

Representation Learning by Detecting Incorrect Location Embeddings

Sepehr Sameni¹, Simon Jenni², Paolo Favaro¹

¹ Computer Vision Group, University of Bern, Switzerland

² Adobe Research

sepehr.sameni@unibe.ch, jenni@adobe.com, paolo.favaro@unibe.ch

Abstract

In this paper, we introduce a novel self-supervised learning (SSL) loss for image representation learning. There is a growing belief that generalization in deep neural networks is linked to their ability to discriminate object shapes. Since object shape is related to the location of its parts, we propose to detect those that have been artificially misplaced. We represent object parts with image tokens and train a ViT to detect which token has been combined with an incorrect positional embedding. We then introduce sparsity in the inputs to make the model more robust to occlusions and to speed up the training. We call our method DILEMMA, which stands for Detection of Incorrect Location EMBeddings with MAsked inputs. We apply DILEMMA to MoCoV3, DINO and SimCLR and show an improvement in their performance of respectively 4.41%, 3.97%, and 0.5% under the same training time and with a linear probing transfer on ImageNet-1K. We also show full fine-tuning improvements of MAE combined with our method on ImageNet-100. We evaluate our method via fine-tuning on common SSL benchmarks. Moreover, we show that when downstream tasks are strongly reliant on shape (such as in the YOGA-82 pose dataset), our pre-trained features yield a significant gain over prior work.¹

Introduction

In computer vision, deep learning models trained on small labeled datasets can benefit greatly from supervised pre-training on datasets such as ImageNet (Girshick et al. 2014). Even more surprisingly, (He et al. 2020) showed that it is possible to pre-train with unlabeled data (with MoCo) and outperform pre-training with supervised learning on several downstream tasks. This led to the rapid development of several Self-Supervised Learning (SSL) methods, such as (Caron et al. 2020; Chen et al. 2020a,b; He et al. 2020).

Representations obtained via SSL have the ability to *generalize* to downstream tasks such as object classification, detection, and segmentation (Deng et al. 2009; Everingham et al. 2009; Krizhevsky 2009). Recent work suggests that representations with a *shape bias* generalize better to these tasks than those with a texture bias (Geirhos et al. 2018; Tartaglini, Vong, and Lake 2022).

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹source code: <https://github.com/Separius/DILEMMA>




Figure 1: Yoga₈₂ is a dataset with a shape-based task. As these examples show, texture alone is not sufficiently indicative of the pose. Thus, models that perform well in this task may demonstrate a strong shape discriminability.

In particular, the expectation is that image representations may do better in the transfer learning to a shape-based task, such as pose classification of Yoga₈₂ (see Fig. 1). Thus, we investigate whether adding a regularization loss that is sensitive to shape to a state-of-the-art SSL method, might lead to better representation learning.

We propose DILEMMA, which is short for Detection of Incorrect Location EMBeddings with MAsked inputs. In our experiments we integrated DILEMMA with SimCLR (Chen et al. 2020a), DINO (Caron et al. 2021) and MoCoV3 (Chen, Xie, and He 2021), and find a consistent improvement across the majority of downstream tasks.

With DILEMMA, the image representation is encouraged to differentiate shapes thanks to two main components: 1) A binary classification loss to detect the correct/incorrect positions of object parts, and 2) the use of randomized input sparsity, so that every subset of object parts contributes to the whole image representation. The first component is a concept already proposed in the context prediction (Doersch, Gupta, and Efros 2015) and jigsaw puzzle SSL methods (Noroozi and Favaro 2016). It takes also inspiration from ELECTRA (Clark et al. 2020a), where some text tokens are replaced by a weak generator and a discriminator is trained to detect them. The second component, is also a concept that has been exploited in VATT (Akbari et al. 2021) and MAE (He et al. 2021) to reduce the computational workload of training with ViTs (Dosovitskiy et al. 2020).

More in detail, as shown in Fig. 2, we split an image into a grid of tiles, map them to tokens, and combine them with positional embeddings. Then, we corrupt the positional embeddings of a fraction of the tokens before we feed them to a ViT. In our DILEMMA loss, we classify the tokens into

those with correct and incorrect positional embeddings. The *sparsification* of the input can be implemented in a ViT simply by discarding a randomized percentage of tokens.

We also use a teacher-student architecture as in MoCoV3 (Chen, Xie, and He 2021). We sparsify only the input to the student network and instead feed all the tiles to the teacher network. Because it is used only in evaluation mode, it does not have a significant impact on storage and computing resources. Moreover, the use of a complete set of tiles (a setting that we call *dense*) and without corrupted positional embeddings, allows the teacher to build a better reference for the student network.

Our contributions can be summarized as follows:

- We introduce DILEMMA, a novel SSL regularization loss that enhances the shape discriminability of image representations; it is based on the detection of misplaced positional embeddings with a ViT and the use of sparsity in the input;
- We propose to randomly sparsify the inputs and to use a student-teacher architecture to: 1) reduce the memory storage, 2) close the gap between training and test data, 3) speed up the training;
- DILEMMA boosts the performance of MoCoV3, SimCLR, DINO, and MAE under the same computational budget.

Related Work

Self-Supervised Learning for Image Representations.

Self-supervised learning gained popularity as a form of unsupervised learning, where pretext tasks leverage supervision signals obtained without human labor. Some classic examples are the classification of image patch locations (Doersch, Gupta, and Efros 2015; Noroozi and Favaro 2016), the reconstruction of color channels (Zhang, Isola, and Efros 2016) or image patches (Pathak et al. 2016), or the recognition of various image transformations (Gidaris, Singh, and Komodakis 2018; Jenni and Favaro 2018). While prior patch-based methods inspired our approach of detecting wrongly placed image patches, ours is both simpler and performs better in transfer experiments. Furthermore, due to the input representation of ViTs (disjoint image patches) and our random sparse patch sampling, our approach suffers less from domain gaps between pre-training and transfer.

Contrastive Learning. Efforts to scale up and improve instance discrimination (Dosovitskiy et al. 2015; Wu et al. 2018a) as a self-supervised pre-training task have established contrastive learning (Chen et al. 2020a; He et al. 2020; Oord, Li, and Vinyals 2018) as the most popular SSL approach in computer vision today. Several modifications of the basic recipe, *i.e.*, learning to discriminate training instances up to data augmentations, have been proposed since. For example, some methods leverage momentum encoded samples for positive and negative sampling (He et al. 2020; Chen et al. 2020b), some remove the need for explicit negative pairs (Grill et al. 2020; Chen and He 2020), and others extend the set of positives beyond data-augmentation through clustering (Caron et al. 2020) or nearest-neighbors in feature space (Dwivedi et al. 2021). Another line of work

considers contrastive pre-training strategies tailored to dense prediction tasks (O Pinheiro et al. 2020; Wang et al. 2021; Xiao et al. 2021; Xie et al. 2021b; Li et al. 2021b; Liu et al. 2021a). More recently, contrastive methods leverage vision transformer architectures (Dosovitskiy et al. 2020; Liu et al. 2021b), *e.g.*, by adapting existing approaches (Chen, Xie, and He 2021; Xie et al. 2021a), tailoring architectures (Li et al. 2021a), or novel objectives (Caron et al. 2021). In our approach, we show that several well-established contrastive baselines (Chen, Xie, and He 2021; Caron et al. 2021; Chen et al. 2020a) can be improved through the addition of a spatial reasoning task and by extending the set of image augmentations through randomized patch dropping.

Self-Supervised Pre-Training of Transformers. The success of the transformer architecture (Vaswani et al. 2017) in natural language is to a great extent due to large-scale self-supervised pre-training tasks. Successful pre-training strategies from NLP like masked token prediction (Devlin et al. 2018) have recently also been adapted to the image domain (Bao, Dong, and Wei 2021; Zhou et al. 2021; He et al. 2021; Zhou et al. 2021). Our patch misplacement detection is similar to another type of pretext task in NLP, where the goal is to detect corrupted tokens, *i.e.*, words replaced by an imperfect masked language model (Clark et al. 2020a,b). However, a key difference in our approach is that we only tamper with the spatial position of the tokens and thus do not require a separate masked token prediction model. In parallel work, Fang et al. (Fang et al. 2022) use BEiT (Bao, Dong, and Wei 2021) for that purpose. The method of DABS (Tamkin et al. 2021) also uses the idea of patch misplacement, but it does not have a way to handle degenerate learning and it does not show performance improvements. MP3 (Zhai et al. 2022) also predicts the position of all the tokens like jigsaw (Noroozi and Favaro 2016) with a ViT. A technique that has proven very beneficial to improve the training efficiency of vision transformers is token dropping (Akbari et al. 2021; He et al. 2021; El-Nouby et al. 2021; Chen et al. 2022). We extend this technique by randomizing the token dropping amount and including the case of no dropping to narrow the domain gap between pre-training and transfer.

Training DILEMMA

Let us define an image sample as $x \in \mathbb{R}^{H \times W \times C}$, *i.e.*, x has $H \times W$ pixels and C color channels. We apply two data augmentations (Grill et al. 2020) to x and obtain \hat{x}_1 and \hat{x}_2 . Similarly to ViT, each input \hat{x}_1 and \hat{x}_2 is divided in 14×14 tiles, flattened and projected to N tokens $t_{1,i}, t_{2,i} \in \mathbb{R}^D$, $\forall i \in U \doteq \{1, \dots, N\}$, through a linear projection. We then combine each token $t_{\cdot,i}$ with a positional embedding $p_i \in \mathbb{R}^D$, which can be either learned or fixed.

As in MoCoV3 (Chen, Xie, and He 2021), we define a *Student S* and a *Teacher T* ViTs (Dosovitskiy et al. 2020), where the Teacher, also called *momentum encoder*, is obtained through the exponential moving average (EMA) of the Student’s weights (thus, it is not trained). The Teacher receives as input all the tokens $t_{1,1}, \dots, t_{1,N}$ with the corresponding positional embeddings p_1, \dots, p_N . The Student instead receives as input a sparse set $M \subset U$ of tokens $t_{2,i}$,

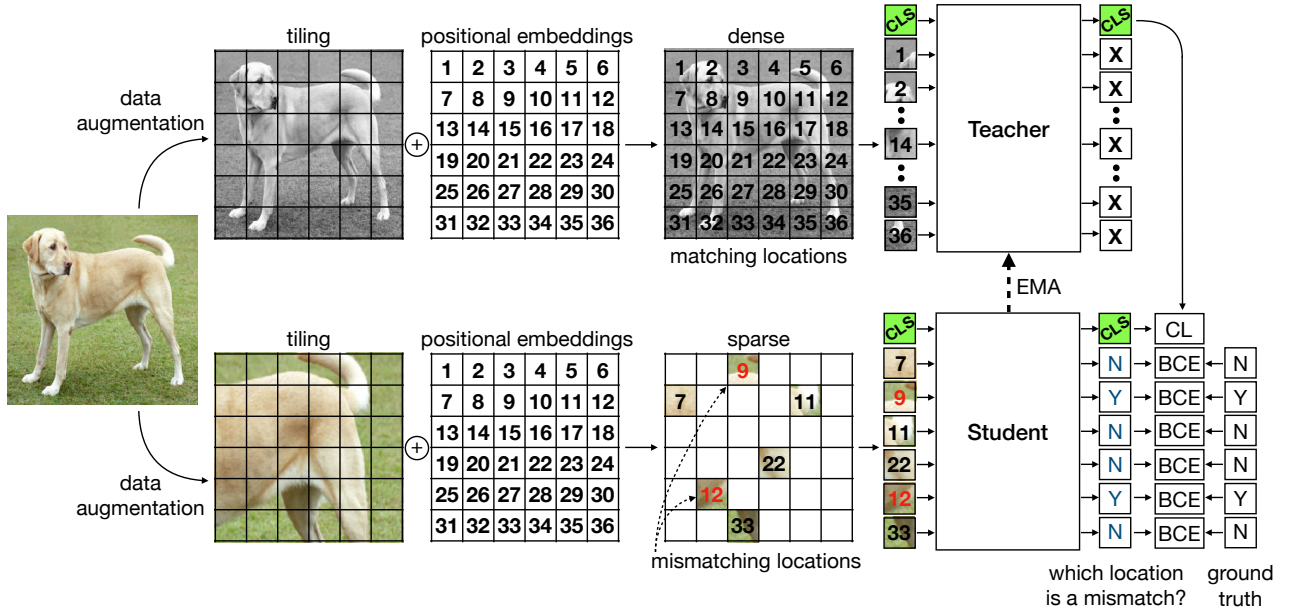


Figure 2: Training with DILEMMA: A sample image is augmented twice and split into tiles (we use a 14×14 grid). The Teacher network takes the complete set of tiles as input (dense) and without mismatches in the positional embeddings for each token. The Student takes only a subset of the tiles as input (sparse) and some tiles have incorrect positional embeddings. The Student is then trained under two losses: one is the contrastive loss of the class tokens (CLS) between the Teacher and the Student, and the other is the DILEMMA binary cross-entropy for each token.

$i \in M$. For a randomized fraction of these tokens $B \subset M$ the corresponding positional embeddings q_i , $i \in M$ are incorrect, i.e., $q_i \doteq p_i$ if $i \in M \setminus B$ and $q_i \doteq p_j$ with $j \in U \setminus M$, if $i \in B$. We call the ratio $\theta = |B|/|M| \in [0, 1]$, between the cardinalities of B and M , the *probability of a positional embedding mismatch*. We choose a different M and B sets for each sample at each iteration. We define a set of ground truth labels $y_i = 0$ (N) if $i \in M \setminus B$ and $y_i = 1$ (Y) if $i \in B$. The i -th output token from the Student is denoted with $S_i(\{q_j \oplus t_{2,j}\}_{j \in M})$. We indicate the extra classification token with $i = 0$ both at the input and output. Also, $q_0, p_0 = 0$, i.e., no location encoding.

Now we are ready to introduce the DILEMMA loss (see also the whole training method in Fig. 2)

$$\mathcal{L}_{\text{DILEMMA}} = \mathbb{E}_x \left[\sum_{i \in M} y_i \log(\sigma(Y S_i(\{q_j \oplus t_{2,j}\}_{j \in M \cup \{0\}}))) + (1 - y_i) \log(1 - \sigma(Y S_i(\{q_j \oplus t_{2,j}\}_{j \in M \cup \{0\}}))) \right] \quad (1)$$

where $\mathbb{E}[\cdot]$ is the expectation over image samples, σ is the sigmoid function and Y is a linear projection.

Because of the sparsity in the input to the Student network, we also obtain a computational benefit. When we increase the sparsity of the input, we can also increase the mini batch size to fully utilize the GPU RAM. This is particularly significant with ViTs, because of their quadratic scaling with the number of tokens (the memory usage is $O(N^2)$). The fact that we can significantly increase the mini batch size is particularly effective with contrastive learners. Moreover, in this way it is also faster to train our model, because the av-

erage mini batch size is much larger than when using dense inputs (in our case it is $2.5 \times$ more).

Combining DILEMMA and Contrastive Learning

The DILEMMA loss can be integrated with other SSL losses. Here we describe the integration with the contrastive loss, but other choices follow an identical procedure.

The contrastive loss is defined as

$$\mathcal{L}_{\text{CNT}} = \mathbb{E}_x \left[L_{\text{CE}}(S_0(\{q_j \oplus t_{2,j}\}_{j \in M \cup \{0\}}), T_0(\{p_j \oplus t_{1,j}\}_{j=0, \dots, N})) \right], \quad (2)$$

where

$$L_{\text{CE}}(A, V) = -2\tau \sum_n z_n \log \text{softmax} \left(\frac{A_n^\top V}{\tau} \right) \quad (3)$$

and A and V are $G \times m$ matrices, with m the minibatch size and G the vector size after the projection Y (see eq. (1)), z_j is the one-hot vector with 1 at the j -th position and the index n indicates the class token within the minibatch.

When we combine both the DILEMMA and the contrastive losses into a single cost we obtain

$$\mathcal{L}_{\text{UNION}} = \lambda_{\text{DILEMMA}} \mathcal{L}_{\text{DILEMMA}} + \mathcal{L}_{\text{CNT}}, \quad (4)$$

which we minimize and where $\lambda_{\text{DILEMMA}} > 0$ is a hyperparameter which we always set to 0.4.

Implementation

Architecture. We use Vision Transformers (ViT) (Dosovitskiy et al. 2020) with a patch size of 16×16 pixels and an input image size of 224×224 pixels, which gives a total of

$(224/16)^2 = 196$ tokens. Due to computational limitations, we mostly use the small variant of the Vision Transformer (ViT-S) which has 12 transformer blocks and 384 channels. For the three baselines: 1) For MoCoV3 (Chen, Xie, and He 2021) experiments, we use 12 attention heads in each attention layer as specified in the official implementation. This is different from most ViT-S implementations, which use 6 heads. This does not change the total number of parameters of the model, but incurs a speed penalty. We use a 3-layer MLP for the projection and prediction heads with synchronized batch normalization. We also freeze the weights of the patch embedding layer for better stability; 2) SimCLR (Chen et al. 2020a) experiments are also conducted with the exact same settings, but without a teacher network and instead both augmentations are sparsified, misplaced and then fed to the student network; 3) For DINO (Caron et al. 2021) we used the official implementation and, whenever multi-crop is used, we have disabled random sparsity and used constant sparsity for the large crops and no sparsity for the small crops (96×96 images).

Pre-training Setup. For our main model, we pre-train DILEMMA on ImageNet-1K (Deng et al. 2009) with the exact same hyper-parameters of MoCoV3 using three GeForce RTX 3090 GPUs for 100 epochs with a base batch size of 345. We set λ_{DILEMMA} to 0.4 and the probability of positional embedding mismatch $\theta = 0.2$. We use sparsity ratios of 0%, 40%, 55%, 65% with $1\times$, $2\times$, $3\times$, $4\times$ base batch size and disable the DILEMMA loss when the input is dense.

To show the compatibility of the proposed method with other SSL methods, we also added two short runs for SimCLR and DINO with multi-cropping. For the DINO experiments we used ViT-Base to show that DILEMMA scales to larger models. Since input sparsity allows for faster training, we also report results of DILEMMA variants with equal training time as the baselines.

Linear Probing. To evaluate the pre-trained features for image classification, we train a simple linear layer on top of frozen features, without any data augmentation (Linear_F). Note that it is different from the standard linear probing, and we opt to use this method for its simplicity and speed. It is also more aligned with the end goal of representation learning. In all the linear probing experiments, we use the embedding of the CLS token of the last layer and perform a coarse grid search over learning rates, batch sizes and whether to normalize the data before feeding it to the linear layer or not (similarly to the added BatchNorm layer (Ioffe and Szegedy 2015) in MAE (He et al. 2021)). In contrast, DINO (Caron et al. 2021), obtains its representation by concatenating the CLS token of the last four attention layers of the network.

Experiments

We evaluate the use of DILEMMA on several datasets, compare it to state-of-the-art (SotA) SSL baselines, and perform ablations to show the role of each loss component. In each table, where we compare to an SSL baseline, we indicate the baseline with a method name (*e.g.*, MoCoV3 (Chen, Xie,

Method	Epochs	Time	BS	k -NN	Linear _F	Linear
SimCLR	30	15.7h	512	41.46	50.21	-
+Sparsity	30	12.2h	512	41.11	49.73	-
+DILEMMA	30	12.2h	512	41.90	50.71	-
DINO* [†]	45	120.9h	192	61.35	65.46	-
+Sparsity* [†]	45	90.7h	192	62.33	68.49	-
+DILEMMA* [†]	45	90.7h	192	62.48	68.55	-
+DILEMMA* ^{††}	60	121.0h	192	63.74	69.43	-
MoCoV3	100	102.8h	345	59.68	63.62	65.1
+Sparsity	100	68.4h	345	61.64	65.16	-
+DILEMMA	100	68.4h	345	61.97	65.62	66.6
+Sparsity [†]	150	102.6h	345	63.27	67.07	-
+DILEMMA [†]	150	102.6h	345	64.69	68.03	-
MoCoV3	300	-	4096	67.90	72.72	73.2
DINO	300	-	1024	67.9	-	72.5
DINO [†]	800	-	1024	74.30	75.74	77.0
Supervised	300	-	1024	-	-	79.8

Table 1: ImageNet-1K Classification Transfer. The evaluation uses k -NN and linear probing with a ViT-S/16 or, where indicated, a ViT-Base/16 architecture. The [†] models are trained for a number of epochs, such that the total training time (see column Time) is the same as for the baseline methods. BS stands for Batch Size. [†] models are trained with multi-crop. * indicates ViT-Base/16 models

and He 2021)) and use a $+\{\text{DILEMMA/sparsity}\}$ to indicate that the baseline immediately above is combined with just sparsity or with the DILEMMA loss, which includes sparsity. We compare these two cases to show the added benefit of the DILEMMA positional classification loss over the lone sparsity.

Classification on ImageNet-1K

We show that DILEMMA leads to better representations for ImageNet-1K than prior SotA methods. Since this dataset has been used as a reference in SSL, it allows an easy comparison with previous work. In all tested cases, DILEMMA shows a consistent and significant improvement over the baseline it has been integrated with. Notice that the improvement due to the positional loss, relative to the use of sparsity, becomes more significant with a longer training.

k -NN and Linear Probing. In Table 1, we evaluate the quality of the ImageNet-1K pre-trained features. We either use a weighted k nearest neighbor classifier (we always use $k = 20$) (Wu et al. 2018b) or a simple linear layer on top of a frozen backbone and frozen features. Since the use of sparsity has the added benefit of reducing the computational load at each iteration, we also show the actual training time. For example, with a ViT-Base/16 model and multi-crop, DINO + DILEMMA (denoted with the [†] symbol) trains for 60 epochs in about the same time DINO trains for 45 epochs. This gives a significant advantage in performance. Furthermore, DILEMMA outperforms the baseline methods even if trained for the same number of epochs. The improvement under the same number of epochs is about 1 – 2% due to sparsity and 0.15 – 0.33% due to the positional classification

	Aircraft	Caltech ₁₀₁	Cars	CIFAR ₁₀	CIFAR ₁₀₀	DTD	Flowers ₁₀₂	Food ₁₀₁	INat ₁₉	Pets	STL ₁₀	SVHN	Yoga ₈₂	Avg.
MoCoV3	38.70	87.35	28.72	91.97	75.09	64.63	91.67	67.97	33.30	84.63	95.76	64.56	56.41	67.75
-Position	16.29	60.79	5.88	31.36	13.96	50.59	60.51	31.64	18.12	53.34	70.61	20.02	15.25	34.49
+Sparsity	43.29	89.25	40.14	92.28	77.30	65.05	93.25	71.39	42.38	86.35	96.09	64.60	62.29	71.05
+DILEMMA	44.43	89.55	42.32	93.03	77.56	64.47	93.41	71.94	43.76	85.88	96.08	65.41	63.76	71.66
+Sparsity [†]	44.64	89.89	40.21	93.53	78.88	65.48	94.21	72.96	42.84	88.93	96.90	64.52	63.24	72.02
+DILEMMA [†]	46.02	90.29	43.44	94.20	80.05	65.37	94.47	74.10	44.13	88.53	96.75	66.30	64.90	72.97
DINO	45.66	88.30	47.07	90.61	74.06	66.22	94.54	73.00	47.36	85.83	96.62	53.74	59.90	70.99
+Sparsity	46.83	89.16	48.45	92.47	77.00	68.40	94.70	74.97	51.32	85.45	96.85	66.65	61.90	73.40
+DILEMMA	46.86	89.58	49.04	92.46	77.22	67.50	95.20	74.90	52.51	85.91	97.06	70.01	62.87	73.93
+DILEMMA [†]	48.60	89.66	50.85	93.39	78.85	68.30	95.38	75.48	52.71	86.54	97.41	71.47	64.00	74.82

Table 2: Transfer learning for image classification on 13 datasets. The [†] variants are trained for the same duration as the corresponding (non-sparse) baselines. — **Position** is the case of MoCoV3, where the input tokens do not have the corresponding positional embeddings

Method	ImageNet-1%		ImageNet-10%		
	k -NN	Linear _F	k -NN	Linear _F	
Single-Crop	DINO	40.60	45.24	52.95	58.35
	MoCoV3	38.48	43.69	50.83	56.08
	+Sparsity	40.02	45.44	52.56	59.06
	+DILEMMA	41.64	47.95	53.15	60.00
	+Sparsity [†]	42.42	48.34	54.62	61.29
	+DILEMMA [†]	45.62	51.58	56.66	62.61
Multi-Crop	DINO	41.79	46.88	53.00	59.48
	+Sparsity	42.36	48.65	53.61	62.33
	+DILEMMA	42.73	48.81	53.81	62.32
	+DILEMMA [†]	43.87	50.45	55.29	63.36

Table 3: Low-shot learning on ImageNet-1K. The [†] variants are trained for the same duration as the corresponding (non-sparse) baselines. In the Single-Crop case, DINO is shown only as a reference

loss for the k -NN evaluation. Similarly, it is about 2 – 3% due to sparsity and 0.06 – 0.46% due to the positional classification loss for our linear probing. Notice that the boost due to the positional classification loss becomes more significant with more epochs (e.g, for MoCoV3 and under the same running time, the k -NN evaluation shows a boost of 3.59% due to sparsity and an additional 1.42% due to the positional classification).

For the sake of completeness, we have also included best reported numbers for ViT-S/16 with significantly larger batch sizes and more training epochs.

Low-shot learning. In Table 3, we simulate transfers to small datasets. With reference to ImageNet, we use the model pre-trained on the whole unlabeled dataset, train a linear layer on top of the frozen features of the 1% or 10% subsets (Chen et al. 2020a) and then evaluate the results on the whole validation set. The results show that adding DILEMMA to MoCoV3 or DINO yields a more label-efficient representation than with the corresponding baselines. Notice that in this implementation DILEMMA is based on MoCoV3, which, as was observed in DINO (Caron et al. 2021), has a consistently worse k -NN accuracy than DINO. Nonetheless, the addition of DILEMMA can more

Method	Seg. w/ Lin.			Seg. w/ UPerNet		
	mIoU	mAcc	aAcc	mIoU	mAcc	aAcc
MoCoV3	12.44	15.91	65.95	32.13	43.37	76.79
+Sparsity	15.77	19.87	67.70	33.66	45.27	77.44
+DILEMMA	16.81	21.05	67.84	33.79	45.33	77.68
+Sparsity [†]	15.93	20.03	67.87	34.03	45.90	77.47
+DILEMMA [†]	17.11	21.48	67.98	34.98	46.73	77.97
DINO	23.51	30.42	68.73	30.64	43.90	74.52
+Sparsity	26.75	34.20	71.61	33.82	47.01	76.56
+DILEMMA	27.78	35.63	72.63	34.11	47.50	76.73
+DILEMMA [†]	28.72	36.72	72.79	34.87	47.95	77.61

Table 4: Semantic Segmentation on ADE20K. The [†] variants are trained for the same duration as the corresponding (non-sparse) baselines

than compensate for the performance gap.

Downstream Tasks

We evaluate DILEMMA on several datasets to assess its generalization capability across different classification and detection tasks. While DILEMMA improves the performance over the baselines in all the datasets, the most significant improvement seems to occur for more shape-based tasks, such as pose classification. The evaluation on object segmentation, which is a dense downstream task, illustrates the representation captured by the non-CLS tokens.

Transfer Learning. In Table 2, we evaluate the transfer capability of our representations for image classification on several datasets. We use: Aircraft (Maji et al. 2013), Caltech₁₀₁ (Fei-Fei, Fergus, and Perona 2004), Cars (Krause et al. 2013), CIFAR₁₀ (Krizhevsky 2009), CIFAR₁₀₀ (Krizhevsky 2009), DTD (Cimpoi et al. 2014), Flowers₁₀₂ (Nilsback and Zisserman 2008), Food₁₀₁ (Bossard, Guillaumin, and Van Gool 2014), INat₁₉ (iNaturalist 2019), Pets (Parkhi et al. 2012), STL₁₀ (Coates, Ng, and Lee 2011), SVHN (Netzer et al. 2011), and Yoga₈₂ (Verma et al. 2020). We train a linear layer on top of the frozen features to accelerate the process. DILEMMA performs well in transfer learning across all datasets and significantly more on

Method	$(\mathcal{J}\&\mathcal{F})_m$	DAVIS		VOC12
		\mathcal{J}_m	\mathcal{F}_m	Jac. _{sim.}
MoCoV3	58.28	57.46	59.09	46.50
+Sparsity	58.94	57.05	60.83	45.34
+DILEMMA	60.00	57.99	62.02	48.89
+Sparsity [†]	58.03	56.74	59.33	45.93
+DILEMMA [†]	59.84	57.98	61.69	49.36
DINO	57.01	55.13	58.90	41.60
+Sparsity	56.83	54.84	58.81	40.03
+DILEMMA	57.60	55.39	59.80	39.71
+DILEMMA [†]	57.25	55.18	59.31	44.14

Table 5: Unsupervised object segmentation. We show the mean region similarity \mathcal{J}_m and the mean contour-based accuracy \mathcal{F}_m for DAVIS, and the Jaccard similarity for VOC12. The [†] variants are trained for the same duration as the corresponding (non-sparse) baselines

datasets with shape-based tasks, such as Yoga₈₂ (Verma et al. 2020) (for yoga position classification).

We also try to measure approximately how much shape matters in each dataset. We evaluate MoCoV3 with tokens without their position embedding. For simplicity, we use the same pre-trained MoCoV3 used throughout the experiments (although one should ideally use a MoCoV3 trained without position embeddings). We indicate this case with **—Position** in Table 2. Without position embedding these features are equivalent to a bag of features. We can see that the improvement due to DILEMMA relative to the baseline MoCoV3 follows the corresponding relative degradation due to the bag of features representation. This suggests that DILEMMA tends to generalize better on datasets with shape-based tasks.

Semantic Segmentation on ADE20K. In Table 4, we show the evaluation of DILEMMA on semantic segmentation. This is a task that strongly relates to the shape of objects. Thus, we expect to see a significant improvement from a boost in the shape discriminability. The semantic segmentation capability of self-supervised methods is usually evaluated by fine-tuning the model with an extra decoder. For that we use UPerNet (Xiao et al. 2018) on the ADE20K (Zhou et al. 2017) dataset and train the model for 160K iterations with a batch size of 2 for ViT-Base and 8 for ViT-Small. We also follow the evaluation protocol of iBOT (Zhou et al. 2021) and just train a linear layer (for 160K iterations and a batch size of 16) for semantic segmentation with a frozen backend to directly assess the per-token representation. The results show that DILEMMA is also better than the baseline models for dense classification tasks. It yields remarkable mIoU improvements of 4.6% against MoCoV3 and of 5.2% against DINO in the linear settings and under the same training time.

Unsupervised Object Segmentation. In Table 5, we evaluate the single frame object segmentation task. We use the mask generated from the attention of the CLS token

Method	Shape Accuracy	Texture Accuracy
MoCoV3	80.78	82.66
+Sparsity	82.55	81.78
+DILEMMA	83.58	82.11
+Sparsity [†]	82.72	82.77
+DILEMMA [†]	83.52	83.82
DINO	80.84	79.47
+Sparsity	83.18	81.01
+DILEMMA	83.58	80.79
+DILEMMA [†]	83.64	81.45

Table 6: Humanoid Vision Engine Benchmark (Ge et al. 2022) results. The [†] models are trained for a number of epochs, such that the total training time is the same as for the baseline methods

(thresholded to keep 0.9 of the mass) as in DINO (Caron et al. 2021), and report the Jaccard similarity between the ground truth and the mask evaluated on the validation set of PASCAL-VOC12 (Everingham et al. 2009). For the videos we use the DAVIS-2017 video instance segmentation benchmark (Pont-Tuset et al. 2017) and by following the protocol introduced in Space-time by Jabri et al. (Jabri, Owens, and Efros 2020) we segment scenes via the nearest neighbor propagation of the mask. In these evaluations, the role of the positional classification loss seems to be more important than sparsity alone.

Humanoid Vision Engine Benchmark. We also use the newly introduced HVE (Ge et al. 2022) to evaluate our shape bias in Table 6. In HVE Shape dataset, the input images are only the depth map of the foreground object which only contains shape information. We see that DILEMMA outperforms the base model which confirms our hypothesis that DILEMMA can focus on shape. For the HVE Texture, only four grey scaled random crops of the foreground object are concatenated and fed as input, so predicting the right class requires high texture discriminability. Results on HVE Texture show that DILEMMA’s better shape understanding did not harm the texture discriminability.

Robustness against Background Change. Following the background challenge evaluation metric (Xiao et al. 2020), we compute the classification accuracy of the model on a subset of ImageNet (IN-9) by changing the background and foreground. As shown in Table 7, in O/N.F. (Only/No Foreground), M.S/R/N. (Mixed Same/Random/Next), where the foreground is visible or accurately masked out, we outperform the base model. When the foreground is not visible (O.BB. (Only Background with foreground box Blacked out) and O.BT. (Only Background with foreground replaced with Tiled background)) the model performs correctly and does not just rely on the background for image classification.

Method	Background Change							Clean
	<i>M.N.</i> (\uparrow)	<i>M.R.</i> (\uparrow)	<i>M.S.</i> (\uparrow)	<i>N.F.</i> (\uparrow)	<i>O.BB.</i> (\downarrow)	<i>O.BT.</i> (\downarrow)	<i>O.F.</i> (\uparrow)	IN-9(\uparrow)
MoCoV3	64.52	65.68	78.57	38.69	9.41	10.67	77.80	91.65
+Sparsity	65.53	67.75	80.25	38.72	9.48	10.40	78.15	92.52
+DILEMMA	65.19	68.37	79.63	39.19	8.42	9.68	78.37	92.00
+Sparsity \uparrow	66.25	69.60	81.26	40.25	10.99	10.25	80.10	92.77
+DILEMMA \uparrow	68.86	71.16	81.85	40.40	8.69	10.64	82.42	93.46
DINO	65.56	68.94	79.95	33.28	9.70	9.90	80.99	92.17
+Sparsity	67.68	71.28	82.10	35.23	8.89	11.11	83.26	93.43
+DILEMMA	69.58	71.85	82.89	36.02	9.31	10.47	83.75	93.06
+DILEMMA \uparrow	69.38	73.75	82.94	38.54	8.81	9.90	84.69	93.93

Table 7: Robustness of pre-trained models against background changes. The \uparrow models are trained for a number of epochs, such that the total training time is the same as for the baseline methods

Sampling Method	<i>k</i> -NN	Linear
Importance Based	71.88	76.76
Random	73.98	77.78

Table 8: Token dropping policy. Results are evaluated on IN100

Sparsity	IN100		IN-1K	
	<i>k</i> -NN	Linear	<i>k</i> -NN	Linear
0% (Dense)	76.16	77.50	53.27	58.20
75%	73.98	77.78	52.99	57.90
Random	74.46	78.82	55.71	59.55

Table 9: Random Dropping Ratio. Results on the left are evaluated on IN100 and on the right on IN-1K

Ablations

In these experiments, we want to validate empirically a number of choices: 1) we ask how much the trained model is robust to occlusions (sparsity) and positional errors; 2) whether the selection of tokens should be random or guided; 3) whether the ratio of dropped tokens should remain constant in time or instead vary; 4) what the relevance of the positional classification loss is; 5) the impact of the number of positional errors used during training; 6) whether other design variations are more effective than DILEMMA.

Ablation studies are conducted either on ImageNet100 (IN100) or ImageNet-1K (IN-1K). For the smaller dataset we train the dense models for 300 epochs and the sparse models for 450 epochs (with the same hardware and time settings). For IN-1K experiments we train all models for 50 epochs with MoCoV3 unless stated otherwise.

Token Dropping Policy. In Table 8, we compare the case of dropping the tokens that are less important based on the attention of the teacher network (Li et al. 2021c) compared to randomly dropping the tokens. Results show that simple random dropping works well and there is no need to introduce extra complexity to the policy.

	IN-1K		Yoga82		MD Acc.
	<i>k</i> -NN	Linear	<i>k</i> -NN	Linear	
MoCoV3	53.27	58.20	31.60	51.27	-
+MD	54.18	58.78	35.78	54.53	100.00
+Sparsity	55.71	59.55	32.73	50.90	-
+Both	55.63	59.84	35.94	57.26	96.21

Table 10: Mismatch Detection (MD). Detecting misplaced tokens for dense inputs is easily solved, but still improves the model’s performance on shape based tasks. Note that adding both MD and Sparsity to the base model is the same as DILEMMA

θ	IN-1K		Yoga82	
	<i>k</i> -NN	Linear	<i>k</i> -NN	Linear
0.3	55.34	59.79	34.95	56.63
0.2	55.63	59.84	35.94	57.26

Table 11: Mismatch Probability. Too much mismatch hurts performance

Task	IN-1K		Yoga82	
	<i>k</i> -NN	Linear	<i>k</i> -NN	Linear
None (MoCoV3)	53.27	58.20	31.60	51.27
Pos. Correction	54.77	58.95	35.74	56.15
Partial Jigsaw	55.72	59.19	34.77	56.79
Flip Detection	55.69	59.59	35.09	55.00
DILEMMA	55.63	59.84	35.94	57.26

Table 12: Variants of the loss. Although the variants improve the performance wrt the baseline, DILEMMA is the most effective one

Randomized Dropping Ratio. In Table 9, we verify that a randomized dropping ratio is better than a constant one. We conducted two experiments: one on IN100 and one on IN-1K. The results show that a randomized dropping ratio performs better than a constant one. On the more difficult IN-1K dataset, just applying sparsity is worse than using the dense model. Only with a random drop-

Method	BatchSize	EpochTime	MaxMem(GB)
SimCLR	640	21:35	23.65
+DILEMMA	1680	18:20 ($\times 0.85$)	23.17
MoCoV3	656	49:08	23.41
+DILEMMA	1664	32:11 ($\times 0.65$)	23.55
DINO	576	37:57	22.73
+DILEMMA	1184	24:13 ($\times 0.64$)	22.79
DINO(MC) [†]	144	3:07:21	22.85
+DILEMMA [†]	216	2:15:52 ($\times 0.72$)	23.69

Table 13: Timing and memory usage of training ViT-Small models with four RTX Geforce 3090 GPUs. MC stands for Multi-Crop. [†] models use ViT-Base

ping ratio can the sparse model outperform the dense model.

Position Classification Loss. In Table 10, we verify that the position classification loss helps, by training a dense model with position mismatch detection. Surprisingly, even though the Mismatch Detection (MD) (*i.e.*, the average classification accuracy of the token locations – see “MD Acc.” in Table 10) is easily solved (it achieves 100% in the dense case), the dense model can still improve the performance of the model on a downstream task. The performance improvement for a task like in Yoga₈₂, which requires a better understanding of shape, is quite significant both with the dense and randomized sparsity inputs.

Mismatch Probability. The probability of a positional embedding mismatch θ is one of the hyper-parameters of DILEMMA. Early in our experiments, we found out that 20% is much better than 15% (which is used by Electra (Clark et al. 2020a)), probably due to the higher information redundancy in images compared to text. In Table 11, we show that $\theta = 30\%$ yields worse performance than the default $\theta = 20\%$.

DILEMMA Variants. We also tried some variants of DILEMMA. Instead of just detecting the misplaced tokens, we predict the right position (as a classification task of 196 classes). The other variant, which we call *Partial Jigsaw*, is to feed some tokens without position encoding and ask the network to predict their position given the other (sparse) correctly position-encoded tokens. Lastly, instead of corrupting the position, one can corrupt the content of a patch. Instead of using complex methods like inpainting we simply horizontally flip some of the patches and use the binary cross-entropy as our loss. Table 12 shows that even though all of these methods do help in terms of shape discrimination, DILEMMA is the one with the best performance both on IN-1K and Yoga₈₂.

Timing. To show the efficiency of the proposed method, we ran SimCLR, MoCoV3, DINO with and without multi-crop on 4 GPUs and reported the epoch times in table 13.

Combining with MAE. To show the general applicability of our proposed method to masked models, we misplaced

Method	Linear	Finetune
MAE	37.30	82.60
+DILEMMA	39.06	83.30

Table 14: Combining DILEMMA with MAE. Results are evaluated on IN100 after pretraining for 200 epochs and using ViT-Base

Method	300 Epochs	1000 Epochs
MoCoV3	77.50	79.76
+DILEMMA	78.82	81.26

Table 15: Longer Pretraining on IN-100. Linear accuracies on IN100 after pretraining for 300 and 1000 epochs

some of the MAE (He et al. 2021) inputs and added DILEMMA loss to the *encoder* of MAE in addition to the reconstruction loss of the decoder. Both MAE and DILEMMA are trained for 200 epochs on IN-100 (using the exact same hyperparameters of the official repository) and results in table 14 show that we can outperform MAE both in terms of linear probe and finetuning.

Longer Pretraining. We pretrain MoCoV3 and DILEMMA for 1000 epochs on IN-100 and evaluate their linear performance to see whether the benefits of DILEMMA still hold for longer pretrainings. Results in table 15 show that indeed DILEMMA always performs better than the baseline even with longer pretraining.

Weaker Data Augmentations. One of the most important factors for the performance of contrastive learners is the data augmentation. In this short experiment (50 epochs of pre-training, and 70 epochs of linear training) we only used random resized cropping (like MAE (He et al. 2021)) on IN-1K for both MoCoV3 and DILEMMA. Linear probe accuracy of DILEMMA is **44.48%** and for MoCoV3 it is **29.65%** (Note that a 100 epoch pretrained ResNet-50 (He et al. 2016) with SimCLR (Chen et al. 2020a) gets 33.1% accuracy). This huge gap shows that DILEMMA is a generic method for representation learning and does not completely depend on the contrastive component of the loss.

Conclusions

We introduced a novel SSL method based on a position classification pseudo-task and a contrastive loss. We showed that awareness of the relative location of tiles of the input image is important for generalization and in particular when finetuning on shape-based downstream tasks. Since our method is based on the ViT architecture, we introduce sparsity in the input (*i.e.*, dropping image tiles), to both speed up the training and also to avoid trivial degenerate learning.

Acknowledgments

This work was supported by grant 200020.188690 of the Swiss National Science Foundation (SNSF) and an Adobe award.

References

- Akbari, H.; Yuan, L.; Qian, R.; Chuang, W.-H.; Chang, S.-F.; Cui, Y.; and Gong, B. 2021. Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text. *Advances in Neural Information Processing Systems*, 34.
- Bao, H.; Dong, L.; and Wei, F. 2021. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*.
- Bossard, L.; Guillaumin, M.; and Van Gool, L. 2014. Food-101 – Mining Discriminative Components with Random Forests. In *European Conference on Computer Vision*.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33: 9912–9924.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9650–9660.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020a. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Chen, X.; Ding, M.; Wang, X.; Xin, Y.; Mo, S.; Wang, Y.; Han, S.; Luo, P.; Zeng, G.; and Wang, J. 2022. Context Autoencoder for Self-Supervised Representation Learning. *arXiv preprint arXiv:2202.03026*.
- Chen, X.; Fan, H.; Girshick, R.; and He, K. 2020b. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*.
- Chen, X.; and He, K. 2020. Exploring Simple Siamese Representation Learning. *arXiv preprint arXiv:2011.10566*.
- Chen, X.; Xie, S.; and He, K. 2021. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9640–9649.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; ; and Vedaldi, A. 2014. Describing Textures in the Wild. In *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Clark, K.; Luong, M.-T.; Le, Q. V.; and Manning, C. D. 2020a. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*.
- Clark, K.; Luong, M.-T.; Le, Q. V.; and Manning, C. D. 2020b. Pre-training transformers as energy-based cloze models. *arXiv preprint arXiv:2012.08561*.
- Coates, A.; Ng, A.; and Lee, H. 2011. An Analysis of Single-Layer Networks in Unsupervised Feature Learning. In *AISTATS*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Doersch, C.; Gupta, A.; and Efros, A. A. 2015. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, 1422–1430.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Dosovitskiy, A.; Fischer, P.; Springenberg, J. T.; Riedmiller, M.; and Brox, T. 2015. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(9): 1734–1747.
- Dwivedi, D.; Aytar, Y.; Tompson, J.; Sermanet, P.; and Zisserman, A. 2021. With a little help from my friends: Nearest-neighbor contrastive learning of visual representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9588–9597.
- El-Nouby, A.; Izacard, G.; Touvron, H.; Laptev, I.; Jegou, H.; and Grave, E. 2021. Are Large-scale Datasets Necessary for Self-Supervised Pre-training? *arXiv preprint arXiv:2112.10740*.
- Everingham, M.; Gool, L. V.; Williams, C. K. I.; Winn, J. M.; and Zisserman, A. 2009. The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision*, 88: 303–338.
- Fang, Y.; Dong, L.; Bao, H.; Wang, X.; and Wei, F. 2022. Corrupted Image Modeling for Self-Supervised Visual Pre-Training. *ArXiv*, abs/2202.03382.
- Fei-Fei, L.; Fergus, R.; and Perona, P. 2004. Learning Generative Visual Models from Few Training Examples: An Incremental Bayesian Approach Tested on 101 Object Categories. *2004 Conference on Computer Vision and Pattern Recognition Workshop*, 178–178.
- Ge, Y.; Xiao, Y.; Xu, Z.; Wang, X.; and Itti, L. 2022. Contributions of Shape, Texture, and Color in Visual Recognition. *arXiv preprint arXiv:2207.09510*.
- Geirhos, R.; Rubisch, P.; Michaelis, C.; Bethge, M.; Wichmann, F. A.; and Brendel, W. 2018. ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. *arXiv preprint arXiv:1811.12231*.
- Gidaris, S.; Singh, P.; and Komodakis, N. 2018. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*.
- Girshick, R.; Donahue, J.; Darrell, T.; and Malik, J. 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 580–587.
- Grill, J.-B.; Strub, F.; Altch’e, F.; Tallec, C.; Richemond, P. H.; Buchatskaya, E.; Doersch, C.; Pires, B. Á.; Guo, Z. D.; Azar, M. G.; Piot, B.; Kavukcuoglu, K.; Munos, R.; and Valko, M. 2020. Bootstrap Your Own Latent: A New Approach to Self-Supervised Learning. *ArXiv*, abs/2006.07733.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2021. Masked autoencoders are scalable vision learners. *arXiv preprint arXiv:2111.06377*.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9729–9738.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- iNaturalist. 2019. iNaturalist 2019 competition dataset. https://github.com/visipedia/inat_comp/tree/master/2019. Accessed: 2023-03-04.
- Ioffe, S.; and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, 448–456. PMLR.
- Jabri, A.; Owens, A.; and Efros, A. 2020. Space-time correspondence as a contrastive random walk. *Advances in neural information processing systems*, 33: 19545–19560.

- Jenni, S.; and Favaro, P. 2018. Self-supervised feature learning by learning to spot artifacts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2733–2742.
- Krause, J.; Stark, M.; Deng, J.; and Fei-Fei, L. 2013. 3D Object Representations for Fine-Grained Categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*. Sydney, Australia.
- Krizhevsky, A. 2009. Learning multiple layers of features from tiny images. Technical report, University of Toronto.
- Li, C.; Yang, J.; Zhang, P.; Gao, M.; Xiao, B.; Dai, X.; Yuan, L.; and Gao, J. 2021a. Efficient self-supervised vision transformers for representation learning. *arXiv preprint arXiv:2106.09785*.
- Li, X.; Zhou, Y.; Zhang, Y.; Zhang, A.; Wang, W.; Jiang, N.; Wu, H.; and Wang, W. 2021b. Dense semantic contrast for self-supervised visual representation learning. In *Proceedings of the 29th ACM International Conference on Multimedia*, 1368–1376.
- Li, Z.; Chen, Z.; Yang, F.; Li, W.; Zhu, Y.; Zhao, C.; Deng, R.; Wu, L.; Zhao, R.; Tang, M.; et al. 2021c. Mst: Masked self-supervised transformer for visual representation. *Advances in Neural Information Processing Systems*, 34.
- Liu, Y.; Sanginetto, E.; Bi, W.; Sebe, N.; Lepri, B.; and Nadai, M. 2021a. Efficient Training of Visual Transformers with Small Datasets. *Advances in Neural Information Processing Systems*, 34.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021b. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10012–10022.
- Maji, S.; Rahtu, E.; Kannala, J.; Blaschko, M. B.; and Vedaldi, A. 2013. Fine-Grained Visual Classification of Aircraft. *ArXiv*, abs/1306.5151.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; and Ng, A. Y. 2011. Reading Digits in Natural Images with Unsupervised Feature Learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*.
- Nilsback, M.-E.; and Zisserman, A. 2008. Automated Flower Classification over a Large Number of Classes. *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, 722–729.
- Noroosi, M.; and Favaro, P. 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, 69–84. Springer.
- O Pinheiro, P. O.; Almahairi, A.; Benmalek, R.; Golemo, F.; and Courville, A. C. 2020. Unsupervised learning of dense visual representations. *Advances in Neural Information Processing Systems*, 33: 4489–4500.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Parkhi, O. M.; Vedaldi, A.; Zisserman, A.; and Jawahar, C. V. 2012. Cats and Dogs. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Pathak, D.; Krähenbühl, P.; Donahue, J.; Darrell, T.; and Efros, A. 2016. Context Encoders: Feature Learning by Inpainting. In *CVPR*.
- Pont-Tuset, J.; Perazzi, F.; Caelles, S.; Arbeláez, P.; Sorkine-Hornung, A.; and Van Gool, L. 2017. The 2017 DAVIS Challenge on Video Object Segmentation. *arXiv:1704.00675*.
- Tamkin, A.; Liu, V.; Lu, R.; Fein, D.; Schultz, C.; and Goodman, N. 2021. DABS: A Domain-Agnostic Benchmark for Self-Supervised Learning. *arXiv preprint arXiv:2111.12062*.
- Tartaglini, A. R.; Vong, W. K.; and Lake, B. M. 2022. A Developmentally-Inspired Examination of Shape versus Texture Bias in Machines. *arXiv preprint arXiv:2202.08340*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Verma, M.; Kumawat, S.; Nakashima, Y.; and Raman, S. 2020. Yoga-82: a new dataset for fine-grained classification of human poses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 1038–1039.
- Wang, X.; Zhang, R.; Shen, C.; Kong, T.; and Li, L. 2021. Dense contrastive learning for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3024–3033.
- Wu, Z.; Xiong, Y.; Yu, S. X.; and Lin, D. 2018a. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3733–3742.
- Wu, Z.; Xiong, Y.; Yu, S. X.; and Lin, D. 2018b. Unsupervised Feature Learning via Non-parametric Instance Discrimination. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3733–3742.
- Xiao, K.; Engstrom, L.; Ilyas, A.; and Madry, A. 2020. Noise or signal: The role of image backgrounds in object recognition. *arXiv preprint arXiv:2006.09994*.
- Xiao, T.; Liu, Y.; Zhou, B.; Jiang, Y.; and Sun, J. 2018. Unified perceptual parsing for scene understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 418–434.
- Xiao, T.; Reed, C. J.; Wang, X.; Keutzer, K.; and Darrell, T. 2021. Region similarity representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10539–10548.
- Xie, Z.; Lin, Y.; Yao, Z.; Zhang, Z.; Dai, Q.; Cao, Y.; and Hu, H. 2021a. Self-supervised learning with swin transformers. *arXiv preprint arXiv:2105.04553*.
- Xie, Z.; Lin, Y.; Zhang, Z.; Cao, Y.; Lin, S.; and Hu, H. 2021b. Propagate yourself: Exploring pixel-level consistency for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16684–16693.
- Zhai, S.; Jaitly, N.; Ramapuram, J.; Busbridge, D.; Likhomanenko, T.; Cheng, J. Y.; Talbott, W. A.; Huang, C.; Goh, H.; and Susskind, J. M. 2022. Position Prediction as an Effective Pretraining Strategy. In *ICML*.
- Zhang, R.; Isola, P.; and Efros, A. A. 2016. Colorful image colorization. In *European Conference on Computer Vision*, 649–666. Springer.
- Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 633–641.
- Zhou, J.; Wei, C.; Wang, H.; Shen, W.; Xie, C.; Yuille, A.; and Kong, T. 2021. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*.