
Head2Toe: Utilizing Intermediate Representations for Better Transfer Learning

Utku Evci¹ Vincent Dumoulin¹ Hugo Larochelle¹ Michael C. Mozer¹

Abstract

Transfer-learning methods aim to improve performance in a data-scarce target domain using a model pretrained on a data-rich source domain. A cost-efficient strategy, *linear probing*, involves freezing the source model and training a new classification head for the target domain. This strategy is outperformed by a more costly but state-of-the-art method—*fine-tuning* all parameters of the source model to the target domain—possibly because fine-tuning allows the model to leverage useful information from intermediate layers which is otherwise discarded by the previously trained later layers. We explore the hypothesis that these intermediate layers might be directly exploited. We propose a method, *Head-to-Toe probing* (HEAD2TOE), that selects features from all layers of the source model to train a classification head for the target domain. In evaluations on the **Visual Task Adaptation Benchmark (VTAB)**, Head2Toe matches performance obtained with fine-tuning on average while reducing training and storage cost a hundred fold or more, but critically, for out-of-distribution transfer, Head2Toe outperforms fine-tuning¹.

1. Introduction

Transfer learning is a widely used method for obtaining strong performance in a variety of tasks where training data is scarce (e.g., Zhu et al., 2020; Alyafeai et al., 2020; Zhuang et al., 2020). A well-known recipe for transfer learning involves the supervised or unsupervised pretraining of a model on a *source* task with a large training dataset (also referred to as *upstream training*). After pretraining,

¹Google Research, Brain Team. Correspondence to: Utku Evci <evcu@google.com>.

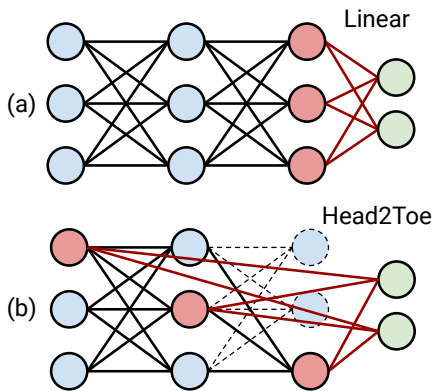


Figure 1. (a) Whereas **LINEAR** utilizes only the last layer for transfer learning, (b) **HEAD2TOE** selects the most useful features from the entire network and trains a linear head on top.

the model’s output head is discarded, and the rest of the network is used to obtain a feature embedding, i.e., the output of what was formerly the penultimate layer of the network. When transferring to a *target* task, a new output head is trained on top of the feature extractor (*downstream training*). This approach makes intuitive sense: if a linear combination of embedding features performs well on the source task and the source and target domains are similar, we would expect a different linear combination of features to generalize to the target domain.

This approach of training a new output head, which we refer to as **LINEAR**, often yields significant improvements in performance on the target task over training the network from scratch (Kornblith et al., 2019). An alternative to **LINEAR** is *fine-tuning* (**FINETUNING**), which uses target-domain data to adapt all weights in the feature extractor together with the new output head. This procedure requires running forward and backward passes through the entire network at each training step and therefore its per-step cost is significantly higher than **LINEAR**. Furthermore, since the entire network is fine-tuned, the entire set of new weights needs to be stored for every target task, making **FINETUNING** impractical when working on edge devices or with a large number of target tasks. However, **FINETUNING** is often preferred over **LINEAR** since it consistently leads to better

performance on a variety of target tasks even when data is scarce (Zhai et al., 2019).

FINETUNING's superior generalization in the low-data regime is counterintuitive given that the number of model parameters to be adapted is often large relative to the amount of available training data. How does **FINETUNING** learn from few examples successfully? We conjecture that **FINETUNING** better leverages existing internal representations rather than discovering entirely new representations; **FINETUNING** exposes existing features buried deep in the net for use by the classifier. Under this hypothesis, *features needed for transfer are already present in the pretrained network and might be identified directly without fine-tuning the backbone itself*. In Section 3.1, we argue that **FINETUNING** can be approximated by a linear probe operating on the intermediate features of a network, thus enabling state-of-the-art transfer performance with significantly less cost.

In this work, we propose and explore methods for selecting useful features from *all* layers of a pretrained net, including the embedding, and then applying the **LINEAR** transfer approach to the constructed representation. We compare the standard approach (Figure 1a) to our approach, called **HEAD2TOE** (Figure 1b). **HEAD2TOE** shows significant improvements over **LINEAR** (Figure 2) and matches **FINETUNING** performance on average. Our key contributions are as follows:

1. We observe a strong correlation between the degree to which a target domain is **out-of-distribution** (OOD) with respect to the source domain and the benefit of incorporating intermediate representations in **LINEAR** (Figure 2 and Section 3.2), corroborating observations made in Adler et al. (2020).
2. We introduce **HEAD2TOE**, an efficient transfer learning method for selecting relevant features from intermediate representations (Section 3.3).
3. On the **VTAB** collection of data sets, we show that **HEAD2TOE** outperforms **LINEAR** and matches the performance of the more computationally costly **FINETUNING** with only 0.6% of the training FLOPs and 1% of the storage cost (Section 4).
4. Critically, **HEAD2TOE** outperforms **FINETUNING** on **OOD** target domains. If a practitioner can make an educated guess about whether a target domain is **OOD** with respect to a source, using **HEAD2TOE** improves on the state-of-the-art for transfer learning.

2. Preliminaries

Source domain and backbone models. In our experiments, we use source models pretrained on ImageNet-2012 (Russakovsky et al., 2015), a large scale image classification benchmark with 1000 classes and over 1M natural

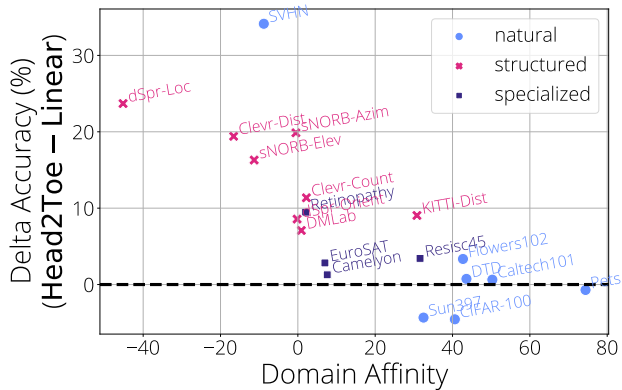


Figure 2. Improvement in accuracy due to **HEAD2TOE** over **LINEAR** for the **VTAB** target domains, arranged on the abscissa by domain affinity (Section 2), proxy for OOD-ness of the target relative to the source domain (ImageNet). The more OOD domains (left side of the plot) benefit most from **HEAD2TOE**.

images. We benchmark **HEAD2TOE** using convolutional (ResNet-50, Wu et al., 2018) and attention-based (ViT-B/16, Dosovitskiy et al., 2021) architectures pretrained on ImageNet-2012.

Target domains. In this work, we focus on target tasks with few examples (i.e., *few-shot*) and use **Visual Task Adaptation Benchmark-1k** (Zhai et al., 2019) to evaluate different methods. **VTAB-1k** consists of 19 different classification tasks, each having between 2 to 397 classes and a total of 1000 training examples. The domains are grouped into three primary categories: (1) natural images (*natural*), (2) specialized images using non-standard cameras (*specialized*), and (3) rendered artificial images (*structured*).

Characterizing out-of-distribution (far) domains. Adler et al. (2020) use the difference in Fréchet inception distance (FID) between two domains to characterize how far OOD domains are. However, we need a metric which reflects not only changes in $p(x)$ —as FID does—but also changes in the task $p(y | x)$ itself. Consider the relationship between source and target domains. If the domain overlap is high, then features extracted for linear classification in the source domain should also be relevant for linear classification in the target domain, and **LINEAR** should yield performance benefits. If the domain overlap is low, then constraining the target domain to use the source domain embedding may be harmful relative to training a model from scratch on the target domain. Therefore, we might quantify the source-target distribution shift in terms of how beneficial **LINEAR** is relative to training a model from scratch (denoted **SCRATCH**):

$$\text{DomainAffinity} = \text{ACC}_{\text{LINEAR}} - \text{ACC}_{\text{SCRATCH}}.$$

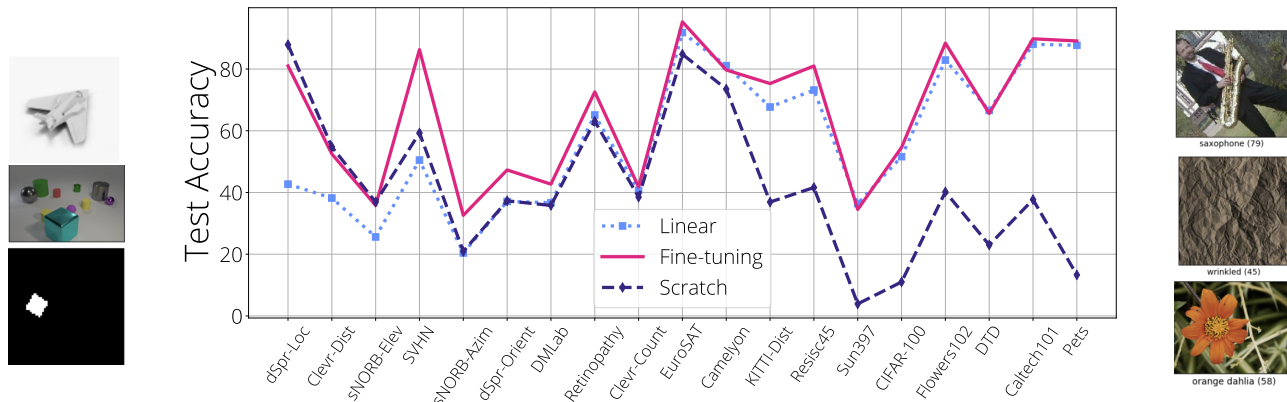


Figure 3. **Characterizing out-of-distribution (far) domains** Generalization performance of various baselines on the VTAB-1k benchmark using ResNet-50 architecture and 240 image size. The architecture is pretrained on ImageNet-2012 for the TL baselines. Datasets (and the three groups of the benchmark) are ordered according to their Domain Affinity scores in ascending order from left to right. Examples from the left- and right-most datasets are also shown on corresponding sides.

In Figures 2 and 3, the 19 VTAB-1k target tasks are arranged from low to high by their domain affinity to ImageNet-2012 for a pretrained ResNet-50 backbone. The left and right ends of Figure 3 show examples of the three target domains with the least and most similar distributions, respectively. These examples seem consistent with intuitive notions of distribution shift.

The domain affinity calculated from other pretrained backbones provides a similar ordering. In Appendix F, we show that calculating domain affinity using the ViT-B/16 backbone obtains a high correlation (Spearman, 0.907) with the original scores calculated using the ResNet-50 backbone. Furthermore, we did additional investigations using the 15 representation learning methods presented on the VTAB-leaderboard², where we calculated median percentage-improvement over scratch training for each task, and similarly observed a high Spearman correlation (0.803).

Baselines. Figure 3 also presents transfer test accuracy of LINEAR, FINETUNING, and SCRATCH baselines. Consistent with the literature, FINETUNING performs as well or better than the other two baselines. For in-distribution targets (right side of graph), LINEAR matches FINETUNING; for OOD targets (left side of graph), LINEAR is worse than FINETUNING. With distribution mismatch, the source network may filter out information available in lower layers because it is not needed for the source task but is essential for the target task. Observing FINETUNING performs better than SCRATCH even in OOD tasks, we hypothesize that intermediate features are key for FINETUNING, since if learning novel features was possible using limited data, training the network from scratch would work on-par with FINETUNING. Motivated by this observation, HEAD2TOE

probes the intermediate features of a network directly and aims to match the fine-tuning results without modifying the backbone itself.

3. HEAD2TOE Transfer of Pretrained Models

3.1. Taylor Approximation of Fine-Tuning

Maddox et al. (2021) and Mu et al. (2020) observe that the parameters of a pre-trained backbone change very little during fine-tuning and that a linearized approximation of the fine-tuned model obtains competitive results in transfer learning. Following a similar motivation, we argue that the linearized fine-tuning solution should be well captured by a linear combination of the intermediate activations.

To demonstrate our intuition, consider a multi-layer, fully-connected neural network with input x and scalar output $F(x; w)$, parameterized by weights w . We denote the individual elements of w by w_{ij} , where the indices reflect the neurons connected such that the activations of neurons are given by $z_j = \sum_i w_{ij} h_i$ and $h_i = f(z_i)$ where f is the activation function used in the network. Then, we can write the fine-tuned neural network, parameterized by optimized parameters w^* , using the first-order Taylor approximation:

$$F(x; w^*) \approx F(x; w) + \sum_{i,j} \frac{\partial F(x; w)}{\partial w_{ij}} \Delta w_{ij}$$

where $\Delta w_{ij} = w_{ij}^* - w_{ij}$ reflects the displacement of updated weights during fine-tuning. Expanding the gradient term using the chain rule and rearranging the summations, the linearized solution found by FINETUNING can be writ-

²https://google-research.github.io/task_adaptation/benchmark

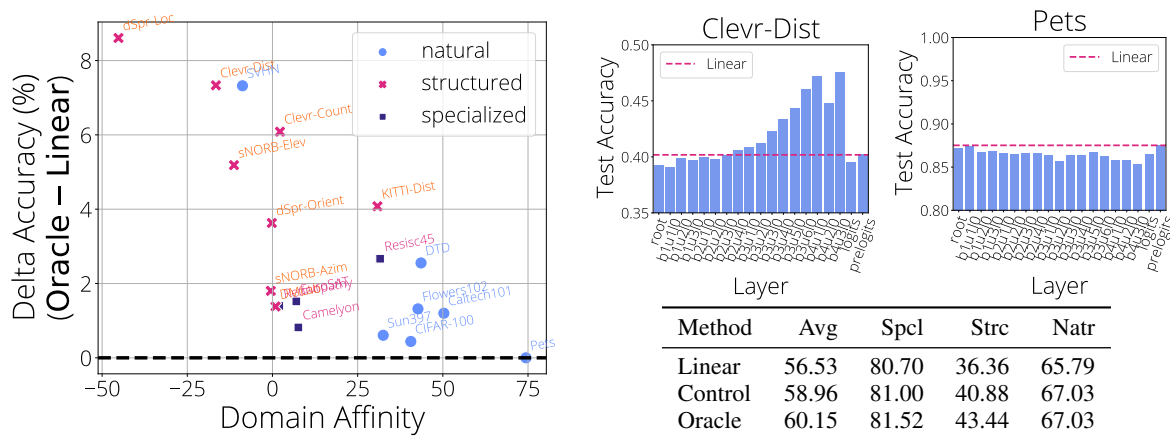


Figure 4. (left) Accuracy gains when prelogits are augmented with an additional layer correlates negatively (-0.745, Spearman) with domain affinity. (right-top) effect of using features from intermediate layers for Clevr-Dist (low domain affinity) and Pets (high domain affinity) tasks (right-bottom) Test accuracies of various baselines on VTAB-1k. *Linear* uses only prelogits, *Oracle* averages are obtained by using the layer that gives best generalization (test) for each task. *Control* experiment uses a second feature embedding from a second pretrained network trained using a different seed. We use ResNet-50 models pretrained on ImageNet-2012.

ten as a linear combination of the intermediate activations:

$$\begin{aligned}
 F(x; w^*) &\approx F(x; w) + \sum_{i,j} h_i \frac{\partial F(x; w)}{\partial z_j} \Delta w_{ij} \\
 &\approx F(x; w) + \sum_i h_i \sum_j \frac{\partial F(x; w)}{\partial z_j} \Delta w_{ij} \\
 &\approx F(x; w) + \sum_i h_i c_{i,x}
 \end{aligned}$$

Thus, so long as fine-tuning produces small displacements Δw_{ij} , the **FINETUNING** solution for a given input x can be approximated by a linear combination of all features in the network. More broadly, if the coefficients of the most relevant features are robust to input x , the **FINETUNING** solution has an approximate equivalence to **LINEAR** when trained on features selected from all layers of the network. This conclusion supports **HEAD2TOE**'s use of intermediate activations as a bridge between **FINETUNING** and **LINEAR**.

3.2. Your Representations are Richer Than You Think

In this section, we conduct a simple experiment to demonstrate the potential of using representations from intermediate layers. We concatenate the feature embedding of a pretrained ResNet-50 backbone (features from the penultimate layer) with features from *one* additional layer and train a linear head on top of the concatenated features. When extracting features from convolutional layers, we reduce the dimensionality of the convolutional stack of feature maps using strided average pooling, with a window size chosen so that the resulting number of features is similar to the number of embedding features (2048 for ResNet-50).

To estimate an upper bound on performance improvement

over **LINEAR** by including a single intermediate layer, we use an oracle to select the layer which yields the largest boost in test performance, separately for each target task. Percentage improvement over **LINEAR** using this **ORACLE** is shown in Figure 4-left. We observe a Spearman correlation of -0.75 between the domain affinity of a target task and the accuracy gain. In accordance with our hypothesis, adding intermediate representations does not improve in-domain generalization because the feature embedding already contains the most useful features. In contrast, generalization on out-of-domain tasks are improved significantly when intermediate features are used.

In Figure 4-top-right, we show test accuracy for two domains as a function of the ResNet-50 layer whose internal representation is used to augment the feature embeddings. Figures for the remaining tasks can be found in Appendix G. Different tasks on the VTAB-1k benchmark benefit from the inclusion of different layers to obtain the best generalization accuracy, which emphasizes the importance of domain-specific selection of the appropriate set of features for optimal generalization. Overall, the **ORACLE** that selects the layer with best test performance for each task yields an average of 3.5% improvement on the VTAB-1k benchmark. One possible explanation for the improvement in performance with the augmented representation is simply that it has more degrees of freedom (4096 features instead of 2048). To demonstrate that the improvement is due to inclusion of intermediate layers and not simply due to increased dimensionality, Figure 4-bottom-right compares the **ORACLE** to a **CONTROL** condition whose representation is matched in dimensionality but formed by concatenating a feature embedding obtained from a second ResNet-50 backbone pretrained on ImageNet-2012. Note that this ex-

periment bears similarity to *ensembling* (Zhou et al., 2002), which is known to bring significant accuracy gains on its own (Mustafa et al., 2020). Using a second backbone doubles the amount of resources required yet falls 1% shy of ORACLE performance, showing the extent to which intermediate representations can be helpful for generalization.

3.3. Head2Toe

Motivated by our observations in the previous section, we hypothesize that we can attain—or possibly surpass—the performance of FINE TUNING without modifying the backbone itself by using LINEAR augmented with well-chosen intermediate activations. Our investigation leads us to HEAD2TOE, an efficient transfer learning algorithm based on utilizing intermediate activations using feature selection.

Notation. Our method applies to any network with any type of layers, but here, for simplicity and without loss of generality, we consider a network with L fully connected layers, each layer receiving input from the layer below:

$$\mathbf{z}_\ell = \mathbf{h}_{\ell-1} \mathbf{W}_\ell \quad ; \quad \mathbf{h}_\ell = f(\mathbf{z}_\ell) \quad (1)$$

where the subscript denotes a layer index, $\mathbf{h}_0 = \mathbf{x}$ is the input, f is the activation function, \mathbf{W}_ℓ is the weight matrix of layer ℓ , and \mathbf{z}_L is the logit vector used for classification.

When transferring a pretrained network to a target task using LINEAR, we discard the last layer of the pretrained network and train a new set of linear weights, \mathbf{W}'_L , such that predictions (logits) for the new task are obtained by $\mathbf{z}'_L = \mathbf{h}_{L-1} \mathbf{W}'_L$.

Head2Toe. Consider a simple scheme that augments the backbone embedding with activations from all layers of the network, such that:

$$\mathbf{z}'_L = \mathbf{h}_{all} \mathbf{W}_{all} \quad ; \quad \mathbf{h}_{all} = [a_1(\mathbf{h}_1), \dots, a_L(\mathbf{h}_L)] \quad (2)$$

where $a_\ell(\cdot)$ denotes a fixed function to reduce the dimensionality of the activation vector and normalize at a given layer ℓ . Such functions are valuable for network architectures like convolutional networks that generate many intermediate features. Though better aggregation schemes may exist, we simply perform one- or two-dimensional strided average pooling to reduce dimensionality. After aggregation, we normalize features coming from each layer to a unit norm. This scaling preserves the relative magnitude of features within a layer while accounting for inter-layer differences. It works better than normalizing each feature separately or not normalizing at all.

Even with dimensionality reduction, \mathbf{h}_{all} can exceed a million elements, and \mathbf{W}_{all} is underconstrained by the training data, leading to overfitting. Further, \mathbf{W}_{all} may become so

large as to be impractical for deploying this model.³ We can address these issues by selecting a subset of features before training the target-domain classifier.

Feature selection based on group lasso. Group lasso (Yuan & Lin, 2006) is a popular method for selecting relevant features in multi-task adaptation settings (Argyriou et al., 2007; Nie et al., 2010). When used as a regularizer on a weight matrix \mathbf{W} , the group-lasso regularizer encourages the ℓ_2 norm of the rows of the matrix to be sparse, and is defined as:

$$|\mathbf{W}|_{2,1} = |\mathbf{s}|_1 = \sum_i |s_i| \quad ; \quad s_i = \sqrt{\sum_j w_{ij}^2} \quad . \quad (3)$$

To determine which features are most useful for the task, the linear head is trained with group-lasso regularization on \mathbf{W}_{all} . In contrast to the approach taken by Argyriou et al. (2007) and Nie et al. (2010), which use costly matrix inversions, we incorporate the regularization as a secondary loss and train with stochastic gradient descent. Following training, a *relevance score* s_i is computed for each feature i . We select a fraction F of the features with the largest relevance and train a new linear head to obtain the final logit mapping. Feature selection alone provides strong regularization, therefore during the final training we do not use any additional regularization.

We make two remarks here. First, because the initial round of training \mathbf{W}_{all} with the group-lasso regularizer is used only to rank features by importance, the method is robust to the regularization coefficient; it simply needs to be large enough to distinguish the contributions of individual features. Second, interpreting s_i as the importance of feature i depends on all features having the same dynamic range. This constraint is often satisfied naturally due to the normalization done after aggregation as explained previously.

Selecting F . The fraction F determines the total number of features retained. One would expect the optimal value to depend on the target task. Therefore, we select F for each task separately by cross-validation on the training set. This validation procedure is inexpensive compared to the cost of the initial phase of the algorithm (i.e., training of \mathbf{W}_{all} to obtain \mathbf{s}) due to the reduced number of features in the second step. Overall, HEAD2TOE together with its validation procedure requires 18% more operations compared to training \mathbf{W}_{all} alone (details shared in Appendix B).

Cost of HEAD2TOE. HEAD2TOE’s use of a fixed backbone means that as we search for features to include, the

³For example, using a pooling size of 2, ResNet-50 generates 1.7 million features and storing \mathbf{W}_{all} requires 6.6×10^8 parameters (2.6GB for float32) for SUN-397.

actual features values are fixed. Consequently, we can calculate them once and re-use as necessary, instead of recalculating at every step, as required for FINE-TUNING. Furthermore, since the backbone is frozen, the storage required for each target task includes only the final output head and the binary mask indicating the features selected. Due to these properties, the cost of HEAD2TOE follows the cost of LINEAR closely while being significantly less than FINE-TUNING.

4. Evaluating Head2Toe

We evaluate HEAD2TOE on the VTAB-1k benchmark using two popular vision architectures, ResNet-50 (Wu et al., 2018) and ViT-B/16 (Dosovitskiy et al., 2021), both pretrained on ImageNet-2012. ResNet-50 consists of 50 convolutional layers. To utilize the intermediate convolutional features, we aggregate spatial information using average pooling on non-overlapping patches, as explained in Section 3.3. We adjust the pooling size to target a fixed dimensionality of the representation. For example, for a feature map of shape $20 \times 20 \times 128$ and a *target size* of 512, we average-pool disjoint patches of size 5×5 , resulting in 4 features per channel and 1024 features in total. This helps us to balance different layers in terms of the number features they contribute to the concatenated embedding. ViT-B/16 consists of 12 multi-headed self-attention layers. When we aggregate the output of the self-attention layers, we perform 1-D average pooling over the patch/token dimension choosing the pooling size to match a target number of features as before. Given that token dimension is permutation invariant, 1-D average pooling is unlikely to be the best choice here and a better aggregation function should provide further gains. We share a detailed list of intermediate representations utilized for each architecture in Appendix C.

HEAD2TOE selects a subset of features and trains a linear classifier without regularization on top of the selected features. We compare HEAD2TOE with regularization baselines that utilize all features. These baselines are denoted as +All- ℓ_1 , +All- ℓ_2 and +All- $\ell_{2,1}$ according to the regularizer norm they use.

We perform five-fold cross validation for each task and method in order to pick the best hyperparameters. All methods search over the same learning rates and training steps (two values of each). Methods that leverage intermediate features (i.e., regularization baselines and HEAD2TOE) additionally search over regularization coefficients and the *target size* of the aggregated representation at each layer. The FINE-TUNING baseline searches over 4 hyperparameters; thus the comparison of HEAD2TOE, which searches over 24 values, to fine-tuning might seem unfair. However, this was necessary due to fine-tuning being significantly more costly (about 200x on average) than training a linear classifier on

intermediate features. We repeat each evaluation using 3 different seeds and report median values and share standard deviations in Appendix D. More details on hyperparameter selection and values used are shared in Appendix A.

4.1. ResNet-50

The top half of Table 1 presents results on the 19 VTAB-1k target domains when transferring from a pretrained ResNet-50 architecture. On average, HEAD2TOE slightly outperforms all other methods, including FINE-TUNING (see right-most column of Table 1). HEAD2TOE, along with the regularization baselines that use intermediate layers, is far superior to LINEAR, indicating the value of the intermediate layers. And HEAD2TOE is superior to the regularization baselines, indicating the value of explicit feature selection. Among the three categories of tasks, HEAD2TOE excels relative to the other methods for the *specialized* category, but does not outperform FINE-TUNING for the *natural* and *structured* categories.

Does HEAD2TOE select different features for each task? Which layers are used more frequently? In Appendix I, we show the distribution of features selected across different layers and the amount of intersection among features selected for different tasks and seeds. We observe high variation across tasks, motivating the importance of performing feature selection for each task separately. In addition to the fact that HEAD2TOE outperforms FINE-TUNING, it requires only 0.5% of FLOPs during training on average. Similarly, the cost of storing the adapted model is reduced to 1% on average. We discuss HEAD2TOE’s computational and storage costs in detail in Appendix B.

4.2. ViT-B/16

Results for ViT-B/16 are shared in the bottom half of Table 1. As with the ResNet-50 architecture, HEAD2TOE achieves the best accuracy among methods that keep the backbone fixed: HEAD2TOE improves accuracy over LINEAR by about 10% on average (in absolute performance), and HEAD2TOE outperforms the regularization baselines that include intermediate features but that do not explicitly select features. Similarly to the ResNet-50 experiments, HEAD2TOE matches the performance of FINE-TUNING. We share the distribution of features selected over layers in Appendix J.

HEAD2TOE-FT. Our goal in this work is to show that state-of-the-art performance can be obtained efficiently *without* changing the backbone itself. However, we expect to see further gains if the backbone is unfrozen and trained together with the final set of selected features. We performed some initial experiments to demonstrate the potential of such an approach. After selecting features with

Head2Toe: Utilizing Intermediate Representations for Better Transfer Learning

	Natural							Specialized				Structured							Mean	
	CIFAR-100	Caltech101	DTD	Flowers102	Pets	SVHN	Sun397	Camelyon	EuroSAT	Resisc45	Retinopathy	Clevr-Count	Clevr-Dist	DMLab	KITTI-Dist	dSpr-Loc	dSpr-Ori	sNORB-Azim		sNORB-Elev
ResNet-50 backbone																				
Linear	48.5	86.0	67.8	84.8	87.4	47.5	34.4	83.2	92.4	73.3	73.6	39.7	39.9	36.0	66.4	40.4	37.0	19.6	25.5	57.0
+All- ℓ_2	44.7	87.0	67.8	84.2	86.1	81.1	31.9	82.6	95.0	76.5	74.5	50.0	56.3	38.3	65.5	59.7	44.5	37.5	40.0	63.3
+All- ℓ_1	50.8	88.6	67.4	84.2	87.7	84.2	34.6	80.9	94.9	75.6	74.7	49.9	57.0	41.8	72.9	59.0	44.8	37.5	40.8	64.6
+All- $\ell_{2,1}$	49.1	86.7	68.5	84.2	88.0	84.4	34.8	81.5	94.9	75.7	74.3	48.3	58.4	42.0	74.4	58.8	45.2	37.8	34.4	64.3
Head2Toe	47.1	88.8	67.6	85.6	87.6	84.1	32.9	82.1	94.3	76.0	74.1	55.3	59.5	43.9	72.3	64.9	51.1	39.6	43.1	65.8
Scratch*	11.0	37.7	23.0	40.2	13.3	59.3	3.9	73.5	84.8	41.6	63.1	38.5	54.8	35.8	36.9	87.9	37.3	20.9	36.9	42.1
Fine-tuning	33.2	84.6	54.5	85.2	79.1	87.8	16.6	82.0	92.5	73.3	73.5	54.6	63.7	46.3	72.1	94.8	47.1	35.0	33.3	63.6
Head2Toe-FT	16.3	87.7	63.1	84.3	66.9	82.5	24.3	82.6	11.0	76.7	73.5	54.8	69.1	44.7	69.2	94.2	51.0	33.3	44.4	59.5
Head2Toe-FT+	46.9	88.9	66.6	84.0	87.3	84.4	32.4	84.2	94.4	76.7	74.1	55.8	69.1	45.3	74.7	94.4	51.0	39.7	42.6	68.0
ViT-B/16 backbone																				
Linear	55.0	81.0	53.6	72.1	85.3	38.7	32.3	80.1	90.8	67.2	74.0	38.5	36.2	33.5	55.7	34.0	31.3	18.2	26.3	52.8
+All- ℓ_2	57.3	87.0	64.3	82.8	84.0	75.7	32.4	82.0	94.7	79.7	74.8	47.4	57.8	41.4	62.8	46.6	33.3	31.0	38.8	61.8
+All- ℓ_1	58.4	87.3	64.9	83.3	84.6	80.0	34.4	82.3	95.6	79.6	73.6	47.9	57.7	42.2	65.1	44.5	33.4	32.4	38.4	62.4
+All (Group)	59.6	87.1	64.9	85.2	85.4	79.5	35.3	82.0	95.3	80.6	74.2	47.9	57.8	40.7	64.9	46.7	33.6	31.9	39.0	62.7
Head2Toe	58.2	87.3	64.5	85.9	85.4	82.9	35.1	81.2	95.0	79.9	74.1	49.3	58.4	41.6	64.4	53.3	32.9	33.5	39.4	63.3
Scratch	7.6	19.1	13.1	29.6	6.7	19.4	2.3	71.0	71.0	29.3	72.0	31.6	52.5	27.2	39.1	66.1	29.7	11.7	24.1	32.8
Fine-tuning	44.3	84.5	54.1	84.7	74.7	87.2	26.9	85.3	95.0	76.0	70.4	71.5	60.5	46.9	72.9	74.5	38.7	28.5	23.8	63.2
Head2Toe-FT	43.9	82.3	53.5	84.9	76.7	86.5	24.5	79.9	95.9	77.5	74.3	68.0	70.9	48.2	72.4	76.1	44.8	32.1	42.5	65.0
Head2Toe-FT+	57.3	87.1	63.8	83.7	84.8	86.8	35.1	80.2	96.1	79.9	74.1	69.9	71.2	47.8	72.8	77.4	45.9	33.9	43.0	67.9

Table 1. Median test accuracy over 3 seeds on the VTAB-1k benchmark using pretrained backbones. Regularization baselines that use all layers are indicated with the +All prefix. "*" indicates results obtained from Zhai et al. (2019). Fine-tuning results for ViT-B/16 are obtained using the checkpoints provided by Dosovitskiy et al. (2021).

HEAD2TOE, we fine-tuned the backbone together with the final output layer (**HEAD2TOE-FT**) and observed around 2% increase in accuracy for ViT-B/16. We use the same lightweight validation procedure used for **FINETUNING** to pick the learning rate and training steps for the final fine-tuning steps. **HEAD2TOE-FT** provides significant gains over **HEAD2TOE** in the *structured* category, but performs poorly when transferring to some of the *natural* category tasks. Next, we determined whether or not to fine-tune the backbone by examining validation-set accuracy. This procedure, denoted **HEAD2TOE-FT+**, provides an additional 3% improvement and results in 67.9% accuracy over all VTAB-1k tasks, without any additional training or storage costs compared to **FINETUNING**⁴.

4.3. Understanding **HEAD2TOE**

HEAD2TOE selects individual features from a pre-trained backbone for each task separately and ignores the layer structure of the features. In this section we investigate different parts of the **HEAD2TOE** algorithm and compare them with some alternatives. Appendix I and Appendix J include further experiments on the effect of the support set size, choice of intermediate activations and the effectiveness of

⁴**HEAD2TOE-FT+** is not a simple max of **HEAD2TOE-FT** and **HEAD2TOE**. This is due to the variance during the re-runs and the mismatch between validation and test accuracies.

the relevance scores.

Selecting features or selecting layers? **HEAD2TOE** selects individual features independent of the layer in which they are embedded. We compare this *feature-wise* strategy to selecting layers as whole (i.e., selecting all features in a layer or none). One might expect *layer-wise* selection to be superior because it involves fewer decisions and therefore less opportunity to overfit to the validation set used for selecting the fraction F . Further, layer-wise selection may be superior because the relevance of one feature in a layer may indicate the relevance of others. To obtain a layer-wise relevance score, we compute the mean relevance score of all features in a layer and then rank layers by relevance. We also run an alternative layer selection algorithm, in which we group weights originating from the same layer together and select layers using the ℓ_2 norm of the groups, referred to as *layer-wise (group)*. Figure 5-left compares feature-wise and layer-wise selection methods, matched for number of features retained. Feature-wise selection performs better than layer-wise selection on average and the best performance is obtained when around 1000 features are kept. Figure 5-right shows the distribution of features selected from each layer by the feature-wise and layer-wise strategies for the SVHN transfer task. We hypothesize that combining features across different layers provide better predictions, while including only the most important fea-

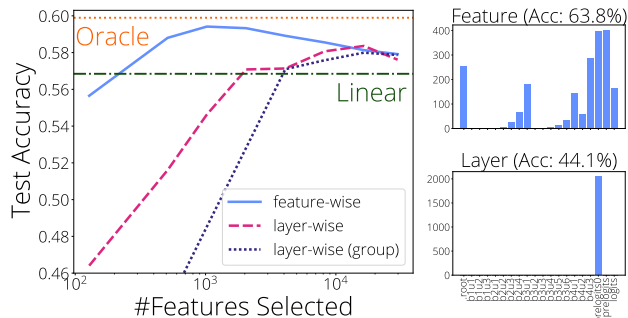


Figure 5. (left) Average accuracy over all VTAB tasks as a function of the number of features included. In this experiment, we only use features from the first layer of each of the 18 ResNet blocks and adjust pooling to have around 2048 features for each layer, totalling 29800 features. Results show that selecting layers performs worse than selecting features when adapting to a target domain. ORACLE results are explained in Section 3.2. (right) Distribution of 2048 intermediate features retained from a ResNet-50 when using feature-wise and layer-wise scores on the SVHN transfer task.

tures from each layer reduces over-fitting. We share figures for the 19 individual tasks in Appendix H.

Transfer across tasks. In practice, and in the literature, tasks are evaluated in isolation, i.e., the datasets for other tasks are not available. Nonetheless, in Figure 6-left, we investigate how features selected using some task i performs when evaluated on a different task j . Each cell in the array represents the average accuracy over three seeds. For each column, we subtract the diagonal term (i.e., self-transfer, $i = j$) to obtain delta accuracy for each task. For most tasks, using a separate task for feature selection hurts performance. Results in Flowers-102 and Pets get better when other tasks like sNORB-Elev are used. Crucially, no single task (rows of Figure 6-left) yields features that are universally optimal, which highlights the importance of doing the feature selection during adaptation, not beforehand.

Increasing the number of candidate features brings better results. HEAD2TOE often selects a few thousand features of the over one million features available. The high number of candidate features is critical for obtaining the best performance: in Figure 6-right, we vary the number of intermediate features used for each of our pretrained backbones, indicated by line style. We share scaling curves for the ViT-B/16 backbone in Appendix J and scaling curves for individual tasks in Appendix K. We observe that including all layers always performs better on average. However when varying the number of target features for each layer, we observed two distinct sub-groups of target tasks that behave differently as the number of features increases, indicated by the red and blue lines. This observation informed our decision to include both small and large target sizes in

our validation hyperparameter search. Given the positive slope of the scaling curves, further increasing the number of available features for selection is a promising research direction.

5. Related Work

Transfer Learning is studied extensively in the literature and used widely in practice. To our knowledge, the utility of combining intermediate layers of a deep neural network is first shown in speech domain (Choi et al., 2017; Lee & Nam, 2017). Since then many have explored the utility of intermediate representations. Shor et al. (2020; 2021) observed earlier layers of an unsupervised speech model to transfer better. ELMo (Peters et al., 2018) averaged two LSTM embeddings using a learned linear combination (a softmax). Lee et al. (2018) used intermediate features to improve the detection of OOD and adversarial inputs. Tang & de Sa (2020) used random projections to reduce the cost of combining different layers, whereas Adler et al. (2020) ensembled hebbian learners trained on intermediate layers and observed larger gains in far target domains. Similar approaches are also used recently to combine multiple backbones (Guo et al., 2020; Lopes et al., 2021) or to improve calibration (Khalifa & Alabdulmohsin, 2022). Algorithms that combine layers with all features require embeddings to be same size. Thus, they are most similar to the suboptimal layer-selection baseline shown in Figure 5-left. Most similar to our work is the work of (Dalvi et al., 2019; 2020), which proposes a five-step method to select token representations from multiple locations in a pretrained model. They only consider the representation of the [CLS] token after the second MLP of the self-attention block, which makes the feature selection problem significantly smaller and as shown in Appendix J results in sub-optimal transfer.

Given the ever increasing size and never saturating performance of pretrained models, the importance of reducing the cost of FINETUNING models is stated in the literature regularly. Methods like feature-wise transformations (Bilen & Vedaldi, 2017; Dumoulin et al., 2018), residual adapters (Houlsby et al., 2019; Rebuffi et al., 2017; Puigcerver et al., 2021), Diff-pruning (Guo et al., 2021) and selective fine-tuning (Guo et al., 2019; Fu et al., 2021) are proposed in order to reduce the cost of storing fine-tuned models. However, none of these methods match the simplicity of training (and storing) a linear classifier (i.e. HEAD2TOE) and they can be applied in conjunction with HEAD2TOE. Teerapittayanon et al. (2016), Kaya et al. (2019), and Zhou et al. (2020) studied intermediate representations to reduce “overthinking” and thus provide better early-exit criteria. Similarly Baldock et al. (2021) showed a correlation between early classification of a sample and how easy its classification is. Intermediate features of a

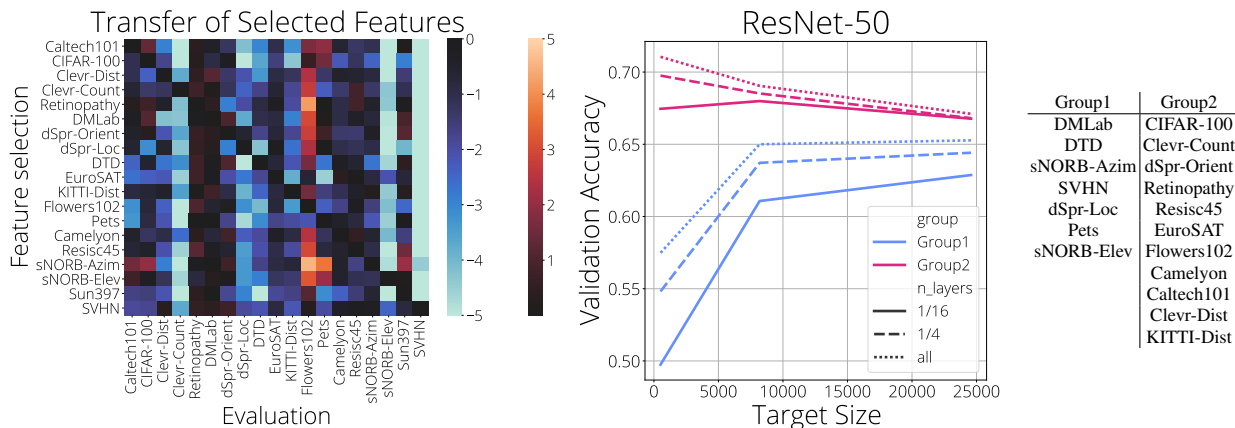


Figure 6. (left) Change in accuracy when features selected for a different task are used for adaptation. Most tasks get their best accuracy when the same task is also used for feature selection. (right) Effect of increasing the number of intermediate features that HEAD2TOE uses from the ResNet-50 backbone. The abscissa of the graph indicates the dimensionality of the representation extracted from each layer of the backbone (*target size*). The tasks are split into two groups (see right side of the figure), which show different behavior. The solid, dashed, and dotted lines indicate the fraction of layers selected for forming the representation used by HEAD2TOE: 1/16, 1/4, and 1, respectively.

pretrained backbone are also used in object detection (Har-
 iharan et al., 2015; Bell et al., 2016; Lin et al., 2017) and
 recently observed to improve the training of generative-
 adversarial networks (Sauer et al., 2021). Multiple feature
 representations are also used by approaches that use multi-
 domain training as an inductive bias, either without (Dvornik
 et al., 2020; Triantafillou et al., 2021; Li et al., 2021b;a)
 or with (Liu et al., 2021) meta-learning. However, in the
 large-scale setting FINETUNING remains a top-performing
 approach to few-shot classification (Dumoulin et al., 2021).

Feature selection approaches can be grouped according to
 whether labeled data is used—supervised (Nie et al., 2010)
 or unsupervised (Ball & Hall, 1965; Hart, 1968; He et al.,
 2005; Balin et al., 2019; Atashgahi et al., 2020)—or what
 high-level approach is taken—filter methods (Blum & Lan-
 gley, 1997), wrapper methods (Kohavi & John, 1997), or
 embedded methods (Yuan & Lin, 2006). Most relevant to
 our work are embedded supervised methods as they have
 good scaling properties which is vital in our setting with over
 a million features. Embedded supervised feature selection
 methods use a cost function to iteratively refine the subset of
 features selected and popular approaches include forward se-
 lection (Viola & Jones, 2001; Borboudakis & Tsamardinos,
 2019), backward selection (pruning) (Mozer & Smolensky,
 1989; Guyon et al., 2004) and regularization/feature-rank-
 ing based methods (Yuan & Lin, 2006; Blum & Langley,
 1997; Zhao et al., 2010). Most relevant to our work is Argyriou
 et al. (2007); Nie et al. (2010), both of which uses $\ell_{2,1}$
 regularization to select features, however their approach
 requires matrix inversions which is not practical in our set-
 ting. We point interested readers to Gui et al. (2017) and
 Boln-Canedo et al. (2015) for a detailed discussion.

6. Conclusion

In this work, we introduce HEAD2TOE, an approach that
 extends linear probing (LINEAR) by selecting the most rel-
 evant features among a pretrained network’s intermediate
 representations. We motivate this with a first-order Tay-
 lor series approximation to FINETUNING and show that
 the approach greatly improves performance over LINEAR
 and attains performance competitive with—and in some
 cases superior to—FINETUNING at much lower space and
 time complexity. Our findings challenge the conventional
 belief that FINETUNING is required to achieve good per-
 formance on OOD tasks. While more work is needed before
 HEAD2TOE can match the computational efficiency and
 simplicity of linear probing, our work paves the way for ap-
 plying new and more efficient feature selection approaches
 and for experimenting with HEAD2TOE probing in other
 domains such as regression, video classification, object de-
 tection, reinforcement learning, and language modelling.

Acknowledgments

We like to thank members of the Google Brain team for
 their useful feedback. Specifically we like to thank Cristina
 Vasconcelos, Eleni Triantafillou, Hossein Mobahi, Ross
 Goroshin for their feedback during the team meetings. We
 thank Joan Puigcerver, Fabian Pedregosa, Robert Gower,
 Laura Graesser, Rodolphe Jenatton and Timothy Nguyen for
 their feedback on the preprint. We thank Lucas Beyer and
 Xiaohua Zhai for creating the compact table for reporting
 VTAB results.

References

- Adler, T., Brandstetter, J., Widrich, M., Mayr, A., Kreil, D. P., Kopp, M., Klambauer, G., and Hochreiter, S. Cross-domain few-shot learning by representation fusion. *ArXiv*, abs/2010.06498, 2020.
- Alyafeai, Z., AlShaibani, M. S., and Ahmad, I. A survey on transfer learning in natural language processing. *arXiv preprint arXiv:2007.04239*, 2020.
- Argyriou, A., Evgeniou, T., and Pontil, M. Multi-task feature learning. In Schölkopf, B., Platt, J., and Hoffman, T. (eds.), *Advances in Neural Information Processing Systems*. MIT Press, 2007.
- Atashgahi, Z., Sokar, G., van der Lee, T., Mocanu, E., Mocanu, D. C., Veldhuis, R. N. J., and Pechenizkiy, M. Quick and robust feature selection: the strength of energy-efficient sparse training for autoencoders. *ArXiv*, abs/2012.00560, 2020.
- Baldock, R., Maennel, H., and Neyshabur, B. Deep learning through the lens of example difficulty. *ArXiv*, abs/2106.09647, 2021.
- Baln, M. F., Abid, A., and Zou, J. Concrete autoencoders: Differentiable feature selection and reconstruction. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.
- Ball, G. and Hall, D. ISODATA, a novel method of data analysis and classification. Technical report, Technical report, Stanford University, USA, 1965.
- Beattie, C., Leibo, J. Z., Teplyashin, D., Ward, T., Wainwright, M., Küttler, H., Lefrancq, A., Green, S., Valdés, V., Sadik, A., et al. Deepmind lab. *arXiv preprint arXiv:1612.03801*, 2016.
- Bell, S., Zitnick, C. L., Bala, K., and Girshick, R. B. Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2874–2883, 2016.
- Bilen, H. and Vedaldi, A. Universal representations: The missing link between faces, text, planktons, and cat breeds. *arXiv preprint arXiv:1701.07275*, 2017.
- Blum, A. L. and Langley, P. Selection of relevant features and examples in machine learning. *Artificial Intelligence*, 1997.
- Boln-Canedo, V., Snchez-Maroo, N., and Alonso-Betanzos, A. *Feature Selection for High-Dimensional Data*. Springer Publishing Company, Incorporated, 1st edition, 2015. ISBN 3319218573.
- Borboudakis, G. and Tsamardinos, I. Forward-backward selection with early dropping. *J. Mach. Learn. Res.*, 20: 8:1–8:39, 2019.
- Cheng, G., Han, J., and Lu, X. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 2017.
- Choi, K., Fazekas, G., Sandler, M. B., and Cho, K. Transfer learning for music classification and regression tasks. In *ISMIR*, 2017.
- Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., and Vedaldi, A. Describing textures in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2014.
- Dalvi, F., Durrani, N., Sajjad, H., Belinkov, Y., Bau, A., and Glass, J. R. What is one grain of sand in the desert? analyzing individual neurons in deep nlp models. In *AAAI*, 2019.
- Dalvi, F., Sajjad, H., Durrani, N., and Belinkov, Y. Analyzing redundancy in pretrained transformer models. In *EMNLP*, 2020.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houslsby, N. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- Dumoulin, V., Perez, E., Schucher, N., Strub, F., Vries, H. d., Courville, A., and Bengio, Y. Feature-wise transformations. *Distill*, 3(7):e11, 2018.
- Dumoulin, V., Houslsby, N., Evci, U., Zhai, X., Goroshin, R., Gelly, S., and Larochelle, H. Comparing transfer and meta learning approaches on a unified few-shot classification benchmark. *arXiv preprint arXiv:2104.02638*, 2021.
- Dvornik, N., Schmid, C., and Mairal, J. Selecting relevant features from a multi-domain representation for few-shot classification. In *European Conference on Computer Vision*, pp. 769–786. Springer, 2020.
- Fu, C., Huang, H., Chen, X., Tian, Y., and Zhao, J. Learn-to-share: A hardware-friendly transfer learning framework exploiting computation and parameter sharing. In *ICML*, 2021.
- Geiger, A., Lenz, P., Stiller, C., and Urtasun, R. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research*, 2013.

- Gui, J., Sun, Z., Ji, S., Tao, D., and Tan, T. Feature selection based on structured sparsity: A comprehensive study. *IEEE Transactions on Neural Networks and Learning Systems*, 28:1490–1507, 2017.
- Guo, D., Rush, A. M., and Kim, Y. Parameter-efficient transfer learning with diff pruning. In *ACL/IJCNLP*, 2021.
- Guo, Y., Shi, H., Kumar, A., Grauman, K., Simunic, T., and Feris, R. S. Spottune: Transfer learning through adaptive fine-tuning. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4800–4809, 2019.
- Guo, Y., Codella, N. C., Karlinsky, L., Codella, J. V., Smith, J. R., Saenko, K., Rosing, T., and Feris, R. A broader study of cross-domain few-shot learning. In *ECCV*, 2020.
- Guyon, I., Weston, J., Barnhill, S. D., and Vapnik, V. N. Gene selection for cancer classification using support vector machines. *Machine Learning*, 46:389–422, 2004.
- Hariharan, B., Arbeláez, P., Girshick, R. B., and Malik, J. Hypercolumns for object segmentation and fine-grained localization. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 447–456, 2015.
- Hart, P. The condensed nearest neighbor rule (corresp.). *IEEE transactions on information theory*, 14(3):515–516, 1968.
- He, X., Cai, D., and Niyogi, P. Laplacian score for feature selection. In *Proceedings of the 18th International Conference on Neural Information Processing Systems*, 2005.
- Helber, P., Bischke, B., Dengel, A., and Borth, D. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2019.
- Houlsby, N., Giurghi, A., Jastrzebski, S., Morrone, B., de Laroussilhe, Q., Gesmundo, A., Attariyan, M., and Gelly, S. Parameter-efficient transfer learning for nlp. In *ICML*, 2019.
- Johnson, J., Hariharan, B., van der Maaten, L., Fei-Fei, L., Lawrence Zitnick, C., and Girshick, R. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- Kaggle and EyePacs. Kaggle diabetic retinopathy detection, July 2015. URL <https://www.kaggle.com/c/diabetic-retinopathy-detection/data>.
- Kaya, Y., Hong, S., and Dumitras, T. Shallow-deep networks: Understanding and mitigating network overthinking. In *ICML*, 2019.
- Khalifa, A. and Alabdulmohsin, I. Improving the post-hoc calibration of modern neural networks with probe scaling, 2022. URL <https://openreview.net/forum?id=PO-32ODWng>.
- Kohavi, R. and John, G. H. Wrappers for feature subset selection. *Artif. Intell.*, 1997.
- Kornblith, S., Shlens, J., and Le, Q. V. Do better imagenet models transfer better? *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2656–2666, 2019.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- LeCun, Y., Huang, F. J., and Bottou, L. Learning methods for generic object recognition with invariance to pose and lighting. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2004.
- Lee, J. and Nam, J. Multi-level and multi-scale feature aggregation using pretrained convolutional neural networks for music auto-tagging. *IEEE Signal Processing Letters*, 24:1208–1212, 2017.
- Lee, K., Lee, K., Lee, H., and Shin, J. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *Advances in Neural Information Processing Systems*, 2018. URL <https://proceedings.neurips.cc/paper/2018/file/abdeb6f575ac5c6676b747bca8d09cc2-Paper.pdf>.
- Li, F.-F., Fergus, R., and Perona, P. One-shot learning of object categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2006.
- Li, W.-H., Liu, X., and Bilen, H. Improving task adaptation for cross-domain few-shot learning. *arXiv preprint arXiv:2107.00358*, 2021a.
- Li, W.-H., Liu, X., and Bilen, H. Universal representation learning from multiple domains for few-shot classification. *arXiv preprint arXiv:2103.13841*, 2021b.
- Lin, T.-Y., Dollár, P., Girshick, R. B., He, K., Hariharan, B., and Belongie, S. J. Feature pyramid networks for object detection. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 936–944, 2017.

- Liu, L., Hamilton, W., Long, G., Jiang, J., and Larochelle, H. A universal representation transformer layer for few-shot image classification. In *International Conference on Learning Representations*, 2021.
- Lopes, R. G., Dauphin, Y., and Cubuk, E. D. No one representation to rule them all: Overlapping features of training methods. *ArXiv*, abs/2110.12899, 2021.
- Maddox, W., Tang, S., Moreno, P. G., Wilson, A. G., and Damianou, A. C. Fast adaptation with linearized neural networks. In *AISTATS*, 2021.
- Matthey, L., Higgins, I., Hassabis, D., and Lerchner, A. dsprites: Disentanglement testing sprites dataset. <https://github.com/deepmind/dsprites-dataset/>, 2017.
- Mozer, M. C. and Smolensky, P. Skeletonization: A technique for trimming the fat from a network via relevance assessment. In *Advances in Neural Information Processing Systems*, 1989. URL <https://proceedings.neurips.cc/paper/1988/file/07e1cd7dca89a1678042477183b7ac3f-Paper.pdf>.
- Mu, F., Liang, Y., and Li, Y. Gradients as features for deep representation learning. *ArXiv*, abs/2004.05529, 2020.
- Mustafa, B., Riquelme, C., Puigcerver, J., and Andr'e Susano Pinto, Keysers, D., and Houlsby, N. Deep ensembles for low-data transfer learning. *ArXiv*, abs/2010.06866, 2020.
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*, 2011.
- Nie, F., Huang, H., Cai, X., and Ding, C. Efficient and robust feature selection via joint $\ell_{2,1}$ -norms minimization. In Lafferty, J., Williams, C., Shawe-Taylor, J., Zemel, R., and Culotta, A. (eds.), *Advances in Neural Information Processing Systems*, 2010.
- Nilsback, M.-E. and Zisserman, A. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics and Image Processing*, Dec 2008.
- Parkhi, O. M., Vedaldi, A., Zisserman, A., and Jawahar, C. V. Cats and dogs. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2012.
- Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. Deep contextualized word representations. In *NAACL*, 2018.
- Puigcerver, J., Riquelme, C., Mustafa, B., Renggli, C., Pinto, A. S., Gelly, S., Keysers, D., and Houlsby, N. Scalable transfer learning with expert models. *ArXiv*, abs/2009.13239, 2021.
- Rebuffi, S.-A., Bilen, H., and Vedaldi, A. Learning multiple visual domains with residual adapters. *arXiv preprint arXiv:1705.08045*, 2017.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 2015.
- Sauer, A., Chitta, K., Müller, J., and Geiger, A. Projected gans converge faster. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Shor, J., Jansen, A., Maor, R., Lang, O., Tuval, O., de Chaumont Quiry, F., Tagliasacchi, M., Shavitt, I., Emanuel, D., and Haviv, Y. A. Towards learning a universal non-semantic representation of speech. *ArXiv*, abs/2002.12764, 2020.
- Shor, J., Jansen, A., Han, W., Park, D., and Zhang, Y. Universal paralinguistic speech representations using self-supervised conformers. *ArXiv*, abs/2110.04621, 2021.
- Tang, S. and de Sa, V. R. Deep transfer learning with ridge regression. *ArXiv*, abs/2006.06791, 2020.
- Teerapittayanon, S., McDanel, B., and Kung, H. T. Branchynet: Fast inference via early exiting from deep neural networks. *2016 23rd International Conference on Pattern Recognition (ICPR)*, pp. 2464–2469, 2016.
- Triantafillou, E., Larochelle, H., Zemel, R., and Dumoulin, V. Learning a universal template for few-shot dataset generalization. In *International Conference on Machine Learning*, 2021.
- Veeling, B. S., Linmans, J., Winkens, J., Cohen, T., and Welling, M. Rotation equivariant cnns for digital pathology. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2018.
- Viola, P. A. and Jones, M. J. Rapid object detection using a boosted cascade of simple features. *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, 1:I–I, 2001.
- Wu, S., Zhong, S., and Liu, Y. Deep residual learning for image steganalysis. *Multimedia Tools and Applications*, 2018.
- Xiao, J., Hays, J., Ehinger, K. A., Oliva, A., and Torralba, A. Sun database: Large-scale scene recognition from

abbey to zoo. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2010.

Yuan, M. and Lin, Y. Model selection and estimation in regression with grouped variables. *Journal of The Royal Statistical Society Series B-statistical Methodology*, 68: 49–67, 2006.

Zhai, X., Puigcerver, J., Kolesnikov, A., Ruysen, P., Riquelme, C., Lucic, M., Djolonga, J., Pinto, A. S., Neumann, M., Dosovitskiy, A., Beyer, L., Bachem, O., Tschannen, M., Michalski, M., Bousquet, O., Gelly, S., and Houlsby, N. The visual task adaptation benchmark. *ArXiv*, abs/1910.04867, 2019.

Zhao, Z., Wang, L., and Liu, H. Efficient spectral feature selection with minimum redundancy. In *AAAI*, 2010.

Zhou, W., Xu, C., Ge, T., McAuley, J., Xu, K., and Wei, F. Bert loses patience: Fast and robust inference with early exit. *ArXiv*, abs/2006.04152, 2020.

Zhou, Z.-H., Wu, J., and Tang, W. Ensembling neural networks: Many could be better than all. *Artificial Intelligence*, 2002. URL <https://www.sciencedirect.com/science/article/pii/S000437020200190X>.

Zhu, Z., Lin, K., and Zhou, J. Transfer learning in deep reinforcement learning: A survey. *arXiv preprint arXiv:2009.07888*, 2020.

Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., and He, Q. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76, 2020.

Author Contributions

- Utku: Proposed/planned/led the project, wrote the majority of the code, performed most experiments, wrote the initial draft of the paper.
- Vincent: Participated in weekly meetings for the project, reviewed code, contributed to framing the findings in terms of refuting the hypothesis that a pre-trained network lacks the features required to solve OOD classification tasks, helped with paper writing, helped run evaluations.
- Hugo: Helped identify and frame the research opportunity, attended regular meetings, contributed to analysis discussions, minor contributions to the writing. Also made this publication possible by pointing out that the paper was over the page-limit just before the submission.
- Mike: Participated in weekly meetings, confused matters due to his unfamiliarity with the literature, argued that the central idea of the paper was never going to work, contributed to framing research and helped substantially with the writing.

A. Hyperparameter Selection

We pick hyperparameters for each VTAB task separately by doing a 5-fold cross validation on the training data. For all methods, we chose the learning rate and the total number of training steps using the grid $lr = 0.1, 0.01$ and $steps = 500, 5000$, following the lightweight hyperparameter sweep recommended by the VTAB benchmark (Zhai et al., 2019).

For regularization baselines ℓ_1 , ℓ_2 and $\ell_{2,1}$ we search for regularization coefficients using (0.00001, 0.0001, 0.001). We include an extra value in this setting in order to account for the overhead introduced by HEAD2TOE.

For HEAD2TOE we choose $\ell_{2,1}$ regularization coefficients from (0.001, 0.00001) and target feature sizes from (1024, 16384, 40000) for ResNet-50 and (768, 15360, 32448) for ViT-B/16. After calculating feature scores HEAD2TOE validates the following fractions: (0.0005, 0.001, 0.002, 0.005, 0.01, 0.02, 0.05, 0.1) and thus requires 18% more operations compared to other regularization baselines. Note that this is because initial training to obtain feature scores s_i is performed once and therefore searching for optimal number of features has a small overhead. Hyper parameters selected by HEAD2TOE for each VTAB task are shared in Table 2. Next we explain how this overhead is estimated.

Dataset	T	F	LR	Steps	$\ell_{2,1}$	T	F	LR	Steps	$\ell_{2,1}$
	ResNet-50					ViT-B/16				
Caltech101	8192	0.010	0.01	5000	0.00001	768	0.050	0.01	5000	0.00100
CIFAR-100	512	0.200	0.01	500	0.00001	768	0.020	0.01	500	0.00001
Clevr-Dist	8192	0.001	0.01	500	0.00100	15360	0.002	0.01	500	0.00100
Clevr-Count	512	0.005	0.10	5000	0.00100	768	0.050	0.10	5000	0.00001
Retinopathy	8192	0.200	0.01	500	0.00001	768	0.010	0.01	500	0.00100
DMLab	8192	0.020	0.01	500	0.00001	32448	0.005	0.01	500	0.00001
dSpr-Orient	512	0.200	0.01	500	0.00001	768	0.100	0.01	5000	0.00001
dSpr-Loc	8192	0.005	0.10	500	0.00100	32448	0.002	0.10	500	0.00100
DTD	24576	0.005	0.01	5000	0.00001	768	0.100	0.01	500	0.00100
EuroSAT	512	0.100	0.01	500	0.00001	768	0.100	0.01	500	0.00100
KITTI-Dist	8192	0.020	0.01	500	0.00001	32448	0.050	0.01	5000	0.00100
Flowers102	512	0.100	0.01	5000	0.00001	768	0.020	0.01	500	0.00001
Pets	8192	0.002	0.01	5000	0.00001	768	0.020	0.01	5000	0.00100
Camelyon	512	0.020	0.10	500	0.00100	768	0.100	0.01	500	0.00100
Resisc45	8192	0.020	0.01	500	0.00001	768	0.050	0.01	5000	0.00001
sNORB-Azim	24576	0.002	0.01	500	0.00001	32448	0.010	0.01	500	0.00001
sNORB-Elev	8192	0.050	0.01	500	0.00100	15360	0.200	0.01	500	0.00100
Sun397	512	0.100	0.01	5000	0.00100	768	0.050	0.01	5000	0.00100
SVHN	24576	0.005	0.01	500	0.00001	32448	0.005	0.01	500	0.00001

Table 2. Hyper parameters selected for the VTAB-1k benchmark tasks when using pretrained ResNet-50 and ViT-B/16 backbones. **T**: target sizes of features included from each layer, **F**: fraction of features kept, **LR**: learning rate, **Steps**: Training Steps, $\ell_{2,1}$: regularization coefficient.

Head2Toe: Utilizing Intermediate Representations for Better Transfer Learning

Dataset	F	N	C	FLOPs (vs FINE TUNING)	Size (vs FINE TUNING)	Size (vs LINEAR)
Caltech101	0.010	467688	102	0.009675	0.020750	2.353167
CIFAR-100	0.200	30440	100	0.005792	0.025743	2.977301
Clevr-Dist	0.001	467688	6	0.005747	0.000741	1.417419
Clevr-Count	0.005	30440	8	0.000568	0.000092	0.132278
Retinopathy	0.200	467688	5	0.005657	0.020531	47.099634
DMLab	0.020	467688	6	0.005747	0.003011	5.756287
dSpr-Orient	0.200	30440	16	0.005302	0.004183	3.001686
dSpr-Loc	0.005	467688	16	0.006644	0.002212	1.587624
DTD	0.005	1696552	47	0.015823	0.019157	4.692396
EuroSAT	0.100	30440	10	0.005267	0.001336	1.532776
KITTI-Dist	0.020	467688	4	0.005567	0.002215	6.350983
Flowers102	0.100	30440	102	0.001117	0.013146	1.490882
Pets	0.002	467688	37	0.003842	0.002089	0.649417
Camelyon	0.020	30440	2	0.005220	0.000092	0.529114
Resisc45	0.020	467688	45	0.009247	0.018474	4.725480
sNORB-Azim	0.002	1696552	18	0.011069	0.004851	3.094923
sNORB-Elev	0.050	467688	9	0.006016	0.009578	12.210897
Sun397	0.100	30440	397	0.002839	0.049781	1.487498
SVHN	0.005	1696552	10	0.008464	0.005865	6.730334
Average				0.006295	0.010729	5.674742

Table 3. Relative cost of HEAD2TOE when compared with FINE TUNING and LINEAR. F is the fraction of features kept, N is the total number of features and C is the number of classes. On average HEAD2TOE requires 0.5% of the FLOPs required by FINE TUNING during the adaptation. Cost of storing each adapted model is also small: requiring only 1% of the FINE TUNING and only 5.7x more than LINEAR. See main text for details on how the numbers are calculated.

B. Cost of HEAD2TOE

We evaluate different values of F and pick the value with best validation performance. Cost of HEAD2TOE consists of three parts: (1) C_I : Cost of calculating the representations using the pretrained backbone. (2) C_{all} : cost of training the initial head W_{all} (in order to obtain s_i 's) (3) $\sum_f C_{F=f}$: total cost of validating different values of F . Cost of validating a fraction value f , assuming equal number of training steps, is equal to $C_f = C_{all} * f$. Therefore relative cost of searching for F is equal to the sum of fractions validated (in comparison to the initial training of (s)).

In Table 3, we compare cost of running HEAD2TOE adaptation with FINE TUNING. HEAD2TOE uses the backbone once in order to calculate the representations and then trains the W_{all} , whereas FINE TUNING requires a forward pass on the backbone at each step. Therefore we calculate the cost of fine-tuning for t steps as $C_I \cdot t$. Similarly, cost of HEAD2TOE is calculated as $C_I + C_{all} \cdot t$. The overall relative cost of C_{all} increases with number of classes C and number of features considered N . As shown in Table 1-top, HEAD2TOE obtains better results than FINE TUNING, yet it requires 0.5% of the FLOPs needed on average.

After adaptation, all methods require roughly same number of FLOPs for inference due to all methods using the same backbone. The cost of storing models for each task becomes important when the same pre-trained backbone is used for many different tasks. In Table 3 we compare the cost of storing different models found by different methods. A fine-tuned model requires all weights to be stored which has the same cost as storing the original network, whereas LINEAR and HEAD2TOE requires storing only the output head: W_{linear} . HEAD2TOE also requires to store the indices of the features selected using a bitmap. Even though HEAD2TOE considers many more features during adaptation, it selects a small subset of the features and thus requires a much smaller final classifier (on average 1% of the FINE TUNING). Note that hyper parameters are selected to maximize accuracy, not the size of the final model. We expect to see greater savings with an efficiency oriented hyperparameter selection.

C. Details of Intermediate Representations Included

ResNet-50 has 5 stages. First stage includes a single convolutional layer (root) followed by pooling. We include features after pooling. Remaining stages include 3,4,6 and 3 bottleneck units (v2) each. Bottleneck units start with normalization and activation and has 3 convolutional layers. We includes features right after the activation function is applied, resulting in 3 intermediate feature sets per unit. Output after 5 stages are average-pooled and passed to the output head. We include features before and after pooling. with the output of the final layer (logits), total number of locations where features are extracted makes 52.

ViT-B/16 model consists of multiple encoder units. Each encoder unit consists of a self attention module followed by 2 MLPs. For each of these units, we include features (1) after layer-norm (but before self-attention) (2) features after self-attention (3-4) features after MLP layers (and after gelu activation function). Additionally we use patched and embedded image (i.e. tokenized image input), pre-logits (final CLS-embedding) and logits.

D. Standard Deviations for Table 1

In Table 4, we share the standard deviations for the median test accuracies presented in Table 1. On average we observe **LINEAR** obtains lower variation as expected, due to the limited adaptation and convex nature of the problem. Head2Toe seem to have similar (or less) variation then the other regularization baselines that use all features.

	Natural							Specialized				Structured							Mean	
	CIFAR-100	Caltech101	DTD	Flowers102	Pets	SVHN	Sun397	Camelyon	EuroSAT	Resisc45	Retinopathy	Clevr-Count	Clevr-Dist	DMLab	KITTI-Dist	dSpr-Loc	dSpr-Ori	sNORB-Azim		sNORB-Elev
Linear	0.09	0.08	0.14	0.06	0.08	0.17	0.06	0.06	0.03	0.12	0.08	0.42	0.1	0.03	0.21	0.14	0.07	0.0	0.21	0.11
+All- ℓ_2	0.09	0.78	0.44	0.29	0.22	0.02	0.04	0.02	0.06	0.05	0.02	0.09	0.2	0.1	1.31	0.32	0.08	0.03	0.35	0.24
+All- ℓ_1	0.14	0.11	0.11	0.11	0.11	0.12	0.03	0.06	0.04	0.2	0.02	0.47	0.33	0.18	0.2	0.34	0.07	0.06	0.44	0.17
+All- $\ell_{2,1}$	0.09	1.0	0.13	0.1	0.1	0.18	0.23	0.15	0.07	0.09	0.05	0.03	0.08	0.11	0.41	0.51	0.24	0.16	0.07	0.2
Head2Toe	0.14	0.25	0.08	0.08	0.24	0.24	0.16	0.23	0.06	0.06	0.03	0.18	0.23	0.13	0.43	0.3	0.06	0.4	0.08	0.18
Fine-tuning	0.53	1.05	0.18	0.62	0.16	0.24	0.86	2.2	0.52	0.57	0.37	2.39	2.49	0.43	1.06	0.54	0.33	0.39	0.73	0.83
Head2Toe-FT	0.23	50.52	0.47	0.49	0.74	0.05	13.78	0.39	0.0	0.64	0.0	3.02	1.53	0.9	1.11	0.83	0.43	0.46	0.76	4.02
Head2Toe-FT+	0.29	0.15	0.45	0.08	0.28	0.25	0.12	1.61	0.08	0.15	0.18	0.27	0.67	1.0	0.76	0.96	0.1	0.38	2.09	0.52

	Natural							Specialized				Structured							Mean	
	CIFAR-100	Caltech101	DTD	Flowers102	Pets	SVHN	Sun397	Camelyon	EuroSAT	Resisc45	Retinopathy	Clevr-Count	Clevr-Dist	DMLab	KITTI-Dist	dSpr-Loc	dSpr-Ori	sNORB-Azim		sNORB-Elev
Linear	0.1	0.23	0.19	0.16	0.06	0.06	0.08	0.09	0.08	0.06	0.0	0.06	0.02	0.07	0.92	0.21	0.08	0.08	0.03	0.14
+All- ℓ_2	0.13	0.17	0.0	0.66	0.36	0.04	0.56	0.02	0.59	0.01	0.04	0.23	0.13	0.12	1.24	0.28	0.09	0.84	0.2	0.3
+All- ℓ_1	0.06	0.11	0.15	0.08	0.19	0.31	0.08	0.08	0.06	0.12	0.05	0.28	0.17	0.23	0.87	0.34	0.07	0.22	0.28	0.2
+All- $\ell_{2,1}$	1.55	0.03	0.12	0.04	0.06	0.41	0.13	0.18	0.08	0.1	0.04	0.09	0.13	0.15	0.87	0.66	0.1	0.23	0.22	0.27
Head2Toe	0.29	0.16	0.26	0.5	0.19	0.14	0.07	0.13	0.04	0.09	0.09	0.04	0.44	0.14	1.34	0.21	0.01	0.31	0.08	0.24
Scratch	0.25	0.29	0.43	0.68	0.16	1.09	0.2	0.57	0.79	1.03	1.63	0.13	0.94	0.52	1.8	4.73	3.17	0.61	1.39	1.08
Fine-tuning	1.71	0.89	0.13	0.83	0.81	3.41	0.33	1.59	0.45	0.43	1.1	1.99	0.47	0.25	1.24	0.98	3.52	1.33	1.8	1.22
Head2Toe-FT	0.68	0.39	1.05	0.05	0.65	0.45	0.3	0.24	0.1	0.76	0.29	1.78	1.33	0.21	0.81	2.69	0.45	0.64	1.65	0.76
Head2Toe-FT+	0.25	0.08	0.16	0.24	0.28	0.21	0.06	0.55	0.13	0.14	0.03	1.84	1.03	0.5	1.29	2.32	1.98	0.19	0.39	0.61

Table 4. Standard deviation of test accuracy over 3 seeds on the VTAB-1k benchmark using pretrained (**top**) ResNet-50 and (**bottom**) ViT-B/16 backbones. The mean column averages the standard deviations for each dataset.

E. Details of Datasets used in VTAB-1k benchmark

Table 5 include datasets used in VTAB-1k benchmark.

Dataset	Classes	Reference
Caltech101	102	(Li et al., 2006)
CIFAR-100	100	(Krizhevsky, 2009)
DTD	47	(Cimpoi et al., 2014)
Flowers102	102	(Nilsback & Zisserman, 2008)
Pets	37	(Parkhi et al., 2012)
Sun397	397	(Xiao et al., 2010)
SVHN	10	(Netzer et al., 2011)
EuroSAT	10	(Helber et al., 2019)
Resisc45	45	(Cheng et al., 2017)
Patch Camelyon	2	(Veeling et al., 2018)
Retinopathy	5	(Kaggle & EyePacs, 2015)
Clevr	8	(Johnson et al., 2017)
dSprites	16	(Matthey et al., 2017)
SmallNORB	18	(LeCun et al., 2004)
DMLab	6	(Beattie et al., 2016)
KITTI/distance	4	(Geiger et al., 2013)

Table 5. Datasets used in VTAB-1k benchmark.

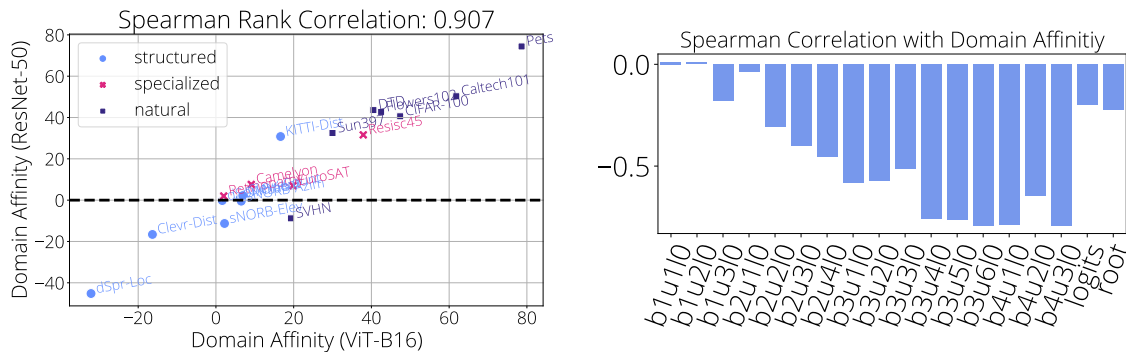


Figure 7. (left) Domain affinity scores calculated using a ResNet-50 and a ViT-B/16 show a clear correlation. (right) For each layer, we calculate the the percentage improvement over LINEAR for each task. We report Spearman’s rank correlation between the percentage improvement and our domain affinity metric. All layers have a near-zero or negative correlation: the benefit from incorporating intermediate features diminishes with increased domain affinity, especially for the later layers.

F. Domain Affinity Metric

Domain affinity metric aims to capture the similarity between two supervised learning task. If two tasks are similar, we expect the representations learned in one task to transfer better using a simple linear probe on the last layer compared to training the network from scratch in a low-data regime. Thus we define our domain affinity as the difference between linear probe and scratch accuracies. Here we demonstrate the robustness of the domain affinity metric to different backbone architectures and pretraining algorithms.

We expect the domain affinity calculated using different pretrained backbones to provide similar orderings. We demonstrate this in Figure 7-left, where we calculate domain affinity using the ViT-B/16 backbone and observe a high correlation (Spearman, 0.907) with the original scores calculated using the ResNet-50 backbone. Furthermore we did additional investigations using the 15 representation learning methods presented on the VTAB-leaderboard⁵, where we calculated median percentage-improvement over scratch training for each task, and similarly observed a high Spearman correlation (0.803).

G. Additional Plots for Experiments using Single Additional Layer

Test accuracies when using a single additional intermediate layer from a pretrained ResNet-50 backbone are shown in Figure 8. Natural datasets (except SVHN) are highly similar to upstream dataset (ImageNet-2012) and thus adding an

⁵https://google-research.github.io/task_adaptation/benchmark

extra intermediate layer doesn't improve performance much. However performance on OOD tasks (mostly of the tasks in structured category) improves significantly when intermediate representations are used, which correlates strongly with datasets in which HEAD2TOE exceeds FINETUNING performance. This is also demonstrated in Figure 7-right, where we observe a negative correlation between domain similarity and accuracy gains for most of the intermediate layers when included during transfer.

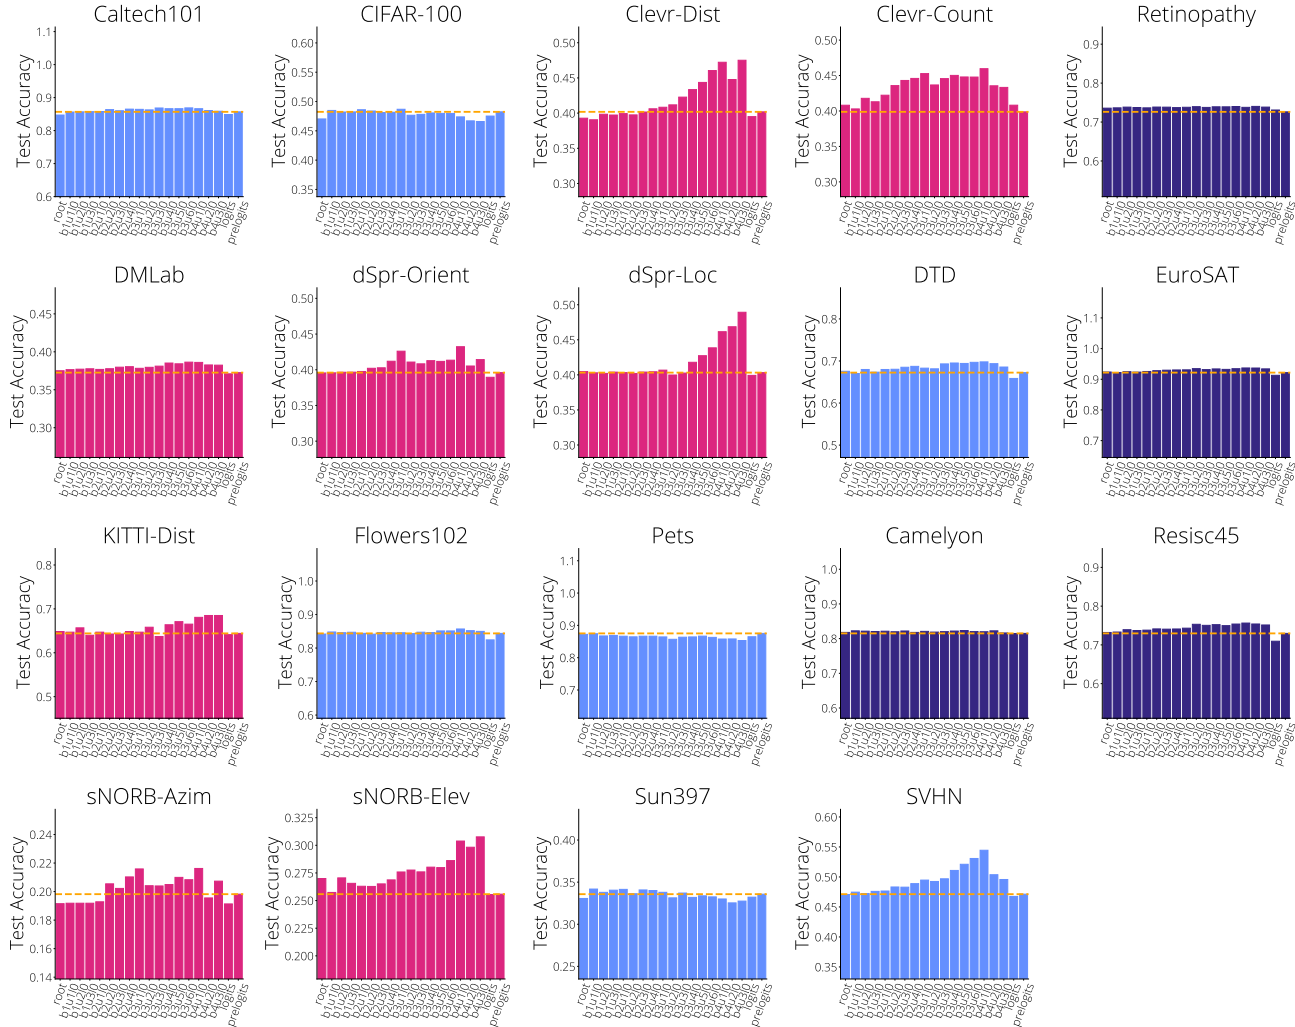


Figure 8. Test accuracies when using a single additional intermediate layer from a pretrained ResNet-50 backbone.

H. Additional Plots for Layer/Feature-wise Selection Comparison

We compare layer-wise selection strategy discussed in Section 4.3 to HEAD2TOE in Figure 9. For almost all datasets, feature-wise selection produces better results. Retinopathy and Flowers-102 are the only two datasets where the layer-wise strategy performs better.

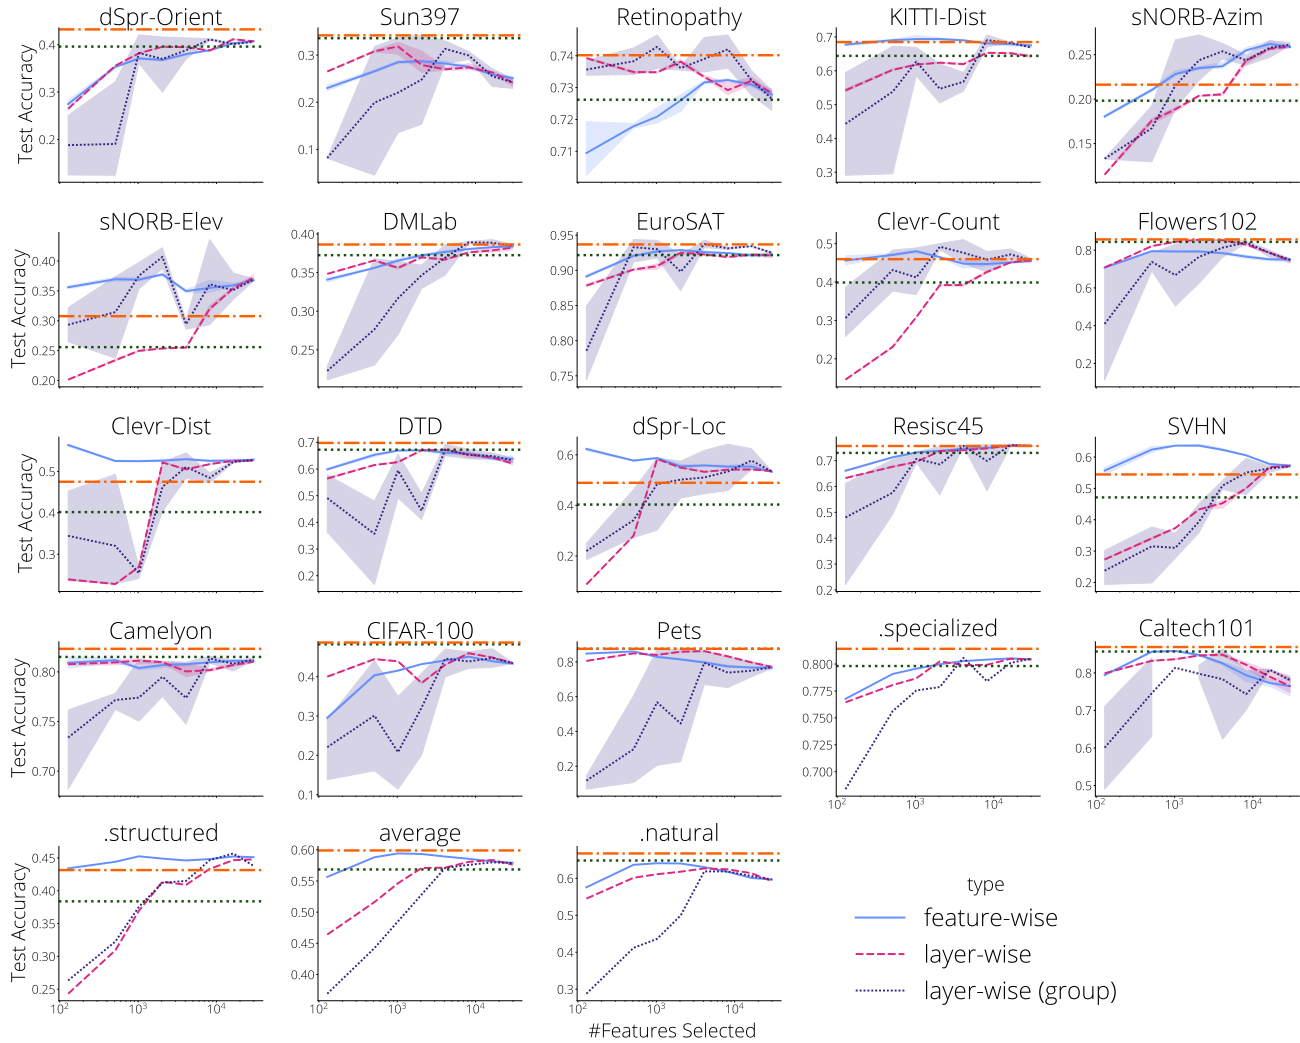


Figure 9. Test accuracies when varying the number of features selected for HEAD2TOE using a pretrained ResNet-50 backbone.

I. Additional Results for ResNet-50

Improvement of HEAD2TOE over fine-tuning test accuracy for ResNet-50 backbone is shown in Figure 10-left. Similar to earlier plots, we observe a clear trend between being OOD and improvement in accuracy: HEAD2TOE obtains superior few-shot generalization for most of OOD tasks. We also share the distribution of features selected for each task in Figure 11. Since different tasks might have different number features selected, we normalize each plot such that bars add up to 1. Overall, features from later layers seem to be preferred more often. Early layers are preferred, especially for OOD tasks like Clevr and sNORB. We observe a diverse set of distributions for selected features, which reflects the importance of selecting features from multiple layers. Even when distributions match, HEAD2TOE can select different features from the same layer for two different tasks. Therefore, next, we compare the indices of selected features directly to measure the diversity of features selected across tasks and seeds.

Head2Toe: Utilizing Intermediate Representations for Better Transfer Learning

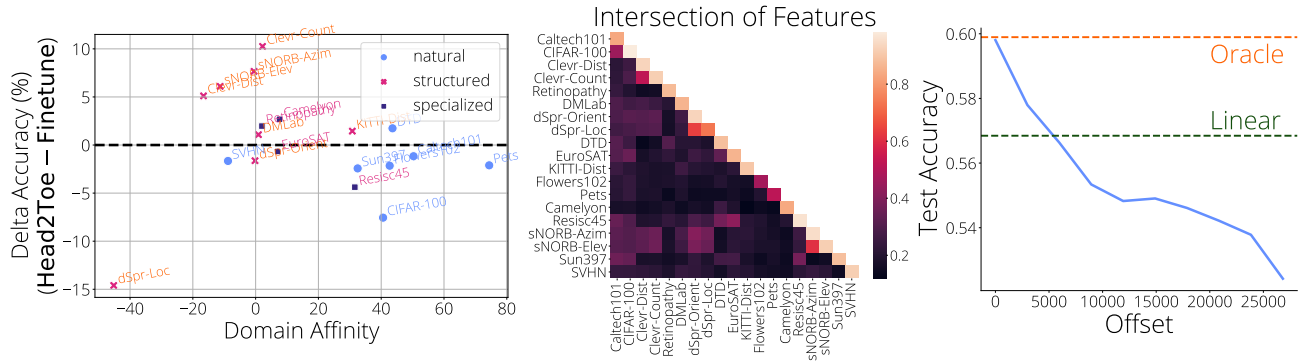


Figure 10. (left) Accuracy improvement of HEAD2TOE compared to FINETUNING. (center) Intersection of features when selecting 2048 features from 29800 (same settings as in Figure 5). The intersection is calculated as the fraction of features selected in two different runs. Values are averaged over 9 pairs of runs (3 seeds for each datasets), except the diagonal terms for which we remove identical pairs resulting in 6 pairs. (right) Over all VTAB tasks, average accuracy of HEAD2TOE when selecting 2048 consecutive features sorted by their relevance score, starting with an index specified by *offset*.



Figure 11. Distribution of selected features over different ResNet-50 layers for VTAB-1k tasks for results presented in Table 1. We group the layers in each block (group of 3 layers) to reduce numbers of bars.

Similarity of features selected. In Figure 10-center we investigate the intersection of features selected by HEAD2TOE. We select 2048 features for each task from the pool of 29800 features (same experiments as in Figure 5). For each task HEAD2TOE selects a different subset of features. We calculate the fraction of features shared between each subset. For each target task we run 3 seeds resulting in 3 sets of features kept. When comparing the similarity across seeds (diagonal terms),

we average over 6 possible combinations among these 3 sets. Similarly, when comparing two different tasks we average over 9 possible combinations. Both dSprites and sNORB images are generated artificially and have similar characteristics. Results show that such complementary tasks have the highest overlap, around 40%. Apart from a small fraction of tasks, however, most tasks seem to share less than 20% of the features, which highlights, again, the importance of doing the feature selection for each target task separately.

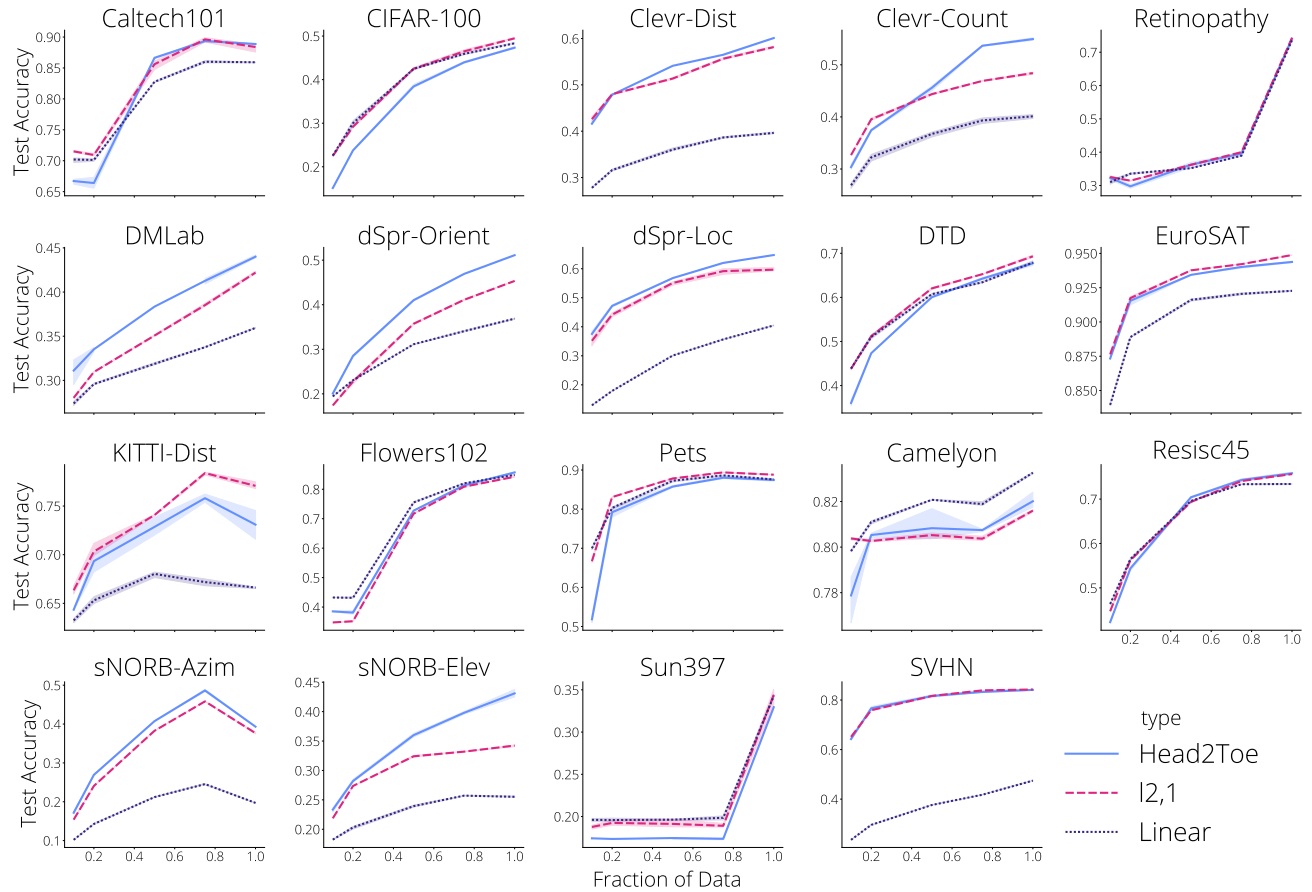


Figure 12. Effect of data available during training to the test accuracy. Fraction=1 indicates original tasks with 1000 training samples. Overall we observe the performance of HEAD2TOE improves with amount of data available, possibly due to the reduced noise in feature selection.

Effect of Training Data In Figure 12, we compare the performance of HEAD2TOE with other baselines using reduced training data. Fraction (f_d)=1 indicates original tasks with 1000 training samples. For other fractions we calculate number of shots for each class as $\text{int}(1000 * f_d / C)$ where C is the number of classes and then sample a new version of the task accordingly. For SUN-397 task, number of shots are capped at 1 and thus fractions below 0.75 lead to 1-shot tasks and thus results are all the same. Overall we observe the performance of HEAD2TOE improves with amount of data available, possibly due to the reduced noise in feature selection.

Pre-activations or Activations? Taylor approximation of the fine-tuning presented in Section 3.1 suggests that the activations at every layer should be used. As an ablation we compare this approach with other alternatives in Table 6. Using only the CLS tokens at every layer without pooling (*Only CLS*) or appending the CLS token to the pooled representations (*Original+CLS*) didn't improve the results.

Including pooled input as a candidate for feature selection We also try providing (possibly pooled) input to HEAD2TOE. As shown in Table 6, this didn't improved the results and we got around 0.8% lower accuracy on average.

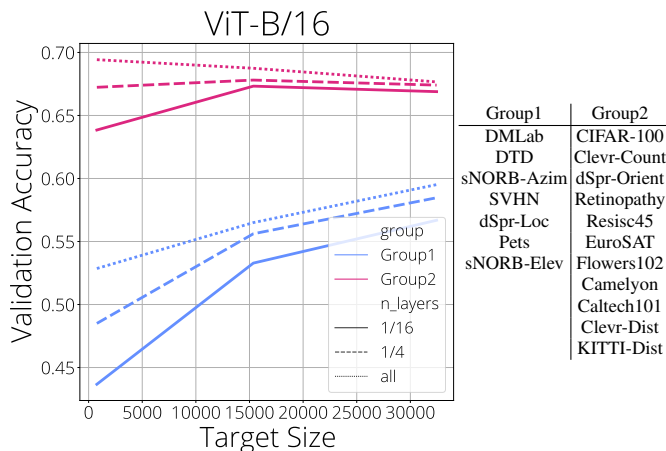


Figure 13. Effect of increasing the number of intermediate features that HEAD2TOE uses from the ViT-B/16 backbone. The abscissa of the graph indicates the dimensionality of the representation extracted from each layer of the backbone (*target size*). The tasks are split into two groups (see right side of Figure), which show different behavior. The solid, dashed, and dotted lines indicate the fraction of layers selected for forming the representation used by HEAD2TOE: 1/16, 1/4, and 1, respectively. Scaling curves for individual tasks can be found in Appendix K.

	Average	Natural	Specialized	Structured
Original	65.8	70.5	81.6	53.7
Pre-Normalization	63.5	67.4	81.1	51.3
Pre-Activation	63.5	67.4	81.3	51.3
Adding Pooled Input	65.0	69.7	81.2	52.9

Table 6. Ablation of choice of activations used on HEAD2TOE results for ResNet-50 backbone. Using features before the activation (*Pre-Activation*) or the normalization (*Pre-Normalization*) or adding the inputs to the pool of features considered bring worse results.

Relevance Scores. In Figure 10-right, we demonstrate the effectiveness of group lasso on identifying relevant intermediate features of a ResNet-50 trained on ImageNet. We rank all features by their relevance score, s_i , and select groups of 2048 consecutive features beginning at a particular offset in this ranking. Offset 0 corresponds to selecting the features with largest relevance. We calculate average test accuracies across all VTAB tasks. As the figure shows, test accuracy decreases monotonically with the offset, indicating that the relevance score predicts the importance of including a feature in the linear classifier.

J. Additional Results for ViT-B/16

Handling of CLS token Class (CLS) tokens in vision transformers are often added to the input and the classification layer is trained on top of the final representation of this token. Given that the representation of each token changes slowly, one might expect therefore the CLS representations along different layer to have more discriminate features. HEAD2TOE treats all tokens same and pools them together and in Table 7 we compare our approach with few other alternative approaches. Using only the CLS token at every layer without pooling (*Only CLS*) or appending the CLS token to the pooled representations (*Original+CLS*) doesn’t improve the results.

Scaling of ViT-B/16 We repeat the scaling plot for ResNet-50 (Figure 6-right) for ViT-B/16 at Figure 13 and observe a quite similar scaling behaviour for the two groups of datasets.

	Average	Natural	Specialized	Structured
Original	63.3	71.3	82.6	46.7
Only CLS	57.3	65.1	81.4	38.4
Original+CLS	63.4	72.0	82.0	46.6

Table 7. Alternatives treatments for the CLS token of the ViT-B/16 architecture. Using only the CLS tokens at every layer without pooling (*Only CLS*) or appending the CLS token to the pooled representations (*Original+CLS*) doesn't improve the results.

Distribution of features across layers We share the distribution of features selected for each task in Figure 14. As before, we normalize each plot such that bars add up to 1. Similar to ResNet-50, features from later layers selected more often for tasks from natural category. Early layers are preferred, especially for OOD tasks. In general distributions are more balanced, in other words, more intermediate features are used compared to ResNet-50.



Figure 14. Distribution of selected features over different ViT-B/16 layers for VTAB-1k tasks for results presented in Table 1. We group the layers in encoder (group of 4 activations) to reduce numbers of bars.

K. Additional Plots for Scaling Behaviour of Head2Toe

Scaling behaviour of HEAD2TOE when using different feature target size and number of layers over 19 VTAB-1k tasks is shown in Figure 15.

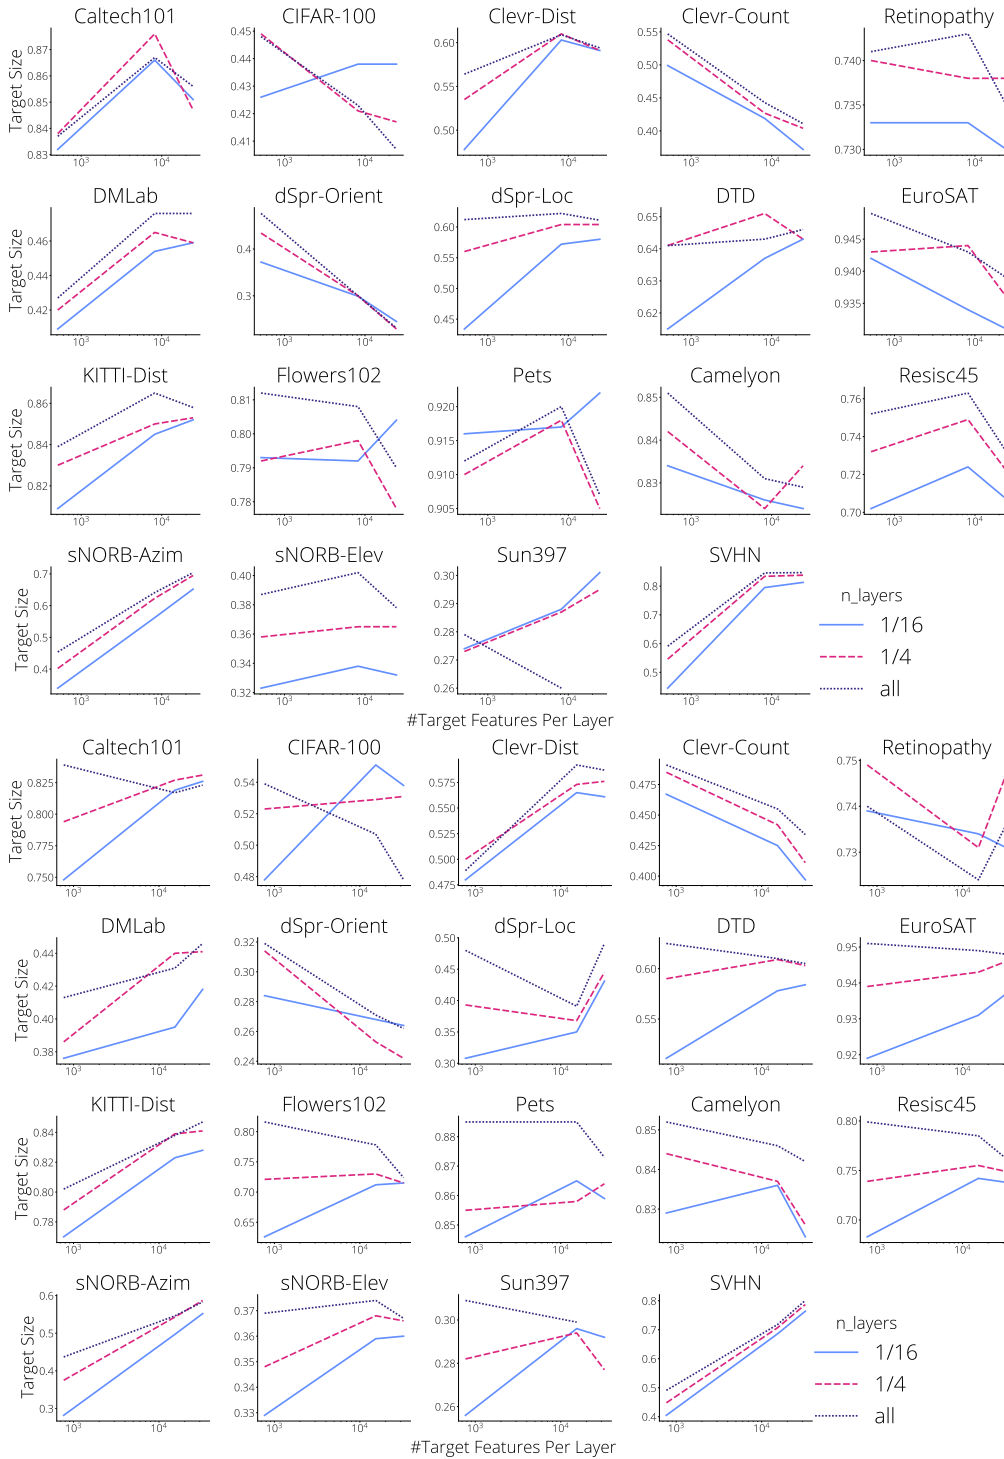


Figure 15. Scaling Behaviour over 19 VTAB-1k tasks when varying feature target size and number of layers utilized for (top) ResNet-50 and (bottom) ViT-B/16.