
MetaMask: Revisiting Dimensional Confounder for Self-Supervised Learning

Jiangmeng Li*, Wenwen Qiang*, Yanan Zhang

University of Chinese Academy of Sciences
Institute of Software Chinese Academy of Sciences
Southern Marine Science and Engineering Guangdong Laboratory (Guangzhou)
{jiangmeng2019, wenwen2018, yanan2018}@iscas.ac.cn

Wenyi Mo

Gaoling School of Artificial Intelligence
Renmin University of China
2022101010@ruc.edu.cn

Changwen Zheng

Institute of Software Chinese Academy of Sciences
Southern Marine Science and Engineering Guangdong Laboratory (Guangzhou)
changwen@iscas.ac.cn

Bing Su[†]

Gaoling School of Artificial Intelligence
Renmin University of China
Beijing Key Laboratory of Big Data Management and Analysis Methods
subingats@gmail.com

Hui Xiong

Thrust of Artificial Intelligence
The Hong Kong University of Science and Technology (Guangzhou)
Guangzhou HKUST Fok Ying Tung Research Institute
xionghui@ust.hk

Abstract

As a successful approach to self-supervised learning, contrastive learning aims to learn invariant information shared among distortions of the input sample. While contrastive learning has yielded continuous advancements in sampling strategy and architecture design, it still remains two persistent defects: the interference of task-irrelevant information and sample inefficiency, which are related to the recurring existence of trivial constant solutions. From the perspective of dimensional analysis, we find out that the *dimensional redundancy* and *dimensional confounder* are the intrinsic issues behind the phenomena, and provide experimental evidence to support our viewpoint. We further propose a simple yet effective approach *MetaMask*, short for the *dimensional Mask learned by Meta-learning*, to learn representations against dimensional redundancy and confounder. MetaMask adopts the redundancy-reduction technique to tackle the dimensional redundancy issue

*Equal contributions.

[†]Corresponding author.

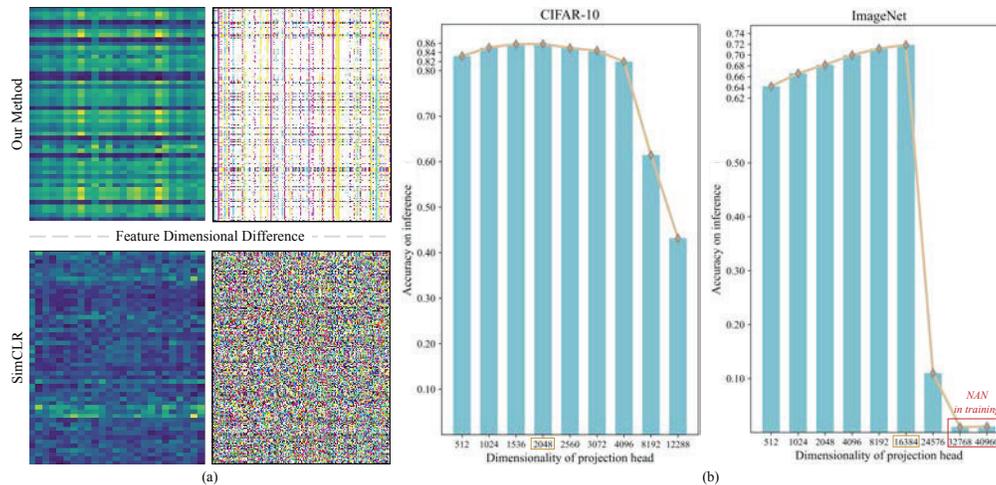


Figure 1: (a) The visualization of the representations learned by SimCLR [1] and our method using the redundancy-reduction technique [2] on CIFAR-10, respectively. The learned features are projected into a color image in RGB format, where different colors represent different types of features. The abscissa axis represents the feature dimensions, and the ordinate axis represents samples of different classes. The greater the color contrast, the lower the dimensional feature similarity. The left plots present the contribution of different dimensions to a specific category classification, and the right plots present the similarity between dimension features within a batch. Compared with SimCLR, our method using redundancy-reduction can indeed learn representations with decoupled dimensions. (b) The experimental results obtained by Barlow Twins on CIFAR-10 and ImageNet datasets. We intuitively increase the dimensionality of the projection head, which leads the dimensions of representations to model more different information. Note that we adopt the official code of Barlow Twins and train on 8 GPUs of NVIDIA Tesla V100.

and innovatively introduces a dimensional mask to reduce the gradient effects of specific dimensions containing the confounder, which is trained by employing a meta-learning paradigm with the objective of improving the performance of masked representations on a typical self-supervised task. We provide solid theoretical analyses to prove MetaMask can obtain tighter risk bounds for downstream classification compared to typical contrastive methods. Empirically, our method achieves state-of-the-art performance on various benchmarks.¹

1 Introduction

A fundamental idea behind self-supervised learning is to learn discriminative representations from the input data without relying on human annotations. Recent advances in visual self-supervised learning [3, 4, 5, 6, 7] demonstrate that unsupervised approaches can achieve competitive performance over supervised approaches by introducing sophisticated self-supervised tasks. A representative learning paradigm is contrastive learning [8, 1, 9, 10, 11, 12, 13], which aims to learn invariant information from different views (generated by data augmentations) by performing instance-level contrast, i.e., pulling views of the same sample together while pushing views of different samples away. Various [14, 15, 16, 17] techniques have been explored to address the problem of trivial solutions to self-supervised contrastive learning, e.g., constant representations. However, state-of-the-art methods still suffer from two crucial defects, including the interference of task-irrelevant information and sample inefficiency. The trivial solution problem is related to such defects, e.g., a small number of negative samples or a large proportion of task-irrelevant information may lead the model to learn constant representations. We revisit the current self-supervised learning paradigm from the perspective of dimensional analysis and argue that the intrinsic issues behind the defects of self-supervised learning are the *dimensional redundancy* and *dimensional confounder*².

¹The implementation is available at <https://github.com/lionellee9089/MetaMask>

²Throughout this paper, the term *confounder* is used in its idiomatic sense rather than the specific statistical sense in Structural Causal Models [18].

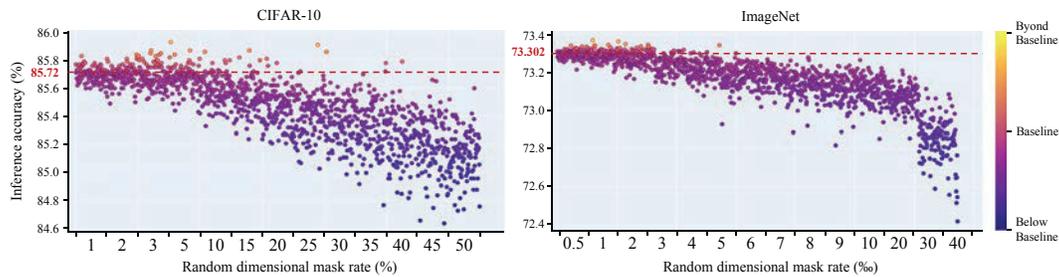


Figure 2: Experimental scatter diagrams obtained by Barlow Twins [2] with randomly masked dimensions on CIFAR-10 and ImageNet datasets, where *Baseline* and the *red dashed lines* denote the performance achieved by the unmasked representation of Barlow Twins. Every single point denotes an independent experimental result achieved by randomly masking the original representation at the dimensional level with a specific mask rate. Note that the original representation is consistently fixed. Following the experimental protocol of [2, 16], we evaluate the performance on CIFAR-10 and ImageNet by adopting KNN prediction and linear probing, respectively, which further proves the general existence of dimensional confounders.

From the perspective of information theory, each dimension holds a subset of the representation's information entropy. The dimensional redundancy denotes that multiple dimensions hold overlapping information entropy. To tackle this issue, Barlow Twins [2] takes advantage of a neuroscience approach, i.e., redundancy-reduction [19], in self-supervised learning. It makes the cross-correlation matrix of twin embeddings close to the identity matrix, which is conceptually simple yet effective to address the dimensional redundancy issue. Figure 1 (a) supports the superiority of redundancy-reduction in addressing the dimensional redundancy problem. Since Barlow Twins encourages the dimensions of learned representations to model decoupled information, it naturally avoids the collapsed trivial solution that outputs the same vector for all images.

However, only reducing dimensional redundancy still suffers from the problem of dimensional confounder. The dimensional confounder indicates a set of dimensions that contains "harmful" information and further degenerates the performance of the model, e.g., the dimensions simply capturing chaotic and worthless background information from the input. To prove the existence of such dimensional confounders, we conduct motivating experiments from two perspectives. It is stated in [2] that Barlow Twins keeps improving as the dimensionality of the projection head increases, since the backbone network acts as a dimensionality bottleneck to constrain the dimensionality of the representation so that the learning paradigm of Barlow Twins can promote the representation capture more task-irrelevant information with the same dimensionality. Counterintuitively, our observation from the motivating experiments is in stark contrast with the conclusion in [2]. As shown in Figure 1 (b), the performance of Barlow twins as a function of the dimensionality presents a convex pattern on either CIFAR-10 or ImageNet. After reaching the peak, the performance gradually decreases as the dimensionality increases until it collapses. We reckon that since the task-relevant semantic information is limited, as the information captured in different dimensions is gradually decoupled, more task-irrelevant noisy information is encoded into the representation, i.e., the dimensional confounder increases, so the performance of the learned representation is getting worse until collapse.

To further validate our inference, we randomly mask several dimensions by frankly setting the masked values to 0, and then we evaluate the performance of the masked representation. The experimental results are demonstrated in Figure 2, where every single point denotes the result of a sampled mask. The representations with some specific masked dimensions can indeed achieve better performance than the unmasked representation, and such an observation is consistent on a small-scaled dataset, e.g., CIFAR-10, and a large-scale dataset, e.g., ImageNet. The dimensional mask rate that can improve the performance of the model on ImageNet is on average lower than that on CIFAR-10. We understand that this phenomenon depends on an ingredient: the representation learned on CIFAR-10 models discriminative information for 10 categories in 512 dimensions, while the representation uses 2048 dimensions to model information for 1000 categories on ImageNet, so the dimension-to-class ratio on CIFAR-10 (51.2) is much higher than that on ImageNet (2.048). Therefore, the presence of dimensional confounders naturally remains at a low level on ImageNet since supporting 1000-category classification requires a large amount of heterogeneous discriminative information. This empirical evidence further supports our assumption of the dimensional confounder.

To this end, we intuitively propose a simple yet effective approach *MetaMask*, short for the *Meta-learned dimensional Mask*, to adjust the impacts of different dimensions in a self-paced manner. In order to learn representations against dimensional redundancy and confounder, we regard the cross-correlation matrix constraint of Barlow Twins as a dimensional regularization term and further introduce a dimensional mask, which simultaneously enhances the gradient effect of the dimensions containing discriminative information and reduces that of the dimensions containing the confounder during training. We employ the meta-learning paradigm to train the learnable dimensional mask with the optimization objective of improving the performance of masked representations on typical self-supervised tasks, e.g., contrastive learning. Consequently, the dimensional mask incessantly adjusts the weights of different dimensions with respect to their gradient contribution to the optimization, so as to make the model focus on the acquisition of discriminative information, which can address the trivial solution caused by treating all dimensions equally in optimization, especially for high-dimensional representations. Empirically, we conduct head-to-head comparisons on various benchmark datasets, which prove the effectiveness of MetaMask. Concretely, the **contributions** of this paper are four-fold:

- We present a timely study on the adaptation of the state-of-the-art self-supervised models on downstream applications and identify two critical issues associated with the learned representation, i.e., dimensional redundancy and dimensional confounder.
- We propose a novel approach to tune the gradient effects of dimensions in a self-paced manner during optimization according to their improvements on a self-supervised task.
- Theoretically, we provide theorems and proofs to support that MetaMask achieves relatively tighter risk bounds for downstream classification compared to typical contrastive methods.
- We conduct extensive experiments to demonstrate the superiority of MetaMask over benchmark methods on downstream tasks.

2 Related Work

Self-supervised learning achieved impressive success in the field of representation learning without human annotation by constructing auxiliary self-supervised tasks. Typically, self-supervised approaches can be categorized as either generative or discriminative [20, 1, 14]. Conventional generative approaches bridge the input data and latent embedding by building a distribution over them, and the learned embeddings are treated as image representations [21, 22, 23, 24, 25, 26]. However, such a learning paradigm learns the distribution at the pixel-level, which is computationally expensive. The detailed information required for image generation may be redundant for representation learning.

For discriminative methods, contrastive learning achieves state-of-the-art performance on various downstream tasks [6, 27, 28, 29, 30, 31, 32, 33, 15]. Deep InfoMax [3] explores to maximize the mutual information between the input and output of an encoder. CPC [5] adopts noise-contrastive estimation [34] as the contrastive loss to train the model by maximizing the mutual information of multiple views, which is deduced by using the Kullback-Leibler divergence [35]. CMC [8] and AMDIM [4] employs contrastive learning on the multi-view data. SimCLR [1] and MoCo [9, 36] use large batch or memory bank to enlarge the amount of available negative features. To sample informative negatives, DebiasedCL [10] and HardCL [11] design sampling strategies by using positive-unlabeled learning methods [37, 38]. Motivated by [39], [40] provides an information theoretical framework for self-supervised learning and proposes an information bottleneck-based approach. However, these methods rely on a large number of negative samples to learn valuable representations while avoiding trivial solutions, i.e., constant representation. SimSiam [16] and BYOL [14] jointly waive negatives and avoid training degradation by using the techniques of asymmetric siamese network and stop-gradient. SSL-HSIC [41] provides a unified framework for negative-free contrastive learning based on Hilbert Schmidt Independence Criterion [42]. The representation learning paradigm employed by current contrastive methods is contrary to the coding criteria proposed by [19] and [43], respectively. Therefore, Barlow Twins [2] proposes to constrain the cross-correlation matrix between the features to be close to the identity matrix, which can effectively reduce the dimensional redundancy. Yet, such an approach still suffers from the dimensional confounder issue, and we thus propose a new meta-learning-based method to attenuate the over-learning of dimensional confounders by learning a mask to filter out irrelevant information.

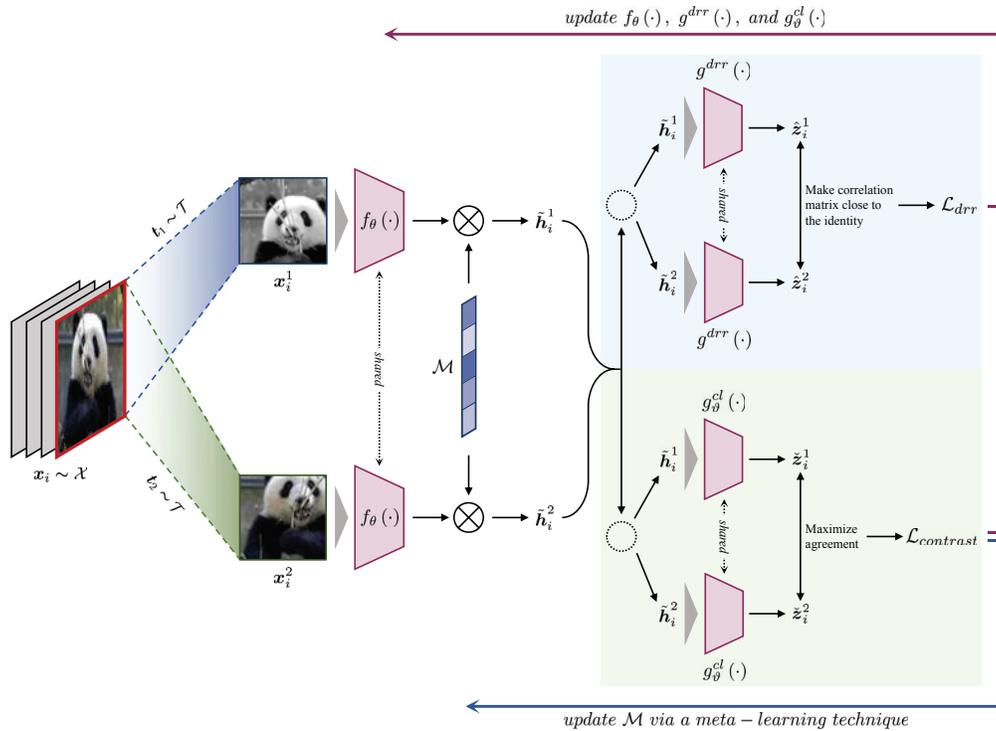


Figure 3: The architecture of MetaMask. The solid red line pointing backwards represents the regular training step, where the encoder and projection heads are trained by jointly back-propagating \mathcal{L}_{drr} and $\mathcal{L}_{contrast}$. The solid blue line pointing backwards represents the meta-learning step, where the dimensional mask \mathcal{M} is trained by performing the second-derivative technique on $\mathcal{L}_{contrast}$. All learnable networks are trained until convergence.

3 Preliminaries

Formally, we suppose the input multi-view dataset as $\mathcal{X} = \{x_i^j | i \in \llbracket 1, N^* \rrbracket, j \in \llbracket 1, M \rrbracket\}$, where N^* denotes the number of samples, and M denotes the number of views. Therefore, x_i^j represents the j -th view of the i -th sample following $x_i^j \sim \mathcal{X}$. The multiple views are built by performing identically distributed data augmentations [8, 1, 2], i.e., $\{t_j\}_{j=1}^M \sim \mathcal{T}$.

3.1 Contrastive Learning

We recap the preliminaries of contrastive learning [1], which aims to learn an embedding that jointly maximizes agreement between views of the same sample and separates the views of different samples in the latent space. Given a minibatch of N examples sampled from the multi-view dataset \mathcal{X} , *positives* represent the pairs $\{x_i^j, x_i^{j'}\}$ where $j, j' \in \llbracket 1, M \rrbracket$, and *negatives* represent $\{x_i^j, x_{i'}^{j'}\}$ where $i \neq i'$. Conventional contrastive methods feed the input x_i^j into a view-shared encoder $f_\theta(\cdot)$ to learn the representation h_i^j , which is mapped into a feature z_i^j by a projection head $g_\vartheta(\cdot)$, where θ and ϑ are the network parameters. $f_\theta(\cdot)$ and $g_\vartheta(\cdot)$ are trained by using the contrastive loss [5, 1]:

$$\mathcal{L}_{contrast} = - \sum_{i \in \llbracket 1, N \rrbracket} \sum_{\substack{j \in \llbracket 1, M-1 \rrbracket \\ j^+ \in \llbracket j+1, M \rrbracket}} \log \frac{\exp \left[d \left(z_i^j, z_i^{j^+} \right) / \tau \right]}{\sum_{i' \in \llbracket 1, N \rrbracket} \sum_{j' \in \llbracket 1, M \rrbracket} \mathbb{1}_{[i \neq i' \vee j \neq j']} \cdot \exp \left[d \left(z_i^j, z_{i'}^{j'} \right) / \tau \right]} \quad (1)$$

where $d(\cdot)$ is a discriminating function to measure the similarity of the feature pairs, $\mathbb{1}_{[i \neq i' \vee j \neq j']}$ denotes an indicator function equalling to 1 if $i \neq i'$ or $j \neq j'$, and τ is a temperature parameter valued by following [1]. In inference, the projection head $g_\vartheta(\cdot)$ is discarded, and the representation h_i^j is directly used for downstream tasks.

3.2 Dimensional Redundant Reduction

Following the principle proposed by [19], [2] introduces a redundancy-reduction objective function:

$$\mathcal{L}_{drr} = \sum_{\substack{j \in \llbracket 1, M-1 \rrbracket \\ j^+ \in \llbracket j+1, M \rrbracket}} \left[\sum_{k \in \llbracket 1, D \rrbracket} \left(1 - C_{kk}^{j, j^+}\right)^2 + \lambda \sum_{\substack{k, k' \in \llbracket 1, D \rrbracket \\ k \neq k'}} \left(C_{kk'}^{j, j^+}\right)^2 \right] \quad (2)$$

where λ is a positive hyper-parameter valued by following [2], and C is the cross-correlation matrix computed between the outputs of the encoder along the batch dimension:

$$C_{kk'}^{j, j^+} = \frac{\sum_{i \in \llbracket 1, N \rrbracket} z_{i, k}^j \cdot z_{i, k'}^{j^+}}{\sqrt{\sum_{i \in \llbracket 1, N \rrbracket} \left(z_{i, k}^j\right)^2} \cdot \sqrt{\sum_{i \in \llbracket 1, N \rrbracket} \left(z_{i, k'}^{j^+}\right)^2}} \quad (3)$$

where k and k' index the dimension of features. C is a square matrix with the size of D , which denotes the dimensionality of the projected feature z . The value range of the elements in C is $[-1, 1]$, i.e., complete anti-correlation to complete correlation.

4 Methodology

Our goal is to learn discriminative representations against dimensional redundancy and dimensional confounder from multiple views by performing self-supervised learning. To this end, we propose MetaMask to automatically determine the optimal dimensionality of the representation by decoupling information into different dimensions and filtering out those dimensions with irrelevant information. MetaMask addresses the dimensional redundancy issue by adopting a redundancy-reduction regularization. To tackle the dimensional confounder problem, MetaMask learns a dimensional mask to adjust the gradient weights of dimensions according to their contributions to improving a specific self-supervised task during optimization, which is achieved in a meta-learning manner.

MetaMask's architecture is shown in Figure 3. Specifically, the sample x_i^j is first fed into the encoder to obtain its representation h_i^j . Then, we build a learnable dimensional mask $\mathcal{M} = \{\omega_k | k \in \llbracket 1, D \rrbracket\}$, which is introduced to assign a *weight* to each dimension of the representation h_i^j by

$$\tilde{h}_i^j = h_i^j \otimes \mathcal{M} \quad (4)$$

where \tilde{h}_i^j denotes the masked representation, and $\cdot \otimes \cdot$ is an element-wise Hadamard product function. \tilde{h}_i^j is then fed into two separate projection heads $g^{drr}(\cdot)$ and $g^{cl}(\cdot)$, respectively. $g^{cl}(\cdot)$ plays the same role as in conventional contrastive learning methods, i.e., the contrastive loss $\mathcal{L}_{contrast}$ in Eq. (1) is applied to the output features of $g^{cl}(\cdot)$. For the projection head $g^{drr}(\cdot)$, the redundancy-reduction loss \mathcal{L}_{drr} in Eq. (2) is performed to its output features to tackle the dimensional redundancy.

In training, the encoder $f_\theta(\cdot)$, the projection heads $g^{drr}(\cdot)$ and $g^{cl}(\cdot)$, as well as the dimensional mask \mathcal{M} are trained in a meta-learning manner consisting of two steps. In the first regular training step, we train $f_\theta(\cdot)$, $g^{drr}(\cdot)$ and $g^{cl}(\cdot)$ by jointly minimizing the redundancy-reduction and contrastive losses, which is formalized by

$$\mathcal{L}_{regular} = \mathcal{L}_{drr} + \alpha \cdot \mathcal{L}_{contrast} \quad (5)$$

where α is a coefficient that controls the balance between \mathcal{L}_{drr} and $\mathcal{L}_{contrast}$, and we evaluate the influence of this hyper-parameter in Section 6.

In the second meta-learning-based step, we update \mathcal{M} by using the second-derivative technique [44] to solve a bi-level optimization problem. We encourage \mathcal{M} to *mask* (reduce the gradient weights of) the specific dimensions containing task-irrelevant information, which are regarded as dimensional confounders, so that the encoders can better explore the discriminative information in training. To this end, \mathcal{M} can be updated by computing its gradients with respect to the performance of $f_\theta(\cdot)$ and $g^{cl}(\cdot)$. The corresponding performance is measured by using the gradients of $f_\theta(\cdot)$ and $g^{cl}(\cdot)$ during the back-propagation of contrastive loss. Formally, we update the dimensional mask \mathcal{M} by

$$\arg \min_{\mathcal{M}} \mathcal{L}_{contrast} \left(g_{\theta_{trial}}^{cl} \left(f_{\theta_{trial}} \left(\mathcal{X}' \otimes \mathcal{M} \right) \right) \right) \quad (6)$$

Algorithm 1 MetaMask

Input: Multi-view dataset \mathcal{X} , minibatch size N , and a hyper-parameter α .

Initialize The neural network parameters: θ for $f_\theta(\cdot)$, ϑ for $g_\vartheta^{cl}(\cdot)$, ϑ^{drr} for $g^{drr}(\cdot)$, and $\mathcal{M} = \{\omega_k | k \in \llbracket 1, D \rrbracket\}$. The learning rates: ℓ_θ and ℓ_ϑ , etc.

repeat

for t -th training iteration **do**

 Iteratively sample a minibatch \mathcal{X}' with N examples from \mathcal{X} .

 # regular training step, fix \mathcal{M}

$\arg \min_{\theta, \vartheta, \vartheta^{drr}} \mathcal{L}_{drr} (g_{\vartheta^{drr}}^{drr} (f_\theta (\mathcal{X}') \otimes \mathcal{M})) + \alpha \cdot \mathcal{L}_{contrast} (g_\vartheta^{cl} (f_\theta (\mathcal{X}') \otimes \mathcal{M}))$

 # compute trial weights and retain computational graph, fix θ and ϑ

$\theta_{trial} = \theta - \ell_\theta \nabla_\theta \mathcal{L}_{contrast} (g_\vartheta^{cl} (f_\theta (\mathcal{X}') \otimes \mathcal{M}))$

$\vartheta_{trial} = \vartheta - \ell_\vartheta \nabla_\vartheta \mathcal{L}_{contrast} (g_\vartheta^{cl} (f_\theta (\mathcal{X}') \otimes \mathcal{M}))$

 # meta training step using second derivative

$\arg \min_{\mathcal{M}} \mathcal{L}_{contrast} (g_{\vartheta_{trial}}^{cl} (f_{\theta_{trial}} (\mathcal{X}') \otimes \mathcal{M}))$

end for

until θ , ϑ , ϑ^{drr} , and \mathcal{M} converge.

where \mathcal{X}' represents a minibatch sampled from the training dataset \mathcal{X} , ϑ_{trial} and θ_{trial} represent the trial weights of the encoders and projection heads after one gradient update using the contrastive loss defined in Equation 1, respectively. We formulate the updating of such trial weights as follows:

$$\theta_{trial} = \theta - \ell_\theta \nabla_\theta \mathcal{L}_{contrast} (g_\vartheta^{cl} (f_\theta (\mathcal{X}') \otimes \mathcal{M})), \vartheta_{trial} = \vartheta - \ell_\vartheta \nabla_\vartheta \mathcal{L}_{contrast} (g_\vartheta^{cl} (f_\theta (\mathcal{X}') \otimes \mathcal{M})) \quad (7)$$

where ℓ_θ and ℓ_ϑ are learning rates. The intuition behind such a behavior is that we perform a derivative over the derivative (Hessian matrix) of the combination $\{\theta, \vartheta\}$ to update \mathcal{M} , i.e., the second-derivative trick. Specifically, we compute the derivative with respect to \mathcal{M} by using a retained computational graph of $\{\theta, \vartheta\}$ and then update \mathcal{M} by back-propagating the derivative as Equation 6.

We train \mathcal{M} only based on the performance of contrastive learning for two reasons: 1) the dimensional redundancy reduction is a regularization term with strong constraints and generally stable convergence, and it is independent of downstream tasks; 2) such a meta-learning-based training approach that promotes contrastive learning performance empowers our method to constrain the upper and lower bounds of the cross-entropy loss on downstream tasks, which ensures that \mathcal{M} can partially mask the gradient contributions of dimensions containing task-irrelevant information and further promote the encoder to focus on learning task-relevant information, which is demonstrated in Section 5.

The two steps for updating $f_\theta(\cdot)$, $g^{drr}(\cdot)$, and $g_\vartheta^{cl}(\cdot)$ and updating \mathcal{M} are iteratively repeated until convergence. The training pipeline is detailed by Algorithm 1.

5 Theoretical Analyses

Current contrastive methods follow the fundamental assumption of multi-view learning (Assumption 1 in [39], Assumption 1 in [40], and Assumption 4.1 in [45]), i.e., views (or distortions) of a same sample holds the invariant discriminative label-relevant information. Theorem 4.2 in [45] further proposes the guarantees for general encoders in the learning paradigm of regular contrastive learning, which proves that the contrastive loss in the self-supervised learning stage can constrain the upper and lower bounds of the cross-entropy loss in the supervised learning stage for downstream tasks. To elaborate the behavior of our MetaMask, we extend this Theorem to build a connection between the proposed masked representation and downstream performance by

Theorem 5.1. (Connecting Masked Representations to Downstream Cross-Entropy Loss). Under the general assumption of self-supervised multi-view learning, when \mathcal{M} is optimal, for any coupled $\mathbf{h}, \tilde{\mathbf{h}} \in \mathbb{R}$, $\mathcal{L}_{crossentropy}(\mathbf{h})$ for downstream classification can be bounded by $\mathcal{L}_{contrast} (g_\vartheta^{cl}(\tilde{\mathbf{h}}))$

Table 1: Comparison of different methods on classification accuracy (top 1). We follow the experimental setting of [8] and adopt *conv* and *fc* as the backbones in the experiments.

Model	IN-200 [46]		STL-10 [47]		CIFAR-10 [46]		CIFAR-100 [46]		
	conv	fc	conv	fc	conv	fc	conv	fc	
Barlow Twins [2]	39.81	40.34	80.97	81.43	76.63	78.49	52.80	52.95	
+ MetaMask	+4.95	+3.10	+1.72	+3.99	+4.70	+8.61	+8.41	+5.22	+3.83
SwAV [6]	39.67	39.02	74.65	75.30	72.19	72.36	46.58	46.75	
+ MetaMask	+5.00	+0.61	+1.03	+6.69	+6.89	+8.18	+4.86	+5.49	+6.23
SimCLR [1]	36.24	39.83	75.57	77.15	80.58	80.07	50.03	49.82	
+ MetaMask	+3.60	+4.32	+0.99	+5.80	+5.75	+2.27	+1.57	+3.67	+4.42
CMC [8]	41.58	40.11	83.03	85.06	81.31	83.28	58.13	56.72	
+ MetaMask	+1.73	+1.33	+1.95	+1.93	+1.07	+3.93	+3.62	-0.09	+0.06
BYOL [14]	41.59	41.90	81.73	81.57	77.18	80.01	53.64	53.78	
+ MetaMask	+1.68	+1.13	+1.39	+1.28	+1.87	+2.08	+2.72	+2.04	+0.94
SimSiam [16]	41.03	41.27	80.91	81.88	78.14	81.13	52.55	53.52	
+ MetaMask	+1.44	+0.79	+1.40	+1.44	+2.08	+1.89	+0.26	+2.15	+1.50
DCL [10]	38.79	40.26	77.09	78.39	80.89	80.93	51.38	51.09	
+ MetaMask	+1.83	+1.49	+1.43	+3.85	+2.67	+0.44	+0.41	+1.67	+2.65
HCL [11]	40.05	41.23	79.86	80.20	82.13	82.76	52.69	53.13	
+ MetaMask	+1.14	+0.38	-0.22	+1.66	+1.89	+1.52	+1.19	+1.34	+1.13

in self-supervised learning:

$$\begin{aligned}
 \mathcal{L}_{contrast} \left(g_{\vartheta}^{cl}(\tilde{\mathbf{h}}) \right) &= \sqrt{\Phi \left(g_{\vartheta}^{cl}(\tilde{\mathbf{h}}) \mid \mathbf{y} \right)} - \frac{1}{2} \cdot \sum_{k \in [1, D]} \sqrt{\Phi \left(\lfloor g_{\vartheta}^{cl}(\tilde{\mathbf{h}}) \rfloor^k \mid \mathbf{y} \right)} - \mathcal{O} \left(M^{-\frac{1}{2}} \right) \\
 &\leq \mathcal{L}_{crossentropy}(\mathbf{h}) + \log \left(\frac{M}{D} \right) \leq \mathcal{L}_{contrast} \left(g_{\vartheta}^{cl}(\mathbf{h}) \right) + \sqrt{\Phi \left(g_{\vartheta}^{cl}(\mathbf{h}) \mid \mathbf{y} \right)} + \mathcal{O} \left(M^{-\frac{1}{2}} \right)
 \end{aligned} \tag{8}$$

where M denotes the quantity of negative samples, D denotes the dimensionality of the representation, \mathbf{h} denotes the unmasked representation, $\tilde{\mathbf{h}}$ denotes the masked representation obtained by Equation 4, \mathbf{y} denotes the target label, $\Phi(\cdot \mid \mathbf{y}) = \mathbb{E}_{\mathcal{P}(\cdot \mid \mathbf{y})} [\mathbb{E}_{\mathcal{P}(\cdot \mid \mathbf{y})} \delta(\cdot, \mathbb{E}_{\mathcal{P}(\cdot \mid \mathbf{y})})]$ represents the conditional variance function where $\delta(\cdot, \cdot)$ is a discriminating function to measure the difference between the terms, e.g., $\delta(\cdot, \cdot) = \|\cdot - \cdot\|^2$ in low-dimensional embedding space or $\delta(\cdot, \cdot) = -\log(\cos(\cdot, \cdot)) = -\log \frac{|\cdot| \times |\cdot|}{|\cdot| \times |\cdot|}$ in high-dimensional embedding space, $\lfloor \cdot \rfloor^k$ is a function acquiring k -th dimension vector, $\log \left(\frac{M}{D} \right)$ is a constant, and $\mathcal{O} \left(M^{-\frac{1}{2}} \right)$ is the approximation error's order.

Theorem 5.1 states that reducing the risk of contrastive learning can improve the performance on downstream tasks, which further supports our intuition to make the model focus on the acquisition of discriminative information by learning a dimensional mask with the objective of improving the performance of masked representations on contrastive learning is theoretically sound. Refer to Appendix A.1 for the elaboration of MetaMask's behavior to mask the gradient contribution of specific dimensions containing confounder information. Hence, we derive

Theorem 5.2. (Guarantees for Reduced Conditional Variance of Masked Representations). When \mathcal{M} is optimal, for any coupled $\mathbf{h}, \tilde{\mathbf{h}} \in \mathbb{R}$, given label \mathbf{y} , the conditional variance of $\tilde{\mathbf{h}}$ is reduced:

$$\Phi \left(g_{\vartheta}^{cl}(\tilde{\mathbf{h}}) \mid \mathbf{y} \right) \leq \Phi \left(g_{\vartheta}^{cl}(\mathbf{h}) \mid \mathbf{y} \right), \text{ yet } \Phi \left(\lfloor g_{\vartheta}^{cl}(\tilde{\mathbf{h}}) \rfloor^k \mid \mathbf{y} \right) \cong \Phi \left(\lfloor g_{\vartheta}^{cl}(\mathbf{h}) \rfloor^k \mid \mathbf{y} \right). \tag{9}$$

Theorem 5.2 states that given the label \mathbf{y} , the masked representation $\tilde{\mathbf{h}}$ has smaller conditional variance than \mathbf{h} in contrastive learning. We bring Theorem 5.2 into Theorem 5.1 to derive a conclusion: our approach can better bound the downstream classification risk, i.e., the upper and lower bounds of supervised cross-entropy loss obtained by MetaMask are tighter than typical contrastive learning methods. Please refer to Appendix A.2 for the corresponding proofs.

6 Experiments

6.1 Benchmarking MetaMask with Various Backbones

Implementations. For the comparisons demonstrated in Table 1, we uniformly set the batch size as 64, and we adopt a network with the 5 convolutional layers in AlexNet [48] as *conv* and a network with 2 additional fully connected layers as *fc*. We straightforwardly adopt the same data augmentations in [8] and [49], and the memory bank [50] is adopted to facilitate calculations by retrieving 4,096 past negative features during training. Following the experimental principle of linear probing, we remove the projection head and re-train an MLP as the classification head for the learned representation.

For the experiments in Table 2, the batch size is valued by 512, and ResNet-18 [51] is used as the backbone encoder. We adopt the data augmentation and other experimental settings following [2]. Note that the KNN classifier is adopted for evaluation.

Discussion of the comparisons using conv and fc encoders. Table 1 shows the comparisons on four benchmark datasets, and the corresponding positive or negative margins are reported. We observe that MetaMask beats the best prior methods by a significant margin on most downstream tasks, which demonstrates that MetaMask can efficiently improve the performance of the learned representation when the self-supervision is insufficient, i.e., the limited batch size and weak encoder.

Discussion of the comparisons using the ResNet encoder. We further perform classification comparisons using a strong encoder, i.e., ResNet. Table 2 shows that MetaMask can consistently improve the state-of-the-art methods, which indicates that the proposed method has strong adaptability to different encoders. The experimental evidence proves our assumption of the dimensional confounder. Thus, given a fixed dimensionality of the projection head, MetaMask can indeed improve the discriminativeness of the learned representation by *masking* the dimensional confounder.

6.2 Validation of the Robustness of MetaMask against Dimensional Confounder

We conduct experiments to prove that MetaMask is able to mitigate the negative impact of the dimensional confounder on the discriminativeness of the learned representations, and the results are shown in Figure 4. The intriguing outcome demonstrates two merits of the proposed method: 1) even if the dimensionality of the projection head changes, MetaMask is consistently robust to the dimensional confounder, i.e., the model maintains relatively consistent performance; 2) MetaMask can alleviate the negative impact of *curse of dimensionality* [52, 53], e.g., given a fixed number of examples, as the dimensionality of the representation increases during training, the performance of the model on downstream tasks instead degenerates. Concretely, Figure 2 demonstrates that regardless of whether representation learning suffers from the curse of dimensionality, dimensionality confounding factors are widely present in learned representations, and hence, the defined dimensional confounder is a more generalized issue of the curse of dimensionality. Refer Appendix A.3 for the further discussion. Moreover, we observe a surprising result that MetaMask* achieves better performance. We provide an understandable explanation: the essence of the curse of dimensionality is that a network with a large number of parameters is prone to over-fitting to the training data. Therefore, compared with the annealing strategy, a fixed and relatively large learning rate can prevent the model from falling into a local optimum (over-fitting) during optimization. Our further experiments show that MetaMask* collapses abruptly at larger dimensional settings, e.g., 16384, which demonstrates that the trick of using a fixed learning rate can only unsteadily improve the performance of MetaMask, but this behavior is *not* essential. Figure 4 shows that under the different experimental settings, vanilla MetaMask consistently counteracts the negative effects of the dimensional confounder.

7 Conclusion and Future Discussion

From the dimensional perspective, we first clarify that the dimensional redundancy and confounder issues are the essence of the current self-supervised learning paradigm's defects. Multiple motivating experiments are conducted to support our viewpoint, and the intriguing results are opposite to the traditional conclusions. To jointly tackle the dimensional redundancy and confounder issues, we propose a heuristic learning approach, called MetaMask, which leverages a second-derivative technique to learn a dimensional mask in order to adjust the gradient weights of different dimensions during training. The proposed theoretical analyses demonstrate that MetaMask obtains relatively

Table 2: Comparison of different methods on classification accuracy (top 1). We adopt *ResNet-18* as the backbones.

Model	CIFAR-10	CIFAR-100	STL-10	IN-200
Barlow Twins [2]	85.72	60.31	73.62	45.10
+ MetaMask ⁺	87.53 ^{+1.81}	61.42 ^{+1.11}	73.97 ^{+0.35}	46.81 ^{+1.71}
SimCLR [1]	81.73	58.27	70.09	44.46
+ MetaMask ⁺	86.01 ^{+4.28}	61.03 ^{+2.76}	74.90 ^{+4.81}	45.87 ^{+1.41}
BYOL [14]	88.05	60.94	72.04	46.72
+ MetaMask ⁺	87.53 ^{-0.52}	61.42 ^{+0.48}	73.97 ^{+1.93}	46.81 ^{+0.11}
DCL [10]	85.61	59.29	71.18	45.77
+ MetaMask ⁺	86.47 ^{+0.86}	60.82 ^{+1.53}	74.79 ^{+3.61}	46.35 ^{+0.58}
HCL [11]	85.27	61.21	71.92	46.90
+ MetaMask ⁺	85.97 ^{+0.70}	61.63 ^{+0.42}	74.03 ^{+2.11}	46.06 ^{+0.55}
NNCLR [54]	81.44	61.64	71.89	47.10
+ MetaMask ⁺	85.78 ^{+4.34}	61.89 ^{+0.25}	74.02 ^{+2.13}	46.90 ^{-0.20}

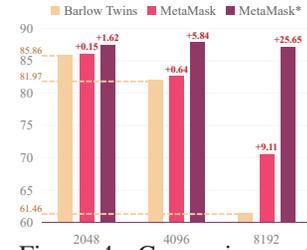


Figure 4: Comparisons of varying the dimensionality of the projection head on CIFAR-10 by using ResNet-18. * denotes MetaMask using a trick of fixed learning rate instead of the cosine annealing strategy.

tighter risk bounds for downstream classification compared to conventional contrastive methods. From the experiments, we observe a consistent improvement in the performance of our method over state-of-the-art methods on benchmark datasets.

One of our critical contributions is proposing a heuristic learning paradigm for self-supervised learning: regarding the dimensional redundancy reduction as a regularization term and further exploring to mask dimensions containing the confounder. In the current methodology, we train the learnable dimensional mask with the objective of improving the performance of masked representations on a specific *self-supervised task*, i.e., contrastive learning, and provide the theoretical analysis to prove that such a learning approach can indeed prompt the model focus on acquiring discriminative information. However, contrastive learning is immutable for the *self-supervised task* in MetaMask. It is commonly acknowledged that masked image modeling (MIM) methods [55, 56, 57] achieve impressive performance in self-supervised learning. Although the theoretical contribution bridging the gap between MIM loss and cross-entropy loss on downstream tasks has not yet emerged, we still deem combining MIM approaches with our proposed paradigm is attractive to the community.

Acknowledgements

The authors would like to thank the anonymous reviewers for their valuable comments. This work is supported in part by National Natural Science Foundation of China No. 61976206 and No. 61832017, Key Special Project for Introduced Talents Team of Southern Marine Science and Engineering Guangdong Laboratory (Guangzhou) No. GML2019ZD0603, National Key Research and Development Program of China No. 2019YFB1405100, Foshan HKUST Projects (FSUST21-FYTRI01A, FSUST21-FYTRI02A), Beijing Outstanding Young Scientist Program NO. BJJWZYJH012019100020098, Beijing Academy of Artificial Intelligence (BAAI), the Fundamental Research Funds for the Central Universities, the Research Funds of Renmin University of China 21XNLG05, and Public Computing Cloud, Renmin University of China. This work is also supported in part by Intelligent Social Governance Platform, Major Innovation & Planning Interdisciplinary Platform for the “Double-First Class” Initiative, Renmin University of China, and Public Policy and Decision-making Research Lab of Renmin University of China.

References

- [1] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, Proceedings of Machine Learning Research. PMLR, 2020.
- [2] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, Proceedings of Machine Learning Research. PMLR, 2021.

- [3] R. Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Philip Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- [4] Philip Bachman, R. Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. In *NeurIPS 2019*, 2019.
- [5] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [6] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [7] Jiangmeng Li, Wenwen Qiang, Changwen Zheng, Bing Su, Farid Razzak, Ji-Rong Wen, and Hui Xiong. Modeling multiple views via implicitly preserving global consistency and local complementarity. *IEEE Transactions on Knowledge and Data Engineering*, pages 1–17, 2022.
- [8] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019.
- [9] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*. Computer Vision Foundation / IEEE, 2020.
- [10] Ching-Yao Chuang, Joshua Robinson, Yen-Chen Lin, Antonio Torralba, and Stefanie Jegelka. Debaised contrastive learning. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [11] Joshua David Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. Contrastive learning with hard negative samples. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [12] Jiangmeng Li, Wenwen Qiang, Changwen Zheng, Bing Su, and Hui Xiong. Metaug: Contrastive learning via meta feature augmentation. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato, editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 12964–12978. PMLR, 2022.
- [13] Wenwen Qiang, Jiangmeng Li, Changwen Zheng, Bing Su, and Hui Xiong. Interventional contrastive learning with meta semantic regularizer. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato, editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 18018–18030. PMLR, 2022.
- [14] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Ávila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent - A new approach to self-supervised learning. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [15] Aleksandr Ermolov, Aliaksandr Siarohin, Enver Sangineto, and Nicu Sebe. Whitening for self-supervised representation learning. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 3015–3024. PMLR, 2021.
- [16] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*. Computer Vision Foundation / IEEE, 2021.

- [17] Hang Gao, Jiangmeng Li, Wenwen Qiang, Lingyu Si, Fuchun Sun, and Changwen Zheng. Bootstrapping informative graph augmentation via A meta learning approach. In Luc De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 3001–3007. ijcai.org, 2022.
- [18] Madelyn Glymour, Judea Pearl, and Nicholas P Jewell. *Causal inference in statistics: A primer*. John Wiley & Sons, 2016.
- [19] H. B Barlow. Possible principles underlying the transformations of sensory messages. In *Sensory Communication*. The MIT Press, 1961.
- [20] Carl Doersch, Abhinav Gupta, and Alexei A. Efros. Unsupervised visual representation learning by context prediction. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*. IEEE Computer Society, 2015.
- [21] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Machine Learning, Proceedings of the Twenty-Fifth International Conference (ICML 2008), Helsinki, Finland, June 5-9, 2008*. ACM, 2008.
- [22] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.
- [23] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, 2014.
- [24] Jeff Donahue and Karen Simonyan. Large scale adversarial representation learning. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 2019.
- [25] Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Alex Lamb, Martín Arjovsky, Olivier Massopietto, and Aaron C. Courville. Adversarially learned inference. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [26] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- [27] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning? In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [28] Olivier J. Hénaff. Data-efficient image recognition with contrastive predictive coding. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 4182–4192. PMLR, 2020.
- [29] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*. Computer Vision Foundation / IEEE, 2020.
- [30] Junnan Li, Pan Zhou, Caiming Xiong, and Steven C. H. Hoi. Prototypical contrastive learning of unsupervised representations. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [31] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E. Hinton. Big self-supervised models are strong semi-supervised learners. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [32] Vikas Verma, Thang Luong, Kenji Kawaguchi, Hieu Pham, and Quoc V. Le. Towards domain-agnostic contrastive learning. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, *Proceedings of Machine Learning Research*. PMLR, 2021.

- [33] Tete Xiao, Xiaolong Wang, Alexei A. Efros, and Trevor Darrell. What should not be contrastive in contrastive learning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [34] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In Yee Whye Teh and D. Mike Titterton, editors, *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Chia Laguna Resort, Sardinia, Italy, May 13-15, 2010*, volume 9 of *JMLR Proceedings*, pages 297–304. JMLR.org, 2010.
- [35] J. Goldberger, S. Gordon, and H. Greenspan. An efficient image similarity measure based on approximations of kl-divergence between two gaussian mixtures. In *IEEE International Conference on Computer Vision*, 2003.
- [36] Xinlei Chen, Haoqi Fan, Ross B. Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *CoRR*, 2020.
- [37] Charles Elkan and Keith Noto. Learning classifiers from only positive and unlabeled data. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Las Vegas, Nevada, USA, August 24-27, 2008*. ACM, 2008.
- [38] Marthinus Christoffel du Plessis, Gang Niu, and Masashi Sugiyama. Analysis of learning from positive and unlabeled data. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, 2014.
- [39] K. Sridharan and S. M. Kakade. An information theoretic framework for multi-view learning. *Conference on Learning Theory*, 2008.
- [40] Yao-Hung Hubert Tsai, Yue Wu, Ruslan Salakhutdinov, and Louis-Philippe Morency. Self-supervised learning from a multi-view perspective. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [41] Yazhe Li, Roman Pogodin, Danica J. Sutherland, and Arthur Gretton. Self-supervised learning with kernel dependence maximization. *CoRR*, 2021.
- [42] Gretton A., Bousquet O., Smola A., and B. Schölkopf. Measuring statistical dependence with hilbert-schmidt norms. *ALT, Springer*, 2005.
- [43] Yaodong Yu, Kwan Ho Ryan Chan, Chong You, Chaobing Song, and Yi Ma. Learning diverse and discriminative representations via the principle of maximal coding rate reduction. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [44] Shikun Liu, Andrew J. Davison, and Edward Johns. Self-supervised generalisation with meta auxiliary learning. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 1677–1687, 2019.
- [45] Yifei Wang, Qi Zhang, Yisen Wang, Jiansheng Yang, and Zhouchen Lin. Chaos is a ladder: A new theoretical understanding of contrastive learning via augmentation overlap. *arXiv preprint arXiv:2203.13457*, 2022.
- [46] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [47] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 2011.
- [48] A. Krizhevsky, I. Sutskever, and G. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [49] Jiangmeng Li, Wenwen Qiang, Changwen Zheng, Bing Su, and Hui Xiong. Metaug: Contrastive learning via meta feature augmentation. *arXiv preprint arXiv:2203.05119*, 2022.
- [50] Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 3733–3742. Computer Vision Foundation / IEEE Computer Society, 2018.

- [51] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [52] R. Bellman and R. Karush. Dynamic programming: A bibliography of theory and application., *Dynamic Programming A Bibliography of Theory and Application*, 1964.
- [53] Vincent Spruyt. The curse of dimensionality in classification. 2014.
- [54] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. With a little help from my friends: Nearest-neighbor contrastive learning of visual representations. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 9568–9577. IEEE, 2021.
- [55] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.
- [56] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. *arXiv preprint arXiv:2111.06377*, 2021.
- [57] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. *arXiv preprint arXiv:2111.09886*, 2021.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? **[Yes]**
 - (b) Did you describe the limitations of your work? **[Yes]** See Section 7.
 - (c) Did you discuss any potential negative societal impacts of your work? **[No]** Since this work presents a general technique to tackle the dimensional confounder for self-supervised learning, we did not see any particular foreseeable ethical consequence.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]** See Section 5.
 - (b) Did you include complete proofs of all theoretical results? **[Yes]** See Appendix A.1 and Appendix A.2.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[Yes]** See the supplementary files.
 - (b) Did you specify all the training details (e.g., data splits, hyper-parameters, how they were chosen)? **[Yes]** See Section 6, Appendix A.4, and the supplementary files.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[No]** The experimental results provided by benchmark methods do not contain such error bars. We also follow this principle, and we reckon the reason behind such a behavior is that the comparisons are conducted in balanced datasets, so the difference between results is not excessively large.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[Yes]** See Section 1 and Appendix A.4.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? **[Yes]** See Section 1, Section 2, Section 5, and Section 6.
 - (b) Did you mention the license of the assets? **[No]** The assets we adopted are totally established and public to the community.
 - (c) Did you include any new assets either in the supplemental material or as a URL? **[Yes]** See the supplementary files.

- (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No] The datasets are all public to the community.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No] The datasets have been discussed, validated, and used by many researchers.
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]