
You Don't Need Domain-Specific Data Augmentations When Scaling Self-Supervised Learning

Théo Moutakanni

FAIR at Meta
MICS, Université Paris-Saclay
theomoutakanni@meta.com

Maxime Oquab

FAIR at Meta

Marc Szafraniec

FAIR at Meta

Maria Vakalopoulou

MICS, CentraleSupélec,
Université Paris-Saclay

Piotr Bojanowski

FAIR at Meta

Abstract

Self-Supervised learning (SSL) with Joint-Embedding Architectures (JEA) has led to outstanding performances. All instantiations of this paradigm were trained using strong and well-established hand-crafted data augmentations, leading to the general belief that they are required for the proper training and performance of such models. On the other hand, generative reconstruction-based models such as BEIT and MAE or Joint-Embedding Predictive Architectures such as I-JEPA have shown strong performance without using data augmentations except masking. In this work, we challenge the importance of invariance and data-augmentation in JEAs at scale. By running a case-study on a recent SSL foundation model – DINOv2 – we show that strong image representations can be obtained with JEAs and only cropping without resizing provided the training data is large enough, reaching state-of-the-art results and using the least amount of augmentation in the literature. Through this study, we also discuss the impact of compute constraints on the outcomes of experimental deep learning research, showing that they can lead to very different conclusions.

1 Introduction

Self-supervised learning (SSL) has significantly improved performance across various tasks. Despite not requiring supervision, SSL heavily depends on carefully selected data augmentations [9, 11–13, 35]. Previous experimental literature has shown that removing even one augmentation such as random rescaling or color jittering reduces linear evaluation performance on ImageNet1k [15, 41] by at least 10 points [1, 11, 39]. But reliance on image-specific data augmentations limits the generalizability of self-supervised approaches. How can this be applied to graph, time-series or, for example, medical imaging with totally different channels and characteristics? According to multiple studies [1, 5, 21, 37], the best data augmentations are dependent on the target tasks, and no set of hand-crafted data augmentations can solve them all. But is this the case?

The problem is the following: SSL methods that do not build on such augmentations are based on image reconstruction, and are able to achieve competitive results, though only through fine-tuning ([4, 25]). On the other hand, joint-embedding architecture (JEA) methods, delivering strong results without fine-tuning and producing linearly separable features, seem largely reliant on data augmentations in light of the seminal SimCLR [11, 12] study. I-JEPA [2] was the first method to challenge this assumption, but the authors still show a significant performance gap when compared to their augmentation-based alternatives. Therefore, one can reasonably believe that the core principles driving JEA methods towards great performance are inherently dependent on data augmentation,

Table 1: **Comparison of our model** trained using RandomCrop, without random resizing nor other photometric augmentations against SSL models that do not leverage hand-crafted augmentations. All other models are reconstruction based, in the pixel space or in the latent space, and *use more augmentations* than our setup. We only use RandomCrop without resizing, and masking.

Method	Data	Arch.	ImageNet-1k			Other Classif.		Segm.	Depth
			val	V2	real	Places205	iNat18	ADE20k	NYUd ↓
AIM	DFN-2B+	H/14	70.0	43.7	64.4	51.7	64.0	13.2	0.862
BEIT	INet-1k	L/16	73.5	59.0	78.3	57.9	40.8	6.4	0.544
MAE	INet-1k	H/14	78.0	66.5	83.7	56.0	42.2	33.3	0.482
I-JEPA	INet-1k	H/14	79.3	66.0	83.8	58.4	47.6	31.1	0.537
I-JEPA	INet-22k	g/16	71.3	59.7	78.6	59.1	55.3	32.0	0.521
Ours	INet-22k	L/14	84.0	74.8	88.6	65.1	75.1	43.5	0.411

and the underlying invariance learning that it entails. As a side-effect, JEA setups have been extensively tuned to work with augmentations [9, 13, 35], cementing this belief in the community. In this context, the present paper successfully disproves the assumption that these augmentations are necessary. We show that it is possible to train **state-of-the-art** joint embedding architectures without augmentations other than crops **without random rescaling**, and optionally masking. To the best of our knowledge, we train such model using the **least amount of data-augmentations in the literature**. Data cropping without rescaling, along with masking, might be considered two of the only universal data augmentations that can be applied on a broad number of modalities by just using the properties of the data. We show that building such SoTA model simply requires a larger amount of diverse data, a longer training schedule, and more careful optimization.

For most methods, data augmentations prevent model collapse on trivial solutions such as color histogram [1, 7]. They can also enhance performance by artificially increasing the diversity and size of datasets using multiple crops per image [9, 10]. While preliminary work has attempted to explain the theory behind data augmentations and self-supervised learning or have conducted experimental analysis of their impact and necessity in SSL [1, 5, 11, 21, 37], such analysis has **not been conducted on large-scale SSL pretraining** with newer methods, restricting the conclusions that can be drawn from their results.

In this work, we aim to determine the role of data augmentations at scale by training models using DINOv2 [35] on multiple datasets of varying sizes (from 1M to 140M images), on multiple model sizes (from Small to Large), and using diverse data augmentation combinations.

Contributions. By doing a rigorous experimental analysis through this paper, we are the first to prove two counter-intuitive claims that go against most previous assumptions:

- (i) We show that the impact of data-augmentations invariance enforcement in self-supervised learning is secondary to the impact of their artificial increase in dataset size and distribution.
- (ii) We explore when and why we can remove hand-crafted data augmentations with minor impact on performance by looking at all scaling aspects of deep learning: data, compute and model size.

In the process, we also provide two by-products on top of our claims:

- (iii) We show through our experimental setup that scaling laws are ethereal and that practitioners optimizing for different compute and data budget might find very different conclusions.
- (iv) We achieve state-of-the-art performance on extensive evaluations for a model trained without hand-crafted augmentations. This result is an **existence proof** that joint-embedding architecture methods do not inherently rely on learning invariances defined by data augmentations, challenging the prior beliefs in the community and prompting a new theoretical understanding of SSL.

2 Related Work

Data augmentation in SSL practice. Data augmentations have been a key component of most self-supervised works, including the early Exemplar-CNN [18]. The network would be trained

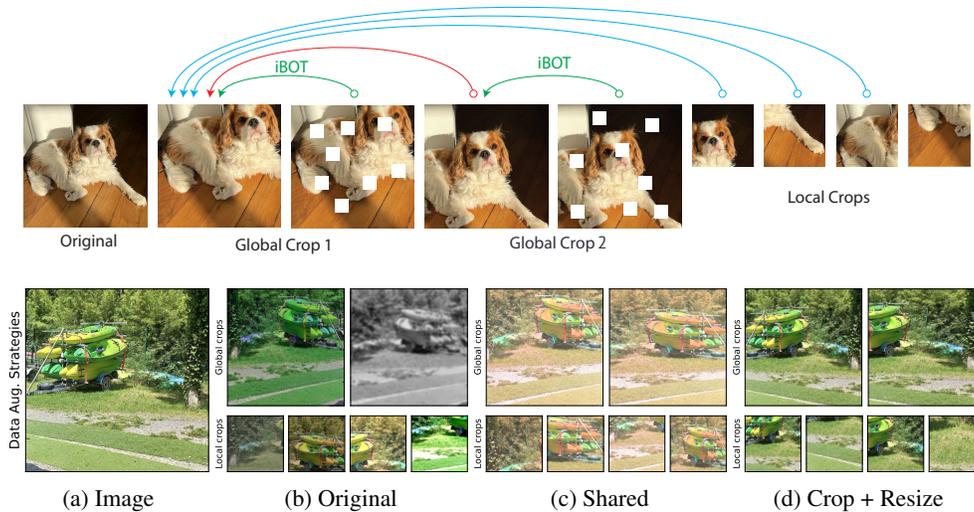


Figure 1: **Top: Visual description of pretraining losses.** In blue: the local to global DINO loss, in red: the global to global DINO loss and in green the latent masked token prediction (iBOT) loss. **Bottom: Our different augmentation strategies.** 'Original' uses several augmentations (RandomResizedCrop, ColorJitter, RandomGrayscale, GaussianBlur, RandomHorizontalFlip and RandomSolarize), 'Shared' uses the same augmentations but shares them between each view of the same image obtained with RandomResizedCrop. The 'Crop + Resize' setting only uses RandomResizedCrop. We also introduce a 'Crop' setup which uses RandomCrop without random rescaling and that is visually similar to 'Crop + Resize'.

using a classification objective, with each class corresponding to data augmentations of one image in the dataset. In the context of early pretext-task-based SSL, when training convolutional neural networks by predicting relative location of patches, Doersch, Gupta, and Efros [16] have observed trivial solutions. The network was overfitting to chromatic aberrations, and the authors used data augmentations to counter this effect. In a similar way, Bojanowski and Joulin [7] have observed that their discriminative model would emphasize trivial factors of variation such as colour. The solution proposed in this work was to discard color information and take the output of Sobel filtering as an input to the algorithm, along with image jitter and flipping. In 2018, the community observed a jump in performance when evaluating on Pascal VOC 2007 image classification [8, 22], the implementation of both works leveraging RandomResizedCrop instead of RandomCrop. The authors of SimCLR [11] carried out the first extensive ablation of the impact of data augmentation on performance, and many of the conclusions from this work *serve as standard in follow-up SSL work* [10, 24, 26, 35]: most modern joint-embedding self-supervised models learn using at least one loss term that pushes together the representation of two augmented views of the same image. By applying different data augmentations to the two views, they force the model to map the augmented characteristics to the same point, enforcing invariance in the model representation.

Sensitivity of learning to data augmentation. With the widely accepted view of the importance of invariance learning for SSL models, several works have studied the impact of the choice of data augmentations on downstream performance: Xiao et al. [52] show that the wrong choice of augmentations can have a detrimental effect; Tian et al. [46] consider data augmentations as a way to reduce the mutual information between different views of an image to improve results. Bendidi et al. [5] show through experiments that the choice, amplitude and combination of transformations effects the learnt representation performance in downstream tasks and that the benefit of data-augmentations depend on each specific class. Purushwalkam and Gupta [37] show that a poor choice of cropping parameters can harm performances.

At the same time, self-supervised learning algorithms based on image reconstruction don't use many data augmentations. Models trained to reconstruct missing patches directly in the pixel space [4, 25, 33] have led to significantly lower linear-evaluation performance but did not leverage any photometric or hand-crafted augmentations. El-Nouby et al. [32] show that BEIT is more robust to smaller dataset size. Recently, I-JEPA [2], based on reconstruction in the feature space obtained strong performance

without using photometric augmentations. The authors split the comparison against other SSL frameworks according to the type of augmentations used (see Table 1 in [2]), suggesting data augmentations make for an unfair advantage in SSL training. It is worth noting that *reconstruction based methods still use random resizing* in their augmentations.

Theoretical studies. Furthermore, a series of theoretical works have been studying the *apparently critical* data augmentation for SSL algorithms. For example, Kùgelgen et al. [30] prove, under some assumptions, that data augmentations allow isolating content from style, when using InfoNCE as an objective, as in SimCLR. Eastwood et al. [19] follow up on this work, and propose using a structured data augmentation framework in order to disentangle content from style in the learnt representations; the goal is to avoid discarding style information during the learning process. Zhang and Ma [53] propose splitting the multiple operations involved in data augmentation pipelines in order to produce a hierarchy of augmentations: invariance to each augmentation is translated into a loss applied at different stages of the backbone, leading to accuracy improvements. Notably, they mention that “*The most indispensable process in contrastive learning is the data augmentation module.*” Saunshi et al. [42] propose seeing data augmentations as inductive biases applied to the learning algorithm, and derive theorems valid in the case of linear representations, in order to further understand the principles driving the success of contrastive learning.

In this work, we provide a critical result regarding the importance of data augmentation, that dramatically improves our understanding of the core principles of self-supervised learning, prompting renewed theory: **data augmentations are not necessary and only cropping without resizing is sufficient to train strong SSL joint-embedding architecture models.**

3 Approach

The goal of this work is to study the importance of data augmentation for training joint embedding architectures. In this section, we provide a quick recap of the SSL training algorithm that we use for our study. We then provide a precise description of the experimental setup, with details about the training regimes and data augmentation strategies.

3.1 Short description of the DINOv2 algorithm

We tailor our study around the algorithm that was used for training the recently proposed DINOv2 [35] models. We hypothesize that learnings from this study should apply to other SSL methods, but we chose to focus our experimental work on the best modern JEA method. That way, we can provide a strong counter-argument to the assumptions used in previous literature. We briefly summarize the components of the DINOv2 learning algorithm, and for more details, we point to the original publication [35]. Figure 1 provides a schematic diagram of the loss terms used for training the model.

Global DINO loss. The DINOv2 training loop is largely based on the DINO algorithm [9]. Given a student (s) and a teacher (t) network, and two views of the same image v_1 and v_2 , the student network is trained by minimizing a loss between the two representations:

$$\ell(s(v_1), t(v_2)). \quad (1)$$

The teacher network, is taken as a moving average of the student. We refer the reader to [9] for more details about the definition of ℓ .

Local reconstruction-based iBOT loss. DINOv2 models were trained using an additional local masked-image reconstruction loss. As opposed to MAE or BEIT, the reconstruction task in iBOT happens in the latent space, not in the pixel input space. Given a view v_1 as used in the DINO loss, we create an impaired version \tilde{v}_1 by replacing a fixed percentage of patches with a [MASK] token. The target is to reconstruct the value of masked features using the remaining context:

$$\ell(s(\tilde{v}_1), t(v_1)). \quad (2)$$

If several views are used, the iBOT loss can be applied to every such view. The strategy for choosing the masked patches is a hyper-parameter of the method.

Multicrop and data augmentations. In practice, one does not use only two views v_1 and v_2 . Following [10], DINOv2 uses Multicrop, with two global crops and eight local crops. Global crops are obtained using RandomResizedCrop with a `scale` parameter of $[0.32, 1.0]$, then resized to 224 pixels. For local crops we used $[0.05, 0.32]$ resized to 98 pixels instead. Each crop, be it global

or local, undergoes a series of photometric augmentations: ColorJitter, RandomGrayscale, GaussianBlur, RandomHorizontalFlip and RandomSolarize. We refer to the official DINOv2 repository for the exact implementation details.¹

3.2 Experimental setup

Hyperparameters and training regimes. The original DINOv2 repository proposes two sets of hyperparameters for training SSL models. The first set (that we refer to as the *low-compute regime*) corresponds to a setup for fast experimental iterations, designed to run for 100 epochs (125k iterations) with a batch size of 2048. This setup is optimized for performance on the ImageNet-1k dataset (corresponding to the *low-data regime*). The second set (*high-compute regime*) is designed for longer training runs of 500 epochs (625k iterations) and is optimized for performance on larger datasets, such as ImageNet-22k (the *high-data regime*). These two sets differ in many values, including notably the learning rate schedule and warmup, weight decay, batch size, patch size. In our experiments, we use both sets to compare the behavior of the algorithm against data augmentation. We provide a detailed empirical discussion of the differences in Sec. 4.3.

Modifications applied on data augmentation. In our study, we apply different configurations of data augmentations. We provide here the details of the four configurations that we consider.

- **Original:** we use all the data augmentations, exactly like in the original DINOv2 work. Each crop is passed through a series of random photometric augmentations independently.
- **Shared:** each view of an image is obtained using a different crop, but apply the exact same photometric augmentation to all views. Two images in the batch will have different photometric augmentations applied to them.
- **Crop + Resize:** we create views with no photometric augmentations at all, only `RandomResizedCrop`.
- **Crop:** we create views without any photometric augmentations at all. We also don't use `RandomResizedCrop`. Instead, we use a `Resize` to 256×256 pixels and `Crop` to 224 or 98 pixels, leading to views with the exact same aspect ratio and zoom, with no random resizing. This setup is very close to a masking strategy where different sets of tokens would be dropped, up to pixel alignment to patch boundaries.

Figure 1 provides an illustration of the images that we obtain by applying these configurations. We omit the 'Crop' setting from the illustration given that the 'Crop' and 'Crop+Resize' provide images that are hard to qualitatively differentiate with a naked eye, even if they differ in term of final performances. In the end, our 'Crop' strategy only applies two augmentations between generated views: cropping without random resizing and random patch masking, which are two very similar augmentations (the cropping drops pixels outside of a field-of-view, and masking inside of it). We chose to keep those two specific augmentations as we think that they can be generalized to nearly all kind of data, from graphs to 3D medical images and time-series.

Implementation details. For our study, we use the standard ImageNet-1k [41] and ImageNet-22k [15] datasets, as well as the LVD-142M dataset originally used in DINOv2 [35]. The selection of these datasets allows us to benchmark the impact of their size in the trainings. The pre-training code performs 100 epochs in 27 hours on 5 compute nodes of 8 A100-80GB GPUs each, where 1 epoch is set to 1250 iterations. For evaluations, we follow DINOv2 linear evaluation protocol on multiple different tasks including classification (ImageNet-1k [41], Places205[55], ImageNet-v2 [38], ImageNet-Real [6], iNaturalist'18 [49], FlickrLogos32 [40], GTSRB [45], RP2K [36], Products-10k [3], FireRisk [43], RESISC [14], MIMIC [28], CheXpert [27], VinDr [31], NIH-14 [51]), segmentation (ADE20k [55]) and depth estimation (NYU-Depth v2 [44]).

Table 2: **New domains classification results** of DINOv2 ViT-L trained on ImageNet-22k when varying data augmentations. None of those domains were used in the pretraining data of the models.

Domain (Metric) Dataset	Remote Sensing (Acc.)		Medical Imaging (AUROC)			
	FireRisk	RESISC	MIMIC	CheXpert	VinDr	NIH-14
Original	59.0 =	91.9 =	75.3 =	83.3 =	80.0 =	72.5 =
Shared	60.8 \uparrow 1.8	91.6 \downarrow 0.2	75.8 \uparrow 0.5	85.6 \uparrow 2.3	81.0 \uparrow 1.0	74.4 \uparrow 1.9
Crop+Resize	60.7 \uparrow 1.7	91.1 \downarrow 0.8	75.8 \uparrow 0.5	84.1 \uparrow 0.8	80.8 \uparrow 0.8	73.7 \uparrow 1.2
Crop	60.9 \uparrow 1.9	92.3 \uparrow 0.4	75.5 \uparrow 0.2	83.0 \downarrow 0.3	82.5 \uparrow 2.5	74.2 \uparrow 1.7

4 Experiments and discussion

4.1 Why do we need augmentations?

The first question we want to answer is: why do we need data augmentations? Are they systematically needed, because they constitute a core component of the modeling? Are they a trick that facilitates the optimisation of deep neural networks, exactly like in supervised learning?

The role of data-augmentations. In the context of supervised learning, data augmentations have been used as a means of virtually augmenting the training dataset [21]. That way, additional training samples could be obtained, reducing the risks of overfitting and improving the robustness of the models trained. In self-supervised learning however, data-augmentations have been usually modeled as a way to enforce representations to be invariant toward specific image’s characteristics, such as color or scale [46, 52, 53].

Invariance can be harmful. Most of the self-supervised methods and their data-augmentations have been optimised on ImageNet. While DINOv2 has probably been optimised on more benchmarks, some tasks and domain are missing. In Table 2 we show the results of removing hand-crafted data-augmentations on new domains (not used for training) and in Table 3 (left) on new tasks (in the training set but not used for hyperparameter search). We can see that data-augmentations are actually harmful on most of those, meaning that understanding the real impact of augmentations and invariance while removing them is very important to increase SSL robustness.

Disambiguating the impact of data augmentation in self-supervised learning. We developed the ‘Shared’ configuration that still applies augmentations to artificially increase the number of training samples but *doesn’t enforce invariance* to those augmentations. We compare invariance in Table 3 (right) and show that the ‘Shared’ setup has effectively lower invariance than the ‘Original’ one. We quantify it using the average cosine similarity of 16 augmented views of the same image, repeated on 100 images from different classes. Higher cosine similarity means higher invariance. We also report a mean/std normalized similarity, computing the metrics using negative pairs’ cosine similarity.

Enforcing invariance losses its usefulness at scale. By comparing the ‘Shared’ and ‘Original’ results in Fig. 2, we can see that both setups reach way above 80% linear-probe accuracy on ImageNet1k when trained on all three datasets. While the ‘Original’ setup gets better performances, we can see

¹<https://github.com/facebookresearch/dinov2> (Apache-2.0 license)

Table 3: **(left): New task classification results** of DINOv2 ViT-L trained on ImageNet-22k when varying data augmentations. None of those tasks were used to tune DINOv2’s hyperparameters. **(right): Measure of invariance toward augmentation.** Higher cosine similarity means higher invariance as the model embeds multiple augmentations of the same image to closer vectors.

Task Dataset	Logos	Signs	Products		Metric	Cosine Similarity	Normalized Cos. Sim.
	FlickrLogos32	GTSRB	RP2K	Products-10k			
Original	80.1 =	80.3 =	92.8 =	66.5 =	Original	0.69	2.69
Shared	78.2 \downarrow 1.9	78.9 \downarrow 1.4	93.4 \uparrow 0.6	68.1 \uparrow 1.6	Shared	0.64	2.62
C+R	80.8 \uparrow 0.7	76.3 \downarrow 4.0	93.3 \uparrow 0.5	67.7 \uparrow 1.2	C+R	0.58	2.47
Crop	76.9 \downarrow 3.2	85.1 \uparrow 4.8	93.4 \uparrow 0.6	67.3 \uparrow 0.8	Crop	0.56	2.44

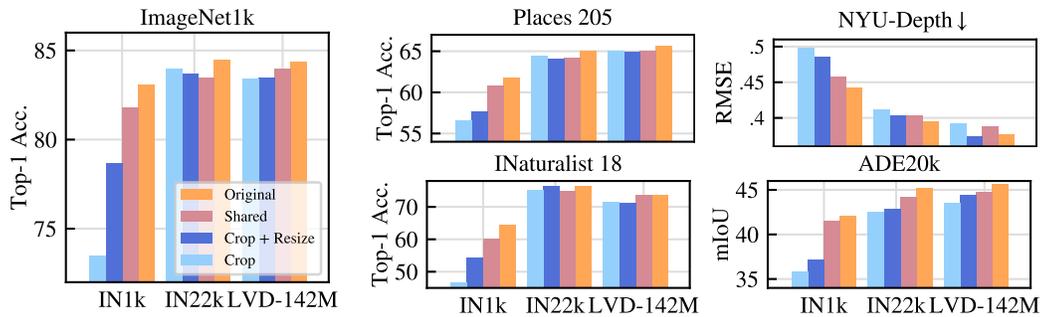


Figure 2: **Impact of dataset size when varying data augmentations.** Results of ViT-L on linear evaluation benchmarks, including classification (ImageNet1k, Places 205 and INaturalist18), depth estimation (NYU-Depth) and segmentation (ADE20k). Cropping without resizing ('Crop') reaches very high performances on a wide variety of benchmarks, given that the dataset size is large enough.

that the gap decreases with the amount of data, from -1.3% when trained on ImageNet1k and -1.0% with ImageNet22k, to -0.4% with LVD-142M. As we show in Fig. 3 (left), the gap between the 'Shared' and 'Original' setups also decreases with the number of training epochs. According to experiments in prior literature [1, 11, 39], data augmentations are required at the sample level to enforce invariance and obtain competitive performance. For example, when trained and evaluated on ImageNet1k, SimCLR and BYOL respectively lose 27.6 and 13.1 points of accuracy in the 'Crop+Resize' setup [39]. In contrast, our 'Shared' setup only loses 1.2%, disproving this view.

Increasing the sample count is the key to strong SSL. We know that reconstruction-based models like BEIT, MAE or I-JEPA don't need handcrafted data-augmentations on ImageNet1k, and according to El-Nouby et al. [32], "*denoising autoencoders are more sample efficient than joint embedding techniques*". Those findings led us to conjecture that the real impact of data-augmentation is to artificially increase the number of samples and allow JEA to reach good performance with less data.

To assess this, we compare the 'Crop + Resize' and the 'Shared' settings. Both settings do not enforce invariance, but they differ because the 'Shared' method artificially increases the sample count of the pretraining dataset. We show in Fig. 2 that the gap with respect to the 'Original' setting is way bigger for 'Crop + Resize' setting the compared to the 'Shared' one when we use ImageNet1k to pretrain our models. Without the sample count increase of data-augmentations, DINOv2 reaches lower performance on all five benchmarks when pretrained on a small dataset (Fig. 2).

Data-augmentations play the same role in supervised and self-supervised learning. All those results lead to the same conclusion: invariance is useful but not necessary for the learning of JEA, and augmentations artificially increase the number of samples in the pretraining dataset. This also explains why reconstruction-based methods do not need as many augmentations: they are more data efficient [32]. But this is not their only role, as data-augmented setups perform better in shorter trainings (Fig. 3, left). Since setups eventually converge to the same result with sufficient data, this

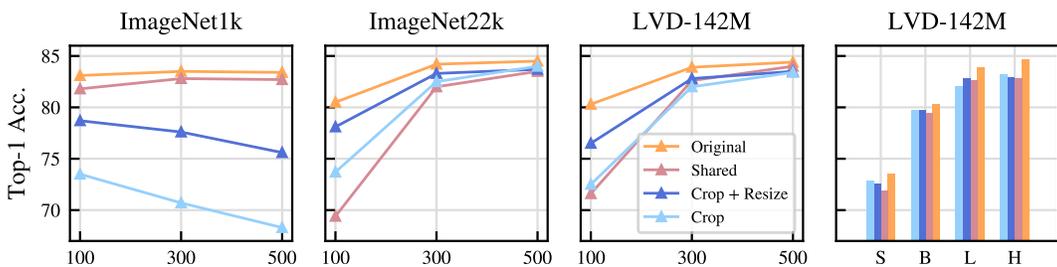


Figure 3: **Impact of data augmentations when we scale the number of training epoch (left) or the ViT architecture size (right)** on the accuracy of a linear probe on ImageNet1k for a ViT-L when pretraining on ImageNet1k, ImageNet22k and LVD-142M. The 'Original' and 'Shared' setups scale with the number of epochs for all datasets, but the 'Crop' and 'Crop+Resize' setups only scale with larger datasets.

suggests data augmentations are likely linked to easier optimization rather than to a core learning signal. They are just another hyper-parameter.

4.2 Can we remove hand-crafted augmentations totally?

We focused on what was the impact of data augmentations in modern self-supervised learning using the 'Original', 'Shared' and 'Crop + Resize' settings. We will now show that we can reach state-of-the-art performance without hand-crafted data-augmentations and resizing.

Scaling dataset size. Given that DINOv2 is more robust than other methods when training on ImageNet1k without handcrafted data-augmentations [39], and knowing from Sec. 4.1 that data augmentations mainly increase artificially the dataset size, we hypothesize that a self-supervised JEA can reach SOTA performance by increasing the size of the pretraining dataset used, even when removing the random resizing. We can see in Fig. 2 that increasing the dataset size has a notable property: all four settings converge to the same performances and the gap almost disappears between settings with and without hand-crafted augmentations, displaying a non-decreasing scaling curve with the number of samples. The linear-probe accuracy gap decreases below 1% on ImageNet1k and remains small on other benchmarks.

Scaling pretraining epochs. The highest performances were obtained when we pretrained our models for 500 epochs on larger datasets and 100 on the smaller one. However, a question arises: what happens if we train for longer on ImageNet1k, and faster on ImageNet22k and LVD-142M. In Fig. 3 (left) we show that the behavior of low and high data regimes differ. While the larger datasets show non-decreasing scaling curves with respect to the number of epochs, ImageNet1k has an uncoupling between the settings with and without data-augmentations. We notice that training longer on smaller scale dataset is harmful for performances when we don't use the dataset size increase property of data-augmentations. This is probably a consequence of overfitting and another proof of our first claim in Sec. 4.1. What is more compelling is that, following the 'Shared' setting scaling curve, the 'Crop+Resize' and 'Crop' settings' performance increases with longer pretraining, reducing the gap with the 'Original' one and converging to the same performance.

Data augmentation and model capacity. In Fig. 3 (right), we report ImageNet1k linear results when varying both data-augmentation strategies and model size. We can see that the gap between the best setup without hand-crafted data-augmentations and the classical setup ('Crop+Resize' or 'Crop' versus 'Original') increases when we use bigger model size, going from -0.6% for a ViT-S and ViT-B to -1.1% for a ViT-L and -1.5% for a ViT-H. This means that as we scale our model, we need more augmentation to increase the amount of data and reduce overfitting. This is in line with the most common understanding of scaling laws in experimental deep learning [29]: we need to scale compute, model size and data at the same time. This also explains why previous work [11, 39] had a hard time scaling self-supervised JEA without data-augmentations. They needed larger dataset, more compute and larger models, but all previous experimental studies on the impact of data-augmentations in SSL were mainly applied to ImageNet1k with smaller ResNets. This also explains why data-augmentation strategies became more aggressive during the evolution of self-supervised methods from its inception to modern methods like DINOv2, when people were scaling models but not datasets.

Role of the 'reconstruction-based' iBOT loss. While DINO and iBOT use the exact same loss, they differ in terms of target. With DINO, the teacher and the student are looking at two different views of the same image, with either different data augmentations in the 'Original' setting, the exact same in the 'Shared' one, or no augmentation at all with the 'Crop' one. In the iBOT case however, the student and the teacher see the same exact augmented view, but the student has some masked tokens.

Table 4: **Impact of the iBOT loss** on linear evaluation for multiple datasets for a ViT-L trained for 500 epochs on LVD-142M. We compare results with and without using masking and the iBOT loss.

		IN1k		INat18		ADE20k		NYU-Depth ↓	
iBOT + masking		✗	✓	✗	✓	✗	✓	✗	✓
Data Aug.	Original	84.0	84.4	74.7	73.6	41.4	45.6	0.430	0.413
	Shared	82.6	84.0	74.1	73.6	40.7	44.8	0.448	0.414
	Crop+Resize	82.6	83.5	71.3	71.1	40.3	44.4	0.437	0.415
	Crop	82.1	83.4	69.8	71.4	40.6	43.5	0.445	0.417

It means that regardless of the data-augmentation strategy we use, both the teacher and the student see the same thing in the larger crops. Both settings are similar, and despite iBOT not using a conditional latent variable, it is possible to see it as a masked image modeling task in latent space, which is more akin to joint embedding predictive architecture (JEPA) than joint embedding architecture (JEA). To remove all confounding factors, we ablate the iBOT loss by setting its weight to zero and by removing the masking strategy in Table 4. We can see that it is still possible to train a very good model without iBOT and masking even if we use a stricter definition of JEA. The gap between the 'Original' and 'Crop' settings increases slightly, eg. from 1.0% with iBOT to 1.9% without, for ImageNet1k linear probing, or from 0.04 to 0.015 RMSE for NYUd, but the performances remain competitive; more importantly, it is in fact possible to train DINO alone without hand-crafted data-augmentations.

Comparison against other models. We compare our model trained on ImageNet22k with the 'Crop' setting against other SOTA models trained without hand-crafted augmentations in Table 1. It's worth noting that our setup uses the smaller amount of augmentations: AIM [34] uses 'Crop+Resize' and RandomHorizontalFlip, BEIT [4, 50] uses 'Crop+Resize' and ColorJitter, MAE [25] uses 'Crop+Resize' and I-JEPA [2] uses 'Crop+Resize'. Our model outperforms alternatives on a wide variety of tasks by a significant margin without hand-crafted data-augmentations.

4.3 Is scaling all you need in research?

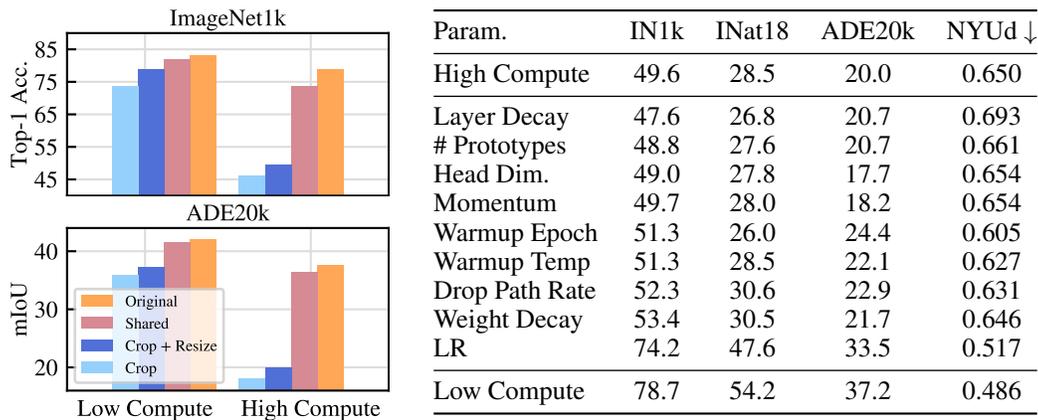


Figure 4: **(left): Impact of hyper-parameter optimisation's target compute** on the accuracy of a linear probe on ImageNet1k and ADE20k for models trained on ImageNet1k. We can see that optimising for high compute leads to poor performances on the 'Crop' and 'Crop+Resize' settings, which is the opposite of our findings when we optimize for low compute. **(right): Impact of hyperparameter tuning** using 'Crop+Resize' on ImageNet1k for 100 epochs. For each line, we switch only one hyper-parameter from the configuration optimized on ImageNet22k for 500 epochs (High compute) to the one optimized for ImageNet1k for 100 epochs (Low).

Same experiment, different outcomes. During this study, we used *two sets of hyper-parameters optimized for different use cases*. One optimized for training our models on the smaller scale dataset ImageNet1k during 100 epochs (low compute) that we used for all trainings on ImageNet1k, and one for training those models on the larger datasets ImageNet22k and LVD-142M during 500 epochs (high compute) that we used for all trainings on ImageNet22k and LVD-142M. Indeed, training on ImageNet1k using hyper-parameters optimized on the larger datasets gives different conclusions.

We show in Fig. 4 (left) the ImageNet1k linear-probe accuracy and ADE20k linear segmentation mIoU obtained after training a ViT-L on ImageNet1k when we use the 'Low Compute' and 'High Compute' sets of hyper-parameters. Interestingly, DINOv2 doesn't work when trained on ImageNet1k using the 'Crop' or 'Crop+Resize' settings with the 'High Compute' hyper-parameters. It reaches only 46% linear accuracy, despite obtaining SoTA performances when trained on ImageNet22k and LVD-142M with the exact same hyper-parameters, providing nice scaling curves in the process. However, when we optimize the hyper-parameters for the specific scale of data (ImageNet1k) and compute (100 epochs), we can see that it reaches 73.5% with a linear-probe, as presented before.

With the 'High Compute' hyperparameters, going from the 'Original' to 'Crop+Resize' setups makes DINOv2 lose 29.1% on the linear-probe evaluation, which is in phase with previous literature [39]. This means that if our only goal was to get a scaling curve leading to SoTA models, without optimizing for lower compute and data, we wouldn't have been able to disprove the assumption from previous studies that pretraining a strong JEA without data-augmentations with minimal loss in performances is impossible when using ImageNet1k [11, 39]).

The issue with high-scale experimental studies. In Fig. 4 (right), we show that it's not trivial to go from one hyper-parameter setup to the other, gains being non linear when tweaking parameters one at a time. Such issue already arose previously in the literature. For example, it was originally thought that ViT needed hundreds of millions of images to reach good performances [17], but such claim has been proven false in later work [20, 48]. The same happened with Large Language Models, with work scaling them to hundred of billions of parameters like OPT [54] reaching 175B before newer methods showed improved performances with smaller sizes as in [23, 47].

5 Conclusion

In this work, we challenged the common belief that joint-embedding architectures require data augmentations to deliver strong performance. The underlying interpretation of that belief, followed by theorists, is that the core learning signal in these methods relies on mapping differently augmented samples to the same embedding in the latent space. In particular, this reasoning implies expert knowledge in the design of data augmentations. Also, the photometric augmentations of images (such as blurring and color jittering) cannot be easily mapped to other modalities (such as speech or text). As such, this belief restricts the scope of high-performance JEA methods to images.

Our experiments show that this is not true, as it is possible to achieve performance similar to data-augmented pipelines with the DINOv2 SSL algorithm without using what is referred to, in prior literature, as *hand-crafted data augmentations*. Therefore, our results prove that joint-embedding architectures do not necessarily require domain knowledge during training to deliver state-of-the-art performance. The data augmentations merely influence training by increasing the dataset size, as shown by our 'Shared' setup, which does not enforce the learning of invariances.

Additional experiments with even weaker data augmentations, such as the 'Crop' setup, show that given enough training iterations and data, the models will converge to comparable performance as their stronger data-augmented counterparts. Our results suggest that prior conclusions largely depended on the smaller scale of experiments. However, we note that strong data augmentations lead to stronger performance in the case of shorter training, suggesting that they also help ease the optimization process during learning. We hope future work can shed more light on that observation.

Limitations. Given the cost of pretraining models, we focused on the leading self-supervised joint-embedding architecture, DINOv2, that provides strong performance at different data scales. While the results might be different for other algorithms, in principle, the existence proof that we offer in this work still holds. We noted that data augmentations tend to ease optimization, but the underlying mechanism is not known at this point. We merely claim that handcrafted photometric augmentations are optional for successful learning. While it has a measurable effect on some setups, this does not invalidate our conclusions. For large architectures, there is still a tiny gap in performance between augmented and non-augmented setups for long training. While we hypothesize that this might be due to slight variations in the optimal hyperparameters, there could also be a deeper reason that is not covered within the scope of this study.

Statement of Broader Impact. The total compute cost for this study was approximately 150k GPU-hours. We used strong models to blur all human faces in our web-based image data pool.

References

- [1] Y. M. Asano, C. Rupprecht, and A. Vedaldi. *A critical analysis of self-supervision, or what we can learn from a single image*. 2020. arXiv: 1904.13132 [cs.CV].
- [2] M. Assran et al. “Self-Supervised Learning from Images with a Joint-Embedding Predictive Architecture”. In: *arXiv preprint arXiv:2301.08243* (2023).
- [3] Y. Bai et al. *Products-10K: A Large-scale Product Recognition Dataset*. 2020. arXiv: 2008.10545 [cs.CV].
- [4] H. Bao, L. Dong, and F. Wei. “BEiT: BERT Pre-Training of Image Transformers”. In: *arXiv preprint arXiv:2106.08254* (2021).
- [5] I. Bendidi et al. *No Free Lunch in Self Supervised Representation Learning*. 2023. arXiv: 2304.11718 [cs.CV].
- [6] L. Beyer et al. “Are we done with ImageNet?” In: *CoRR abs/2006.07159* (2020). arXiv: 2006.07159.
- [7] P. Bojanowski and A. Joulin. “Unsupervised learning by predicting Noise”. In: *ICML*. 2017.
- [8] M. Caron et al. “Deep clustering for unsupervised learning of visual features”. In: *ECCV*. 2018.
- [9] M. Caron et al. “Emerging properties in self-supervised vision transformers”. In: *ICCV*. 2021.
- [10] M. Caron et al. “Unsupervised learning of visual features by contrasting cluster assignments”. In: *NeurIPS*. 2020.
- [11] T. Chen et al. “A simple framework for contrastive learning of visual representations”. In: *preprint arXiv:2002.05709* (2020).
- [12] T. Chen et al. “Big self-supervised models are strong semi-supervised learners”. In: *Advances in neural information processing systems* 33 (2020), pp. 22243–22255.
- [13] X. Chen et al. “Improved baselines with momentum contrastive learning”. In: *preprint arXiv:2003.04297* (2020).
- [14] G. Cheng, J. Han, and X. Lu. “Remote Sensing Image Scene Classification: Benchmark and State of the Art”. In: *Proceedings of the IEEE* 105.10 (Oct. 2017), pp. 1865–1883. ISSN: 1558-2256. DOI: 10.1109/jproc.2017.2675998.
- [15] J. Deng et al. “ImageNet: A Large-Scale Hierarchical Image Database”. In: *CVPR09*. 2009.
- [16] C. Doersch, A. Gupta, and A. A. Efros. “Unsupervised visual representation learning by context prediction”. In: *ICCV*. 2015.
- [17] A. Dosovitskiy et al. “An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale”. In: *preprint arXiv:2010.11929* (2020).
- [18] A. Dosovitskiy et al. “Discriminative unsupervised feature learning with exemplar convolutional neural networks”. In: *TPAMI* (2016).
- [19] C. Eastwood et al. “Self-Supervised Disentanglement by Leveraging Structure in Data Augmentations”. In: *arXiv preprint arXiv:2311.08815* (2023).
- [20] H. Gani, M. Naseer, and M. Yaqub. *How to Train Vision Transformer on Small-scale Datasets?* 2022. arXiv: 2210.07240 [cs.CV].
- [21] J. Geiping et al. *How Much Data Are Augmentations Worth? An Investigation into Scaling Laws, Invariance, and Implicit Regularization*. 2023. arXiv: 2210.06441 [cs.LG].
- [22] S. Gidaris, P. Singh, and N. Komodakis. *Unsupervised Representation Learning by Predicting Image Rotations*. 2018. arXiv: 1803.07728 [cs.CV].
- [23] F. Gloeckle et al. *Better & Faster Large Language Models via Multi-token Prediction*. 2024. arXiv: 2404.19737 [cs.CL].
- [24] J.-B. Grill et al. “Bootstrap your own latent: A new approach to self-supervised learning”. In: *NeurIPS*. 2020.
- [25] K. He et al. “Masked autoencoders are scalable vision learners”. In: *arXiv preprint arXiv:2111.06377* (2021).
- [26] K. He et al. “Momentum contrast for unsupervised visual representation learning”. In: *CVPR*. 2020.
- [27] J. Irvin et al. *CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison*. 2019. arXiv: 1901.07031 [cs.CV].

- [28] A. E. W. Johnson et al. “MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports”. In: *Scientific Data* 6.1 (Dec. 2019), p. 317. ISSN: 2052-4463. DOI: 10.1038/s41597-019-0322-0.
- [29] J. Kaplan et al. *Scaling Laws for Neural Language Models*. 2020. arXiv: 2001.08361 [cs.LG].
- [30] J. von Kügelgen et al. *Self-Supervised Learning with Data Augmentations Provably Isolates Content from Style*. 2022. arXiv: 2106.04619 [stat.ML].
- [31] H. Q. Nguyen et al. *VinDr-CXR: An open dataset of chest X-rays with radiologist’s annotations*. 2020. arXiv: 2012.15029 [eess.IV].
- [32] A. El-Nouby et al. *Are Large-scale Datasets Necessary for Self-Supervised Pre-training?* 2021. arXiv: 2112.10740 [cs.CV].
- [33] A. El-Nouby et al. “Scalable Pre-training of Large Autoregressive Image Models”. In: *arXiv preprint arXiv:2401.08541* (2024).
- [34] A. El-Nouby et al. *Scalable Pre-training of Large Autoregressive Image Models*. 2024. arXiv: 2401.08541 [cs.CV].
- [35] M. Oquab et al. “Dinov2: Learning robust visual features without supervision”. In: *arXiv preprint arXiv:2304.07193* (2023).
- [36] J. Peng, C. Xiao, and Y. Li. *RP2K: A Large-Scale Retail Product Dataset for Fine-Grained Image Classification*. 2021. arXiv: 2006.12634 [cs.CV].
- [37] S. Purushwalkam and A. Gupta. *Demystifying Contrastive Self-Supervised Learning: Invariances, Augmentations and Dataset Biases*. 2020. arXiv: 2007.13916 [cs.CV].
- [38] B. Recht et al. “Do ImageNet Classifiers Generalize to ImageNet?” In: *Proceedings of the 36th International Conference on Machine Learning*. Ed. by K. Chaudhuri and R. Salakhutdinov. Vol. 97. Proceedings of Machine Learning Research. PMLR, June 2019, pp. 5389–5400.
- [39] P. H. Richemond et al. “BYOL works even without batch statistics”. In: *preprint arXiv:2010.10241* (2020).
- [40] S. Romberg et al. “Scalable logo recognition in real-world images”. In: Apr. 2011, p. 25. DOI: 10.1145/1991996.1992021.
- [41] O. Russakovsky et al. “Imagenet large scale visual recognition challenge”. In: *IJCV* (2015).
- [42] N. Saunshi et al. “Understanding contrastive learning requires incorporating inductive biases”. In: *International Conference on Machine Learning*. PMLR. 2022, pp. 19250–19286.
- [43] S. Shen et al. *FireRisk: A Remote Sensing Dataset for Fire Risk Assessment with Benchmarks Using Supervised and Self-supervised Learning*. 2023. arXiv: 2303.07035 [cs.CV].
- [44] N. Silberman et al. “Indoor segmentation and support inference from rgb-d images”. In: *ECCV*. 2012.
- [45] J. Stallkamp et al. “Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition”. In: *Neural Networks* 0 (2012), pp. -. ISSN: 0893-6080. DOI: 10.1016/j.neunet.2012.02.016.
- [46] Y. Tian et al. “What makes for good views for contrastive learning?” In: *Advances in neural information processing systems* 33 (2020), pp. 6827–6839.
- [47] H. Touvron et al. “LLaMA: Open and Efficient Foundation Language Models”. In: *arXiv preprint arXiv:2302.13971* (2023).
- [48] H. Touvron et al. *Training data-efficient image transformers & distillation through attention*. 2021. arXiv: 2012.12877 [cs.CV].
- [49] G. Van Horn et al. “The inaturalist species classification and detection dataset”. In: *CVPR*. 2018.
- [50] W. Wang et al. “Image as a Foreign Language: BEiT Pretraining for All Vision and Vision-Language Tasks”. In: *arXiv preprint arXiv:2208.10442* (2022).
- [51] X. Wang et al. “ChestX-Ray8: Hospital-Scale Chest X-Ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases”. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, July 2017. DOI: 10.1109/cvpr.2017.369.
- [52] T. Xiao et al. *What Should Not Be Contrastive in Contrastive Learning*. 2021. arXiv: 2008.05659 [cs.CV].

- [53] J. Zhang and K. Ma. “Rethinking the augmentation module in contrastive learning: Learning hierarchical augmentation invariance with expanded views”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2022, pp. 16650–16659.
- [54] S. Zhang et al. “Opt: Open pre-trained transformer language models”. In: *arXiv preprint arXiv:2205.01068* (2022).
- [55] B. Zhou et al. “Scene parsing through ade20k dataset”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017, pp. 633–641.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We provide an existence proof that it is possible to train a self-supervised joint-embedding architecture (based on DINOv2) that reaches state-of-the-art performance without relying on invariance to data augmentations, as was previously deemed necessary. Table 1 compares with alternative methods not using domain-specific data augmentations, while Figures 2 show that the performance almost matches the data-augmented baseline.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: A "Limitations" paragraph is present at the end of the conclusion section. The main limitation is in the generality of the claim regarding alternative methods. The existence proof still holds true though, as the example we provide is sufficient to disprove prior community belief that SSL joint-embedding architectures are based on learning invariance between different domain-specific (*hand-crafted*) data augmentations.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best

judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This is an experimental paper.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The hyperparameter sets and code are available publicly in the official DINOv2 repository. The ImageNet-1k and Imagenet-22k experiments are sufficient to support all the claims in the paper, and these datasets are available publicly as well as the linear-probing evaluation code and documentation to reproduce results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.

- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Code, hyperparameters, and data sufficient to support our claims are available for the main experimental results as mentioned above.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We employ the base settings provided in the DINOv2 repository, that we used to produce the results in this paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: The claims following our experimental results rely on the presence, or not, of large differences in performance between the different setups that we used. During the study, the authors did not observe variability in the evaluation results large enough to justify the cost of running multiple model trainings to provide such bars, and believe that this tradeoff is acceptable.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide a description of the setup (5 compute nodes of 8 A100-80GB GPUs each) as well as its speed (125k iterations in 27 hours) in the Experimental Setup section 3.2

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.

- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: The authors believe this work does not have societal impact, as it focuses on the understanding of SSL algorithms and builds on prior literature and code that is already publicly available.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We don't release such high risk information.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [No]

Justification: We did not include the licenses of the datasets. However, proper citations are referenced for all datasets and code; in particular the code is Apache-2.0 as mentioned in the corresponding footnote.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We don't release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We don't do such thing.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: NA.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.