

MixMask: Revisiting Masking Strategy for Siamese ConvNets

Kirill Vishniakov¹
kirill.vishniakov@mbzuai.ac.ae

Eric Xing^{1,2}
eric.xing@mbzuai.ac.ae

Zhiqiang Shen¹
zhiqiang.shen@mbzuai.ac.ae

¹ Mohamed bin Zayed University of
Artificial Intelligence

² Carnegie Mellon University

Abstract

The recent progress in self-supervised learning has successfully combined Masked Image Modeling (MIM) with Siamese Networks, harnessing the strengths of both methodologies. Nonetheless, certain challenges persist when integrating conventional erase-based masking within Siamese ConvNets. Two primary concerns are: (1) The continuous data processing nature of ConvNets, which doesn't allow for the exclusion of non-informative masked regions, leading to reduced training efficiency compared to ViT architecture; (2) The misalignment between erase-based masking and the contrastive-based objective, distinguishing it from the MIM technique. To address these challenges, this work introduces a novel filling-based masking approach, termed **MixMask**. The proposed method replaces erased areas with content from a different image, effectively countering the information depletion seen in traditional masking methods. Additionally, we unveil an adaptive loss function that captures the semantics of the newly patched views, ensuring seamless integration within the architectural framework. We empirically validate the effectiveness of our approach through comprehensive experiments across various datasets and application scenarios. The findings underscore our framework's enhanced performance in areas such as linear probing, semi-supervised and supervised finetuning, object detection and segmentation. Notably, our method surpasses the MSCN, establishing MixMask as a more advantageous masking solution for Siamese ConvNets. Our code and models are publicly available at github.com/kirill-vish/MixMask.

1 Introduction

Self-supervised learning is a popular method for deriving representations from data without requiring human-labeled annotations. One common approach within this paradigm is the Siamese Network, characterized by its dual-branch structure. The distance or relationship between these two branches is typically assessed using similarity loss [1, 2], contrastive loss [3, 4, 5], or distillation loss [6]. A recent innovation, termed Masked Image Modeling (MIM) [7, 8, 9], has shown promise in enhancing representation learning. To integrate the strengths of both masked design and Siamese networks, the Masked Siamese Network (MSN) built on ViT was introduced [10]. Parallely, convolutional networks such

as ConvNeXt [24] and its V2 iteration [60] have also been making strides, exhibiting strong performance in both labeled and unlabeled learning contexts, underlining the potential of ConvNets. A more recent advance is the Masked Siamese ConvNets (MSCN) [40], which is an extension of MSN but utilizes a ConvNet-based encoder. This method, however, takes the masking strategy directly from MIM [0, 19], without customizing it for the unique attributes of Siamese ConvNet. Additionally, both MSN and MSCN lean heavily on the multicrop strategy, as detailed in [0, 9], to counterbalance performance limitations stemming from the erasure function. As an example, while MSN produces ten additional perspectives with a focal mask in each cycle, MSCN creates two supplementary views. This results in an increased computational load during training.

Generally, erase-based masking works particularly well in tandem with the image patchify mechanism seen in ViT [12]. This mechanism creates image patches that the encoder processes independently. By doing so, masked patches can easily be omitted, subsequently reducing computational expenses [19]. In contrast, ConvNets employ a continuous data processing approach. This is not suitable for omitting masked patches because it would maintain the same processing time, even though the masked patches no longer carry significant semantic value. Another challenge arises from numerous contrastive frameworks employing embedding loss [0, 18]. This loss is not structured to restore the removed regions but rather to distinguish between different inputs. Moreover, multiple studies [0, 10] have shown that contrastive loss tends to capture more general, coarse-grained features. Masking operations compromise these overarching features, subsequently slowing down a model’s convergence rate. In contrast, vision transformer-based MIM methods [0, 19] aim to restore the masked sections by working directly within the pixel domain using a reconstruction loss.

We have pinpointed two primary limitations associated with Masked Siamese ConvNets: (1) The standard erase-based masking operation impinges upon global features crucial for the contrastive objective as referred in [0, 10]. Plainly put, conventional masking excises significant semantic data from the input, and such loss is irretrievable through subsequent processing. To illustrate, masking 25% of the image means a corresponding 25% information loss during training, thereby hindering the training’s efficacy. This highlights the pressing need for a refined masking approach to optimize the learning of representations. (2) The inherent “symmetric” semantic distance loss in Siamese networks does not adjust for the shift in semantic distance across varying views of an identically masked image. Notably, the global semantic essence of an image is altered when random segments are erased. Take, for example, an image of a dog in the center. The semantics of this image diverge markedly from a version where distinct portions of the dog are masked and thus hidden. Such uniform loss configuration has a marginal impact when both branches undergo masking. Nonetheless, issues arise when solely one branch gets masked. Given that asymmetric Siamese networks, as proposed in [30], are identified to be superior learners, a more adaptable loss system utilizing soft distance becomes indispensable to truly capture the semantic gap between the duo of asymmetric branches in masked Siamese convnets.

To counter the limitations mentioned, we introduce a filling-based masking technique as an alternative to the erasing method, ensuring that there is no information loss. Rather than randomly removing regions of an image, we opt to randomly select another image, harnessing its pixel values to populate the masked segments. Importantly, in contrast to the discrete random masking employed in [19], we discovered that a block-wise approach yields superior results. This method maintains crucial global information, essential for contrastive loss. Paired with our mask-filling strategy, this block-wise masking effectively replaces blank spaces with comprehensive viewpoints from other images within a batch. In recognition




Figure 1: Illustration of the different mask patterns with a mask grid size of 8. (a) and (b) are input images. (c) is the discrete/random mask pattern, and (d) and (e) are mixed images using this mask. (f) is the blocked mask pattern, and (g) and (h) are mixed images with a blocked mask. Discrete masking breaks (c) – (e) the completeness of an object which is important for the contrastive loss because it operates on the global object level. On the other hand, blocked masking (f) – (h) preserves important global features leading to superior performance.

of the semantic shift post-filling, we have incorporated the asymmetric loss design as detailed in [B]. Our filling-based technique, compared to MSCN’s erase-based strategy, supplants nondescript erased areas with regions that offer richer semantic insights. Notably, this method does away with the need for multicrops, only demanding one supplementary view per image, marking a significant efficiency leap over MSN and MSCN that utilize ten and two cropping pairs respectively. Such invaluable information derived from our filling-based approach can be harnessed by Siamese networks, amplifying the training representation. Our ablation studies confirm its superior efficacy for the Masked Siamese ConvNets.

Our proposed approach has been exhaustively validated across the CIFAR-100, Tiny-ImageNet, and ImageNet-1K datasets and seamlessly incorporated into various Siamese ConvNets, including MoCo, BYOL, SimCLR, and SimSiam. Across all datasets, our method substantially elevated the performance of multiple baseline models. We also noted consistent improvement in our refined models across semi-supervised and supervised fine-tuning tasks, including object detection and segmentation. Our contributions in this work are as follows:

- We reveal that *Siamese ConvNets* experience reduced convergence when using the conventional erase-based masking scheme. In response, we introduce an effective filling-based masking strategy, which better complements the image-level contrastive objective of self-supervised Siamese ConvNets.
- We integrate a flexible loss structure with a soft distance to harmonize the combined *masking* and *Siamese* architecture. This ensures there are no discrepancies between the transformed input and objectives in Masked Siamese ConvNets.
- Extensive experiments are performed on various datasets and Siamese frameworks to solidify the efficacy and broad applicability of our approach. Our results consistently indicate that our method outperforms MSCN in areas such as linear probing, semi-supervised fine-tuning, and downstream tasks like detection and segmentation.

2 Related Work

Masked Siamese Networks. Recent developments in MIM and siamese self-supervised learning have prompted research efforts to combine these two techniques. One such approach, Masked Siamese Networks (MSN), was proposed in [I]. MSN generates an anchor view and a target view and applies masking operations exclusively to the anchor branch. To assign masked anchor representations to the same cluster as the unmasked target, the method

employs prototype cluster assignment. By contrast, CNNs are less compatible with masked images as they operate at the pixel level, which can cause image continuity to be disrupted by masking operations. To address this, Masked Siamese ConvNets were introduced in [20].

Self-supervised Learning. Self-supervised learning (SSL) is a technique for representation learning that leverages vast amounts of unlabeled data. Early SSL approaches in the field of computer vision were based on pre-text tasks [24, 25, 29, 35]. A pivotal advancement in SSL was the introduction of a simple contrastive learning framework by [5, 9], which utilized a siamese architecture in conjunction with an InfoNCE [26] contrastive loss. MoCo [18] implemented a memory bank to store negative samples, but several studies have shown that negative pairs are not always necessary. BYOL [28] employed an asymmetric architecture with EMA, where the online network attempts to predict the representations of the target network. SimSiam [7] provided a straightforward siamese framework that incorporated a stop-gradient operation on one branch and an additional predictor step on the other. More recently, SSL has begun to adopt the use of vision transformers [12]. DINO [9] combined a distillation loss with vision transformers in a siamese framework.

3 Approach

In this section, we first introduce each component of our framework elaborately, including: (1) a filling-based masking strategy; (2) an asymmetric loss formulation with soft distance to match the proposed mix-masking scheme. Then, we provide an overview of the proposed architecture comparing to the basic model and differences from other counterparts.

3.1 Masking Strategy

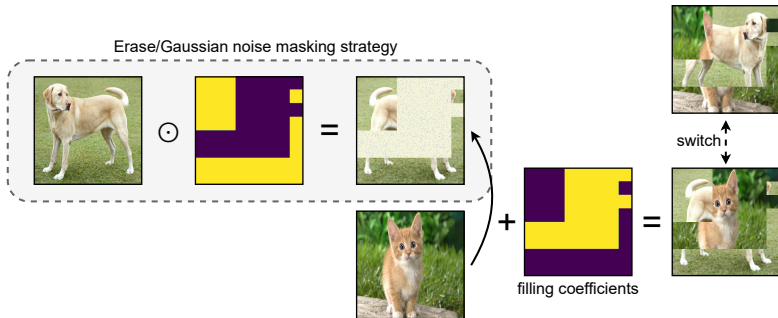


Figure 2: Illustration of the proposed filling-based masking strategy. The gray dashed box shows Erase/Gaussian noise [20] masking strategy. A formal definition of a switch image in the case of reverse permutation is given in Eq. 2.

Masking Scheme: Erasing or Filling? In Fig. 2, the prior *Erase-based Gaussian noise masking* strategy is illustrated inside the dashed box. In such strategy, erased regions can be filled with a Gaussian noise [20]. Different from Masked Image Modeling (MIM), which is to reconstruct the masked contents for learning good representation, the *Masked Siamese Networks* will not predict the information in removed areas, so erasing will only lose information and is not desired in the learning procedure. In contrast to erase-based masking, our filling-based strategy will repatch the removed areas using an auxiliary image, as shown in the right part of Fig. 2. After that, we will *switch* the content between the main and auxiliary

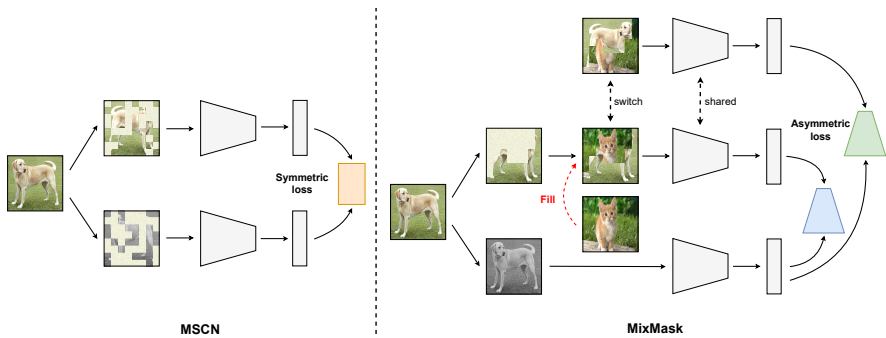


Figure 3: Illustration of the Masked Siamese ConvNets (left) and our proposed framework (right). MixMask branch incorporates asymmetry into the loss function design by generating images with different rates of similarity to the images in the original branch. In MixMask branch image of the truck is presented twice with different levels of similarity to the image in the original branch due to the regions masked with contents of another image.

images to generate a new image for information completeness of the two input images. For a given original image I_i we define its mixture as a mix of the pair (I_i, I_{n-i}) and *switch* image of I_i as a mixture of the pair (I_{n-i}, I_i) . Although the mixture contains information from another image, which may have different semantics, it will still be beneficial by cooperating with the soft loss design, discussed in Section 3.2, which takes into account the degree of each individual image included in the mixture.

Masking Pattern: Blocked or Discrete? Masking pattern determines the difficulty for the model to generate representations in siamese networks, even if the masking ratio is the same, the representation will still be different for different masking patterns. It will directly influence the information of input to further affect the latent representations. *Our observation on Masked Siamese ConvNets is opposite to that in MIM methods, which found discrete/random masking is better [49, 63].* From our empirical experiments, on CIFAR-100, blocked and discrete masking patterns achieve similar accuracy, and discrete is slightly better, however, on Tiny-ImageNet and ImageNet-1K, blocked mask clearly shows superiority over discrete/random. We explain this as that if the input size is small, the mask pattern is not so important since the semantic information of the object is still preserved. On larger datasets like Tiny-ImageNet and ImageNet-1K, discrete/random masking will entirely destroy the completeness of the object in an image, as shown in Fig. 1 (d, e), while this is crucial for ConvNet to extract a meaningful representation of the object. Blocked masking also better synergizes with the contrastive objective, which is biased to learn global features [6, 10]. Therefore, blocked masking shows superior ability on the MSCN and is a better choice than random discrete masking.

3.2 Distance in Siamese Networks

Objective Calculation: Inflexible or Soft? It has been observed [10] that different pre-text and data processes (e.g., masking, mixture) will change the semantic distance of two branches in the siamese networks, hence the default symmetric loss will no longer be aligned to reflect the true similarity of the representations. It thus far has not attracted enough attention for such a problem in this area. In this work, we introduce a soft objective calculation method that can fit the filling-based masking strategy in a better way. To calculate the soft distance, we start by generating a binary mask with a fixed grid size that will later be used

to mix a batch of images, denoted as \mathbf{I} , from a single branch. In case when we use a reverse permutation to obtain the mixture, each image in the batch with index i is mixed according to the mask with the image in the same batch but with index $n - i$ as described in Eq. 1:

$$\text{mix}_i = \text{mix}(\mathbf{I}_i, \mathbf{I}_{n-i}) = m \odot \mathbf{I}_i + (1 - m) \odot \mathbf{I}_{n-i} \quad (1)$$

$$\text{switch}_i = \text{mix}(\mathbf{I}_{n-i}, \mathbf{I}_i) = m \odot \mathbf{I}_{n-i} + (1 - m) \odot \mathbf{I}_i, \quad (2)$$

where n is batch size and m is mask. The mixed image contains parts from both \mathbf{I}_i and \mathbf{I}_{n-i} whose spatial locations in the mixture are defined by the contents of the binary mask. Using m and $1 - m$ ensures that each region in the mixture will contain pixels from exactly one image and that there will be no empty unfilled regions. Furthermore, to reflect the contribution of each image in the mixture, we calculate a mixture coefficient λ , which is equal to the ratio of the masked area to the total area of the image using the Eq. 3:

$$\lambda = \frac{\sum_{x,y} 1[\text{mask}(x,y) = 1]}{\text{width} \cdot \text{height}}, \quad (3)$$

where 1 is the indicator function that measures the masked area of an image.

Loss Function. The final loss is defined as a summation of the original loss and mixture loss: $L = L_{\text{Orig}} + L_{\text{MixMask}}$. Here we take a contrastive loss from MoCo [18] as an example, where q and k are two randomly augmented non-mixed views of the input batch. The notation stems from queries and keys from MoCo nomenclature, the L_{Orig} will be: $L_{\text{Orig}} = -\log \frac{\exp(q \cdot k / \tau)}{\sum_{i=0}^K \exp(q \cdot k_i / \tau)}$. The mixture term L_{MixMask} will contain two terms which are scaled with λ coefficient:

$$L_{\text{MixMask}} = \lambda L_{\uparrow} + (1 - \lambda) L_{\downarrow} = -\left(\lambda \log \frac{\exp(q_{\uparrow} \cdot k / \tau)}{\sum_{i=0}^K \exp(q_{\uparrow} \cdot k_i / \tau)} + (1 - \lambda) \log \frac{\exp(q_{\downarrow} \cdot k / \tau)}{\sum_{i=0}^K \exp(q_{\downarrow} \cdot k_i / \tau)} \right), \quad (4)$$

where q_{\uparrow} and q_{\downarrow} are normal and reverse orders of mixed queries in a mini-batch, k is the unmixed single key, λ is calculated using the Eq. 3 and τ is the temperature.

3.3 Framework Overview

Our framework overview is shown in Fig. 3. In this figure, the left is the conventional Masked Siamese ConvNets (MSCN), right is our proposed MixMask with asymmetric distance loss. The motivation behind this design is that directly erasing regions will lose a significant proportion of information in the Siamese ConvNets, which cannot be recovered by post-training. This is quite different from the mechanism of Masked Autoencoders (MAE) [19] that predict masked areas to learn good representations. According to this, we propose a filling-based scheme to overcome the drawback. The soft distance loss is designed to fit the true similarity of the two branches. We empirically show that with the integrality of mix-masking and objective, we can learn more robust and generalized representations from the masked input. **Differences from Prior Counterparts.** MSN and MSCN utilize regular erase-based masking and Un-Mix approach integrates Mixup and CutMix techniques into siamese networks for self-supervised learning. In contrast, our proposed MixMask method is a generalized masking approach specifically designed for siamese networks that allows for arbitrary mask area shapes. Our empirical study demonstrates that MixMask exhibits stronger representation learning capabilities. Interestingly, MixMask is also compatible with Un-Mix [20] that can be employed jointly to further improve performance and achieve state-of-the-art accuracy. SparK and ConvNeXt V2 propose a way of doing MIM with single branch ConvNet and reconstruction loss, while our work is focused on the Siamese contrastive framework.

4 Experiments

Base Models. In our experimental section, we use base models including: MoCo V1&V2 [8, 18], Un-Mix [30], SimCLR [6], BYOL [28] and SimSiam [7].

Datasets and Training Settings. We conduct experiments on Tiny-ImageNet [22], CIFAR-100 [21], and ImageNet-1K [11]. For CIFAR-100 and Tiny-ImageNet, we train each framework for 1000 epochs with ResNet-18 [16]. For ImageNet-1K, we pretrain ResNet-50 for 200 epochs and then finetune a linear classifier on top of the frozen features for 100 epochs and report Top-1 accuracy. We use MoCo (and MoCo V2 for ImageNet-1K) as a base framework unless stated otherwise. We report Top-1 on linear evaluation, except the case of MoCo on CIFAR-100 and Tiny-ImageNet for which we provide k -NN accuracy. * indicates that we build our method upon Un-Mix. All our experiments are conducted on A100 (40G) GPUs.

4.1 Ablation Study

In the experiments, we use a blocked masking strategy, masking ratio of 0.5, and grid size of 2, 4, 8 for CIFAR-100, Tiny-ImageNet, and ImageNet-1K if not stated otherwise.

Masking Strategy. We first explore how different masking strategies affect the final result. We consider three different hyperparameters: grid size, grid strategy, and masking ratio. To make our experiments more reliable, we consider three datasets with different image sizes: CIFAR-100: 32×32 , Tiny-ImageNet: 64×64 , and ImageNet-1K: 224×224 .

The first parameter grid size specifies the granularity of the $n \times n$ square grid of the image. We consider the following values of $n = 2, 4, 8, 16, 32, 48$. However, we have different upper bounds for different datasets depending on the spatial size of the images. From our experiments, we can conclude that a very large grid size completely disrupts the semantic features of the image and leads to poor performance. Our results indicate that optimal grid size increases proportionally to the input size of the image. We obtain the optimal grid size for CIFAR-100 is 2, for Tiny-ImageNet is 4, and for ImageNet-1K is 8. The masks with very small and very large grid sizes show bad performance. We think this happens because a small grid size does not provide enough variance in the mask structure whilst a large grid size destroys the important semantic features of the image.

We consider two different strategies for the random mask generation, discrete/random mask and blocked mask. We generate a blocked mask according to the algorithm described in [2]. A discrete mask does not have any underlying structure, whilst a blocked mask is generated in a way to preserve global spatial continuity and, thus, is more suitable for capturing global semantic features. For CIFAR-100, we observe a negligible difference between blocked and discrete masks. On the other hand, for Tiny-ImageNet and ImageNet-1K, which have larger spatial sizes of the image, blocked mask performs better than discrete. This observation is different from MIM-based approaches [19, 63], where a discrete mask achieves

	CIFAR-100	Tiny-ImageNet	ImageNet-1K
Input size	32×32	64×64	224×224
Grid Size			
2×2	68.1	46.4	68.4
4×4	67.6	47.4	69.0
8×8	67.5	46.5	69.2
16×16	–	46.1	68.5
32×32	–	–	68.7
48×48	–	–	68.8
Masking Strategy			
blocked	67.8	47.4	69.2
discrete	68.1	46.4	68.4
Masking Ratio			
0.25 0.75	66.8	45.6	68.5
0.5	68.1	47.4	69.2
uniform(0, 1)	67.6	45.5	68.7
Switch Mixture			
Yes	68.1	47.4	69.2
No	67.3	45.6	67.8

Table 1: Ablation study on masking and switch strategies using MoCo V1/V2 with original and MixMask branches. k -NN accuracy averaged over 3 runs is reported.

whilst a large grid size destroys the important

	CIFAR-100				Tiny-INet		INet-1K
	MoCo	SimSiam	BYOL	SimCLR	MoCo	BYOL	MoCoV2
Vanilla	65.7	66.7	66.8	67.1	42.8	51.4	67.5
Un-Mix [60]	68.6	70.0	70.0	69.8	45.3	52.9	68.5
MixMask	68.4	69.8	71.9	69.2	47.4	50.8	69.2
MixMask*	70.0	70.6	72.1	70.0	47.9	53.7	69.5

Table 2: Ablation on different base frameworks. * means we use our method with Un-Mix.

better performance and highlights the importance of maintaining global features when generating an image mixture as opposed to the case when erasing parts of the image by performing a vanilla masking operation. This also reflects the different mechanisms and requirements of reconstruction-based and siamese-based self-supervised learning approaches.

For the masking ratio, we consider two cases with constant values of 0.5 and 0.25/0.75 (as there is no difference between these two due to the loss function design) and two cases when the value is sampled from the uniform distribution with different bounds. We obtain the best results for masking ratio 0.5 under different settings, conjecturing that this value causes the blocked mask to generate consistent global views for both of the images being mixed. Sampling masking ratio from a uniform distribution with bounds 0.25 and 0.75 yields better results than using 0 and 1 as bounds. We believe this shows that extreme values close to either 0 or 1 generate a mixture where one image heavily dominates over the other when in the optimal mixture, areas of each image should be roughly proportional.

Switch Mixture. We also examine the efficacy of the second term in the mixture loss in Eq. 4, which is computed using switch images and multiplied with soft coefficient $(1 - \lambda)$. Certainly, having two terms in the mixture loss part is beneficial and gives better results in Table 1 on all datasets as it provides more training signal.

Training Budgets. We test our method with different training budgets of 200, 400, 600, 800, and 1000 epochs on CIFAR-100 using MoCo, SimCLR, and SimSiam. The results are shown in Fig. 5. We can observe our method achieves consistent improvement over various frameworks.

Results for Different Base Frameworks. In Table 2 we consider the generalizability of our method by applying it on top of four different self-supervised learning frameworks. When applying our method upon Un-Mix, we yield superior performance in all cases. Plain MixMask also provides a competitive performance in all the experiments.

Method	Top-1 200 ep	COCO AP _{det}	COCO AP _{seg}	Top-1 1% labels
MoCo v2	67.5	39.3	34.4	50.3
MSCN	68.2	39.1	34.2	54.0
MixMask	69.2	39.8	34.8	54.5




Table 3: MixMask shows better performance than MoCo v2 on linear probing, COCO detection and Figure 4: MixMask outperforms MSCN segmentation, and 1% label supervised finetuning, on ImageNet-1K by 1%.




Figure 5: Results with different training budgets and base frameworks on CIFAR-100. MixMask consistently performs better than the baseline for every configuration.

Method	MS COCO 2017						Pascal VOC 2007		
	Object detection			Segmentation			Object detection		
	AP	AP ₅₀	AP ₇₅	AP	AP ₅₀	AP ₇₅	AP	AP ₅₀	AP ₇₅
MoCo V2 [8]	39.3	58.9	42.5	34.4	55.8	36.5	57.4	82.5	64.0
SwAV [9]	38.4	58.6	41.3	33.8	55.2	35.9	56.1	82.6	62.7
Barlow Twins [10]	39.2	59.0	42.5	34.3	56.0	36.5	56.8	82.6	63.4
MSCN [11]	39.1	59.1	42.1	34.2	55.7	36.4	57.5	83.0	64.4
MixMask	39.8	60.0	43.0	34.8	56.3	37.0	57.9	83.0	64.5

Table 4: MixMask gives the best performance on all metrics on COCO and Pascal VOC.

4.2 Comparison and Superiority to MSCN

We present evidence demonstrating that our approach, MixMask, outperforms MSCN [11]. Specifically, Table 3 shows the results of linear probing for 200 epoch pretraining and downstream evaluation for object detection/segmentation on COCO [12] and supervised fine-tuning on 1% of ImageNet-1K. MixMask achieves superior performance by proposing a method to fill the erased regions with meaningful semantics on all of the benchmarks. Additional results for semi-supervised fine-tuning and object detection are available in Table 4 and Table 5. The results demonstrate that MixMask outperforms MSCN and even surpasses the performance of SwAV in semi-supervised fine-tuning. Moreover, MixMask achieves better results than MSCN on all evaluation metrics for detection and segmentation. Notably, our approach is conceptually simpler and computationally more efficient than MSCN, which requires combining multiple masking strategies, adding Gaussian noise, and multicrops.

4.3 Results on Semi- and Supervised Finetuning

For the semi and supervised finetuning on ImageNet-1K, we follow the protocol described in [11]. We explore three different data regimes of 1%, 10%, and 100% of labels and finetune all models for 20 epochs. The results are given in Table 5, MixMask outperforms MSCN as well as MoCoV2 and SwAV on 1% data regime and shows competitive performance to SwAV and Barlow Twins on 10%.

Method	1% labels	10% labels	100% labels
	Top-1 Acc	Top-1 Acc	Top-1 Acc
MoCo V2 [8]	50.3	66.8	76.7
SwAV [9]	53.9	70.2	-
Barlow Twins [10]	55.0	69.7	-
MSCN [11]	54.0	-	-
MixMask	54.5	69.3	77.7

Table 5: Results of semi-supervised and supervised finetuning on ImageNet-1K.

4.4 Results on Object Detection and Segmentation

We test our method on the downstream task of object detection and segmentation. For that, we finetune a Faster-RCNN [13], and Mask-RCNN [14] models from Detectron2 [15] on Pascal VOC 2007 [16], and MS COCO 2017 [17] datasets. For VOC 2007, we follow the standard evaluation protocol in [18] with 24k training iterations. For COCO, we use 1× schedule from [15]. The results in Table 4 verify the superiority of the proposed MixMask to all other frameworks, including MSCN.

5 Conclusion

We have introduced MixMask, a new masking approach tailored for self-supervised learning on ConvNets. This method harmoniously merges an asymmetric loss with a filling-based masking procedure within the framework of the Masked Siamese ConvNets. Our experiments consistently demonstrate that MixMask surpasses the conventional symmetric loss and erase-centric masking techniques used in MSCN. We have demonstrated that our distinct asymmetric loss function outshines the standard symmetric loss associated with MSCN. Moreover, we addressed the challenge of diminished learning efficiency inherent in traditional erase-centric masking by adopting a filling-centric approach. In this strategy, an alternative image serves to occupy the voided regions. We have also laid out a comprehensive guideline for seamlessly infusing MixMask into other blending techniques, elaborating on how varying masking parameters influence the caliber of the resultant representations. The results unequivocally highlight our method’s superiority over MSCN in linear probing and multiple downstream applications such as object detection and segmentation.

References

- [1] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Michael Rabbat, and Nicolas Ballas. Masked siamese networks for label-efficient learning. *arXiv preprint arXiv:2204.07141*, 2022.
- [2] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.
- [3] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. 2020.
- [4] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021.
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [6] Ting Chen, Calvin Luo, and Lala Li. Intriguing properties of contrastive losses. *Advances in Neural Information Processing Systems*, 34:11834–11845, 2021.
- [7] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021.
- [8] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020.
- [9] Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *2005 IEEE Computer Society Conference*

- on *Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 539–546. IEEE, 2005.
- [10] Elijah Cole, Xuan Yang, Kimberly Wilber, Oisín Mac Aodha, and Serge Belongie. When does contrastive visual representation learning work? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14755–14764, 2022.
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [13] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010.
- [14] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations, 2018.
- [15] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [17] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [18] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020.
- [19] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022.
- [20] Li Jing, Jiachen Zhu, and Yann LeCun. Masked siamese convnets. *arXiv preprint arXiv:2206.07700*, 2022.
- [21] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [22] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.

- [23] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [24] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11976–11986, 2022.
- [25] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, pages 69–84. Springer, 2016.
- [26] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [27] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015.
- [28] Pierre H Richemond, Jean-Bastien Grill, Florent Alché, Corentin Tallec, Florian Strub, Andrew Brock, Samuel Smith, Soham De, Razvan Pascanu, Bilal Piot, et al. Byol works even without batch statistics. *arXiv preprint arXiv:2010.10241*, 2020.
- [29] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, and Sergey Levine. Time-contrastive networks: Self-supervised learning from video, 2018.
- [30] Zhiqiang Shen, Zechun Liu, Zhuang Liu, Marios Savvides, Trevor Darrell, and Eric Xing. Un-mix: Rethinking image mixtures for unsupervised visual representation learning. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2022.
- [31] Sanghyun Woo, Shoubhik Debnath, Ronghang Hu, Xinlei Chen, Zhuang Liu, In So Kweon, and Saining Xie. Convnext v2: Co-designing and scaling convnets with masked autoencoders. *arXiv preprint arXiv:2301.00808*, 2023.
- [32] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- [33] 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9653–9663, 2022.
- [34] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR, 2021.
- [35] Richard Zhang, Phillip Isola, and Alexei A. Efros. Colorful image colorization, 2016.