

Self-Supervised Representation Learning for Adversarial Attack Detection

Yi Li[✉], Plamen Angelov[✉], and Neeraj Suri[✉]

Computing and Communications, Lancaster University, UK
y.li154@lancaster.ac.uk

Abstract. Supervised learning-based adversarial attack detection methods rely on a large number of labeled data and suffer significant performance degradation when applying the trained model to new domains. In this paper, we propose a self-supervised representation learning framework for the adversarial attack detection task to address this drawback. Firstly, we map the pixels of augmented input images into an embedding space. Then, we employ the prototype-wise contrastive estimation loss to cluster prototypes as latent variables. Additionally, drawing inspiration from the concept of memory banks, we introduce a discrimination bank to distinguish and learn representations for each individual instance that shares the same or a similar prototype, establishing a connection between instances and their associated prototypes. We propose a parallel axial-attention (PAA)-based encoder to facilitate the training process by parallel training over height- and width-axis of attention maps. Experimental results show that, compared to various benchmark self-supervised vision learning models and supervised adversarial attack detection methods, the proposed model achieves state-of-the-art performance on the adversarial attack detection task across a wide range of images.

Keywords: Self-supervised learning, adversarial attack detection, prototype, contrastive learning, discrimination bank

1 Introduction

Given an image potentially perturbed by an attack algorithm, the goal of adversarial attack detection is to distinguish between adversarial and normal samples using the differences between them. Adversarial attack detection is an important security topic applicable in real-world applications such as autonomous driving systems, object detection, medical image processing, and robotics [28] [39] [46] [48] among many others. With recent advancements in deep learning, several neural-network-based approaches have been proposed for adversarial attack detection [19] [24]. These networks and approaches are predominantly trained in a supervised manner, where a large number of labeled adversarial and normal samples are provided as input to neural networks. The model is then trained to reconstruct the corresponding clean sample and compare it with the input

sample to provide the detection result. Consequently, supervised learning-based adversarial attack detection approaches suffer from three main drawbacks.

Firstly, human-imperceptible adversarial attacks on images are challenging to label manually. This process can be time-consuming and may introduce errors, particularly when the annotator lacks familiarity with the task. Secondly, the trained adversarial attack detection models may need to be deployed in previously unseen conditions, including novel attack algorithms and datasets. Consequently, there is a strong likelihood of a mismatch between the training and testing conditions. In such cases, we lack the ability to leverage recorded test data to improve the model’s performance in the unseen test setting. Thirdly, prototype-based adversarial attack detection methods [36] [37] estimate an object’s category (e.g., cats or dogs) as the prototype. These methods calculate the degree of similarity between new data samples and autonomously chosen prototypes to classify images as adversarial or normal samples. However, each prototype may potentially consists of multiple instance samples, which often leads to a neglect of the rich intrinsic semantic relationships between prototypes of individual objects in images. For example, while the model may be trained on some tank images, it may struggle to classify new tanks or entirely new classes of objects when faced with previously unseen types of tanks.

To overcome these drawbacks, our contributions are summarized as follows:

- We propose a self-supervised representation learning framework aimed at extracting feature representations for the downstream task, i.e., adversarial attack detection. Building upon pixel mapping 3.2 and contrastive estimation in 3.3, in 3.4, we propose a discrimination bank to distinguish individual instances for each prototype from the embedding space. We demonstrate that the instance-wise feature maps capture richer information compared to the prototype-based approach, resulting in performance improvements.
- In 3.5, we propose a parallel axial-attention (PAA)-based encoder to split the 2-D attention map into two 1-D sub-attention maps, one for height and one for width. Unlike the original axial-attention approach [44], PAA can be simultaneously trained on two GPU devices, enabling parallel calculations of the attention maps and facilitating the training process.
- We demonstrate the effectiveness of our proposed methods by comparing them to state-of-the-art pre-trained models and existing adversarial attack detection methods across a diverse range of images.

2 Related Works

2.1 Self-Supervised Learning

Self-supervised learning aims to develop effective feature representations without the need for large annotated datasets, thereby addressing the annotation bottleneck, which is one of the primary challenges in the practical deployment of deep learning today. Recent studies have shown a growing interest in self-supervised learning, particularly after Yann LeCun’s keynote address at the AAAI 2020

conference [22]. Generally, self-supervised learning can be categorized into three main approaches: predictive [42], generative [25], and contrastive learning [4]. For instance, He et al. proposed masked autoencoders (MAE) to mask random patches of input images and reconstruct the missing pixels [13].

2.2 Contrastive Learning

Contrastive learning aims to develop low-dimensional representations of data by contrasting similar and dissimilar samples [4]. It encourages the learning of feature representations with both inter-class separability and intra-class compactness, which can be highly beneficial for network learning.

One type of contrastive loss function used for self-supervised learning, noise-contrastive estimation (InfoNCE), employs logistic regression to distinguish the target data from noise [33]. Ding et al. use prototypes derived from contextual information to better explore the intrinsic semantics of relations [10]. However, the application of contrastive learning in pixel-level change detection remains a challenging and relatively unexplored area.

2.3 Axial-attention

Wang et al. introduce axial-attention in their work [44], which splits 2D self-attention into two 1D self-attentions. This approach reduces computational complexity and enables attention operations within larger or even global regions. As another variety, Li et al. generate time and frequency sub-attention maps by calculating attention maps along the time and frequency axes of speech spectra [26]. To more efficiently highlight local foreground information, Li et al. propose group parallel axial-attention (GPA) to solve medical image segmentation task [23]. It's worth noting that these studies are primarily conducted in supervised settings, which necessitate labeled data during model training.

3 Self-Supervised Representation Learning

Our proposed solution is built on a progression of novel contributions that combine to build the detection framework. As depicted in Fig 1, the innovations consist of a proposal for pixel mapping (3.2) that facilitates mapping input images into the embedding space through a pair of data transformations. Section 3.3 develops a contrastive estimation approach for adversarial attack detection. As the core idea of our contributions, Section 3.4 provides the discrimination bank, which preserves each instance corresponding to its associated prototype. Finally, we provide an overview of each network component, with a particular focus on our proposed PAA in 3.5.

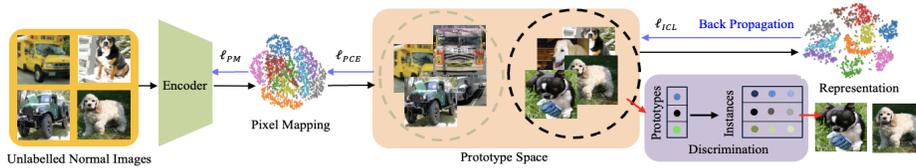


Fig. 1: Self-supervised representation learning framework

3.1 Preliminaries

Given a training set $X = \{x_1, x_2, \dots, x_n\}$ of n image samples, we aim to learn the instance-wise feature representation $Z = \{z_1, z_2, \dots, z_n\}$ of unlabelled normal images for the downstream task, i.e., adversarial attack detection. As a contrastive learning method, InfoNCE [14] achieves this objective by optimizing a contrastive loss function, defined as:

$$\mathcal{L}_{\text{InfoNCE}} = \sum_{i=1}^n -\log \frac{\exp(v_i \cdot v'_i / \tau)}{\sum_{j=0}^r \exp(v_i \cdot v'_j / \tau)} \quad (1)$$

where $V = \{v_1, v_2, \dots, v_n\}$ is the embedding vectors for n instances, and v'_i is the i -th positive embedding. Moreover, v'_j includes one positive embedding and r negative embeddings for other instances, and τ is a dynamic hyper-parameter. In this work, we estimate the prototypes by replacing v'_i in equation (1) with prototype p_i^+ , enabling the proposed discrimination bank establish the connection between the prototype space and each instance.

3.2 Pixel Mapping

As the first major component of the encoder, a PAA-based network with parameter θ is exploited to transform X to feature vectors $V = \{v_1, v_2, \dots, v_I\}$, such that V best describes X . Different from previous work (e.g., InfoNCE loss), we propose a novel pixel mapping loss with data augmentation, \mathcal{L}_{PM} , to learn an invariant representation of x_i by minimizing the risk $\sum_i \mathcal{L}(x_i, v_i; \theta)$. To achieve that, we use a pair of transformations, denoted as t and s , in some set of transformations \mathcal{T} (e.g. geometric transformations, color transformations, etc.) to x_i , to produce the augmentation as $x_i^{t_i}$ and $x_i^{s_i}$. We define this process as $V = f_{\text{PM}}(X)$ with the loss as:

$$\mathcal{L}_{\text{PM}} = -\log \frac{\exp\left(f_{\text{PM}}(x_i^{t_i})^T \cdot f_{\text{PM}}(x_i^{s_i}) / \tau\right)}{\sum_{b=1}^B \exp\left(f_{\text{PM}}(x_b^{t_b})^T \cdot f_{\text{PM}}(x_i^{s_i}) / \tau\right)} \quad (2)$$

where T and B are the transpose symbol and batch size, respectively. It is highlighted that all the embeddings in the loss function are L2-normalized [47]. While previous data augmentation studies [6] have shown that the choice of transformation techniques plays an important part in self-supervised representation

learning, most previous works do not give much consideration to the individual choice of t_i and s_i on pairs of images, which are simply uniformly sampled over \mathcal{T} . Therefore, in the proposed pixel mapping technique, we aim to overcome this limitation and select the optimal transformation algorithm for each sample x_i . To achieve this, we select transformation algorithms that maximize the risk defined by the loss \mathcal{L}^{PM} :

$$\{t_i, s_i\} = \arg \max_{\{t_i, s_i\} \in \mathcal{T}} \sum_{i=1}^n \mathcal{L}_{\text{PM}}(x_i^{t_i}, x_i^{s_i}; \theta, \mathcal{T}) \quad (3)$$

In the proposed pixel mapping technique, we prioritize the difference between t_i and s_i for each image over their absolute values.

3.3 Prototype-wise Contrastive Estimation

We assume that the observed data x_i are related to latent variable $P = \{p_i\}$ which denotes the prototypes of the data. We aim to find a network parameter that maximizes the log-likelihood function of the observed n samples by a prototype-wise contrastive estimation (PCE). To achieve that, we use the local peaks of the density [2] as the prototype, in other words, the most representative data samples of X . The loss, namely \mathcal{L}_{PCE} , is defined as:

$$\mathcal{L}_{\text{PCE}} = \frac{1}{|\mathcal{M}|} \sum_{p_i^+ \in \mathcal{M}} -\log \frac{\exp(v_i \cdot p_i^+ / \gamma)}{\sum_{p_i^- \in \mathcal{N}} \exp(v_i \cdot p_i^- / \gamma)} \quad (4)$$

where \mathcal{M}_i and \mathcal{N}_i are prototype collections of the positive and negative samples, respectively. As aforementioned, inspired from previous supervised learning work [11] [41], we find different levels of concentration distributes around each prototype embeddings. Therefore, we exploit γ as the concentration level around the prototype p^m within the m -th cluster as:

$$\gamma = \frac{\sum_{i=1}^n \|p^m - v_i^m\|_2}{n \log(n + \beta)} \quad (5)$$

where the momentum features are denoted as $\{v_i^m\}_{i=1}^n$ within the same cluster as a prototype p . We set a smooth parameter β to ensure that small clusters do not have an overly-large γ . In the proposed prototype clustering, γ acts as a scaling factor on the similarity between an embedding v and its prototype p . With the proposed γ , the similarity in a loose cluster (larger γ) are down-scaled, pulling embeddings closer to the prototype. On the contrary, embeddings in a tight cluster (smaller γ) have an up-scaled similarity, thus less encouraged to approach the prototype. Therefore, learning with PCE yields more balanced clusters with similar concentration.

3.4 Instance-Wise Contrastive Learning

The core of our method lies in establishing a connection between prototype and instance features to facilitate instance clustering. This approach enables accurate

classification of unseen and attacked samples based on the learned representation. To accomplish this, drawing inspiration from the recent success of memory banks [49], we introduce a discrimination bank for clustering instances that share a common prototype. Initially, we create K independent discrimination banks to enhance instance discrimination across clusters. Similar to a memory bank, the discrimination bank aids in contrastive learning, leveraging extensive data to acquire robust representations. We assume a contrastive set J_i for the t -th bank A_t as:

$$J_i = \{z'_i \mid z'_i \in A_t \forall t \in [1, C]\} \quad (6)$$

where z'_i is the estimated representation of x_i . Specifically, for each training batch with B samples and M prototypes, our discrimination memory is built with size $M \times B \times D$, where D is the dimension of pixel embeddings. The (p^m, b) -th element in the discrimination memory is a D -dimensional feature vector obtained by average pooling all the embeddings of pixels labeled as p^m prototype in the b -th batch. To update the discrimination bank, we enqueue each instance to the nearest prototype and add the new one in each back propagation cycle:

$$\mathcal{L}_{\text{ICL}} = \frac{\exp(\cos(v_i, z_i) \cdot \cos(v_i, p_i^m / \phi))}{\sum_{z' \in A_t} \sum_{j=0}^r \exp(\cos(v_i, z'_j) \cdot \cos(v_i, p_j^m / \phi)) \cdot J_i} \quad (7)$$

where $\cos(\cdot, \cdot)$ is the cosine similarity between a pair of representations. The concentration level of \mathcal{L}_{ICL} is presented as ϕ and estimated similar as γ in (4) but we replace v'_c to z'_c . With the loss, we discriminate representations belongs to the same bank. To discover the underlying concepts with unique visual characteristics, we infer their decision boundaries by reducing the visual redundancy among clusters, namely maximising the visual similarity of samples within the same clusters and minimising that between clusters. Concretely, as the representation of samples with different pseudo labels are stored independently in the discrimination bank, they can be taken as anchors to describe their corresponding clusters. The overall cost-function used to train the MAE is now a combination of the above loss terms with hyper-parameters λ_1 and λ_2 :

$$\mathcal{L} = \mathcal{L}_{\text{PM}} + \lambda_1 \cdot \mathcal{L}_{\text{PCE}} + \lambda_2 \cdot \mathcal{L}_{\text{ICL}} \quad (8)$$

3.5 Parallel Axial-attention-Based Encoder

Our encoder has four major components for each learning objective:

1. For \mathcal{L}_{PM} in Eq. (2), we set the temperature τ as 0.1. To transform X into embeddings V , we use a ResNet-50 [15] as the backbone. However, different from conventional ResNet families, we transform it to a parallel axial-attention (PAA)-ResNet by replacing the convolutional layer in the residual bottleneck block by two parallel multi-head axial-attention layers (one for height-axis and the other for width-axis). Additionally, two Conv1D layers are included to shuffle the features, as illustrated in Fig. 2.

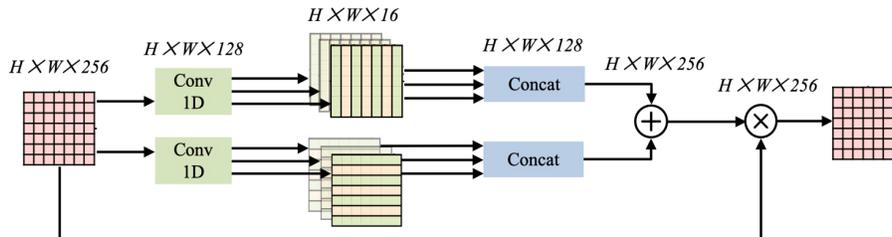


Fig. 2: Parallel axial-attention (PAA) block.

2. Prototype-wise contrastive estimation (PCE) is achieved by two Conv1D layers with ReLU. PCE is only applied during training and is removed at inference time. Thus it does not introduce any changes to the segmentation network or extra computational cost in deployment.

3. Discrimination bank consists of two parts that store prototype and instance embeddings, respectively. For each prototype, we set the maximum size of the instance queue as 10. The discrimination bank is discarded after training.

4. Instance-wise contrastive learning is achieved by two Conv1D layers with ReLU and faiss [18] for efficient instance clustering. Similar as PCE, ICE is removed in the test stage.

4 Experiments

4.1 Datasets and Attacks

We randomly select 50,000 images from ImageNet [9] and 10,000 images from ImageNet Large Scale Visual Recognition Challenge (ILSVRC) [40] for the training and validation, respectively. As aforementioned, we evaluate the competitor and proposed models with unseen datasets. In the test stage, we extensively perform experiments on several public datasets, including ImageNet-R [16], Canadian Institute For Advanced Research-10 (CIFAR-10) [20], and Microsoft Common Objects in Context (COCO) [27]. In each above-mentioned dataset, we randomly select 10,000 images to evaluate the detection performance of models.

We select seven attack algorithms [30] [12] [32] [21] [3] [35] [29] in the test stage because they are robust to novel adversarial attack detection and defense techniques. Parameters of all the seven attacks are shown in Table 1.

4.2 Backbones

As aforementioned we use ResNet-50 as the encoder’s backbone without pre-training. Moreover, various backbones (i.e., HRNet [45] and Xception [7]) in supplementary experiments to validate the proposed algorithm and compare to ResNet-50.

Table 1: Parameters of seven adversarial attacks

| Attack | Parameters |
|----------|--|
| FGSM | $\epsilon=0.008$ |
| PGD | $\epsilon=0.01, \alpha=0.02, \text{Steps}=40$ |
| SSAH | $\alpha=0.01$ |
| DeepFool | $\text{Steps}=20$ |
| BIM | $\epsilon=0.03, \alpha=0.01, \text{Steps}=10$ |
| CW | $C, \text{Kappa}=2, \text{Steps}=500, \text{learning rate}=0.01$ |
| JSMA | $\gamma=0.02$ |

4.3 Implementation Details

In the experiment, we build the PAA-ResNet-S from ResNet-50 [15] whose last fully-connected layer outputs a 128-D and L2-normalized feature. As aforementioned, to achieve that, we replace the convolutional layer in the residual bottleneck block by two parallel multi-head axial-attention layers (one for height-axis and the other for width-axis). We multiply all the channels by 1.5 and 2, resulting in PAA-ResNet-M, L, respectively. We always use 8 heads in multi-head attention blocks [43]. In order to avoid careful initialization of weights (W_Q, W_K, W_V) and location vectors (r^q, r^k, r^v), we use batch normalizations [17] in all attention layers. To evaluate and compare the adversarial attack detection accuracy, we use the detection rate (DR).

The proposed model is trained by using the SGD optimizer with a weight decay of 0.0001, a momentum of 0.9, and a batch size of 256. We train the networks for 200 epochs, where we warm-up the network in the first 20 epochs by only using the pixel-mapping loss. The initial learning rate is 0.03, and is multiplied by 0.1 at 120 and 160 epochs. In terms of the hyper-parameters, we set $\tau = 0.1, \beta = 10, r = 16000, \lambda_1 = 1$ and $\lambda_2 = 1$. All the experiments are run on the High End Computing (HEC) Cluster with Tesla V100 GPUs.

5 Results

5.1 ImageNet

The learned representation is evaluated on adversarial attack detection task over the ImageNet-R dataset. Table 2 shows the results, each of them is the average of 70,000 experiments (10,000 images \times 7 attacks).

Table 2 shows the averaged adversarial attack detection performance of the proposed method as compared to [41] [5] [38] [19] [51] [13] [50] [1] [34] [8] [31] using the ImageNet-R dataset. From Table 2, it can be observed that: (1) In all the evaluated models, the proposed PAA-ResNet-L achieves 93.8% and 86.8% for clean and attacked images detection, respectively, which offers the best effectiveness. (2) Compared to state-of-the-art models, the proposed PAA-ResNet family requires the least computational cost and inference latency in the training stage because the model utilizes a parallel computation, which makes it feasible in real-world applications.

Table 2: Comparison on the ImageNet-R dataset.

| Method | Configuration | | Computational Cost | | DR (%) | |
|---------------------|-------------------------|---|--------------------|--------------|-------------|-------------|
| | Supervised Pre-training | | Para. | Latency (ms) | Clean | Attacked |
| TiCo [51] | ✗ | ✓ | 178.0 M | 160.3 | 79.5 | 65.9 |
| MAE [13] | ✗ | ✓ | 307.8 M | 104.5 | 88.4 | 72.0 |
| Mugs [50] | ✗ | ✓ | 303.5 M | 185.4 | 89.2 | 72.1 |
| Unicom [1] | ✗ | ✓ | 303.5 M | 268.9 | 91.4 | 80.5 |
| DINOv2 [34] | ✗ | ✗ | 1.0 B | 572.0 | 92.5 | 82.8 |
| ESMAF [5] | ✓ | ✗ | 171.0 M | 196.2 | 69.8 | 53.7 |
| TS [19] | ✓ | ✗ | 36.5 M | 45.5 | 78.0 | 57.6 |
| sim-DNN [41] | ✓ | ✓ | 663.9 M | 227.1 | 79.2 | 60.3 |
| DTBA [38] | ✓ | ✓ | 554.2 M | 208.4 | 84.2 | 62.3 |
| TLC [8] | ✗ | ✓ | 27.5 M | 6.9 | 83.5 | 71.7 |
| SimCat [31] | ✗ | ✓ | 27.8 M | 7.2 | 85.1 | 72.9 |
| <i>PAA-ResNet-S</i> | ✗ | ✗ | 26.8 M | 6.2 | 92.0 | 83.1 |
| <i>PAA-ResNet-M</i> | ✗ | ✗ | 35.4 M | 8.7 | 93.5 | 85.2 |
| <i>PAA-ResNet-L</i> | ✗ | ✗ | 45.9 M | 10.1 | 93.8 | 86.8 |

5.2 COCO & CIFAR-10

We assess the learned representation over CIFAR-10 and COCO. Tables 3 & 4 show the results, each of them is the average of 70,000 experiments (10,000 images×7 attacks).

Table 3: Comparison on CIFAR-10.

| Models | Clean (%) | Attacked (%) |
|---------------------|-------------|--------------|
| TiCo [51] | 81.4 | 78.0 |
| MAE [13] | 89.9 | 74.2 |
| Mugs [50] | 90.5 | 73.7 |
| Unicom [1] | 92.6 | 84.1 |
| DINOv2 [34] | 94.3 | 86.7 |
| ESMAF [5] | 73.8 | 56.4 |
| TS [19] | 89.7 | 59.5 |
| sim-DNN [41] | 82.0 | 65.7 |
| DTBA [38] | 87.0 | 74.1 |
| TLC [8] | 84.9 | 72.4 |
| SimCat [31] | 88.0 | 77.3 |
| <i>PAA-ResNet-S</i> | 92.7 | 84.4 |
| <i>PAA-ResNet-M</i> | 94.1 | 87.8 |
| <i>PAA-ResNet-L</i> | 94.8 | 89.0 |

Table 4: Comparison on COCO.

| Models | Clean (%) | Attacked (%) |
|---------------------|-------------|--------------|
| TiCo [51] | 78.9 | 67.3 |
| MAE [13] | 88.9 | 73.5 |
| Mugs [50] | 89.0 | 73.3 |
| Unicom [1] | 90.2 | 82.8 |
| DINOv2 [34] | 91.7 | 83.9 |
| ESMAF [5] | 75.4 | 55.6 |
| TS [19] | 76.7 | 56.8 |
| sim-DNN [41] | 80.6 | 62.2 |
| DTBA [38] | 85.3 | 68.8 |
| TLC [8] | 80.8 | 71.5 |
| SimCat [31] | 82.6 | 70.1 |
| <i>PAA-ResNet-S</i> | 90.9 | 83.7 |
| <i>PAA-ResNet-M</i> | 91.5 | 84.9 |
| <i>PAA-ResNet-L</i> | 91.7 | 85.6 |

On both datasets, our models show strong detection performance: accuracy improves considerably with the proposed algorithm. Additionally, our results

outperforms both the self-supervised and supervised results by large margins on clean images detection.

5.3 Diagnostic Experiment

In this section, we conduct the ablation study to evaluate the effectiveness of each contribution.

Contrastive Learning In this section, we compare the effectiveness of each proposed contrastive learning loss to previous contrastive learning techniques and conduct the ablation study on the ImageNet-R dataset. Moreover, we compare our contrastive losses to similar losses in [33] [10]. Each result is the average of 70,000 experiments (10,000 images \times 7 attacks).

Table 5: Ablation study of the three contributions in the proposed method.

| Ablation Settings | Clean (%) | Attacked (%) |
|---|-------------|--------------|
| PAA-ResNet-L | 60.1 | 52.3 |
| + \mathcal{L}_{NCE} [33] | 72.0 | 63.5 |
| +Data Augmentation [6] | 65.6 | 50.1 |
| + \mathcal{L}_{PM} | 76.7 | 68.9 |
| + \mathcal{L}_{S2Z} [10] | 79.7 | 68.6 |
| + \mathcal{L}_{PCE} | 82.5 | 72.4 |
| + \mathcal{L}_{PM} + \mathcal{L}_{PCE} | 90.4 | 82.8 |
| + \mathcal{L}_{PCE} + \mathcal{L}_{ICL} | 89.5 | 82.0 |
| + \mathcal{L}_{PM} + \mathcal{L}_{PCE} + \mathcal{L}_{ICL} | 93.8 | 86.8 |

We first perform experiments to validate the design of our pixel mapping with a PAA-ResNet-L as the baseline. As shown in Table 3, additionally considering data augmentation leads to a substantial performance gain (i.e., 5.4 %), compared with \mathcal{L}_{NCE} .

We next investigate the effectiveness of \mathcal{L}_{PCE} . On the one hand, \mathcal{L}_{PCE} boosts the performance based on \mathcal{L}_{PM} (i.e., 68.9% \rightarrow 82.8%). On the other hand, when we replace \mathcal{L}_{PCE} to a prototype-based loss \mathcal{L}_{S2Z} [10], the proposed \mathcal{L}_{PM} + \mathcal{L}_{PCE} achieves a higher score (i.e., 82.8%) is achieved, which confirms the efficiency of the proposed loss terms.

We then presents a comprehensive examination of \mathcal{L}_{ICL} . Compared with \mathcal{L}_{PM} + \mathcal{L}_{PCE} , \mathcal{L}_{ICL} brings a substantial performance gain (i.e., 4.4 %). Moreover, the performance slightly drops (i.e., 82.8% \rightarrow 82.0%) when \mathcal{L}_{ICL} replace \mathcal{L}_{PM} within combining \mathcal{L}_{PCE} .

Parallel Axial-attention We conducted experiments to demonstrate the trade-off between performance improvement and network depth, specifically, varying the number of proposed PAA blocks. It’s important to note that each PAA block consists of two parallel sub-blocks, i.e., height and width attentions. Furthermore, we compared the performance of different backbones, such as ResNet-50,

HRNetV2-W48 [45], and Xception-71 [7]. Additionally, we assessed the performance improvements resulting from the inclusion of axial-attention [43] and PAA. Fig. 3 (a) and (b) present these results, with each data point being an average of 70,000 experiments (10,000 images of ImageNet-R \times 7 attacks).

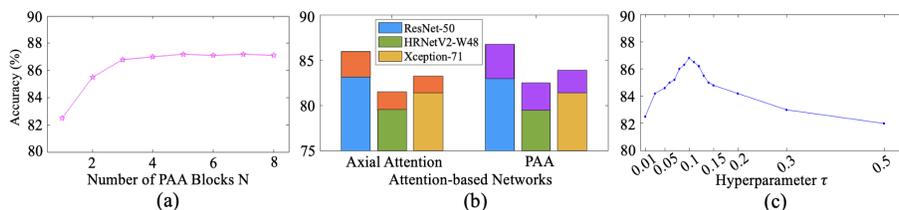


Fig. 3: Ablation study for (a) number of parallel axial-attention (PAA) blocks and (b) attention blocks to different backbones (c) hyperparameter τ . The red and purple rectangles represent accuracy improvements when axial attention and PAA are added to the backbones, respectively.

Fig. 3(a) compares the number of PAA blocks against detection accuracy on ImageNet-R. The results indicate that $N=3$ offers the best trade-off, validating the chosen implementation setting. Fig. 3(b) presents the experiment results of networks with the axial-attention and PAA. It can be observed that: (1) Among the three backbones without added attention blocks, ResNet-50 outperforms HRNetV2-W48 and Xception-71. (2) Both axial-attention and PAA enhance detection performance based on three backbones. Particularly, the proposed PAA provides a greater performance boost compared to the axial-attention block, with an improvement from 2.1% to 2.9% on average across the three backbones.

Hyper-parameter An ablation study of hyper-parameters is conducted using the PAA-ResNet-L on the ImageNet-R dataset. Fig. 3(c) shows downstream attack detection accuracy when τ varies from 0 to 2. Besides, β impact is evaluated in Fig. 4(a). More ablation studies are performed to evaluate the loss term hyper-parameters λ_1 and λ_2 in Fig. 4(b). Each data point being an average of 70,000 experiments (10,000 images \times 7 attacks).

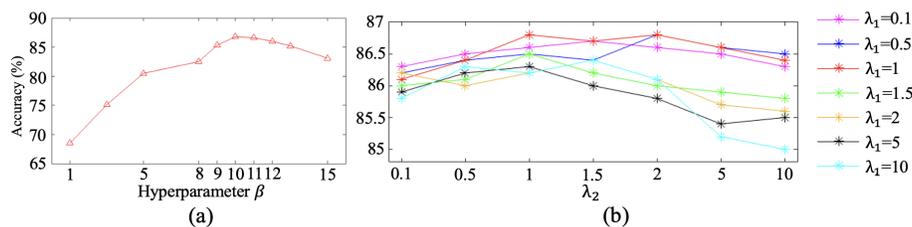


Fig. 4: Ablation study for (a) hyper-parameter β and (b) λ_1 & λ_2 .

As Fig. 3(c) shows, detection accuracy starts to increase with $\tau = 0.01$ and reaches its peak around $\tau = 0.1$, but performance is fairly stable for $0.03 \leq \tau \leq 0.2$. Additionally, Fig. 4(a) suggests the hyper-parameter $\beta = 10$. Fig. 4(b) presents detection accuracy against λ_1 and λ_2 . There is no significant accuracy drop even when the importance of loss terms is significantly weighted, such as by as much as tenfold that of \mathcal{L}_{PM} ($\lambda_1, \lambda_2 = 10$). This demonstrates that the features derived from each loss terms contribute positively to the learning process.

5.4 Robustness Evaluation

We perform experiments to evaluate the robustness of the supervised and self-supervised models on adversarial attack detection. For fair comparison, we fine-tune pre-trained self-supervised models [1] [13] [34] [50] using the same training data as the supervised models on the corresponding dataset. Table 6 shows the detection accuracy results (in %), each of them is the average of 70,000 experiments (10,000 images \times 7 attacks). Apart from aforementioned datasets, we introduce one more dataset, i.e., Canadian Institute For Advanced Research-100 (CIFAR-100) [20].

Table 6: Adversarial attack detection performance (clean / attacked images) on seen and unseen datasets.

| Training | ImageNet-R | | ILSVRC | | CIFAR-100 | |
|---------------------|--------------------|--------------------|--------------------|--------------------|--------------------|--------------------|
| Test | ImageNet-R | CIFAR-10 | ILSVRC | CIFAR-100 | CIFAR-100 | ImageNet-R |
| MAE [50] | 89.4 / 74.1 | 89.0 / 73.4 | 91.0 / 79.2 | 89.1 / 73.0 | 90.4 / 74.3 | 85.4 / 70.5 |
| Mugs [50] | 89.8 / 74.0 | 89.1 / 73.6 | 91.8 / 78.1 | 89.4 / 74.5 | 90.9 / 75.0 | 86.2 / 71.1 |
| Unicom [1] | 91.9 / 82.7 | 91.0 / 80.4 | 94.7 / 88.5 | 92.0 / 81.1 | 93.3 / 82.7 | 89.3 / 77.9 |
| DINOv2 [34] | 93.4 / 84.5 | 92.4 / 81.7 | 96.2 / 90.0 | 93.4 / 82.6 | 95.1 / 84.0 | 90.5 / 79.4 |
| ESMAF [5] | 88.1 / 72.3 | 75.5 / 59.7 | 90.1 / 78.7 | 77.8 / 62.6 | 89.5 / 72.6 | 74.2 / 59.8 |
| TS [19] | 90.2 / 75.8 | 79.6 / 66.2 | 92.6 / 79.8 | 82.3 / 68.1 | 91.2 / 79.8 | 83.7 / 66.6 |
| sim-DNN [41] | 90.1 / 77.4 | 81.5 / 70.6 | 93.4 / 82.5 | 81.9 / 71.2 | 92.8 / 82.4 | 83.6 / 65.9 |
| DTBA [38] | 92.2 / 85.2 | 85.3 / 76.9 | 96.0 / 90.3 | 86.8 / 78.2 | 94.7 / 83.1 | 88.2 / 69.9 |
| <i>PAA-ResNet-L</i> | 93.5 / 87.9 | 92.9 / 85.7 | 97.1 / 90.5 | 94.2 / 87.0 | 96.0 / 87.6 | 92.1 / 83.4 |

It can be observed that the adversarial attack detection accuracy on unseen datasets is lower compared to the seen dataset utilized in both the training and test stages, primarily due to differences in distributions of datasets. However, when compared to supervised learning-based methods [5] [19] [38] [41], the proposed SSL representation learning method experiences relatively less performance degradation.

5.5 Visualization

As qualitative analysis, Fig. 5 presents the t-distributed stochastic neighbour embedding (t-SNE) visualisation of PAA-ResNet-L trained with different losses. Compared to the representation learned by \mathcal{L}_{PM} , the representation learned by

\mathcal{L}_{ICL} forms more separated clusters, which also suggests representation of lower entropy. In Fig. 5(b), it can be observed that the feature embeddings within a single prototype are not separable. However, when the discrimination bank is added in Fig. 5(c), individual instances become separated. This demonstrates that the proposed methods can learn discriminative feature representations that generalize well for adversarial attack detection across various attack algorithms.

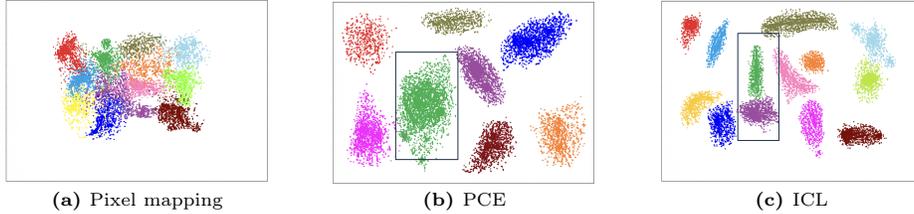


Fig. 5: t-SNE feature visualization of the model with (a) \mathcal{L}_{PM} ; (b) $\mathcal{L}_{PM} + \mathcal{L}_{PCE}$; (c) $\mathcal{L}_{PM} + \mathcal{L}_{PCE} + \mathcal{L}_{ICL}$.

6 Discussion and Conclusion

We enumerate these advantages of this work below:

1. In the training stage, we only require access to unlabelled normal image samples. Therefore, unlike supervised adversarial attack detection methods, we do not need paired training data, i.e., adversarial and normal images. Moreover, unlike supervised learning models, the proposed model extracts the underlying sub-class structure from the augmented data distribution and encodes it into the embeddings. This guarantees the robustness of the downstream task.

2. The proposed loss \mathcal{L}_{PM} considers the difference between two transformation techniques for each image x_i more than their absolute values. These transformed images create more challenging positive pairs and \mathcal{L}_{PM} produces harder negative pairs by accounting for all images in a batch during the optimization, enabling the learning of richer features.

Adversarial attacks result in a deviation of the constructed graphs of adversarial examples from the prototypes of their correct classes, creating challenges for accurate classification based on representations. Despite this, the distributions between the reconstructed image and the prototypes of target classes remain significantly different. The proposed \mathcal{L}_{PCL} correctly encourages representations to be closer to their assigned prototypes because the encoder would learn the shared information among prototypes, and ignore the individual noise that exists in each prototype. The shared information is more likely to capture higher-level feature knowledge.

By using the proposed discrimination bank, the encoder preserves the connection between the instance and the associated prototype. This enables the

learning of richer features from samples of a single class. Therefore, the proposed method outperforms other contrastive learning methods in the literature.

3. PAA blocks capture richer feature information, enabling models to seamlessly integrate local and global feature representations, thereby enhancing their ability to provide local-global feature information for downstream tasks. Additionally, PAA blocks simplify computational complexity by allowing height and width attentions to be simultaneously trained on parallel GPU devices to facilitate the training. Furthermore, PAA blocks can be easily implemented on various backbones and consistently demonstrate promising results.

While the concept of parallel axial-attention structure has been introduced in previous work [23], it’s important to note that the algorithm in our work is fundamentally different from the literature. In [23], the feature maps are reshaped to $(H \times C) \times W$ and $(W \times C) \times H$ in two parallel attentions blocks, respectively. However, in this work, we only focus on one dimension, either width or height, in each sub-PAA block. Therefore, the computational cost is further reduced.

In this paper, we have proposed a self-supervised representation learning approach for adversarial attack detection, offering an effective alternative to traditional supervised pipelines. We establish a connection between prototype and instance features through the use of a discrimination bank, thereby enriching the information available to enhance the proposed model’s ability to detect adversarial attacks. Our evaluation with different datasets and attacks has demonstrated the robust performance of the proposed method on unseen datasets. Additionally, PAA-ResNet models offer faster inference speeds compared to competitive pre-trained models, making them feasible for potential real-world applications.

Acknowledgment

This work is supported by the UKRI Trustworthy Autonomous Systems Node in Security/EP SRC Grant EP/V026763/1.

This work is also supported by ELSA – European Lighthouse on Secure and Safe AI funded by the European Union under grant agreement No. 101070617. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Commission. Neither the European Union nor the European Commission can be held responsible.

References

1. An, X., Deng, J., Yang, K., Li, J., Feng, Z., Guo, J., Yang, J., Liu, T.: Unicom: universal and compact representation learning for image retrieval. *Proceedings of International Conference on Learning Representations (ICLR) (2023)* [8](#), [9](#), [12](#)
2. Angelov, P., Soares, E.: Towards explainable deep neural networks (xDNN). *Neural Networks* **130**, 185–194 (2020) [5](#)
3. Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. *IEEE Symposium on Security and Privacy (2017)* [7](#)
4. Chang, J., Zhang, J., Xu, Y., Li, J., Ma, S., Gao, W.: Consistency-contrast learning for conceptual coding. *Proceedings of the 30th ACM International Conference on Multimedia (2022)* [3](#)
5. Chen, J., Yu, T., Wu, C., Zheng, H., Zhao, W., Pang, L., Li, H.: Adversarial attack detection based on example semantics and model activation features. *Proceedings of International Conference on Data Science and Information Technology (DSIT) (2022)* [8](#), [9](#), [12](#)
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. *Proceedings of International Conference on Machine Learning (ICML) (2020)* [4](#), [10](#)
7. Chollet, F.: Xception: deep learning with depthwise separable convolutions. *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2017)* [7](#), [11](#)
8. Chyou, C.C., Su, H.T., Hsu, W.H.: Unsupervised adversarial detection without extra model: training loss should change. *Proceedings of International Conference on Machine Learning (ICML) (2023)* [8](#), [9](#)
9. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2009)* [7](#)
10. Ding, N., Wang, X., Fu, Y., Xu, G., Wang, R., Xie, P., Shen, Y., Huang, F., Zheng, H.T., Zhang, R.: Prototypical representation learning for relation extraction. *Proceedings of International Conference on Learning Representations (ICLR) (2021)* [3](#), [10](#)
11. Gong, Y., Wang, S., Jiang, X., Yin, L., Sun, F.: Adversarial example detection using semantic graph matching. *Applied Soft Computing* **141**, 110317 (2023) [5](#)
12. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. *Proceedings of International Conference on Learning Representations (ICLR) (2015)* [7](#)
13. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)* [3](#), [8](#), [9](#), [12](#)
14. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)* [4](#)
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2016)* [6](#), [8](#)
16. Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., Song, D., Steinhardt, J., Gilmer, J.: The many faces of robustness: a critical analysis of out-of-distribution generalization. *Proceedings of IEEE/CVF International Conference on Computer Vision (ICCV) (2021)* [7](#)

17. Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. *Proceedings of International Conference on Machine Learning (ICML)* (2015) [8](#)
18. Johnson, J., Douze, M., Jegou, H.: Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data* **7**, 4249–4260 (2021) [7](#)
19. Kiani, S., Awan, S., Lan, C., F. Li, B.L.: Two souls in an adversarial image: towards universal adversarial example detection using multi-view inconsistency. *Asia-Pacific Computer Systems Architecture Conference (APCSAC)* (2021) [1](#), [8](#), [9](#), [12](#)
20. Krizhevsky, A.: Learning multiple layers of features from tiny images. Master’s thesis (2009) [7](#), [12](#)
21. Kurakin, A., Goodfellow, I., Bengio, S.: Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533* (2016) [7](#)
22. LeCun, Y.: Self-supervised learning. *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI)* (2017) [3](#)
23. Li, C., Wang, L., Li, Y.: Transformer and group parallel axial attention co-encoder for medical image segmentation. *Scientific Reports* **12**, 16117 (2022) [3](#), [14](#)
24. Li, Y., Angelov, P., Suri, N.: Domain generalization and feature fusion for cross-domain imperceptible adversarial attack detection. *Proceedings of the International Joint Conference on Neural Networks (IJCNN)* (2023) [1](#)
25. Li, Y., Angelov, P., Suri, N.: Rethinking self-supervised learning for cross-domain adversarial sample recovery. *Proceedings of the International Joint Conference on Neural Networks (IJCNN)* (2024) [3](#)
26. Li, Y., Sun, Y., Wang, W., Naqvi, S.M.: U-shaped Transformer with frequency-band aware attention for speech enhancement. *IEEE/ACM Transactions on Audio, Speech and Language Processing* **31**, 1511–1521 (2023) [3](#)
27. Lin, T.Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C.L., Dollár, P.: Microsoft COCO: common objects in context. *Proceedings of European Conference on Computer Vision (ECCV)* (2014) [7](#)
28. Liu, J., Levine, A., Lau, C.P., Chellappa, R., Feizi, S.: Segment and complete: defending object detectors against adversarial patch attacks with robust patch detection. *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022) [1](#)
29. Luo, C., Lin, Q., Xie, W., Wu, B., Xie, J., Shen, L.: Frequency-driven imperceptible adversarial attack on semantic similarity. *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022) [7](#)
30. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. *Proceedings of International Conference on Machine Learning (ICML)* (2017) [7](#)
31. Moayeri, M., Feizi, S.: Sample efficient detection and classification of adversarial attacks via self-supervised embeddings. *Proceedings of IEEE/CVF International Conference on Computer Vision (ICCV)* (2021) [8](#), [9](#)
32. Moosavi-Dezfooli, S.M., Fawzi, A., Frossard, P.: Deepfool: a simple and accurate method to fool deep neural networks. *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2016) [7](#)
33. Oord, A., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. *arXiv preprint arXiv: 1807.03748* (2018) [3](#), [10](#)
34. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Assran, M., Ballas, N., Galuba, W., Howes, R., Huang, P.Y., Li, S.W., Misra, I., Rabbat, M., Sharma, V., Synnaeve, G., Xu, H., Jegou, H., Mairal, J., Labatut, P., Joulin, A., Bojanowski, P.:

- Dinov2: learning robust visual features without supervision. arXiv preprint arXiv:2304.07193 (2023) [8](#), [9](#), [12](#)
35. Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, Z.B., Swami, A.: The Limitations of Deep Learning in Adversarial Settings. *IEEE Symposium on Security and Privacy* (2016) [7](#)
 36. Pellicier, A.L., Giatgong, K., Li, Y., Suri, N., Angelov, P.: UNICAD: A unified approach for attack detection, noise reduction and novel class identification. *Proceedings of the International Joint Conference on Neural Networks (IJCNN)* (2024) [2](#)
 37. Pellicier, A.L., Li, Y., Angelov, P.: PUDD: Towards Robust Multi-modal Prototype-based Deepfake Detection. *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2024) [2](#)
 38. Qi, P., Jiang, T., Wang, L., Yuan, X., Li, Z.: Detection tolerant black-box adversarial attack against automatic modulation classification with deep learning. *IEEE Transactions on Reliability* **71**(2), 674–686 (2022) [8](#), [9](#), [12](#)
 39. Raina, V., Gales, M.: Residue-based natural language adversarial attack detection. *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)* (2022) [1](#)
 40. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Li, F.: Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)* (2015) [7](#)
 41. Soares, E., Angelov, P., Suri, N.: Similarity-based deep neural network to detect imperceptible adversarial attacks. *Proceedings of IEEE Symposium Series on Computational Intelligence (SSCI)* (2022) [5](#), [8](#), [9](#), [12](#)
 42. Tsai, Y.H., Ma, M.Q., Yang, M., Zhao, H., Morency, L.P., Salakhutdinov, R.: Self-supervised representation learning with relative predictive coding. *Proceedings of International Conference on Learning Representations (ICLR)* (2021) [3](#)
 43. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *International Conference on Neural Information Processing Systems (NeurIPS)* (2017) [8](#), [11](#)
 44. Wang, H., Zhu, Y., Green, B., Adam, H., Yuille, A., Chen, L.C.: Axial-DeepLab: stand-alone axial-attention for panoptic segmentation. *Proceedings of European Conference on Computer Vision (ECCV)* (2020) [2](#), [3](#)
 45. Wang, J., Sun, K., Cheng, T., Jiang, B., Deng, C., Zhao, Y., Liu, D., Mu, Y., Tan, M., Wang, X., Liu, W., Xiao, B.: Deep high-resolution representation learning for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **43**, 3349–3364 (2021) [7](#), [11](#)
 46. Wang, X., Li, S., Liu, M., Wang, Y., Roy-Chowdhury, A.: Multi-expert adversarial attack detection in person re-identification using context inconsistency. *Proceedings of IEEE/CVF International Conference on Computer Vision (ICCV)* (2021) [1](#)
 47. Yang, M., Meng, Z., King, I.: FeatureNorm: L2 feature normalization for dynamic graph embedding. *Proceedings of IEEE International Conference on Data Mining (ICDM)* (2020) [4](#)
 48. Yang, Y., Yang, S., Xie, J., Si, Z., Guo, K., Zhang, K., Liang, K.: Multi-head uncertainty inference for adversarial attack detection. *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2023) [1](#)
 49. Yokoo, S.: Contrastive learning with large memory bank and negative embedding subtraction for accurate copy detection. arXiv preprint arXiv:2112.04323 (2021) [6](#)

50. Zhou, P., Zhou, Y., Si, C., Yu, W., Ng, T.K., Yan, S.: Mugs: a multi-granular self-supervised learning framework. arXiv preprint arXiv: 2203.14415 (2022) [8](#), [9](#), [12](#)
51. Zhu, J., Moraes, R., Karakulak, S., Sobol, V., Canziani, A., LeCun, Y.: Tico: transformation invariance and covariance contrast for self-supervised visual representation learning. arXiv preprint arXiv: 2203.14415 (2022) [8](#), [9](#)