

# Extending and Analyzing Self-Supervised Learning Across Domains

Bram Wallace and Bharath Hariharan

Cornell University

`bw462@cornell.edu` `bharathh@cs.cornell.edu`

**Abstract.** Self-supervised representation learning has achieved impressive results in recent years, with experiments primarily coming on ImageNet or other similarly large internet imagery datasets. There has been little to no work with these methods on other smaller domains, such as satellite, textural, or biological imagery. We experiment with several popular methods on an unprecedented variety of domains. We discover, among other findings, that Rotation is the most semantically meaningful task, while much of the performance of Jigsaw is attributable to the nature of its induced distribution rather than semantic understanding. Additionally, there are several areas, such as fine-grain classification, where all tasks underperform. We quantitatively and qualitatively diagnose the reasons for these failures and successes via novel experiments studying pretext generalization, random labelings, and implicit dimensionality. Code and models are available at [https://github.com/BramSW/Extending\\_SSRL\\_Across\\_Domains/](https://github.com/BramSW/Extending_SSRL_Across_Domains/).

## 1 Introduction

A good visual representation is key to all visual recognition tasks. However, in current practice, one needs large labeled training sets to train such a representation. Unfortunately, such datasets can be hard to acquire in many domains, such as satellite imagery or the medical domain. This is often either because annotations require expertise and experts have limited time, or the images themselves are limited (as in medicine). To bring the benefits of visual recognition to these disparate domains, we need powerful representation learning techniques that do not require large labeled datasets.

A promising direction is to use self-supervised representation learning (SSRL), which has gained increasing interest over the last few years [15,34,58,18,23,59]. However, past work has primarily evaluated these techniques on general category object recognition in internet imagery (e.g. ImageNet classification) [42]. There has been very little attention on how (and if) these techniques extend to other domains, be they fine-grained classification problems or datasets in biology and medicine. Paradoxically, these domains are often most in need of such techniques precisely because of the lack of labeled training data.

As such, a key question is whether conclusions from benchmarks on self-supervised learning [18,23] which focused on internet imagery, carry over to

this broader universe of recognition problems. In particular, does one technique dominate, or are different pretext tasks useful for different types of domains (**Sec. 5.1**)? Are representations from an ImageNet classifier still the best we can do (**Sec. 5.1**)? Do these answers change when labels are limited (**Sec. 5.1**)? Are there problem domains where all proposed techniques currently fail (**Sec. 5.2**)?

A barrier to answering these questions is our limited understanding of self-supervised techniques themselves. We have seen their *empirical* success on ImageNet, but when they do succeed, what is it that drives their success (**Sec. 5.3**)? Furthermore, what does the space of learned representations look like, for instance in terms of the dimensionality (**Sec. 6.1**) or nearest neighbors (**Sec. 6.2**)?

In this work, we take the first steps towards answering these questions. We evaluate and analyze multiple self-supervised learning techniques (Rotation[15], Instance Discrimination[54] and Jigsaw[34]) on the broadest benchmark yet of 16 domains spanning internet, biological, satellite, and symbolic imagery. We find that Rotation has the best overall accuracy (reflective of rankings on ImageNet), but is outperformed by Instance Discrimination on biological domains (**Sec. 5.1**). When labels are scarce, pretext methods outperform ImageNet initialization and even full supervision on numerous tasks (**Sec. 5.1**). A prominent failure case for SSRL is fine-grained classification problems, due to important cues such as color being discarded during training (**Sec. 5.2**). Finally, when SSRL techniques do succeed, their reason for success varies: Rotation relies more on the semantic nature of the pretext task, compared to Jigsaw and Instance Discrimination (**Sec. 5.3**). Perhaps as a consequence, the representations of Rotation having comparatively higher implicit dimensionality (**Sec. 6.1**).

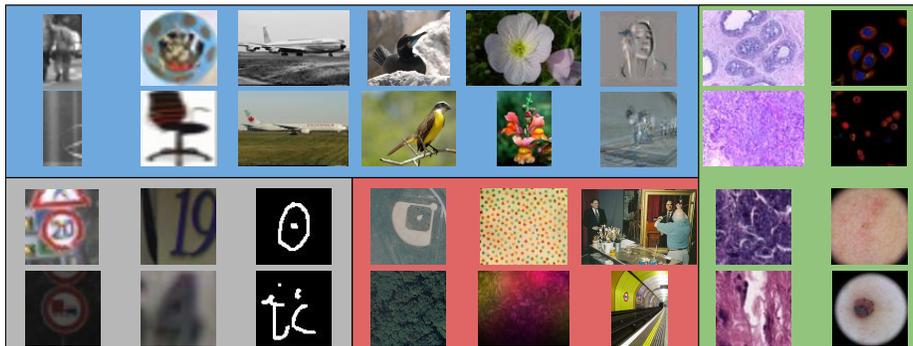
## 2 Datasets

We include 16 datasets in our experiments, significantly more than all prior work. Dataset samples are shown in Figure 1. We group these datasets into 4 categories: **Internet**, **Symbolic**, **Scenes & Textures**, and **Biological**. A summary is shown in Table 1. Some of the datasets in the first three groups are also in the Visual Domain Decathlon (VDD)[40], a multi-task learning benchmark.

**Internet Object Recognition:** This group consists of object recognition problems on internet imagery. We include both coarse-grained (CIFAR100, Daimler Pedestrians) and fine-grained (FGVC-Aircraft, CUB, VGG Flowers) object classification tasks. Finally, we include the “dynamic images” of UCF101, a dataset that possesses many of the same qualitative attributes of the group.

**Symbolic:** We include three well-known symbolic tasks: Omniglot, German Traffic Signs (GTSRB), and Street View House Numbers (SVHN). Though the classification problems might be deemed simple, these offer domains where classification is very different from natural internet imagery: texture is not generally a useful cue and classes follow strict explainable rules.

**Scenes & Textures:** These domains, UC Merced Land Use (satellite imagery), Describable Textures, and Indoor Scenes, all require holistic understand-



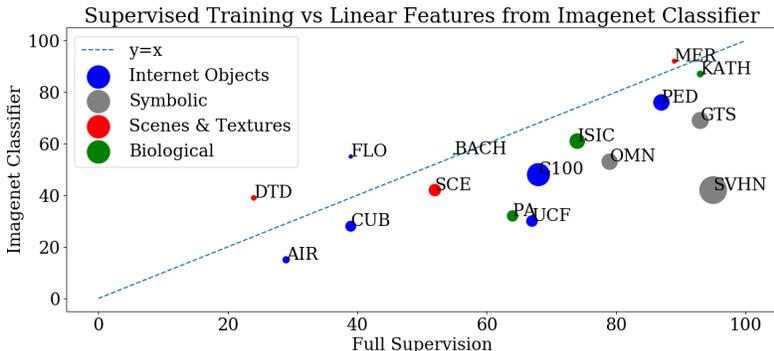
**Fig. 1.** Samples from all datasets. Top rows: Daimler Pedestrians, CIFAR100, FGVC-Aircraft, CU Birds, VGG-Flowers, UCF101, BACH, Protein Atlas. Bottom rows: GT-SRB, SVHN, Omniglot, UC Merced Land Use, Describable Textures, Indoor Scenes, Kather, ISIC. Color coding is by group: **Internet** **Symbolic** **Scenes & Textures** **Biological**

**Table 1.** Summary of the 16 datasets included in our experiments: encompassing fine-grain, symbolic, scene, textural, and biological domains. This is the first work exploring self-supervised representation learning on almost all of these tasks

Name	Type	Size (Train)	Coarse/Fine	Abbreviation
Daimler Pedestrians[31]	Road Object	20k	Coarse	PED
CIFAR100[24]	Internet Object	40k	Coarse	C100
FGVC-Aircraft[28]	Internet Object	3.3k	Fine	AIR
Caltech-UCSD Birds[53]	Internet Object	8.3k	Fine	CUB
VGG-Flowers[33]	Internet Object	1k	Fine	FLO
UCF101[45,3]	Pseudo-Internet Action	9.3k	Coarse	UCF
German Traffic Signs [46]	Symbolic	21k	Coarse	GTS
Street View House Numbers[32]	Symbolic	59k	Coarse	SVHN
Omniglot[25]	Symbolic	19k	Fine	OMN
UC Merced Land Use[55]	Aerial Scene	1.5k	Coarse	MER
Describable Textures[9]	Texture	1.9k	Fine	DTD
Indoor Scene Recognition [39]	Natural Scene	11k	Coarse	SCE
ICIAR BACH[1]	Biological	240	Coarse	BACH
Kather[22]	Biological	3k	Coarse	KATH
Protein Atlas[35]	Biological	9k	Fine	PA
ISIC[10,49]	Biological	17k	Coarse	ISIC

ings, none having an overarching definition of object/symbol. Indoor Scenes does contain internet imagery as in our first group, but is not object-focused.

**Biological:** BACH and Kather consist of histological (microscopic tissue) images of breast and colon cancer respectively, with the classes being the condition/type of cancer. Protein Atlas is microscopy images of human cells, with the goal being classification of the cell part/structure shown. Finally, ISIC is a dermatology dataset consisting of photographs of different types of skin lesions.



**Fig. 2.** Test accuracy of a fully supervised network vs. a linear classifier on top of an ImageNet-classification frozen feature extractor. Marker area indicates dataset size

Before evaluating self-supervision, we study the datasets themselves in terms of difficulty and similarity to ImageNet. To do so, we compare the accuracy of a network trained from scratch with that of a linear classifier operating on an ImageNet-pretrained network (Figure 2). The higher of the two numbers measures the difficulty, while their relationship quantifies the similarity to ImageNet.

We find that small datasets in the Internet domain tend to be the hardest, while large Symbolic datasets are the simplest. The symbolic tasks also have the largest gap between supervision and feature extraction, suggesting that these are the farthest from ImageNet. Overall, the ImageNet feature extractor performance is strongly linearly correlated to that of the fully supervised model ( $p = 0.004$ ). This is expected for the Internet domain, but the similar utility of the feature extractor for the Biological domains is surprising. Dataset size also plays a role, with the pretrained feature extractor working well for smaller datasets.

For fine-grained classification (AIR, CUB, FLO), the supervised models perform comparably, while the ImageNet classifier’s performance varies widely. In addition to the small size of VGG-Flowers, this case is also partly explainable by how these datasets overlap with ImageNet. Almost half of the classes in ImageNet are animals or plants, including flowers, making pretraining especially useful.

### 3 Methods

#### 3.1 Self-Supervised Learning Techniques

In this paper we look at three popular methods, Rotation, Jigsaw, and Instance Discrimination. We also look at the classical technique of Autoencoders as a baseline for the large variety of autoencoder-based pretexts[59,58,38]. We briefly describe each method below, please view the cited works for detailed information.

**Learning by Rotation:** A network is trained to classify the angle of a rotated image among the four choices of 0, 90, 180, or 270 degrees[15].

**Learning by Solving Jigsaw Puzzles:** The image is separated into a 3x3 grid of patches, which are then permuted and fed through a siamese network which must identify the original layout[34,5]. We use 2000 training permutations in our experiments, finding this offered superior performance to 100 or 10,000 under our hyperparameters.

**Instance Discrimination:** Instance Discrimination (ID) maps images to features on the unit sphere with each image being considered as a separate class under a non-parametric softmax classifier[54].

**Autoencoders:** Autoencoders were one of the earliest methods of self-supervised learning[59,50,38,2,20]. An autoencoder learns an encoder-decoder pair of networks that reconstruct the input image.

### 3.2 Architecture & Evaluation

We resize inputs to  $64 \times 64$  and use a ResNet26, as in Rebuffi et al. [40,41]. The lower resolution eases computational burden as well as comparison and future adaptation to the VDD. For Autoencoding, a simple convolutional decoder is used. Features maps of size  $256 \times 8 \times 8$  are extracted before the final pooling layer and average pooled to 256, 4096 ( $256 \times 4 \times 4$ ), or 9216 ( $256 \times 6 \times 6$ ), with 256 being the default. A linear classifier is trained on these features. Training/architecture details are in the Supplementary.

## 4 Related Work

There are three pertinent recent surveys of self-supervision. The first is by Kolesnikov et al. who evaluate several methods on ImageNet and Places205 classification across network architectures[23]. We focus on many domains, an order of magnitude more than their work. The second relevant survey is by Goyal et al., who introduce a benchmark suite of tasks on which to test models as feature extractors. While they scale on a variety of downstream tasks, the pretraining datasets are all internet imagery, either ImageNet or YFCC variants[18,48]. Our work includes a much wider variety of both pretraining and downstream datasets. VTAB tests *pretrained feature extractors* on a variety of datasets, performing self-supervised learning only on ImageNet[57]. Finally, a concurrent paper evaluates these self-supervised techniques as an auxilliary loss for few-shot learning[47].

One trend of inquiry concerns classifiers that perform well on multiple domains while sharing most of the parameters across datasets[40,41,52]. The pre-eminent examples of this are by Rebuffi et al. [40,41] who present approaches on the VDD across 10 different domains. We use these datasets (and more) in our training, but evaluate *self-supervised approaches* in *single-domain* settings.

There has also been prior work using problem/domain-specific SSRL methods, such as [21,27,7,14] in the biological and medical fields or [44] for aerial imagery. Many of these approaches use variations of autoencoding as the pretext task; we include autoencoding in our evaluation. In contrast to these, our focus is on the cross-dataset applicability of these pretexts.

Other SSRL methods include generative models[12,13,16], colorization [26,58], video-based techniques[36,43,37,30,17,51], or generic techniques[6,4,19,29,8]. It is very possible that a subset of these methods could offer improved performance on some of the domains we work with, however in this work we focus on the popular fundamental methods of Rotation, Jigsaw, and Instance Discrimination as well as Autoencoding as a representative set. Doersch and Zisserman use multiple pretexts simultaneously to improve representation quality [11]. This approach could be complementary to the work featured here. Semi-supervised approaches such as S4L[56] are relevant to Section 5.1.

## 5 Downstream task Performance Analysis

### 5.1 Downstream task accuracy

A summary of downstream testing accuracies is shown in Figure 3. The obvious question to ask in this investigation is “which pretext is best”? On ImageNet, per the respective original works, the ordering (best to worst) was Rotation, Instance Discrimination, Jigsaw, Autoencoding. We see that this ranking is not universal: while Rotation is the best pretext on the first three groups, it lags behind even random initialization on the Biological domains. Furthermore, the relative rankings of ID and Jigsaw vary between groups. We investigate what powers the performance of each of these methods in Section 5.3.

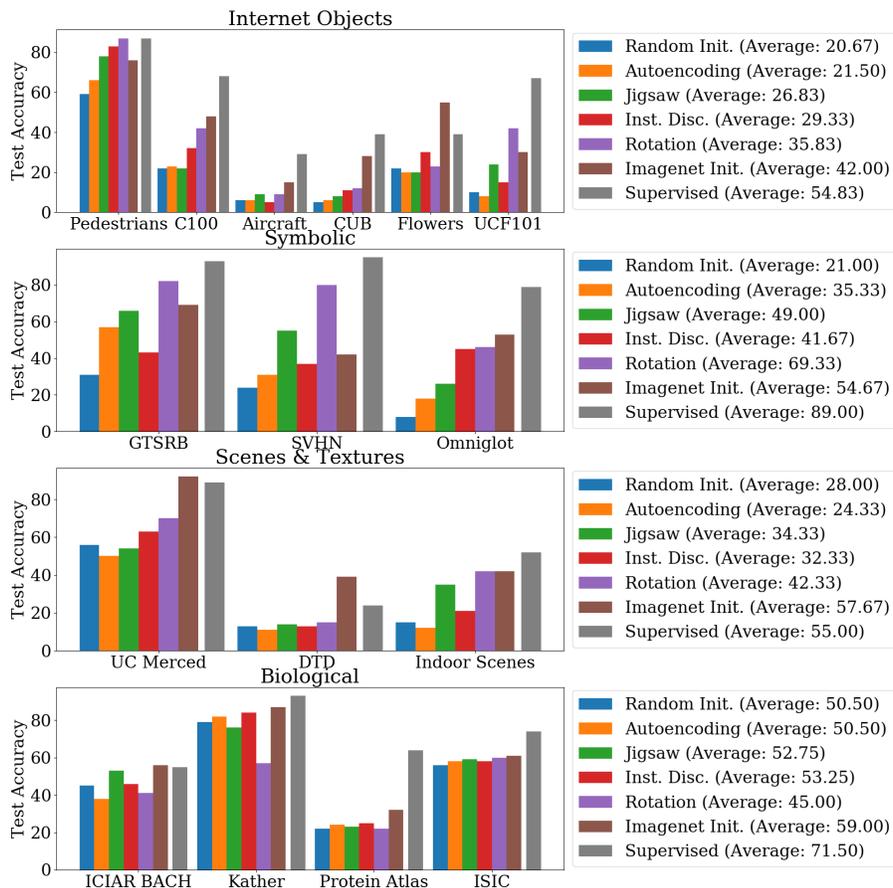
**Limited Label Training** Self-supervised learning has made an impact in the field of *semi*-supervised learning, where only a subset of the dataset given is labeled, as in the work of Zhai et al.[56]. In that work on ImageNet, a self-supervised feature extractor followed by a supervised linear head falls well short of the purely supervised baseline (40% accuracy compared to 80% when 10% of labels are available). We find that this conclusion does not hold on our domains when 10% of labels are used, with a pretext + linear setup outperforming the fully supervised models on some datasets/groups (Figure 4).<sup>1</sup>

On the Internet and Scenes & Textures groups, Rotation matches/outperforms full supervision. On ISIC, both Instance Discrimination and Jigsaw match Supervision. Interestingly, Autoencoding performs well in the Biological domains. Given the difficulty of expert annotation in the medical field, this is a valuable finding to encourage Autoencoder-based methods in these label-scarce domains, vindicating choices in past work[21,27,7,14]. The unlabeled 90% of the data being available for SSRL methods is critical: when only the labeled subset is used for pretraining the performance drops an average of approximately 10%.

### 5.2 Inspecting Failure Modes

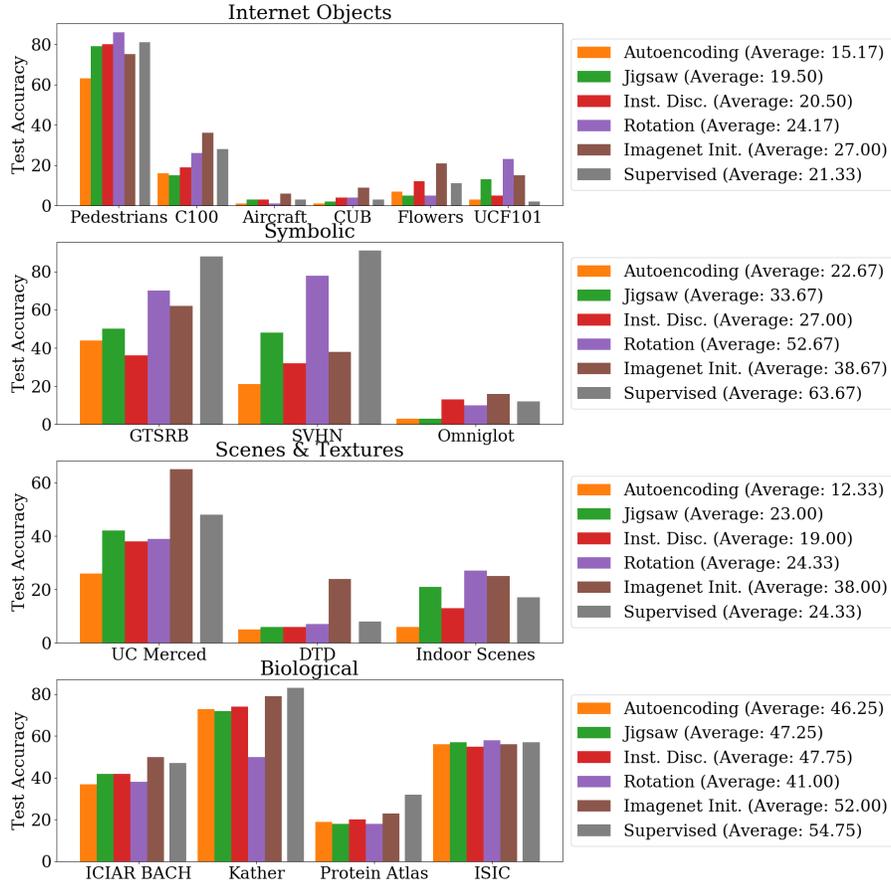
We call specific attention to the problems where self-supervised techniques do not achieve even *half* of the supervised accuracy (Aircraft, CUB, Textures and Protein Atlas). These seem to involve two kinds of problems.

<sup>1</sup> Note that [56] performs extensive hyperparameter tuning for the supervised baseline.



**Fig. 3.** Downstream classification accuracies for each pretext, as well as a randomly initialized feature extractor, a frozen ImageNet classifier, and fully supervised training. Rotation achieves the highest accuracies on the first three groups, but fails on the Biological domain, where Instance Discrimination performs best. The relative rankings of Jigsaw and ID vary between groups and are practically equal in overall average

**Textures and Protein Atlas:** In both these datasets, the entities being classified are not objects of recognizable shape, which is true of most *object recognition* datasets where the self-supervised techniques were developed. The images also do not have a canonical orientation, unlike in internet imagery where gravity and photographers’ biases provide such an orientation. We hypothesize that the lack of orientability hobbles Rotation, and the textural nature of both problems results in there being little to distinguish a patch in the Jigsaw pretext from a complete image, meaning Jigsaw can do little besides cue off of low-level properties (such as chromatic aberrations [34]) As such, modulo low-level dataset biases, the Rotation and Jigsaw pretext tasks *cannot even be solved in these domains*.



**Fig. 4.** Pretext feature extractors trained on the entire dataset, then a linear head trained on 10% of the labels. The fully supervised method simply trains on 10% of the dataset. We observe Rotation matching/outperforming Supervision for Internet and Scenes & Textures imagery. Autoencoding performs very well on the Biological domains in the semi-supervised setting, a novel result that indicates the potential for future development (given that this Autoencoder implementation is the simplest possible)

**Fine-Grained Classification on Aircraft and CUB:** In these datasets, the Rotation and Jigsaw pretexts are solvable, involving objects captured in a canonical orientation. Even so, neither pretext learns a good representation. While the *domains* are favorable to the pretexts, the *tasks* are not: both tasks involve fine-grained distinctions in color or texture, and subtle differences in shape. One hypothesis is that modeling these subtle differences is not necessary for solving the pretext task, causing the learnt representation to be *invariant* to these vital distinctions. Indeed, when we look at nearest neighbors in CUB for



**Fig. 5.** The 10 nearest neighbors (in order) of the far-left image in the CUB validation set measure by the unpooled feature space of the Rotation pretext model. We see that coloring plays almost no role in determining neighbors, but pose is a very strong signal

our Rotation model, we see birds of completely different colors but in the same pose, suggesting that the representation has captured pose, but ignored color (Figure 5). We quantitatively evaluate this hypothesis further in Section 5.3.

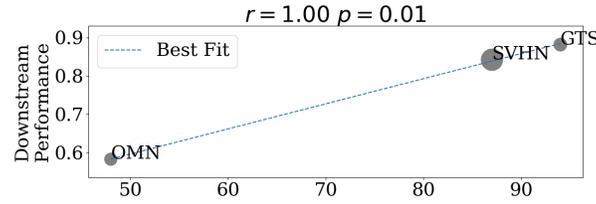
Note that we have not discussed the failures of Instance Discrimination in both domains. Instance Discrimination must learn to distinguish between individual images, which should be solvable in any domain as long as the images are distinct. As such, its failure and general performance is much more unpredictable and less interpretable, as we elaborate on in the next section.

### 5.3 Reasons for Success

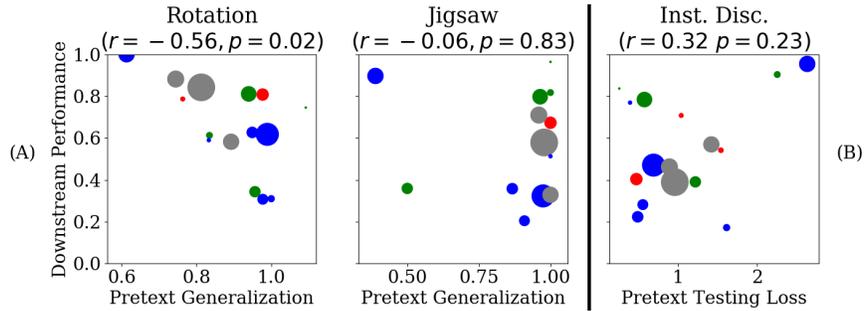
The flip side of why they fail is why they succeed. In prior work on the domain of internet imagery, this has mostly been answered by intuition. Rotation and Jigsaw were engineered as tasks that ostensibly require semantic understanding to solve. Instance Discrimination’s success is similarly intuitive: spreading points in a constrained space will naturally cluster similar images. Given the failures above, neither of these intuitions endure without modification in the varied domains of our work, and a more nuanced reasoning is necessary.

*Semantic understanding:* As discussed above, for some domains such as Aircraft or CUB, the pretexts do not produce a good semantic representation perhaps because they do not require a semantic understanding to solve. As an additional example, Jigsaw classifies permutations on Omniglot with 77% accuracy, but performs poorly on the classification task: line-matching is not understanding. Another example is Kather, where Rotation picks up on some hidden cue to fully solve the pretext problem without attaining any semantic knowledge. On the other side of the spectrum, we observe on the Symbolic domains in Figure 6 that Rotation *is* implicitly performing semantic classification, making it a near-optimal pretext task. *When* is semantic understanding a prerequisite for a pretext task, and how can we test this?

We propose that semantic understanding is a prerequisite for a pretext *when the solution method is class-dependent*: in this case, the network must necessarily implicitly classify the image to perform the pretext well. For example, in the fine-grained classification of airplanes and birds, determining the orientation of the object is class-*independent*, so a network trained on the Rotation pretext does not need to learn features indicative of class. One way of testing if the pretext solution is class independent is to see if a network trained to solve the



**Fig. 6.** Rotation pretext testing accuracy vs. unnormalized downstream performance. Almost perfect correlation is evident



**Fig. 7.** (A) Generalization of the pretexts to novel classes is plotted on the x-axis, calculated as  $\frac{PretextAcc_{Test}}{PretextAcc_{Val0.5}}$ , where Train/Val0.5 contain only half the classes of their usual counterparts. On the y-axis is downstream accuracy normalized by Supervision accuracy for the regularly trained model, e.g.  $\frac{ClassAcc_{Rot}}{ClassAcc_{Sup}}$ . We see a strong correlation for Rotation, but none for Jigsaw. (B) Instance Discrimination pretext loss vs. downstream classification accuracy, no significant correlations. Marker area indicates dataset size and color the domain grouping.

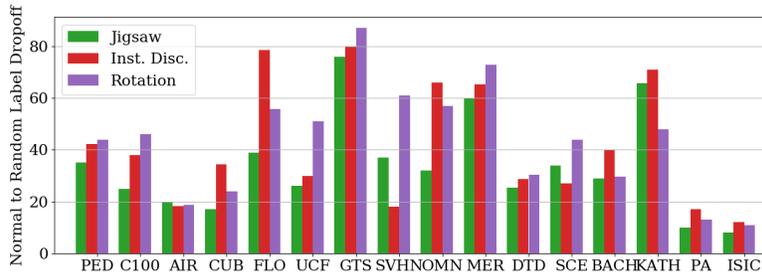
pretext on one set of classes solves the pretext on *unseen* classes. We call this test *pretext generalization* (Figure 7): we train and validate pretexts on only half of the available classes, and then evaluate the pretext on the entire test set. We predict that worse pretext generalization should imply better downstream performance.

Our prediction is correct for Rotation, affirming our hypothesis, but Jigsaw models seem to show high pretext generalization in almost all cases. This lack of correlation might be because Jigsaw is relying on low-level cues that do generalize *in addition to* the semantic information. We contrast these findings with a concurrent paper [47] which speculates that Rotation is not as useful for few-shot learning on FGVC-Aircraft or VGG-Flowers due to the relative difficulty of the pretext task.

For Instance Discrimination, counterintuitively, downstream accuracy is *not* correlated with pretext loss (Figure 7B). Our proposal is thus not the complete picture.

**Table 2.** Pearson correlations and p-values for training accuracy with random labels vs. normal. Jigsaw and ID have significant correlations, while Rotation trails substantially

	Jigsaw	Inst. Disc.	Rotation.
$(r, p)$	0.59 0.02	0.46 0.07	0.45 0.08

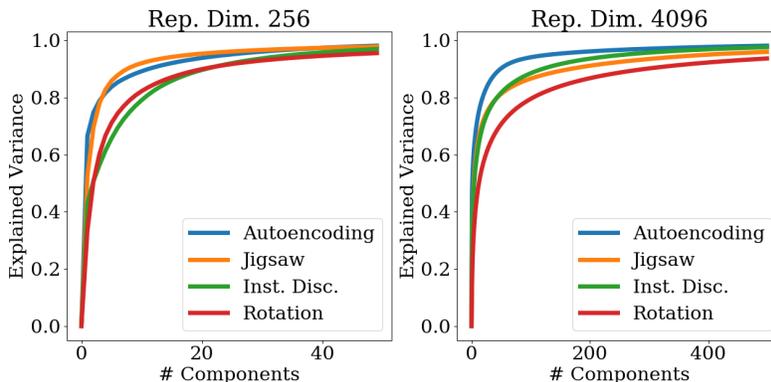


**Fig. 8.** Difference between normal and random label training accuracy. Note that even though Instance Discrimination had lower average *validation* accuracy than Rotation across all datasets, they actually had similar (normal label) average *training* accuracy at 56% (Jigsaw had 47%).

*Linear separability:* The above discussion assumes that the relationship between the pretext task and the downstream task is the key. But what if the downstream task is immaterial? Pretext tasks might be succeeding by simply enabling *all* tasks, by creating a feature space where *many* labelings of the data points are expressible as linear classifiers. To test this hypothesis, we retrain linear classifiers with randomly shuffled training labels and compare this (training) accuracy to the training accuracy with correct labels. We find that these quantities are more correlated for Jigsaw than Instance Discrimination or Rotation (Table 2), per task decreased are shown in Figure 8. This means that Jigsaw succeeds more by learning generic feature descriptors of images rather than by capturing semantics specific to the downstream task, while this is less true for Rotation or Instance Discrimination. Note that here we are talking of the training accuracy, or the empirical risk; this experiment does not reveal how or why pretext methods generalize to data not in the training set.

## 6 Feature Space Exploration

The discussion above suggests the virtue of analyzing the learnt representations independent of the downstream task. Below, we look at the *intrinsic dimensionality* of the learned representations, and the resulting notions of similarity.



**Fig. 9.** Plot of the average fraction of explained variance vs. number of principal components (validation sets); a higher number means greater explanation, and thus relatively *lower* implicit dimension. Bach is omitted due to its extremely small size

### 6.1 Implicit Dimensionality of the Representations

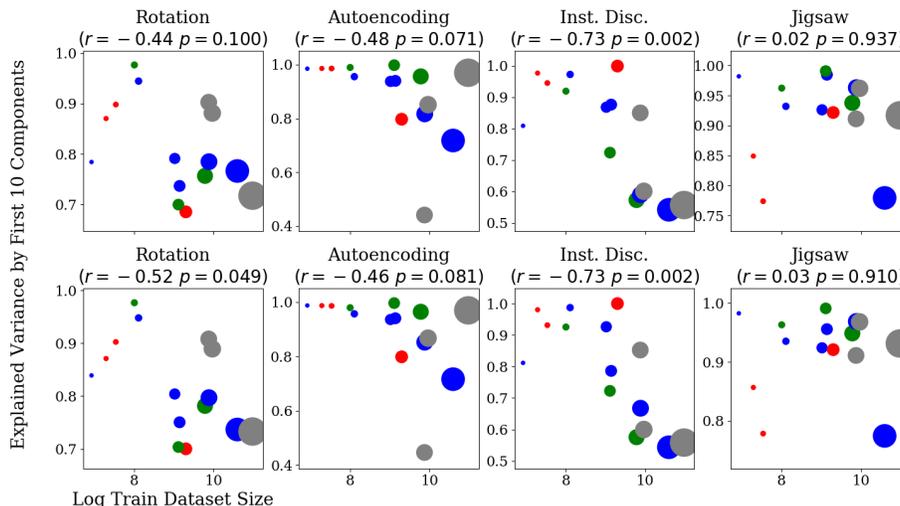
The dimensionality of the representation is largely regarded as a hyperparameter to be tuned. Kolesnikov et al. show that a larger representation is always more useful when operating on large datasets, but this comes with a tradeoff of memory/storage[23]. What has not been studied is the *implicit* dimensionality of the representations, such as via Principal Component Analysis (PCA). Intuitively, one would expect the 4-way Rotation task to produce compact representations, with Instance Discrimination using all available dimensions to spread out the points as the loss demands. We find that this is *not* the case, via performing PCA on the representations and summing the explained variance of the first  $n$  values averaged across all datasets (Figure 9). We also find that the implicit dimensionality induced by each pretext varies considerably across domains (Figure 10).

Surprisingly, Instance Discrimination and Rotation have extremely similar implicit dimensionality in the 256-dimensional case. Also note that Instance Discrimination clearly is not fully utilizing the available latent space, as the first 40 components explain over 95% of the variance. In the higher-dimensional example, Rotation has by far the largest implicit dimensionality.

We next investigate the effect of dataset size on implicit dimensionality. Intuition says that a larger dataset would demand a more expressive representation on every task, and we see in Figure 10 that this indeed holds true for *all* tasks except for Jigsaw. We hypothesize the lack of dimensionality increase for Jigsaw is because it exploits relatively low-level attributes instead of semantic knowledge.

### 6.2 Nearest Neighbors

As seen in Figure 5, the nearest neighbors in feature space can yield great insight into the inner workings of our models. We repeat this experiment for all of



**Fig. 10.** The fraction of variance explained by the first 10 components vs. the log of the training set size. Top row is training PCA, bottom row is validation PCA. Moderate to strong correlations exist for all pretext tasks besides Jigsaw. Marker area indicates dataset size and the colors are domain groupings.

our main methods on three datasets (CUB, Omniglot, Kather) in Figure 11 (additional examples in the supplementary).

**CUB:** The strong pose-capturing of Rotation is again observed. Jigsaw exhibits similar phenomenon slightly more weakly, and also favors similarly cluttered edge-heavy backgrounds as in the base image. Instance Discrimination appears to have reasonably matched the species in the first and tenth neighbor, but in between has many unrelated birds.

**Omniglot:** The poor performance of Autoencoding is reflected in its matching ability. Rotation exhibits strong retrieval, with Jigsaw offering similar but weaker results. Instance Discrimination is a similar case to CUB, where a discernable trend is hard to spot, but several correct characters are matched.

**Kather:** No method improved much on random initialization (79% accuracy). Reflectively, Autoencoding, Jigsaw, and ID all simply match reasonably visually similar results. Rotation, however, failed catastrophically (<60% accuracy), despite high pretext accuracy (95% testing). We hypothesized that the pretext was exploiting a spurious cue, which is affirmed by radically varying neighbors.

## 7 Conclusion

In this work, we have for the first time explored the notion of self-supervised learning in small-scale images in novel domains, identifying three domains where all current self-supervised methods have need of development: fine-grained, textural, and biological domains. In addition, we have revealed intriguing properties of



**Fig. 11.** Each block of images comes from a single dataset, with each row corresponding to the pretext method. The farthest left image is the base image from which the distance in feature space is calculated to all the others in the validation set. Left-to-right is increasing distance in feature space among the 10 nearest neighbors

the pretexts and the corresponding learnt representations, whose impact deserves further study. We hope that the release of our codes, models, and formatted dataset splits will help aid progress on all of these fronts.

**Acknowledgements:** This work was funded by a DARPA LwLL grant.

## References

1. Aresta, G., Araújo, T., Kwok, S., Chennamsetty, S.S., Safwan, M., Alex, V., Marami, B., Prastawa, M., Chan, M., Donovan, M., et al.: Bach: Grand challenge on breast

- cancer histology images. *Medical image analysis* **56**, 122–139 (2019)
2. Bengio, Y., et al.: Learning deep architectures for ai. *Foundations and trends® in Machine Learning* **2**(1), 1–127 (2009)
  3. Bilen, H., Fernando, B., Gavves, E., Vedaldi, A., Gould, S.: Dynamic image networks for action recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3034–3042 (2016)
  4. Bojanowski, P., Joulin, A.: Unsupervised learning by predicting noise. In: *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. pp. 517–526. *JMLR. org* (2017)
  5. Bromley, J., Guyon, I., LeCun, Y., Säckinger, E., Shah, R.: Signature verification using a” siamese” time delay neural network. In: *Advances in neural information processing systems*. pp. 737–744 (1994)
  6. Caron, M., Bojanowski, P., Joulin, A., Douze, M.: Deep clustering for unsupervised learning of visual features. In: *European Conference on Computer Vision* (2018)
  7. Chen, L., Bentley, P., Mori, K., Misawa, K., Fujiwara, M., Rueckert, D.: Self-supervised learning for medical image analysis using image context restoration. *Medical image analysis* **58**, 101539 (2019)
  8. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709* (2020)
  9. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., Vedaldi, A.: Describing textures in the wild. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3606–3613 (2014)
  10. Codella, N., Rotemberg, V., Tschandl, P., Celebi, M.E., Dusza, S., Gutman, D., Helba, B., Kalloo, A., Liopyris, K., Marchetti, M., et al.: Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic). *arXiv preprint arXiv:1902.03368* (2019)
  11. Doersch, C., Zisserman, A.: Multi-task self-supervised visual learning. In: *The IEEE International Conference on Computer Vision (ICCV) (Oct 2017)*
  12. Donahue, J., Krähenbühl, P., Darrell, T.: Adversarial feature learning. *arXiv preprint arXiv:1605.09782* (2016)
  13. Dumoulin, V., Belghazi, I., Poole, B., Mastropietro, O., Lamb, A., Arjovsky, M., Courville, A.: Adversarially learned inference. *arXiv preprint arXiv:1606.00704* (2016)
  14. Esser, P., Sutter, E., Ommer, B.: A variational u-net for conditional appearance and shape generation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 8857–8866 (2018)
  15. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: *International Conference on Learning Representations (2018)*, <https://openreview.net/forum?id=S1v4N210->
  16. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Advances in neural information processing systems*. pp. 2672–2680 (2014)
  17. Goroshin, R., Mathieu, M.F., LeCun, Y.: Learning to linearize under uncertainty. In: *Advances in Neural Information Processing Systems*. pp. 1234–1242 (2015)
  18. Goyal, P., Mahajan, D., Gupta, A., Misra, I.: Scaling and benchmarking self-supervised visual representation learning. In: *The IEEE International Conference on Computer Vision (ICCV) (October 2019)*
  19. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722* (2019)
  20. Hinton, G.E., Salakhutdinov, R.R.: Reducing the dimensionality of data with neural networks. *science* **313**(5786), 504–507 (2006)

21. Kallenberg, M., Petersen, K., Nielsen, M., Ng, A.Y., Diao, P., Igel, C., Vachon, C.M., Holland, K., Winkel, R.R., Karssemeijer, N., et al.: Unsupervised deep learning applied to breast density segmentation and mammographic risk scoring. *IEEE transactions on medical imaging* **35**(5), 1322–1331 (2016)
22. Kather, J.N., Weis, C.A., Bianconi, F., Melchers, S.M., Schad, L.R., Gaiser, T., Marx, A., Zöllner, F.G.: Multi-class texture analysis in colorectal cancer histology. *Scientific reports* **6**, 27988 (2016)
23. Kolesnikov, A., Zhai, X., Beyer, L.: Revisiting self-supervised visual representation learning. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2019)
24. Krizhevsky, A., et al.: Learning multiple layers of features from tiny images. *Tech. rep., Citeseer* (2009)
25. Lake, B.M., Salakhutdinov, R., Tenenbaum, J.B.: Human-level concept learning through probabilistic program induction. *Science* **350**(6266), 1332–1338 (2015)
26. Larsson, G., Maire, M., Shakhnarovich, G.: Learning representations for automatic colorization. In: *European Conference on Computer Vision*. pp. 577–593. Springer (2016)
27. Lu, A.X., Kraus, O.Z., Cooper, S., Moses, A.M.: Learning unsupervised feature representations for single cell microscopy images with paired cell inpainting. *PLoS computational biology* **15**(9), e1007348 (2019)
28. Maji, S., Rahtu, E., Kannala, J., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151* (2013)
29. Misra, I., van der Maaten, L.: Self-supervised learning of pretext-invariant representations. *arXiv preprint arXiv:1912.01991* (2019)
30. Misra, I., Zitnick, C.L., Hebert, M.: Shuffle and learn: unsupervised learning using temporal order verification. In: *European Conference on Computer Vision*. pp. 527–544. Springer (2016)
31. Munder, S., Gavrilu, D.M.: An experimental study on pedestrian classification. *IEEE transactions on pattern analysis and machine intelligence* **28**(11), 1863–1868 (2006)
32. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning (2011)
33. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*. pp. 722–729. IEEE (2008)
34. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: *ECCV* (2016)
35. Ouyang, W., Winsnes, C., Hjelmare, M., Cesnik, A., Åkesson, L., Xu, H., Sullivan, D., Dai, S., Lan, J., Jinmo, P., Galib, S.M., Henkel, C., Hwang, K., Poplavskiy, D., Tunguz, B., Wolfinger, R., Gu, Y., Li, C., Xie, J., Lundberg, E.: Analysis of the human protein atlas image classification competition. *Nature Methods* **16**, 1254–1261 (12 2019). <https://doi.org/10.1038/s41592-019-0658-6>
36. Owens, A., Wu, J., McDermott, J.H., Freeman, W.T., Torralba, A.: Ambient sound provides supervision for visual learning. In: *European conference on computer vision*. pp. 801–816. Springer (2016)
37. Pathak, D., Girshick, R., Dollár, P., Darrell, T., Hariharan, B.: Learning features by watching objects move. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2701–2710 (2017)
38. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2536–2544 (2016)

39. Quattoni, A., Torralba, A.: Recognizing indoor scenes. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. pp. 413–420. IEEE (2009)
40. Rebuffi, S.A., Bilen, H., Vedaldi, A.: Learning multiple visual domains with residual adapters. In: Advances in Neural Information Processing Systems (2017)
41. Rebuffi, S.A., Bilen, H., Vedaldi, A.: Efficient parametrization of multi-domain deep neural networks. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
42. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International journal of computer vision* **115**(3), 211–252 (2015)
43. de Sa, V.R.: Learning classification with unlabeled data. In: Advances in neural information processing systems. pp. 112–119 (1994)
44. Saha, S., Bandyopadhyay, S.: Unsupervised pixel classification in satellite imagery using a new multiobjective symmetry based clustering approach. In: TENCON 2008 - 2008 IEEE Region 10 Conference. pp. 1–6 (2008)
45. Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402* (2012)
46. Stallkamp, J., Schlipsing, M., Salmen, J., Igel, C.: Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition. *Neural networks* **32**, 323–332 (2012)
47. Su, J.C., Maji, S., Hariharan, B.: When does self-supervision improve few-shot learning? *arXiv preprint arXiv:1910.03560* (2019)
48. Thomee, B., Shamma, D.A., Friedland, G., Elizalde, B., Ni, K., Poland, D., Borth, D., Li, L.J.: Yfcc100m: The new data in multimedia research. *Communications of the ACM* **59**(2), 64–73 (2016)
49. Tschandl, P., Rosendahl, C., Kittler, H.: The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data* **5**, 180161 (2018)
50. Vincent, P., Larochelle, H., Bengio, Y., Manzagol, P.A.: Extracting and composing robust features with denoising autoencoders. In: Proceedings of the 25th international conference on Machine learning. pp. 1096–1103. ACM (2008)
51. Walker, J., Doersch, C., Gupta, A., Hebert, M.: An uncertain future: Forecasting from static images using variational autoencoders. In: European Conference on Computer Vision. pp. 835–851. Springer (2016)
52. Wang, X., Cai, Z., Gao, D., Vasconcelos, N.: Towards universal object detection by domain attention. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
53. Welinder, P., Branson, S., Mita, T., Wah, C., Schroff, F., Belongie, S., Perona, P.: Caltech-ucsd birds 200 (2010)
54. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3733–3742 (2018)
55. Yang, Y., Newsam, S.: Bag-of-visual-words and spatial extensions for land-use classification. In: Proceedings of the 18th SIGSPATIAL international conference on advances in geographic information systems. pp. 270–279. ACM (2010)
56. Zhai, X., Oliver, A., Kolesnikov, A., Beyer, L.: S4l: Self-supervised semi-supervised learning. In: Proceedings of the IEEE international conference on computer vision. pp. 1476–1485 (2019)
57. Zhai, X., Puigcerver, J., Kolesnikov, A., Ruysen, P., Riquelme, C., Lucic, M., Djolonga, J., Pinto, A.S., Neumann, M., Dosovitskiy, A., et al.: The visual task adaptation benchmark. *arXiv preprint arXiv:1910.04867* (2019)

58. Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: European conference on computer vision. pp. 649–666. Springer (2016)
59. Zhang, R., Isola, P., Efros, A.A.: Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1058–1067 (2017)