_i} , excessive noise in test samples can still destabilize adaptation. To enhance robustness, we integrate widely adopted techniques in the

TTA literature into the final objective. We adopt an entropy-based filtering strategy inspired by prior works [15, 23, 24]:

$$\mathcal{S}_{ENT} = \{x_i \in D_{\text{test}} \mid \mathcal{L}_{ENT}(x_i) < \gamma_{ENT}\}, \quad (10)$$

where γ_{ENT} is a threshold filtering out high-entropy predictions, and \mathcal{L}_{ENT} is the entropy minimization loss:

$$\mathcal{L}_{ENT}(x_i) = - \sum_{j=0}^{K-1} p_j \log p_j. \quad (11)$$

To jointly incorporate entropy and alignment confidence, we define a dynamic weight:

$$\lambda_i = \frac{1}{\exp(\mathcal{L}_{ENT}(x_i) - \tau_{ENT})} + \frac{\nu}{1 + \eta d_{i\hat{y}_i}}, \quad (12)$$

where $d_{i\hat{y}_i}$ denotes P-FCA distance, and τ_{ENT} , ν , η control sensitivity to entropy and P-FCA. Combining entropy filtering, NC-guided soft alignment loss, and adaptive weighting, the final loss for each test sample is:

$$\mathcal{L}_{\text{total}}(x_i) = \lambda_i \cdot \mathbb{I}_{x_i \in \mathcal{S}_{ENT}} \cdot (\mathcal{L}_{ENT}(x_i) + \mathcal{L}_{NC}(x_i)), \quad (13)$$

where $\mathbb{I}_{(\cdot)}$ is an indicator function that evaluates to 1 if (\cdot) is true and 0 otherwise. This formulation jointly promotes confident predictions and robust alignment between feature embeddings and classifier weights, thereby improving robustness to domain shifts at test time.

5. Experiments

5.1. Experimental Setup

The proposed method is evaluated on five widely used benchmarks: CIFAR-10/100-C [10], ImageNet-C [10], PACS [16] and Waterbirds [30]. Detailed descriptions of these datasets provided in Appendix B.1 and experiment results on CIFAR-100-C are shown in Appendix C. NCTTA is compared against a wide range of state-of-the-art TTA baselines, including no_adapt (no adaptation), Tent [36], BN_adapt [31], NOTE [8], MEMO [40], SAR [24],

Table 1. Classification accuracy (%) on CIFAR-10-C and ImageNet-C under corruption severity level 5. The **bold** entries indicate the highest performance for each corruption type, and the underlined entries indicate the next-highest performance.

CIFAR-10-C (ResNet50)																
Mild	Gauss.	Shot	Impul.	Defoc.	Glass	Motion	Zoom	Snow	Frost	Fog	Brit.	Contr.	Elastic	Pixel	JPEG	Avg.
no_adapt	33.69	40.07	35.74	53.11	49.86	67.04	56.85	73.89	62.59	64.66	<u>89.00</u>	43.57	74.37	41.96	74.43	57.39
BN_adapt	<u>62.43</u> ± 0.02	<u>64.33</u> ± 0.02	<u>57.30</u> ± 0.08	<u>84.48</u> ± 0.00	<u>59.46</u> ± 0.03	<u>82.02</u> ± 0.04	<u>83.25</u> ± 0.03	<u>76.87</u> ± 0.01	<u>74.30</u> ± 0.01	<u>79.29</u> ± 0.00	<u>87.24</u> ± 0.02	<u>83.33</u> ± 0.02	<u>71.72</u> ± 0.01	<u>73.41</u> ± 0.01	<u>71.10</u> ± 0.02	<u>74.03</u>
Tent	64.71 ± 0.03	66.97 ± 0.03	59.61 ± 0.01	84.97 ± 0.03	60.96 ± 0.04	82.48 ± 0.04	83.93 ± 0.01	77.78 ± 0.01	75.42 ± 0.00	80.00 ± 0.04	87.78 ± 0.01	83.54 ± 0.00	72.62 ± 0.05	74.73 ± 0.01	72.32 ± 0.03	75.19
EATA	62.43 ± 0.02	64.33 ± 0.02	57.31 ± 0.08	84.48 ± 0.00	59.46 ± 0.03	82.02 ± 0.04	83.25 ± 0.02	76.88 ± 0.01	74.31 ± 0.01	79.28 ± 0.00	87.24 ± 0.02	83.33 ± 0.02	71.72 ± 0.01	73.42 ± 0.01	71.13 ± 0.02	74.04
SAR	63.74 ± 0.07	65.87 ± 0.06	58.47 ± 0.02	84.70 ± 0.01	60.51 ± 0.02	82.22 ± 0.02	83.62 ± 0.02	77.26 ± 0.03	74.74 ± 0.02	79.74 ± 0.02	87.40 ± 0.01	83.29 ± 0.02	72.24 ± 0.03	74.15 ± 0.00	71.87 ± 0.02	74.67
NOTE	55.17 ± 0.03	60.13 ± 0.15	55.13 ± 0.06	74.53 ± 0.00	59.43 ± 0.08	78.61 ± 0.04	75.72 ± 0.08	79.90 ± 0.00	<u>76.46</u> ± 0.01	78.18 ± 0.05	<u>90.91</u> ± 0.01	70.43 ± 0.04	<u>76.18</u> ± 0.03	58.91 ± 0.12	75.83 ± 0.00	71.03
MEMO	52.22 ± 0.00	58.09 ± 0.00	51.40 ± 0.00	73.17 ± 0.00	58.01 ± 0.00	76.50 ± 0.01	74.61 ± 0.00	77.74 ± 0.00	73.01 ± 0.00	75.47 ± 0.00	90.27 ± 0.00	63.89 ± 0.00	<u>76.95</u> ± 0.00	55.86 ± 0.00	75.62 ± 0.00	68.85
DeYO	<u>67.61</u> ± 0.27	<u>70.15</u> ± 0.02	<u>63.15</u> ± 0.15	<u>85.59</u> ± 0.00	<u>62.98</u> ± 0.03	<u>83.09</u> ± 0.13	<u>84.57</u> ± 0.00	79.05 ± 0.00	76.35 ± 0.02	<u>80.68</u> ± 0.03	88.17 ± 0.00	<u>83.57</u> ± 0.00	74.16 ± 0.04	<u>76.45</u> ± 0.02	<u>74.13</u> ± 0.08	<u>76.65</u>
NCTTA (ours)	70.94 ± 0.14	72.88 ± 0.28	66.18 ± 0.60	86.25 ± 0.11	64.69 ± 0.30	83.88 ± 0.30	85.54 ± 0.43	80.16 ± 0.28	77.73 ± 0.22	81.55 ± 0.24	88.82 ± 0.10	83.81 ± 0.13	75.54 ± 0.09	78.57 ± 0.31	75.92 ± 0.44	78.16

ImageNet-C (ViT-B/16)																
Mild	Gauss.	Shot	Impul.	Defoc.	Glass	Motion	Zoom	Snow	Frost	Fog	Brit.	Contr.	Elastic	Pixel	JPEG	Avg.
no_adapt	34.93	33.67	36.41	32.38	22.22	36.61	31.10	22.15	26.73	53.16	61.17	50.27	31.96	54.52	55.87	38.88
Tent	51.93 ± 0.01	52.77 ± 0.00	53.05 ± 0.01	52.59 ± 0.03	45.53 ± 0.97	54.89 ± 0.01	48.07 ± 0.07	17.33 ± 0.87	19.79 ± 1.10	66.34 ± 0.09	73.19 ± 0.16	65.06 ± 0.01	46.44 ± 12.61	66.57 ± 0.03	64.53 ± 0.01	51.87
EATA	<u>55.72</u> ± 0.01	<u>56.94</u> ± 0.00	<u>56.76</u> ± 0.03	57.71 ± 0.02	<u>55.81</u> ± 0.05	62.07 ± 0.00	<u>60.45</u> ± 0.01	64.98 ± 0.48	63.74 ± 0.04	71.90 ± 0.01	76.71 ± 0.05	<u>67.81</u> ± 0.24	66.29 ± 0.00	72.48 ± 0.01	69.25 ± 0.00	<u>63.91</u>
SAR	51.93 ± 0.01	52.90 ± 0.01	53.06 ± 0.00	52.66 ± 0.00	47.08 ± 0.09	54.88 ± 0.05	48.42 ± 0.01	26.60 ± 0.31	34.87 ± 4.86	66.28 ± 0.13	72.79 ± 0.01	<u>64.67</u> ± 0.00	50.80 ± 0.58	66.40 ± 0.00	64.36 ± 0.00	53.97
NOTE	35.11 ± 0.00	33.86 ± 0.00	36.58 ± 0.00	32.71 ± 0.00	22.39 ± 0.00	36.92 ± 0.00	31.30 ± 0.00	22.33 ± 0.00	26.90 ± 0.00	53.14 ± 0.00	62.25 ± 0.02	51.18 ± 0.00	32.04 ± 0.00	54.64 ± 0.00	55.95 ± 0.00	39.15
MEMO	43.78 ± 0.00	43.25 ± 0.01	45.83 ± 0.00	32.90 ± 0.03	28.46 ± 0.03	45.67 ± 0.07	38.25 ± 0.12	30.08 ± 0.06	34.63 ± 0.15	54.80 ± 0.20	69.99 ± 0.37	56.43 ± 0.16	34.47 ± 0.01	62.56 ± 0.03	59.61 ± 0.04	45.38
DeYO	55.68 ± 0.01	56.92 ± 0.02	56.42 ± 0.03	<u>57.85</u> ± 0.01	55.71 ± 0.33	<u>62.68</u> ± 0.10	46.53 ± 0.38	66.42 ± 0.01	<u>65.46</u> ± 0.01	<u>72.45</u> ± 0.03	<u>78.38</u> ± 0.01	66.52 ± 0.05	<u>67.48</u> ± 0.04	<u>73.63</u> ± 0.00	<u>70.18</u> ± 0.01	63.49
NCTTA (ours)	57.58 ± 0.05	59.03 ± 0.23	58.98 ± 0.16	60.14 ± 0.04	58.90 ± 0.35	65.14 ± 0.04	63.21 ± 0.19	68.73 ± 0.29	67.80 ± 0.02	74.24 ± 0.16	78.75 ± 0.15	69.19 ± 0.15	69.25 ± 0.16	74.57 ± 0.06	71.39 ± 0.01	66.46

Table 2. Waterbirds average and worst-group accuracy (%) on the target domain. **Bold** = best, underline = second-best.

Method	Avg. Acc	Worst-group Acc
no_adapt	92.14	70.87
BN_adapt	91.20 ± 0.00	61.78 ± 0.02
Tent	91.66 ± 0.00	63.50 ± 0.91
EATA	91.53 ± 0.01	63.14 ± 0.34
SAR	91.27 ± 0.00	62.10 ± 0.23
NOTE	90.22 ± 0.03	41.12 ± 0.44
MEMO	92.47 ± 0.01	64.69 ± 0.16
DeYO	92.90 ± 0.07	75.65 ± 2.20
NCTTA (ours)	93.13 ± 0.00	76.56 ± 0.05

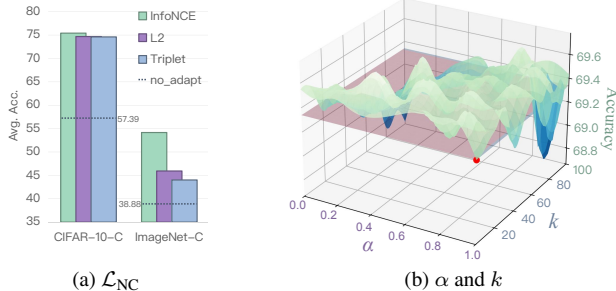


Figure 6. **Ablation of \mathcal{L}_{NC} , α and k .** (a) \mathcal{L}_{NC} is instantiated with three variants: InfoNCE-style, L2-style, and Triplet-style. (b) Sensitivity analysis of α and k , evaluated on ImageNet-C under the Contrast corruption at severity level 5.

EATA [23] and DeYO [15]. Furthermore, we provide the results of EATA-C [34], SAR² [25], COME [41], and AdaDEM [21] on CIFAR-10/100-C and ImageNet-C in the Appendix for completeness. Detailed descriptions of these baselines, along with their hyperparameter configurations, are provided in Appendix B.2. The backbone used in the experiments are ResNet50 [9] for CIFAR-10/100-C, WaterBirds and PACS, and ViT-B/16 [4] for ImageNet-C. Experiments are conducted on the TTAB [43] benchmark, with algorithmic parameters set to their default values unless otherwise specified.

We designed our experiments to answer the following questions: **RQ1:** How does NCTTA perform in comparison to baseline methods across various scenarios, including biased and wild scenarios that resemble real-world situations? **RQ2:** How to select the specific form of \mathcal{L}_{NC} ? And is the proposed NCTTA effective? **RQ3:** What role does the NCTTA play?

5.2. Main Results

To thoroughly evaluate the effectiveness of NCTTA, three challenging experimental setups are designed: 1) Mild Scenario: Each test batch consists of samples from a single corruption type or domain, simulating stable domain shifts. 2) Small Batch Size: The model processes one test sample per iteration, testing the adaptability of TTA methods under extreme resource constraints. 3) Continual Test-Time Adaptation (CTTA): The test data undergoes continual shifts across multiple domains, presenting a dynamic and evolving distribution. These setups are designed to assess NCTTA’s robustness and generalization capabilities under diverse and realistic conditions. The concrete implementation of NCTTA will be detailed in Section 5.3. We adopt Eq. 14 as the default instantiation of \mathcal{L}_{NC} .

In the mild shift scenario, where test samples are corrupted but drawn from a single domain (Table 1, Table 3 and Table 2), NCTTA consistently outperforms prior methods across all benchmarks. On CIFAR-10-C and ImageNet-C, NCTTA achieves the highest average accuracy, ranking first on nearly all corruption types. Similar gains are observed on PACS and Waterbirds, where NCTTA handles spurious correlations and domain shifts more effectively than existing methods. These results demonstrate the benefit of our feature-classifier alignment under domain shifts. Moreover, under certain corruptions, such as Snow and Frost on ImageNet-C, Tent performs even worse than no_adapt. In

Table 3. Classification accuracy (%) on PACS. The **bold** entries indicate the highest performance for each corruption type, and the underlined entries indicate the next-highest performance.

Source Target	Art			Photo			Cartoon			Avg.
	Cartoon	Photo	Sketch	Cartoon	Art	Sketch	Art	Photo	Sketch	
no_adapt	65.61	97.66	57.32	39.55	73.24	43.93	75.63	90.48	72.23	68.41
BN_adapt	74.67 \pm 0.03	96.85 \pm 0.02	69.32 \pm 0.16	64.39 \pm 0.01	77.17 \pm 0.11	45.85 \pm 0.09	82.137 \pm 0.03	94.69 \pm 0.14	73.05 \pm 0.09	75.35
Tent	75.17 \pm 0.02	96.93 \pm 0.02	71.52 \pm 0.08	65.09 \pm 0.11	77.39 \pm 0.21	<u>48.80\pm0.07</u>	83.11 \pm 0.00	<u>94.89\pm0.06</u>	75.00 \pm 0.08	76.43
EATA	74.70 \pm 0.03	96.85 \pm 0.02	69.01 \pm 0.02	64.32 \pm 0.00	77.22 \pm 0.14	45.58 \pm 0.00	82.13 \pm 0.03	94.69 \pm 0.14	73.05 \pm 0.05	75.28
SAR	74.94 \pm 0.03	96.85 \pm 0.02	69.38 \pm 0.05	64.51 \pm 0.03	77.31 \pm 0.22	45.85 \pm 0.08	82.37 \pm 0.03	94.81 \pm 0.10	73.63 \pm 0.12	75.52
NOTE	70.72 \pm 0.16	97.68 \pm 0.00	65.47 \pm 0.14	47.95 \pm 0.06	75.13 \pm 0.07	50.23\pm0.14	79.70 \pm 0.01	93.33 \pm 0.06	76.17\pm0.07	72.93
MEMO	69.51 \pm 0.01	98.00 \pm 0.00	61.37 \pm 0.01	45.42 \pm 0.02	74.67 \pm 0.02	44.28 \pm 0.00	80.16 \pm 0.02	93.87 \pm 0.01	74.44 \pm 0.00	71.30
DeYO	76.82 \pm 0.27	96.95 \pm 0.00	70.75 \pm 0.10	66.77 \pm 0.00	77.74 \pm 0.00	46.31 \pm 0.15	84.30 \pm 0.11	95.07\pm0.08	75.29 \pm 0.09	<u>76.67</u>
NCTTA (ours)	77.21\pm0.99	98.00\pm0.69	71.68\pm0.00	70.94\pm0.00	77.78\pm0.28	46.17 \pm 0.00	84.80\pm0.00	94.41 \pm 0.35	75.26 \pm 0.00	77.36

contrast, NCTTA can mitigate the shortcomings of the \mathcal{L}_{ENT} loss simply by incorporating \mathcal{L}_{NC} , without the need for additional augmentations or carefully designs.

In the wild shift scenario, where test domains change continually (Table 5) or the batch size is restricted to 1 (Table 6), NCTTA demonstrates robust and consistent performance, showcasing its ability to handle challenging real-world conditions. Specifically, in the CTTA setting, NCTTA achieves an impressive accuracy of 71.32%, significantly surpassing all other methods, which underscores its effectiveness in adapting to continuously evolving test domains. Furthermore, in the batch size 1 setting, a scenario that limits the availability of test samples and poses a severe challenge for most methods, NCTTA attains an accuracy of 59.55%, outperforming Tent by more than 10.36%. This substantial margin highlights the superiority of NCTTA in scenarios with limited batch sizes, where efficient and reliable adaptation is critical. Overall, these findings affirm the stability, adaptability, and scalability of NCTTA across diverse and dynamically shifting test-time conditions, making it a robust solution for real-world deployment.

Table 4. Ablation study of the proposed components. The table reports average classification accuracy (%) across the 15 corruption types of CIFAR-10-C and ImageNet-C at severity level 5. **Bold** numbers denote the best performance.

	\mathcal{L}_{ENT}	\mathcal{L}_{NC}	S_{ENT}	λ	CIFAR-10-C	ImageNet-C
(1)Tent	✓				75.19	51.13
(2)SAR	✓		✓		74.81	52.27
(3)		✓			75.31	53.17
(4)		✓	✓		74.95	51.39
(5)	✓	✓			76.25	54.84
(6)	✓	✓	✓		75.64	57.34
(7)		✓	✓	✓	76.87	53.97
(8)	✓		✓	✓	76.55	53.20
(9)NCTTA(ours)	✓	✓	✓	✓	78.10	66.50

5.3. Ablation Studies

In this work, we instantiate three variants of \mathcal{L}_{NC} , InfoNCE-style, L2-style, and Triplet-style, defined as:

$$\mathcal{L}_{\text{NC}}^{\text{InfoNCE}}(x_i) = -\log \frac{\frac{1}{|\mathcal{T}|} \sum_{j \in \mathcal{T}} \exp\left(\left\langle \frac{\mathbf{h}_i}{\|\mathbf{h}_i\|_2}, \frac{w_j}{\|w_j\|_2} \right\rangle\right)}{\sum_{j=0}^{K-1} \exp\left(\left\langle \frac{\mathbf{h}_i}{\|\mathbf{h}_i\|_2}, \frac{w_j}{\|w_j\|_2} \right\rangle\right)}, \quad (14)$$

$$\mathcal{L}_{\text{NC}}^{\text{L2}}(x_i) = \frac{1}{|\mathcal{T}|} \sum_{j \in \mathcal{T}} \left\| \frac{\mathbf{h}_i}{\|\mathbf{h}_i\|_2} - \frac{w_j}{\|w_j\|_2} \right\|_2 - \frac{1}{K - |\mathcal{T}|} \sum_{j \notin \mathcal{T}} \left\| \frac{\mathbf{h}_i}{\|\mathbf{h}_i\|_2} - \frac{w_j}{\|w_j\|_2} \right\|_2, \quad (15)$$

$$\mathcal{L}_{\text{NC}}^{\text{Triplet}}(x_i) = \max\left(0, \frac{1}{|\mathcal{T}|} \sum_{j \in \mathcal{T}} \left\| \frac{\mathbf{h}_i}{\|\mathbf{h}_i\|_2} - \frac{w_j}{\|w_j\|_2} \right\|_2 - \min_{j \notin \mathcal{T}} \left\| \frac{\mathbf{h}_i}{\|\mathbf{h}_i\|_2} - \frac{w_j}{\|w_j\|_2} \right\|_2 + \tau_{\text{margin}}\right), \quad (16)$$

where $\langle \cdot, \cdot \rangle$ denotes the similarity measure, instantiated as cosine similarity in our experiments. τ_{margin} is the margin hyperparameter, which we set to 1.0. We evaluate the three formulations on the 15 corruption types of CIFAR-10-C and ImageNet-C, setting $k = 1$ for each. As shown in Figure 6(a), all variants consistently improve performance by alleviating sample-wise misalignment, underscoring the importance of realigning feature embeddings and the corresponding classifier weights. Among them, the InfoNCE-style loss delivers the most stable and substantial gains. Therefore, we adopt it as the default instantiation of \mathcal{L}_{NC} in all subsequent experiments.

Table 5. Classification accuracy (%) on ImageNet-C under the CTTA scenario across severity levels 1 to 5. The **bold** entries indicate the highest performance for each corruption type, and the underlined entries indicate the next-highest performance.

CTTA	Level 1	Level 2	Level 3	Level 4	Level 5	Avg.
no_adapt	70.78	63.65	58.85	50.44	38.82	56.51
Tent	76.01	71.47	68.42	62.27	53.17	66.27
EATA	77.85	74.13	71.35	66.20	<u>58.65</u>	69.64
SAR	75.68	71.09	67.90	61.91	53.53	66.02
NOTE	63.82	63.68	58.83	50.61	39.18	55.22
MEMO	74.41	69.32	65.05	56.93	45.55	62.25
DeYO	78.40	74.49	71.75	66.28	57.88	69.76
NCTTA (ours)	78.82	75.44	72.84	68.15	61.11	71.27

To assess the efficacy of the proposed hybrid targets and the sensitivity of NCTTA to their associated parameters α and k , we conduct a controlled study summarized in Figure 6(b). Properly tuning α and k yields superior performance compared to the degenerate setting that solely constrains the P-FCA distance (i.e., $\alpha = 1.0$, $k = 1$). This result indicates that hybrid targets effectively mitigate the unreliability of pseudo-labels caused by sample-wise misalignment in adaptation.

Table 6. Classification accuracy (%) on ImageNet-C under the batch size = 1 scenario across severity level 5. The **bold** entries indicate the highest performance for each corruption type, and the underlined entries indicate the next-highest performance.

batch size = 1	Gauss.	Shot	Impul.	Defoc.	Glass	Motion	Zoom	Snow	Frost	Fog	Brit.	Contr.	Elastic	Pixel	JPEG	Avg.
no_adapt	34.93	33.67	36.41	32.38	22.22	36.61	31.10	22.15	26.73	53.16	61.17	50.27	31.96	54.52	55.87	38.88
Tent	47.15 \pm 0.67	46.40 \pm 1.03	47.51 \pm 0.02	45.19 \pm 1.04	34.01 \pm 0.36	47.57 \pm 0.33	39.52 \pm 0.22	37.29 \pm 4.59	40.03 \pm 2.24	58.97 \pm 0.05	69.43 \pm 0.65	60.61 \pm 0.06	39.88 \pm 0.19	62.76 \pm 2.00	61.49 \pm 0.23	49.19
EATA	48.56 \pm 0.16	48.17 \pm 0.28	49.62 \pm 0.09	48.17 \pm 0.80	39.10 \pm 0.84	51.34 \pm 0.38	44.37 \pm 1.44	45.01 \pm 2.01	46.00 \pm 3.08	64.84 \pm 0.11	72.59 \pm 0.52	61.45 \pm 0.43	47.65 \pm 3.65	63.97 \pm 1.56	64.03 \pm 0.12	52.99
SAR	40.32 \pm 0.04	39.79 \pm 0.41	42.77 \pm 1.08	38.73 \pm 0.08	27.73 \pm 0.08	45.13 \pm 0.21	36.85 \pm 0.36	30.86 \pm 3.35	36.04 \pm 0.93	58.06 \pm 0.14	68.43 \pm 0.42	56.47 \pm 0.44	37.16 \pm 1.27	63.34 \pm 19.33	60.95 \pm 1.43	45.51
NOTE	35.25 \pm 0.16	33.56 \pm 0.06	36.13 \pm 0.88	32.53 \pm 0.20	22.42 \pm 0.25	36.82 \pm 0.10	30.91 \pm 0.49	22.47 \pm 0.21	27.08 \pm 0.45	53.43 \pm 0.02	61.41 \pm 0.29	50.15 \pm 0.33	31.87 \pm 0.45	54.39 \pm 0.10	55.83 \pm 1.21	38.95
MEMO	37.31 \pm 0.34	35.19 \pm 0.40	39.04 \pm 0.01	32.13 \pm 0.27	22.57 \pm 0.40	38.44 \pm 0.16	31.81 \pm 0.00	23.78 \pm 0.26	28.87 \pm 0.00	52.91 \pm 0.02	62.27 \pm 0.06	52.30 \pm 0.06	31.96 \pm 0.23	54.77 \pm 0.06	56.27 \pm 0.28	39.97
DeYO	50.72 \pm 0.48	50.58 \pm 0.13	52.45 \pm 1.54	51.47 \pm 0.43	42.75 \pm 7.59	54.10 \pm 0.07	45.91 \pm 0.25	52.73 \pm 0.28	51.84 \pm 1.90	66.36 \pm 0.16	73.10 \pm 0.09	62.42 \pm 0.50	52.11 \pm 1.54	66.14 \pm 0.93	64.12 \pm 0.58	55.79
NCTTA (ours)	52.32 \pm 0.18	53.03 \pm 0.33	55.01 \pm 0.06	53.09 \pm 0.24	48.53 \pm 0.91	57.50 \pm 0.97	53.35 \pm 0.16	59.64 \pm 0.17	56.79 \pm 0.12	69.83 \pm 0.71	75.59 \pm 0.61	65.09 \pm 0.38	57.55 \pm 0.23	70.34 \pm 0.55	67.70 \pm 0.23	59.69

Table 4 reports ablations on CIFAR-10-C and ImageNet-C at severity level 5. Introducing \mathcal{L}_{NC} substantially improves robustness and complements the widely used entropy regularization \mathcal{L}_{ENT} . The full NCTTA configuration outperforms all partial variants, highlighting the synergistic effect of each component in strengthening feature-classifier alignment and increasing resilience to severe corruptions.

5.4. Role of NCTTA

To empirically validate the effectiveness of NCTTA, we systematically computed the average G-FCA distance \bar{d}_{iy_i} for each TTA batch and compared the results against baseline methods SAR and Tent. As illustrated in Figure 7, NCTTA consistently achieves lower \bar{d}_{iy_i} values compared to the baseline approaches across both CIFAR-10-C and ImageNet-C under Gaussian noise with severity level 5. This indicates that NCTTA effectively mitigates feature-classifier misalignment, thereby enhancing the model’s robustness to severe domain shifts.




Figure 7. **Comparison of G-FCA distance under Gaussian noise (ImageNet-C, ViT-B/16).** The plot compares d_{iy_i} for Tent, SAR and NCTTA on ImageNet-C under Gaussian noise with severity level 5. The results show that NCTTA consistently achieves lower d_{iy_i} compared to Tent and SAR, demonstrating better feature-classifier alignment and enhanced robustness.

Furthermore, t-SNE visualizations [35] presented in Figure 8 substantiate these quantitative findings. The feature embeddings processed with NCTTA form more distinct and well-separated clusters compared to those obtained using Tent. This enhanced clustering demonstrates NCTTA’s superior capability in maintaining discriminative feature spaces, which is crucial for accurate classification in OOD scenarios. Collectively, these results underscore that NCTTA not only improves feature-classifier alignment but also effectively distinguishes features under various distributional shifts, reinforcing the method’s robustness.

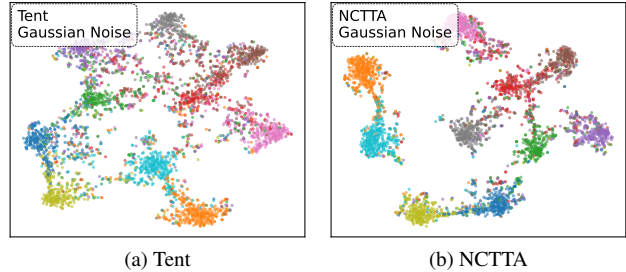


Figure 8. **t-SNE comparison under Gaussian noise.** The figure compares t-SNE visualizations of feature representations for Tent and NCTTA on CIFAR-10-C under Gaussian noise with severity level 5. NCTTA forms more distinct and well-separated clusters compared to Tent, enhancing the model’s discriminative power under severe corruption.

6. Conclusion and Limitation

In this work, we bridge NC theory and TTA by extending NC3 to the sample-wise regime, formulating NC3+. We theoretically and empirically show that, during the TPT, each sample’s feature embedding converges toward its corresponding classifier weight, and we reveal that domain shifts primarily harm performance by inducing sample-wise misalignment. Building on these insights, we propose NCTTA, an NC-guided adaptation framework that realigns feature embeddings to classifier weights using hybrid targets blending geometric proximity (FCA distance) and predictive confidence.

NCTTA consistently improves robustness across CIFAR-10-C/100-C, ImageNet-C, PACS, and Waterbirds, including challenging settings such as CTTA and batch size = 1, demonstrating strong generality and stability. Notably, NCTTA yielded significant improvements in the G-FCA distance d_{iy_i} , alongside enhanced feature clustering, as shown in t-SNE visualizations. Our work leaves several aspects unexplored. Specifically, further research is required to harness the potential of NC properties for computing or learning class prototypes in mini-batch TTA settings. Exploring the applicability of NC in vision-language models (VLMs) [3, 11, 14, 27, 32, 42] also remains an open direction for future investigation. This work provides a new perspective for advancing both the theoretical understanding and practical development of TTA methods.

Acknowledgment

This work was supported by the National Natural Science Foundation of China (Grant No. 92467204 and 62472249), the Shenzhen Science and Technology Program (Grant No. JCYJ20220818101014030, KJZD20240903102300001, and JCYJ20250604145014018), and the Natural Science Foundation of Top Talent of SZTU (Grant No. GDRC202413). We thank the anonymous reviewers for their efforts, helping improve the quality of this paper.

References

- [1] Mouin Ben Ammar, Nacim Belkhir, Sebastian Popescu, Antoine Manzanera, and Gianni Franchi. Neco: Neural collapse based out-of-distribution detection. *arXiv preprint arXiv:2310.06823*, 2023.
- [2] Xiao Chen, Qihui Zhang, and Yan Wang. Cliff: Leveraging ambiguous samples for enhanced test-time adaptation. In *27th European Conference on Artificial Intelligence*, pages 642–649, 2024.
- [3] Xiao Chen, Jiazhen Huang, Zhiming Liu, Qinting Jiang, Fanding Huang, Jingyan Jiang, and Zhi Wang. Test-time distillation for continual model adaptation. *arXiv preprint arXiv:2506.02671*, 2025.
- [4] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021.
- [5] Xiao Fan, Jingyan Jiang, Zhaoru Chen, Fanding Huang, Xiao Chen, Qinting Jiang, Bowen Zhang, Xing Tang, and Zhi Wang. Moetta: Test-time adaptation under mixed distribution shifts with moe-layernorm. *arXiv preprint arXiv:2511.13760*, 2025.
- [6] Cong Fang, Hangfeng He, Qi Long, and Weijie J Su. Layer-peeled model: Toward understanding well-trained deep neural networks. *arXiv preprint arXiv:2101.12699*, 4, 2021.
- [7] Cong Fang, Hangfeng He, Qi Long, and Weijie J Su. Exploring deep neural networks via layer-peeled model: Minority collapse in imbalanced training. *Proceedings of the National Academy of Sciences*, 118(43):e2103091118, 2021.
- [8] Taesik Gong, Jongheon Jeong, Taewon Kim, Yewon Kim, Jinwoo Shin, and Sung-Ju Lee. NOTE: Robust continual test-time adaptation against temporal correlation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [10] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*, 2019.
- [11] Fanding Huang, Jingyan Jiang, Qinting Jiang, Hebei Li, Faisal Nadeem Khan, and Zhi Wang. Cosmic: Clique-oriented semantic multi-space integration for robust clip test-time adaptation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 9772–9781, 2025.
- [12] Yusuke Iwasawa and Yutaka Matsuo. Test-time classifier adjustment module for model-agnostic domain generalization. In *Advances in Neural Information Processing Systems*, pages 2427–2440. Curran Associates, Inc., 2021.
- [13] Qinting Jiang, Chuyang Ye, Dongyan Wei, Bingli Wang, Yuan Xue, Jingyan Jiang, and Zhi Wang. Feature-based instance neighbor discovery: Advanced stable test-time adaptation in dynamic world. *arXiv preprint arXiv:2506.06782*, 2025.
- [14] Adilbek Karmanov, Dayan Guan, Shijian Lu, Abdulmoteleb El Saddik, and Eric Xing. Efficient test-time adaptation of vision-language models. *The IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- [15] Jonghyun Lee, Dahyun Jung, Saehyung Lee, Junsung Park, Juhyeon Shin, Uiwon Hwang, and Sungroh Yoon. Entropy is not enough for test-time adaptation: From the perspective of disentangled factors. *arXiv preprint arXiv:2403.07366*, 2024.
- [16] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Deeper, broader and artier domain generalization, 2017.
- [17] 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 (ICML)*, pages 6028–6039, 2020.
- [18] Jian Liang, Ran He, and Tieniu Tan. A comprehensive survey on test-time adaptation under distribution shifts. *arXiv preprint arXiv:2303.15361*, 2023.
- [19] Xiaofeng Liu, Chaehwa Yoo, Fangxu Xing, Hyejin Oh, Georges El Fakhri, Je-Won Kang, Jonghye Woo, et al. Deep unsupervised domain adaptation: A review of recent advances and perspectives. *APSIPA Transactions on Signal and Information Processing*, 11(1), 2022.
- [20] Jianfeng Lu and Stefan Steinerberger. Neural collapse under cross-entropy loss. *Applied and Computational Harmonic Analysis*, 59:224–241, 2022.
- [21] Jing Ma, Hanlin Li, and Xiang Xiang. Decoupled entropy minimization. *arXiv preprint arXiv:2511.03256*, 2025.
- [22] R Thomas McCoy, Ellie Pavlick, and Tal Linzen. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. *arXiv preprint arXiv:1902.01007*, 2019.
- [23] Shuaicheng Niu, Jiayang Wu, Yifan Zhang, Yafo Chen, Shijian Zheng, Peilin Zhao, and Mingkui Tan. Efficient test-time model adaptation without forgetting. In *International conference on machine learning*, pages 16888–16905. PMLR, 2022.
- [24] Shuaicheng Niu, Jiayang Wu, Yifan Zhang, Zhiquan Wen, Yafo Chen, Peilin Zhao, and Mingkui Tan. Towards stable test-time adaptation in dynamic wild world. *arXiv preprint arXiv:2302.12400*, 2023.
- [25] Shuaicheng Niu, Guohao Chen, Deyu Chen, Yifan Zhang, Jiayang Wu, Zhiquan Wen, Yafo Chen, Peilin Zhao, Chunyan Miao, and Mingkui Tan. Adapt in the wild: Test-time

- entropy minimization with sharpness and feature regularization. *arXiv preprint arXiv:2509.04977*, 2025.
- [26] Vardan Papyan, XY Han, and David L Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, 2020.
- [27] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021.
- [28] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. *Learning Internal Representations by Error Propagation*. 1985.
- [29] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Sathesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [30] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. In *International Conference on Learning Representations*, 2019.
- [31] Steffen Schneider, Evgenia Rusak, Luisa Eck, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. Removing covariate shift improves robustness against common corruptions. *CoRR*, abs/2006.16971, 2020.
- [32] Manli Shu, Weili Nie, De-An Huang, Zhiding Yu, Tom Goldstein, Anima Anandkumar, and Chaowei Xiao. Test-time prompt tuning for zero-shot generalization in vision-language models. *Advances in Neural Information Processing Systems*, 35:14274–14289, 2022.
- [33] Vivienne Sze, Yu-Hsin Chen, Tien-Ju Yang, and Joel S Emer. Efficient processing of deep neural networks: A tutorial and survey. *Proceedings of the IEEE*, 105(12):2295–2329, 2017.
- [34] Minghui Tan, Guohao Chen, Jiayang Wu, Yifan Zhang, Yafo Chen, Peilin Zhao, and Shuaicheng Niu. Uncertainty-calibrated test-time model adaptation without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.
- [35] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [36] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. *arXiv preprint arXiv:2006.10726*, 2020.
- [37] 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*, pages 7201–7211, 2022.
- [38] E Weinan and Stephan Wojtowytsch. On the emergence of simplex symmetry in the final and penultimate layers of neural network classifiers. In *Mathematical and Scientific Machine Learning*, pages 270–290. PMLR, 2022.
- [39] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? *Advances in neural information processing systems*, 27, 2014.
- [40] Marvin Zhang, Sergey Levine, and Chelsea Finn. Memo: Test time robustness via adaptation and augmentation. *Advances in neural information processing systems*, 35:38629–38642, 2022.
- [41] Qingyang Zhang, Yatao Bian, Xinke Kong, Peilin Zhao, and Changqing Zhang. Come: Test-time adaption by conservatively minimizing entropy. *arXiv preprint arXiv:2410.10894*, 2024.
- [42] Taolin Zhang, Jinpeng Wang, Hang Guo, Tao Dai, Bin Chen, and Shu-Tao Xia. Boostadapter: Improving test-time adaptation via regional bootstrapping. *arXiv preprint arXiv:2410.15430*, 2024.
- [43] Hao Zhao, Yuejiang Liu, Alexandre Alahi, and Tao Lin. On pitfalls of test-time adaptation. *arXiv preprint arXiv:2306.03536*, 2023.
- [44] Zhisheng Zhong, Jiequan Cui, Yibo Yang, Xiaoyang Wu, Xiaojuan Qi, Xiangyu Zhang, and Jiaya Jia. Understanding imbalanced semantic segmentation through neural collapse. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19550–19560, 2023.
- [45] Jinxin Zhou, Chong You, Xiao Li, Kangning Liu, Sheng Liu, Qing Qu, and Zhihui Zhu. Are all losses created equal: A neural collapse perspective. *Advances in Neural Information Processing Systems*, 35:31697–31710, 2022.
- [46] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4396–4415, 2022.
- [47] Didi Zhu, Zexi Li, Min Zhang, Junkun Yuan, Jiashuo Liu, Kun Kuang, and Chao Wu. Neural collapse anchored prompt tuning for generalizable vision-language models. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4631–4640, 2024.
- [48] Zhihui Zhu, Tianyu Ding, Jinxin Zhou, Xiao Li, Chong You, Jeremias Sulam, and Qing Qu. A geometric analysis of neural collapse with unconstrained features. *Advances in Neural Information Processing Systems*, 34:29820–29834, 2021.