ON THE ROBUSTNESS OF SELF-SUPERVISED REPRESENTATIONS FOR MULTI-VIEW OBJECT CLASSIFICATION

David Torpey
School of Computer Science and Applied Mathematics
University of the Witwatersrand
Johannesburg, South Africa
`674425@students.wits.ac.za`

Richard Klein
School of Computer Science and Applied Mathematics
University of the Witwatersrand
Johannesburg, South Africa
`richard.klein@wits.ac.za`

August 2, 2022

ABSTRACT

It is known that representations from self-supervised pre-training can perform on par, and often better, on various downstream tasks than representations from fully-supervised pre-training. This has been shown in a host of settings such as generic object classification and detection, semantic segmentation, and image retrieval. However, some issues have recently come to the fore that demonstrate some of the failure modes of self-supervised representations, such as performance on non-ImageNet-like data, or complex scenes. In this paper, we show that self-supervised representations based on the instance discrimination objective lead to better representations of objects that are more robust to changes in the viewpoint and perspective of the object. We perform experiments of modern self-supervised methods against multiple supervised baselines to demonstrate this, including approximating object viewpoint variation through homographies, and real-world tests based on several multi-view datasets. We find that self-supervised representations are more robust to object viewpoint and appear to encode more pertinent information about objects that facilitate the recognition of objects from novel views.

Keywords Deep learning · Self-supervised learning · Representation learning

In recent times, self-supervised learning (SSL) has come to the fore as a viable alternative to fully-supervised models as a way to perform large-scale pre-training that can transfer successfully to downstream tasks [Chen et al., 2020a,b, Caron et al., 2021, 2020]. It is known that self-supervised (SS) pre-training results in performance that is on par, and often superior, to supervised pre-training [Ericsson et al., 2021] (usually by pre-training on the ImageNet [Deng et al., 2009] dataset). However, some of the failure modes are 1) poor visual grounding resulting in decreased performance on scene images versus object-centric images [Selvaraju et al., 2020], and 2) less robustness than supervised models when transferring to data with a distribution shift from ImageNet [Ericsson et al., 2021], such as medical imaging or satellite imagery. However, very little is known about self-supervised representations, their pros and cons, what they do and do not encode, and potential ways to improve them. It is of crucial importance to understand these models and when/where they perform well so that we can lay a path forward to improve on them.

On a high level, we attempt to delve deeper into the efficacy of SS representations, and what they potentially encode differently from supervised representations - specifically in terms of robustness and invariance to viewpoint. Viewpoint invariance is a well-studied area of research, and there is biological evidence that suggests that it is a core trait of the human visual system [Newell and Findlay, 1992]. Firstly, we perform an analysis of SSL versus supervised learning in the context of controlled changes in object view by applying homographies of varying strengths to the images. We perform this analysis on datasets with different properties (e.g. fine-grained/generic, high/low cardinality label set). These homographies approximate possible object transformations in the real world and thus serve as a good proxy to test this hypothesis on large scale, diverse datasets that are not inherently multi-view.

We then perform a set of experiments on real-world, multi-view scenarios. These datasets cover a variety of object types, scenes, and scenarios that could be seen in the wild. The primary goal of this analysis is to empirically demonstrate that SS representations encode information that provides for a better ability to generalise to novel object views in realistic

settings (i.e. datasets that are inherently multi-view). The secondary goal of this analysis is to verify the finding of the previous experiment in which the different object views are synthetically generated through homographies. This analysis is achieved through nearest neighbour searches directly in the respective networks' feature spaces (i.e. there is no form of weight adaptation to the backbone network). We perform ablations to show that the results hold even with these adaptations. Our contributions can be summarised as follows.

(1) Performing an analysis comparing SS representations to supervised representations in terms of robustness to viewpoint changes using common image classification benchmark datasets from the literature. This is done by approximating changes to object view with homographies, and it enables the quantifying of performance in a *controlled environment*. (2) Performing a similar analysis to the above using multiple real-world, multi-view datasets to quantify performance in the wild. These are the more important experiments as they demonstrate realistic multi-view scenarios. (3) Demonstrate that SS representations are more robust to both A) approximate/synthetic viewpoint variations (homography) and B) real-world viewpoint variations through multiple multi-view datasets. (4) Demonstrate that SS representations require less context than multiple supervised baselines to identify an object using the embeddings directly from the backbone networks.

The rest of the paper is formatted as follows. Section 1 presents related work in the domain, including modern SS techniques and previous work analysing SS representations. Section 2 presents our research methodology. Section 3 presents an experimental analysis, including experiments on large-scale datasets using homographies to approximate viewpoint changes as well as tests on various real-world multi-view datasets. Finally, Section 4 concludes the research with some final remarks.

1   Related Work

A wealth of previous work exists in the realm of SSL [Caron et al., 2020, He et al., 2020, Chen et al., 2020a, Grill et al., Hsieh et al., 2015, Doersch et al., 2015]. We discuss the methods pertinent to modern SSL, and where our work fits into the domain. It should be noted that approaches to SSL can largely be cast as either generative (e.g. generative adversarial networks [Goodfellow et al., 2014] and variational autoencoders Kingma and Welling [2014]) or discriminative (pretext tasks and instance discrimination (ID)/contrastive learning (CL)). This review of prior work focuses on discriminative approaches, as these are the most popular techniques currently.

1.1   Pretext Tasks

SSL can be characterised as a way of obtaining a supervision signal from the input data itself. One way to achieve this is by defining a so-called *pretext task*, which is a proxy task for which labels are derived from the input images. A huge variety of such pretext tasks have been proposed including relative patch prediction [Doersch et al., 2015], predicting rotations [Feng et al., 2019, Gidaris et al., 2018], jigsaw puzzles [Noroozi and Favaro, 2016], colourisation [Larsson et al., 2016, 2017, Zhang et al., 2016], inpainting [Pathak et al., 2016], and others [Zhang et al., 2017, Noroozi et al., 2017].

For example, Doersch et al. [2015] define a proxy task of predicting the position on a neighbouring image patch relative to the central patch on a 3×3 grid of patches. Noroozi and Favaro [2016] train a network to solve a jigsaw puzzle by separating the image into numbered patches, shuffling these patches, and training the network to reassemble the patches into the correct order. One of the more successful pretext tasks is colourisation [Larsson et al., 2016, 2017], in which a network is trained to predict the colourised version of an input grayscale image.

Many early approaches to SSL were based on such hand-crafted pretext tasks. However, defining a pretext task that can learn generic, transferable image representations for downstream tasks is difficult. More recently, approaches based on *instance discrimination* have seen more use and gained more popularity from the research community, for various reasons including improved scalability, and better-learned representations that transfer better to various downstream tasks.

1.2   Instance Discrimination

The core idea of ID is to make each image its own class, thereby tasking the network with discriminating between individual images (i.e. instances), instead of the more common discrimination between classes. In its naive formulation, this would introduce a natural bottleneck when the dataset grows too large, as the number of classes would grow linearly with it. Thus, it would not be possible to use the typical parametric cross-entropy loss. To make this problem tractable, Wu et al. [2018] introduce a non-parametric softmax classifier that is trained to be instance-discriminative.

Noise-contrastive estimation (NCE) [Gutmann and Hyvärinen, 2010] is used to make computing the non-parametric softmax more efficient. Cao et al. [2020] manage to perform parametric ID for SSL.

van den Oord et al. [2018] propose a similar loss function based on NCE known as InfoNCE. This loss has been widely used in the literature [Chen et al., 2020a,b]. Their model is based on the idea of predictive coding [Elias, 1955], and predicts the future in latent space using an autoregressive model.

Many SSL methods fall into a class known as *contrastive learning* [Chen et al., 2020a,b, van den Oord et al., 2018, He et al., 2020, Chen et al., 2020c]. Chen et al. [2020a] propose a general framework for SSL based on CL. Two networks, in a Siamese-like fashion, operate on two distinct views of the same input image. These views are generated by random data augmentation, such as cropping, colour dropping, colour jitter, and blurring. This random view generation through augmentation is a common approach in the literature. The InfoNCE between the embeddings of the two views is then minimised. However, such an architecture requires a very large batch size (e.g. 4096+) to be effective, since sufficiently many negative samples need to be present for the CL to work.

Various architectures that overcome the need for such large batch sizes have been proposed [Grill et al., Zbontar et al., 2021, Caron et al., 2020]. Grill et al. propose a student-teacher paradigm in which the student network predicts the embedding of the teacher network. Only the weights of the student are updated via backpropagation, whereas teacher weights are updated using an exponential moving average.

Many other ID approaches have been proposed, including specialised architectures such as memory banks [He et al., 2020, Chen et al., 2020c], clustering-based approaches [Caron et al., 2020, 2018], mutual information maximisation [Hjelm et al., 2019], vision transformer-based models [Caron et al., 2021], and combinations of pretext tasks and CL [Misra and van der Maaten, 2020]. However, many of those have bottlenecks inhibiting their scalability and widespread use, including complicated architectures, training regimes, and very large compute requirements.

1.3 Issues with Self-Supervised Representations

It has been shown that for many downstream tasks (i.e. image classification, object detection, semantic segmentation, surface normal estimation, and image retrieval), SS representations consistently outperform supervised representations [Ericsson et al., 2021]. However, there are some issues with SSL and the associated learned representations. Firstly, it has been shown that their efficacy when the downstream task has a large distribution shift from ImageNet is much lower, and their comparative performance versus a similarly pre-trained supervised alternative is poor.

Secondly, it has been shown SS models have poor visual grounding [Selvaraju et al., 2020], resulting in worse performance on images that are not object-centric, such as scene images with many objects. A GradCAM-based [Selvaraju et al., 2017] loss is proposed to overcome this poor visual grounding.

Interestingly, Purushwalkam and Gupta [2020] find that SS models fail to capture viewpoint and category instance invariance. However, this does not affect our findings of SSL models being more robust to viewpoint, since there are a few key differences from our work. Firstly, we quantify performance in a different way (we measure based on the common linear evaluation paradigm, whereas they use an explicit measure of invariance based on the proposal in [Goodfellow et al., 2009]). Secondly, we test on more varied and realistic datasets specifically tailored to viewpoint. Lastly, we study a different class of models (i.e. SWaV [Caron et al., 2020] and DINO [Caron et al., 2021] instead of MoCo [He et al., 2020] and PIRL [Misra and van der Maaten, 2020]). These are completely different architectures trained in very different ways, resulting in inherently different properties. Furthermore, the findings for the slightly worse performance of SSL versus supervised in terms of viewpoint invariance can partially be attributed to the nature of the datasets used (e.g. ALOI [Geusebroek et al., 2005]), which contains images that are very dissimilar to ImageNet, which has previously been shown to result in poor performance in SSL [Ericsson et al., 2021]. In fact, we too observe similar findings for datasets with large distribution shift from ImageNet.

2 Methodology

2.1 Measuring Robustness to Viewpoint Variation

Consider functions $f : \mathcal{X} \to \mathbb{R}^n$ and $g : \mathcal{X} \to \mathbb{R}^n$, and a sample space of images $\mathcal{X}$. These are a supervised pre-trained model, and a self-supervised pre-trained model, respectively.

We aim to analyse the efficacy and representational power of embeddings $f(x), g(x) \in \mathbb{R}^n$ in terms of robustness to viewpoint variation on a host of different datasets. We assume that both f and g have been estimated by pre-training on ImageNet (as this is the common paradigm in the literature). We then employ a transfer learning approach to analyse the pre-trained models and resultant features.

Figure 1: Homographies of varying strengths (from left: $H_{0.0}$, $H_{0.2}$, $H_{0.4}$, $H_{0.6}$, and $H_{0.8}$). The red boxes depict our *bounded homography*.

Essentially, we aim to demonstrate empirically that the SS representations produced by g are more robust to viewpoint changes than those from f. Mathematically, this can be formalised as follows. Consider a function $V : \mathcal{X} \to \mathcal{X}$ that is tasked with changing an object's viewpoint. This could be a homography, affine transformation, or in general any transformation that can perturb an object along one or more of its axes of variation (such as those that would occur in the wild). Then, a function g is more robust to V than a function f, if:

$$\mathbb{E}[L(f(x), f(V(x)))] \geq \mathbb{E}[L(g(x), g(V(x)))] \tag{1}$$

for some loss function L, and for all $x \in \mathcal{X}$. *This is the criterion we use to measure and compare models' multi-view recognition performance, and implicitly, their viewpoint invariance.* Ideally, we would want full invariance to V, but this is not realistic in practice. However, there are scenarios where full invariance is not ideal for other classes of perturbations besides viewpoint. For example, invariance to colour will make a downstream task of colour classification difficult.

2.2 Synthetic Viewpoint Variation Analysis

To analyse the robustness of these supervised and SS representations to viewpoint variation, we opt for a two-stage approach. We first investigate the viewpoint invariance in a controlled environment, by synthetically varying the view of an object by applying a random homography *to the testing images*. We apply two kinds of homographies: a regular homography where any new background created is left in the image (the default for many libraries), and a *bounded homography*, whereby we crop the maximum-area inscribed axis-aligned rectangle from the resulting polygon. This bounded homography enables the ability to test whether the black background in a default homography affects performance and if the models are biased toward this black background (see Fig. 1 for a comparison).

We represent a homography as $H_\alpha : \mathcal{X} \to \mathcal{X}$, where $\alpha \in [0, 1]$ is a factor controlling the strength of the homography. The goal of this set of experiments is to have full control of the amount of viewpoint variation through the strength of the homography, which allows for a systematic analysis of the two modelling paradigms.

For the linear evaluation experiments, we train a multinomial logistic regression model on the representations $f(x)$ and $g(x)$ (using the L-BFGS optimiser) on a training dataset $X_{\text{tr}} \subset \mathcal{X}$ with labels Y. We then compute accuracy on the test set with a homography applied to each image: $\{H_\alpha(x) | x \in X_{\text{te}}\}$. The application of this homography serves as a mechanism to obtain novel object views, as homographies can be used to approximate object transformations seen in the real world.

2.3 Real-World Viewpoint Variation Analysis

Motivated by the results on synthetic data, we also perform experiments on real-world, multi-view datasets using the representations directly from f and g (i.e. without training a classifier on top of f or g). Importantly, these experiments serve as the main contribution of the paper, as they measure the efficacy of the 2 modelling paradigms in a realistic scenario through the use of real-world datasets. The viewpoint invariance is measured in the form of a k-nearest neighbour search, with $k = 1$. This paradigm allows for better empirical evidence as to whether the bare representations (without a classifier learning from them) are encoding viewpoint in the way we hypothesise. For these real-world experiments, we use a suite of datasets that were explicitly curated to test multi-view performance, and that cover a wide range of object and task types.

3 Experiments

3.1 Experimental Setup

3.1.1 Datasets

We use the following datasets for the deformation experiments in which the homography is applied to the testing images: Aircraft [Maji et al., 2013], Birdsnap [Berg et al., 2014], Color [Hsieh et al., 2015], Caltech101 [Fei-Fei et al., 2004], CIFAR10 [Krizhevsky et al., a], CIFAR100 [Krizhevsky et al., b], Food [Bossard et al., 2014]. These datasets are used to be consistent with previous work on benchmarking SS representations for downstream image classification.

For the real-world, multi-view experiments, we use multiple realistic datasets, including the multi-view car (MVC) dataset [Ozuysal et al., 2009], a multi-view stereo dataset of generic objects (Recon3D) [Kolev et al., 2010], a multi-view, multi-class dataset [Roig et al., 2011] (MVMC), the Amsterdam library of object images [Geusebroek et al., 2005] (ALOI), and a multi-view stereo face dataset [Fransens et al., 2005] (Stereo Face). The latter three of these datasets are from the EPFL CVLAB data repository[1]. These datasets present a diverse range of multi-view scenarios, object and background types, and difficulty for the models under assessment.

3.1.2 Models

In order to effectively evaluate the efficacy of SS and supervised representations for invariance to viewpoint, we test a suite of a total of 15 different models. We separate these 15 models into 3 groups: 1) SSL (ID) - ID-based SSL, 2) SSL (PT) - pretext task-based SSL, and 3) supervised. The models within each group are: **SSL (ID)**: SWaV (RN50) and SWaV (RN50w2) [Caron et al., 2020], SimCLR (RN50) and SimCLR (RN50w2) [Chen et al., 2020a,b], DINO (ViT) and DINO (RN50) [Caron et al., 2021], MoCo v2 (RN50) [Chen et al., 2020c], Barlow Twins (RN50) [Zbontar et al., 2021]; **SSL (PT)**: RotNet (RN50) [Gidaris et al., 2018], Jigsaw (RN50) [Noroozi and Favaro, 2016], Colorization (RN50) [Larsson et al., 2016]; **Supervised**: RN50 [He et al., 2016], ViT [Kolesnikov et al., 2021], EfficientNet [Tan and Le, 2019], RegNet [Radosavovic et al., 2020]. This suite of models provides a wide variety of SSL methods, backbone architectures, and SSL paradigms (ID, CL, pre-text tasks). All models are pre-trained on ImageNet, and available in PyTorch Hub[2] and the VISSL Model Zoo[3].

3.2 Robustness to Deformation

We first conduct experiments on common large-scale downstream benchmark datasets to quantify the performance of the models in the multi-view case by approximating different object views and perspectives by applying a homography to each testing image. We use the linear evaluation paradigm [Chen et al., 2020a, Grill et al., He et al., 2020, Chen et al., 2020b] from the literature to compute our metrics for this set of experiments.

Table 1: Linear evaluation performance on common image classification benchmarks for increasing homography strengths (top 3 best-performing models). Results are averaged over 10 random trials.

	$H_{0.0}$			$H_{0.2}$			$H_{0.4}$			$H_{0.6}$			$H_{0.8}$		
CIFAR10	Supervised (ViT)	Supervised	92.29	Supervised (ViT)	Supervised	91.68	Supervised (ViT)	Supervised	90.94	Supervised (ViT)	Supervised	88.7	Supervised (ViT)	Supervised	79.94
	SWaV (RN50w2)	SSL (ID)	91.52	DINO (ViT)	SSL (ID)	89.04	DINO (ViT)	SSL (ID)	83.67	Supervised EffNet	Supervised	74.87	Supervised EffNet	Supervised	67.83
	DINO (ViT)	SSL (ID)	91.12	Barlow Twins (RN50)	SSL (ID)	84.01	Barlow Twins (RN50)	SSL (ID)	79.92	DINO (ViT)	SSL (ID)	74.15	DINO (ViT)	SSL (ID)	64.65
CIFAR100	Supervised (ViT)	Supervised	75.43	Supervised (ViT)	Supervised	73.99	Supervised (ViT)	Supervised	72.67	Supervised (ViT)	Supervised	68.71	Supervised (ViT)	Supervised	59.73
	SWaV (RN50w2)	SSL (ID)	74.65	DINO (ViT)	SSL (ID)	68.66	DINO (ViT)	SSL (ID)	62.82	DINO (ViT)	SSL (ID)	55.26	Supervised EffNet	Supervised	48.22
	DINO (ViT)	SSL (ID)	74.43	Barlow Twins (RN50)	SSL (ID)	63.03	Supervised EffNet	Supervised	58.25	Supervised EffNet	Supervised	54.36	DINO (ViT)	SSL (ID)	48.2
Color	SWaV (RN50w2)	SSL (ID)	94.04	DINO (ViT)	SSL (ID)	91.94	Supervised (ViT)	Supervised	90.75	Supervised (ViT)	Supervised	88.64	Supervised (ViT)	Supervised	81.9
	DINO (ViT)	SSL (ID)	93.78	SWaV (RN50w2)	SSL (ID)	91.66	DINO (ViT)	SSL (ID)	89.79	DINO (ViT)	SSL (ID)	86.03	DINO (ViT)	SSL (ID)	80.46
	DINO (RN50)	SSL (ID)	93.55	Supervised (ViT)	Supervised	91.66	SWaV (RN50w2)	SSL (ID)	88.47	SWaV (RN50w2)	SSL (ID)	84.09	SWaV (RN50w2)	SSL (ID)	77.28
Caltech101	DINO (ViT)	SSL (ID)	96.45	Supervised (ViT)	Supervised	94.05	DINO (ViT)	SSL (ID)	93.31	DINO (ViT)	SSL (ID)	91.01	Supervised (ViT)	Supervised	83.21
	SWaV (RN50w2)	SSL (ID)	95.34	DINO (ViT)	SSL (ID)	93.87	Supervised (ViT)	Supervised	92.22	Supervised (ViT)	Supervised	90.32	DINO (ViT)	SSL (ID)	81.66
	DINO (RN50)	SSL (ID)	95.02	Supervised EffNet	Supervised	93.4	Supervised EffNet	Supervised	91.86	Supervised EffNet	Supervised	88.99	Supervised EffNet	Supervised	81.52
Aircraft	DINO (ViT)	SSL (ID)	59.65	DINO (ViT)	SSL (ID)	53.95	DINO (ViT)	SSL (ID)	48.06	DINO (ViT)	SSL (ID)	41.95	DINO (ViT)	SSL (ID)	36.12
	DINO (RN50)	SSL (ID)	59.29	DINO (RN50)	SSL (ID)	47.72	Barlow Twins (RN50)	SSL (ID)	41.62	Barlow Twins (RN50)	SSL (ID)	35.8	DINO (RN50)	SSL (ID)	32.33
	SWaV (RN50w2)	SSL (ID)	56.14	Barlow Twins (RN50)	SSL (ID)	47.58	DINO (RN50)	SSL (ID)	40.94	DINO (RN50)	SSL (ID)	35.57	Barlow Twins (RN50)	SSL (ID)	31.76
Birdsnap	DINO (ViT)	SSL (ID)	59.55	DINO (ViT)	SSL (ID)	56.37	DINO (ViT)	SSL (ID)	53.54	DINO (ViT)	SSL (ID)	48.76	DINO (ViT)	SSL (ID)	41.18
	Supervised (ViT)	Supervised	49.39	Supervised (ViT)	Supervised	47.83	Supervised (ViT)	Supervised	45.53	Supervised (ViT)	Supervised	42.56	Supervised (ViT)	Supervised	36.1
	Supervised RegNet	Supervised	43.74	Supervised RegNet	Supervised	39.97	Supervised RegNet	Supervised	38.12	Supervised RegNet	Supervised	34.34	Supervised RegNet	Supervised	29.28
Food	SWaV (RN50w2)	SSL (ID)	75.86	SWaV (RN50w2)	SSL (ID)	71.75	DINO (ViT)	SSL (ID)	68.06	DINO (ViT)	SSL (ID)	62.65	DINO (ViT)	SSL (ID)	54.52
	DINO (ViT)	SSL (ID)	74.47	DINO (ViT)	SSL (ID)	71.33	SWaV (RN50w2)	SSL (ID)	67.01	SWaV (RN50w2)	SSL (ID)	60.2	SWaV (RN50w2)	SSL (ID)	51.51
	SWaV (RN50)	SSL (ID)	72.27	DINO (RN50)	SSL (ID)	69.24	DINO (RN50)	SSL (ID)	65.59	DINO (RN50)	SSL (ID)	59.53	DINO (RN50)	SSL (ID)	50.93

Table 1 shows the linear evaluation performance for varying homography strengths for the benchmark datasets (see Appendix A.2.1 for more results). The supervised baselines perform best on CIFAR10 and CIFAR100 - two of the most common benchmarks in computer vision. However, the SSL ID models dominate the remaining 5 benchmark datasets. These 5 benchmarks contain more variety, and arguably serve as a more comprehensive evaluation than the

---

[1] https://www.epfl.ch/labs/cvlab/data/

[2] https://pytorch.org/hub/

[3] https://github.com/facebookresearch/vissl/blob/main/MODEL_ZOO.md

CIFAR datasets, due to the much higher resolution inputs, higher cardinality label sets, and higher frequency differences between classes. For example, in the most challenging benchmark, Aircraft, there is not a single supervised model in the top 3 for any of the homography strengths.

We note that backbones based on the ViT [Kolesnikov et al., 2021] architecture (both SS and supervised) perform well, and are the best performing models for the majority of the datasets. This motivates further research into ViT-based SS models. Further, we note that wider ResNets (ResNet50w2) tend to perform better than narrower ones (ResNet50), as evidenced in the Color, Caltech101, and Food datasets. This is a known result in SSL for the usual linear evaluation strategy for these benchmarks, and our findings here suggest that this result generalises to the synthetic multi-view setting. It is also interesting to note that no pretext-task-based SSL model (SSL PT) appears in the top 3 for any of the datasets. This further motivates the recent focus of the research community on SSL ID methods, as they typically outperform SSL PT methods on all evaluated tasks. Interestingly, this general trend of SSL (ID) methods performing best, followed by Supervised (with SSL (PT) performing worst) holds for the *bounded* homography as well (refer to Appendix A.2.2 for the full table of results).

Table 2: Relative decrease in performance from $H_{0.0}$ to $H_{0.2,0.4,0.6,0.8}$ for both the default and bounded homographies by model type. Results are averaged over the 7 benchmark datasets.

Decrease Strength	Type	Default Homography	Bounded Homography
$H_{0.0} \rightarrow H_{0.2}$	SSL (ID)	-0.105	**-0.0337**
	SSL (PT)	-0.2547	-0.1483
	Supervised	**-0.0725**	-0.0704
$H_{0.0} \rightarrow H_{0.4}$	SSL (ID)	-0.1797	**-0.0743**
	SSL (PT)	-0.3363	-0.2646
	Supervised	**-0.1226**	-0.1168
$H_{0.0} \rightarrow H_{0.6}$	SSL (ID)	-0.2644	**-0.1974**
	SSL (PT)	-0.3806	-0.4109
	Supervised	**-0.1932**	-0.2252
$H_{0.0} \rightarrow H_{0.8}$	SSL (ID)	-0.3485	**-0.3933**
	SSL (PT)	-0.4134	-0.5642
	Supervised	**-0.2829**	-0.4118

Table 2 shows the **relative decrease** in performance by model type for both the default and bounded homographies. Interestingly, it seems the black background affects the results significantly because Supervised models perform best on all decrease strengths with the default homography (i.e. with the black background), but SSL (ID) models consistently perform best for all decrease strengths with the bounded homography (where no black background is present to bias results). This suggests that the supervised models are somewhat reliant on the black background of the homography to retain the performance as the homography gets stronger. However, without the black background, the supervised models struggle to retain their performance with stronger homographies. In contrast, the SSL (ID) models retain accuracy better when no black background is present to potentially bias the model. Unsurprisingly, SSL (PT) perform worst in both cases.

These results suggest that the SS techniques are more robust, in general, to changes in the perspective/viewpoint of the object *as approximated by a homography*. This motivates further analysis of these methods on real-world, multi-view datasets to test in a more realistic scenario where viewpoint variation of the objects is more natural than a homography. We conduct these experiments and present the results in the next section.

3.3 Multi-View Performance in-the-Wild

Our motivation for this section is to benchmark the performance of SS representations against supervised alternatives for multi-view performance *in-the-wild*. Previous work has not focused on the multi-view efficacy of SSL on such varied and real-world datasets, nor in this detail. We do this by running nearest neighbour searches (NNS) directly in the respective network's embedding spaces. We posit this is a very effective method of gauging the network embedding's multi-view invariance and robustness to viewpoint variation in the wild (i.e without adapting the weights or fine-tuning to a particular dataset). This set of experiments essentially serves as a realistic setting for analysing and measuring real-world viewpoint invariance.

Table 3 shows the top 3 models for each of the real-world datasets (see Appendix B.1 for more results). Clearly, SS representations dominate in terms of overall performance on these datasets. Interestingly, the ViT models are not as effective in this setting as compared to the synthetic setting with the homographies. Further, the pretext-task-based

Table 3: Top 3 performing models for each real-world, multi-view dataset. Results are ranked.

Dataset	Model	Type	Accuracy
ALOI	SWaV (RN50w2)	SSL (ID)	96.8
	SWaV (RN50)	SSL (ID)	96.22
	DINO (RN50)	SSL (ID)	96.16
MVMC	RotNet (RN50)	SSL (PT)	93.14
	MoCov2 (RN50)	SSL (ID)	92.3
	Supervised (ViT)	Supervised	90.89
MVC	SWaV (RN50w2)	SSL (ID)	97.2
	Jigsaw (RN50)	SSL (PT)	95.11
	Supervised (ViT)	Supervised	94.56
Recon3D	Supervised (RN50)	Supervised	94.94
	Supervised RegNet	Supervised	93.2
	DINO (ViT)	SSL (ID)	93.2
Stereo Face	SWaV (RN50w2)	SSL (ID)	97.72
	DINO (ViT)	SSL (ID)	97.71
	DINO (RN50)	SSL (ID)	97.2

SSL methods appear more frequently in this setting, which suggests that the representations learned with this paradigm contain information useful for encoding object viewpoint (more so than can be leveraged in the synthetic setup in the previous section).

It should be noted, however, that the supervised models perform best on the Recon3D dataset. This is somewhat expected since this dataset contains images with a significant distribution shift from ImageNet. It has been shown in previous work that SSL models are less robust than supervised models (pre-trained on ImageNet) when evaluated on datasets with a large distribution shift from ImageNet [Ericsson et al., 2021] (e.g. non-natural images), and this experiment replicates these findings in the multi-view setting.

Table 4: Real-world multi-view dataset performance by model type.

	ALOI	MVC	MVMC	Recon3D	Stereo Face
SSL (ID)	**95.66**	**92.73**	**90.41**	88.34	**94.36**
SSL (PT)	84.04	88.47	85.21	63.26	70.83
Supervised	93.58	89.32	88.34	**90.72**	91.56

Table 4 provides a summary of performance on these datasets by model type. These results are computed by averaging over all trials for each dataset. SSL (ID) models perform best overall except on Recon3D. Supervised performs 2nd best on average for all datasets, and SSL (PT) models are 3rd by some distance on average. We delve further into these results below, going into more granular detail for each dataset. Please note that for all plots below, SSL (ID), SSL (PT), and Supervised models are visualised by red, green, and grey tones, respectively.

Figure 2 shows the performance of such an NNS on a multi-view car dataset for the various models at different azimuthal angle differences. To obtain these metrics, we average the testing accuracy over all possible pairs of images that are $30n$ degrees apart ($n \in \{1, 2, \ldots, 11\}$). In this case, for each trial, the training and testing dataset both have a cardinality equal to the number of classes (i.e. in this case, the number of unique cars in the dataset). The reported results are an average of 30 randomised trials of different images at these angle differences.

On average, SSL (ID) models perform best, with the SWaV models performing particularly well. The Jigsaw SSL (PT) model performs surprisingly well on this task. It is clear that as the azimuthal angle increases, performance decreases for all models. Moreover, at the large azimuthal angle differences (i.e. > 150, when the task is more difficult), the performance improvement of SSL models vs the supervised baselines is notably larger. The Supervised (ViT) model consistently performs best out of all the supervised baselines by a large margin, and indeed performs similarly to some of the better SSL models.

Figure 3 shows the results of an NNS on a real-world, multi-view, multi-class dataset [Roig et al., 2011]. We vary the number of support samples - that is, the number of images of each class in the training set. Each point represents a mean of 100 random trials. We see that with as few as *one* image per class, the SS techniques outperform the supervised baselines significantly. We note that with a lower number of support samples (e.g. 1-2 samples), the variance of the performance metric is fairly high, which motivates the large number of trials. This is due to the fact that the performance depends on whether that single support sample per class is a canonical example of that class, or a poor representative sample that makes classification difficult. Evaluating with a small number of support samples enables us to truly test the efficacy of the embeddings for multi-view object recognition.
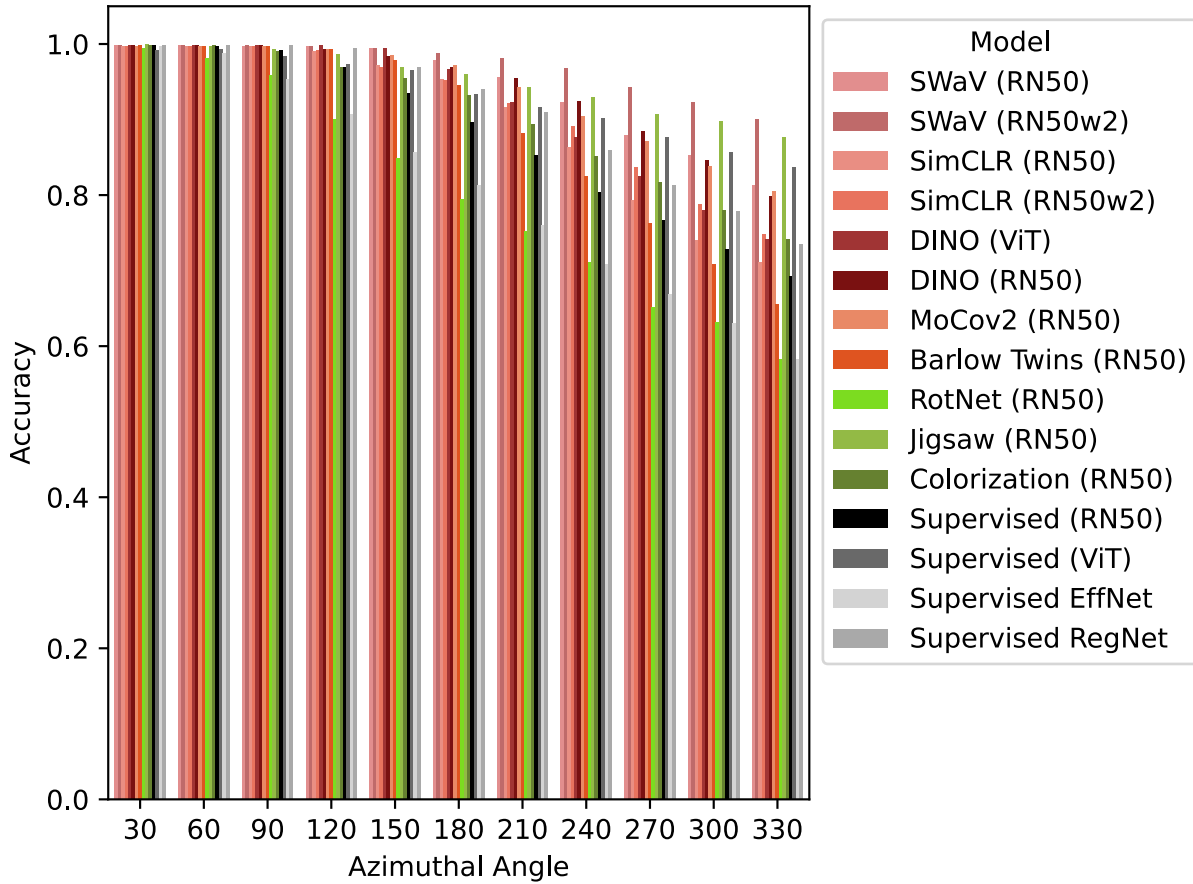
Figure 2: NNS results on the MVC dataset. The x-axis represents the difference in azimuthal angle between the training view and testing view.

As the number of support samples increases, the performance gap between SSL and supervised naturally decreases. This shows that when enough context is available, the supervised baseline improves - most likely due to the support samples illustrating a large enough diversity of different views, and at that stage the viewpoint robustness becomes unnecessary. However, the superior performance of SSL when the amount of context is very low suggests that SSL embeddings encode more information relevant to multi-view tasks. We posit that this is due to both the better occlusion invariance of SSL techniques and the nature of the ID objective in general.

Finally, we report results on the ALOI dataset (Fig. 4). We evaluate by increasing the test split percentage (up to 95% of the data for testing), essentially giving the models less and less context when performing the nearest neighbour search to test how effective the representations are at detecting the objects at different viewpoints. Similarly to the MVMC dataset above, when the amount of context provided to the model is reduced, the performance difference between SS and supervised models increases on average. Furthermore, SSL (PT) methods are consistently outperformed by the two other model types. This, as well as the results of SSL (PT) on the other real-world, multi-view dataset, suggests that the pretext tasks in the literature are not conducive to learning representations that are useful for reliable multi-view recognition, and do not encode viewpoint as reliable as the SSL ID paradigm.

3.4   Self-Supervised Representations vs. Supervised Representations

It is clear from Sections 3.2 and 3.3 that representations learned through self-supervision perform better, on average, than representations learned with supervision for the task of multi-view detection (when the models are pre-trained on ImageNet). The results suggests that this claim holds for both controlled distortion to object viewpoint through homographies (Section 3.2), as well as in real-world experiments on datasets that are inherently multi-view (Section
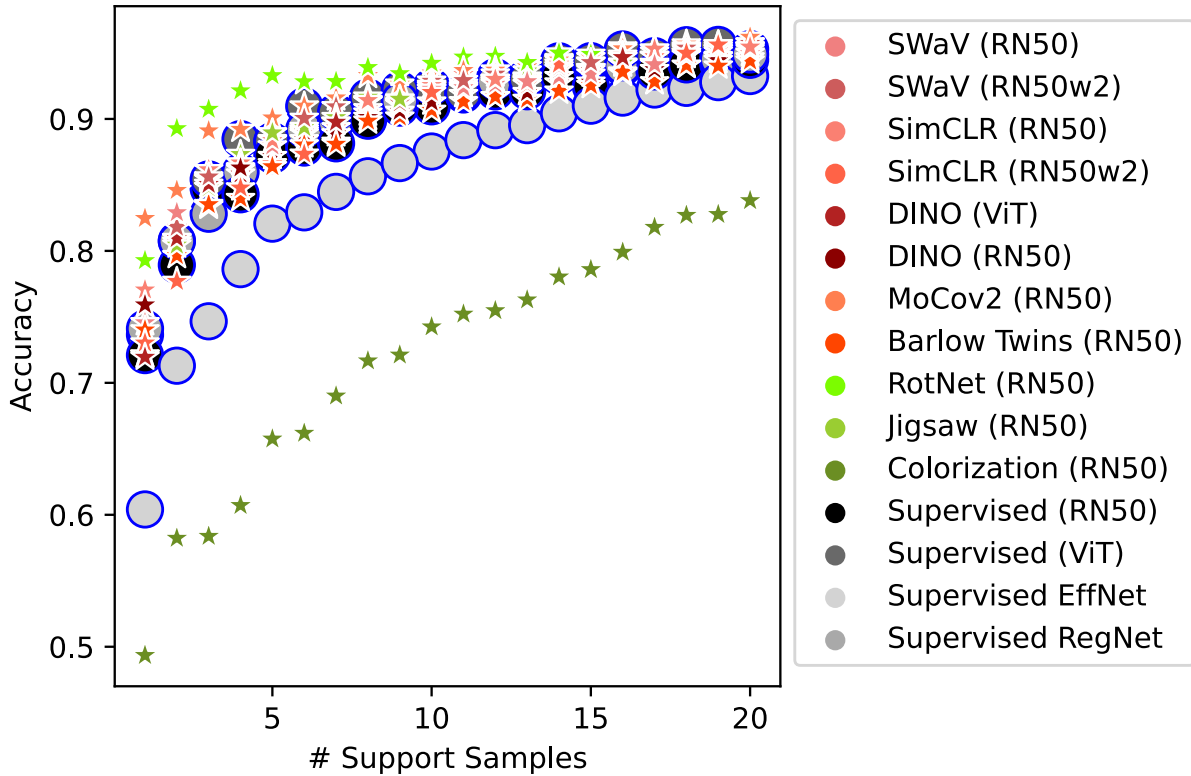
Figure 3: NNS results MVMC dataset. The stars denote SSL models (both ID and PT), and circles with blue borders denote supervised models.

3.3). In particular, SSL (ID) models are particularly effective, which mirrors similar findings outside the multi-view domain in the SSL literature [Ericsson et al., 2021].

We posit that ID models are so effective at multi-view object recognition for one main reason. A typical ID loss is of the following form: $L(\mathbf{x}, \mathbf{y})$, where $\mathbf{x}$ and $\mathbf{y}$ are the embeddings from the network for two distinct *views* of the **same** input image. These views are typically generated through a series of data augmentations, such as random cropping, colour distortion, and random flipping. As the network is trained, these two vectors are learned to be represented by the same vector. Since the two views represent different parts (i.e. crops) of the image, these may be different parts of the same object (particularly in object-centric datasets that are the norm in the SSL literature). This essentially means that the network is explicitly trained to be able to recognise an object from potentially very little context such as a random crop. This may explain the superior performance of SSL in general, particularly for datasets such as MVC and MVMC, as well as why these SSL (ID) models work especially well with very little context when compared to the various supervised baselines. Similar intuitions have been found in previous work in measuring invariances of SSL models, where they were shown to be more occlusion invariant than supervised alternatives [Purushwalkam and Gupta, 2020]. It should be noted, that the performance benefits of SS learning decreases as the amount of context provided to the networks increases (see Figures 4, 2, and 3). This is expected, as the networks have sufficient information about the objects when a large number of embeddings are available to perform the NNS or classification. It is, however, particularly useful to have good performance in a low-data regime, when little context is available. This is the main advantage of SS representations for multi-view object recognition.

Interestingly, some known results from outside the multi-view domain still hold. Performance of SSL models on data that has a large distribution shift from ImageNet has been shown to be lower than supervised learning [Ericsson et al., 2021], and those findings generalise to our multi-view setting, as evidenced by our results on the Recon3D dataset. This is likely due to the benefits of the ID objective discussed above not being able to generalise beyond ImageNet-like images, and not being able outweigh the benefits of a strong supervision signal from human-generated labels in this setting.

9
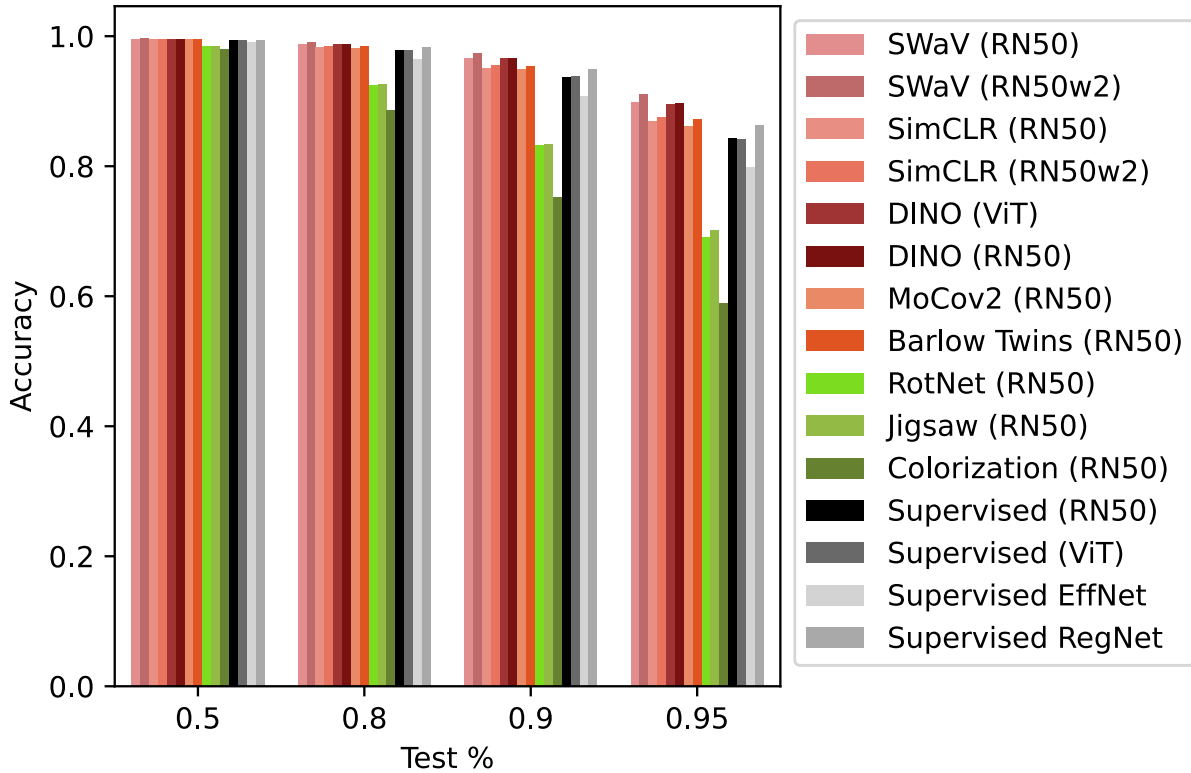
Figure 4: NNS results on the ALOI dataset.

Lastly, ViT based models have shown promising results in SSL literature [Caron et al., 2021], and it is clear from our results that ViTs perform particularly well at multi-view recognition (both SS and supervised versions). We suggest that this is due to the fact (in the SS case) that the SS transformers automatically discover objects (as shown by Caron et al. [2021] by visualising the attention maps of the ViT-based DINO, where the model learns an approximate semantic segmentation of the object). We posit that this property of ViTs assists in multi-view object recognition, and encourage the representations learned by these models to be more viewpoint invariant.

4   Conclusion

We observe that in multiple different scenarios, representations learned through modern self-supervision are more robust than supervised learning with respect to viewpoint changes. We find this holds both for approximate viewpoint variation with homographies, as well as real-world viewpoint variation on the numerous inherently multi-view datasets. Interestingly, even with very little context (see Figures 2, 3, and 4), or in some cases with large context (see Table 1), SSL consistently outperforms supervised learning in the multi-view setting.

We posit that these findings suggest that SS representations encode information more pertinent to object parts, which enables improved robustness to viewpoint. This is likely a byproduct of the fact that most modern SSL models (such as SWaV, DINO, and SimCLR) are trained with strong data augmentation (e.g. random crops) and ID, which encourages this property of being able to recognise objects from a comparatively smaller context such as object parts.

The use cases for these models and associated results are diverse. Models based on SSL should be preferred to supervised alternatives when annotation is a bottleneck/prohibitive. Further, SSL models are preferable when the application has a need to be able to viewpoint invariant, or to recognise a set of distinct objects at different viewpoints - *particularly in low-data regimes*. Many such applications exist, since viewpoint invariance is a core component of effective real-world computer vision [Purushwalkam and Gupta, 2020]. The supervised models should be preferred when there is no guarantee that the downstream application will consist of images similar to that of the pretraining data.

Both previous work, and this work, have shown that supervised models outperform SS models in different contexts with these sorts of data distribution shifts.

References

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. *ArXiv*, abs/2002.05709, 2020a.

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

Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021.

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

Linus Ericsson, Henry Gouk, and Timothy M. Hospedales. How Well Do Self-Supervised Models Transfer? In *CVPR*, 2021.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.

Rampraasath R. Selvaraju, Karan Desai, Justin Johnson, and Nikhil Naik. CASTing your model: Learning to localize improves self-supervised representations. *arXiv preprint arXiv:2012.04630*, 2020.

Fiona Newell and John M. Findlay. Viewpoint invariance in object recognition. *The Irish Journal of Psychology*, 13(4): 494–507, 1992.

Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap Your Own Latent: A new approach to self-supervised learning. In *Neural Information Processing Systems*.

Jun-Wei Hsieh, Li-Chih Chen, Sin-Yu Chen, Duan-Yu Chen, Salah Alghyaline, and Hui-Fen Chiang. Vehicle color classification under different lighting conditions through color correction. *IEEE Sensors Journal*, 15(2):971–983, 2015.

Carl Doersch, Abhinav Gupta, and Alexei A. Efros. Unsupervised visual representation learning by context prediction. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1422–1430, 2015.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27, 2014.

Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.

Zeyu Feng, Chang Xu, and Dacheng Tao. Self-supervised representation learning by rotation feature decoupling. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10356–10366, 2019.

Spyros Gidaris, Praveer Singh, and N. Komodakis. Unsupervised representation learning by predicting image rotations. *ArXiv*, abs/1803.07728, 2018.

M. Noroozi and P. Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016.

Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Learning representations for automatic colorization. In *European Conference on Computer Vision (ECCV)*, 2016.

Gustav Larsson, M. Maire, and Gregory Shakhnarovich. Colorization as a proxy task for visual understanding. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 840–849, 2017.

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

Deepak Pathak, Philipp Krähenbühl, Jeff Donahue, Trevor Darrell, and Alexei Efros. Context encoders: Feature learning by inpainting. 2016.

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

M. Noroozi, H. Pirsiavash, and P. Favaro. Representation learning by learning to count. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 5899–5907, 2017.

Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 3733–3742. Computer Vision Foundation / IEEE Computer Society, 2018.

Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 297–304. PMLR, 13–15 May 2010.

Yue Cao, Zhenda Xie, Bin Liu, Yutong Lin, Zheng Zhang, and Han Hu. Parametric instance classification for unsupervised visual feature learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 15614–15624. Curran Associates, Inc., 2020.

Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *ArXiv*, abs/1807.03748, 2018.

P. Elias. Predictive coding–i. *IRE Transactions on Information Theory*, 1(1):16–24, 1955.

Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020c.

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

Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.

Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Philip Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *ICLR 2019*. ICLR, April 2019.

Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *CVPR*, pages 6706–6716, 06 2020.

Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.

Senthil Purushwalkam and Abhinav Gupta. Demystifying contrastive self-supervised learning: Invariances, augmentations and dataset biases. In *Advances in Neural Information Processing Systems*, volume 33, pages 3407–3418. Curran Associates, Inc., 2020.

Ian Goodfellow, Honglak Lee, Quoc Le, Andrew Saxe, and Andrew Ng. Measuring invariances in deep networks. In *Advances in Neural Information Processing Systems*, volume 22, 2009.

J. M. Geusebroek, G. J. Burghouts, and A. W. M. Smeulders. The amsterdam library of object images. *International Journal of Computer Vision*, 61(1):103–112, 2005.

S. Maji, J. Kannala, E. Rahtu, M. Blaschko, and A. Vedaldi. Fine-grained visual classification of aircraft. Technical report, 2013.

Thomas Berg, Jiongxin Liu, Seung Woo Lee, Michelle L. Alexander, David W. Jacobs, and Peter N. Belhumeur. Birdsnap: Large-scale fine-grained visual categorization of birds. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2019–2026, 2014.

Li Fei-Fei, R. Fergus, and P. Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *2004 Conference on Computer Vision and Pattern Recognition Workshop*, pages 178–178, 2004.

Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research). a.

Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-100 (canadian institute for advanced research). b.

Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – mining discriminative components with random forests. In *European Conference on Computer Vision*, 2014.

Mustafa Ozuysal, Vincent Lepetit, and Pascal Fua. Pose estimation for category specific multiview object localization. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 778–785, 2009. doi:10.1109/CVPR.2009.5206633.

Kalin Kolev, Thomas Pock, and Daniel Cremers. Anisotropic minimal surfaces integrating photoconsistency and normal information for multiview stereo. In *Computer Vision – ECCV 2010*, pages 538–551, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

Gemma Roig, Xavier Boix, Horesh Ben Shitrit, and Pascal Fua. Conditional random fields for multi-camera object detection. In *2011 International Conference on Computer Vision*, pages 563–570, 2011.

R. Fransens, C. Strecha, and L. Van Gool. Parametric stereo for multi-pose face recognition and 3d modelling. In *ICCV*, 2005.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.

Alexander Kolesnikov, Alexey Dosovitskiy, Dirk Weissenborn, Georg Heigold, Jakob Uszkoreit, Lucas Beyer, Matthias Minderer, Mostafa Dehghani, Neil Houlsby, Sylvain Gelly, Thomas Unterthiner, and Xiaohua Zhai. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.

Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.

Ilija Radosavovic, Raj Prateek Kosaraju, Ross B. Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.

A Robustness Deformation

A.1 Data Augmentation Details

For NNS and linear evaluation experiments, the training transformation involves resizing the image such that the smaller side has length 224, and then performing a 224×224 centre crop (hereafter referred to as the RCC). The testing transform for these experiments is the RCC transformation, followed by an optional homography H_α. For the fine-tuning experiments, the training transform is a random-resized crop, and the testing transform is the same as the NNS and linear evaluation paradigms.

A.2   Linear Evaluation

A.2.1   Default Homography

Table 5: Linear evaluation performance on common image classification benchmarks for increasing homography strengths. Results are averaged over 10 random trials, and models are ranked.

Dataset	$H_{0.0}$ Model	Type	Score	$H_{0.2}$ Model	Type	Score	$H_{0.4}$ Model	Type	Score	$H_{0.6}$ Model	Type	Score	$H_{0.8}$ Model	Type	Score
CIFAR10	Supervised (ViT)	Supervised	92.29	Supervised (ViT)	Supervised	91.68	Supervised (ViT)	Supervised	90.94	Supervised (ViT)	Supervised	88.7	Supervised EffNet	Supervised	79.94
	SWaV (RN50w2)	SSL (ID)	91.52	DINO (ViT)	SSL (ID)	89.04	DINO (ViT)	SSL (ID)	83.67	Supervised EffNet	Supervised	74.87	Supervised (ViT)	Supervised	67.83
	DINO (ViT)	SSL (ID)	91.12	Barlow Twins (RN50)	SSL (ID)	84.01	Barlow Twins (RN50)	SSL (ID)	79.92	DINO (ViT)	SSL (ID)	74.15	DINO (ViT)	SSL (ID)	64.65
	SWaV (RN50)	SSL (ID)	88.47	Supervised EffNet	Supervised	80.81	Supervised EffNet	Supervised	78.98	Barlow Twins (RN50)	SSL (ID)	69.47	Barlow Twins (RN50)	SSL (ID)	58.34
	DINO (RN50)	SSL (ID)	86.46	DINO (RN50)	SSL (ID)	80.39	DINO (RN50)	SSL (ID)	71.6	DINO (RN50)	SSL (ID)	62.21	SWaV (RN50w2)	SSL (ID)	55.82
	Barlow Twins (RN50)	SSL (ID)	86.4	SWaV (RN50w2)	SSL (ID)	77.06	SWaV (RN50w2)	SSL (ID)	65.15	SWaV (RN50w2)	SSL (ID)	58.86	DINO (RN50)	SSL (ID)	55.52
	SimCLR (RN50w2)	SSL (ID)	84.6	SWaV (RN50)	SSL (ID)	75.15	SWaV (RN50)	SSL (ID)	64.49	SimCLR (RN50w2)	SSL (ID)	57.47	SimCLR (RN50w2)	SSL (ID)	53.61
	SimCLR (RN50)	SSL (ID)	83.94	SimCLR (RN50w2)	SSL (ID)	69.58	Supervised RegNet	Supervised	62.31	SWaV (RN50)	SSL (ID)	56.92	SWaV (RN50)	SSL (ID)	53.46
	MoCov2 (RN50)	SSL (ID)	81.86	Supervised RegNet	Supervised	67.65	SimCLR (RN50w2)	SSL (ID)	61.46	MoCov2 (RN50)	SSL (ID)	54.93	Supervised (RN50)	Supervised	50.38
	Supervised EffNet	Supervised	81.84	SimCLR (RN50)	SSL (ID)	63.88	MoCov2 (RN50)	SSL (ID)	59.14	Supervised RegNet	Supervised	54.52	MoCov2 (RN50)	SSL (ID)	49.87
	Supervised (RN50)	Supervised	79.71	Supervised (RN50)	Supervised	62.82	Supervised (RN50)	Supervised	56.89	Supervised (RN50)	Supervised	53.1	SimCLR (RN50)	SSL (ID)	49.47
	Supervised RegNet	Supervised	78.57	MoCov2 (RN50)	SSL (ID)	62.42	SimCLR (RN50)	SSL (ID)	56.51	SimCLR (RN50)	SSL (ID)	52.37	Supervised RegNet	Supervised	48.8
	RotNet (RN50)	SSL (PT)	75.12	RotNet (RN50)	SSL (PT)	58.89	RotNet (RN50)	SSL (PT)	54.04	RotNet (RN50)	SSL (PT)	48.6	RotNet (RN50)	SSL (PT)	45.36
	Jigsaw (RN50)	SSL (PT)	57.98	Jigsaw (RN50)	SSL (PT)	44.87	Jigsaw (RN50)	SSL (PT)	41.86	Jigsaw (RN50)	SSL (PT)	38.55	Jigsaw (RN50)	SSL (PT)	35.75
	Colorization (RN50)	SSL (PT)	30.69	Colorization (RN50)	SSL (PT)	24.57	Colorization (RN50)	SSL (PT)	22.99	Colorization (RN50)	SSL (PT)	22.38	Colorization (RN50)	SSL (PT)	21.54
CIFAR100	Supervised (ViT)	Supervised	75.43	Supervised (ViT)	Supervised	73.99	Supervised (ViT)	Supervised	72.67	Supervised (ViT)	Supervised	68.71	Supervised (ViT)	Supervised	59.73
	SWaV (RN50w2)	SSL (ID)	74.65	DINO (ViT)	SSL (ID)	68.66	DINO (ViT)	SSL (ID)	62.82	DINO (ViT)	SSL (ID)	55.26	Supervised EffNet	Supervised	48.22
	DINO (ViT)	SSL (ID)	74.43	Barlow Twins (RN50)	SSL (ID)	63.03	Supervised EffNet	Supervised	58.25	Supervised EffNet	Supervised	54.36	DINO (ViT)	SSL (ID)	48.2
	SWaV (RN50)	SSL (ID)	70.44	Supervised EffNet	Supervised	60.46	Barlow Twins (RN50)	SSL (ID)	57.92	Barlow Twins (RN50)	SSL (ID)	50.42	Barlow Twins (RN50)	SSL (ID)	43.82
	DINO (RN50)	SSL (ID)	67.78	SWaV (RN50w2)	SSL (ID)	58.73	SWaV (RN50w2)	SSL (ID)	48.87	SWaV (RN50w2)	SSL (ID)	42.86	SWaV (RN50w2)	SSL (ID)	40.41
	Barlow Twins (RN50)	SSL (ID)	67.66	SWaV (RN50)	SSL (ID)	57.01	SWaV (RN50)	SSL (ID)	47.91	SWaV (RN50)	SSL (ID)	42.28	SWaV (RN50)	SSL (ID)	38.74
	SimCLR (RN50w2)	SSL (ID)	63.91	DINO (RN50)	SSL (ID)	56.43	Supervised RegNet	Supervised	47.04	Supervised RegNet	Supervised	41.95	DINO (RN50)	SSL (ID)	37.24
	SimCLR (RN50)	SSL (ID)	62.8	Supervised RegNet	Supervised	50.82	DINO (RN50)	SSL (ID)	45.55	DINO (RN50)	SSL (ID)	39.68	Supervised RegNet	Supervised	37.14
	Supervised EffNet	Supervised	61.57	SimCLR (RN50w2)	SSL (ID)	48.09	Supervised (RN50)	Supervised	40.46	Supervised (RN50)	Supervised	36.83	Supervised (RN50)	Supervised	34.07
	MoCov2 (RN50)	SSL (ID)	60.49	Supervised (RN50)	Supervised	44.71	SimCLR (RN50w2)	SSL (ID)	39.59	SimCLR (RN50w2)	SSL (ID)	35.56	SimCLR (RN50w2)	SSL (ID)	33.97
	Supervised (RN50)	Supervised	60.17	MoCov2 (RN50)	SSL (ID)	43.5	MoCov2 (RN50)	SSL (ID)	38.88	MoCov2 (RN50)	SSL (ID)	35.18	SimCLR (RN50)	SSL (ID)	33.06
	Supervised RegNet	Supervised	58.75	SimCLR (RN50)	SSL (ID)	42.66	SimCLR (RN50)	SSL (ID)	36.68	SimCLR (RN50)	SSL (ID)	34.05	MoCov2 (RN50)	SSL (ID)	33.03
	RotNet (RN50)	SSL (PT)	49.59	RotNet (RN50)	SSL (PT)	35.28	RotNet (RN50)	SSL (PT)	31.39	RotNet (RN50)	SSL (PT)	28.73	RotNet (RN50)	SSL (PT)	27.11
	Jigsaw (RN50)	SSL (PT)	30.8	Jigsaw (RN50)	SSL (PT)	20.19	Jigsaw (RN50)	SSL (PT)	18.82	Jigsaw (RN50)	SSL (PT)	17.38	Jigsaw (RN50)	SSL (PT)	16.56
	Colorization (RN50)	SSL (PT)	11.97	Colorization (RN50)	SSL (PT)	8.74	Colorization (RN50)	SSL (PT)	7.41	Colorization (RN50)	SSL (PT)	7.07	Colorization (RN50)	SSL (PT)	6.85
Color	SWaV (RN50w2)	SSL (ID)	94.04	DINO (ViT)	SSL (ID)	91.94	Supervised (ViT)	Supervised	90.75	Supervised (ViT)	Supervised	88.64	Supervised (ViT)	Supervised	81.9
	DINO (ViT)	SSL (ID)	93.78	SWaV (RN50w2)	SSL (ID)	91.66	DINO (ViT)	SSL (ID)	89.79	DINO (ViT)	SSL (ID)	86.03	DINO (ViT)	SSL (ID)	80.46
	DINO (RN50)	SSL (ID)	93.55	Supervised (ViT)	Supervised	91.66	SWaV (RN50w2)	SSL (ID)	88.47	SWaV (RN50w2)	SSL (ID)	84.09	SWaV (RN50w2)	SSL (ID)	77.28
	SWaV (RN50)	SSL (ID)	92.97	DINO (RN50)	SSL (ID)	91.48	DINO (RN50)	SSL (ID)	87.85	DINO (RN50)	SSL (ID)	83.76	DINO (RN50)	SSL (ID)	77.02
	SimCLR (RN50w2)	SSL (ID)	92.38	Barlow Twins (RN50)	SSL (ID)	90.29	Barlow Twins (RN50)	SSL (ID)	87.76	Barlow Twins (RN50)	SSL (ID)	83.06	Barlow Twins (RN50)	SSL (ID)	75.72
	SWaV (RN50)	SSL (ID)	92.18	SimCLR (RN50w2)	SSL (ID)	89.87	Supervised EffNet	Supervised	86.21	Supervised EffNet	Supervised	82.34	Supervised EffNet	Supervised	75.59
	Barlow Twins (RN50)	SSL (ID)	91.35	SWaV (RN50)	SSL (ID)	88.75	SimCLR (RN50w2)	SSL (ID)	84.7	SimCLR (RN50w2)	SSL (ID)	77.29	MoCov2 (RN50)	SSL (ID)	73.23
	SimCLR (RN50)	SSL (ID)	91.22	Supervised EffNet	Supervised	88.17	SimCLR (RN50)	SSL (ID)	83.22	MoCov2 (RN50)	SSL (ID)	77.23	SWaV (RN50)	SSL (ID)	72.42
	Supervised RegNet	Supervised	90.37	SimCLR (RN50)	SSL (ID)	87.24	Supervised RegNet	Supervised	82.91	Supervised RegNet	Supervised	77.05	Supervised RegNet	Supervised	70.99
	Supervised EffNet	Supervised	89.33	Supervised RegNet	Supervised	87.02	SWaV (RN50)	SSL (ID)	82.11	SWaV (RN50)	SSL (ID)	76.38	SimCLR (RN50)	SSL (ID)	68.73
	MoCov2 (RN50)	SSL (ID)	89.07	MoCov2 (RN50)	SSL (ID)	86.63	MoCov2 (RN50)	SSL (ID)	82.03	SimCLR (RN50w2)	SSL (ID)	75.22	Supervised (RN50)	Supervised	65.63
	Supervised (RN50)	Supervised	88.66	Supervised (RN50)	Supervised	84.67	Supervised (RN50)	Supervised	74.92	Supervised (RN50)	Supervised	68.5	SimCLR (RN50w2)	SSL (ID)	64.68
	RotNet (RN50)	SSL (PT)	82.48	RotNet (RN50)	SSL (PT)	62.04	RotNet (RN50)	SSL (PT)	57.07	RotNet (RN50)	SSL (PT)	56.95	RotNet (RN50)	SSL (PT)	56.84
	Jigsaw (RN50)	SSL (PT)	73.81	Colorization (RN50)	SSL (PT)	61.45	Jigsaw (RN50)	SSL (PT)	52.97	Jigsaw (RN50)	SSL (PT)	52.29	Jigsaw (RN50)	SSL (PT)	52.46
	Colorization (RN50)	SSL (PT)	69.37	Jigsaw (RN50)	SSL (PT)	55.59	Colorization (RN50)	SSL (PT)	46.23	Colorization (RN50)	SSL (PT)	40.75	Colorization (RN50)	SSL (PT)	39.43
Caltech101	DINO (ViT)	SSL (ID)	96.45	Supervised (ViT)	Supervised	94.05	DINO (ViT)	SSL (ID)	93.31	DINO (ViT)	SSL (ID)	91.01	Supervised (ViT)	Supervised	83.21
	SWaV (RN50w2)	SSL (ID)	95.34	DINO (ViT)	SSL (ID)	93.87	Supervised (ViT)	Supervised	92.22	Supervised (ViT)	Supervised	90.32	DINO (ViT)	SSL (ID)	81.66
	DINO (RN50)	SSL (ID)	95.02	Supervised EffNet	Supervised	93.4	Supervised EffNet	Supervised	91.86	Supervised EffNet	Supervised	88.99	Supervised EffNet	Supervised	81.52
	Supervised EffNet	Supervised	94.72	SWaV (RN50w2)	SSL (ID)	91.97	Barlow Twins (RN50)	SSL (ID)	89.72	Supervised RegNet	Supervised	86.37	Supervised RegNet	Supervised	76.92
	Supervised (ViT)	Supervised	94.63	Barlow Twins (RN50)	SSL (ID)	91.29	Supervised RegNet	Supervised	89.72	Barlow Twins (RN50)	SSL (ID)	84.07	Barlow Twins (RN50)	SSL (ID)	71.61
	Barlow Twins (RN50)	SSL (ID)	94.29	DINO (RN50)	SSL (ID)	91.21	SWaV (RN50w2)	SSL (ID)	88.44	SWaV (RN50w2)	SSL (ID)	79.46	DINO (RN50)	SSL (ID)	68.77
	SWaV (RN50)	SSL (ID)	94.28	Supervised RegNet	Supervised	90.99	DINO (RN50)	SSL (ID)	87.23	DINO (RN50)	SSL (ID)	79.34	SWaV (RN50w2)	SSL (ID)	66.58
	SimCLR (RN50w2)	SSL (ID)	94.21	SWaV (RN50)	SSL (ID)	90.25	SWaV (RN50)	SSL (ID)	87.1	SimCLR (RN50w2)	SSL (ID)	77.81	SimCLR (RN50w2)	SSL (ID)	65.78
	Supervised RegNet	Supervised	94.08	SimCLR (RN50w2)	SSL (ID)	89.85	SimCLR (RN50w2)	SSL (ID)	86.08	SWaV (RN50)	SSL (ID)	76.59	Supervised (RN50)	Supervised	64.77
	SimCLR (RN50)	SSL (ID)	92.72	Supervised (RN50)	Supervised	88.83	Supervised (RN50)	Supervised	82.77	Supervised (RN50)	Supervised	74.95	SWaV (RN50)	SSL (ID)	63.82
	Supervised (RN50)	Supervised	92.67	SimCLR (RN50)	SSL (ID)	87.78	SimCLR (RN50)	SSL (ID)	81.75	SimCLR (RN50)	SSL (ID)	72.72	SimCLR (RN50)	SSL (ID)	63.31
	MoCov2 (RN50)	SSL (ID)	90.37	MoCov2 (RN50)	SSL (ID)	84.92	MoCov2 (RN50)	SSL (ID)	81.29	MoCov2 (RN50)	SSL (ID)	71.37	MoCov2 (RN50)	SSL (ID)	60.73
	RotNet (RN50)	SSL (PT)	79.45	RotNet (RN50)	SSL (PT)	66.53	RotNet (RN50)	SSL (PT)	58.29	RotNet (RN50)	SSL (PT)	51.73	RotNet (RN50)	SSL (PT)	45.35
	Jigsaw (RN50)	SSL (PT)	71.46	Jigsaw (RN50)	SSL (PT)	57.21	Jigsaw (RN50)	SSL (PT)	49.04	Jigsaw (RN50)	SSL (PT)	42.29	Jigsaw (RN50)	SSL (PT)	38.46
	Colorization (RN50)	SSL (PT)	42.9	Colorization (RN50)	SSL (PT)	31.59	Colorization (RN50)	SSL (PT)	27.06	Colorization (RN50)	SSL (PT)	24.98	Colorization (RN50)	SSL (PT)	24.16
Aircraft	DINO (ViT)	SSL (ID)	59.65	DINO (ViT)	SSL (ID)	53.95	DINO (ViT)	SSL (ID)	48.06	DINO (ViT)	SSL (ID)	41.95	DINO (ViT)	SSL (ID)	36.12
	DINO (RN50)	SSL (ID)	59.29	DINO (RN50)	SSL (ID)	47.72	Barlow Twins (RN50)	SSL (ID)	41.62	Barlow Twins (RN50)	SSL (ID)	35.8	DINO (RN50)	SSL (ID)	32.33
	SWaV (RN50w2)	SSL (ID)	56.14	SWaV (RN50w2)	SSL (ID)	47.58	DINO (RN50)	SSL (ID)	40.94	DINO (RN50)	SSL (ID)	35.57	Barlow Twins (RN50)	SSL (ID)	31.76
	Barlow Twins (RN50)	SSL (ID)	55.32	SWaV (RN50w2)	SSL (ID)	44.9	SWaV (RN50w2)	SSL (ID)	38.45	SWaV (RN50w2)	SSL (ID)	33.69	SWaV (RN50w2)	SSL (ID)	31.31
	SWaV (RN50)	SSL (ID)	54.73	SWaV (RN50)	SSL (ID)	43.88	SWaV (RN50)	SSL (ID)	38.31	SWaV (RN50)	SSL (ID)	33.62	SWaV (RN50)	SSL (ID)	30.15
	Supervised RegNet	Supervised	47.19	Supervised EffNet	Supervised	41.2	Supervised (ViT)	Supervised	37.17	Supervised (ViT)	Supervised	32.63	Supervised RegNet	Supervised	27.67
	SimCLR (RN50w2)	SSL (ID)	45.28	Supervised (ViT)	Supervised	40.31	Supervised EffNet	Supervised	36.82	Supervised RegNet	Supervised	31.43	Supervised (ViT)	Supervised	27.41
	Supervised EffNet	Supervised	45.09	Supervised RegNet	Supervised	39.84	Supervised RegNet	Supervised	36.75	Supervised EffNet	Supervised	31.34	Supervised EffNet	Supervised	26.69
	Supervised (ViT)	Supervised	43.77	SimCLR (RN50w2)	SSL (ID)	37.74	SimCLR (RN50w2)	SSL (ID)	32.95	SimCLR (RN50w2)	SSL (ID)	28.96	SimCLR (RN50w2)	SSL (ID)	25.6
	Supervised (RN50)	Supervised	43.39	Supervised (RN50)	Supervised	33.57	Supervised (RN50)	Supervised	29.25	Supervised (RN50)	Supervised	25.78	Supervised (RN50)	Supervised	23.63
	MoCov2 (RN50)	SSL (ID)	41.44	MoCov2 (RN50)	SSL (ID)	31.0	MoCov2 (RN50)	SSL (ID)	27.08	MoCov2 (RN50)	SSL (ID)	24.57	MoCov2 (RN50)	SSL (ID)	22.21
	SimCLR (RN50)	SSL (ID)	40.2	MoCov2 (RN50)	SSL (ID)	29.44	SimCLR (RN50)	SSL (ID)	26.81	SimCLR (RN50)	SSL (ID)	23.5	SimCLR (RN50)	SSL (ID)	21.63
	RotNet (RN50)	SSL (PT)	25.33	RotNet (RN50)	SSL (PT)	16.18	RotNet (RN50)	SSL (PT)	14.7	RotNet (RN50)	SSL (PT)	13.93	RotNet (RN50)	SSL (PT)	13.53
	Colorization (RN50)	SSL (PT)	14.89	Colorization (RN50)	SSL (PT)	9.36	Colorization (RN50)	SSL (PT)	8.28	Colorization (RN50)	SSL (PT)	7.85	Colorization (RN50)	SSL (PT)	8.11
	Jigsaw (RN50)	SSL (PT)	12.79	Jigsaw (RN50)	SSL (PT)	7.75	Jigsaw (RN50)	SSL (PT)	7.54	Jigsaw (RN50)	SSL (PT)	7.39	Jigsaw (RN50)	SSL (PT)	7.2
Birdsnap	DINO (ViT)	SSL (ID)	59.55	DINO (ViT)	SSL (ID)	56.37	DINO (ViT)	SSL (ID)	53.54	DINO (ViT)	SSL (ID)	48.76	DINO (ViT)	SSL (ID)	41.18
	Supervised (ViT)	Supervised	49.39	Supervised (ViT)	Supervised	47.83	Supervised (ViT)	Supervised	45.53	Supervised (ViT)	Supervised	42.56	Supervised (ViT)	Supervised	36.1
	Supervised RegNet	Supervised	43.74	Supervised RegNet	Supervised	39.97	Supervised RegNet	Supervised	38.12	Supervised RegNet	Supervised	34.34	Supervised RegNet	Supervised	29.28
	Supervised EffNet	Supervised	42.69	Supervised EffNet	Supervised	39.8	Supervised EffNet	Supervised	37.73	Supervised EffNet	Supervised	33.07	Supervised EffNet	Supervised	28.18
	DINO (RN50)	SSL (ID)	42.24	DINO (RN50)	SSL (ID)	38.32	DINO (RN50)	SSL (ID)	34.25	DINO (RN50)	SSL (ID)	30.02	DINO (RN50)	SSL (ID)	26.41
	SWaV (RN50)	SSL (ID)	41.4	SWaV (RN50)	SSL (ID)	36.99	Barlow Twins (RN50)	SSL (ID)	32.48	SWaV (RN50)	SSL (ID)	28.03	SWaV (RN50)	SSL (ID)	24.33
	SWaV (RN50w2)	SSL (ID)	39.42	Barlow Twins (RN50)	SSL (ID)	36.48	SWaV (RN50)	SSL (ID)	32.32	Barlow Twins (RN50)	SSL (ID)	27.64	SWaV (RN50w2)	SSL (ID)	23.59
	Barlow Twins (RN50)	SSL (ID)	38.59	SWaV (RN50w2)	SSL (ID)	34.48	SWaV (RN50w2)	SSL (ID)	30.76	SWaV (RN50w2)	SSL (ID)	27.11	Barlow Twins (RN50)	SSL (ID)	23.13
	Supervised (RN50)	Supervised	36.27	Supervised (RN50)	Supervised	32.85	Supervised (RN50)	Supervised	29.73	Supervised (RN50)	Supervised	25.17	Supervised (RN50)	Supervised	21.74
	SimCLR (RN50w2)	SSL (ID)	34.42	SimCLR (RN50w2)	SSL (ID)	32.06	SimCLR (RN50w2)	SSL (ID)	28.92	SimCLR (RN50w2)	SSL (ID)	24.76	SimCLR (RN50w2)	SSL (ID)	20.84
	SimCLR (RN50)	SSL (ID)	30.96	SimCLR (RN50)	SSL (ID)	27.95	SimCLR (RN50)	SSL (ID)	24.97	SimCLR (RN50)	SSL (ID)	20.99	SimCLR (RN50)	SSL (ID)	17.91
	MoCov2 (RN50)	SSL (ID)	24.72	MoCov2 (RN50)	SSL (ID)	21.51	MoCov2 (RN50)	SSL (ID)	19.9	MoCov2 (RN50)	SSL (ID)	16.96	MoCov2 (RN50)	SSL (ID)	14.61
	RotNet (RN50)	SSL (PT)	6.99	RotNet (RN50)	SSL (PT)	4.36	RotNet (RN50)	SSL (PT)	4.03	RotNet (RN50)	SSL (PT)	3.8	RotNet (RN50)	SSL (PT)	3.67
	Jigsaw (RN50)	SSL (PT)	3.86	Jigsaw (RN50)	SSL (PT)	2.75	Jigsaw (RN50)	SSL (PT)	2.47	Jigsaw (RN50)	SSL (PT)	2.49	Jigsaw (RN50)	SSL (PT)	2.13
	Colorization (RN50)	SSL (PT)	1.76	Colorization (RN50)	SSL (PT)	1.3	Colorization (RN50)	SSL (PT)	1.02	Colorization (RN50)	SSL (PT)	1.14	Colorization (RN50)	SSL (PT)	1.09
Food	SWaV (RN50w2)	SSL (ID)	75.86	SWaV (RN50w2)	SSL (ID)	71.75	DINO (ViT)	SSL (ID)	68.06	DINO (ViT)	SSL (ID)	62.65	DINO (ViT)	SSL (ID)	54.52
	DINO (ViT)	SSL (ID)	74.47	DINO (ViT)	SSL (ID)	71.33	SWaV (RN50w2)	SSL (ID)	67.01	SWaV (RN50w2)	SSL (ID)	60.2	SWaV (RN50w2)	SSL (ID)	51.51
	SWaV (RN50)	SSL (ID)	72.27	DINO (RN50)	SSL (ID)	69.24	DINO (RN50)	SSL (ID)	65.59	DINO (RN50)	SSL (ID)	59.53	DINO (RN50)	SSL (ID)	50.93
	DINO (RN50)	SSL (ID)	72.25	SWaV (RN50)	SSL (ID)	68.3	SimCLR (RN50w2)	SSL (ID)	63.5	Supervised (ViT)	Supervised	56.09	Supervised (ViT)	Supervised	47.77
	SimCLR (RN50w2)	SSL (ID)	71.78	SimCLR (RN50w2)	SSL (ID)	67.67	SWaV (RN50)	SSL (ID)	62.25	SimCLR (RN50w2)	SSL (ID)	55.83	SWaV (RN50)	SSL (ID)	46.53
	Barlow Twins (RN50)	SSL (ID)	67.75	Barlow Twins (RN50)	SSL (ID)	63.85	Supervised (ViT)	Supervised	60.8	SWaV (RN50)	SSL (ID)	54.2	SimCLR (RN50w2)	SSL (ID)	46.48
	SimCLR (RN50)	SSL (ID)	67.58	SimCLR (RN50)	SSL (ID)	63.28	Barlow Twins (RN50)	SSL (ID)	60.24	Barlow Twins (RN50)	SSL (ID)	53.97	Barlow Twins (RN50)	SSL (ID)	45.74
	Supervised (ViT)	Supervised	64.84	Supervised (ViT)	Supervised	63.08	SimCLR (RN50)	SSL (ID)	58.26	Supervised RegNet	Supervised	52.31	SimCLR (RN50)	SSL (ID)	44.83
	Supervised RegNet	Supervised	62.5	MoCov2 (RN50)	SSL (ID)	59.01	Supervised RegNet	Supervised	56.95	SimCLR (RN50)	SSL (ID)	50.81	Supervised RegNet	Supervised	42.94
	MoCov2 (RN50)	SSL (ID)	62.47	Supervised RegNet	Supervised	58.77	MoCov2 (RN50)	SSL (ID)	55.21	Supervised EffNet	Supervised	48.57	Supervised EffNet	Supervised	41.95
	Supervised EffNet	Supervised	56.67	Supervised EffNet	Supervised	54.1	Supervised EffNet	Supervised	52.26	MoCov2 (RN50)	SSL (ID)	48.38	MoCov2 (RN50)	SSL (ID)	41.12
	Supervised (RN50)	Supervised	55.69	Supervised (RN50)	Supervised	51.41	Supervised (RN50)	Supervised	46.63	Supervised (RN50)	Supervised	40.71	Supervised (RN50)	Supervised	35.42
	RotNet (RN50)	SSL (PT)	33.02	RotNet (RN50)	SSL (PT)	27.39	RotNet (RN50)	SSL (PT)	24.39	RotNet (RN50)	SSL (PT)	21.08	RotNet (RN50)	SSL (PT)	18.88
	Jigsaw (RN50)	SSL (PT)	25.73	Jigsaw (RN50)	SSL (PT)	20.11	Jigsaw (RN50)	SSL (PT)	18.27	Jigsaw (RN50)	SSL (PT)	15.48	Jigsaw (RN50)	SSL (PT)	14.24
	Colorization (RN50)	SSL (PT)	6.13	Colorization (RN50)	SSL (PT)	5.33	Colorization (RN50)	SSL (PT)	4.83	Colorization (RN50)	SSL (PT)	4.24	Colorization (RN50)	SSL (PT)	3.89

A.2.2 Bounded Homography

Table 6: Linear evaluation performance on common image classification benchmarks for increasing homography strengths. Models are ranked.

	$H_{0.2}$			$H_{0.4}$			$H_{0.6}$			$H_{0.8}$		
CIFAR10	Supervised (ViT)	Supervised	91.01	Supervised (ViT)	Supervised	88.76	Supervised (ViT)	Supervised	81.44	Supervised (ViT)	Supervised	65.78
	SWaV (RN50w2)	SSL (ID)	90.06	DINO (ViT)	SSL (ID)	87.17	DINO (ViT)	SSL (ID)	78.66	DINO (ViT)	SSL (ID)	61.9
	DINO (ViT)	SSL (ID)	89.62	SWaV (RN50w2)	SSL (ID)	85.67	Barlow Twins (RN50)	SSL (ID)	72.97	SWaV (RN50w2)	SSL (ID)	54.81
	SWaV (RN50)	SSL (ID)	85.84	Barlow Twins (RN50)	SSL (ID)	83.24	SWaV (RN50w2)	SSL (ID)	72.56	Barlow Twins (RN50)	SSL (ID)	53.88
	Barlow Twins (RN50)	SSL (ID)	85.72	DINO (RN50)	SSL (ID)	80.58	Supervised EffNet	Supervised	68.03	Supervised EffNet	Supervised	52.6
	DINO (RN50)	SSL (ID)	85.67	SimCLR (RN50w2)	SSL (ID)	78.94	SimCLR (RN50w2)	SSL (ID)	67.88	SimCLR (RN50w2)	SSL (ID)	51.84
	SimCLR (RN50w2)	SSL (ID)	82.36	SimCLR (RN50)	SSL (ID)	78.11	MoCov2 (RN50)	SSL (ID)	66.55	MoCov2 (RN50)	SSL (ID)	47.99
	SimCLR (RN50)	SSL (ID)	81.92	SWaV (RN50)	SSL (ID)	78.69	SimCLR (RN50)	SSL (ID)	65.77	Supervised (RN50)	Supervised	47.18
	MoCov2 (RN50)	SSL (ID)	79.86	MoCov2 (RN50)	SSL (ID)	77.17	DINO (RN50)	SSL (ID)	64.58	Supervised RegNet	Supervised	46.61
	Supervised EffNet	Supervised	79.81	Supervised EffNet	Supervised	76.43	Supervised RegNet	Supervised	64.19	SimCLR (RN50)	SSL (ID)	46.45
	Supervised (RN50)	Supervised	77.47	Supervised (RN50)	Supervised	74.21	Supervised (RN50)	Supervised	63.77	DINO (RN50)	SSL (ID)	46.28
	Supervised RegNet	Supervised	76.83	Supervised RegNet	Supervised	73.97	SWaV (RN50)	SSL (ID)	60.09	SWaV (RN50)	SSL (ID)	43.76
	RotNet (RN50)	SSL (PT)	72.43	RotNet (RN50)	SSL (PT)	67.95	RotNet (RN50)	SSL (PT)	56.28	RotNet (RN50)	SSL (PT)	40.65
	Jigsaw (RN50)	SSL (PT)	55.86	Jigsaw (RN50)	SSL (PT)	51.69	Jigsaw (RN50)	SSL (PT)	43.22	Jigsaw (RN50)	SSL (PT)	33.19
	Colorization (RN50)	SSL (PT)	29.05	Colorization (RN50)	SSL (PT)	26.56	Colorization (RN50)	SSL (PT)	23.16	Colorization (RN50)	SSL (PT)	20.2
CIFAR100	Supervised (ViT)	Supervised	73.28	Supervised (ViT)	Supervised	70.33	Supervised (ViT)	Supervised	61.6	Supervised (ViT)	Supervised	45.72
	DINO (ViT)	SSL (ID)	72.27	DINO (ViT)	SSL (ID)	68.25	DINO (ViT)	SSL (ID)	57.51	DINO (ViT)	SSL (ID)	40.74
	SWaV (RN50w2)	SSL (ID)	72.12	SWaV (RN50w2)	SSL (ID)	65.48	Barlow Twins (RN50)	SSL (ID)	52.03	Barlow Twins (RN50)	SSL (ID)	35.38
	SWaV (RN50)	SSL (ID)	66.76	Barlow Twins (RN50)	SSL (ID)	62.95	SWaV (RN50w2)	SSL (ID)	50.81	SWaV (RN50w2)	SSL (ID)	34.15
	Barlow Twins (RN50)	SSL (ID)	66.52	DINO (RN50)	SSL (ID)	59.65	Supervised EffNet	Supervised	45.67	Supervised EffNet	Supervised	32.25
	DINO (RN50)	SSL (ID)	65.83	SWaV (RN50)	SSL (ID)	57.38	SimCLR (RN50w2)	SSL (ID)	45.13	SimCLR (RN50w2)	SSL (ID)	30.82
	SimCLR (RN50w2)	SSL (ID)	61.42	SimCLR (RN50w2)	SSL (ID)	56.92	Supervised RegNet	Supervised	44.32	Supervised RegNet	Supervised	30.36
	SimCLR (RN50)	SSL (ID)	60.78	SimCLR (RN50)	SSL (ID)	55.98	DINO (RN50)	SSL (ID)	44.11	Supervised RegNet	Supervised	29.99
	Supervised EffNet	Supervised	58.76	MoCov2 (RN50)	SSL (ID)	55.08	MoCov2 (RN50)	SSL (ID)	43.92	MoCov2 (RN50)	SSL (ID)	28.37
	MoCov2 (RN50)	SSL (ID)	58.59	Supervised EffNet	Supervised	54.35	MoCov2 (RN50)	SSL (ID)	43.8	DINO (RN50)	SSL (ID)	28.25
	Supervised (RN50)	Supervised	57.72	Supervised (RN50)	Supervised	54.04	SimCLR (RN50)	SSL (ID)	42.06	SimCLR (RN50)	SSL (ID)	26.25
	Supervised RegNet	Supervised	57.33	Supervised RegNet	Supervised	53.61	SWaV (RN50)	SSL (ID)	38.52	SWaV (RN50)	SSL (ID)	23.93
	RotNet (RN50)	SSL (PT)	46.62	RotNet (RN50)	SSL (PT)	41.23	RotNet (RN50)	SSL (PT)	29.98	RotNet (RN50)	SSL (PT)	18.52
	Jigsaw (RN50)	SSL (PT)	29.1	Jigsaw (RN50)	SSL (PT)	25.4	Jigsaw (RN50)	SSL (PT)	19.55	Jigsaw (RN50)	SSL (PT)	13.48
	Colorization (RN50)	SSL (PT)	10.98	Colorization (RN50)	SSL (PT)	9.45	Colorization (RN50)	SSL (PT)	7.39	Colorization (RN50)	SSL (PT)	5.55
Color	DINO (ViT)	SSL (ID)	93.57	DINO (ViT)	SSL (ID)	92.71	DINO (ViT)	SSL (ID)	90.14	DINO (ViT)	SSL (ID)	81.18
	SWaV (RN50w2)	SSL (ID)	92.69	SWaV (RN50w2)	SSL (ID)	91.33	Supervised (ViT)	Supervised	87.34	Barlow Twins (RN50)	SSL (ID)	77.44
	DINO (RN50)	SSL (ID)	91.94	Supervised (ViT)	Supervised	90.76	Barlow Twins (RN50)	SSL (ID)	85.76	SWaV (RN50w2)	SSL (ID)	77.29
	SWaV (RN50)	SSL (ID)	91.91	DINO (RN50)	SSL (ID)	90.4	DINO (RN50)	SSL (ID)	85.17	Supervised (ViT)	Supervised	77.23
	Supervised (ViT)	Supervised	91.55	SWaV (RN50)	SSL (ID)	90.34	SWaV (RN50)	SSL (ID)	85.03	SimCLR (RN50w2)	SSL (ID)	75.31
	Barlow Twins (RN50)	SSL (ID)	90.62	Barlow Twins (RN50)	SSL (ID)	89.35	SWaV (RN50w2)	SSL (ID)	84.81	SWaV (RN50w2)	SSL (ID)	74.38
	SimCLR (RN50w2)	SSL (ID)	90.37	SimCLR (RN50w2)	SSL (ID)	88.52	SimCLR (RN50w2)	SSL (ID)	84.36	DINO (RN50)	SSL (ID)	74.36
	SimCLR (RN50)	SSL (ID)	89.98	SimCLR (RN50w2)	SSL (ID)	88.42	SimCLR (RN50)	SSL (ID)	83.64	SimCLR (RN50w2)	SSL (ID)	73.91
	MoCov2 (RN50)	SSL (ID)	88.27	MoCov2 (RN50)	SSL (ID)	88.17	MoCov2 (RN50)	SSL (ID)	80.4	MoCov2 (RN50)	SSL (ID)	70.78
	Supervised RegNet	Supervised	87.58	Supervised RegNet	Supervised	84.93	Supervised (RN50)	Supervised	79.08	Supervised (RN50)	Supervised	70.26
	Supervised EffNet	Supervised	87.23	Supervised (RN50)	Supervised	84.03	Supervised EffNet	Supervised	76.23	Supervised EffNet	Supervised	64.16
	Supervised (RN50)	Supervised	86.26	Supervised EffNet	Supervised	84.01	Supervised RegNet	Supervised	76.19	Supervised RegNet	Supervised	61.88
	RotNet (RN50)	SSL (PT)	77.28	RotNet (RN50)	SSL (PT)	69.46	Jigsaw (RN50)	SSL (PT)	56.99	Jigsaw (RN50)	SSL (PT)	49.08
	Jigsaw (RN50)	SSL (PT)	65.85	Jigsaw (RN50)	SSL (PT)	62.93	RotNet (RN50)	SSL (PT)	55.89	RotNet (RN50)	SSL (PT)	44.62
	Colorization (RN50)	SSL (PT)	65.19	Colorization (RN50)	SSL (PT)	55.53	Colorization (RN50)	SSL (PT)	44.24	Colorization (RN50)	SSL (PT)	35.7
Caltech101	SWaV (RN50w2)	SSL (ID)	94.55	DINO (ViT)	SSL (ID)	93.13	DINO (ViT)	SSL (ID)	90.77	DINO (ViT)	SSL (ID)	80.4
	DINO (ViT)	SSL (ID)	94.18	SWaV (RN50w2)	SSL (ID)	92.38	Supervised (ViT)	Supervised	89.14	Supervised (ViT)	Supervised	79.16
	Supervised (ViT)	Supervised	93.24	Supervised (ViT)	Supervised	92.08	SWaV (RN50w2)	SSL (ID)	88.1	Barlow Twins (RN50)	SSL (ID)	76.21
	Supervised EffNet	Supervised	93.02	Barlow Twins (RN50)	SSL (ID)	91.81	Barlow Twins (RN50)	SSL (ID)	87.75	SWaV (RN50w2)	SSL (ID)	75.16
	SimCLR (RN50w2)	SSL (ID)	92.95	SimCLR (RN50w2)	SSL (ID)	91.7	SimCLR (RN50w2)	SSL (ID)	87.62	SimCLR (RN50w2)	SSL (ID)	73.68
	Barlow Twins (RN50)	SSL (ID)	92.91	SWaV (RN50)	SSL (ID)	90.86	SWaV (RN50)	SSL (ID)	86.89	SWaV (RN50)	SSL (ID)	71.39
	SWaV (RN50)	SSL (ID)	92.82	Supervised EffNet	Supervised	90.44	DINO (RN50)	SSL (ID)	85.58	SimCLR (RN50)	SSL (ID)	71.29
	Supervised RegNet	Supervised	91.52	DINO (RN50)	SSL (ID)	89.37	Supervised EffNet	Supervised	84.86	DINO (RN50)	SSL (ID)	70.07
	DINO (RN50)	SSL (ID)	91.51	SimCLR (RN50)	SSL (ID)	89.0	SimCLR (RN50)	SSL (ID)	84.36	Supervised EffNet	Supervised	69.8
	SimCLR (RN50)	SSL (ID)	91.24	Supervised RegNet	Supervised	88.73	Supervised RegNet	Supervised	82.37	Supervised RegNet	Supervised	66.37
	Supervised (RN50)	Supervised	89.63	Supervised (RN50)	Supervised	86.93	Supervised (RN50)	Supervised	81.03	Supervised (RN50)	Supervised	65.97
	MoCov2 (RN50)	SSL (ID)	88.89	MoCov2 (RN50)	SSL (ID)	86.52	MoCov2 (RN50)	SSL (ID)	79.84	MoCov2 (RN50)	SSL (ID)	62.94
	RotNet (RN50)	SSL (PT)	76.24	RotNet (RN50)	SSL (PT)	70.62	RotNet (RN50)	SSL (PT)	57.64	RotNet (RN50)	SSL (PT)	37.29
	Jigsaw (RN50)	SSL (PT)	63.13	Jigsaw (RN50)	SSL (PT)	55.13	Jigsaw (RN50)	SSL (PT)	40.59	Jigsaw (RN50)	SSL (PT)	26.3
	Colorization (RN50)	SSL (PT)	36.68	Colorization (RN50)	SSL (PT)	30.2	Colorization (RN50)	SSL (PT)	22.2	Colorization (RN50)	SSL (PT)	16.55
Aircraft	DINO (ViT)	SSL (ID)	57.47	DINO (ViT)	SSL (ID)	54.08	DINO (ViT)	SSL (ID)	47.19	DINO (ViT)	SSL (ID)	33.81
	DINO (RN50)	SSL (ID)	54.85	DINO (RN50)	SSL (ID)	50.25	DINO (RN50)	SSL (ID)	39.61	SWaV (RN50w2)	SSL (ID)	24.52
	Barlow Twins (RN50)	SSL (ID)	52.51	Barlow Twins (RN50)	SSL (ID)	47.61	Barlow Twins (RN50)	SSL (ID)	38.04	Barlow Twins (RN50)	SSL (ID)	24.51
	SWaV (RN50w2)	SSL (ID)	51.18	SWaV (RN50w2)	SSL (ID)	47.48	SWaV (RN50w2)	SSL (ID)	37.43	DINO (RN50)	SSL (ID)	23.74
	SWaV (RN50)	SSL (ID)	49.21	SWaV (RN50)	SSL (ID)	46.45	SWaV (RN50)	SSL (ID)	34.61	SWaV (RN50)	SSL (ID)	21.64
	Supervised EffNet	Supervised	39.91	Supervised EffNet	Supervised	36.19	Supervised (ViT)	Supervised	30.1	Supervised (ViT)	Supervised	20.47
	Supervised (ViT)	Supervised	36.98	Supervised (ViT)	Supervised	34.51	Supervised EffNet	Supervised	29.53	Supervised EffNet	Supervised	18.91
	SimCLR (RN50)	SSL (ID)	36.92	SimCLR (RN50)	SSL (ID)	32.9	MoCov2 (RN50)	SSL (ID)	23.54	MoCov2 (RN50)	SSL (ID)	14.46
	MoCov2 (RN50)	SSL (ID)	36.45	MoCov2 (RN50)	SSL (ID)	32.21	SimCLR (RN50)	SSL (ID)	22.16	Supervised (RN50)	Supervised	13.86
	Supervised RegNet	Supervised	35.02	Supervised RegNet	Supervised	29.98	Supervised RegNet	Supervised	21.87	SimCLR (RN50w2)	SSL (ID)	13.51
	SimCLR (RN50w2)	SSL (ID)	34.95	SimCLR (RN50w2)	SSL (ID)	29.4	Supervised (RN50)	Supervised	20.65	Supervised RegNet	Supervised	13.36
	Supervised (RN50)	Supervised	32.81	Supervised (RN50)	Supervised	28.34	SimCLR (RN50w2)	SSL (ID)	20.49	SimCLR (RN50)	SSL (ID)	12.93
	RotNet (RN50)	SSL (PT)	18.86	RotNet (RN50)	SSL (PT)	15.27	RotNet (RN50)	SSL (PT)	11.69	RotNet (RN50)	SSL (PT)	7.9
	Colorization (RN50)	SSL (PT)	10.02	Colorization (RN50)	SSL (PT)	6.8	Colorization (RN50)	SSL (PT)	4.35	Colorization (RN50)	SSL (PT)	3.37
	Jigsaw (RN50)	SSL (PT)	6.88	Jigsaw (RN50)	SSL (PT)	5.7	Jigsaw (RN50)	SSL (PT)	4.28	Jigsaw (RN50)	SSL (PT)	3.11
Birdsnap	DINO (ViT)	SSL (ID)	57.45	DINO (ViT)	SSL (ID)	57.18	DINO (ViT)	SSL (ID)	54.59	DINO (ViT)	SSL (ID)	41.57
	Supervised (ViT)	Supervised	47.27	Supervised (ViT)	Supervised	46.48	Supervised (ViT)	Supervised	42.36	Supervised (ViT)	Supervised	31.04
	DINO (RN50)	SSL (ID)	40.91	SWaV (RN50)	SSL (ID)	40.27	SWaV (RN50)	SSL (ID)	34.97	SWaV (RN50w2)	SSL (ID)	24.81
	SWaV (RN50)	SSL (ID)	40.24	DINO (RN50)	SSL (ID)	39.95	SWaV (RN50w2)	SSL (ID)	34.82	SWaV (RN50)	SSL (ID)	24.46
	SWaV (RN50w2)	SSL (ID)	39.05	SWaV (RN50w2)	SSL (ID)	38.97	DINO (RN50)	SSL (ID)	34.51	DINO (RN50)	SSL (ID)	22.39
	Supervised EffNet	Supervised	38.49	Barlow Twins (RN50)	SSL (ID)	37.68	Barlow Twins (RN50)	SSL (ID)	32.68	Barlow Twins (RN50)	SSL (ID)	22.36
	Barlow Twins (RN50)	SSL (ID)	38.04	Supervised EffNet	Supervised	36.66	Supervised EffNet	Supervised	31.43	Supervised EffNet	Supervised	21.08
	Supervised RegNet	Supervised	36.12	Supervised RegNet	Supervised	32.85	SimCLR (RN50w2)	SSL (ID)	26.13	SimCLR (RN50w2)	SSL (ID)	18.87
	SimCLR (RN50w2)	SSL (ID)	32.56	SimCLR (RN50w2)	SSL (ID)	30.94	Supervised RegNet	Supervised	25.35	Supervised RegNet	Supervised	16.7
	Supervised (RN50)	Supervised	31.08	Supervised (RN50)	Supervised	28.98	SimCLR (RN50)	SSL (ID)	23.72	SimCLR (RN50)	SSL (ID)	15.61
	SimCLR (RN50)	SSL (ID)	29.45	SimCLR (RN50)	SSL (ID)	28.55	Supervised (RN50)	Supervised	23.37	Supervised (RN50)	Supervised	14.74
	MoCov2 (RN50)	SSL (ID)	23.8	MoCov2 (RN50)	SSL (ID)	23.14	MoCov2 (RN50)	SSL (ID)	18.56	MoCov2 (RN50)	SSL (ID)	12.12
	RotNet (RN50)	SSL (PT)	6.27	RotNet (RN50)	SSL (PT)	5.81	RotNet (RN50)	SSL (PT)	4.61	RotNet (RN50)	SSL (PT)	3.21
	Jigsaw (RN50)	SSL (PT)	2.6	Jigsaw (RN50)	SSL (PT)	1.92	Jigsaw (RN50)	SSL (PT)	1.72	Jigsaw (RN50)	SSL (PT)	1.23
	Colorization (RN50)	SSL (PT)	1.49	Colorization (RN50)	SSL (PT)	1.08	Colorization (RN50)	SSL (PT)	1.02	Colorization (RN50)	SSL (PT)	0.68
Food	SWaV (RN50w2)	SSL (ID)	74.42	SWaV (RN50w2)	SSL (ID)	73.08	DINO (ViT)	SSL (ID)	66.48	DINO (ViT)	SSL (ID)	53.85
	DINO (ViT)	SSL (ID)	72.51	DINO (ViT)	SSL (ID)	71.21	SWaV (RN50w2)	SSL (ID)	66.30	SWaV (RN50w2)	SSL (ID)	51.77
	SWaV (RN50)	SSL (ID)	71.16	SWaV (RN50)	SSL (ID)	69.71	SWaV (RN50)	SSL (ID)	62.52	SWaV (RN50)	SSL (ID)	47.52
	DINO (RN50)	SSL (ID)	70.2	DINO (RN50)	SSL (ID)	68.32	DINO (RN50)	SSL (ID)	60.84	Barlow Twins (RN50)	SSL (ID)	45.23
	SimCLR (RN50w2)	SSL (ID)	68.37	SimCLR (RN50w2)	SSL (ID)	65.97	Barlow Twins (RN50)	SSL (ID)	58.86	DINO (RN50)	SSL (ID)	44.53
	Barlow Twins (RN50)	SSL (ID)	66.91	Barlow Twins (RN50)	SSL (ID)	65.47	SimCLR (RN50w2)	SSL (ID)	58.29	Supervised (ViT)	Supervised	43.61
	SimCLR (RN50)	SSL (ID)	65.76	SimCLR (RN50)	SSL (ID)	63.79	SimCLR (RN50)	SSL (ID)	55.05	SimCLR (RN50w2)	SSL (ID)	42.52
	MoCov2 (RN50)	SSL (ID)	61.26	MoCov2 (RN50)	SSL (ID)	59.43	Supervised (ViT)	Supervised	54.35	SimCLR (RN50)	SSL (ID)	40.07
	Supervised (ViT)	Supervised	61.13	Supervised (ViT)	Supervised	59.22	MoCov2 (RN50)	SSL (ID)	51.63	MoCov2 (RN50)	SSL (ID)	37.04
	Supervised RegNet	Supervised	55.44	Supervised RegNet	Supervised	51.91	Supervised RegNet	Supervised	44.71	Supervised EffNet	Supervised	33.47
	Supervised EffNet	Supervised	52.51	Supervised EffNet	Supervised	49.57	Supervised EffNet	Supervised	44.03	Supervised RegNet	Supervised	33.31
	Supervised (RN50)	Supervised	50.72	Supervised (RN50)	Supervised	48.31	Supervised (RN50)	Supervised	42.71	Supervised (RN50)	Supervised	31.44
	RotNet (RN50)	SSL (PT)	28.61	RotNet (RN50)	SSL (PT)	25.85	RotNet (RN50)	SSL (PT)	20.87	RotNet (RN50)	SSL (PT)	14.39
	Jigsaw (RN50)	SSL (PT)	18.08	Jigsaw (RN50)	SSL (PT)	15.58	Jigsaw (RN50)	SSL (PT)	11.87	Jigsaw (RN50)	SSL (PT)	8.4
	Colorization (RN50)	SSL (PT)	4.97	Colorization (RN50)	SSL (PT)	4.35	Colorization (RN50)	SSL (PT)	3.48	Colorization (RN50)	SSL (PT)	2.73

B   Real-World Datasets

B.1   Overall Performance

Table 7: KNN performance on the real-world, multi-view benchmarks. Models are ranked.

	Model	Type	Accuracy
ALOI	SWaV (RN50w2)	SSL (ID)	96.8
	SWaV (RN50)	SSL (ID)	96.22
	DINO (RN50)	SSL (ID)	96.16
	DINO (ViT)	SSL (ID)	96.12
	SimCLR (RN50w2)	SSL (ID)	95.27
	Barlow Twins (RN50)	SSL (ID)	95.1
	SimCLR (RN50)	SSL (ID)	94.91
	Supervised RegNet	Supervised	94.74
	MoCov2 (RN50)	SSL (ID)	94.69
	Supervised (ViT)	Supervised	93.79
	Supervised (RN50)	Supervised	93.75
	Supervised EffNet	Supervised	91.52
	Jigsaw (RN50)	SSL (PT)	86.16
	RotNet (RN50)	SSL (PT)	85.79
	Colorization (RN50)	SSL (PT)	80.18
MVMC	RotNet (RN50)	SSL (PT)	93.14
	MoCov2 (RN50)	SSL (ID)	92.3
	Supervised (ViT)	Supervised	90.89
	SimCLR (RN50)	SSL (ID)	90.83
	SWaV (RN50w2)	SSL (ID)	90.69
	Jigsaw (RN50)	SSL (PT)	90.48
	DINO (ViT)	SSL (ID)	90.39
	SWaV (RN50)	SSL (ID)	90.37
	Supervised RegNet	Supervised	90.14
	DINO (RN50)	SSL (ID)	89.81
	SimCLR (RN50w2)	SSL (ID)	89.77
	Supervised (RN50)	Supervised	89.35
	Barlow Twins (RN50)	SSL (ID)	89.11
	Supervised EffNet	Supervised	85.25
	Colorization (RN50)	SSL (PT)	72.01
MVC	SWaV (RN50w2)	SSL (ID)	97.2
	Jigsaw (RN50)	SSL (PT)	95.11
	Supervised (ViT)	Supervised	94.56
	SWaV (RN50)	SSL (ID)	94.45
	DINO (RN50)	SSL (ID)	94.11
	MoCov2 (RN50)	SSL (ID)	93.67
	DINO (ViT)	SSL (ID)	91.83
	SimCLR (RN50w2)	SSL (ID)	91.73
	Supervised RegNet	Supervised	90.88
	SimCLR (RN50)	SSL (ID)	90.27
	Colorization (RN50)	SSL (PT)	90.24
	Barlow Twins (RN50)	SSL (ID)	88.56
	Supervised (RN50)	Supervised	87.58
	Supervised EffNet	Supervised	80.59
	RotNet (RN50)	SSL (PT)	80.08
Recon3D	Supervised (RN50)	Supervised	94.94
	Supervised RegNet	Supervised	93.2
	DINO (ViT)	SSL (ID)	93.2
	Supervised (ViT)	Supervised	91.91
	DINO (RN50)	SSL (ID)	91.08
	Barlow Twins (RN50)	SSL (ID)	90.21
	SimCLR (RN50w2)	SSL (ID)	89.47
	SWaV (RN50)	SSL (ID)	87.82
	Supervised EffNet	Supervised	86.8
	SWaV (RN50w2)	SSL (ID)	86.21
	SimCLR (RN50)	SSL (ID)	85.24
	MoCov2 (RN50)	SSL (ID)	83.49
	Jigsaw (RN50)	SSL (PT)	66.94
	Colorization (RN50)	SSL (PT)	66.25
	RotNet (RN50)	SSL (PT)	56.6
Stereo Face	SWaV (RN50w2)	SSL (ID)	97.72
	DINO (ViT)	SSL (ID)	97.71
	DINO (RN50)	SSL (ID)	97.2
	SWaV (RN50)	SSL (ID)	96.65
	Supervised RegNet	Supervised	93.91
	SimCLR (RN50w2)	SSL (ID)	93.42
	Supervised EffNet	Supervised	92.81
	Barlow Twins (RN50)	SSL (ID)	92.08
	SimCLR (RN50)	SSL (ID)	90.44
	Supervised (ViT)	Supervised	90.32
	MoCov2 (RN50)	SSL (ID)	89.66
	Supervised (RN50)	Supervised	86.9
	Colorization (RN50)	SSL (PT)	79.64
	Jigsaw (RN50)	SSL (PT)	77.68
	RotNet (RN50)	SSL (PT)	55.16