

Assessing Pre-Trained Models for Transfer Learning Through Distribution of Spectral Components

Tengxue Zhang¹, Yang Shu^{1*}, Xinyang Chen^{2*}, Yifei Long¹, Chenjuan Guo¹, Bin Yang¹

¹School of Data Science and Engineering, East China Normal University

²School of Computer Science and Technology, Harbin Institute of Technology (Shenzhen)

{txzhang, yflong}@stu.ecnu.edu.cn, {yshu, cjguo, byang}@dase.ecnu.edu.cn, chenxinyang@hit.edu.cn

Abstract

Pre-trained model assessment for transfer learning aims to identify the optimal candidate for the downstream tasks from a model hub, without the need of time-consuming fine-tuning. Existing advanced works mainly focus on analyzing the intrinsic characteristics of the entire features extracted by each pre-trained model or how well such features fit the target labels. This paper proposes a novel perspective for pre-trained model assessment through the Distribution of Spectral Components (DISCO). Through singular value decomposition of features extracted from pre-trained models, we investigate different spectral components and observe that they possess distinct transferability, contributing diversely to the fine-tuning performance. Inspired by this, we propose an assessment method based on the distribution of spectral components which measures the proportions of their corresponding singular values. Pre-trained models with features concentrating on more transferable components are regarded as better choices for transfer learning. We further leverage the labels of downstream data to better estimate the transferability of each spectral component and derive the final assessment criterion. Our proposed method is flexible and can be applied to both classification and regression tasks. We conducted comprehensive experiments across three benchmarks and two tasks including image classification and object detection, demonstrating that our method achieves state-of-the-art performance in choosing proper pre-trained models from the model hub for transfer learning.

1 Introduction

With the widespread use of deep neural networks, numerous pre-trained models are readily available in the artificial intelligence community, such as those found on HuggingFace (Wolf et al. 2019) and TensorFlowHub (Abadi et al. 2016). Thanks to being trained on large-scale datasets, these models provide fundamental and general knowledge for downstream tasks. Fine-tuning is a successful paradigm that leverages this knowledge to enhance performance on target tasks, without having to start the training process from scratch. Despite the remarkable capabilities of pre-trained models, such as those pre-trained on large-scale images for computer vision (Russakovsky et al. 2015; He et al. 2016; Dosovitskiy

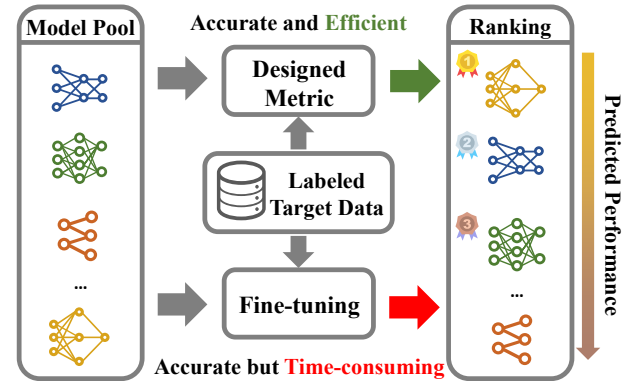


Figure 1: The overall framework for assessing pre-trained models for transfer learning. Given a model pool, a designed metric quickly predicts and ranks model performance on the target dataset. The predicted rankings are expected to strongly correlate with the ground-truth fine-tuning results.

et al. 2020), and large corpus for natural language processing (Devlin et al. 2018; Brown et al. 2020), there is no single pre-trained model that excels in all downstream tasks across various domains. Consequently, selecting the best pre-trained model to apply to downstream tasks has become a challenging problem. Fine-tuning each pre-trained model on the target dataset is a straightforward way to obtain the ground-truth performance and select the best model. However, this procedure can be time-consuming with a large-scale pool of pre-trained models. As a result, recent works have focused on designing a metric to quickly predict the performance of models on the target dataset for pre-trained model selection (Ding et al. 2024), as illustrated in Figure 1.

Most of these works calculate the transferability metric based on the original features of the target dataset extracted by pre-trained models. Some primarily consider statistical information (Nguyen et al. 2020), such as the expected conditional probability of target labels given the pre-training labels, to evaluate how well the extracted features fit the target labels. Others focus on the intrinsic characteristics of the entire extracted features, such as inter-class separation or intra-class and inter-class structure (Wang et al. 2023).

One significant limitation of previous methods is their re-

*Corresponding author

liance on the original extracted features, often overlooking the fine-grained changes occurring during the fine-tuning process. For cases where the difference between the source and target domains is substantial, the extracted features obtained from fine-tuned models may differ considerably from those before fine-tuning. In such situations, a transferability metric based solely on the original extracted features of pre-trained models and the fitness between such features and the target labels can be unreliable. Despite some works considering the dynamics in the fine-tuning process, their metrics are designed based on hypotheses specific to class relationships, making them primarily only applicable to classification tasks. For example, the class separability hypothesis of the fine-tuning process is explored in methods like GBC (Pándy et al. 2022), which directly calculates the pairwise class overlap using the Bhattacharyya coefficients and SFDA (Shao et al. 2022), which projects the extracted features into a more class-separable Fisher space.

In this work, we consider the fine-grained changes in extracted features from a novel perspective of the **D**istribution of **S**pectral **C**omponents (**DISCO**) and propose a flexible framework for assessing models for transfer learning. First, we resort to singular value decomposition (SVD) to investigate the changes in the extracted features before and after fine-tuning. We observe that, during the fine-tuning process, different spectral components are under varying degrees of change and exhibit distinct transferability. After fine-tuning, the distribution of singular values tends to become more concentrated on more transferable components. Leveraging these observations and considering the potential effect of different spectral components on the fine-tuning process, we design a framework to evaluate pre-trained models through the distribution of the spectral components. Furthermore, taking into account the label information of the downstream task, we design different metrics based on this framework for two general task types: classification and regression tasks. The successful implementation in both the image classification and the object detection tasks demonstrates that our method is more flexible and can be customized with specific metrics for different downstream tasks. In summary, the major contributions of our work are as follows:

- We investigate the fine-grained changes in extracted features during the fine-tuning process through the perspective of spectral components. We propose an assessment framework for pre-trained models based on the distribution of spectral components.
- Our proposed framework is highly flexible. We design assessment metrics tailored to the two general tasks: classification and regression. These metrics are then successfully applied in image classification and object detection.
- Extensive experiments on three benchmarks and two specific tasks, demonstrate that our method achieves state-of-the-art performance in assessing pre-trained models for transfer learning.

2 Related Work

Transfer Learning Transfer learning leverages transferring knowledge from the source domains to enhance perfor-

mance in the target domain. One prominent branch of transfer learning is inductive transfer, commonly known as fine-tuning (Erhan et al. 2010; Yosinski et al. 2014), which utilizes pre-trained models to improve target task performance, especially when labeled data is limited. Various strategies exist for fine-tuning, ranging from Full Fine-Tuning (FFT) (Donahue et al. 2014; Radford et al. 2021) to Parameter-Efficient Fine-Tuning (PEFT) (Kumar et al. 2022; Jia et al. 2022; Liu et al. 2023; Dettmers et al. 2024). However, this process can be time-consuming due to the iterative nature of training and the need to select appropriate hyperparameters. Despite the availability of numerous pre-trained models, choosing the most effective pre-trained model for a specific target task without requiring a substantial time investment remains a significant challenge.

Pre-trained Model Selection Pre-trained model selection seeks to quickly identify the best model for downstream tasks from a model zoo without extensive fine-tuning. Previous works like NCE (Tran, Nguyen, and Hassner 2019) and LEEP (Nguyen et al. 2020) have focused on probabilistic methods based on the expected empirical distribution of the target labels. However, these metrics are not suitable for self-supervised pre-trained models due to their reliance on pre-trained classifiers. To address this limitation, LogME (You et al. 2021) estimates maximum label marginalized likelihood, and NLEEP (Nguyen et al. 2020) extends LEEP by replacing the output layer with a Gaussian mixture model. GBC (Pándy et al. 2022) calculates the overlap between pairwise target classes in extracted features. SFDA improves class separability using a Fisher space and then applies a Bayesian classifier. PED (Li et al. 2023) stimulates dynamics during fine-tuning through the lens of potential energy and integrates into existing model selection metrics. TMI (Xu and Kang 2023) uses intra-class feature variance as a performance indicator, positing that lower variance reflects tighter class feature clustering. Inspired by neural collapse phenomena, NCTI (Wang et al. 2023) develops metrics to measure the distance from the current status of pre-trained models to their hypothetical fine-tuned state. Different from these methods that primarily rely on the entire extracted features, we propose a novel perspective for pre-trained model selection through the distribution of spectral components.

3 Method

3.1 Problem Formulation

Given a pool of M pre-trained models $\{\phi^m\}_{m=1}^M$ and a target dataset $D = \{x_i, y_i\}_{i=1}^N$ with N samples, pre-trained model selection aims to efficiently assess and select the optimal model for transfer learning. Typically, in the fine-tuning paradigm, we initialize a predictor head on the backbone of the pre-trained model and then fine-tune the entire network on the target dataset. By fine-tuning each model in a brute-force manner, we obtain the ground-truth performance $\{P^m\}_{m=1}^M$ for the model hub. To avoid this time-consuming process, model selection methods compute an assessment score S^m for each pre-trained model ϕ^m . A higher S^m suggests that the model ϕ^m is likely to perform well when transferred to the target dataset. Ideally, the calculated scores

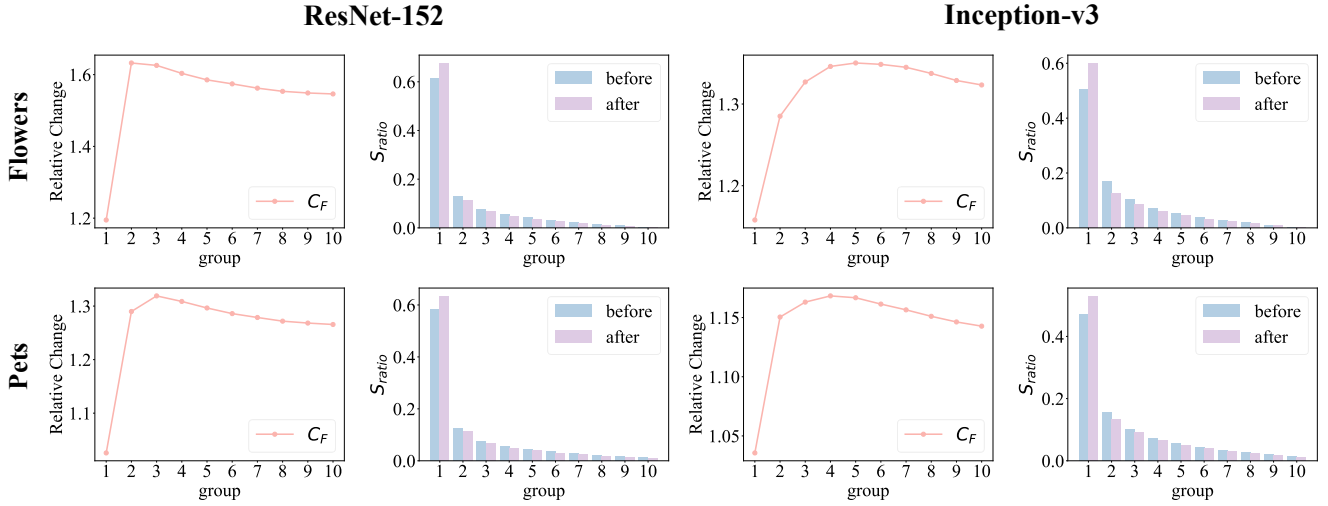


Figure 2: The relative changes of Frobenius norm C_F and the proportion of singular values S_{ratio} in different spectral components of extracted features before and after fine-tuning.

$\{S^m\}_{m=1}^M$ for the model hub should strongly correlate with their actual fine-tuning performance $\{P^m\}_{m=1}^M$, allowing us to identify the best model for transfer.

3.2 Transferability Assessment Framework through Spectral Component Distribution

To explore fine-grained changes in extracted features during fine-tuning, we employ singular value decomposition (SVD) for in-depth analysis. We fine-tune ResNet-152 (He et al. 2016), DenseNet-201 (Huang et al. 2017), and Inception-v3 (Szegedy et al. 2016) on Caltech101 (Fei-Fei, Fergus, and Perona 2004), Oxford-102 Flowers (Nilsback and Zisserman 2008), and Pets (Parkhi et al. 2012) respectively. Let ϕ^m represent the m -th pre-trained model and ϕ'^m the fine-tuned model. The features extracted by ϕ^m and ϕ'^m are denoted as $\mathbf{Z} \in \mathbb{R}^{N \times d}$ and $\mathbf{Z}' \in \mathbb{R}^{N \times d}$ respectively, where N is the target dataset size and d is the feature dimension. We apply SVD to extracted features \mathbf{Z} and \mathbf{Z}' as follow:

$$\mathbf{Z} = \mathbf{U}\Sigma\mathbf{V}^T, \quad \mathbf{Z}' = \mathbf{U}'\Sigma'\mathbf{V}'^T \quad (1)$$

where \mathbf{U} and \mathbf{U}' are orthogonal matrices of left singular vectors, while \mathbf{V} and \mathbf{V}' are orthogonal matrices of right singular vectors, Σ and Σ' are diagonal matrices of non-negative singular values arranged in descending order. We divide the original feature \mathbf{Z} into different spectral components by grouping the singular values in descending order into G groups, each containing s consecutive singular values. Similar operations are performed on \mathbf{U} and \mathbf{V} . The g -th spectral component is defined as the product of the matrices corresponding to the g -th group's columns of \mathbf{U}_g , the diagonal singular values Σ_g , and rows of \mathbf{V}_g^T . In this way, we obtain different spectral components \mathbf{Z}_g and \mathbf{Z}'_g of the original feature \mathbf{Z} as truncated SVD reconstruction. We perform the same process for \mathbf{Z}' for consistency:

$$\mathbf{Z}_g = \mathbf{U}_g\Sigma_g\mathbf{V}_g^T, \quad \mathbf{Z}'_g = \mathbf{U}'_g\Sigma'_g\mathbf{V}'_g^T \quad (2)$$

We investigate the changes in different spectral components of the extracted features before and after fine-tuning from two perspectives. The first perspective is to compute the relative change of different spectral components using the Frobenius norm, which can be calculated as:

$$C_F^g = \frac{\|\mathbf{Z}_g - \mathbf{Z}'_g\|_F}{\|\mathbf{Z}_g\|_F} \quad (3)$$

We set $G = 10$ in our experiments. Partial results are shown in Figure 2, with the complete results available in Appendix B.1. We observe that the spectral component with the larger singular value generally undergoes the least variation while variation increases in subsequent groups and then decreases slightly for components with smaller singular values. These differences suggest varying levels of transferability during fine-tuning, with the components with larger singular values often exhibiting higher transferability.

Another perspective is to compute the singular value ratio S_{ratio}^g for singular values σ^g of each spectral component:

$$S_{\text{ratio}}^g = \frac{\sum_{j=1}^s \sigma_j^g}{\sum_{i=1}^d \sigma_i} \quad (4)$$

where σ_i is the i -th singular value in Σ . By analyzing the proportional alterations of singular values across different spectral components, we observe that after fine-tuning, the larger singular values increase in proportion, while smaller ones diminish. This suggests that the distribution of singular values tends to concentrate on components with larger singular values during fine-tuning, indicating that more transferable components play a greater role in fine-tuned features.

Inspired by our empirical observations of the fine-grained changes during fine-tuning, the distribution of spectral components influences the transfer learning process and results. We first propose an intuitive solution to assess a model's transferability. The key insight is that a concentration of singular values on more transferable spectral components

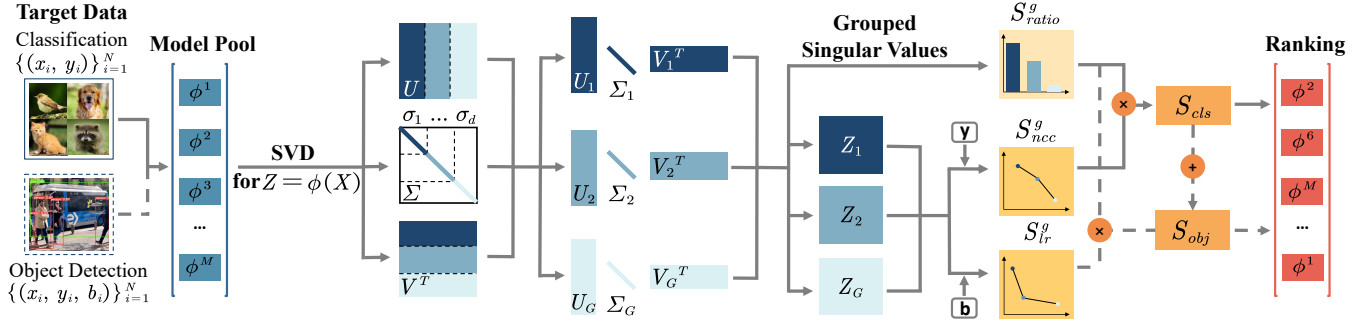


Figure 3: Overview of DISCO’s framework (better viewed in color). S_{ratio}^g represents the singular value ratio of the g -th spectral component, while S_{ncc}^g and S_{lr}^g are task-specific scores for classification and regression tasks, respectively. The overall transferability of the entire feature is calculated through the perspective of the distribution of spectral components.

suggests greater overall transferability of entire features. To quantify this, we calculate the ratio of top- k groups’ singular values to all singular values for each pre-trained model:

$$S_{\text{topk}} = \frac{\sum_{g=1}^k \sum_j^s \sigma_j^g}{\sum_{i=1}^d \sigma_i} \quad (5)$$

However, in practice, the distribution of singular values varies across different models, making it difficult to draw a hard boundary of k , where these top- k components are considered transferable and the rest are deemed not. Focusing solely on the top- k components ignores the potential contributions of other ones. Moreover, without task-specific evaluation, the actual performances of different spectral components on target tasks are uncertain, particularly in cases with significant domain gaps. Therefore, to quantitatively assess the transferability and enable comparisons between different pre-trained models, it is crucial to leverage the labels from the target dataset to obtain the performance on downstream tasks. Based on these considerations, we propose our DISCO framework that quantifies the overall transferability of each pre-trained model through the perspective of the distribution of spectral components in their extracted features:

$$S_{\text{DISCO}} = \sum_{g=1}^G S_{\text{task}}^g * S_{\text{ratio}}^g \quad (6)$$

Here, S_{task}^g represents the quantitative transferability of the g -th spectral component on specific downstream tasks, while S_{ratio}^g is its corresponding singular value ratio. S_{task}^g captures how well this component performs when transferred to new tasks, and S_{ratio}^g reflects the importance of this component within the entire feature. The summation considers each spectral component’s contribution to the overall transferability of the entire extracted feature. Pre-trained models with features concentrated in more transferable components, and with such components demonstrating superior downstream performance, are considered more suitable candidates for transfer learning.

3.3 Task-Specific Performance Assessment for Spectral Components

Classification and regression are foundational tasks in machine learning, serving as the basis for many specific applications. In this section, we extend our framework by leveraging target labels and designing specific metrics for S_{task}^g . The overview of our framework is shown in Figure 3. By tailoring to both classification and regression tasks, our framework is flexible and applicable to various downstream tasks, including image classification and object detection.

Classification Score Due to the domain gap, pre-trained classifiers from the source domain are not directly applicable to the target task. Therefore, the most straightforward way to evaluate the performance of different spectral components is by assessing their ability to discriminate between classes. Under a commonly used assumption that the class distribution follows a multivariate Gaussian distribution, we use the performance of the nearest centroid classifier as a measure. Given a sample x_i with corresponding feature \mathbf{z}_i^g from the g -th spectral component, the posterior probability that it belongs to the c -th class using Mahalanobis distance can be calculated as:

$$P(y = c | \mathbf{z}_i^g) = -\frac{1}{2} (\mathbf{z}_i^g - \boldsymbol{\mu}_c)^\top \boldsymbol{\Lambda}_c^{-1} (\mathbf{z}_i^g - \boldsymbol{\mu}_c) + \log(\pi_c) \quad (7)$$

where $\boldsymbol{\Lambda}_c$ and $\boldsymbol{\mu}_c$ represent the covariance matrix and mean of the c -th class. The prior probability π_c is given by N_c/N , where N_c is the number of the samples in class c , and N is the total number of samples. We apply the softmax function to normalize the posterior probability of the i -th sample \mathbf{z}_i^g predicted by the g -th spectral component as:

$$\hat{y}_{i,c}^g = \frac{\exp(\log P(y = c | \mathbf{z}_i^g))}{\sum_{k=1}^C \exp(\log P(y = k | \mathbf{z}_i^g))} \quad (8)$$

where $\hat{y}_{i,c}^g$ is the final prediction probability that the i -th sample belongs to class c . As such, we quantify the average confidence of the nearest centroid classifier on the g -th spectral component as the score for classification performance:

$$S_{\text{ncc}}^g = \frac{1}{N} \sum_{i=1}^N \mathbf{y}_i^\top \hat{\mathbf{y}}_i^g \quad (9)$$

where \mathbf{y}_i denotes the one-hot ground truth label of x_i and $\hat{\mathbf{y}}_i^g$ is the C dimensional logits predicted by the g -th spectral component, with each dimension representing one class.

For image classification, the goal is to correctly classify images. Therefore, we use S_{ncc}^g as the performance assessment for the g -th spectral component. Based on the proposed framework, we derive the final assessment criterion for each pre-trained model in image classification as follows:

$$S_{\text{cls}} = \sum_{g=1}^G S_{\text{ncc}}^g * S_{\text{ratio}}^g \quad (10)$$

Regression Score Regression is a task for predicting continuous numerical values. To stimulate the target task fine-tuning and avoid significant time expenditure, we propose a linear approximation between extracted features and their labels. In linear regression, our objective is to find the coefficients β to minimize ordinary least-square (OLS) distance $\|y - \mathbf{Z}\beta\|^2$. However, directly solving this may be problematic due to the potential non-invertibility of \mathbf{Z} . Singular value decomposition (SVD) offers a robust and numerically stable method for matrix inversion. Since SVD is used to divide spectral components, it adds no extra time cost for the regression score. Therefore, we employ SVD to decompose the extracted features \mathbf{Z} and compute its pseudo-inverse \mathbf{Z}^\dagger as:

$$\mathbf{Z} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top, \quad \mathbf{Z}^\dagger = \mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{U}^\top \quad (11)$$

where \mathbf{U} and \mathbf{V} are orthogonal matrices, and $\mathbf{\Sigma}$ is a diagonal matrix containing the singular values of \mathbf{Z} . The coefficient β can be calculated using the pseudo-inverse: $\beta = \mathbf{Z}^\dagger y$.

In the object detection task, where the goal is to predict bounding box coordinates b along with category labels y , we calculate the approximated bounding box coordinates as $\hat{b} = \mathbf{Z}\hat{\mathbf{Z}}^\dagger b$. To evaluate the quality of these approximations, we employ the mean squared error (MSE) between the ground truth and approximated bounding box labels for the g -th spectral component. Given a total of K bounding boxes in the target dataset, the regression score S_{lr}^g for the g -th spectral component is defined as:

$$S_{\text{lr}}^g = -\frac{1}{K \times 4} \sum_{k=1}^K \sum_{j=1}^4 (b_k^{(j)} - \hat{b}_k^{(j)})^2 \quad (12)$$

where $b_k^{(j)}$ and $\hat{b}_k^{(j)}$ are the j -th ground-truth and approximated coordinate of the k -th bounding boxes, respectively. Similarly, based on the proposed framework, we use S_{lr}^g to assess the performance of the g -th spectral component. The final assessment for each pre-trained model in the regression task is as follows:

$$S_{\text{reg}} = \sum_{g=1}^G S_{\text{lr}}^g * S_{\text{ratio}}^g \quad (13)$$

Since both classification and regression scores are required in object detection and they have different scales, we use min-max normalization to standardize each score to a $[0, 1]$ range, avoiding the need for weighted adjustments. We sum the normalized S_{cls} and S_{reg} to obtain the final model selection score for each pre-trained models for object detection:

$$S_{\text{obj}} = S_{\text{cls}} + S_{\text{reg}} \quad (14)$$

3.4 Hard-example Selection

The naive implementation of singular value decomposition on extracted features $\mathbf{Z} \in \mathbb{R}^{N \times d}$ has a computational complexity of $\mathcal{O}(N^3 d)$, which becomes impractical as N increases. To address this issue, we propose a hard-example selection method to sample a subset of the datasets, thereby lowering N and the overall time complexity. The method leverages the insight that a pre-trained model’s performance on hard examples better reflects its transferability. We first project the features into a class-separable space using Linear Discriminant Analysis (LDA) with a complexity of $\mathcal{O}(Nd^2)$ when $N > d$, and then select the N' data points with the lowest confidence scores from its classifier as hard examples (detailed in Appendix A.3). This approach reduces time complexity, without causing a significant drop in performance. In practice, one can consider the trade-off between time and performance when choosing the sampling ratio.

4 Experiments

4.1 Experimental Setup

Evaluation Protocol To quantify the correlation between estimated assessment scores $\{S^m\}_{m=1}^M$ and actual fine-tuning results $\{P^m\}_{m=1}^M$, we use the weighted Kendall’s τ_ω (see Appendix A.1), which measures ranking agreement with higher weights for higher ranks. A larger τ_ω indicates a more effective metric for ranking models. We supplement this with the top- k probability, $Pr(\text{top-}k)$, which measures the probability that the best ground-truth model is among the top k estimated models. While τ_ω evaluates overall ranking agreement, $Pr(\text{top-}k)$ is particularly useful for identifying the best pre-trained model in practical scenarios.

Feature Construction We compute our DISCO over the features extracted by pre-trained models on the target dataset. For image classification, we perform once forward propagation to obtain the extracted feature $\mathbf{Z} \in \mathbb{R}^{N \times d}$. For object detection, varying object counts and additional bounding boxes complicate feature construction. We apply adaptive average pooling to unify the dimensions of box-specific features to a consistent size \hat{h} . The final feature matrix, $\mathbf{Z} \in \mathbb{R}^{K \times \hat{h}}$, is created by concatenating all normalized box-specific features, where K is the total number of bounding boxes in the target dataset (detailed in Appendix A.2).

4.2 Image Classification

Datasets We adopt 11 widely-used datasets in classification tasks, including FGVC Aircraft (Maji et al. 2013), Stanford Cars (Krause et al. 2013), Food101 (Bossard, Guillaumin, and Van Gool 2014), Oxford-IIIT Pets (Parkhi et al. 2012), Oxford-102 Flowers (Nilsback and Zisserman 2008), Caltech101 (Fei-Fei, Fergus, and Perona 2004), CIFAR-10 (Krizhevsky, Hinton et al. 2009), CIFAR-100 (Krizhevsky, Hinton et al. 2009), VOC2007 (Everingham et al. 2010), SUN397 (Xiao et al. 2010), DTD (Cimpoi et al. 2014). The fine-tuning performances (see Appendix B.4) of both supervised CNN and self-supervised CNN models are obtained from (Shao et al. 2022).

	Aircraft	Caltech	Cars	Cifar10	Cifar100	DTD	Flowers	Food	Pets	SUN	VOC	Average
NCE (2019)	-0.161	0.465	<u>0.685</u>	0.709	0.723	0.302	-0.482	0.627	0.772	0.760	0.571	0.452
LEEP (2020)	-0.277	0.605	0.367	0.824	0.677	0.486	-0.291	0.434	0.389	0.658	0.413	0.390
LogME (2021)	0.439	0.463	0.605	0.852	0.725	0.700	0.147	0.385	0.411	0.511	0.695	0.539
NLEEP (2021)	-0.531	0.614	0.489	0.825	0.731	0.820	0.054	0.529	0.955	0.848	0.699	0.548
SFDA (2022)	<u>0.614</u>	0.696	0.518	0.949	0.866	0.575	0.514	<u>0.815</u>	0.522	0.558	0.671	0.663
Etran (2023)	-0.091	0.440	0.246	<u>0.887</u>	0.900	0.303	<u>0.580</u>	0.829	0.713	0.708	0.667	0.562
NCTI (2023)	0.496	<u>0.679</u>	0.647	0.843	0.879	0.704	0.541	0.773	<u>0.924</u>	<u>0.756</u>	<u>0.741</u>	<u>0.726</u>
DISCO	0.652	0.661	0.795	0.823	<u>0.895</u>	<u>0.764</u>	0.712	0.678	0.575	<u>0.773</u>	0.802	0.739

Table 1: Method comparison of their correlation strength with ground-truth fine-tuning accuracy on supervised models. Weighted Kendall’s τ on 11 target datasets and their average are listed above. For each column, the best, and second-best results are in bold, and underlined, respectively. Our method achieves the best overall weighted Kendall’s τ_ω .

	Aircraft	Caltech	Cars	Cifar10	Cifar100	DTD	Flowers	Food	Pets	SUN	VOC	Average
LogME (2021)	0.021	0.075	0.627	0.417	0.146	0.743	0.763	0.686	0.738	0.260	0.181	0.423
NLEEP (2021)	-0.286	0.662	0.595	0.108	0.374	0.779	0.598	0.716	0.864	<u>0.880</u>	-0.091	0.473
SFDA (2022)	0.167	0.674	0.683	0.846	0.789	<u>0.882</u>	0.897	0.837	0.564	0.831	<u>0.621</u>	0.708
NCTI (2023)	0.036	0.811	0.796	0.758	<u>0.811</u>	0.796	0.762	0.945	<u>0.805</u>	0.774	0.606	<u>0.719</u>
DISCO	<u>0.063</u>	<u>0.691</u>	<u>0.768</u>	<u>0.825</u>	0.835	0.918	<u>0.886</u>	<u>0.898</u>	0.542	0.916	0.642	0.726

Table 2: Method comparison of their correlation strength with ground-truth fine-tuning accuracy on self-supervised models. All the settings and notations are the same as Table 1. Our method achieves the best overall weighted Kendall’s τ_ω .

Results and Analysis on Supervised Model

Pre-trained Models We construct a pool of 11 widely-used models including ResNet-34 (He et al. 2016), ResNet-50 (He et al. 2016), ResNet-101 (He et al. 2016), ResNet-152 (He et al. 2016), DenseNet-121 (Huang et al. 2017), DenseNet-169 (Huang et al. 2017), DenseNet-201 (Huang et al. 2017), MNet-A1 (Tan et al. 2019), MobileNet-v2 (Sandler et al. 2018), GoogleNet (Szegedy et al. 2015) and Inception-v3 (Szegedy et al. 2016). All these supervised source models were pre-trained on ImageNet and were downloaded from the PyTorch repository.

Performance Comparison We evaluate DISCO against established model selection metrics, including NCE (Tran, Nguyen, and Hassner 2019), LEEP (Nguyen et al. 2020), LogME (You et al. 2021), SFDA (Shao et al. 2022), NLEEP (Li et al. 2021), ETran (Gholami et al. 2023) and NCTI (Wang et al. 2023). Results in Table 1 show that DISCO achieves SOTA with an average τ_ω of 0.739, a 1.8% relative improvement over the second-best NCTI. DISCO’s superior performance is attributed to its ability to consider the fine-grained contribution of different spectral components during the fine-tuning process. This leads to strong results on fine-grained classification datasets like Aircraft, Cars, and Flowers, where the relative improvements are 31.5%, 22.9%, and 31.6% compared to the NCTI, respectively.

Results and Analysis on Self-Supervised Model

Pre-trained Models To further assess the effectiveness of our metric in self-supervised learning (SSL), we select 10 SSL models based on ResNet-50 architecture (He et al.

2016), including BYOL (Grill et al. 2020), Deepclusterv2 (Caron et al. 2018), Infomin (Tian et al. 2020), MoCo-v1 (He et al. 2020), MoCo-v2 (Chen et al. 2020), Instance Discrimination (Wu et al. 2018), PCL-v1 (Li et al. 2020), PCL-v2 (Li et al. 2020), Sela-v2 (Asano, Rupprecht, and Vedaldi 2019) and SWAV (Caron et al. 2020).

Performance Comparison Since NCE and LEEP require pre-trained classifiers, they are no longer suitable for ranking self-supervised models, which typically lack classifiers. We compare DISCO with NLEEP, LogME, SFDA and NCTI, as shown in Table 2. DISCO performs consistently well in self-supervised model selection, while conventional metrics collectively tend to exhibit lower correlation scores. The DISCO leads and secures the highest average weighted Kendall’s τ_ω of 0.726 ranking correlation among 11 datasets. This superior performance highlights DISCO’s ability to capture intricate details from spectral components, unlike traditional metrics that rely on overall feature characteristics. By leveraging this rich information, DISCO provides more effective model assessment for transfer learning.

4.3 Object Detection

Datasets We select five datasets that span different domains and sizes for object detection to evaluate our model selection metric, including Blood (Roboflow 2022), Fork-Lift (Traore 2022), NFL (home 2022), Valorant Video Game (Magonis 2022), and CSGO Video Game (ASD 2022).

Pre-trained Models To demonstrate the generality of DISCO for pre-trained models on object detection, we assemble a pool of six models: YOLOv5s, YOLOv5m,

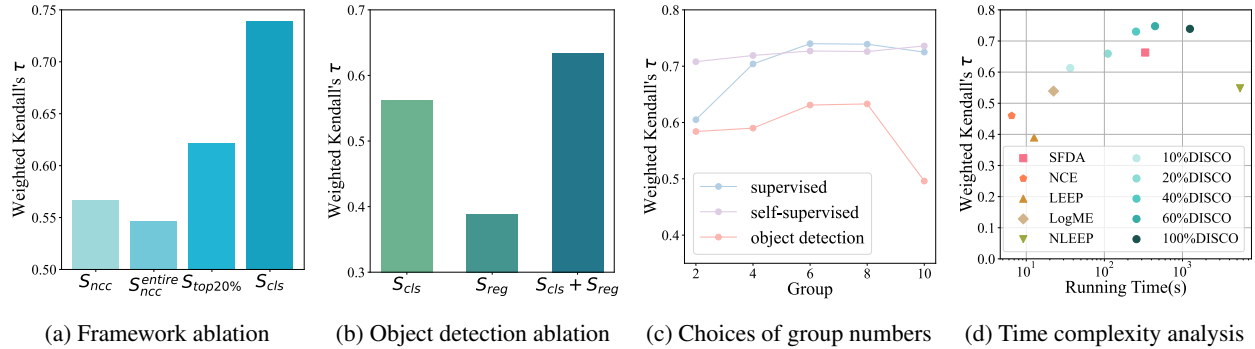


Figure 4: (a) Ablation study on the framework and (b) on the object detection benchmark. (c) The average τ_ω with different group numbers on three benchmarks. (d) Method comparison w.r.t average running time (seconds) and τ_ω on 11 datasets.

Method	Pr(top1)	Pr(top2)	Pr(top3)	τ_ω
LogME (2021)	0.600	0.600	0.800	0.374
PACTran (2022)	0.400	0.400	0.600	0.140
Linear (2022)	0.200	0.400	0.600	0.214
SFDA (2022)	0.400	0.600	1.000	0.312
LogME+ S_{lmr}	0.600	0.600	0.800	0.400
PACTran+ S_{lmr}	0.200	0.400	0.800	0.322
Linear+ S_{lmr}	0.200	0.200	0.800	0.306
SFDA+ S_{lmr}	0.200	0.400	0.800	0.202
ETran (2023)	0.600	0.800	1.000	0.522
DISCO	1.000	1.000	1.000	0.634

Table 3: Results on object detection benchmark. Weighted Kendall’s τ_ω is the average of five object detection datasets.

YOLOv5n (Jocher et al. 2023), YOLOv8s, YOLOv8m, and YOLOv8n (Jocher, Chaurasia, and Qiu 2023), all pre-trained on the COCO dataset (Lin et al. 2014). We obtain the ground-truth fine-tuned results for these source models from the HuggingFace (Wolf et al. 2019).

Performance Comparison We compare our DISCO with naive LogME, SFDA, PACTran (Ding et al. 2022), and Linear method (Ding et al. 2022) using classification-only score, as well as their variants incorporating the regression score S_{lmr} baseline proposed in Etran. The experimental comparison results of τ_ω and $Pr(\text{top-k})$ are presented in Table 3. Generally, variants with the additional regression score baseline perform better than the naive estimation scores and Etran surpasses the regression baseline. However, our DISCO achieves the SOTA result of average weighted Kendall τ_ω of 0.634, reflecting a 21% improvement over Etran’s τ_ω of 0.522. The superior results show that our method is not only applicable in classification tasks but also excels in object detection, indicating the generality of our approach across a broad range of applications.

4.4 Further Analysis

Ablation Study Figure 4a shows the ablation results of S_{ncc} ($\sum_{g=1}^G S_{ncc}^g$, w/o S_{ratio}^g compared to our framework),

S_{ncc}^{entire} (S_{ncc} when $G = 1$), $S_{top20\%}$ (top20% singular values ratio) compared to our proposed S_{cls} for image classification. These results highlight the effectiveness of DISCO in accounting for different spectral components with different contributions. Figure 4b presents the ablation results for each term in object detection, where their combination achieves the highest τ_ω , demonstrating the indispensability of the regression score in object detection and the flexibility of our framework. More details are in Appendix B.2.

Choices of Group Numbers The number of groups for spectral components is a crucial factor in determining the performance of our method. As shown in Figure 4c, small groups fail to differentiate between spectral components’ contributions, while large groups cause each group to have too little information, reducing accuracy in downstream tasks. We found that $G = 6$ or $G = 8$ offer the best balance and are ideal for practical use.

Time Complexity We compare the running time vs. average τ_ω for 11 datasets on supervised models using different model selection metrics in Figure 4d. DISCO with the full dataset delivers the best τ_ω but has high time costs, albeit still far below fine-tuning. Using hard-example selection to sample 60%, 40%, 20% and 10% of the dataset reduces time complexity with minimal performance loss. Notably, at 20% sampling, DISCO matches SFDA’s performance with lower time costs, and at 10%, it outperforms LogME with a similar time cost, demonstrating the efficiency of DISCO and the hard-example selection method. Detailed analyses are provided in Appendix B.3.

5 Conclusion

In this paper, we investigate the fine-grained changes in extracted features during fine-tuning and propose an assessment framework for pre-trained models based on the Distribution of Spectral Components (DISCO). We design tailored metrics for classification and regression, which successfully apply to image classification and object detection. Extensive experiments demonstrate that our method achieves SOTA in assessing pre-trained models for transfer learning.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (62406112, 62306085, 62372179) and Shenzhen College Stability Support Plan (GXWD20231130151329002).

References

- Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G. S.; Davis, A.; Dean, J.; Devin, M.; et al. 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*.
- Asano, Y. M.; Rupprecht, C.; and Vedaldi, A. 2019. Self-labelling via simultaneous clustering and representation learning. *arXiv preprint arXiv:1911.05371*.
- ASD. 2022. wlots Dataset.
url <https://universe.roboflow.com/asd-culfr/wlots> . Accessed: 2023-01-27.
- Bossard, L.; Guillaumin, M.; and Van Gool, L. 2014. Food-101—mining discriminative components with random forests. In *ECCV*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. In *NeurIPS*.
- Caron, M.; Bojanowski, P.; Joulin, A.; and Douze, M. 2018. Deep clustering for unsupervised learning of visual features. In *ECCV*.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*.
- Chen, X.; Fan, H.; Girshick, R.; and He, K. 2020. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing textures in the wild. In *CVPR*.
- Dettmers, T.; Pagnoni, A.; Holtzman, A.; and Zettlemoyer, L. 2024. Qlora: Efficient finetuning of quantized llms. In *NeurIPS*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Ding, N.; Chen, X.; Levinboim, T.; Changpinyo, S.; and Soriccut, R. 2022. Pactran: Pac-bayesian metrics for estimating the transferability of pretrained models to classification tasks. In *ECCV*.
- Ding, Y.; Jiang, B.; Yu, A.; Zheng, A.; and Liang, J. 2024. Which Model to Transfer? A Survey on Transferability Estimation. *arXiv preprint arXiv:2402.15231*.
- Donahue, J.; Jia, Y.; Vinyals, O.; Hoffman, J.; Zhang, N.; Tzeng, E.; and Darrell, T. 2014. Decaf: A deep convolutional activation feature for generic visual recognition. In *ICML*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Erhan, D.; Courville, A.; Bengio, Y.; and Vincent, P. 2010. Why does unsupervised pre-training help deep learning? In *AISTATS*.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *IJCV*, 88: 303–338.
- Fei-Fei, L.; Fergus, R.; and Perona, P. 2004. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *CVPR*.
- Gholami, M.; Akbari, M.; Wang, X.; Kamranian, B.; and Zhang, Y. 2023. Etran: Energy-based transferability estimation. In *ICCV*.
- Grill, J.-B.; Strub, F.; Althé, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; et al. 2020. Bootstrap your own latent—a new approach to self-supervised learning. In *NeurIPS*.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *CVPR*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- home. 2022. NFL-competition Dataset.
url <https://universe.roboflow.com/home-mxzv1/nfl-competition> . Accessed: 2023-01-18.
- Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *CVPR*.
- Jia, M.; Tang, L.; Chen, B.-C.; Cardie, C.; Belongie, S.; Hariharan, B.; and Lim, S.-N. 2022. Visual prompt tuning. In *ECCV*.
- Jocher, G.; Chaurasia, A.; and Qiu, J. 2023. YOLO by Ultralytics.
- Jocher, G.; et al. 2023. YOLOv5 by Ultralytics. 2020. *URL* <https://github.com/ultralytics/yolov5>.
- Krause, J.; Deng, J.; Stark, M.; and Fei-Fei, L. 2013. Collecting a large-scale dataset of fine-grained cars.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Kumar, A.; Raghunathan, A.; Jones, R.; Ma, T.; and Liang, P. 2022. Fine-tuning can distort pretrained features and underperform out-of-distribution. *arXiv preprint arXiv:2202.10054*.
- Li, J.; Zhou, P.; Xiong, C.; and Hoi, S. C. 2020. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*.
- Li, X.; Hu, Z.; Ge, Y.; Shan, Y.; and Duan, L.-Y. 2023. Exploring model transferability through the lens of potential energy. In *ICCV*.
- Li, Y.; Jia, X.; Sang, R.; Zhu, Y.; Green, B.; Wang, L.; and Gong, B. 2021. Ranking Neural Checkpoints. In *CVPR*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *ECCV*.

- Liu, X.; Zheng, Y.; Du, Z.; Ding, M.; Qian, Y.; Yang, Z.; and Tang, J. 2023. GPT understands, too. *AI Open*.
- Magonis, D. 2022. valorant Dataset.
url <https://universe.roboflow.com/daniels-magonis-Opjzx/valorant-9ufcp> . Accessed: 2023-01-27.
- Maji, S.; Rahtu, E.; Kannala, J.; Blaschko, M.; and Vedaldi, A. 2013. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*.
- Nguyen, C.; Hassner, T.; Seeger, M.; and Archambeau, C. 2020. Leep: A new measure to evaluate transferability of learned representations. In *ICML*.
- Nilsback, M.-E.; and Zisserman, A. 2008. Automated flower classification over a large number of classes. In *2008 Sixth Indian conference on computer vision, graphics & image processing*. IEEE.
- Pándy, M.; Agostinelli, A.; Uijlings, J.; Ferrari, V.; and Mensink, T. 2022. Transferability estimation using bhat-tacharyya class separability. In *CVPR*.
- Parkhi, O. M.; Vedaldi, A.; Zisserman, A.; and Jawahar, C. 2012. Cats and dogs. In *CVPR*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*.
- Roboflow, T. 2022. Blood Cell Detection Dataset.
url <https://universe.roboflow.com/team-roboflow/blood-cell-detection-1ekwu> . Accessed: 2023-01-18.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *IJCV*, 115: 211–252.
- Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; and Chen, L.-C. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*.
- Shao, W.; Zhao, X.; Ge, Y.; Zhang, Z.; Yang, L.; Wang, X.; Shan, Y.; and Luo, P. 2022. Not all models are equal: Predicting model transferability in a self-challenging fisher space. In *ECCV*.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions. In *CVPR*.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the inception architecture for computer vision. In *CVPR*.
- Tan, M.; Chen, B.; Pang, R.; Vasudevan, V.; Sandler, M.; Howard, A.; and Le, Q. V. 2019. Mnasnet: Platform-aware neural architecture search for mobile. In *CVPR*.
- Tian, Y.; Sun, C.; Poole, B.; Krishnan, D.; Schmid, C.; and Isola, P. 2020. What makes for good views for contrastive learning? In *NeurIPS*.
- Tran, A. T.; Nguyen, C. V.; and Hassner, T. 2019. Transferability and hardness of supervised classification tasks. In *ICCV*.
- Traore, M. 2022. Forklift Dataset.
url <https://universe.roboflow.com/mohamed-traore-2ekkp/forklift-dsitv> . Accessed: 2023-01-15.
- Wang, Z.; Luo, Y.; Zheng, L.; Huang, Z.; and Baktashmotlagh, M. 2023. How far pre-trained models are from neural collapse on the target dataset informs their transferability. In *ICCV*.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Wu, Z.; Xiong, Y.; Yu, S. X.; and Lin, D. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*.
- Xiao, J.; Hays, J.; Ehinger, K. A.; Oliva, A.; and Torralba, A. 2010. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*.
- Xu, H.; and Kang, U. 2023. Fast and accurate transferability measurement by evaluating intra-class feature variance. In *ICCV*.
- Yosinski, J.; Clune, J.; Bengio, Y.; and Lipson, H. 2014. How transferable are features in deep neural networks? In *NeurIPS*.
- You, K.; Liu, Y.; Wang, J.; and Long, M. 2021. Logme: Practical assessment of pre-trained models for transfer learning. In *ICML*.