

# DCAC: Dynamic Class-Aware Cache Creates Stronger Out-of-Distribution Detectors

Yanqi Wu<sup>1,2,3</sup>, Qichao Chen<sup>4</sup>, Runhe Lai<sup>1,2,3</sup>, Xinhua Lu<sup>1,2,3</sup>,  
Jia-Xin Zhuang<sup>5\*</sup>, Zhilin Zhao<sup>1,3</sup>, Wei-Shi Zheng<sup>1,2,3</sup>, Ruixuan Wang<sup>1,2,3\*</sup>

<sup>1</sup>School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou, China

<sup>2</sup>Peng Cheng Laboratory, Shenzhen, China

<sup>3</sup>Key Laboratory of Machine Intelligence and Advanced Computing, MOE, Guangzhou, China

<sup>4</sup>University of Nottingham Malaysia, Semenyih, Malaysia

<sup>5</sup>Hong Kong University of Science and Technology, Hong Kong, China

wuyq268@mail2.sysu.edu.cn, hcxc1@nottingham.my, {lairh5,luxh55}@mail2.sysu.edu.cn, jzhuangad@cse.ust.hk,  
zhaozhlin@mail.sysu.edu.cn, wszheng@ieee.org, wangruix5@mail.sysu.edu.cn

## Abstract

Out-of-distribution (OOD) detection remains a fundamental challenge for deep neural networks, particularly due to overconfident predictions on unseen OOD samples during testing. We reveal a key insight: OOD samples predicted as the same class, or given high probabilities for it, are visually more similar to each other than to the true in-distribution (ID) samples. Motivated by this class-specific observation, we propose DCAC (Dynamic Class-Aware Cache), a training-free, test-time calibration module that maintains separate caches for each ID class to collect high-entropy samples and calibrate the raw predictions of input samples. DCAC leverages cached visual features and predicted probabilities through a lightweight two-layer module to mitigate overconfident predictions on OOD samples. This module can be seamlessly integrated with various existing OOD detection methods across both unimodal and vision-language models while introducing minimal computational overhead. Extensive experiments on multiple OOD benchmarks demonstrate that DCAC significantly enhances existing methods, achieving substantial improvements, i.e., reducing FPR95 by 6.55% when integrated with ASH-S on ImageNet OOD benchmark.

**Code** — <https://github.com/wyqstan/DCAC>

## Introduction

In recent years, deep neural networks (DNNs) have demonstrated remarkable performance on many challenging tasks. However, their reliability drops significantly in open-world settings when encountering previously unseen classes, often producing overconfident predictions for OOD inputs. This poses significant safety risks in critical applications like autonomous driving and medical diagnosis. Existing methods typically fall into two groups: designing new scoring functions (Hendrycks and Gimpel 2017; Hendrycks et al. 2022; Liu et al. 2020; Ming et al. 2022), or training with auxiliary outliers (Chen et al. 2024b; Hendrycks, Mazeika, and Dietterich 2019; Wu et al. 2025). The latter is often more effective due to the usage of explicit outlier knowledge.

\*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: The t-SNE visualization of normalized image features for test samples predicted as the same class (left) and OOD detection performance using CLIP-B/16 (right).

However, if the OOD samples encountered during testing differ greatly from the outliers seen during training, the model may still produce overconfident predictions on these unseen OOD inputs. To address this limitation, test-time OOD detection has emerged as a promising direction. Existing methods (Fan et al. 2024; Yang et al. 2023, 2025) leverage test samples to perform immediate calibration or parameter updates to the model, improving robustness against previously unseen OOD inputs. Most of these methods rely on the predictive confidence, typically measured by the softmax probability assigned to the predicted class. However, they often overlook the rich information embedded in the predicted labels assigned to the OOD samples. This raises a natural question:

*Do OOD samples sharing the same predicted labels exhibit any underlying relationships?*

To answer this question, we analyzed samples classified into the same class and uncovered an interesting phenomenon. As illustrated on the left side of Fig. 1, OOD samples that are predicted as the same class exhibit significantly higher internal visual similarity. To further investigate this phenomenon, we categorize these OOD samples into overconfident and unconfident groups based on their predictive entropy, finding that visual similarity between unconfident and overconfident OOD samples tends to be higher than that between unconfident OOD and ID samples. In addition, we observe that

some OOD samples, although not predicted as this class, are assigned high predicted probabilities for it and also share a relatively high degree of visual similarity with the aforementioned OOD samples predicted as this class.

Based on this class-specific observation, we propose collecting unconfident OOD samples predicted as each ID class and using their visual similarity to overconfident OOD samples as a reference to calibrate the predictions of the overconfident ones. Motivated by this, we propose DCAC (Dynamic Class-Aware Cache), a training-free test-time calibration module. DCAC maintains separate caches for each ID class to collect unconfident samples and uses their visual features and predicted probabilities to calibrate overconfident predictions through a lightweight two-layer module inspired by Tip-Adapter (Zhang et al. 2022). For each test sample, its visual features are processed through this module to generate calibration logits that are combined with original predictions for enhanced OOD detection.

DCAC offers several practical advantages: it is a plug-and-play module that is architecture-agnostic, requires no additional training, and seamlessly integrates with diverse OOD detection methods across both unimodal and vision-language models (VLMs), further enhancing their OOD detection performance. Comprehensive evaluations demonstrate that DCAC consistently enhances existing baselines with substantial performance gains. The main contributions of this work are summarized below.

- We reveal that OOD samples predicted as the same class or assigned relatively high probabilities for that class tend to share higher visual similarity with each other than with true ID samples, and propose leveraging unconfident samples to calibrate overconfident OOD predictions.
- We introduce DCAC, a training-free and architecture-agnostic module that calibrates predictions using the visual features and probabilities from test samples.
- We show that DCAC can be flexibly fused with existing OOD detection methods, achieving consistent performance improvements across multiple benchmarks.

## Related Works

**Traditional Unimodal OOD Detection.** OOD detection has attracted growing research attention in recent years. Many methods have emerged to address the issue, which can generally be categorized into two groups: post-hoc methods and training-based methods. Post-hoc methods typically design an OOD scoring function by analyzing model outputs such as logits, feature representations or layer-wise statistics. MSP (Hendrycks and Gimpel 2017) is the first baseline using the maximum softmax probability as the OOD score. ODIN (Liang, Li, and Srikant 2018) enhances MSP by applying temperature scaling and input perturbation. Energy score (Liu et al. 2020) is introduced as a new scoring function for OOD detection. ReAct (Sun, Guo, and Li 2021) suppresses high activation values in the penultimate layer of visual features, while ASH (Djurisic et al. 2023) discards most activations in that layer. OptFs (Zhao et al. 2024) optimizes feature shaping via a closed-form solution. CADRef (Ling et al. 2025) enhances feature discrimination by decoupling

class-aware relative features based on sign alignment with model weights. On the other hand, training-based methods often use outlier knowledge during training. MOS (Huang and Li 2021) partitions ID classes into semantically distinct groups and treats samples from other groups as OOD during training. TagFog (Chen et al. 2024b) creates OOD samples by shuffling ID image patches, and FodFom (Chen et al. 2024a) utilizes Stable Diffusion to synthesize training OOD samples that closely resemble ID images.

**VLM-based OOD Detection.** With rapid development of VLMs such as CLIP (Radford et al. 2021), more and more researchers have begun to focus on the OOD detection problem in VLMs. MCM (Ming et al. 2022) was the first method to adapt CLIP for OOD detection by combining MSP with temperature scaling. GLMCM (Miyai et al. 2023b) further incorporates the maximum probability of local image patches. NegLabel (Jiang et al. 2024), CSP (Chen, Gao, and Xu 2024) and CMA (Lee et al. 2025) introduce text information that is semantically unrelated to ID classes to improve OOD discrimination. These are post-hoc methods built on vanilla CLIP. Many recent methods apply finetuning to adapt CLIP to downstream tasks. LoCoOp (Miyai et al. 2023a) and SCT (Yu et al. 2024) use entropy maximization strategies to push OOD visual embeddings away from ID-aligned text features in the multimodal space. Local-Prompt (Zeng et al. 2025) freezes the global prompt while learning local prompts with region-level regularization and negative sample enhancement. OspCoOp (Xu, Xiang, and Liang 2025) leverages ID background as auxiliary training OOD data.

**Test-Time OOD Detection.** Very recently, increasing attention has been given to leveraging information from real test-time scenarios to assist OOD detection. AUTO (Yang et al. 2023) performs test-time adaptation by applying stochastic gradient descent directly to the model, aiming to reduce confidence on potential OOD samples. Unlike AUTO that update the entire model that suffers from performance degradation under noisy test streams, AdaND (Cao et al. 2025) freezes the model and just trains a lightweight noise detector to filter out OOD samples. RTL++ (Fan et al. 2024) assumes a linear correlation between OOD scores and features. However, this simple assumption is challenged by the overlap in feature distributions between hard OOD and ID samples. TULIP (Zhang et al. 2025) introduces priors of ID and OOD distributions at test time to calibrate model uncertainty. Our proposed method calibrates model prediction based on the image features and the raw predictions of test samples, without relying on any external prior. OODD (Yang et al. 2025) maintains a priority queue to accumulate representative OOD features and leveraging them to calibrate the OOD score of test samples. Unlike OODD that updates a queue solely based on OOD scores, this study designs and updates the cache based on the entropy and predicted label of each test sample. Furthermore, our method focuses on calibrating output logits, whereas OODD adjusts the final OOD scores.

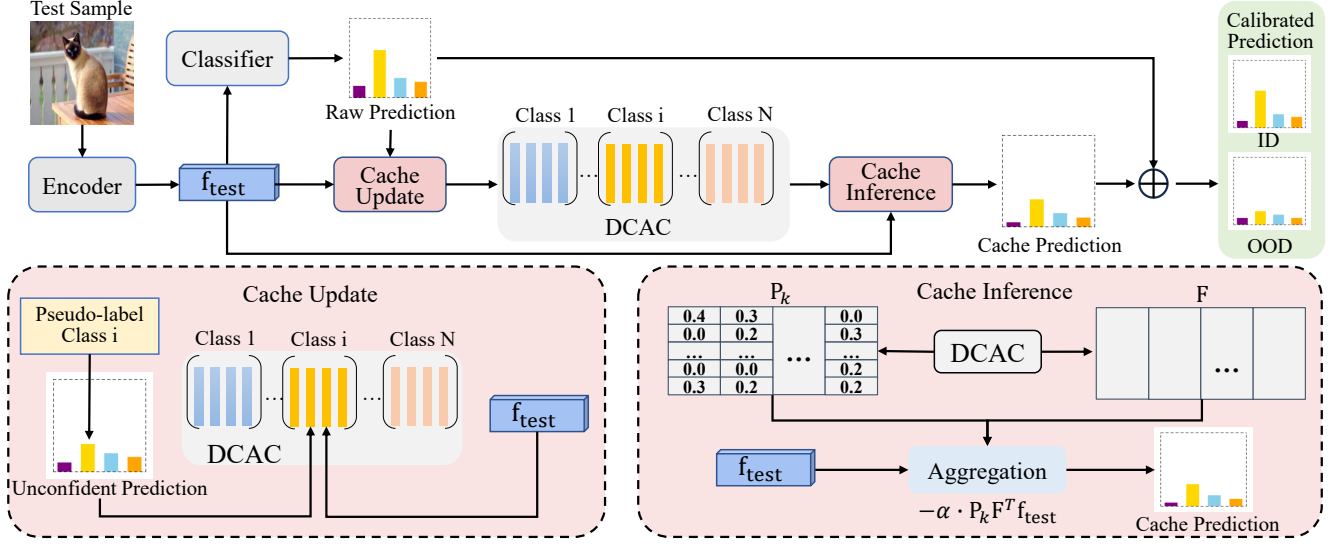


Figure 2: Overview of the Dynamic Class-aware Cache (DCAC) framework. DCAC maintains caches for each ID class to collect unconfident samples during testing guided by their entropy and the cache capacity. For calibration, cache samples generate a signal for each test sample which is combined with raw prediction to produce calibrated outputs for OOD detection.

## Methodology

### Preliminaries

**OOD Detection.** Consider a  $C$ -class classification problem where a deep neural network is trained on an ID training dataset  $\mathcal{D}_{\text{train}}$ . For an input image  $\mathbf{x}$ , the network produces a feature vector  $\mathbf{f} = \mathcal{F}(\mathbf{x}) \in \mathbb{R}^d$  through feature extractor  $\mathcal{F}$ , and class logits  $\mathbf{z}$  through a linear classifier:

$$\mathbf{z} = \mathbf{W}^\top \mathbf{f} + \mathbf{b}, \quad (1)$$

where  $\mathbf{W} \in \mathbb{R}^{d \times C}$  and  $\mathbf{b} \in \mathbb{R}^C$  are learnable parameters.

OOD detection determines whether the input  $\mathbf{x}$  belongs to the training distribution. A scoring function  $\mathcal{S}(\mathbf{z})$  is designed often based on the logits to assess distributional membership. Given threshold  $\tau$ , the detector outputs:

$$\mathcal{G}(\mathbf{x}) = \begin{cases} \text{ID}, & \text{if } \mathcal{S}(\mathbf{z}) \geq \tau \\ \text{OOD}, & \text{otherwise} \end{cases}. \quad (2)$$

### Overview

We propose a test-time prediction calibration method (Fig. 2) based on a key observation: for samples predicted as the same class, unconfident OOD samples show greater visual similarity to overconfident OOD samples than to true ID samples within that class. Our approach leverages this property by collecting unconfident samples during testing to calibrate overconfident predictions. The method has two components. First, we maintain dynamic class-aware caches that collect high-entropy samples and adapt to distribution shifts (Section 3.3). Second, we calibrate predictions using cached features and probability distributions as adaptive weights (Section 3.4). The calibrated predictions are then processed by OOD detection scoring functions  $\mathcal{S}(\cdot)$ .

### Dynamic Class-Aware Caches

To calibrate predictions of overconfident OOD samples, we first need to collect unconfident samples with similar characteristics. We maintain a separate cache for each ID class, initially as an empty set, that collects high-entropy samples during testing. Higher entropy indicates higher prediction uncertainty, making these samples likely candidates for OOD data that can inform calibration. We implement an entropy-based gating mechanism that admits only samples whose entropy exceeds threshold  $\delta$ . Since the true OOD entropy distribution is unknown at test time, we estimate  $\delta$  as the  $\beta$ -th percentile of the entropy distribution from ID training data, with  $\beta = 95$  in our experiments. This method ensures we preferentially collect potential OOD samples over confident ID samples.

To maintain computational efficiency, each cache has fixed capacity and follows a first-in-first-out (FIFO) update policy. When a high-entropy sample arrives, we store its visual features  $\mathbf{f}$  and predicted probability distribution  $\mathbf{p}$  in the cache for its predicted class. If the cache is full, we remove the oldest sample before storing a new one. This dynamic updating captures recent distribution patterns and enables rapid adaptation to evolving test conditions, providing fresh knowledge for calibrating subsequent predictions.

### Calibrate Model Outputs

Whenever a test sample is processed by the model and the cache has been updated accordingly, we immediately utilize all the information currently stored in the cache to calibrate the raw predictions of this sample. Let  $\mathbf{F} = [\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_N] \in \mathbb{R}^{d \times N}$  denote the cached feature matrix and  $\mathbf{P} = [\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N] \in \mathbb{R}^{C \times N}$  denote the corresponding probability matrix, where  $d$  is the visual feature dimension,  $N$  is the number of cached samples across all classes, and

$C$  is the number of ID classes. Based on  $\mathbf{F}$  and  $\mathbf{P}$ , a novel two-layer linear transformation module is designed. Specifically, the first layer is fully connected with weight parameters  $\mathbf{F}$ . This layer can be interpreted as computing the cosine similarity between the visual feature of the test sample and those of the cached samples. Similarly, the second layer is also fully connected with the weight parameters  $\mathbf{P}$ , which functions as a weighted aggregation on the similarity scores obtained from the first layer. Note that the weight parameters of the two-layer module are not learned, but set directly.

To reduce the influence of less relevant predictions from the cached samples, we apply a sparsification to the cached probability matrix  $\mathbf{P}$ . Specifically, for each cached sample, we retain only the top- $k$  largest values in its probability distribution and set the rest entries to zero. This operation can be formulated as:

$$\mathbf{P}_k = \mathbf{P} \odot \mathbf{M}, \quad (3)$$

where  $\mathbf{M} \in \{0, 1\}^{C \times N}$  is a binary mask that retains only the top- $k$  entries in each column of  $\mathbf{P}$ , and  $\odot$  denotes element-wise multiplication.

Ultimately, the cache prediction vector, which shares the same dimension as the original logit output, is computed as:

$$\mathbf{z}_{\text{cache}} = -\mathbf{P}_k \mathbf{F}^\top \mathbf{f}_{\text{test}}, \quad (4)$$

where  $\mathbf{f}_{\text{test}} \in \mathbb{R}^d$  denotes the visual feature of the test sample. We take the negative of  $\mathbf{z}_{\text{cache}}$  before combining it with the raw prediction  $\mathbf{z}$ . Since overconfident OOD samples yield high values for the predicted class in both  $\mathbf{z}$  and  $\mathbf{z}_{\text{cache}}$ , introducing the negative sign helps suppress excessive confidence in the final calibrated prediction. The original logit output is finally calibrated as:

$$\hat{\mathbf{z}} = \mathbf{z} + \alpha \cdot \mathbf{z}_{\text{cache}}, \quad (5)$$

where  $\alpha$  is a hyperparameter controlling calibration strength.  $\hat{\mathbf{z}}$  replaces  $\mathbf{z}$  in Eq. 2 to be fed into various scoring functions for OOD detection. In particular, the prediction of some recent methods (Chen, Gao, and Xu 2024; Jiang et al. 2024) includes both ID and OOD classes, our method only calibrates the ID-relevant components. The complete algorithmic process is detailed in Algorithm 1.

## Experiments

### Experimental Setup

**Datasets.** Following prior work (Huang and Li 2021), we use ImageNet-1K as ID dataset and evaluate on six OOD datasets: four far-OOD datasets including iNaturalist (Van Horn et al. 2018), SUN (Xiao et al. 2010), Places (Zhou et al. 2017), and Textures (Cimpoi et al. 2014), and two near-OOD datasets, NINCO (Bitterwolf, Mueller, and Hein 2023) and SSB-hard (Vaze et al. 2022). We further validate our method on five additional ID datasets: ImageNet100 (Ming et al. 2022), UCF101 (Soomro, Zamir, and Shah 2012), FGVC-Aircraft (Maji et al. 2013), CIFAR10 and CIFAR100 (Krizhevsky, Hinton et al. 2009).

**Implementation Details.** Our method is implemented in PyTorch and all experiments are conducted on NVIDIA GeForce A30 GPUs with a batch size of 512. To simulate a realistic deployment scenario, all experiments are conducted

---

### Algorithm 1: Test-time Calibration for a Single Sample

---

**Require:** Test sample  $\mathbf{x}$ , model  $f$ , cache  $\mathcal{C}$ , entropy threshold  $\delta$ , top- $k$ , correction factor  $\alpha$ , max cache size  $m$

- 1:  $\mathbf{z}, \mathbf{f}_{\text{test}} \leftarrow f(\mathbf{x})$
- 2:  $\mathbf{p} \leftarrow \text{softmax}(\mathbf{z}), \quad H \leftarrow H(\mathbf{p}), \quad \hat{y} \leftarrow \arg \max(\mathbf{p})$
- 3: **if**  $H > \delta$  **then**
- 4:    $\mathcal{B} \leftarrow \mathcal{C}[\hat{y}]$                             {Select buffer for predicted class}
- 5:   **if**  $|\mathcal{B}| \geq m$  **then**
- 6:     Remove the oldest element from  $\mathcal{B}$
- 7:   **end if**
- 8:   Append  $(\mathbf{f}_{\text{test}}, \mathbf{p})$  to  $\mathcal{B}$
- 9: **end if**
- 10: **Calibration:**
- 11:  $\mathbf{F} \leftarrow [\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_N] \in \mathbb{R}^{d \times N}$
- 12:  $\mathbf{P} \leftarrow [\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N] \in \mathbb{R}^{C \times N}$
- 13: Construct mask  $\mathbf{M} \in \{0, 1\}^{C \times N}$  such that only top- $k$  entries in each column of  $\mathbf{P}$  are 1
- 14:  $\mathbf{P}_k \leftarrow \mathbf{P} \odot \mathbf{M}$                             {Element-wise sparsification}
- 15:  $\mathbf{z}_{\text{cache}} \leftarrow -\mathbf{P}_k \mathbf{F}^\top \mathbf{f}_{\text{test}}$
- 16:  $\hat{\mathbf{z}} \leftarrow \mathbf{z} + \alpha \cdot \mathbf{z}_{\text{cache}}$
- 17: **return**  $\hat{\mathbf{z}}$

---

**Ensure:**  $\hat{\mathbf{z}}$

---

by mixing ID and OOD data together and shuffling them. For each experiment, we use five different random seeds for shuffling. For the ImageNet-1K benchmark, we set the cache size per class  $m$  to 20, calibration strength  $\alpha$  to 0.9, and top  $k$  to 20. Test samples with entropy above the 95th percentile ( $\beta = 95$ ) of ID training set are selected for caching.

**Metric.** The evaluation metrics include the false positive rate (F, FPR95) of OOD samples when the true positive rate of ID samples is at 95%, and the area under the receiver operating characteristic curve (A, AUROC).

**Baseline.** We choose various competitive methods as baseline for comparison. For unimodal methods, we select MSP (Hendrycks and Gimpel 2017), Energy (Liu et al. 2020), ReAct (Sun, Guo, and Li 2021), ASH-S (Djurisic et al. 2023), DDCS (Yuan et al. 2024), LINE (Ahn, Park, and Kim 2023), OODD (Yang et al. 2025) as baseline methods. For VLM-based method, we select MCM (Ming et al. 2022), CSP (Chen, Gao, and Xu 2024), CMA (Lee et al. 2025), AdaND (Cao et al. 2025), OODD, OSPCoOp (Xu, Xiang, and Liang 2025) as baseline methods. Specifically, OODD and AdaND are test-time methods and AdaND achieves state-of-art performance on ImageNet-1K far-OOD benchmark. To ensure fair comparison, we retain the original hyperparameter settings of the baseline methods.

### Performance on Benchmarks

**ImageNet-1K benchmarks.** As shown in Tab. 1, we compare the performance on the far-OOD benchmark using ResNet-50 and CLIP-B/16 as backbones. When combined with ASH-S and CSP, DCAC achieves superior performance, demonstrating the effectiveness of our method. Additional results of DCAC integrated with other OOD detection methods are illustrated in Fig. 3. Besides, we conduct compatibility experiments on the near-OOD benchmark, combining DCAC with various OOD detection methods (Tab. 2). All results support efficacy of our method.

Architecture	Method	OOD Datasets								Average	
		iNaturalist		SUN		Textures		Places			
		F↓	A↑	F↓	A↑	F↓	A↑	F↓	A↑	F↓	A↑
ResNet50	MSP	52.73	88.42	68.58	81.75	66.15	80.46	71.59	80.63	64.76	82.82
	Energy	53.96	90.59	58.28	86.73	52.3	86.73	65.43	84.12	57.49	87.04
	ReAct	19.55	96.39	24.01	94.41	45.83	90.45	33.45	91.93	30.71	93.30
	OptFS	16.79	96.88	35.31	93.13	23.08	95.74	44.78	90.42	29.99	94.04
	CADRef	16.08	96.9	39.23	91.26	12.60	97.14	51.12	87.8	29.76	93.28
	OODD	<u>7.13</u>	<b>98.71</b>	42.36	92.07	16.16	97.01	53.59	87.13	29.81	93.73
	ASH-S	11.49	97.87	27.96	94.02	<u>11.97</u>	<u>97.60</u>	39.83	90.98	22.81	95.12
	LINE	12.26	97.56	19.48	95.26	22.54	94.44	<b>28.52</b>	92.85	20.70	95.03
	DACS	11.63	97.85	<u>18.63</u>	<u>95.68</u>	18.40	95.77	<u>28.78</u>	<u>92.89</u>	<u>19.36</u>	<u>95.55</u>
	<b>ASH-S+Ours</b>	<b>6.69</b>	<u>98.54</u>	<b>17.47</b>	<b>96.21</b>	<b>11.70</b>	<b>97.74</b>	29.18	<b>93.11</b>	<b>16.26</b>	<b>96.40</b>
CLIP-B/16	MCM	31.95	94.16	37.22	92.55	58.35	85.83	42.98	90.10	42.63	90.66
	SCT	13.94	95.86	20.55	95.33	41.51	89.06	29.86	92.24	26.47	93.37
	CMA	8.88	98.19	29.03	93.32	51.26	87.00	27.60	93.87	29.19	93.10
	AdaND	4.19	98.91	17.08	95.86	<u>21.76</u>	93.01	<b>20.95</b>	<b>94.55</b>	<u>16.00</u>	95.58
	OODD	2.22	99.36	21.49	95.01	30.69	93.27	44.76	87.10	24.79	93.69
	OSPCoOp	15.25	97.13	18.26	96.74	41.26	91.13	25.74	<u>94.01</u>	25.13	94.75
	CSP	<u>1.54</u>	99.60	<u>13.66</u>	96.66	25.52	93.86	29.32	92.90	17.51	95.76
	<b>CSP+Ours</b>	<b>1.20</b>	<b>99.62</b>	<b>10.18</b>	<b>97.58</b>	<b>20.62</b>	<b>95.27</b>	<u>25.29</u>	93.94	<b>14.32</b>	<b>96.60</b>

Table 1: Comparison with state-of-the-art methods on the far-OOD benchmarks. The best results are in bold, while the second are underlined.  $\uparrow$  indicates larger values are better and  $\downarrow$  indicates smaller values are better.




Figure 3: DCAC integrated with existing methods.

**Other ID benchmarks.** DCAC also provides effective prediction calibration on other ID datasets, as shown in Tab. 3. For the ImageNet100, UCF101 and FGVC-Aircraft datasets, their OOD datasets are the same as the far-OOD datasets used for ImageNet-1K. For the CIFAR datasets, their OOD datasets include SVHN (Netzer et al. 2011), LSUN-R (Wang et al. 2016), LSUN-C, iSUN (Pan and Nieto 2015), Textures (Cimpoi et al. 2014) and Places365 (Zhou et al. 2017).

### Analysis of the Proposed Method

For consistency, all subsequent experiments are conducted using the combination of DCAC and MCM as a representative example for in-depth analysis.

**Impact of Cache Initialization.** Following OODD (Yang et al. 2025), we present the results on the ImageNet-1K benchmark using DCAC with different cache initialization strategies in Tab. 4, including (1) Empty: initializing the cache as empty, (2) C-Out: using ID training data with random cropping and initializing the cache with high-entropy samples, (3) D-Out: initializing the cache with OOD samples from OpenImage dataset (different from the test OOD

Architecture	Method	NINCO		ssb-hard	
		F↓	A↑	F↓	A↑
ResNet50	MSP	75.94	79.97	84.54	72.16
	<b>MSP+Ours</b>	<b>70.58</b>	<b>82.61</b>	<b>74.49</b>	<b>77.38</b>
	CADRef	64.89	85.36	78.79	74.58
	<b>CADRef+Ours</b>	<b>63.64</b>	<b>85.46</b>	<b>74.32</b>	<b>78.71</b>
	ASH-S	65.02	82.77	82.53	70.49
	<b>ASH-S+Ours</b>	<b>63.36</b>	<b>83.89</b>	<b>69.95</b>	<b>81.32</b>
CLIP/B-16	MCM	79.50	74.57	89.90	62.95
	<b>MCM+Ours</b>	<b>75.34</b>	<b>75.72</b>	<b>83.40</b>	<b>69.41</b>
	LoCoOp	78.27	73.58	87.92	64.10
	<b>LoCoOp+Ours</b>	<b>73.82</b>	<b>74.53</b>	<b>81.82</b>	<b>66.59</b>
	CSP	69.01	77.69	82.74	71.93
<b>CSP+Ours</b>	<b>66.71</b>	<b>80.52</b>	<b>78.90</b>	<b>76.12</b>	

Table 2: OOD detection performance on near-OOD datasets.

data), and (4) T-Out: initializing the cache with test OOD samples. For each initialization strategy, 800 samples are pre-stored in the cache in advance. In addition, we further explore how different initialization strategies influence the stability during the early stages of testing. We conduct experiments using ImageNet-1K as the ID dataset and Places as the OOD dataset. This experiment is repeated five times, and the AUROC is recorded over the first 10 batch, which are presented in Fig. 4. The results show that all initialization methods exhibit unstable performance during the early stages of testing, with the empty initialization being the most severe, but the performance gradually stabilizes as testing progresses. Although initializing the cache with OOD data




Figure 4: AUROC over the first 10 batches under different cache initialization strategies with mean and standard deviation.

ID Dataset	MCM		MCM+Ours	
	F↓	A↑	F↓	A↑
ImageNet100	30.05	95.02	<b>11.79</b>	<b>97.67</b>
UCF101	23.28	95.01	<b>18.19</b>	<b>95.99</b>
FGVC-Aircraft	13.92	96.70	<b>11.19</b>	<b>97.41</b>
CIFAR10	15.09	96.02	<b>12.73</b>	<b>96.51</b>
CIFAR100	86.11	76.42	<b>68.93</b>	<b>79.76</b>

Table 3: OOD detection performance on other ID datasets.

Strategy	C-Out	D-Out	T-Out	Empty
F↓	25.76	24.73	<b>21.53</b>	24.25
A↑	94.35	94.60	<b>95.28</b>	94.68

Table 4: OOD detection performance on ImageNet far-OOD benchmark with different cache initialized strategies.

can partially alleviate early-stage instability, and T-Out even enhances the final OOD detection performance, obtaining OOD samples requires substantial human and material resources. However, in real-world scenarios, the OOD distribution is inherently unknown, making it difficult to acquire samples that match the true distribution. As shown by the results of D-Out, when the initial samples do not reflect the true OOD distribution, the final performance may even be inferior to that of an empty-cache initialization. The C-Out strategy may introduce ID samples into the cache, which can degrade OOD detection performance.

**Impact of Cache Update Strategies.** We consider three cache update strategies: (1) first-in-first-out (FIFO), (2) replacing the sample with the highest entropy in the cache (RH), and (3) replacing the sample with the lowest entropy in the cache (RL). Considering that the distribution of test data may shift over time in real-world scenarios, we simulate a temporal distribution shift where OOD data appears sequentially: Textures in  $t_0 \rightarrow t_1$ , iNaturalist in  $t_1 \rightarrow t_2$ , SUN in  $t_2 \rightarrow t_3$ , and Places in  $t_3 \rightarrow t_4$ , denoted as  $T \rightarrow i \rightarrow S \rightarrow P$ . The OOD detection performance of the three cache update strategies under this setting is shown in Fig. 5. During the experiments, both RH and RL strategies tend to retain certain low- or high-entropy samples in the cache for extended periods, making them less likely to be replaced. This reduces the overall data turnover within the cache. In contrast, the FIFO strategy updates solely based on




Figure 5: AUROC for the temporal drift OOD scenario settings with different update strategies.




Figure 6: AUROC for the temporal drift OOD scenario settings with different cache construction strategies.

temporal order, ensuring higher cache fluidity and enabling better adaptation to the most recent distribution shifts.

**Impact of Cache Construction Strategies.** DCAC maintains a fixed-capacity cache for each class, making it class-aware. To assess its effectiveness, we compare it against a class-agnostic cache (non-class-aware, NCA), with both caches updated using a FIFO strategy. We simulate a temporal distribution shift in which OOD data appears sequentially, following the ImageNet-1K far-OOD benchmark setup in Fig.5. We set DCAC to store up to 20 entries per ID class, and the NCA cache to hold maximum 20000 entries in total. As shown in Fig. 6, DCAC consistently outperforms the NCA cache under these conditions. We attribute this to class bias in the OOD datasets. For example, when the OOD test set is Places, which contains a large number of landscape images, these samples are more likely to be predicted as ID classes related to scenery. Consequently, the NCA cache tends to accumulate a large number of landscape images. In contrast, DCAC, due to its per-class cache capacity constraint, only stores a limited number of such images. When the OOD distribution shifts abruptly (e.g., from Places to iNaturalist), the large volume of landscape samples in the NCA cache becomes mismatched with the new distribution,




Figure 7: Sensitivity analysis of hyper-parameters. Each dashed line represents the baseline (MCM) performance.




Figure 8: Comparison of the MCM score distribution on iNaturalist before and after calibration.

impairing the calibration of incoming data. By maintaining a balanced per-class cache, We substantially reduces the detrimental impact of the previous distribution on the new one.

**Hyperparameter Sensitivity Analysis.** We evaluate our method through sensitivity experiments. The hyperparameters evaluated and their ranges are as follows: per-class cache size in the range of [1, 30], calibration strength  $\alpha$  in [0.1, 2.5], percentile thresholds  $\beta$  for sample selection in [80, 99], and the number of top confidence scores retained per sample  $k$  in [1, 1000]. As shown in Fig. 7, our method maintains superior performance over the baseline across wide hyperparameter intervals.

**Effectiveness of Calibration.** As shown in Fig. 8, after applying DCAC for prediction calibration, the overlap between ID and OOD samples in MCM score distribution is significantly reduced, further demonstrating calibration efficacy.

**Computation Cost.** During inference, DCAC incurs minimal computational overhead. When integrated into CLIP-B/16, it adds 134 MB of memory usage and slightly reduces inference speed from 523 FPS to 511 FPS. For ResNet-50, the memory overhead is 75 MB, with FPS dropping from 1034 to 1027. These results confirm DCAC maintains efficient inference with only modest resource cost.

### Further Discussion

To validate our observation, we compute and analyze the average visual similarity on both ResNet50 and CLIP-B/16 between unconfident OOD samples predicted as each ID class and two groups: (1) overconfident OOD samples assigned high confidence to the same class, and (2) ID samples predicted to be the same class. As shown in Fig. 9, the average similarity between the two types of OOD samples is higher than that between unconfident OOD samples and ID samples, which verifies our observation.




Figure 9: Visual similarity comparison between ID and OOD samples using CLIP-B/16 (left) and ResNet-50 (right).

This phenomenon may be attributed to the following reasons. Previous studies (Ming, Yin, and Li 2022; Pezeshki et al. 2021) have shown that DNNs trained on ID data often overfits background cues shared among samples of the same class, leading to misclassification of OOD samples possessing similar backgrounds. Meanwhile, in VLMs (e.g., CLIP), we observed a distinct issue: some OOD samples, despite clear visual dissimilarities, exhibit unexpectedly high similarity with textual embeddings of specific ID classes, resulting in overconfident predictions. DCAC addresses both scenarios effectively by storing samples with representative background information and misclassified OOD examples and it significantly reduces false confidence without the need for extensive fine-tuning or additional labeled data.

### Conclusion

This paper tackles the critical problem of overconfident OOD predictions by revealing a key insight: OOD samples with high probability in the same class exhibit greater visual similarity to each other than to true ID samples. Based on this observation, we propose DCAC, a training-free test-time calibration module that maintains class-aware caches to collect unconfident samples and calibrate overconfident predictions. DCAC demonstrates remarkable versatility, seamlessly integrating with diverse OOD detection methods across both unimodal and vision-language models while consistently delivering substantial performance improvements. Our extensive experiments show significant gains on challenging benchmarks with minimal computational overhead. This work provides a practical and effective solution for enhancing model robustness in real-world deployments where distribution shifts are inevitable.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (Grant No. 62571559), the Major Key Project of PCL (Grant No. PCL2025AS209), the Guangdong Excellent Youth Team Program (Grant No. 2023B1515040025), the Research Start-up Fund for Introduced Talents at Sun Yat-sen University (Grant No. 67000-12255004), and the Young Scientist Program of Sun Yat-sen University (Grant No. 67000-13130003).

## References

- Ahn, Y. H.; Park, G.-M.; and Kim, S. T. 2023. Line: Out-of-distribution detection by leveraging important neurons. In *CVPR*.
- Bitterwolf, J.; Mueller, M.; and Hein, M. 2023. In or Out? Fixing ImageNet Out-of-Distribution Detection Evaluation. In *ICML*.
- Cao, C.; Zhong, Z.; Zhou, Z.; Liu, T.; Liu, Y.; Zhang, K.; and Han, B. 2025. Noisy Test-Time Adaptation in Vision-Language Models. arXiv:2502.14604.
- Chen, J.; Deng, L.; Gan, Z.; Zheng, W.-S.; and Wang, R. 2024a. FodFoM: Fake Outlier Data by Foundation Models Creates Stronger Visual Out-of-Distribution Detector. In *ACM MM*.
- Chen, J.; Zhang, T.; Zheng, W.-S.; and Wang, R. 2024b. Tag-fog: textual anchor guidance and fake outlier generation for visual out-of-distribution detection. In *AAAI*.
- Chen, M.; Gao, J.; and Xu, C. 2024. Conjugated semantic pool improves ood detection with pre-trained vision-language models. In *NeurIPS*.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing textures in the wild. In *CVPR*.
- Djurisic, A.; Bozanic, N.; Ashok, A.; and Liu, R. 2023. Extremely Simple Activation Shaping for Out-of-Distribution Detection. In *ICLR*.
- Fan, K.; Liu, T.; Qiu, X.; Wang, Y.; Huai, L.; Shanguan, Z.; Gou, S.; Liu, F.; Fu, Y.; Fu, Y.; et al. 2024. Test-time linear out-of-distribution detection. In *CVPR*.
- Hendrycks, D.; Basart, S.; Mazeika, M.; Mostajabi, M.; Steinhardt, J.; and Song, D. 2022. Scaling out-of-distribution detection for real-world settings. In *ICML*.
- Hendrycks, D.; and Gimpel, K. 2017. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *ICLR*.
- Hendrycks, D.; Mazeika, M.; and Dietterich, T. 2019. Deep anomaly detection with outlier exposure. In *ICLR*.
- Huang, R.; and Li, Y. 2021. MOS: Towards Scaling Out-of-distribution Detection for Large Semantic Space. In *CVPR*.
- Jiang, X.; Liu, F.; Fang, Z.; Chen, H.; Liu, T.; Zheng, F.; and Han, B. 2024. Negative Label Guided OOD Detection with Pretrained Vision-Language Models. In *ICLR*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.(2009).
- Lee, Y.; Cao, X.; Guo, J.; Ye, W.; Guo, Q.; and Chang, Y. 2025. Concept Matching with Agent for Out-of-Distribution Detection. In *AAAI*.
- Liang, S.; Li, Y.; and Srikant, R. 2018. Enhancing the reliability of out-of-distribution image detection in neural networks. In *ICLR*.
- Ling, Z.; Chang, Y.; Zhao, H.; Zhao, X.; Chow, K.; and Deng, S. 2025. CADRef: Robust Out-of-Distribution Detection via Class-Aware Decoupled Relative Feature Leveraging. In *CVPR*.
- Liu, W.; Wang, X.; Owens, J.; and Li, Y. 2020. Energy-based out-of-distribution detection. In *NeurIPS*.
- Maji, S.; Rahtu, E.; Kannala, J.; Blaschko, M.; and Vedaldi, A. 2013. Fine-grained visual classification of aircraft. arXiv:1306.5151.
- Ming, Y.; Cai, Z.; Gu, J.; Sun, Y.; Li, W.; and Li, Y. 2022. Delving into Out-of-Distribution Detection with Vision-Language Representations. In *NeurIPS*.
- Ming, Y.; Yin, H.; and Li, Y. 2022. On the impact of spurious correlation for out-of-distribution detection. In *AAAI*.
- Miyai, A.; Yu, Q.; Irie, G.; and Aizawa, K. 2023a. LoCoOp: Few-Shot Out-of-Distribution Detection via Prompt Learning. In *NeurIPS*.
- Miyai, A.; Yu, Q.; Irie, G.; and Aizawa, K. 2023b. Zero-shot in-distribution detection in multi-object settings using vision-language foundation models. arXiv:2304.04521.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; and Ng, A. Y. 2011. Reading Digits in Natural Images with Un-supervised Feature Learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*.
- Pan, J.; and i Nieto, X. G. 2015. End-to-end Convolutional Network for Saliency Prediction. arXiv:1507.01422.
- Pezeshki, M.; Kaba, O.; Bengio, Y.; Courville, A. C.; Precup, D.; and Lajoie, G. 2021. Gradient starvation: A learning proclivity in neural networks. *NeurIPS*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*.
- Soomro, K.; Zamir, A. R.; and Shah, M. 2012. UCF101: A dataset of 101 human actions classes from videos in the wild. arXiv:1212.0402.
- Sun, Y.; Guo, C.; and Li, Y. 2021. React: Out-of-distribution detection with rectified activations. In *NeurIPS*.
- Van Horn, G.; Mac Aodha, O.; Song, Y.; Cui, Y.; Sun, C.; Shepard, A.; Adam, H.; Perona, P.; and Belongie, S. 2018. The inaturalist species classification and detection dataset. In *CVPR*.
- Vaze, S.; Han, K.; Vedaldi, A.; and Zisserman, A. 2022. Open-Set Recognition: a Good Closed-Set Classifier is All You Need? In *ICLR*.
- Wang, L.; Guo, S.; Huang, W.; Xiong, Y.; and Qiao, Y. 2016. Knowledge Guided Disambiguation for Large-Scale Scene Classification With Multi-Resolution CNNs. *IEEE TIP*, 26: 2055–2068.
- Wu, Y.; Yu, R.; Cheng, X.; He, Z.; and Huang, X. 2025. Pursuing Feature Separation based on Neural Collapse for Out-of-Distribution Detection. In *ICLR*.

Xiao, J.; Hays, J.; Ehinger, K. A.; Oliva, A.; and Torralba, A. 2010. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*.

Xu, Z.; Xiang, X.; and Liang, Y. 2025. Overcoming Shortcut Problem in VLM for Robust Out-of-Distribution Detection. In *CVPR*.

Yang, P.; Liang, J.; Cao, J.; and He, R. 2023. Auto: Adaptive outlier optimization for online test-time ood detection. arXiv:2303.12267.

Yang, Y.; Zhu, L.; Sun, Z.; Liu, H.; Gu, Q.; and Ye, N. 2025. OODD: Test-time Out-of-Distribution Detection with Dynamic Dictionary. In *CVPR*.

Yu, G.; Zhu, J.; Yao, J.; and Han, B. 2024. Self-calibrated tuning of vision-language models for out-of-distribution detection. In *NeurIPS*.

Yuan, Y.; He, R.; Dong, Y.; Han, Z.; and Yin, Y. 2024. Discriminability-Driven Channel Selection for Out-of-Distribution Detection. In *CVPR*.

Zeng, F.; Cheng, Z.; Zhu, F.; Wei, H.; and Zhang, X.-Y. 2025. Local-prompt: Extensible local prompts for few-shot out-of-distribution detection. In *ICLR*.

Zhang, R.; Wei, Z.; Fang, R.; Gao, P.; Li, K.; Dai, J.; Qiao, Y.; and Li, H. 2022. Tip-Adapter: Training-free Adaption of CLIP for Few-shot Classification. In *ECCV*.

Zhang, Y.; Wu, D.; Wada, Y.; and Kanamori, T. 2025. TULiP: Test-time Uncertainty Estimation via Linearization and Weight Perturbation. arXiv:2505.16923.

Zhao, Q.; Xu, M.; Gupta, K.; Asthana, A.; Zheng, L.; and Gould, S. 2024. Towards Optimal Feature-Shaping Methods for Out-of-Distribution Detection. In *ICLR*.

Zhou, B.; Lapedriza, A.; Khosla, A.; Oliva, A.; and Torralba, A. 2017. Places: A 10 million image database for scene recognition. *IEEE TPAMI*, 40(6): 1452–1464.