Guided Recommendation for Model Fine-Tuning

Hao Li, Charless Fowlkes, Hao Yang, Onkar Dabeer, Zhuowen Tu, Stefano Soatto
AWS AI Labs
{haolimax, fowlkec, haoyng, onkardab, ztu, soattos}@amazon.com

Abstract

Model selection is essential for reducing the search cost of the best pre-trained model over a large-scale model zoo for a downstream task. After analyzing recent hand-designed model selection criteria with 400+ ImageNet pre-trained models and 40 downstream tasks, we find that they can fail due to invalid assumptions and intrinsic limitations. The prior knowledge on model capacity and dataset also can not be easily integrated into the existing criteria. To address these issues, we propose to convert model selection as a recommendation problem and to learn from the past training history. Specifically, we characterize the meta information of datasets and models as features, and use their transfer learning performance as the guided score. With thousands of historical training jobs, a recommendation system can be learned to predict the model selection score given the features of the dataset and the model as input. Our approach enables integrating existing model selection scores as additional features and scales with more historical data. We evaluate the prediction accuracy with 22 pre-trained models over 40 downstream tasks. With extensive evaluations, we show that the learned approach can outperform prior hand-designed model selection methods significantly when relevant training history is available.

1. Introduction

Much of the success of deep learning can be ascribed to its flexibility: One can train a neural network on a task, and then use it on a different one, typically after fine-tuning. There are currently two trends for scaling this practice: One is to pre-train a large number of specialized models (a "Model Zoo" [10]) and then select one to fine-tune once the downstream task of interest becomes manifest, typically with a smaller fine-tuning dataset. Another is to pre-train a single "Foundation Model" which is then used to support any and all downstream tasks [47, 57].

Without additional specifications, the second case is a subset of the first, for one can take the Model Zoo and Model Selection (MS) mechanism and call it a single model. For this reason, Foundation Models are characterized as *homogeneous and task-agnostic*, where homogeneity refers to a single neural network architecture, in contrast with the *heterogeneous* collection of models in a zoo. Even with this restriction, the model zoo is more general, for nothing prevents a Foundation Model to be part of a zoo. In addition, selecting a smaller dedicated model pretrained for a task can be much more efficient than using a giant monolithic model For these reasons, we focus on model selection over a large heterogeneous model zoo for fine-tuning as the key solution for scaling inference to a wide variety of downstream tasks.

Brute-force model selection [1, 12] requires fine-tuning each pre-trained model on the task of interest, and then ranking them using the test error on a held-out dataset as a *model selection score*. This is not feasible for large model zoos. Current model-selection methods therefore aim to *predict* the model selection score without actually fine-tuning.

However, current model selection methods do not take into explicit account even basic characteristics of the fine-tuning dataset, such as the number of classes or the number of images, nor of the pre-trained model, such as the model family, the size of the input, the number of parameters and the dataset on which it is pre-trained. While coarse, these features can affect the best model to fine-tune, since a mismatch between fine-tuning dataset size and pre-trained model, or input dimensions, or number of classes, can influence the success of downstream performance.

Instead of proposing yet another model selection score, *we propose re-framing model selection as a recommender system*, and directly predict the selection score and corresponding ranking, from whatever existing model selection scores are readily available, in addition to whatever coarse features a user deems informative – which may be context dependent, as some users may wish to penalize large models, or models that require high-resolution input. Such features help *guide* the model selection using criteria beyond raw downstream validation error. For this reason, we refer to our recommendation approach as *guided*, in addition to *trained*.

We find that incorporating model size, dataset size, cardinality of the hypothesis set and other simple features already improves the prediction of the expected model selection score compared to current model selection methods. Coarse features, such as the index of the model class family (convolutional, fully-connected, residual, attention-based, etc.) can help associate certain architectural inductive biases such as translation vs. permutation invariance, to the best-fitting downstream tasks, for instance object detection vs. image inpainting or segmentation.

Our contribution can be summarized as:

- We conduct comprehensive analysis of existing model selection approaches with a large heterogeneous model zoo and confirmed their limitations. We find feature-based model selection becomes inaccurate when the target dataset is different from the source task and the effect of model initialization diminishes as the number of images grows. The useful meta information and prior knowledge in the training history are often neglected and cannot be easily integrated into existing model selection criteria.

- We convert the model selection problem as model recommendation by learning from past training history. The meta information of both dataset and model are embedded as features and a recommender system can be learned to predict the performance. The existing model selection can be used as additional features and makes the framework comply with existing approaches. We show significant performance improvement over traditional model selection methods when historical training data is available and relevant.

In the next section we formalize the MS problem, and discuss the issues with existing approaches. In section 3 we describe our approach to casting it as a recommendation system and evaluate it in the following section.

2. Background

2.1. Problem Formalization

Let T_i be pre-training candidate tasks, with $i = 1, \dots, M$, encoded in their corresponding datasets $D_i = \{(x_k, y_k)\}_{k=1}^{N_i}$ (the dataset is all a model knows about the task prior to training), used to train a chosen architecture (function class) $\phi_i(\cdot; w_i)$ by minimizing a loss function L_i with respect to the weights w, yielding

$$w_i = \arg\min_w \underbrace{\sum_{(x_k,y_k)\in D_i} \ell(y_k, \phi_i(x_k; w))}_{L_i} \doteq \hat{w}(\phi_i, D_i) \tag{1}$$

where the pre-trained weights w_i are a function $\hat{w}$ of the dataset, the architecture, and the pointwise loss ℓ, which is typically cross-entropy, in addition to the optimization procedure, regularizers, hyperparameters, and other factors that we omit for simplicity since we wish to focus on the relative role of the architecture and the dataset.

When fine-tuning a model ϕ_i for a different task $T_{j\neq i}$, the architecture is conditioned on using ϕ_i, either as a frozen embedding, or as an initialization, so the model to be fine-tuned for the task T_j using the dataset D_j, has the form $\phi_j(\phi_i(\cdot; w_i); w)$ and the fine-tuning loss $\tilde{V}_{ij}$ is

$$\min_w \sum_{(x_k,y_k)\in D_j} \ell(y_k, \phi_j(\phi_i(x_k; w_i); w)) \doteq \tilde{V}(\phi_i|\phi_j, D_i, D_j) \tag{2}$$

$\tilde{V}_{ij}$ is the *empirical model selection score*, corresponding to the training error during fine-tuning. Brute-force model selection consists of solving

$$\hat{\phi}_i = \arg\min_{\phi_i} \tilde{V}(\phi_i \mid \phi_j, D_i, D_j). \tag{3}$$

It is immediate to see that for the function $\hat{\phi}_i$ to be constant with respect to ϕ_j (that is, for the pre-trained representation to be *task-agnostic*) it would either have to be conditioned on all possible tasks (including those with different hypothesis spaces, hence be *non-homogeneous*), or be a *trivial* lossless compression of the data, for the task could turn out to be reproducing an identical copy of the data. This would defer the burden of learning to the fine-tuning phase, annihilating the value of pre-training and undermining the main premise of Foundation Models as homogeneous *and* task-agnostic *and* optimal for fine-tuning. This further reinforces our focus on heterogeneous model selection.

The validation error on a held-out dataset, or ideally the marginal over all possible fine-tuning datasets, is the (expected) *model selection score*

$$\hat{V}_{ij} = \mathbb{E}_{D_j} \tilde{V}(\phi_i|\phi_j, D_i, D_j) \tag{4}$$

which is clearly not computable.

Feature-based Model Selection Most recent MS methods (e.g., LFC [10], PARC [5], and LogME [60]) extract features with each candidate model on the target dataset, and then calculate the MS score. Given a dataset $D = \{\mathbf{x}, \mathbf{y}\}$, let $f_w(x_i)$ denote the feature vector extracted from penultimate layer of pre-trained model ϕ for data x_i. The LFC score is calculated as

$$S_{\mathbf{LFC}}(\mathbf{x}, \mathbf{y}) = f_w(\mathbf{x}) f_w(\mathbf{x})^T \cdot \mathbf{y}\mathbf{y}^T \tag{5}$$

where $(\mathbf{y}\mathbf{y}^T)_{i,j} = 1$ if x_i and x_j have the same label and -1 otherwise. The normalized S_{LFC} can be interpreted as the Pearson Correlation between the features and labels. The assumption is that a better candidate model's initialization usually has similar features for samples with the same labels.

2.2. Model Selection Limitations

Feature-based MS usually fix the candidate model as feature extractors or assume the fine-tuning process does not change the backbone weights much. However, the assumption may not hold in practice and results in failure. Here we evaluate three MS algorithms (LFC [10], PARC [5], and LogME [60]) with two settings and demonstrate the cases that they can fail. a) fine-tuning the 400+ ImageNet pre-trained models on ImageNet. The fine-tuning performance should be consistent with their pre-training performance. b) we select 22 models near the Pareto frontier of the 400+ models and fine-tune them on 40 downstream datasets. More setup details can refer to Sec 4.1.

Difficulty with heterogeneous model zoo Existing MS methods usually validate their approach with a homogeneous model zoo in which models differ only in pre-trained domains. And it is often believed that better ImageNet model also transfer better on downstream tasks [31], which seems to makes the problem of MS with heterogeneous model zoo trivial. However, we find that the optimal architecture or Pareto front models can be task dependent, which relies on both the inductive bias of the model and the characteristics of the dataset. In addition, existing MS algorithms can fail to accommodate new architectures such as ViTs which have much smaller feature dimensions compared to ResNets. As shown in Fig. 1, ViTs are outliers for MS methods without explicit normalization. Normalizing the input features of all architectures can solve this issue improve the Pearson correlation scores. This was observed in [5] and they further improve PARC by applying PCA and adding normalized network depth to incorporate the network capacity. However, the heuristic cannot generalize across architectures, e.g., ViTs that comes without the same depth concept in terms of convolutional layers. More details on the heterogeneous model zoo can be found in Appendix.

Difficulty with dissimilar dataset MS algorithms such as LFC and PARC assume that models with consistent feature similarities and label similarities can generalize better, which is valid for few-shot or linear probing where the majority of weights do not change much. However, the effect of initialization often diminishes as the dataset size grows. When training data is large, a random initialized model with high capacity can yield better performance than a pre-trained simple model. The MS score of the random initialized model can be lower than the pre-trained one, which does not represent the underlying generalization ability of the model. As shown in Fig. 2, we see a clear difference for the MS performance between `dogs` and `aircrafts`. We know that the `dogs` dataset is similar to ImageNet but not `aircrafts`[1]. It verifies that MS algorithms can fail to predict top performing models when a downstream task is very different from the one used for training the source model.

Difficulty of incorporating prior knowledge The failures cases of model selection are mainly due to the lack of proper usage of model's inductive bias for datasets with special characteristics. Note that the inductive biases can be heuristically added to existing MS score e.g., PARC [5] incorporates the model depth. However, adding such heuristics to MS score requires ad-hoc tuning of the scale of the new added score, which is hard to extend to more indicators. On the other hand, the importance of the model inductive biases are often associated with the dataset characteristics, which is hard to integrate manually, e.g., “a random initialized large model generalize better than a small pre-trained model for a large dataset" and “a shallow model perform the same as a deeper model for a simple task". In such cases, the effectiveness of model depth is also determined by the dataset characteristics. Therefore we need a model the connection between characteristics of the task and inductive bias of the model. However, existing MS methods often use a small probe set with fixed number of training images to reduce the computation of cost of MS computation, which neglects the actual dataset size information.

3. Learning to Recommend

Instead of manually designing a model selection criteria, we propose to learn to select models from the training history. Given the historical training results, we can characterize the features of dataset and models, and use fine-tuning performance as the ground truth. The goal is to predict performance on the target dataset. Then a model selector can be learned to select the optimal model for a given task.

3.1. Model Selection as Recommendation

In order to frame model selection as a recommendation system, we represent the pre-trained model ϕ_i with an element of a vector space $\mathbf{v}_i$, and/or a simpler vectorized version of coarse features such as the number of parameters, input dimension, number of classes, index of the architecture family and pre-training dataset, etc. Similarly, we embed the fine-tuning dataset D_j onto a set of features $\mathbf{v}_j$, for instance its cardinality and dimension of the hypothesis space. In addition, we can use any available predictive model selection score U_{ij}. A recommendation system then implements a learnable map that, for each pre-trained model i and down-

[1] Stanford Dogs [30] was built using images and annotation from ImageNet for the task of fine-grained image categorization.

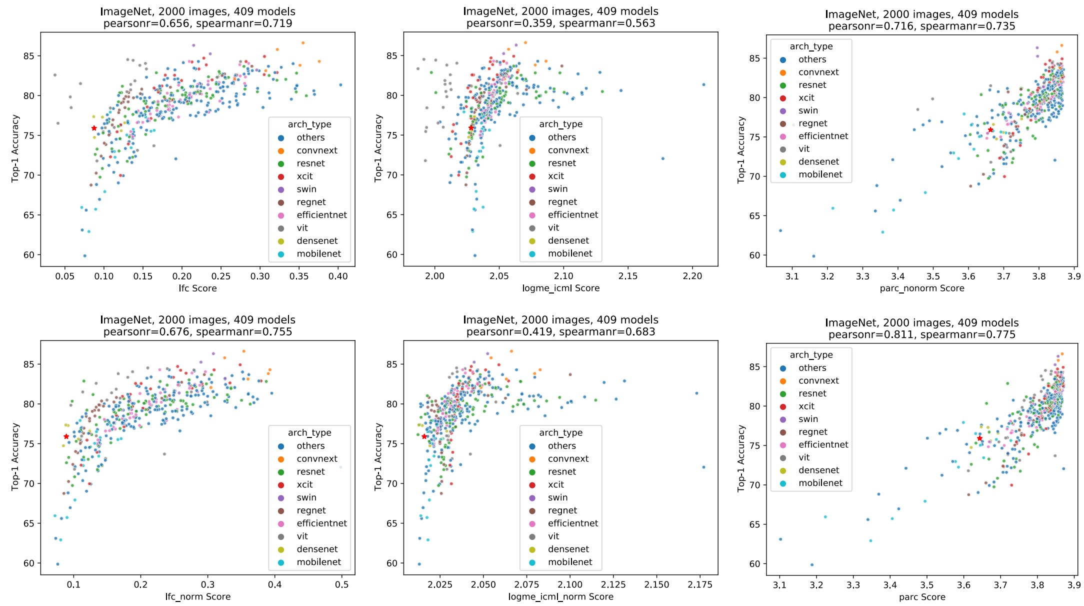

Figure 1. Comparison of MS algorithms with 400+ ImageNet pre-trained models. In the first row, features are not normalized. LFC [10], left, LogME [60] and PARC [5] (w.o. normalization) all treat ViTs (gray points) as outliers. The features are normalized in the 2nd row and the Pearson correlation increases for all methods. The ViTs are not outliers, indicating the importance of feature normalization for heterogeneous models.

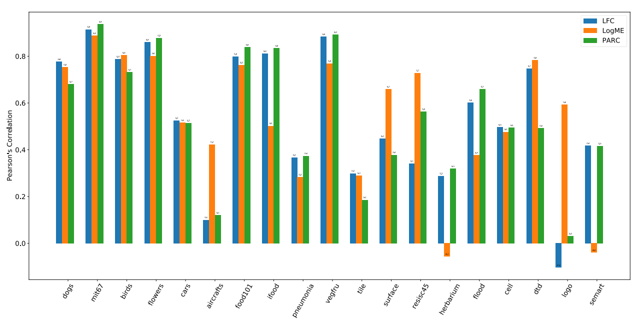

Figure 2. Pearson Correlation of three MS methods on the 19 fine-grained datasets with 22 models. Note that the correlation scores for *aircrafts* and *pneumonia* are much lower than other datasets.

stream task j predicts the expected score $\hat{V}_{ij}$:

$$\begin{aligned} \mathrm{RM} : [1, \ldots, M] \times [1, \ldots, J] \quad &\rightarrow \quad \mathbb{R} \\ (i, j) \quad &\mapsto \quad \hat{V}_{ij} = \mathrm{RM}(i, j) \end{aligned} \tag{6}$$

where $\hat{V}_{ij} = \psi_w(\mathbf{v_i}(\phi_i), \mathbf{v_j}(D_j), U_{ij})$ and ψ_w is a parameterized map, for instance a factorization machine, with learnable parameters w, trained to approximate the validation scores $\hat{V}_{ij}$. By choosing the features $\mathbf{v}_i, \mathbf{v}_j$ a user can factor in additional model selection criteria besides the structure of the data, which is often captured in the model selection scores U_{ij}, to guide the recommendation. Since the analytical expression for the functions $\hat{w}(\cdot)$, $\hat{\phi}(\cdot)$ and $\hat{V}(\cdot)$ are intractable, in this paper we study the problem in (3) empirically in the next section.

3.2. Recommendation Model

There are several options for learning the recommendation model. If the model zoo is fixed, a straight-forward solution is to learn a classifier that directly maps a given task to the best model [46]. The challenge lies at the usual insufficient training samples in comparison with the high dimensions of the dataset representation (e.g., 2048 for ResNet feature). And the fixed model zoo size makes it hard to adapt when new models are added to the model zoo. Therefore, we mainly consider following two options:

- **Linear Regression (LR)** A LR model can be learned to predict performance with the concatenated task and model features. However, LR learns the effect of each feature independently and the interaction among features can not be modeled.
- **Factorization Machine (FM)**. Factorization Machine (FM) [49] is widely used in recommender systems and CTR prediction. FM is often preferred over Linear Regression (LR) as it can learn the correlations among different features via latent embedding, even when there is no data point for the correlation. Given N historical transfer learning results $\{(\mathbf{d}_i, \mathbf{m}_i, \mathbf{y}_i)\}_{i=1}^N$, where $\mathbf{d}_i = \mathbf{v}(D) \in \mathbb{R}^d$ and $\mathbf{m}_i = \mathbf{v}(\phi) \in \mathbb{R}^m$ represents the feature embedding of dataset D and model ϕ for the ith

fine-tuning job, d and m are their embedding length, y_i is the fine-tuning top-1 accuracy on the validation set. Let $\mathbf{z} \in \mathbb{R}^{d+m}$ denote the concatenated features of the pair of dataset and model, the predicted score of FM is

$$S_{\text{FM}}(\mathbf{z}) = w_0 + \sum_{i=1}^{|z|} w_i z_i + \sum_{i=1}^{|z|-1} \sum_{j=i+1}^{|z|} \langle \mathbf{u}_i, \mathbf{u}_j \rangle z_i z_j$$

where $\mathbf{u}_i \in \mathbb{R}^k$ is the latent representation of the ith feature. Note that the first two terms is actually LR. With the third term, FM considers interactions among features in addition to linear combination of features.

In the next sub-section, we will describe the feature embedding for datasets and models in detail.

3.3. Characterizing Datasets and Models

Dataset Embedding We explore the following descriptors for describing a task:

- **Task difficulty**: If a task can be solved with a model's initial weights without much change, then the task is relatively easy for the model, e.g., linear probing (fixed embedding + SVM) is often used as a baseline for transfer learning. If a task gets high performance with linear probing, then it indicates the dataset is relatively easy to solve with a simple linear classifier. A MS score calculated with a fixed backbone (e.g., ResNet-18) can estimate the relative difficulty of the dataset.
- **number of samples**: The dataset size affects the task difficulty and model selection. A few-shot task is generally harder than tasks with large sizes and requires a strong model. The larger the dataset size, the more possibility of choosing a model without a strong initialization. Note that current MS algorithms (e.g., LFC [10]) usually use a prob set with fixed size, while in reality the prob set size could vary significantly.
- **number of classes**: When the total images are fixed, the task difficulty usually increase as the number of classes.

Model Embedding To characterize the model's inductive bias, we use the following features for model embedding:

- **architecture family**: architectures of the same family usually have similar inductive biases as they consist of similar modules, blocks and activation functions. We use the architecture family to categorize the inductive biases of models of the same family, such as ConvNeXt, ViT, Swin-Transformers, EfficientNet and etc.
- **input size**: it is reported higher resolution usually helps for downstream tasks [31], and we see this is true for fine-grained tasks (e.g., EfficientNet-B3 works best for `cars` and `aircrafts` as seen in Appendix). On the other hand, simple cases (e.g., MNIST) may not benefit from higher resolutions.
- **model capacity**: a model with high capacity usually generalizes better with more data. This is measured by the number of parameters.
- **model complexity**: the calculation cost (GMACs) can represent the complexities.
- **pre-trained domain**: the pre-trained domain matters for the downstream task performance. If the source dataset is available, then the domain distance between the source domain and the target domain can be a indicator. However, such information is not always available. We have models pre-trained on ImageNet-1K and ImageNet-22K.

Additional Features The advantage of recommender system is that features related to the prediction can always be added, which makes the solution scalable to new features. Beyond the embedding of datasets and models, we can add additional features that are relevant to the performance prediction. The existing MS scores can be treated as additional features, as it considers the feasibility of the model's initial features (Eq 5). Other features such as the semantic distance between the target task and the model's source task can also be added as additional feature, which could be useful for few-shot or zero-shot learning. We will leave this extension for future works.

4. Experiments

4.1. Settings

Datasets We collected three benchmarks and a total of 40 image classification tasks, including *19 fine-grained datasets*, *DomainNet* [45] and the *VTAB* [61]. Those datasets cover a wide range of domains and applications, such as *scenes*, *objects*, *food*, *texture*, *art* and *medical imaging*. DomainNet consists of 6 datasets of different domains with the same labels. VTAB consists of various tasks which can be categorized into *natural*, *structured* and *special*. More details about the datasets can be found in Appendix.

Models The TIMM model zoo collected more than 550 ImageNet pre-trained models. We evaluated all pre-trained models and keep 409 models that can be fine-tuned with batch size 32 with a single v100 GPU. We evaluated their single image inference latency and identified the Pareto frontier in their latency-accuracy trade-off plot (see Appendix). We select 22 models that are near the Pareto Front curve, which covers a wide range of common architecture families, including ReseNet [17], DenseNet [21], MobileNet [20], EfficientNet [53], ViTs [11], Swin-T [36] and ConvNeXt [37]. The complete model list can be found in Appendix.

Training History The models in the TIMM model zoo usually have ImageNet validation accuracy. Those results are obtained by training the model from scratch. If we reinitialize the last layer of each pre-trained model and fine-tune on ImageNet, we should expect the performance same with the reported results. In addition, we fine-tuned the selected 22 models end-to-end with HPO to obtain their best Top-1 accuracy on the 40 downstream tasks. All pre-trained models are trained with a single V100 GPU with the same range of hyperparameters. More details about the models and fine-tuning settings can be found in Appendix.

Evaluation Metrics We measure MS performance on a given dataset (or probe set) with Pearson correlation coefficient, which measures linear correlation between MS score and oracle transfer performance, which is the covariance of the two variables divided by the product of their standard deviations. We use the mean Pearson correlation over all datasets in a benchmark to for comparison.

Dataset Sub-Sampling For feature-based MS on ImageNet, we sample 2,000 images from the ImageNet training set, with the number of images per class set to 2. We use the fixed sub-sampled prob set for evaluating all models. For MS on downstream tasks, we sample at most 2,000 images with the constraint that no class has more than 25 images.

Recommendation Tasks We consider the following scenarios for model recommendation: a) Learn from the training history of one dataset with a subset of models and evaluate unseen models on the same dataset (e.g., ImageNet). b) Learn from the training history of one dataset and evaluate with the same models on unknown downstream tasks. c) Learning from the training history of both ImageNet and downstream tasks, and evaluate MS with known models on unseen tasks. Note that the diversity and amount of training samples (number of datasets and models) increases in these settings progressively. An illustration of the three settings can be seen in Fig. 3.

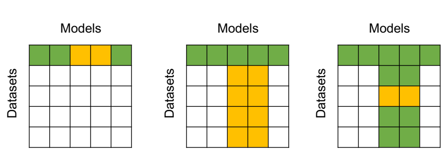

Figure 3. An illustration of the three evaluation settings. The green boxes represent the available training pairs of dataset and model with fine-tuning accuracy, and the yellow boxes indicate the pairs of dataset and model to be predicted for their performance.

4.2. Recommendation Results

4.2.1 Learning from the training history on ImageNet and evaluating unseen models on the same dataset

The ImageNet pre-trained model zoo provides off-the-shelf Top-1 accuracy for 400+ models, which can be used as the groundtruth performance. Note that given the volume of ImageNet, it is expected that models with or without ImageNet pre-training will converge to the same accuracy. To evaluate the learned MS on predicting the performance of unseen architectures on ImageNet, we randomly split the 400+ models with 80% of them for training and the rest 20% for evaluation. In Table 1, we compare the learned MS with different training features with the traditional feature-based MS methods.

With the pre-trained model weights, we see that learned MS (both LR and FM) with only meta features perform better than S_{LogME}. When the MS score is used as additional feature, the performance is better than using simply meta features or the MS score itself. When the models are randomly initialized, the feature-based MS methods fail to rank the models due to the randomness of extracted feature. In contrast, learning based method still get reasonable correlation score and is robust even when the noisy MS score is added as additional feature. The knowledge that larger models usually generalizes better can still be learned.

Note that since all training data are based on ImageNet training history, the dataset features are the same for all training data, and the correlation between dataset feature and model feature cannot be well learned. The learned MS score is mainly determined by the model feature, i.e., models with large capacity has better performance. Thus we see FM does not show advantage over a simple LR model, which is expected. We will see difference when expanding the training set in Sec. 4.2.3.

Table 1. MS learned with only ImageNet training history. The ImageNet benchmark samples 80% of the 409 models as training set while the rest of models are used for evaluation. The experiment is repeated 10 times and the mean/std values are reported.

Methods	Features	ImageNet Pre-trained	ImageNet Random Init.
feature-based	S_{LFC} [10]	0.65 ± 0.07	0.03 ± 0.10
	S_{LogME} [60]	0.35 ± 0.09	0.04 ± 0.08
	S_{PARC} [5]	$\mathbf{0.83} \pm 0.04$	0.08 ± 0.09
LR (ours)	$\mathbf{d}, \mathbf{m}$	0.53 ± 0.07	$\mathbf{0.57} \pm 0.10$
	$\mathbf{d}, \mathbf{m}, S_{\text{LFC}}$	0.73 ± 0.06	0.56 ± 0.10
	$\mathbf{d}, \mathbf{m}, S_{\text{LogME}}$	0.55 ± 0.08	0.56 ± 0.09
	$\mathbf{d}, \mathbf{m}, S_{\text{PARC}}$	$\mathbf{0.85} \pm 0.04$	0.57 ± 0.11
FM (ours)	$\mathbf{d}, \mathbf{m}$	0.54 ± 0.06	$\mathbf{0.57} \pm 0.10$
	$\mathbf{d}, \mathbf{m}, S_{\text{LFC}}$	0.70 ± 0.12	0.56 ± 0.10
	$\mathbf{d}, \mathbf{m}, S_{\text{LogME}}$	0.55 ± 0.09	0.56 ± 0.10
	$\mathbf{d}, \mathbf{m}, S_{\text{PARC}}$	$\mathbf{0.84} \pm 0.05$	0.57 ± 0.11

Table 2. We evaluate the average Pearson correlation of predicted performance and the groundtruth performance of 22 models on each benchmark. The ImageNet column is the MS learned with all 409 ImageNet training jobs. The column of LOO (leave-one-out) denotes MS learned with combined training history of ImageNet jobs and all downstream jobs except jobs on the test dataset.

Methods	Features	**19 fine-grained**		**6 DomainNet**		**15 VTAB**	
feature-based MS	S_{LFC} [10]	0.55		0.63		0.14	
	S_{LogME} [60]	0.54		0.52		0.20	
	S_{PARC} [5]	0.54		0.50		0.13	
		ImageNet	LOO	ImageNet	LOO	ImageNet	LOO
LR (ours)	$\mathbf{d}, \mathbf{m}$	0.53	0.66	0.80	0.82	0.29	0.37
	$\mathbf{d}, \mathbf{m}, S_{\text{LFC}}$	**0.67**	**0.74**	**0.84**	**0.85**	**0.38**	**0.41**
	$\mathbf{d}, \mathbf{m}, S_{\text{LogME}}$	0.54	0.65	0.81	0.84	0.30	0.36
	$\mathbf{d}, \mathbf{m}, S_{\text{PARC}}$	0.54	0.66	0.81	0.85	0.30	0.40
FM (ours)	$\mathbf{d}, \mathbf{m}$	0.53	0.65	0.81	0.85	0.35	0.39
	$\mathbf{d}, \mathbf{m}, S_{\text{LFC}}$	**0.64**	**0.74**	0.82	**0.87**	**0.39**	0.41
	$\mathbf{d}, \mathbf{m}, S_{\text{LogME}}$	0.60	0.67	0.82	0.86	0.31	0.40
	$\mathbf{d}, \mathbf{m}, S_{\text{PARC}}$	0.56	0.69	**0.86**	0.86	0.30	**0.43**

4.2.2 Learning from ImageNet training history and evaluating known models on downstream tasks

To evaluate the transferability of learned MS on new datasets, we use all 400+ ImageNet pre-training history as the training data and predict the performance of 22 models on three benchmarks. As shown in the ImageNet columns of Table 2, the learned MS gets comparable or significantly better mean Pearson correlation than feature-based MS on all benchmarks, especially on DomainNet. We see adding extra MS feature can improve performance over learning with only meta features. Note that the Pearson correlation of all MS methods are relatively low on the VTAB benchmark which consists of *structured* and *special* tasks that are much different from ImageNet. It verifies that feature based MS may not transfer well for tasks that are very different from the source task. Since the training data contains 400+ off-the-shelf models are pre-trained on ImageNet-1K or ImageNet-22K, it is prone to learn MS rules such as bigger models lead to better performance, which is mostly true on ImageNet. Because the training data only consists of ImageNet, the lack of dataset diversity leads to the learning of such inductive biases. FM does not show much advantage over LR.

4.2.3 Learning from all training history and evaluating known models on new tasks.

The power of the recommendation formulation is that its performance will improve as more training data is available. For example, if the training history contains models performance on a similar dataset as the target dataset, it is possible that the model works well on the reference dataset will rank higher for the given task. We further increase the number and diversity of training history for learning based methods. For each dataset in a benchmark, we train the MS model with all available training history except the ones for that dataset, i.e., *leave-one-out* (LOO) training data. Table 2 shows that with more training data added, the LOO results improve significantly over the results learned only from ImageNet. Also both LR and FM learned with only meta features (underline) are comparable with the ones trained with additional MS features.

4.3. Ablation Study

Comparing Feature-based MS and Learned MS To understand which datasets benefit from the learned MS, we compare the learned MS models with feature-based MS (e.g.,

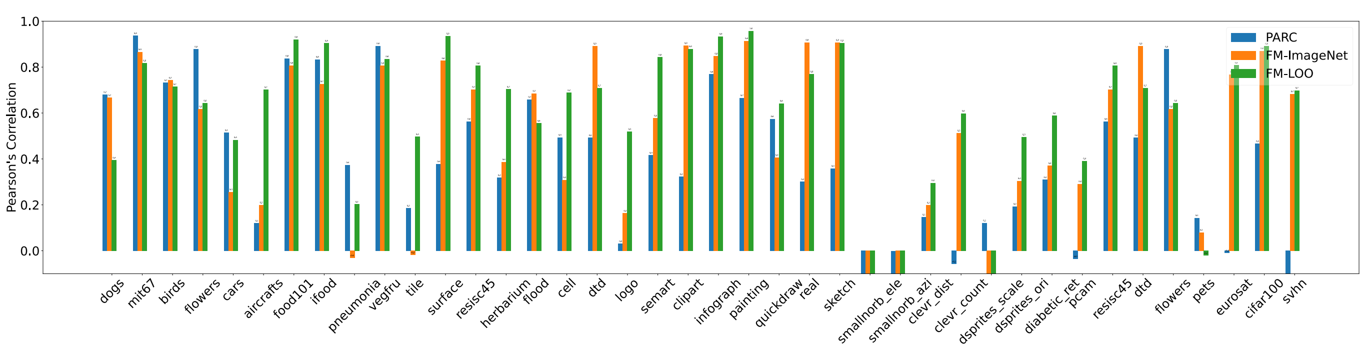

Figure 4. Comparison of learned MS with feature-based MS on each dataset. The FM-ImageNet model is learned with ImageNet only training history and FM-LOO refers to using all datasets except the testing set. Both are trained with only meta feature $\mathbf{d}$ and $\mathbf{m}$.

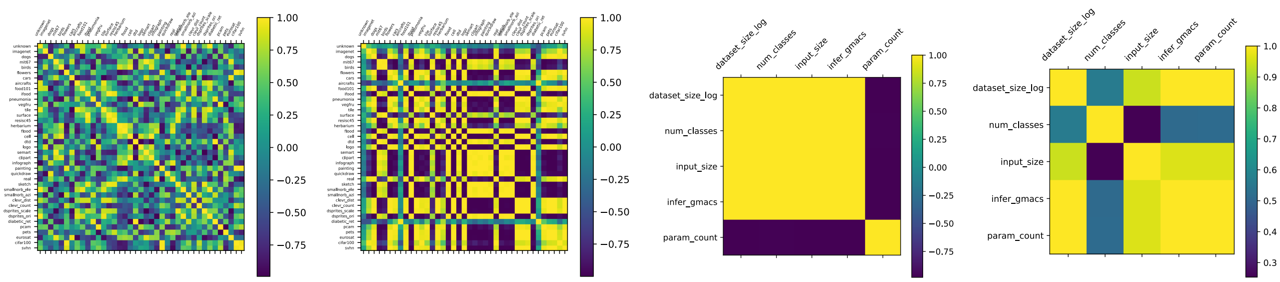

Figure 5. The correlation of selected latent features of FM learned with the ImageNet only training history (a and c) and all training history (b and d). a) and b) select the latent features of dataset IDs, while c) and d) show the correlation of scalar features of $\mathbf{d}$ and $\mathbf{m}$.

PARC) on each dataset. As shown in Fig. 4, the learned FM with LOO outperform PARC on 13/19 fine-grained datasets, 6/6 DomainNet datasets, and 10/15 VTAB datasets. The FM model learned with ImageNet only training history transfers well to benchmarks that are similar to ImageNet (e.g, *dogs, mit67, birds, dtd, clipart, real, and sketch, cifar100, svhn*) but underperforms or fails on dissimilar datasets such as *cars, pneumonia, tile*, and *cell*. On the other hand, with more diverse training data, LOO trained model performs significantly better than PARC or ImageNet-only trained model on datasets such as *aircrafts, tile, logo, smart, dsprites_ori, diabetic_ret* and *resisc45*.

Learned Feature Correlations Fig. 5 visualizes the covariance matrix of learned latent representation $\mathbf{u}$ of selected features in FMs trained with different data. Fig. 5 (a-b) shows the correlation of latent representation of dataset IDs. When only ImageNet history is used, the latent features of other datasets remains random. We see more structured correlation among datasets when more diverse training data is added, i.e., datasets that are similar to each other also have high correlation in their latent representations, such as *dogs* and *pets*. Fig. 5 (c-d) shows clear correlation among scalar features of d and m emerges when more data is used, such as dataset size and MACs/parameters, which indicates that larger dataset and larger models weights more. We can also see less correlation among class number and other features. Note that more advanced algorithms such as field-aware factorization machine (FFM) [26] could further improve the performance, in which the correlation of features belonging to the same filed (e.g., dataset features) are not learned.

5. Related Work

Model Selection MS methods can be categorized based on whether the source dataset is available. When source data is available, models are in the same architecture and differ only in pre-trained domains, the features and labels of source data and target data can be compared with methods such as EMD [9] and NCE [54]. Probabilistic based methods such as H-Score [2], LEEP [42], $\mathcal{N}$LEEP [35] and LogME [60] estimate the likelihood or the marginalized likelihood of labeled target examples, assuming that a linear classifier is added on top of the pre-trained model. Recent TransRate [22] measures the mutual information between the backbone features and the labels, and also extends to layer selection. LFC [10] approximates the fine-tuning dynamics by looking at a linearization of the source model around the pre-trained weights with the assumption that fine-tuned weights tend to remain close to the pre-trained weights. PARC [5] main differs with LFC with the choice of correlation metric. Note that an improved PARC adds model depth with heuristic weight, which is essentially a linear combination of MS score with model's meta feature.

Learning to Recommend There are also learning based methods to recommend dataset, hyperparameters and techniques for a given task. Neural Data Server [59] provides a search engine to find the most useful transfer learning data for the target domain. MS can also be viewed as a hyperparameter selection problem. HyperStar [40] learns to predict the performance of a hyperparameter set for a given image classification task with a end-to-end trained CNN. The work [14] is most relevant to us, in which a general probabilistic model matrix factorization is learned for ML pipeline selection. A learning based approach for MS is [46], which introduced a model routing algorithm for a large number of expert models. Its domain prediction method classifies the expert from the image via an auxiliary network, which is a classification-based approach as we mentioned in Sec 3.2.

6. Conclusion

The nature of long-tailed tasks determines that no single model works best for all tasks, which makes model selection in a model zoo with diverse inductive biases necessary. Instead of manually designing MS criteria, learning the relationship between tasks and models via recommendation models can be more efficient, effective and scalable to new meta features and models, and it can be continuously improved with the growing volume of training history. This makes the framework applicable to other selection problems as well such as selecting optimal models and hyperparameters at the same time.

References

[1] Alessandro Achille, Michael Lam, Rahul Tewari, Avinash Ravichandran, Subhransu Maji, Charless C Fowlkes, Stefano Soatto, and Pietro Perona. Task2vec: Task embedding for meta-learning. In *ICCV*, 2019. 1

[2] Yajie Bao, Yang Li, Shao-Lun Huang, Lin Zhang, Amir R Zamir, and Leonidas J Guibas. An information-theoretic metric of transferability for task transfer learning. In *ICIP*, 2019. 8

[3] Björn Barz, Kai Schröter, Moritz Münch, Bin Yang, Andrea Unger, Doris Dransch, and Joachim Denzler. Enhancing flood impact analysis using interactive retrieval of social media images. *arXiv preprint arXiv:1908.03361*, 2019. 11

[4] Sean Bell, Paul Upchurch, Noah Snavely, and Kavita Bala. Material recognition in the wild with the materials in context database. 2015. 11

[5] Daniel Bolya, Rohit Mittapalli, and Judy Hoffman. Scalable diverse model selection for accessible transfer learning. In *NeurIPS*, 2021. 2, 3, 4, 6, 7, 8

[6] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – mining discriminative components with random forests. In *ECCV*, 2014. 11, 16

[7] Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, 2017. 11

[8] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, , and A. Vedaldi. Describing textures in the wild. In *CVPR*, 2014. 11

[9] Yin Cui, Yang Song, Chen Sun, Andrew Howard, and Serge Belongie. Large scale fine-grained categorization and domain-specific transfer learning. In *CVPR*, 2018. 8

[10] Aditya Deshpande, Alessandro Achille, Avinash Ravichandran, Hao Li, Luca Zancato, Charless Fowlkes, Rahul Bhotika, Stefano Soatto, and Pietro Perona. A linearized framework and a new benchmark for model selection for fine-tuning. *arXiv preprint arXiv:2102.00084*, 2021. 1, 2, 3, 4, 5, 6, 7, 8

[11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 6, 12

[12] Kshitij Dwivedi and Gemma Roig. Representation similarity analysis for efficient task taxonomy & transfer learning. In *CVPR*, 2019. 1

[13] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. *CVPR Workshop*, 2004. 11

[14] Nicolo Fusi, Rishit Sheth, and Melih Elibol. Probabilistic matrix factorization for automated machine learning. *NeurIPS*, 31, 2018. 8

[15] Noa Garcia and George Vogiatzis. How to read paintings: semantic art understanding with multi-modal retrieval. In *ECCV Workshops*, 2018. 11

[16] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013. 11

[17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6, 12

[18] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2019. 11

[19] Saihui Hou, Yushan Feng, and Zilei Wang. Vegfru: A domain-specific dataset for fine-grained visual categorization. In *ICCV*, 2017. 11

[20] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 6, 12

[21] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, 2017. 6, 12

[22] Long-Kai Huang, Junzhou Huang, Yu Rong, Qiang Yang, and Ying Wei. Frustratingly easy transferability estimation. In *ICML*, 2022. 8

[23] Yibin Huang, Congying Qiu, and Kui Yuan. Surface defect saliency of magnetic tile. *The Visual Computer*, 36(1):85–96, 2020. 11

[24] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*, 2017. 11

[25] Alexis Joly and Olivier Buisson. Logo retrieval with a contrario visual query expansion. In *ACM MM*, pages 581–584, 2009. 11

[26] Yuchin Juan, Yong Zhuang, Wei-Sheng Chin, and Chih-Jen Lin. Field-aware factorization machines for ctr prediction. In *Proceedings of the 10th ACM conference on recommender systems*, 2016. 8

[27] Kaggle and EyePacs. Kaggle diabetic retinopathy detection, jul 2015. 11

[28] Parneet Kaur, , Karan Sikka, Weijun Wang, serge Belongie, and Ajay Divakaran. Foodx-251: A dataset for fine-grained food classification. *arXiv preprint arXiv:1907.06167*, 2019. 11

[29] Daniel S Kermany, Michael Goldbaum, Wenjia Cai, Carolina CS Valentim, Huiying Liang, Sally L Baxter, Alex McKeown, Ge Yang, Xiaokang Wu, Fangbing Yan, et al. Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell*, 172(5):1122–1131, 2018. 11

[30] Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Li Fei-Fei. Novel dataset for fine-grained image categorization. In *First Workshop on Fine-Grained Visual Categorization, CVPR*, 2011. 3, 11

[31] Simon Kornblith, Jonathon Shlens, and Quoc V Le. Do better imagenet models transfer better? In *CVPR*, 2019. 3, 5

[32] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013. 11

[33] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 11

[34] Yann LeCun, Fu Jie Huang, and Léon Bottou. Learning methods for generic object recognition with invariance to pose and lighting. In *CVPR*, 2004. 11

[35] Yandong Li, Xuhui Jia, Ruoxin Sang, Yukun Zhu, Bradley Green, Liqiang Wang, and Boqing Gong. Ranking neural checkpoints. In *CVPR*, 2021. 8

[36] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 6, 12, 17

[37] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *arXiv preprint arXiv:2201.03545*, 2022. 6, 12

[38] S. Maji, J. Kannala, E. Rahtu, M. Blaschko, and A. Vedaldi. Fine-grained visual classification of aircraft. Technical report, 2013. 11

[39] Loic Matthey, Irina Higgins, Demis Hassabis, and Alexander Lerchner. dsprites: Disentanglement testing sprites dataset. https://github.com/deepmind/dsprites-dataset/, 2017. 11

[40] Gaurav Mittal, Chang Liu, Nikolaos Karianakis, Victor Fragoso, Mei Chen, and Yun Fu. Hyperstar: Task-aware hyperparameters for deep networks. In *CVPR*, 2020. 8

[41] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011. 11

[42] Cuong V Nguyen, Tal Hassner, Cedric Archambeau, and Matthias Seeger. Leep: A new measure to evaluate transferability of learned representations. *arXiv preprint arXiv:2002.12462*, 2020. 8

[43] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics & Image Processing*. IEEE, 2008. 11

[44] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *CVPR*, 2012. 11

[45] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *ICCV*, 2019. 5, 11

[46] Joan Puigcerver, Carlos Riquelme, Basil Mustafa, Cedric Renggli, André Susano Pinto, Sylvain Gelly, Daniel Keysers, and Neil Houlsby. Scalable transfer learning with expert models. In *ICLR*, 2021. 4, 8

[47] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 1

[48] Sivaramakrishnan Rajaraman, Sameer K Antani, Mahdieh Poostchi, Kamolrat Silamut, Md A Hossain, Richard J Maude, Stefan Jaeger, and George R Thoma. Pre-trained convolutional neural networks as feature extractors toward improved malaria parasite detection in thin blood smear images. *PeerJ*, 6:e4568, 2018. 11

[49] Steffen Rendle. Factorization machines. In *ICDM*, 2010. 4

[50] Ali Sharif Razavian, Hossein Azizpour, Josephine Sullivan, and Stefan Carlsson. Cnn features off-the-shelf: an astounding baseline for recognition. In *CVPR workshops*, pages 806–813, 2014. 11

[51] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015. 12

[52] Kiat Chuan Tan, Yulong Liu, Barbara Ambrose, Melissa Tulig, and Serge Belongie. The herbarium challenge 2019 dataset. *arXiv preprint arXiv:1906.05372*, 2019. 11

[53] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019. 6, 12, 16

[54] Anh T Tran, Cuong V Nguyen, and Tal Hassner. Transferability and hardness of supervised classification tasks. In *ICCV*, 2019. 8

[55] Bastiaan S Veeling, Jasper Linmans, Jim Winkens, Taco Cohen, and Max Welling. Rotation equivariant CNNs for digital pathology. June 2018. 11

[56] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. 11, 16

[57] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. *arXiv preprint arXiv:2202.03052*, 2022. 1

[58] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*, pages 3485–3492, June 2010. 11

[59] Xi Yan, David Acuna, and Sanja Fidler. Neural data server: A large-scale search engine for transfer learning data. In *CVPR*, 2020. 8

[60] Kaichao You, Yong Liu, Jianmin Wang, and Mingsheng Long. Logme: Practical assessment of pre-trained models for transfer learning. In *ICML*, 2021. 2, 3, 4, 6, 7, 8

[61] Xiaohua Zhai, Joan Puigcerver, Alexander Kolesnikov, Pierre Ruyssen, Carlos Riquelme, Mario Lucic, Josip Djolonga, Andre Susano Pinto, Maxim Neumann, Alexey Dosovitskiy, et al. A large-scale study of representation learning with the visual task adaptation benchmark. *arXiv preprint arXiv:1910.04867*, 2019. 5, 11, 17

[62] Chongzhi Zhang, Mingyuan Zhang, Shanghang Zhang, Daisheng Jin, Qiang Zhou, Zhongang Cai, Haiyu Zhao, Shuai Yi, Xianglong Liu, and Ziwei Liu. Delving deep into the generalization of vision transformers under distribution shifts. *arXiv preprint arXiv:2106.07617*, 2021. 17