
What Does Rotation Prediction Tell Us about Classifier Accuracy under Varying Testing Environments?

Weijian Deng¹ Stephen Gould¹ Liang Zheng¹

Abstract

Understanding classifier decision under novel environments is central to the community, and a common practice is evaluating it on labeled test sets. However, in real-world testing, image annotations are difficult and expensive to obtain, especially when the test environment is changing. A natural question then arises: given a trained classifier, can we evaluate its accuracy on varying unlabeled test sets? In this work, we train semantic classification and rotation prediction in a multi-task way. On a series of datasets, we report an interesting finding, *i.e.*, the semantic classification accuracy exhibits a strong linear relationship with the accuracy of the rotation prediction task (Pearson’s Correlation $r > 0.88$). This finding allows us to utilize linear regression to estimate classifier performance from the accuracy of rotation prediction which can be obtained on the test set through the freely generated rotation labels.

1. Introduction

Knowing the generalization ability of a classifier is important for its applications in fields like autonomous driving, medical diagnosis and product recommendation. Many well-established benchmarks in the community, *e.g.*, ImageNet (Deng et al., 2009) and CIFAR (Krizhevsky et al., 2009), provide held-out labeled test sets that allow empirical results to be computed. However, when applying the classifier in real-world scenarios, we often encounter novel distributions and unknown cases, and empirical results calculated from existing benchmarks do not accurately reflect its generalization ability during deployment. To evaluate classifier performance under novel environments, the standard way is to collect new labeled test sets. However, it is prohibitively expensive to label sufficient images in every novel scenario.

This work investigates an interesting problem: how to estimate classifier accuracy on unlabeled test sets? Some existing research studies classifier generalization on *fixed* test sets by developing complexity measurements on model parameters (Arora et al., 2018; Corneanu et al., 2020; Jiang et al., 2018). In this work, we focus on *varying test domains*. In the absence of test labels, we aim to find an unsupervised evaluation criterion that can reflect classifier accuracy. A self-supervised task produces ground truths from data without human annotations. Thus, we can always calculate its accuracy on an unlabeled test set. As such we raise a question: given a network that is good at a self-supervised task on an unlabeled test set, can it also address object recognition task well under the same environment? If such a relationship exists, we can infer classifier performance on an unlabeled test set from the accuracy of the self-supervised task. In light of this question, we study the statistical correlation between semantic classification and rotation prediction under varying environments.

In the community, self-supervised tasks are usually used to learn effective representations for downstream tasks (He et al., 2020; Noroozi & Favaro, 2016; Gidaris et al., 2018). Different from them, we explore the feasibility of using self-supervision to predict classifier accuracy on novel test sets. We pay particular attention to the pretext task of *rotation prediction*, where we classify the rotation angle of an image into $\{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$. Specifically, we train a network jointly for the main classification task and pretext task. To study the accuracy correlation between the two tasks, we construct different setups, including CIFAR-10, MNIST, Tiny-ImageNet, and COCO. Each setup contains many labeled test sets for calculating network performance on the two tasks. By plotting classification accuracy against rotation prediction accuracy, we consistently observe a strong correlation between them (Pearson’s Correlation $r > 0.88$) across the four setups. This shows that if the network can solve the rotation prediction task well, it is likely to achieve good classification accuracy, and vice versa.

Based on this, we learn a linear regression model to predict classification accuracy on unseen test sets. Given an unlabeled test set, we obtain rotation ground truths by manually rotating the images. Then, we use the multi-task network

¹College of Engineering and Computer Science, Australian National University, Canberra, ACT 0200, Australia. Correspondence to: Liang Zheng <liang.zheng@anu.edu.au>.

to predict rotation degrees of the test images and calculate rotation prediction accuracy. Finally, the linear regression model predicts the semantic classification performance from rotation prediction accuracy. Empirical results on three different setups show that our regression model achieves very promising estimation results on unseen test sets.

In summary, this work investigates the problem of unsupervised classifier prediction, *i.e.*, estimating classifier performance on unseen test sets without any human-annotated label. We observe that the accuracy of a simple self-supervised task, rotation prediction (Gidaris et al., 2018), is highly correlated with the classification accuracy under various testing environments.

2. Related Work

Model generalization. Estimating model generalization error on unseen images is a crucial problem. Some works develop *complexity measurements* on trained models and training sets to predict the generalization gap between a certain pair of training-testing set (Eilertsen et al., 2020; Unterthiner et al., 2020; Arora et al., 2018; Corneanu et al., 2020; Jiang et al., 2018; Neyshabur et al., 2017; Jiang et al., 2019). The authors of (Corneanu et al., 2020) use persistent topology measures to predict the performance gap between training and testing error. Jiang *et al.* (Jiang et al., 2018) introduce a measurement of layer-wise margin distributions for generalization ability. Neyshabur *et al.* (Neyshabur et al., 2017) develop bounds on the generalization gap based on the product of norms of the weights across layers. Our work differs significantly. Instead of studying the statistics of model weights, we use rotation prediction accuracy to characterize models. Moreover, we are concerned with various test distributions that have domain gaps with the training one, while the aforementioned works typically assume a fixed test set without a domain gap.

In addition, some methods try to predict accuracy on unlabeled test samples based on the agreement score of several classifiers’ predictions (Madani et al., 2004; Platanios et al., 2016; 2017; Donmez et al., 2010). Platanios *et al.* (Platanios et al., 2017) use a probabilistic soft logic (PSL) model to infer the errors of the classifiers. Recent work also considers the classifier accuracy estimation and proposes to predict accuracy based on the distribution difference (Deng & Zheng, 2021). In this work, we attempt to estimate classifier accuracy on unlabeled test sets by leveraging self-supervision.

Self-Supervised learning. Self-supervised learning aims to learn good representations on unlabeled data for the downstream tasks (He et al., 2020; Noroozi & Favaro, 2016; Gidaris et al., 2018; Caron et al., 2018; Bojanowski & Joulin, 2017; Kolesnikov et al., 2019; Doersch & Zisserman, 2017; Lee et al., 2019; Ren & Jae Lee, 2018). To achieve this,

several pretext tasks are proposed, such as solving jigsaw puzzles (Noroozi & Favaro, 2016), colorization (Larsson et al., 2017), rotation prediction (Gidaris et al., 2018), denoising auto-encoders (Vincent et al., 2008), tracking (Wang & Gupta, 2015), and clustering (Caron et al., 2019; 2018). While these works use self-supervision for learning unsupervised representations, this work investigates its utility for unsupervised accuracy estimation.

Reed *et al.* (Metzger et al., 2020) use rotation prediction is used to automatically select augmentation policies for unsupervised contrastive learning method (*e.g.*, MoCo (He et al., 2020)). Liu (Liu et al., 2020) leverage self-supervised pretext tasks for unsupervised neural architecture search. Self-supervised methods can also be used to improve supervised classifier learning (Hendrycks et al., 2019; Sun et al., 2020; Carlucci et al., 2019; Sun et al., 2019). Hendrycks *et al.* (Hendrycks et al., 2019) show jointly training classification and self-supervised task improves classifier robustness and uncertainty estimation. Sun *et al.* (Sun et al., 2020) achieve test-time training by using self-supervised task on unlabeled test images. Moreover, the supervised tasks can be used to alleviate distribution shifts (Carlucci et al., 2019; Sun et al., 2019; Bucci et al., 2021). This work does not aim to improve supervised classification by using self-supervised tasks. Instead, we study the feasibility of using the rotation prediction task for estimating classifier accuracy.

Out-of-distribution detection. The overall goal of out-of-distribution is to detect test samples from a novel distribution (a distribution different from the training distribution) (Hsu et al., 2020; Hendrycks & Gimpel, 2017; Lee et al., 2018; Liang et al., 2018; Guo et al., 2017; Ovadia et al., 2019; Ren et al., 2019). Hendrycks *et al.* (Hendrycks & Gimpel, 2017) use probability outputs from a softmax classifier as an indicator to find out-of-distribution samples. Liang *et al.* (Liang et al., 2018) use temperature scaling (Guo et al., 2017) and input pre-processing to improve detection accuracy. Ren *et al.* (Ren et al., 2019) investigate deep generative model-based methods for the sample detection. While the above methods aim to capture abnormal samples during the testing, our work focuses on predicting overall accuracy that involves all test samples.

3. Classifier Accuracy Estimation

3.1. Problem Definition

In the problem of unsupervised accuracy estimation, we are provided with a labeled training set and a set of unlabeled datasets for testing. We denote the training set as $\mathcal{D}_{train} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$, where \mathbf{x}_i is the i -th training image, $\mathbf{y}_i \in \{0, 1, \dots, K - 1\}$ is its label, and n is the total number of images. Similarly, we denote each test dataset as $\mathcal{D}_{test}^t = \{\mathbf{x}_j^t\}_{j=1}^{n_t}$, where $t \in \{1, 2, \dots, M\}$ indicates the t -th test




Figure 1. Our multi-task network structure for studying the statistical correlation between the accuracy of rotation prediction and classification under different test sets. The network contains a shared backbone (e.g., ResNet-18) for the two tasks, a fully-connected layer for 4-way rotation prediction, and another fully-connected layer for K -way classification (K is the number of classes). The shared backbone receives losses from both tasks.

set, \mathbf{x}_j^t is the j -th test image, and n_t is the total number of images in \mathcal{D}_{test}^t . Given a classifier trained on \mathcal{D}_{train} , our goal is to estimate its accuracy on every unlabeled test set \mathcal{D}_{test}^t , $t = 1, 2, \dots, M$.

3.2. Correlation Study

3.2.1. NETWORK DESCRIPTION

We learn a multi-task neural network for semantic classification along with an auxiliary self-supervised task. The self-supervised task *should* 1) introduce no learning complexity for the main classification, 2) require minimal network structure change, and 3) not degrade classification accuracy. We choose *rotation prediction* (Gidaris et al., 2018) as the auxiliary task because it meets the above requirements. In the experiment, we observe using Jigsaw also has a strong correlation, but it decreases the classification task accuracy.

Rotation prediction is an effective pretext task initially designed for unsupervised representation learning (Gidaris et al., 2018). It is shown that predicting rotation degrees helps learn effective features for various downstream tasks, such as recognition, detection and segmentation.

Following the practice in (Gidaris et al., 2018), we define the set of rotation transformations $G = \{R_r(\mathbf{x})\}$, where R_r is rotation transformation and $r \in \{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$. Given an image, the goal of rotation prediction is to tell which one of the four rotations this image undergoes; it can be formulated as a 4-way classification task.

Network structure. We use a convolutional neural network (ConvNet), e.g., ResNet-50. We only add one fully-connected (FC) layer for rotation prediction and do not modify any other component. The overall structure is very simple, shown in Fig. 1. The ConvNet has three parts: a shared backbone parameterized by θ_s , a FC layer for clas-

sification parameterized by θ_c , and another FC layer for rotation prediction parameterized by θ_r .

Loss function. Our network is a kind of multi-task model that jointly learns the self-supervised pretext task and semantic classification task. Specifically, the model is trained to simultaneously 1) estimate the geometric transformation applied to an image and 2) classify the semantic category of this image. For rotation prediction, its loss function is written as,

$$\mathcal{L}_{\text{rot}} = \frac{1}{4} \left[\sum_{r \in \{0^\circ, 90^\circ, 180^\circ, 270^\circ\}} \mathcal{L}_{\text{CE}}(\mathbf{y}_r, \theta_r(\theta_s(R_r(\mathbf{x}))) \right], \quad (1)$$

where \mathbf{y}_r is the one-hot label of $r \in \{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$.

For semantic classification, the loss function is:

$$\mathcal{L}_{\text{cls}} = \mathcal{L}_{\text{CE}}(\mathbf{y}_c, \theta_c(\theta_s(\mathbf{x}))), \quad (2)$$

where \mathbf{y}_c is the one-hot class label of image \mathbf{x} . The overall loss function of our multi-task network is: $\mathcal{L} = \mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{rot}}$.

Network inference. After training is complete, we test the multi-task network on unlabeled test images. For the classification task, the network predicts the semantic class for each test image. For rotation prediction task, we apply *all* four rotation transformations G to each test image. The network predicts rotations of all transformed images, from which the accuracy of rotation prediction is calculated.

3.2.2. DATASET SYNTHESIS

With the network introduced above, we now study the accuracy relationship between rotation prediction and semantic classification under various testing environments.

To this end, we need many test sets that 1) have various distributions, 2) contain the same classes with the training set, and 3) have sufficient samples. However, there are very few datasets meeting these requirements. Following (Deng & Zheng, 2021), we synthesize test sets by using various geometric and photometric transforms on top of MNIST, CIFAR-10 (Krizhevsky et al., 2009) and COCO (Lin et al., 2014), creating the corresponding setups.

MNIST setup. MNIST contains 50K training image and 10K test image from 10 categories. We use MNIST training set to train a ConvNet based on LeNet-5. To generate various test sets, we change the backgrounds of MNIST test set (named *original set*). Specifically, we use COCO training images as backgrounds and synthesize 1,000 test sets. For all the setups in this section, because the transformations do not change the main image content, the synthesized test sets *inherit the semantic labels* from the original set.

CIFAR-10 setup. We use DenseNet-40-12 (40 layers with growth rate 12) (Huang et al., 2017) as the backbone. We

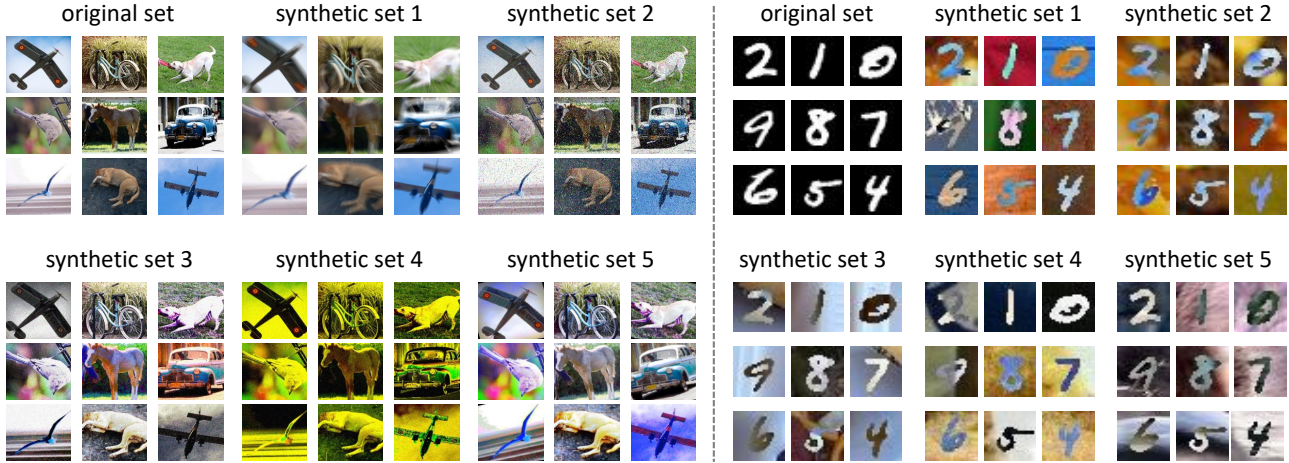


Figure 2. Examples of synthetic datasets used in the correlation study. **Left:** datasets under the COCO setup. Given an original set (labeled, top left), we use geometric and photometric transforms to synthesize more datasets. For example, we adopt corruption methods *zoom blur* and *impulse noise* to generate “synthetic set 1” and “synthetic set 2”, respectively. We combine several image transformations (e.g., sharpness, ColorTemperature) to produce other synthetic sets. **Right:** datasets under the MNIST setup. We change the image background of the original set to synthesize various datasets. The synthetic datasets exhibit distinct visual characteristics, and yet inherit the foreground objects from images in the original set, and thus are fully labeled.

train our network on CIFAR-10 training set (50K images), and test it on CIFAR-10-C (Hendrycks & Dietterich, 2019). CIFAR-10-C was proposed to study the model robustness to common corruptions. It contains 15 types of corruption and each corruption type comes at 5 levels of severity, resulting in 75 corruption datasets. These corruptions include noise, blur, weather, and digital categories. We refer readers to (Hendrycks & Dietterich, 2019) for more details. To synthesize test sets, we apply image transformations. For each dataset, we randomly select three transformations from {Sharpness, Equalize, ColorTemperature, Solarize, Autocontrast, Brightness, Rotate}. Then, we apply three transformations with random magnitudes on images. This practice generates another 425 datasets.

COCO setup. We use 12 classes, *i.e.*, aeroplane, bike, bird, boat, bottle, bus, car, dog, horse, monitor, motorbike, and person, following (Peng et al., 2017). We build the training and test sets for the classification task based on COCO training and validation sets, respectively. Each class has 600 images. We also use the 15 common corruptions on the test set and synthesize 75 corrupted test sets. Moreover, we randomly use three transformations to generate another 425 datasets. Examples of synthetic sets are shown in Fig. 2.

3.2.3. EMPIRICAL ANALYSIS

For each setup introduced above, we compute the accuracy numbers of semantic classification and 4-way rotation prediction on every synthetic test set. We empirically study the statistical relationship between the two tasks’ performance

using Pearson’s Correlation Coefficient (r).

Linear correlation. In Fig. 3, we plot semantic classification accuracy against rotation prediction accuracy on these synthetic test sets. We have an interesting observation: *given a multi-task model, when the testing environment undergoes changes (i.e., testing on different synthesized test sets), the performance of semantic classification has a linear relationship with that of rotation prediction.* This observation is consistent across three different setups. For example, on the CIFAR-10 and COCO, we find the Pearson’s Correlation Coefficient (r) is 0.977 and 0.883, respectively, indicating strong linearity. In other words, if the multi-task network is good at predicting rotations, it is most likely to achieve good object recognition accuracy under the same environment, and vice versa.

3.3. Estimation with Linear Regression

Based on the strong linear relationship between classification and rotation prediction under changing environments, we train a linear regression model to predict classification accuracy. Specifically, we use N synthetic sets as training samples. The i -th ($i = 1, 2, \dots, N$) synthetic set is denoted as (a_i^{rot}, a_i^{cls}) , where a_i^{rot} and a_i^{cls} denote the rotation prediction accuracy and semantic classification accuracy on this set, respectively. The regression model is:

$$a^{cls} = w_1 a^{rot} + w_0, \quad (3)$$

where $w_1, w_0 \in \mathbb{R}$ are linear regression parameters. In practice, we use robust linear regression (Huber, 2004) with




Figure 3. Correlation between classifier accuracy and rotation prediction accuracy on three setups. From left to right: CIFAR-10 with DenseNet-40-12, MNIST with LeNet-5, and COCO with ResNet-50, respectively. Each point in the figure represents a synthetic test set. Across the three setups, we consistently observe a strong linear relationship (Pearson Correlation $r > 0.88$) between the two accuracy numbers. The straight lines are calculated by robust linear regression (Huber, 2004).

Huber loss to learn the regression model. In Section 4, we will present results of using the learned linear regression models for accuracy prediction on unseen test sets.

4. Experiment

4.1. Task Setting and Baseline

Task setting. For the problem of unsupervised classifier evaluation, we experiment on two semantic classification settings, *i.e.*, digit classification and natural image classification. In each setting, we estimate the semantic classification accuracy on *unseen test sets*.

For digit classification, we test LeNet-5 (trained on MNIST) on two unseen real datasets, *i.e.*, USPS (Hull, 1994) and SVHN (Netzer et al., 2011), both with 10 classes. For natural image classification, we use ResNet-50 trained on COCO and estimate classification accuracy on three datasets: PASCAL (Everingham et al., 2007), Caltech (Griffin et al., 2007) and ImageNet (Deng et al., 2009), all with 12 classes. In addition, we use CIFAR-10.1 (Recht et al., 2018) for CIFAR-10 setup: we estimate semantic classification accuracy of DenseNet-40-12 on this unseen test set through ConvNet trained on CIFAR-10.

Metric. To evaluate the correctness of performance estimation, we use root mean squared error (RMSE). RMSE measures the average squared distance between the estimated accuracy and ground-truth accuracy. Small RMSE corresponds to good estimation and vice versa.

Baselines. There are very few studies on the unsupervised accuracy estimation problem. Thus, we introduce two intuitive methods for comparison, which are inspired by the uncertainty estimation methods (Hendrycks & Gimpel, 2017; Kendall & Gal, 2017; Liang et al., 2018).

Prediction score based. Given an input image, its score is

obtained as the maximum softmax output. If the score is greater than a threshold $\tau_1 \in [0, 1]$, this image is regarded as being correctly classified.

Entropy score based. The score is given by the entropy of softmax outputs, normalized by $\log(K)$, where K is the number of classes. Predictions with scores less than a threshold $\tau_2 \in [0, 1]$ are considered correct.

4.2. Experimental Results

In Table 1, we report the performance of our linear regression models on the six real-world test sets described above. We have the following observations.

Score-based methods are sensitive to thresholds. As shown in Table 1, the two intuitive score-based methods could produce good prediction results when using proper thresholds. However, they are sensitive to the threshold value. For example, using $\tau_1 = 0.9$ and $\tau_1 = 0.8$ will lead to a 5.44% difference in RMSE under the COCO setup. Furthermore, different datasets need different thresholds. For instance, for the prediction score based method, $\tau_1 = 0.9$ is suitable at CIFAR-10.1, but is not good for three natural image classification datasets. Notably, it is hard to select the optimal threshold because 1) test labels are unavailable and 2) the testing environment keeps changing. That said, designing a strategy to select thresholds might be a potential way to improve the practicability of two intuitive methods.

Our linear regression model is feasible to predict classifier performance. The linear regression model (Section 3.3) based on rotation prediction accuracy could achieve promising predictions across three setups, shown in Table 1. For example, when predicting the accuracy of a classifier on three test sets under the COCO setup, our regression model has a 2.68% RMSE. In addition, it estimates the accuracy to be 91.89% on CIFAR-10.1 (the ground truth is 88.15%).

Table 1. Results of classifier accuracy prediction using various methods on three groups of unseen test sets: 1) SVHN and USPS; 2) CIFAR-10.1; 3) Pascal, Caltech and ImageNet. The training set for each group is MNIST, CIFAR-10, and COCO, respectively. For each test set, we report the estimated classification accuracy (%) and the ground-truth recognition accuracy (%). RMSE (%) is reported to measure the estimation precision. For “prediction” and “entropy” methods, images with high prediction scores (greater than τ_1) or low entropy scores (less than τ_2) are regarded as being classified correctly, respectively.

Train Set	MNIST			CIFAR-10		COCO			
Unseen Test Set	SVHN	USPS	RMSE↓	CIFAR-10.1	RMSE↓	Caltech	Pascal	ImageNet	RMSE↓
Ground-truth Accuracy	23.06	65.52	-	88.15	-	92.61	86.43	87.83	-
Prediction ($\tau_1 = 0.8$)	33.64	44.34	16.74	91.15	3.00	89.36	83.98	85.17	2.81
Prediction ($\tau_1 = 0.9$)	22.07	30.39	24.85	86.85	1.30	84.30	78.00	79.83	8.25
Entropy ($\tau_2 = 0.2$)	26.63	33.23	22.97	89.20	1.05	86.80	80.14	82.50	5.82
Entropy ($\tau_2 = 0.3$)	40.35	46.87	17.98	93.80	5.65	92.49	86.21	88.50	0.41
Linear Regression	24.84	53.10	8.87	91.89	3.74	90.70	89.29	90.98	2.68




Figure 4. Impact of the number of training samples (synthetic sets) for linear regression. **Left:** RMSE (%) on CIFAR10.1. **Right:** RMSE (%) on Caltech, Pascal, and ImageNet. We observe that the regression model becomes stable when using more than 200 training samples.

Since the training samples (synthetic datasets) of the regression model are diverse enough (e.g., diverse distributions and various visual characteristics), it can learn to achieve reasonably accurate estimations on unseen test sets.

Impact of the number of training samples (synthetic sets). In Fig. 4, we present the RMSE (%) obtained by linear regressors trained using different numbers of synthetic sets. Since the linear regression model only has two parameters, it does not require too many training samples (Huber, 2004). In our problem, when using more than 200 synthetic sets, RMSE is relatively stable. In practice, we speculate that using sufficient and diverse synthetic sets is necessary to improve regression robustness.

Robustness of linear regressor on transformed test sets. We apply various image transformations to edit three natural image datasets (ImageNet, Pascal and Caltech), and test the regression model on them. Specifically, we use *Cutout* and *Posterize* to produce ImageNet-A, Pascal-A and Caltech-A, and use *Pepper* and *FilterSmooth* to produce ImageNet-B, Pascal-B and Caltech-B. Visual examples of these datasets are shown in Fig. 5. Note that these transformations were




Figure 5. Visual examples of transformed natural ImageNet images. For ImageNet-A, we use *Cutout* and *Posterize* to edit original images. For ImageNet-B, we use *Pepper* and *FilterSmooth* to introduce visual changes.

not used when synthesizing datasets to train the regression model. Thus, testing on them allows us to understand the robustness of the regression model.

We have the following observations in Fig. 6. First, the transformations introduce extra visual changes, so the classifier accuracy decreases on the edited dataset. Second, the regression model still makes good estimated results on the edited test sets. For example, it makes good predictions on Caltech-A and Caltech-B, as the absolute errors are 0.01% and 2.95%, respectively. Although the absolute error on Pascal-A is 6.35%, the estimated accuracy still helps us gain a general understanding of classifier behaviour on this dataset. We think the regression precision can be further improved by using other information, such as distribution shift (Deng & Zheng, 2021). Moreover, if we are provided with a small number of labeled test images, it could be helpful for adjusting regression predictions.




Figure 6. Linear regression on transformed test sets (ImageNet, Pascal, and Caltech). The transformed datasets are denoted by “-A” and “-B”. The absolute error (%) between estimated classifier accuracy and the ground truth accuracy (%) (shown below each dataset) is reported. (-) / (+) denotes the predicted accuracy is lower and higher than the ground-truth accuracy, respectively.




Figure 7. The absolute error (%) of predictions on *four natural* test sets under COCO setup. The ground truth (%) is shown under each dataset. (-) / (+) means the predicted accuracy is lower / higher than the ground-truth accuracy, respectively.

Generalization ability of the regression model. To further validate this, we test on *four natural test sets* with different types of domain shifts (Taori et al., 2020) and report estimations in Fig. 7. We observe that our method overall gains good estimations, and its the average RMSE of the four datasets is 3.07%.

Based on the above experimental analysis, we think self-supervised rotation prediction provides a feasible way to predict classifier accuracy on varying testing environments without annotated category labels.

5. Discussion

When trained in a multi-task ConvNet, does rotation prediction affect semantic classification accuracy? To answer this question, we conduct experiments to study whether classification is compromised by the introduction of the rotation prediction head. In Table 2, we use three types

Table 2. Impact of the auxiliary rotation prediction (Rot.) task on the main classification task. We conduct experiment on CIFAR-10 with ResNet26, Dense40-12, and VGG19. Top-1 accuracy (%) is reported. We show that the rotation prediction task *does not* affect the classifier accuracy.

Train w/ Rot.?	CIFAR-10		
	ResNet26	Dense40	VGG19
	92.86±0.10	94.65±0.15	92.31±0.13
✓	92.84±0.10	94.75±0.15	92.51±0.20

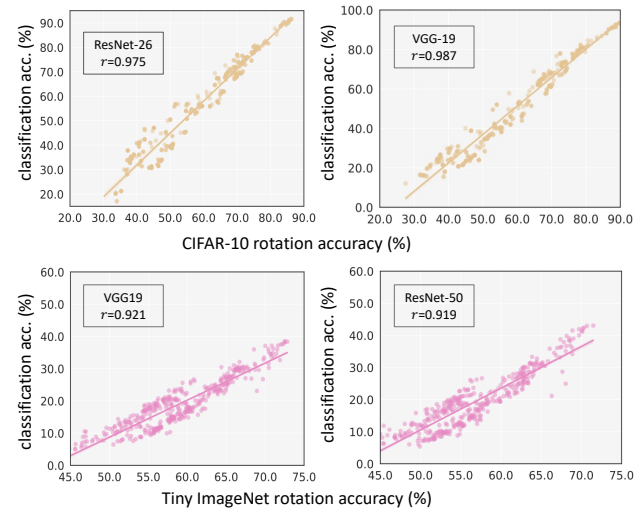


Figure 8. Correlation between semantic classification accuracy and rotation prediction accuracy when using difference types of network structures. On CIFAR-10, we use ResNet-26 and VGG-19. The Pearson Correlation r is 0.975 and 0.987 with ResNet-26 and VGGNet-19, respectively. On Tiny ImageNet, the correlation r is 0.921 and 0.919 using VGG-19 and ResNet-50, respectively.

of networks (ResNet-26, DenseNet-40-12, and VGG-19). We observe that auxiliary rotation prediction task *does not* degrade main classification on the CIFAR-10 test set.

In fact, rotation prediction requires the neural network to capture the rotation-related representations of an image (e.g., particular structures and salient object-level patterns), which are also useful for objection recognition (Gidaris et al., 2018). Moreover, while rotation prediction does not improve accuracy on original or clean testing image, it is reported to improve classifier robustness to input corruptions (Hendrycks et al., 2019). In addition, recent works (Sun et al., 2019; Bucci et al., 2021) have shown that using rotation prediction for training can improve the domain adaptation ability of the classifier. Based on the above discussion, introducing auxiliary rotation prediction task *will not* negatively impact the main classification task.

Table 3. Linear correlation (r) on CIFAR-10 with different backbone structures, including VGG11, ResNet26, VGG19, ResNet44, and Dense40-12. We also include the accuracy (%) of classification (Class. Acc) and rotation prediction (Rot. Acc) on the CIFAR-10 test set. Using different backbones, the two task performance is varying, but their correlation remains at the same level: the “mean \pm std” of correlation (r) is 0.983 ± 0.005 .

	VGG11	VGG19	ResNet26	ResNet44	Dense40
Class. Acc.	92.53	92.51	92.84	93.73	94.75
Rot. Acc.	91.32	92.07	87.84	88.81	91.28
Cor. (r)	0.990	0.987	0.975	0.981	0.981

Does the linear relationship hold when using other backbone structures? To study this, we use different types of backbones for correlation study and report results in Fig. 8. We also include Tiny ImageNet for the correlation study. It has 200 classes, and each class has 500 training images. We follow the COCO setup to synthesize 500 datasets.

For CIFAR-10, when using ResNet-26 and VGGNet-19 as backbones, we observe a strong linear correlation between the two task performance. Specifically, Pearson’s Correlation (r) is 0.975 and 0.987 with ResNet-26 and VGGNet-19, respectively. When using VGG-19 and ResNet-50 on Tiny ImageNet, we also observe the correlation (r) is 0.921 and 0.919, respectively. It also indicates the strong linear correlation between two task accuracy on varying test sets is maintained when using different backbones.

In addition, we show the correlation values on CIFAR-10 setup using different backbones in Table 3. In the experiment, we use five backbones: VGG11, VGG19, ResNet26, ResNet44, and Dense40-12. We have two observations. First, the correlation values are very high ($r > 0.975$) when using different backbones. Another interesting finding is: the performance of two tasks is varying with different backbones, but their correlation on different test sets remains at the same level: the “mean \pm std” of Pearson’s Correlation (r) is 0.983 ± 0.005 .

Impact of the number of categories on correlation. For rotation prediction, it distinguishes between four different classes, but for image classification, the number of categories can be much greater. In Fig. 3 and Fig. 8, we observe the correlation is high ($r > 0.88$) across different datasets with various numbers of categories. For example, the correlation is $r > 0.90$ on Tiny-ImageNet with 200 categories.

We further study on CIFAR-100 setup and report results in Table 4. We follow the CIFAR-10 setup to synthesize 500 datasets for CIFAR-100 setup. Although the correlation r is lower than that on CIFAR-10 (0.918 vs. 0.975 with ResNet26), the value $r = 0.918$ still indicates strong linear correlation under the CIFAR-100 setup. Based on the above

Table 4. Effect of the number of semantic categories on the correlation. Although the rank correlation (ρ) on CIFAR-100 is slightly lower than that on CIFAR-10, it still indicates that the two tasks have a strong linear correlation ($\rho > 0.9$). The accuracy (%) of classification (Class. Acc) and rotation prediction (Rot. Acc) on CIFAR-100 test set is also reported.

Backbone	CIFAR-10	CIFAR-100	
	Cor. (r)	Cor. (r)	Class Acc. Rot. Acc.
ResNet26	0.975	0.918	69.31 73.18
ResNet44	0.981	0.910	71.38 75.60
Dense40	0.981	0.950	74.55 75.20

analysis, we speculate when the number of categories is huge (e.g., 10K (Deng et al., 2010)), the correlation r might decrease but it will still have a high value.

Why is rotation prediction performance closely correlated with classification performance? Our multi-task network jointly learns both tasks. By doing so, they are coupled together and share the feature representations. When testing on images from the novel distribution, the dataset shift will make the presentations not suitable for the test set. Thus, the network performance on both tasks drops simultaneously. That is, different test sets introduce different dataset shifts, leading to varying degrees of performance drop. This might be the reason why their performance shows a positive correlation trend on varying test domains.

Difference from those using self-supervision for classification (Sun et al., 2020; Hendrycks et al., 2019; Carlucci et al., 2019; Zhai et al., 2019). These works present that the *self-supervised* task can help *supervised* classification task in different perspectives, such as robustness (Hendrycks et al., 2019) and domain adaptive ability (Sun et al., 2020). This work differs significantly: our work does not aim to improve supervised classification by using self-supervised tasks. Rather we offer a new and interesting perspective into the relationship between the two tasks, a finding that allows estimation of generalization performance for semantic classification models from easily obtained self-supervised rotation prediction results.

Limitation. We might encounter corner cases (e.g., balls and airplanes) where they are even impossible for a human to tell the image rotation angle. The proposed method, because of the intrinsic limitation of rotation prediction, would fail in such cases. In fact, as an underlying assumption, the self-supervised task is required to be well-defined and non-trivial (Sun et al., 2020). Moreover, we might encounter images from unseen classes in the open world. To improve the estimation under this scenario, a feasible solution is using out-of-distribution detection techniques (Hsu et al., 2020; Hendrycks & Gimpel, 2017; Lee et al., 2018; Liang et al., 2018) to detect and reject such novel images.

Table 5. The extent of linear correlation (r) on CIFAR-10 when using Jigsaw or rotation prediction as the pretext task. The accuracy (%) of classification (Class. Acc) on the CIFAR-10 test set is also reported. Jigsaw and semantic classification also have a strong linear correlation ($r > 0.95$). However, it degrades the classification accuracy when trained in the multi-task network.

Backbone	Rotation		Jigsaw	
	Cor. (r)	Class Acc.	Cor. (r)	Class Acc.
ResNet26	0.975	92.84	0.953	90.90
ResNet44	0.981	93.73	0.958	91.81

Feasibility of other self-supervised tasks than rotation prediction. There are many other self-supervised tasks, such as jigsaw puzzles (Noroozi & Favaro, 2016) and colorization (Larsson et al., 2017). Since our network jointly learns semantic classification and pretext tasks, we select the pretext task based on the criteria listed at the beginning of Section 3.2.1. In the preliminary experiment, we found that *colorization* discards the color information of images and thus makes the learning of semantic classification unstable. In addition, we observe that *Jigsaw* and semantic classification exhibit a linear correlation ($r > 0.950$) under the CIFAR-10 setup (Table 5). Nevertheless, in our multi-task network, Jigsaw compromises the classification accuracy by about 2%, while rotation prediction does not affect semantic classification (shown in Table 2). Moreover, our network needs to learn the semantic classification and pretext task together. It would be interesting to study other potential ways to use pretext tasks for estimating classifier accuracy.

6. Conclusion

In this paper, we try to estimate classifier accuracy under varying testing environments without human-annotated labels. We report a strong linear correlation ($r > 0.88$) between semantic classification accuracy and rotation prediction accuracy through empirical studies. This interesting finding makes it feasible to estimate classifier accuracy on unlabeled test sets using rotation prediction performance which can be easily obtained. The extensive experiments provide supports and insights for using such a pretext task to predict the generalization ability of the classifier.

Acknowledge

This work was supported in part by the Australian Research Council (ARC) Discovery Early Career Researcher Award (DE200101283), ARC Centre of Excellence in Robotic Vision (CE140100016), and the ARC Discovery Project (DP210102801). We thank all anonymous reviewers for their constructive comments in improving this paper.

References

- Arora, S., Ge, R., Neyshabur, B., and Zhang, Y. Stronger generalization bounds for deep nets via a compression approach. In *International Conference on Machine Learning*, pp. 254–263, 2018.
- Bojanowski, P. and Joulin, A. Unsupervised learning by predicting noise. In *International Conference on Machine Learning*, pp. 517–526, 2017.
- Bucci, S., D’Innocente, A., Liao, Y., Carlucci, F. M., Caputo, B., and Tommasi, T. Self-supervised learning across domains. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- Carlucci, F. M., D’Innocente, A., Bucci, S., Caputo, B., and Tommasi, T. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2229–2238, 2019.
- Caron, M., Bojanowski, P., Joulin, A., and Douze, M. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision*, pp. 132–149, 2018.
- Caron, M., Bojanowski, P., Mairal, J., and Joulin, A. Unsupervised pre-training of image features on non-curated data. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2959–2968, 2019.
- Corneanu, C. A., Escalera, S., and Martinez, A. M. Computing the testing error without a testing set. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2677–2685, 2020.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 248–255, 2009.
- Deng, J., Berg, A. C., Li, K., and Fei-Fei, L. What does classifying more than 10,000 image categories tell us? In *Proceedings of the European Conference on Computer Vision*, pp. 71–84, 2010.
- Deng, W. and Zheng, L. Are labels always necessary for classifier accuracy evaluation? In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2021.
- Doersch, C. and Zisserman, A. Multi-task self-supervised visual learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2051–2060, 2017.
- Donmez, P., Lebanon, G., and Balasubramanian, K. Unsupervised supervised learning i: Estimating classification and regression errors without labels. *Journal of Machine Learning Research*, 11(4), 2010.

- Eilertsen, G., Jönsson, D., Ropinski, T., Unger, J., and Ynnerman, A. Classifying the classifier: dissecting the weight space of neural networks. *arXiv preprint arXiv:2002.05688*, 2020.
- Everingham, M., Van Gool, L., Williams, C. K., Winn, J., and Zisserman, A. The pascal visual object classes challenge 2007 (voc2007) results. 2007.
- Gidaris, S., Singh, P., and Komodakis, N. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations*, 2018.
- Griffin, G., Holub, A., and Perona, P. Caltech-256 object category dataset. Technical Report 7694, California Institute of Technology, 2007. URL <http://authors.library.caltech.edu/7694>.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. In *International Conference on Machine Learning*, pp. 1321–1330, 2017.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9729–9738, 2020.
- Hendrycks, D. and Dietterich, T. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2019.
- Hendrycks, D. and Gimpel, K. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *International Conference on Learning Representations*, 2017.
- Hendrycks, D., Mazeika, M., Kadavath, S., and Song, D. Using self-supervised learning can improve model robustness and uncertainty. In *Advances in Neural Information Processing Systems*, pp. 15663–15674, 2019.
- Hsu, Y.-C., Shen, Y., Jin, H., and Kira, Z. Generalized odin: Detecting out-of-distribution image without learning from out-of-distribution data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 10951–10960, 2020.
- Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.
- Huber, P. J. *Robust statistics*, volume 523. John Wiley & Sons, 2004.
- Hull, J. J. A database for handwritten text recognition research. *IEEE Transactions on pattern analysis and machine intelligence*, 16(5):550–554, 1994.
- Jiang, Y., Krishnan, D., Mobahi, H., and Bengio, S. Predicting the generalization gap in deep networks with margin distributions. In *International Conference on Learning Representations*, 2018.
- Jiang, Y., Neyshabur, B., Mobahi, H., Krishnan, D., and Bengio, S. Fantastic generalization measures and where to find them. In *International Conference on Learning Representations*, 2019.
- Kendall, A. and Gal, Y. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in neural information processing systems*, pp. 5574–5584, 2017.
- Kolesnikov, A., Zhai, X., and Beyer, L. Revisiting self-supervised visual representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 1920–1929, 2019.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Larsson, G., Maire, M., and Shakhnarovich, G. Colorization as a proxy task for visual understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6874–6883, 2017.
- Lee, K., Lee, K., Lee, H., and Shin, J. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *Advances in Neural Information Processing Systems*, pp. 7167–7177, 2018.
- Lee, M. A., Zhu, Y., Srinivasan, K., Shah, P., Savarese, S., Fei-Fei, L., Garg, A., and Bohg, J. Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks. In *International Conference on Robotics and Automation*, pp. 8943–8950, 2019.
- Liang, S., Li, Y., and Srikant, R. Enhancing the reliability of out-of-distribution image detection in neural networks. In *International Conference on Learning Representations*, 2018.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft coco: Common objects in context. In *Proceedings of the European Conference on Computer Vision*, pp. 740–755, 2014.
- Liu, C., Dollár, P., He, K., Girshick, R., Yuille, A., and Xie, S. Are labels always necessary for neural architecture search? In *Proceedings of the European Conference on Computer Vision*, pp. 798–813, 2020.

- Madani, O., Pennock, D., and Flake, G. Co-validation: Using model disagreement on unlabeled data to validate classification algorithms. In *Advances in neural information processing systems*, volume 17, pp. 873–880, 2004.
- Metzger, S., Srinivas, A., Darrell, T., and Keutzer, K. Evaluating self-supervised pretraining without using labels. *arXiv preprint arXiv:2009.07724*, 2020.
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. 2011.
- Neyshabur, B., Bhojanapalli, S., McAllester, D., and Srebro, N. Exploring generalization in deep learning. In *Advances in neural information processing systems*, pp. 5947–5956, 2017.
- Noroozi, M. and Favaro, P. Unsupervised learning of visual representations by solving jigsaw puzzles. In *Proceedings of the European Conference on Computer Vision*, pp. 69–84, 2016.
- Ovadia, Y., Fertig, E., Ren, J., Nado, Z., Sculley, D., Nowozin, S., Dillon, J., Lakshminarayanan, B., and Snoek, J. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. In *Advances in Neural Information Processing Systems*, pp. 13991–14002, 2019.
- Peng, X., Usman, B., Kaushik, N., Hoffman, J., Wang, D., and Saenko, K. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017.
- Platanios, E., Poon, H., Mitchell, T. M., and Horvitz, E. J. Estimating accuracy from unlabeled data: A probabilistic logic approach. In *Advances in Neural Information Processing Systems*, pp. 4361–4370, 2017.
- Platanios, E. A., Dubey, A., and Mitchell, T. Estimating accuracy from unlabeled data: A bayesian approach. In *International Conference on Machine Learning*, pp. 1416–1425, 2016.
- Recht, B., Roelofs, R., Schmidt, L., and Shankar, V. Do cifar-10 classifiers generalize to cifar-10? *arXiv preprint arXiv:1806.00451*, 2018.
- Ren, J., Liu, P. J., Fertig, E., Snoek, J., Poplin, R., Deprieto, M., Dillon, J., and Lakshminarayanan, B. Likelihood ratios for out-of-distribution detection. In *Advances in Neural Information Processing Systems*, pp. 14707–14718, 2019.
- Ren, Z. and Jae Lee, Y. Cross-domain self-supervised multi-task feature learning using synthetic imagery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 762–771, 2018.
- Sun, Y., Tzeng, E., Darrell, T., and Efros, A. A. Unsupervised domain adaptation through self-supervision. *arXiv preprint arXiv:1909.11825*, 2019.
- Sun, Y., Wang, X., Liu, Z., Miller, J., Efros, A., and Hardt, M. Test-time training with self-supervision for generalization under distribution shifts. In *International Conference on Machine Learning*, pp. 9229–9248, 2020.
- Taori, R., Dave, A., Shankar, V., Carlini, N., Recht, B., and Schmidt, L. Measuring robustness to natural distribution shifts in image classification. In *Advances in Neural Information Processing Systems*, 2020.
- Unterthiner, T., Keysers, D., Gelly, S., Bousquet, O., and Tolstikhin, I. Predicting neural network accuracy from weights. *arXiv preprint arXiv:2002.11448*, 2020.
- Vincent, P., Larochelle, H., Bengio, Y., and Manzagol, P.-A. Extracting and composing robust features with denoising autoencoders. In *International Conference on Machine Learning*, pp. 1096–1103, 2008.
- Wang, X. and Gupta, A. Unsupervised learning of visual representations using videos. In *Proceedings of the IEEE international conference on computer vision*, pp. 2794–2802, 2015.
- Zhai, X., Oliver, A., Kolesnikov, A., and Beyer, L. S4l: Self-supervised semi-supervised learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1476–1485, 2019.