Masked Siamese Networks for Label-Efficient Learning

Mahmoud Assran[1,2,5*] **Mathilde Caron**[1,3] **Ishan Misra**[1] **Piotr Bojanowski**[1]
Florian Bordes[1,2,4] **Pascal Vincent**[1,2] **Armand Joulin**[1] **Michael Rabbat**[1,2] **Nicolas Ballas**[1]

[1]Facebook AI Research
[2]Mila – Quebec AI Institute
[3]Inria, Univ. Grenoble Alpes
[4]Universite de Montreal, DIRO
[5]McGill University, Dept. of Electrical and Computer Engineering

ABSTRACT

We propose Masked Siamese Networks (MSN), a self-supervised learning framework for learning image representations. Our approach matches the representation of an image view containing randomly masked patches to the representation of the original unmasked image. This self-supervised pre-training strategy is particularly scalable when applied to Vision Transformers since only the unmasked patches are processed by the network. As a result, MSNs improve the scalability of joint-embedding architectures, while producing representations of a high semantic level that perform competitively on low-shot image classification. For instance, on ImageNet-1K, with only 5,000 annotated images, our base MSN model achieves 72.4% top-1 accuracy, and with 1% of ImageNet-1K labels, we achieve 75.7% top-1 accuracy, setting a new state-of-the-art for self-supervised learning on this benchmark. Our code is publicly available at https://github.com/facebookresearch/msn.

1 Introduction

Self-Supervised Learning (SSL) has emerged as an effective strategy for unsupervised learning of image representations, eliminating the need to manually annotate vast quantities of data. By training large models on unlabeled data, SSL aims to learn representations that can be effectively applied to a downstream prediction task with few labels (Chen & He, 2020).

One of the core ideas of SSL is to remove a portion of the input and learn to predict the removed content (Pathak et al., 2016). Auto-regressive models and denoising auto-encoders instantiate this principle in vision by predicting the missing parts at the pixel or token level (Chen et al., 2020a; Vincent et al., 2010; He et al., 2021; Bao et al., 2021; Baevski et al., 2022). Masked auto-encoders in particular, which learn representations by reconstructing randomly masked patches from an input, have been successfully applied in vision (He et al., 2021; Xie et al., 2021; Wei et al., 2021; Bao et al., 2021). However, optimizing a reconstruction loss requires modelling low-level image details that are not necessary for classification tasks involving semantic abstraction. Thus, the resulting representations often need to be fine-tuned for semantic recognition tasks which can lead to overfitting in low-shot settings. Nevertheless, masked auto-encoders have enabled the training of large-scale models and demonstrated state-of-the-art performance when fine-tuning on large labeled datasets, with millions of labels (Bao et al., 2021; He et al., 2021; Xie et al., 2021; Baevski et al., 2022).

Joint-embedding architectures, on the other hand, avoid reconstruction. Approaches such as Siamese Networks (He et al., 2019; Caron et al., 2020; Chen & He, 2020; Grill et al., 2020; Caron et al., 2021; Zbontar et al., 2021; Bardes et al., 2021) learn a representation by training an encoder network to produce similar embeddings for two different views of the same image (Bromley et al., 1993; Dosovitskiy et al., 2014). Here the views are typically constructed by applying different image

---

*correspondence to massran@fb.com

1

Evaluation on 1% ImageNet-1k

Low-Shot Evaluation on ImageNet-1k
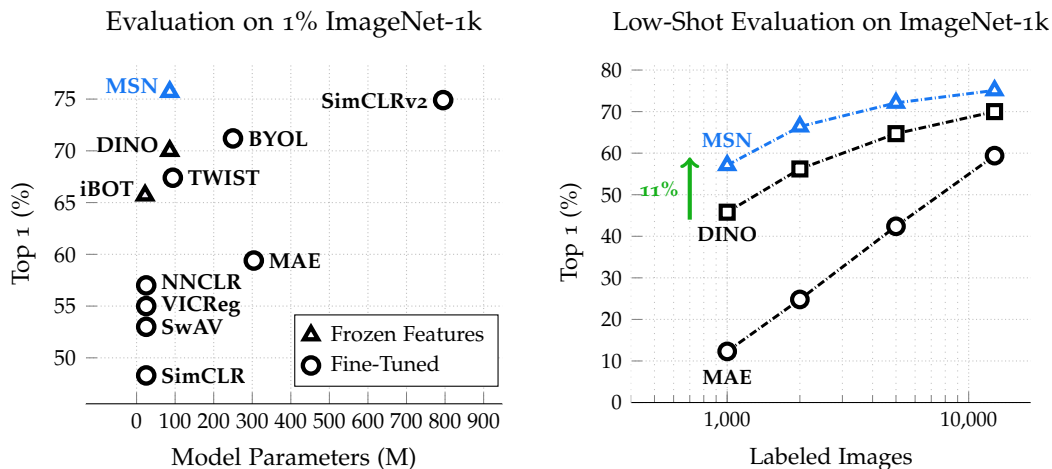
(a) Evaluation using 1% of ImageNet-1K labels (∼13 imgs/class). Evaluation with *Frozen Features* corresponds to freezing the weights and training a logistic regression classifier with the available labeled samples. Evaluation with *Fine-Tuning* corresponds to adding a linear head and fine-tuning the model+head, end-to-end.

(b) Low-shot evaluation comparing MSN (ViT-L/7) to the best publicly available models in low-shot classification for DINO (ViT-B/8) and MAE (ViT-L/16). MSN and DINO use a linear probe, whereas MAE uses partial fine-tuning, where the last block of the pre-trained model along with a linear head are adapted.

Figure 1: **Low-shot Evaluation of self-supervised models, pre-trained on ImageNet-1K.** (Left) MSN surpasses the previous 800M parameter state-of-the-art, while using a model that is $10\times$ smaller. (Right) MSN achieves good classification performance using less labels than current mask-based auto-encoders.

transforms — such as random scaling, cropping, and color jitter — to the input (Wu et al., 2018; Misra & van der Maaten, 2020). The inductive bias introduced by this invariance-based pre-training typically produces strong off-the-shelf representations of a high semantic level (Caron et al., 2021) but often disregards rich local structure that can be helpful to model.

In this work, we propose Masked Siamese Networks (MSNs), a self-supervised learning framework that leverages the idea of mask-denoising while avoiding pixel and token-level reconstruction. Figure 3 shows a schematic of the method. Given two views of an image, MSN randomly masks patches from one view while leaving the other view unchanged. The objective is to train a neural network encoder, parametrized with a vision transformer (ViT) (Dosovitskiy et al., 2020), to output similar embeddings for the two views. In this procedure, MSN does not predict the masked patches at the input level, but rather performs the denoising step implicitly at the representation level by ensuring that the representation of the masked input matches the representation of the unmasked one. Figure 2 qualitatively demonstrates the effectiveness of the MSN denoising process.

Empirically, we demonstrate that MSNs learn strong off-the-shelf representations that excel at low-shot prediction (cf. Figure 1). In particular, MSN achieves good classification performance using $100\times$ fewer labels than current mask-based auto-encoders (He et al., 2021; Xie et al., 2019). In the standard 1% ImageNet low-shot classification task, an MSN-trained ViT-B/4 (using a patch size of $4x4$ pixels) achieves 75.7% top-1 accuracy, outperforming the previous 800M parameter state-of-the-art convolutional network (Chen et al., 2020c) while using nearly $10\times$ fewer parameters (cf. Figure 1a).

Since a good representation should not need many examples to learn about a concept (Goyal et al., 2019), we also consider a more challenging evaluation benchmark for label-efficient low-shot classification (Sohn et al., 2020; Lucas et al., 2021), using from 1 labeled image per class up to 5 images per class (cf. Table 2). MSN also achieves state-of-the-art performance in that regime. For instance, with only 5 labeled images per class, we can pre-train a ViT-L/7 with MSN on ImageNet-1K to achieve 72.1% top-1 accuracy surpassing the previous state-of-the-art method, DINO (Caron et al., 2021), by 8% top-1.

Similar to masked auto-encoders, MSNs also exhibit good computational scaling since only the unmasked patches are processed by the ViT encoder. For example, by randomly masking 70% of the patches, MSN uses half the computation and memory compared to an unmasked joint-embedding

Figure 2: **Visualization of MSN representations.** First column: original image. Second column: image with 70% of patches masked, input to an MSN pre-trained ViT-L/7 encoder to compute representations. Other columns: Samples of a generative model conditioned on the MSN representations (see Appendix F for more details and other samples). Qualities that vary across samples represent information that the pre-trained representation is invariant to; e.g., in this case, MSN discards background, pose, and lighting information. Qualities that are common across samples represent information that the pre-trained representation is not invariant to. In this case, even with a large fraction of the patches corrupted with mask noise, MSN representations still encode semantic information about the object of interest.

baseline. In practice, we pre-train a ViT-L/7 on as few as 18 AWS `p4d-24xlarge` machines. Without masking, the same job requires over 42 machines.

Finally, we also show that MSNs are competitive with prior works on other self-supervised benchmarks that use many labels for evaluation (e.g., fine-tuning, linear-evaluation, transfer learning).

2   Prerequisites

Problem Formulation   Consider a large collection of unlabeled images, $\mathcal{D} = (\mathbf{x}_i)_{i=1}^{U}$, and a small dataset of annotated images, $\mathcal{S} = (\mathbf{x}_{si}, y_i)_{i=1}^{L}$, with $L \ll U$. Here, the images in $\mathcal{S}$ may overlap with the images in the dataset $\mathcal{D}$. Our goal is to learn image representations by first pre-training on $\mathcal{D}$ and then adapting the representation to the supervised task using $\mathcal{S}$.

Siamese Networks   The goal of siamese networks (Becker & Hinton, 1992; Bromley et al., 1993), as they are used in self-supervised learning, is to learn an encoder that produces similar image embeddings for two views of an image. Specifically, given an encoder $f_\theta(\cdot)$ and two views $\mathbf{x_i}$ and $\mathbf{x_i^+}$ of an image, the encoder independently processes each view and outputs representations z_i and z_i^+ respectively, referred to as the anchor representation and the target representation. The objective of siamese networks is to learn an encoder that is not sensitive to differences between views, so the representations z_i and z_i^+ should match. In practice, the encoder $f_\theta(\cdot)$ is usually parameterized as a deep neural network with learnable parameters θ.

The main challenge with siamese architectures is to prevent representation collapse in which the encoder produces a constant image embedding regardless of the input. Several approaches have been investigated in the literature. Contrastive losses explicitly push away embeddings of different
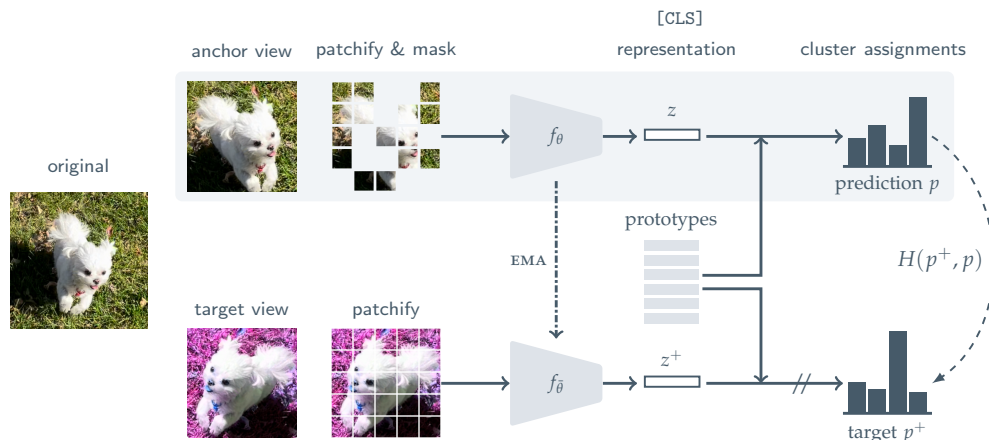
Figure 3: **Masked Siamese Networks.** First use random data augmentations to generate two views of an image, referred to as the anchor view and the target view. Subsequently, a random mask is applied to the anchor view, while the target view is left unchanged. The objective is then to assign the representation of the masked anchor view to the same clusters as the representation of the unmasked target view. A standard cross-entropy loss is used as the criterion to optimize.

images (Bromley et al., 1993; He et al., 2019; Chen & He, 2020). Information maximization approaches try to maximize the entropy of the average prediction (Caron et al., 2021; Assran et al., 2021) or spread out the embeddings uniformly on the surface of a sphere (Caron et al., 2020). Asymmetric approaches rely on an asymmetric architectural choice such as stop-gradient operations and a momentum encoder (Chen & He, 2020; Grill et al., 2020) to prevent collapse. Other approaches try to decorrelate the vector components of the embeddings to minimize redundancy across samples (Zbontar et al., 2021; Bardes et al., 2021).

Vision Transformer We use a standard Vision Transformer (ViT) architecture (Dosovitskiy et al., 2020) as the encoder. Vision Transformers first extract a sequence of non-overlapping patches of resolution $N \times N$ from an image. Next, they apply a linear layer to extract patch tokens, and subsequently add learnable positional embeddings to them. An extra learnable [CLS] token is added to the sequence. This token aims to aggregate information from the full sequence of patches (Dosovitskiy et al., 2020; Caron et al., 2021). The sequence of tokens is then fed to a stack of Transformer layers (Vaswani et al., 2017). A Transformer layer is composed of a self-attention (Vaswani et al., 2017) and a fully-connected layer with skip connections (He et al., 2016). Self-attention uses an attention mechanism (Bahdanau et al., 2014) applied to the entire sequence of elements to update the representation. The output representation associated to the [CLS] token is used as the output of the encoder.

3 Masked Siamese Networks

We now describe the proposed Masked Siamese Network (MSN) training procedure, which combines invariance-based pre-training with mask denoising; see Figure 3 for a schematic. MSNs first use random data augmentations to generate two views of an image, referred to as the anchor view and the target view. Subsequently, a random mask is applied to the anchor view, while the target view is left unchanged. Similar to clustering-based SSL approaches (Caron et al., 2020; 2021; Assran et al., 2021), learning occurs by computing a soft-distribution over a set of prototypes for both the anchor and target views. The objective is then to assign the representation of the masked anchor view to the same prototypes as the representation of the unmasked target view. We use a standard cross-entropy loss to optimize this criterion.

In contrast to previous work on masked image modelling, the mask-denoising process in MSN is discriminative, rather than generative (He et al., 2021; Xie et al., 2021; Wei et al., 2021; Bao et al., 2021; Zhou et al., 2021). MSN architectures do not directly predict pixel values (or tokens) for the

| (a) No Mask | (b) Random Mask | (c) Focal Mask |

Figure 4: **Masking strategies.** When applying a Random Mask, we randomly drop patches across a global view of the image. When applying a Focal Mask, we randomly select a local continuous block of an image, and mask everything around it. We typically leverage both Random and Focal Masking strategies when pre-training with MSNs.

masked patches. Instead, the loss is applied directly to the output corresponding to the [CLS] token of the encoder.

Input Views   In each iteration of pre-training, we sample a mini-batch of $B \geq 1$ images. For an index $i \in [B]$, let $\mathbf{x}_i$ denote the i^{th} image in the mini-batch. For each image $\mathbf{x}_i$, we first apply a random set of data augmentations to generate a target view, denoted $\mathbf{x}_i^+$, and $M \geq 1$ anchor views, denoted $\mathbf{x}_{i,1}, \mathbf{x}_{i,2}, \ldots, \mathbf{x}_{i,M}$.

Patchify and Mask   Next, we "patchify" each view by converting it into a sequence of non-overlapping $N \times N$ patches. After patchifying the anchor view $\mathbf{x}_{i,m}$, we also apply the additional step of masking by randomly dropping some of the patches. We denote by $\hat{\mathbf{x}}_{i,m}$ the sequence of masked anchor patches, and by $\hat{\mathbf{x}}_i^+$ the sequence of unmasked target patches. Because of masking, the anchor sequence $\hat{\mathbf{x}}_{i,m}$ can have a different length than the patchified target sequence $\hat{\mathbf{x}}_i^+$, even if both image views originally have the same resolution.

We investigate two strategies for masking the anchor views, Random Masking and Focal Masking, which are depicted in Figure 4. When applying Random Masking, we randomly drop potentially non-contiguous patches across the sequence. Conversely, when applying Focal Masking, we randomly select a local continuous block of the anchor view and drop all the patches around it.

Encoder   Given a parameterized anchor encoder, denoted $f_\theta(\cdot)$, let $z_{i,m} \in \mathbb{R}^d$ denote the representation computed from the patchified (and masked) anchor view $\hat{\mathbf{x}}_{i,m}$. Similarly, given a parameterized target encoder $f_{\bar{\theta}}(\cdot)$, with a potentially different set of parameters $\bar{\theta}$, let $z_i^+ \in \mathbb{R}^d$ denote the representation computed from the patchified target view $\hat{\mathbf{x}}_i^+$. In MSNs, the parameters $\bar{\theta}$ of the target encoder are updated via an exponential moving average of the anchor encoder parameters (Grill et al., 2020). Both encoders correspond to the trunk of a ViT (Dosovitskiy et al., 2020). We take the output of the network to be the representation corresponding to the [CLS] token.

Similarity Metric and Predictions   Let $\mathbf{q} \in \mathbb{R}^{K \times d}$ denote $K > 1$ learnable prototypes, each of dimension d. To train the encoder, we compute a distribution based on the similarity between these prototypes and each anchor and target view pair, and we penalize the encoder for differences between these distributions. More precisely, for an anchor representation $z_{i,m}$, we compute a "prediction" $p_{i,m} \in \Delta_K$ in the K-dimensional simplex by measuring the cosine similarity to the prototypes matrix $\mathbf{q}$. For L_2-normalized representations and prototypes, the predictions $p_{i,m}$ can be concisely written as

$$p_{i,m} := \text{softmax}\left(\frac{z_{i,m} \cdot \mathbf{q}}{\tau}\right),$$

where $\tau \in (0,1)$ is a temperature. Similarly, for each target representation z_i^+, we generate a prediction $p_i^+ \in \Delta_K$ by measuring the cosine similarity to the same prototypes matrix $\mathbf{q}$. When computing the target predictions, we also use a temperature parameter $\tau^+ \in (0,1)$. Note, we always choose $\tau^+ < \tau$ to encourage sharper target predictions, which implicitly guides the model to produce confident low entropy anchor predictions. As shown in Appendix B, target sharpening coupled with with other regularization like mean-entropy maximization (see below) is provably sufficient to

eliminate collapsing solutions in the MSN framework. Empirically, we have observed that training without sharpening can result in collapsing solutions.

Training Objective   As previously mentioned, to train the encoder, we penalize when the anchor prediction $p_{i,m}$ is different from the target prediction p_i^+. We enforce this criterion using a standard cross-entropy loss $H(p_i^+, p_{i,m})$.

We also incorporate the mean entropy maximization (ME-MAX) regularizer, also used in (Assran et al., 2021; Joulin & Bach, 2012), to encourage the model to utilize the full set of prototypes. Denote the average prediction across all the anchor views by

$$\overline{p} := \frac{1}{MB} \sum_{i=1}^{B} \sum_{m=1}^{M} p_{i,m}.$$

The ME-MAX regularizer simply seeks to maximize the entropy of $\overline{p}$, denoted $H(\overline{p})$, or equivalently, minimize the negative entropy of $\overline{p}$. Thus, the overall objective to be minimized when training the encoder parameters θ and prototypes $\mathbf{q}$ is

$$\frac{1}{MB} \sum_{i=1}^{B} \sum_{m=1}^{M} H(p_i^+, p_{i,m}) - \lambda H(\overline{p}), \tag{1}$$

where $\lambda > 0$ controls the weight of the ME-MAX regularization. Note that when training, we only compute gradients with respect to the anchor predictions $p_{i,m}$, not the target predictions p_i^+.

4   Related Work

Unsupervised pre-training for vision has seen rapid progress with the development of view-invariant representation learning and joint embedding architectures (Wu et al., 2018; He et al., 2019; Chen & He, 2020; Grill et al., 2020; Caron et al., 2021; Bardes et al., 2021). Most similar to our approach is DINO (Caron et al., 2021) which leverages a Siamese Network with a cross-entropy loss and a momentum encoder. DINO also uses multi-crop training, which is a form of focal masking, but it requires an unmasked anchor view during training. MSN can be seen as a generalization of DINO, leveraging both random and focal masking without requiring any unmasked anchor views. Since the cross-entropy loss in equation equation 1 is only differentiated with respect to the anchor predictions, not the target, MSN only backpropagates through the anchor network and only needs to store the activation associated with the masked view. MSN therefore reduces the computational and memory requirements. MSN also differs from DINO in its mechanism for preventing representation collapse (entropy maximization as opposed to centering and sharpening). Our empirical results show that MSN compares favourably to DINO across various degrees of supervision for the downstream task.

A prominent line of work in SSL is to remove a portion of the input and learn to reconstruct the removed content (Devlin et al., 2018). For example, in the field of image recognition, some works have proposed to predict augmented image channels (Zhang et al., 2017b), which can be regarded as a form of image colorization (Zhang et al., 2016; Larsson et al., 2016; 2017). Other approaches propose to remove and learn to regress entire image regions: the seminal Context Encoders of Pathak et al. (2016) train a network to generate missing image patches based on their surroundings. Recent works revisit this idea and investigate the pre-training of ViTs with masked auto-encoders (Chen et al., 2020a; He et al., 2021; Xie et al., 2021; Wei et al., 2021; Bao et al., 2021). These approaches corrupt images with mask-noise and predict missing input values at the pixel level (Dosovitskiy et al., 2020; He et al., 2021; Xie et al., 2019) or using a tokenizer (Bao et al., 2021; Wei et al., 2021). Our approach does not predict the missing value at the input level, but instead performs the denoising step implicitly by ensuring that the global representation of the noisy input matches that of the uncorrupted input.

Some recent approaches have started to explore the combination of joint-embedding architectures and denoising pre-training tasks (El-Nouby et al., 2021; Baevski et al., 2022; Zhou et al., 2021). Those approaches mask an image by replacing the masked patches with a learnable mask token, and output a single vector for each masked patch. The objective is then to directly match each computed patch vector to the equivalent patch token extracted from a target encoder. In addition to the patch-level

Table 1: **Extreme low-shot.** We evaluate the label-efficiency of self-supervised models pretrained on the ImageNet-1K dataset. For evaluation, we use an extremely small number of the ImageNet-1K labels and report the mean top-1 accuracy and standard deviation across 3 random splits of the data.

			Images per Class		
Method	**Architecture**	**Epochs**	1	2	5
iBOT (Zhou et al., 2021)	ViT-S/16	800	40.4 ± 0.5	50.8 ± 0.8	59.9 ± 0.2
	ViT-B/16	400	46.1 ± 0.3	56.2 ± 0.7	64.7 ± 0.3
DINO (Caron et al., 2021)	ViT-S/16	800	38.9 ± 0.4	48.9 ± 0.3	58.5 ± 0.1
	ViT-B/16	400	41.8 ± 0.3	51.9 ± 0.6	61.4 ± 0.2
	ViT-S/8	800	45.5 ± 0.4	56.0 ± 0.7	64.7 ± 0.4
	ViT-B/8	300	45.8 ± 0.5	55.9 ± 0.6	64.6 ± 0.2
MAE (He et al., 2021)	ViT-B/16	1600	8.2 ± 0.3	25.0 ± 0.3	40.5 ± 0.2
	ViT-L/16	1600	12.3 ± 0.2	19.3 ± 1.8	42.3 ± 0.3
	ViT-H/14	1600	11.6 ± 0.4	18.6 ± 0.2	32.8 ± 0.2
MSN (Ours)	ViT-S/16	800	47.1 ± 0.1	55.8 ± 0.6	62.8 ± 0.3
	ViT-B/16	600	49.8 ± 0.2	58.9 ± 0.4	65.5 ± 0.3
	ViT-B/8	600	55.1 ± 0.1	64.9 ± 0.7	71.6 ± 0.3
	ViT-B/4	300	54.3 ± 0.4	64.6 ± 0.7	$\mathbf{72.4 \pm 0.3}$
	ViT-L/7	200	$\mathbf{57.1 \pm 0.6}$	$\mathbf{66.4 \pm 0.6}$	$\mathbf{72.1 \pm 0.2}$

loss, iBOT (Zhou et al., 2021) and SplitMask (El-Nouby et al., 2021) apply a joint-embedding loss to an output representing the global sequence (either the `[CLS]` token or a global average pool of the patch vectors). SplitMask shows that by using a patch-level loss, you can reduce the amount of unlabeled pre-training data. In contrast, we focus on reducing the amount of labeled data available for the downstream prediction task. Data2Vec (Baevski et al., 2022) demonstrates that this approach is suitable for multiple modalities such as vision, speech and text. Different from these approaches, we only match the view representations globally and do not consider a patch level loss. Consequently, we can completely ignore the masked patches, significantly reducing the computational and memory requirements. For example, when training our largest model, a ViT-L/7, we mask over 70% of the input patches, and reduce memory and computational overhead by half.

5 Results

We evaluate MSN representations learned on the ImageNet-1K dataset (Russakovsky et al., 2015). We first consider low-shot evaluation on ImageNet-1K using as few as 1–5 images per class. We also compare with the state-of-the-art in settings where more supervision is available and investigate transfer-learning performance. Finally, we conduct ablation experiments with MSN. By default, we pre-train with a batch-size of 1024 images, generating several anchor views from each image: 1 view with a random mask, and 10 views with focal masks. We find that the optimal masking ratio is model-dependent, with larger models benefiting from more aggressive patch dropping. We describe MSN implementation details in Appendix A.

5.1 Label-Efficient Learning

The premise of SSL is to learn representations on unlabeled data that can be effectively applied to prediction tasks with few labels (Chen et al., 2020c). In this section we explore the performance of self-supervised approaches when very few labeled examples are available.

Extreme Low-Shot We first evaluate the classification performance of unsupervised models that have been pre-trained on ImageNet-1K, by using 1, 2, and 5 labeled images per class for supervised evaluation. We compare MSN to the joint-embedding approach, DINO (Caron et al., 2021), the auto-encoding approach, MAE (He et al., 2021), and the hybrid approach, iBOT (Zhou et al., 2021),

Table 2: Low-shot evaluation on ImageNet-1K using 1% of the labels (approximately 13 images per class). [†]Indicates evaluations we computed using publicly available models.

Method	Architecture	Params.	Top 1
Comparing similar architectures			
Barlow-Tw. (Zbontar et al., 2021)	RN50	24M	55.0
SimCLRv2 (Chen et al., 2020c)	RN50	24M	57.9
PAWS (Assran et al., 2021)	RN50	24M	66.5
DINO (Caron et al., 2021)	ViT-S/16	22M	64.5
iBOT (Zhou et al., 2021)	ViT-S/16	22M	65.9
MSN	ViT-S/16	22M	**67.2**
Comparing larger architectures			
BYOL (Grill et al., 2020)	RN200 (2×)	250M	71.2
SimCLRv2 (Chen et al., 2020c)	RN151+SK (3×)	795M	74.9
iBOT (Zhou et al., 2021)[†]	ViT-B/16	86M	69.7
DINO (Caron et al., 2021)[†]	ViT-B/8	86M	70.0
MSN	ViT-L/7	304M	75.1
	ViT-B/4	86M	**75.7**

which combines a joint-embedding architecture with a token-based patch-level loss. We download the official released models of each related approach for evaluation.

To adapt the joint-embeddings models to the supervised task, we freeze the weights of the pre-trained model and train a linear classifier on top using 1, 2 or 5 labeled samples (see Appendix A). For MAE, we rely on partial fine-tuning (He et al., 2021), except for the 1 image per class setting, and all results with the ViT-H/14 architecture, which use a linear classifier. Partial fine-tuning corresponds to fine-tuning the last block of the pre-trained model along with a linear head. MAE benefits from partial fine-tuning, but for sufficiently large models, such as the ViT-H/14, this leads to significant overfitting in the low-shot regime. We compare both protocols in more detail in Appendix C.

Table 1 reports the extreme low-shot evaluation results. MSN outperforms the other representation learning approaches across all levels of supervision. Moreover, the improvement offered by MSN increases as the amount of available labeled data is decreased. The performance of MSN also benefits from increased model size — settings with less labeled data appear to benefit more from increased model depth and smaller patch sizes.

We also observe that joint-embedding approaches appear to be more robust to the limited availability of downstream supervision than reconstruction-based auto-encoding approaches. To explain this observation, we refer to the Masked Auto-Encoders paper (He et al., 2021) which conjectures that using a pixel reconstruction loss results in encoder representations of a lower semantic level than other methods. Conversely, the inductive bias introduced by invariance-based pre-training appears to be helpful in the low-shot regime.

1% ImageNet-1K Table 2 reports a comparison on the 1% ImageNet-1K task, which is a standard benchmark for low-shot evaluation of self-supervised models (Chen et al., 2020b). For reference, the best reported result in the literature on 1% labeled data is 76.6%, achieved with a multi-stage semi-supervised pipeline, i.e., self-distilling from a fine-tuned ResNet-152 with 3× wider channels and selective kernels (Chen et al., 2020c). Here we focus on comparing to other models trained in a self-supervised setting. Our best MSN model using a ViT-B/4 achieves 75.7% top 1 accuracy, surpassing the previous 800M parameter state-of-the-art convolutional network (Chen et al., 2020c) while using significantly fewer parameters and no fine-tuning. When focusing the comparison on similar architectures (models with similar FLOP counts), MSN also consistently improves upon previous approaches.

Table 3: Linear evaluation on ImageNet-1K using 100% of the labels.

Method	Architecture	Params.	Epochs	Top 1
Comparing similar architectures				
SimCLRv2 (Chen et al., 2020c)	RN50	24M	800	71.7
BYOL (Grill et al., 2020)	RN50	24M	1000	74.4
DINO (Caron et al., 2021)	ViT-S/16	22M	800	77.0
iBOT (Zhou et al., 2021)	ViT-S/16	22M	800	**77.9**
MSN	ViT-S/16	22M	600	76.9
Comparing larger architectures				
MAE (He et al., 2021)	ViT-H/14	632M	1600	76.6
BYOL (Grill et al., 2020)	RN200 (2×)	250M	800	79.6
SimCLRv2 (Chen et al., 2020c)	RN151+SK (3×)	795M	800	79.8
iBOT (Zhou et al., 2021)	ViT-B/16	86M	400	79.4
DINO (Caron et al., 2021)	ViT-B/8	86M	300	80.1
MoCov3 (Chen et al., 2021)	ViT-BN-L/7	304M	300	**81.0**
MSN	ViT-L/7	304M	200	**80.7**

Table 4: End-to-end fine-tuning of a ViT-B/16 encoder on ImageNet-1K using 100% of the labels. MSN obtains competitive performance with both joint-embedding approaches and auto-encoding approaches.

Initialization	Pretrain Epochs	Top 1
DINO (Caron et al., 2021)	800	83.6
BEiT (Bao et al., 2021)	800	83.2
iBOT (He et al., 2021)	800	83.8
MAE (He et al., 2021)	1600	83.6
SimMIM (Xie et al., 2021)	-	83.8
MaskFeat (Wei et al., 2021)	-	84.0
Data2Vec (Baevski et al., 2022)	800	**84.2**
MSN	600	83.4

5.2 Linear Evaluation and Fine-tuning

In this section we compare with the state-of-the-art on standard evaluation benchmarks where more supervised samples are available to adapt the representation. We use the full ImageNet-1K training images with 1.28M labels.

Linear Evaluation We evaluate self-supervised pretrained models by freezing their weights and training a linear classifier. Table 3 reports the linear evaluation results on ImageNet-1K. We observe that MSN performs competitively with the state-of-the-art. The best MSN model achieves 80.7% top-1 accuracy.

Fine-Tuning In this evaluation setting, we finetune all the weights of the self-supervised model using all the labels from the ImageNet-1K training set. We focus on the ViT-B/16 architecture. We adopt the same fine-tuning protocol as (Bao et al., 2021), and provide the details in Appendix A. Table 4 reports the comparison with fine-tuning evaluation using 100% labels on ImageNet-1K. MSN is competitive with joint-embedding approaches, such as DINO, and generative auto-encoding approaches, such as MAE.

5.3 Transfer Learning

We also report transfer learning experiments on the CIFAR10, CIFAR100 and iNaturalist datasets in Tables 5 and 6 when using a self-supervised ViT-B/16 pre-trained on ImageNet-1K. Across all tasks and various levels of supervision MSN either outperforms or achieves similar results to DINO pre-

training. Recall that MSN pre-training is also less computationally expensive than DINO pre-training due to the anchor masking.

Table 5: **Fine-Tuning Transfer Learning** with a ViT-Base/16 pre-trained on ImageNet-1K. Across all tasks, MSN either outperforms or achieves similar results to DINO pre-training. The MSN model is trained with a masking ratio of 0.3; i.e., dropping 30% of patches, and thus reduces the computational cost of pre-training relative to DINO.

| | **Top 1** | | | |
Method	CIFAR10	CIFAR100	iNat18	iNat19
DINO	99.0	90.5	72.0	78.2
MSN	99.0	90.5	72.1	78.1

Table 6: **Linear Eval. Transfer Learning** with a ViT-Base/16 pre-trained on ImageNet-1K. Across both tasks and various levels of supervision, MSN either outperforms or achieves similar results to DINO pre-training. The MSN model is trained with a masking ratio of 0.3; i.e., dropping 30% of patches, and thus reduces the computational cost of pre-training relative to DINO.

| | **Top 1** | | |
| | CIFAR10 | | CIFAR100 |
Method	4000 labels	50000 labels	
DINO	93.2	95.3	82.9
MSN	93.8	95.7	82.8

5.4 Ablations

We now conduct a series of experiments to gain insights into the important design decisions used in MSN such as the masking strategy and the data augmentation strategy. We measure the accuracy of the models by training a logistic regression classifier on the frozen trunk using 1% of ImageNet-1K labels (∼13 imgs/class).

Combining Random and Focal Masking In MSN we apply both random and focal masking to the anchor views. Focal masking corresponds to selecting a small crop from the anchor view. Random masking corresponds to randomly dropping potentially non-contiguous patches from the anchor view.

Table 7: **Masking strategy.** Impact of masking strategy on low-shot accuracy (1% of ImageNet-1K labels) of a ViT-B/16. We only generate one anchor view of each image, except in the last row, where we generate two views, one with a Random Mask and one with a Focal Mask. A random masking ratio of 0.5 is used. Applying a random mask to the anchor view is better than applying no mask. By combining both random and focal masking strategies, we obtain the strongest performance.

Anchor View	**Top 1**
No Mask	49.3
Focal Mask	39.3
Random Mask	52.3
Random Mask + Focal Mask	**59.8**

Table 7 reports the effect on low-shot evaluation when using a) No Masking, b) Focal Masking, c) Random Masking, or d) Random and Focal Masking. Applying a random mask to the anchor view is always better than applying no mask. By contrast, applying only a focal mask degrades the performance, which highlights the importance of maintaining a global view during pre-training. By combining both random and focal masking strategies,we obtain the strongest performance.

Random Masking Ratio Here we explore the relationship between the optimal masking ratio and the model size. Table 8 reports the low-shot learning performance for various random masking ratios as we increase the model size.[1]

Table 8: **Masking ratio.** Impact of pre-training random masking ratio (fraction of randomly dropped patches in each random mask) on ImageNet 1% accuracy. Accuracy of larger models improves when leveraging aggressive masking during pre-training.

	Top 1			
	Random Masking Ratio			
Architecture	0.15	0.3	0.5	0.7
ViT-S/16	**66.3**	66.0	64.8	–
ViT-B/16	68.8	**69.6**	–	–
ViT-L/16	NaN	NaN	**70.1**	69.4

When increasing the model size, we find that increasing the masking ratio (dropping more patches) is helpful for improving low-shot performance. We also find that the ViT-L/16 runs with weak masking are unstable, while the runs with more aggressive masking are quite stable. However, we do not have sufficient evidence to claim that increasing the masking ratio always improves the stability of large ViT pre-training.

Augmentation Invariance and Low-Shot Learning We explore the importance of data-augmentation invariance for low-shot learning. We pretrain a ViT-B/16 with MSN, where the teacher and anchor networks either share the input image view or use different input views; in both cases, the anchor view is always masked. The views are constructed by applying random ColorJitter, Crop, Horizontal Flips, and GaussianBlur to the input image.

Table 9: Impact of view-sharing during pre-training on low-shot accuracy (1% of ImageNet-1K labels) of a ViT-B/16. The target view is constructed by applying random ColorJitter, Crop, Horizontal Flips, and GaussianBlur to the input image. When using the same image view, MSN finds a shortcut solution. Using color jitter prevents this pathological behaviour. Randomly applying additional geometric data transformations to the anchor further improves performance, demonstrating the importance of view invariance in the low-shot setting.

Anchor View Generation	**Top 1**
Target View	7.0
Target View + ColorJitter	48.7
Target View + ColorJitter + Crop + Flip + GaussianBlur	**52.3**

Table 9 reports top-1 accuracy when evaluating with 1% of ImageNet-1K labels. Sharing the view leads to a top-1 accuracy of 7%; MSN finds a shortcut solution relying on color statistics. Using different colors in the input views resolves this pathological behaviour and achieves a top-1 of 48.3%. Further applying the geometric data-augmentations independently to the two views (as opposed to sharing views) further improves the performance to 52.3%, showing the importance of learning view-invariant representations in the low-shot setting.

Random Masking Compute and Memory We look at the effect of the random masking ratio, i.e., the fraction of dropped patches from the global anchor view, on the computational requirements of large model pre-training. In each iteration we also generate 10 focal views (small crops) of each input image; the random masking ratio has no impact on these views.

Table 10 reports the memory consumption and throughput (imgs/s) of a ViT-L/7 model on a single AWS `p4d-24xlarge` machine using a batch-size of 2 images per GPU. As expected, using more aggressive masking of the global view progressively reduces device memory utilization and speeds up training. For example, by randomly masking 70% of the patches, we can use MSN to pre-train a

---

[1]Note that the performance of the ViT-S/16 can be improved by removing the Sinkhorn normalization, as we do in Table 2, however for consistency of evaluation with other models, we keep it in for this this ablation.

Table 10: Impact of random masking ratio on GPU memory usage and runtime when pre-training a ViT-L/7. Measurements are conducted on a single AWS `p4d-24xlarge` machine, containing 8 A100 GPUs, using a batch-size of 2 images per GPU. In each iteration we also generate 10 focal views (small crops) of each input image; the random masking ratio has no impact on these views. Using more aggressive masking of the global view progressively reduces device memory utilization and speeds up training.

Masking Ratio	Mem./GPU	Throughput
0.0	26G	415 imgs/s
0.3	21G	480 imgs/s
0.5	18G	525 imgs/s
0.7	17G	600 imgs/s

full-precision ViT-Large with a patch-size of 7×7 on as few as 18 AWS `p4d-24xlarge` machines. Without masking, the same job requires over 42 machines when using the default batch-size of 1024 images.

6   Conclusion

We propose Masked Siamese Networks (MSNs), a self-supervised learning framework that leverages the idea of mask-denoising while avoiding pixel and token-level reconstruction. We demonstrate empirically that MSNs learn strong off-the-shelf representations that excel at label-efficient learning, while simultaneously improving the scalability of joint-embedding architectures. By relying on view-invariant representation learning, MSN does require the specification of data transformations, and it may be that the optimal transformations and invariances are dataset and task dependant. In future work, we plan to explore more flexible mechanisms to learn those transformations and also explore the use of equivariant representations.

References

Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Armand Joulin, Nicolas Ballas, and Michael Rabbat. Semi-supervised learning of visual features by non-parametrically predicting view assignments with support samples. In *ICCV*, 2021.

Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli. Data2vec: A general framework for self-supervised learning in speech, vision and language. *arXiv preprint arXiv:2202.03555*, 2022.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.

Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.

Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.

Suzanna Becker and Geoffrey E Hinton. Self-organizing neural network that discovers surfaces in random-dot stereograms. *Nature*, 355(6356):161–163, 1992.

Florian Bordes, Randall Balestriero, and Pascal Vincent. High fidelity visualization of what your self-supervised representation knows about. *arXiv preprint arXiv:2112.09164*, 2021.

Jane Bromley, James W Bentz, Léon Bottou, Isabelle Guyon, Yann LeCun, Cliff Moore, Eduard Säckinger, and Roopak Shah. Signature verification using a "siamese" time delay neural network. *International Journal of Pattern Recognition and Artificial Intelligence*, 7(04):669–688, 1993.

Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020.

Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021.

Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pretraining from pixels. In *International Conference on Machine Learning*, pp. 1691–1703. PMLR, 2020a.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *preprint arXiv:2002.05709*, 2020b.

Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029*, 2020c.

Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. *arXiv preprint arXiv:2011.10566*, 2020.

Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. *arXiv preprint arXiv:2104.02057*, 2021.

Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *CVPR*, 2019.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

Alexey Dosovitskiy, Jost Tobias Springenberg, Martin A. Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. *CoRR*, 2014.

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.

Alaaeldin El-Nouby, Gautier Izacard, Hugo Touvron, Ivan Laptev, Hervé Jegou, and Edouard Grave. Are large-scale datasets necessary for self-supervised pre-training? *arXiv preprint arXiv:2112.10740*, 2021.

Priya Goyal, Dhruv Mahajan, Abhinav Gupta, and Ishan Misra. Scaling and benchmarking self-supervised visual representation learning. In *ICCV*, 2019.

Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. In *NeurIPS*, 2020.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*, 2019.

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.

Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*, 2019.

Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. The many faces of robustness: A critical analysis of out-of-distribution generalization. *ICCV*, 2021a.

Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. *CVPR*, 2021b.

Armand Joulin and Francis Bach. A convex relaxation for weakly supervised classifiers. *arXiv preprint arXiv:1206.6413*, 2012.

Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Learning representations for automatic colorization. In *ECCV*, 2016.

Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Colorization as a proxy task for visual understanding. In *CVPR*, 2017.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

Thomas Lucas, Philippe Weinzaepfel, and Gregory Rogez. Barely-supervised learning: semi-supervised learning with very few labeled images. *preprint arXiv:2112.12004*, 2021.

Julien Mairal. Cyanure: An open-source toolbox for empirical risk minimization for python, c++, and soon more. *arXiv preprint arXiv:1912.08165*, 2019.

Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *CVPR*, 2020.

Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016.

Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.

Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *arXiv preprint arXiv:2001.07685*, 2020.

Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pp. 10347–10357. PMLR, 2021.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.

Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, Pierre-Antoine Manzagol, and Léon Bottou. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research*, 11(12), 2010.

Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P Xing. Learning robust global representations by penalizing local predictive power. In *Advances in Neural Information Processing Systems*, pp. 10506–10518, 2019.

Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. *arXiv preprint arXiv:2112.09133*, 2021.

Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018.

Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. Unsupervised data augmentation. *arXiv preprint arXiv:1904.12848*, 2019.

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.

Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019.

Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. *arXiv preprint arXiv:2103.03230*, 2021.

Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017a.

Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, 2016.

Richard Zhang, Phillip Isola, and Alexei A Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *CVPR*, 2017b.

Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. Ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021.

A Implementations Details

In this appendix section we provide the implementation details for MSN pre-training and evaluation.

A.1 MSN Pre-training

We adopt similar hyper-parameter settings that have previously been reported in the self-supervised literature for training Vision Transformers (Caron et al., 2021; Chen & He, 2020). Specifically, for pre-training, we use the AdamW optimizer (Loshchilov & Hutter, 2017) with a batch-size of 1024. We linearly warm up the learning-rate from 0.0002 to 0.001 during the first 15 epochs, and decay it following a cosine schedule thereafter. To construct the different image views, we apply the SimCLR data augmentations of Chen et al. (2020b) to each sampled image; namely random crop, horizontal flip, color distortion, and Gaussian blur. For each sampled image, we generate one large anchor view of size 224×224 pixels, and apply a random mask with a pre-specified masking ratio (0.15 for the ViT-S/16, 0.3 for the ViT-B/16 and ViT-B/8, and 0.7 for the ViT-L/7 and the ViT-B/4). For each sampled image, we also generate 10 small focal anchor views of size 96×96 pixels. We use a temperature of 0.1 for the anchor network, and a temperature of 0.025 for the target network. Following the DINO method of Caron et al. (2021), we update the target network via an exponential moving average of the anchor network with a momentum value of 0.996, and linearly increase this value to 1.0 by the end of training. Similarly, following Caron et al. (2021), weight decay is set to 0.04 and increased to 0.4 throughout training via a cosine schedule. By default, we set the ME-MAX regularization weight λ to 1.0 and apply Sinkhorn normalization to the targets (Caron et al., 2020) to avoid having to tune the ME-MAX regularization weight; however, in general, we observe stronger MSN performance when omitting Sinkhorn normalization (see Appendix C). We train with a 3-layer projection head with output dimension 256 and batch-normalization at the input and hidden layers, and use 1024 prototypes of dimension 256. We observe that using more prototypes has little effect on training, but using too few prototypes can hurt performance (see Appendix C). We discard the projection head during evaluation, and always use the representations computed from the output of the target encoder trunk for evaluation.

A.2 Low-Shot Evaluation

To avoid overfitting, we freeze the weights of the pre-trained model and train a linear classifier on top using 1, 2 or 5 labeled samples per class. Specifically, we take a single center crop of each labeled image, extract its representation using the pre-trained model, and then train a classifier on these representations using L_2-regularized logistic regression. Following (Caron et al., 2021), we use the `cyanure` package (Mairal, 2019) to run logistic regression on the extracted representations. This objective is smooth and strongly-convex (i.e., has a unique minimizer) and can therefore be efficiently solved for using the `cyanure` python numerical solver on a single CPU core. All low-shot evaluations (including the 1% ImageNet-1K evaluation) are computed with this procedure, except for models pre-trained using MAE (He et al., 2021), which benefit from using partial fine-tuning (He et al., 2021).

Partial fine-tuning corresponds to fine-tuning the last block of the pre-trained model along with a linear head. MAE benefits from partial fine-tuning, but for sufficiently large models, such as the ViT-H/14, this leads to significant overfitting in the low-shot regime. Our results in Table 2 and Figure 1 report the best performance across evaluation methods for MAE. In particular, all the MAE results are obtained via partial fine-tuning, except for the 1 image per class setting, and all results with the ViT-H/14 architecture, which use a linear head. We compare both protocols in more detail in Appendix C.

A.3 Linear Evaluation

For linear evaluation, we use a similar procedure as He et al. (2021). Specifically, we use a large batch-size of 16,384 images and train a linear classifier for 100 epochs using a learning rate of 6.4, and decay it following a cosine schedule. We only apply basic data augmentations; namely, random resized crops to a resolution of 224×224 pixels, and random horizontal flips. We also L_2-normalize the representations before feeding them into the linear classifier, and optimize the classifier weights using SGD with Nesterov momentum. We do not apply any weight-decay and do not use any warmup.

A.4 Fine-Tuning Evaluation

We follow the common practice for fine-tuning SSL pre-trained ViT models. Specifically, we follow the setup of (Touvron et al., 2021; Bao et al., 2021; He et al., 2021). We fine-tune a pre-trained ViT model for 100 epochs on the full supervised ImageNet-1K training data set using the AdamW (Loshchilov & Hutter, 2017) optimizer. We use a batch size of 1024 with a learning rate of 0.002. The learning rate is linearly warmed-up during the first 5 epochs and decayed with a cosine schedule thereafter. A layer-wise decay of 0.65 is also applied, along with the data augmentations defined by RandAugment(9, 0.5) (Cubuk et al., 2019). We additionally use label smoothing set to 0.1, mixup (Zhang et al., 2017a) set to 0.8, cutmix (Yun et al., 2019) set to 1.0, and drop path set to 0.2.

A.5 Transfer Learning

A.5.1 Linear Evaluation

When performing linear evaluation for transfer learning, we freeze the weights of the ImageNet-1K pre-trained model and optimize a linear classifier on top. We resize each downstream image to 256×256 pixels, and take a single center crop of size 224×224 pixels. Next, we extract a representation of each image using the pre-trained model, and subsequently train a classifier on top using L_2-regularized logistic regression.

A.5.2 Fine Tuning

When performing end-to-end fine-tuning for transfer learning, we follow the protocol of DeiT and DINO (Touvron et al., 2021; Caron et al., 2021). Models transferred to CIFAR10 and CIFAR100 are fine-tuned for 1000 epochs using a batch size of 768 and a learning rate of 0.000075. Models transferred to iNat18 and iNat19 models are fine-tuned for 300 epochs using a batch size of 1024 and a learning of 0.0001. All transfer fine-tuning experiments use the data augmentations defined by RandAugment(9, 0.5) (Cubuk et al., 2019). We also use label smoothing set to 0.1, mixup (Zhang et al., 2017a) set to 0.8, cutmix (Yun et al., 2019) set to 1.0, and drop path set to 0.1. The learning rate is linearly warmed-up during the 5 first epochs and decayed with a cosine schedule thereafter.

B Theoretical Guarantees

In this section we describe how MSN pre-training provably avoids representation collapse.

Recall that in each iteration of pre-training, we sample a mini-batch of $B \geq 1$ images, and generate $M \geq 1$ anchor views of each image. Here we show that MSN is guaranteed to avoid the trivial collapse of representations under the following assumption.

Assumption 1 (Target Sharpening). *The target p^+ is sharpened, such that it is not equal to the uniform distribution.*

Proposition 1 (Non-Collapsing Representations). *Suppose Assumption 1 holds. If $f_\theta(\cdot)$ is such that the representations collapse, i.e., $z_{i,m} = z_{j,k}$ for all $i, j \in [B]$ and $m, k \in [M]$, then $\left\| \nabla_\theta H(p_i^+, p_{i,m}) \right\| + \left\| \nabla_\theta H(\overline{p}) \right\| > 0$ for all i, m.*

Proof. For L_2-normalized representations and prototypes, the prediction $p_{i,m} \in \Delta_K$ corresponding to the m^{th} view of the i^{th} image in the mini-batch is given by

$$p_{i,m} := \text{softmax}\left(\frac{z_{i,m} \cdot \mathbf{q}}{\tau} \right),$$

where $\mathbf{q} \in \mathbb{R}^{K \times d}$ is the prototype matrix with $K > 1$ learnable prototypes, each of dimension d, and $\tau > 0$ is a scaler temperature. Since $z_{i,m} = z_{j,k}$ for all $i, j \in [B]$ and $m, k \in [M]$, it holds that $z_{i,m} \cdot \mathbf{q} = z_{j,k} \cdot \mathbf{q}$, and therefore $p_{i,m} = p_{j,k}$. Now consider two separate cases.

Case 1: The predictions are equal to the uniform distribution, i.e., $p_{i,m} = \frac{1}{K} \mathbf{1}_K$, where $\mathbf{1_K} \in \mathbb{R}^K$ is the K-dimensional vector with each entry equal to 1. In that case, since, by Assumption 1, the targets

17

p_i^+ are sharpened such that they are not equal to the uniform distribution, it follows that $p_{i,m} \neq p_i^+$, and hence $\left\| \nabla_\theta H(p_i^+, p_{i,m}) \right\| > 0$.

Case 2: The predictions are not equal to the uniform distribution, i.e., $p_{i,m} \neq \frac{1}{K}\mathbf{1}_K$. In that case, we have that the average prediction across all the anchor views $\overline{p} := \frac{1}{MB} \sum_{i=1}^{B} \sum_{m=1}^{M} p_{i,m}$ is also not equal to the uniform distribution; i.e., $\overline{p} \neq \frac{1}{K}\mathbf{1}_K$, and hence $\left\| \nabla_\theta H(\overline{p}) \right\| > 0$. $\qquad\square$

Proposition 1 provides a theoretical guarantee that MSN is immune to the trivial collapse of representations. In short, the underlying principle is that entropy maximization encourages the anchor predictions to utilize the full set of prototypes, thereby preventing collapse to a non-uniform distribution, while target sharpening encourages the anchor predictions to be confident, thereby preventing collapse to the uniform distribution.

Note that the sharpening mechanism defined in Section 3 (i.e., applying a temperature τ^+ in the target network softmax) may not always satisfy Assumption 1, unless one introduces a simple tie-breaking rule. In practice, such a rule is not necessary as the targets never become uniform (since we apply sharpening from the start of the training), although, it is important to use a sufficiently small temperature value in this case.

C   Additional Ablations

C.1   Sinkhorn Normalization

By default, we set the ME-MAX regularization weight λ to 1.0 and apply Sinkhorn normalization on the targets to avoid having to tune the ME-MAX regularization weight. However, we find that tuning the ME-MAX regularization weight and omitting Sinkhorn normalization can result in better performance; cf. Table 11.

Table 11: **Effect of Sinkhorn normalization.** We train a ViT-S/16 with a masking ratio of 0.15, and explore the impact of Sinkhorn normalization during pre-training on low-shot performance with 1% of ImageNet-1K. Tuning the ME-MAX regularization weight and omitting Sinkhorn normalization gives better performance.

Architecture	Target Normalization	ME-MAX **weight** λ	Top 1
	Sinkhorn	1.0	66.4
ViT-S/16	None	1.0	60.8
	None	5.0	**67.2**

C.2   Number of Prototypes

By default we train with 1024 prototypes of dimension 256. In this section we explore the effect of the number of prototypes on low-shot performance. We observe that using more prototypes has little effect on training, but using too few prototypes can hurt performance; cf. Table 12.

Table 12: **Effect of number of prototypes.** We train a ViT-B/16 with a masking ratio of 0.3, and explore the impact of the number of prototypes during pre-training on low-shot performance with 1% of ImageNet-1K. Using more prototypes has little effect on training, but using fewer prototypes can degrade performance.

Architecture	Prototypes	Top 1
	512	67.6
ViT-B/16	1024	**69.5**
	2048	**69.5**

C.3   Masked Auto-Encoder Partial Fine-Tuning

Here we explore the low-shot performance of MAE when relying on alternative evaluation strategies. He et al. (2021) conjecture that using pixel reconstruction in their MAE objective results in encoder

representations of a lower semantic level than other methods, which may explain their difficulty in training a linear classifier on the frozen features. In Table 13 we explore the effect of partial fine-tuning on the low-shot performance of pre-trained MAE models. Partial fine-tuning corresponds to fine-tuning the last block of the pre-trained model along with a linear head on the available labeled samples. As observed in (He et al., 2021), MAE benefits from partial fine-tuning. However, for sufficiently large models, such as the ViT-H/14, this leads to significant overfitting in the low-shot regime, where one must instead resort to linear evaluation. We report the best numbers for MAE across the two low-shot adaptation strategies in Figure 1.

Table 13: **MAE low-shot evaluations.** Top-1 low-shot validation accuracy for different training strategies with MAE pre-trained models. Partial fine-tuning corresponds to fine-tuning the last block of the pre-trained model along with a linear head on the available labeled samples. Linear evaluation corresponds to training a linear classifier on top of the frozen pre-trained encoder. MAE benefits from partial fine-tuning, but for sufficiently large models, such as the ViT-H/14, this leads to significant overfitting in the low-shot regime, where one must instead one must resort to linear evaluation.

| | | **Top 1** | | |
| | | Images per Class | | |
Architecture	Adaptation Strategy	2	5	∼13
ViT-B/16	Partial Fine-Tuning	**25.0**	**40.5**	**51.1**
	Linear Eval.	14.5	25.2	36.6
ViT-L/16	Partial Fine-Tuning	19.3	**42.3**	**59.4**
	Linear Eval.	**22.1**	35.7	48.6
ViT-H/14	Partial Fine-Tuning	rand	rand	rand
	Linear Eval.	**18.6**	**32.8**	**46.7**

D   MSN Representation Robustness

Next we report the performance of MSN-pre-trained models on datasets that have been developed to evaluate the robustness of models trained on the standard ImageNet training set. We consider four datasets: ImageNet-A (Hendrycks et al. (2021b))[2], ImageNet-R (Hendrycks et al. (2021a))[3], ImageNet-Sketch (Wang et al. (2019))[4], and ImageNet-C (Hendrycks & Dietterich (2019))[5].

Table 14 shows results for a ViT-B/16 pre-trained using MSN and fine-tuned using the protocol described in Appendix A. For comparison, we also report the performance of a fine-tuned ViT-B/16 pre-trained using MAE (He et al., 2021), along with a supervised ResNet50 baseline, which is available in the PyTorch Torchvision package[6]. For ImageNet-A, -R, and -Sketch, we report top-1 accuracy on each provided validation set. For ImageNet-C, we use the mean Corruption Error metric proposed in (Hendrycks & Dietterich, 2019), where values are normalized by AlexNet performance on the same validation set.

In each case we find that the performance of an MSN-pretrained ViT-B/16 is comparable or better than that of an MAE-pretrained ViT-B/16. Note also, that larger MAE-pretrained models achieve stronger performance on all four datasets (He et al., 2021).

E   MSN Invariance to Masking

The goal of MSN pretraining is to denoise the input images at the representation level by ensuring that the representation of a masked input matches the representation of the unmasked one. Here, we shows that MSN pretraining learns representations that are robust to patch masking.

---

[2]https://github.com/hendrycks/natural-adv-examples

[3]https://github.com/hendrycks/imagenet-r

[4]https://github.com/HaohanWang/ImageNet-Sketch

[5]https://github.com/hendrycks/robustness

[6]https://github.com/pytorch/vision

Table 14: **Evaluation on alternative ImageNet validation sets.** We evaluate the performance of a fine-tuned ViT-B/16 model on four alternative ImageNet validation sets: ImageNet-A, ImageNet-R, ImageNet-Sketch, and ImageNet-C. The metric used for the first three (-A, -R, and -Sketch) is top-1 accuracy on the validation set. On ImageNet-C, performance is measured in terms of mean Corruption Error (mCE) (Hendrycks & Dietterich, 2019).

	IN-A (top-1 ↑)	IN-R (top-1 ↑)	IN-Sketch (top-1 ↑)	IN-C (mCE ↓)
Supervised ResNet50	0.04	36.11	24.2	76.7
MAE ViT-B/16 (He et al., 2021)	35.9	48.3	34.5	51.7
MSN ViT-B/16	**37.5**	**50.0**	**36.3**	**46.6**

In Table 15, we evaluate the performance of MSN and DINO when masking parts of an image during evaluation. Models are evaluated on 1% of ImageNet-1K using logistic regression on top of frozen features. The logistic regression classifier is trained using masked images, and then evaluated on the standard ImageNet-1K validation set using unmasked images.

If the MSN representations are robust to missing image patches, then a linear classifier should be able to identify generalizable features when training on the representations of masked images. On the other hand, if the representations output by the learned encoder are not robust to missing image patches, then a linear classifier would have difficulty finding generalizable features when training on the representations of masked images.

We observe that masked pre-training results in representations that are more robust to patch removal, suggesting that MSN is performing an image denoising at the representation level. Furthermore, models pre-trained with more aggressive masking exhibit this quality to a higher degree. For example, the low-shot accuracy of ViT-L/7 pre-trained with aggressive masking is almost unaffected when we remove 70% of the patches at test time; 75.1% top-1 without dropping patches during evaluation versus 74.9% top-1 when dropping 70% of the patches during evaluation.

Table 15: **Robustness to missing patches (low-shot).** Evaluating the low-shot accuracy of pre-trained models on 1% of ImageNet-1K when corrupting the annotated images by dropping patches. We train a linear classifier using masked images, and then evaluate on the standard ImageNet-1K validation set using unmasked images. We observe that MSN pre-training leads to representations that are more robust to masking. Moreover, models pre-trained with more aggressive masking exhibit this behaviour to a higher degree.

			Top 1		
			Eval. Masking Ratio		
Alg.	Arch.	Pre-train Masking Ratio	0.0	0.7	Δ
DINO	ViT-B/16	0.0	67.0	63.1	-3.9
MSN	ViT-B/16	0.3	69.5	67.1	-2.4
	ViT-L/7	0.7	75.1	74.9	**-0.2**

We also report the average cosine distance between masked and unmasked representations of the same image in Table 16. As expected, the cosine similarity between masked and unmasked representations of the same image is higher when pre-training with MSN, supporting the observation that masked-pretraining results in representations that are more robust to patch-removal.

F   Qualitative Analysis

We qualitatively investigate the properties of the MSN pre-trained representations. We follow the RCDM framework (Bordes et al., 2021) and train a conditional generative diffusion model, which maps a learned image representation back to pixel space. Specifically, RCDM takes as input random noise and the representation vector of an image computed by an SSL model (either an MSN pre-trained model or a DINO pre-trained model in this analysis), and aims to reconstruct the image as close as possible to the original one through a diffusion process.

Table 16: **Robustness to missing patches (cosine-similarity).** Average Cosine Distance between masked and unmasked representations of the same image. We compare the representations learned with MSN masked pre-training to those learned with DINO when using a ViT-B/16 encoder. The MSN ViT-B/16 is pre-trained with a masking ratio of 0.3. The cosine distances are computed and averaged over the ImageNet-1k validation set. The cosine similarity between masked and unmasked representations of the same image is higher when pre-training with MSN, supporting the observation that masked-pretraining results in representations that are more robust to patch-removal.

Cosine Similarity					
		Eval.	Masking	Ratio	
Alg.	0.15	0.3	0.5	0.7	0.9
DINO	0.98	0.97	0.92	0.81	0.56
MSN	**0.99**	**0.99**	**0.99**	**0.98**	**0.97**

By using RCDM to sample an image based on its SSL representation, we can visualize how different pre-training strategies affect the degree of information contained in the representation. Qualities that vary across RCDM samples represent information that is not contained in the pre-trained representation. Qualities that are semantically common across samples represent information contained in the representation.

F.1 Comparison with DINO

We apply RCDM on top of either a DINO or MSN pre-trained ViT-B/8 encoder to generate images of resolution 128×128 pixels. RCDM is trained using unmasked images processed with the ViT-B/8 encoder. We then use masked images from the validation set at sampling time.

In Figure 5, we generate samples for RCDM when masking 50% of the conditioning images. The first column depicts images from the ImageNet validation set. The second column depicts the same image, but with 50% of the patches masked. The representation of the masked image is used as conditioning for the RCDM diffusion model. The subsequent columns in Figure 5 show various images sampled from the conditioned RCDM diffusion model. We observe that the RCDM samples conditioned on the MSN representations (cf. Figure 5a) preserve the semantic category of the masked images, and remain visually close to the original image, despite the missing patches. By contrast, the samples generated by the RCDM diffusion model conditioned on the DINO representations (cf. Figure 5b) are more blurry and do not preserve as well the semantic category of the masked images.

Figure 6 depicts similar visualizations, but with 80% of the patches masked. In this case, even with 80% of the patches missing, samples generated by RCDM conditioned on MSN representations preserve some of the structure in original images (cf. Figure 6a). On the other hand, conditioning on DINO representations leads to almost uniform background generation (cf. Figure 6b).

F.2 MSN ViT-L/7 Visualizations

We apply RCDM on top of the MSN pre-trained ViT-L/7 encoder to generate images with a resolution of 256×256 pixels. RCDM is trained using images with 70% of patches masked. We then use masked images from the validation set (with various masking ratios) at sampling time, see Figures 7, 8, and 9.

Visualizations show that MSN discards instance-specific information such as background, pose, and lighting, while retaining semantic information about the images, even when a large fraction of the patches are masked.

(a) MSN Representations visualized on ImageNet validation set.

(b) DINO Representations visualized on ImageNet validation set.

Figure 5: **Visualizations of ViT-B/8 pre-trained representations computed from images with 50% of patches masked.** First column: original image. Second column: image with 50% of patches masked used to compute representations of an SSL pre-trained ViT-B/8 encoder. Other columns: RCDM sampling from generative model conditioned on SSL representation of masked image.

(a) MSN representations visualized on ImageNet validation set.

(b) DINO representations visualized on ImageNet validation set.

Figure 6: **Visualizations of ViT-B/8 pre-trained representations computed from images with 80% of patches masked.** First column: original image. Second column: image with 80% of patches masked used to compute representations of an SSL pre-trained ViT-B/8 encoder. Other columns: RCDM sampling from generative model conditioned on SSL representation of masked image.

Figure 7: **Visualizations of MSN pre-trained ViT-L/7 representations computed from unmasked images.** First column: original image. Other columns: RCDM sampling from generative model conditioned on MSN representation using a ViT-L/7 encoder. MSN representations are computed from unmasked images. Qualities that vary across samples represent information that the representation is invariant to; e.g., in this case, MSN discards background, pose, and lighting information. Qualities that are common across samples represent information contained in the pre-trained representation.

Figure 8: **Visualizations of MSN pre-trained ViT-L/7 representations computed from images with 70% of patches masked.** First column: original image. Second column: image with 70% of patches masked used to compute representations of an SSL pre-trained ViT-L/7 encoder. Other columns: RCDM sampling from generative model conditioned on SSL representation of masked image. Qualities that vary across samples represent information that the representation is invariant to; e.g., in this case, MSN discards background, pose, and lighting information. Qualities that are common across samples represent information contained in the pre-trained representation.
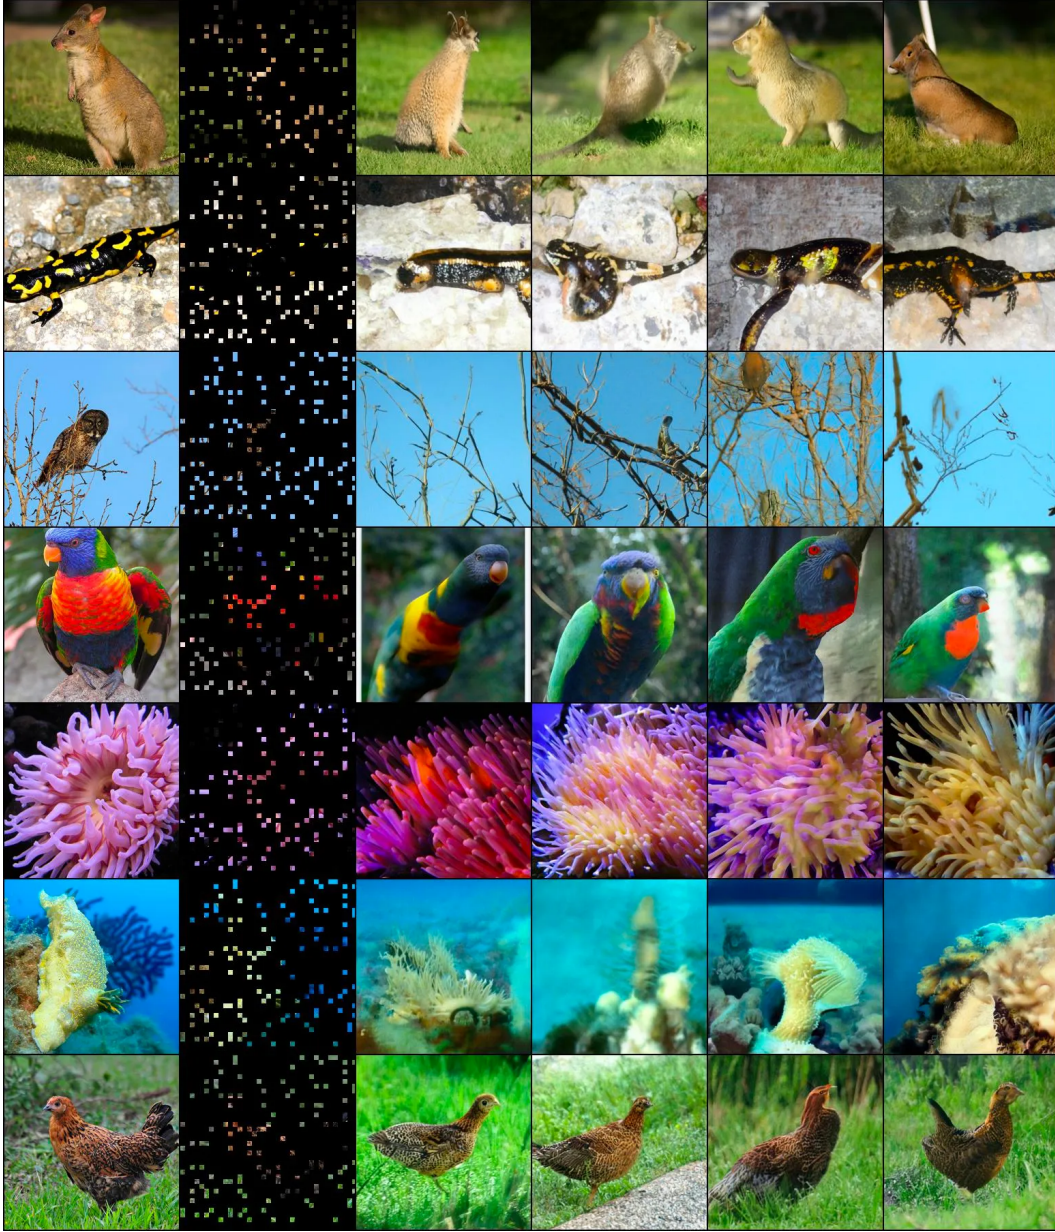
Figure 9: **Visualizations of MSN pre-trained ViT-L/7 representations computed from images with 90% of patches masked.** First column: original image. Second column: image with 90% of patches masked used to compute representations of an SSL pre-trained ViT-L/7 encoder. Other columns: RCDM sampling from generative model conditioned on SSL representation of masked image. Qualities that vary across samples represent information that the representation is invariant to; e.g., in this case, MSN discards background, pose, and lighting information. Qualities that are common across samples represent information contained in the pre-trained representation. Even with high-masking ratio, MSN retains semantic information about the images.