HEAD2TOE: UTILIZING INTERMEDIATE REPRESENTATIONS FOR BETTER OOD GENERALIZATION

Anonymous authors
Paper under double-blind review

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 later pretrained layers. We explore the hypothesis that these intermediate layers might be directly exploited by linear probing. 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, but critically, for out-of-distribution transfer, Head2Toe outperforms fine-tuning.

1 INTRODUCTION

Tranfer learning is a widely used method for obtaining strong performance in a variety of tasks where training data is scarce—see Zhu et al. (2020); Alyafeai et al. (2020); Zhuang et al. (2020) for recent application-specific surveys. 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, 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, we expect a different linear combination of features to generalize to the target domain, provided the source and target tasks are similar.

This approach of training a new output head, referred to as *linear probing (*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 doing 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*.

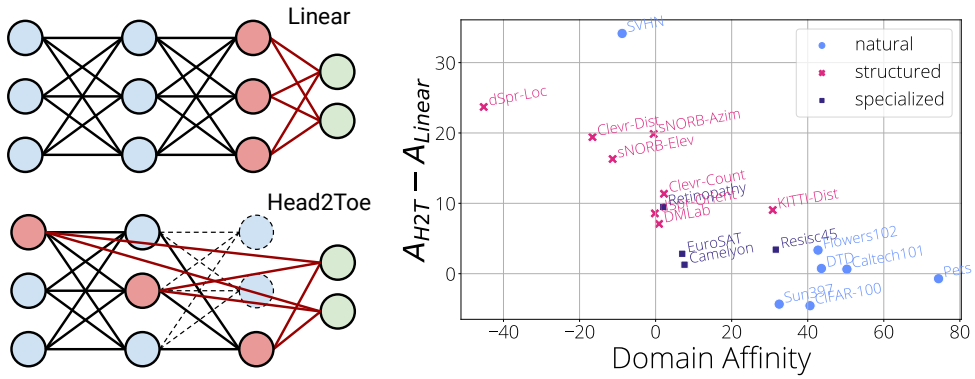

Figure 1: **(left)** HEAD2TOE selects the most useful features from the network and trains a linear head on top. **(right)** HEAD2TOE improves OOD generalization. Domain affinity refers to the difference between linear probe and scratch accuracies (defined in Section 2) and the y-axis represents the improvements (delta accuracy) made by HEAD2TOE over linear probe.

In this work, we propose and explore methods for selecting useful features embedded in internal layers of a pretrained net, concatenating them to the embedding produced by the pretrained net, and then applying the LINEAR transfer approach to the augmented representation (Fig. 1-left). Our approach leads to significant improvements over LINEAR as shown in Fig. 1-right. Our 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.
2. We show that targeted feature selection from intermediate representations is essential to preventing overfitting in transfer learning using few examples (few-shot).
3. Based on the previous observations, we introduce head-to-toe probing (HEAD2TOE), a simple recipe for selecting relevant features for few-shot generalization.
4. On the VTAB collection of data sets, we show that HEAD2TOE outperforms LINEAR and matches the performance of FINETUNING.
5. 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
6. Through ablation studies, we investigate the scaling behaviour of our algorithm and show that increasing the number of candidate features yields better performance.

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 million natural 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. Visual Task Adaptation Benchmark-1k consists of 19 different classification tasks, each having between 2 to 397 classes and a total of 1000 training instances. 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*).

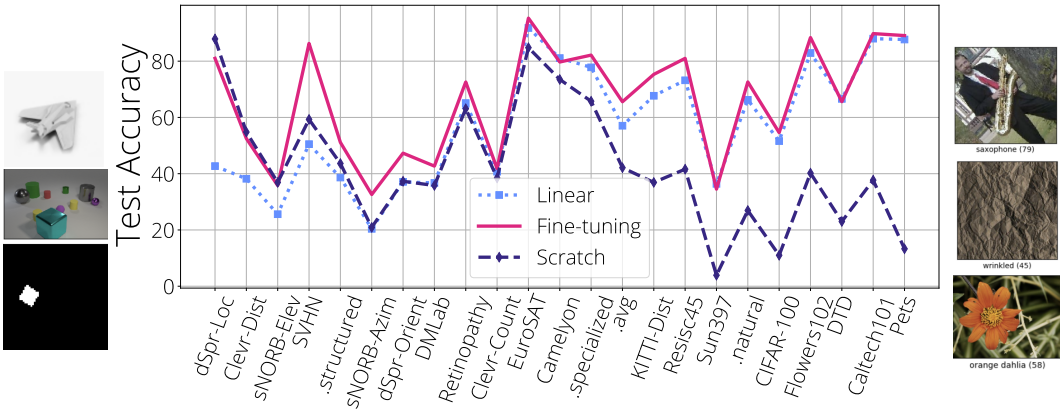

Figure 2: **Characterizing out-of-distribution** Generalization performance of various baselines on the VTAB-1k benchmark using ResNet-50 architecture and 224 image size. The architecture is pretrained on ImageNet-2012 for the transfer learning (TL) baselines. Datasets (and the three groups of the benchmark) 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.

Characterizing out-of-distribution. 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):

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

In Fig. 2, we sort the 19 VTAB-1k target tasks by their domain affinity to ImageNet-2012—from low to high—for a pretrained ResNet-50 backbone. The Figure shows examples of the three target domains with the most and least distribution shift (left and right sides, respectively). These examples seem consistent with intuitive notions of distribution shift.

Baselines. Fig. 2 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 information available in lower layers because it is not needed for the source task but is essential for the target task. FINETUNING makes this information available. The goal of HEAD2TOE is to make this information available to LINEAR by including intermediate features directly.

3 HEAD2TOE UTILIZATION OF PRETRAINED BACKBONES

3.1 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).

In order to estimate an upper bound on how much the performance can be improved over LINEAR by including a single intermediate layer, we use an oracle to select the layer which yields the largest

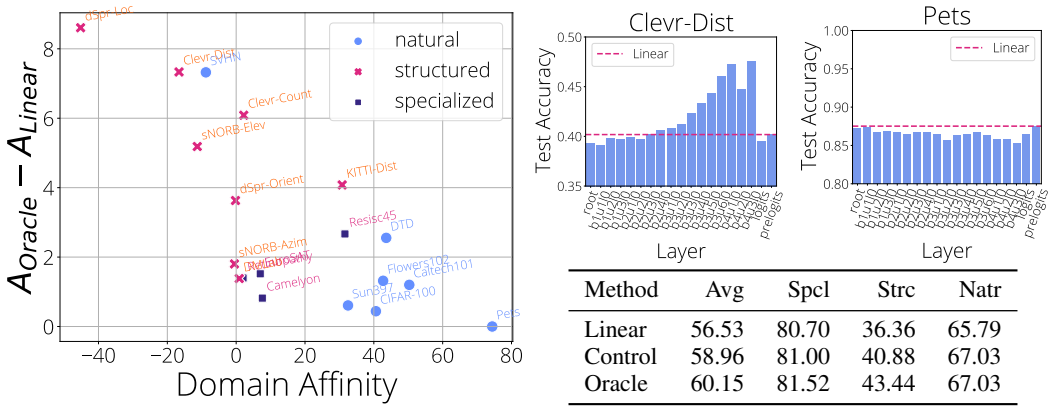

Method	Avg	Spcl	Strc	Natr
Linear	56.53	80.70	36.36	65.79
Control	58.96	81.00	40.88	67.03
Oracle	60.15	81.52	43.44	67.03

Figure 3: **(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 networks trained using a different seed. We use ResNet-50 models pretrained on ImageNet-2012.

boost in test performance (for each target task separately). Percentage improvement over LINEAR using this ORACLE is shown in Fig. 3-left. We observe a negative correlation (-0.745, Spearman) between the domain affinity of a target task and the accuracy gain. As predicted earlier, adding intermediate representations does not improve in-domain generalization because 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 Fig. 3-top-right, we show test accuracy as a function of the ResNet-50 layer whose internal representation is used to augment the feature embeddings to boost LINEAR transfer. Figures for remaining tasks can be found in Appendix B. Different tasks on VTAB benchmark prefer inclusion of different layers to obtain the best generalization accuracy, emphasizing the importance of selecting the right 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 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, we formed a 4096-element feature representation by using a second feature embedding obtained from a second ResNet-50 backbone pretrained on ImageNet-2012 (CONTROL). Note that this experiment bears similarity to *ensembling* (Zhou et al., 2002), which is known to bring significant accuracy gains on its own (Mustafa et al., 2020). Utilizing a second backbone doubles the amount of resources required, yet it falls 1% shy of the Oracle performance demonstrating the extent to which using intermediate representations can be helpful for generalization.

3.2 HEAD2TOE

Motivated by our observations in the previous section, we hypothesize that we can attain—or possibly surpass—the performance of FINETUNING without changing the backbone itself by using LINEAR augmented with well-chosen intermediate activations. However, this approach introduces the problem of selecting features. The number of possible features to include is large relative to the amount of training data available, leading to the potential of overfitting. In this section, we explore various regularization and feature selection methods, which we then incorporate into our method, HEAD2TOE.

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:

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

where the subscript denotes a layer index, $\boldsymbol{h}_0 = \boldsymbol{x}$ is the input, f is the activation function, $\boldsymbol{W}_\ell$ is the weight matrix of layer ℓ, and $\boldsymbol{z}_L$ 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, $\boldsymbol{W}'_L$, such that predictions (logits) for the new task are obtained by $\boldsymbol{z}'_L = \boldsymbol{h}_{L-1}\boldsymbol{W}'_L$.

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

$$\boldsymbol{z}'_L = \boldsymbol{h}_{all}\boldsymbol{W}_{all} \quad ; \quad \boldsymbol{h}_{all} = \text{concat}(a_1(\boldsymbol{h}_1), a_2(\boldsymbol{h}_2), ..., a_L(\boldsymbol{h}_L)) \tag{2}$$

where $a_\ell(.)$ denotes a fixed function to reduce the dimensionality of the activation vector at a given layer ℓ. Such functions become valuable when considering network architectures that generate large number of intermediate features. For example, for convolutional networks, where the representations at each layer are structured as multi-channel spatial arrays of detectors, we perform two-dimensional average pooling as explained in Section 3.1 and Section 5. Even with dimensionality reduction, $\boldsymbol{h}_{all}$ can exceed a million elements, and $\boldsymbol{W}_{all}$ is underconstrained by the training data, leading to overfitting. Further, without care $\boldsymbol{W}_{all}$ may become so large as to be impractical for deploying this model.[1]. We can address these issues by selecting a subset of features before training the target-domain classifier. Equivalently, we can zero out the non-selected rows of $\boldsymbol{W}_{all}$. We will show that appropriate selection of a subset of features leads to better generalization than using all features.

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 $\boldsymbol{W}$, the group-lasso regularizer encourages the ℓ_2 norm of the rows of the matrix to be sparse, and is defined as:

$$|\boldsymbol{W}|_{2,1} = |\boldsymbol{s}|_1 = \sum_i |s_i| \quad ; \quad s_i = \sqrt{\sum_j w_{ij}^2} \tag{3}$$

To determine which features are most useful for the task, the linear head is trained with group-lasso regularization on $\boldsymbol{W}_{all}$. 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, the ℓ_2 norms of the rows of $\boldsymbol{W}_{all}$, $\boldsymbol{s}$, are used to decide which features are most relevant. We refer to s_i as the *relevance score* for feature i. We pick a fraction F of all features with the greatest 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 don't use any additional regularization.

We make two comments on this. First, because the initial round of training $\boldsymbol{W}_{all}$ with the group-lasso regularizer is used only to rank features by importance, the method is quite robust to the regularization coefficient; it simply needs to be large enough to distinguish the contribution of the individual features. Second, interpreting s_i as the importance of feature i depends on all features having the same dynamic range. This constraint will often be satisfied naturally due to the normalization layers found in the architecture itself (we read out representations after normalization). To cover the remaining cases (where no normalization exists between two layers) we normalize the representations of each layer to a unit norm.

Selection of F and the cost of HEAD2TOE. The fraction F determines the total number of features retained. One would expect the optimal value to depend on the target task. We select F based on validation performance. This validation procedure is inexpensive compared to the cost of the initial phase of the algorithm (i.e., training of $\boldsymbol{W}_{all}$ to obtain $\boldsymbol{s}$) 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 $\boldsymbol{W}_{all}$ alone (details shared in Appendix A).

[1]For example, using a pooling size of 2, ResNet-50 generates 1.7 million features and storing $\boldsymbol{W}_{all}$ requires 6.6e8 parameters (2.6GB for float32) for SUN-397.

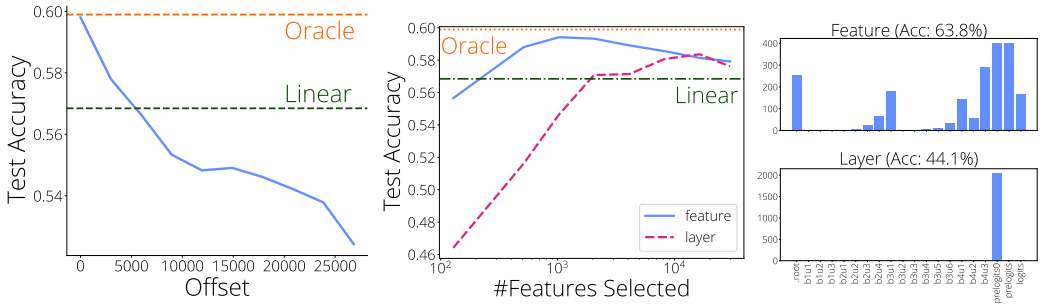

Figure 4: **(left)** Over all VTAB tasks, average accuracy of HEAD2TOE when selecting 2048 consecutive features sorted by their relevance score, starting with an index specified by *offset*. **(center)** Average accuracy over all VTAB tasks as a function of the number of features included. In both experiments, we only use features from the first layer of each of the ResNet blocks and adjust pooling to have around 2048 to features for each layer, totalling 29800 features. **(right)** Distribution of 2048 intermediate features retained from a ResNet-50 when using feature-wise and layer-wise scores on the SVHN transfer task.

Verifying HEAD2TOE. In Fig. 4-left, we demonstrate the effectiveness of group lasso on identifying relevant intermediate features of a ResNet-50 for transfer from ImageNet. We rank all features by their relevance score, s_i, and we 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.

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 that is being used for selection. 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. Fig. 4-center compares feature-wise and layer-wise selection, 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. We share figures for all 19 datasets in Appendix C. Fig. 4-right shows the distribution of features selected over different layers for both strategies.

Note that HEAD2TOE's use of a fixed backbone means that as we search for features to include, the extracted features themselves are not changing. Consequently, we can calculate them once and reuse as necessary, instead of recalculating at every step, as required for FINETUNING. Furthermore, since the backbone is frozen, the only additional weight storage required for each target task is the output head. Such properties are key for transfer-learning methods to be practical when using models with billions of parameters such as GPT-3 (Brown et al., 2020).

4 RELATED WORK

Transfer Learning is studied extensively in the literature and used widely in practice. Peters et al. (2019) compared FINETUNING with LINEAR and reported a small gap in performance for ELMO and BERT models. Kornblith et al. (2019) showed that ImageNet-2012 performance correlates highly with LINEAR and FINETUNING. Using a large pool of pretrained models, Renggli et al. (2020) examined various metrics to decide which backbone to 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), residual adapters (Houlsby et al., 2019; Rebuffi et al., 2017), 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. Zhang et al. (2021) showed that

selecting a sub-network of a pretrained model can increase OOD generalization. However, none of these methods match the simplicity of training (and storing) a linear classifier. 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.

Alternatives to transfer learning exist to learn from few labeled examples. Meta-learning (Thrun & Pratt, 1998; Schmidhuber, 1987; Schmidhuber et al., 1997) is used (among other things) to tackle the few-shot learning problem (i.e learning from limited data) from a learning-to-learn perspective; see Hospedales et al. (2020) for a comprehensive survey. In-between the representation learning and learning-to-learn paradigms lie 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 aims to find the most relevant features from the input and is studied extensively in the machine-learning literature. Approaches can be grouped according to whether labeled data is used—supervised (Nie et al., 2010) or unsupervised (He et al., 2005; Balın et al., 2019; Atashgahi et al., 2020)—or what high-level approach is taken—filter methods (Blum & Langley, 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 while achieving the best performance 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 selection (Viola & Jones, 2001; Borboudakis & Tsamardinos, 2019), backward selection (pruning) (Mozer & Smolensky, 1989; Guyon et al., 2004) and regularization/feature-ranking 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 setting. We point interested readers to the survey of Gui et al. (2017) and book of Boln-Canedo et al. (2015) for a detailed discussion on the topic.

5 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. 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 number features they contribute to the concatenated embedding. 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 and works better than normalizing each feature separately or not normalizing at all. 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 might not be the best choice here and better aggregation function can provide further gains. We share a detailed list of intermediate representations utilized for each architecture in Appendix F.

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 without doing feature selection. These baselines are denoted as +All-ℓ_1, +All-ℓ_2 and +All-$\ell_{2,1}$ according to the norm they use.

We perform five-fold cross validation for each task and method in order to pick the best hyper-parameters. 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 size of the aggregated representation at

	Natural							Structured				Specialized								
	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	Mean
Linear	49.2	86.4	68.2	84.6	88.8	47.6	35.7	**83.3**	92.4	73.2	74.0	41.3	39.0	37.1	66.7	31.8	30.4	17.9	26.7	56.5
+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**
Fine-tuning*	**54.6**	**89.8**	65.6	**88.4**	**89.1**	**86.3**	34.5	79.7	**95.3**	**81.0**	72.6	41.8	52.5	42.7	**75.3**	81.0	47.3	32.6	35.8	65.6
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

	Natural							Structured				Specialized								
	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	Mean
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-$\ell_{2,1}$	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
Fine-tuning	62.6	83.1	61.5	80.4	86.8	83.0	33.7	**83.2**	94.9	78.0	73.6	58.3	59.2	**43.6**	39.8	61.4	**44.8**	27.1	26.2	62.2
Fine-tuning+	**65.2**	82.1	60.5	85.2	**86.9**	**86.7**	**36.9**	82.4	95.2	**81.3**	73.0	**68.2**	**60.6**	41.7	**73.9**	**69.3**	28.0	25.7	36.1	**65.2**

Table 1: Median test accuracy over 3 seeds on the VTAB-1k benchmark using pretrained (**top**) ResNet-50 and (**bottom**) ViT-B/16 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 code and checkpoints provided by Dosovitskiy et al. (2021).

each layer. The FINETUNING baseline searches over 4 hyper-parameters; 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 than training a linear classifier, even with intermediate features. We repeat each evaluation using 3 different seeds and report median values. Standard deviations in our results, more details on hyper-parameter selection and best values are shared in Appendix A.

5.1 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 FINETUNING (see rightmost column of Table). 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 FINETUNING for the *natural* and *structured* categories. More plots and the distribution of features selected from each layer can be found in Appendix D.

5.2 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. The story for HEAD2TOE versus FINETUNING is a bit more complicated. The FINETUNING recipe used by Dosovitskiy et al. (2021) includes many tweaks such as selective data-augmentation, image upsampling, learning rate warm-up, and gradient clipping. In

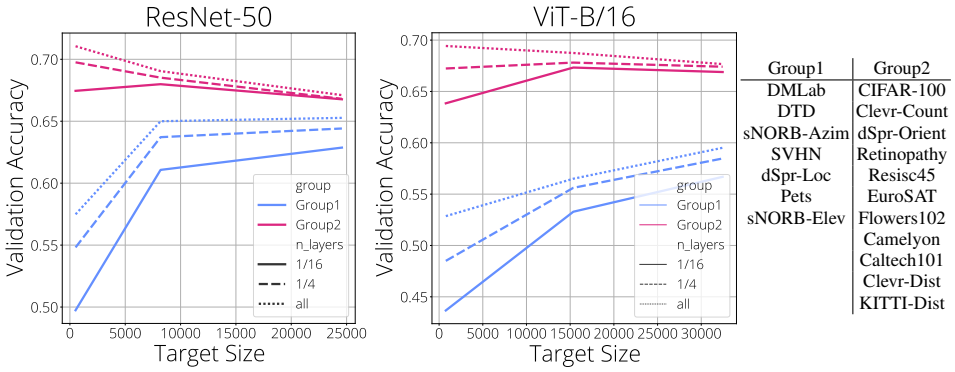

Figure 5: Effect of increasing the number of intermediate features that HEAD2TOE uses from the **(left)** ResNet-50 backbone and **(right)** ViT-B/16 backbone. The abscissa of the graph indicates the dimensionality of the representation extracted from each layer of the backbone (*target size*). The tasks 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 E.

order to make a direct comparison with LINEAR, in the next-to-last row of the Table we report the performance of FINETUNING without all of these tweaks; HEAD2TOE obtains about a 1% absolute accuracy improvement over FINETUNING in this setting. However, when some of the tweaks are incorporated into FINETUNING—in particular gradient clipping—FINETUNING achieves a 2% improvement over HEAD2TOE. That being said, HEAD2TOE retains its competitive advantage in terms of resource costs and its accuracy would likely similarly benefit from the development of more sophisticated tricks (notably moving beyond the naive 1-D pooling of ViT features), much like FINETUNING has gained over time.

5.3 SCALING HEAD2TOE

In Fig. 5 we vary the number of intermediate features used for each of our pretrained backbones. We observe that including all layers always performs better on average. However when varying the number of target features for each layer, we observed 2 distinct sub-group of target tasks that behave differently as the number of features increases. This observation informed our discussion to include both small and large target sizes in our validation hyper parameter search.

6 CONCLUSION

In this work, we introduced HEAD2TOE, an approach that extends linear probing (LINEAR) by selecting the most relevant features among a pretrained network's intermediate representations. We show that doing so greatly improves performance over LINEAR and allows the approach to reach a performance competitive with—and in some cases superior to—FINETUNING. Our findings challenge the conventional belief that FINETUNING is required to achieve good performance on OOD tasks. While more work is needed before HEAD2TOE can realize the full computational benefits of linear probing, our work paves the way for applying new and more efficient feature selection approaches and for experimenting with HEAD2TOE probing in other domains such as regression, reinforcement learning, and language modelling domains.

REFERENCES

Zaid Alyafeai, Maged Saeed AlShaibani, and Irfan Ahmad. A survey on transfer learning in natural language processing. *arXiv preprint arXiv:2007.04239*, 2020.

Andreas Argyriou, Theodoros Evgeniou, and Massimiliano Pontil. Multi-task feature learning. In B. Schölkopf, J. Platt, and T. Hoffman (eds.), *Advances in Neural Information Processing Systems*. MIT Press, 2007.

Zahra Atashgahi, Ghada Sokar, Tim van der Lee, Elena Mocanu, Decebal Constantin Mocanu, Raymond N. J. Veldhuis, and Mykola Pechenizkiy. Quick and robust feature selection: the strength of energy-efficient sparse training for autoencoders. *ArXiv*, abs/2012.00560, 2020.

R. Baldock, Hartmut Maennel, and Behnam Neyshabur. Deep learning through the lens of example difficulty. *ArXiv*, abs/2106.09647, 2021.

Muhammed Fatih Balın, Abubakar Abid, and James Zou. Concrete autoencoders: Differentiable feature selection and reconstruction. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.

Hakan Bilen and Andrea Vedaldi. Universal representations: The missing link between faces, text, planktons, and cat breeds. *arXiv preprint arXiv:1701.07275*, 2017.

Avrim L. Blum and Pat Langley. Selection of relevant features and examples in machine learning. *Artificial Intelligence*, 1997.

Vernica Boln-Canedo, Noelia Snchez-Maroo, and Amparo Alonso-Betanzos. *Feature Selection for High-Dimensional Data*. Springer Publishing Company, Incorporated, 1st edition, 2015. ISBN 3319218573.

Giorgos Borboudakis and I. Tsamardinos. Forward-backward selection with early dropping. *J. Mach. Learn. Res.*, 20:8:1–8:39, 2019.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, T. J. Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeff Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *ArXiv*, abs/2005.14165, 2020.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 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`.

Vincent Dumoulin, Neil Houlsby, Utku Evci, Xiaohua Zhai, Ross Goroshin, Sylvain Gelly, and Hugo Larochelle. Comparing transfer and meta learning approaches on a unified few-shot classification benchmark. *arXiv preprint arXiv:2104.02638*, 2021.

Nikita Dvornik, Cordelia Schmid, and Julien Mairal. Selecting relevant features from a multi-domain representation for few-shot classification. In *European Conference on Computer Vision*, pp. 769–786. Springer, 2020.

Cheng Fu, Hanxian Huang, Xinyun Chen, Yuandong Tian, and Jishen Zhao. Learn-to-share: A hardware-friendly transfer learning framework exploiting computation and parameter sharing. In *ICML*, 2021.

Jie Gui, Zhenan Sun, Shuiwang Ji, Dacheng Tao, and Tieniu Tan. Feature selection based on structured sparsity: A comprehensive study. *IEEE Transactions on Neural Networks and Learning Systems*, 28:1490–1507, 2017.

Demi Guo, Alexander M. Rush, and Yoon Kim. Parameter-efficient transfer learning with diff pruning. In *ACL/IJCNLP*, 2021.

Yunhui Guo, Humphrey Shi, Abhishek Kumar, Kristen Grauman, Tajana Simunic, and Rogério Schmidt Feris. Spottune: Transfer learning through adaptive fine-tuning. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4800–4809, 2019.

Isabelle Guyon, Jason Weston, Stephen D. Barnhill, and Vladimir Naumovich Vapnik. Gene selection for cancer classification using support vector machines. *Machine Learning*, 46:389–422, 2004.

Xiaofei He, Deng Cai, and Partha Niyogi. Laplacian score for feature selection. In *Proceedings of the 18th International Conference on Neural Information Processing Systems*, 2005.

Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. Meta-learning in neural networks: A survey. *arXiv preprint arXiv:2004.05439*, 2020.

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *ICML*, 2019.

Yigitcan Kaya, Sanghyun Hong, and T. Dumitras. Shallow-deep networks: Understanding and mitigating network overthinking. In *ICML*, 2019.

Ron Kohavi and George H. John. Wrappers for feature subset selection. *Artif. Intell.*, 1997.

Simon Kornblith, Jonathon Shlens, and Quoc V. Le. Do better imagenet models transfer better? *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2656–2666, 2019.

Wei-Hong Li, Xialei Liu, and Hakan Bilen. Improving task adaptation for cross-domain few-shot learning. *arXiv preprint arXiv:2107.00358*, 2021a.

Wei-Hong Li, Xialei Liu, and Hakan Bilen. Universal representation learning from multiple domains for few-shot classification. *arXiv preprint arXiv:2103.13841*, 2021b.

Lu Liu, William Hamilton, Guodong Long, Jing Jiang, and Hugo Larochelle. A universal representation transformer layer for few-shot image classification. In *International Conference on Learning Representations*, 2021.

Michael C Mozer and Paul Smolensky. 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`.

Basil Mustafa, Carlos Riquelme, Joan Puigcerver, andAndr'e Susano Pinto, Daniel Keysers, and Neil Houlsby. Deep ensembles for low-data transfer learning. *ArXiv*, abs/2010.06866, 2020.

Feiping Nie, Heng Huang, Xiao Cai, and Chris Ding. Efficient and robust feature selection via joint $\ell_{2,1}$-norms minimization. In J. Lafferty, C. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta (eds.), *Advances in Neural Information Processing Systems*, 2010.

Matthew E. Peters, Sebastian Ruder, and Noah A. Smith. To tune or not to tune? adapting pretrained representations to diverse tasks. In *RepL4NLP@ACL*, 2019.

Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. *arXiv preprint arXiv:1705.08045*, 2017.

Cédric Renggli, André Susano Pinto, Luka Rimanic, J. Puigcerver, Carlos Riquelme, Ce Zhang, and Mario Lucic. Which model to transfer? finding the needle in the growing haystack. *ArXiv*, abs/2010.06402, 2020.

Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 2015.

Jürgen Schmidhuber. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. PhD thesis, Technische Universität München, 1987.

Jürgen Schmidhuber, Jieyu Zhao, and Marco Wiering. Shifting inductive bias with success-story algorithm, adaptive levin search, and incremental self-improvement. *Machine Learning*, 28(1): 105–130, 1997.

Surat Teerapittayanon, Bradley McDanel, and H. T. Kung. Branchynet: Fast inference via early exiting from deep neural networks. *2016 23rd International Conference on Pattern Recognition (ICPR)*, pp. 2464–2469, 2016.

Sebastian Thrun and Lorien Pratt. Learning to learn: Introduction and overview. In *Learning to learn*, pp. 3–17. Springer, 1998.

Eleni Triantafillou, Hugo Larochelle, Richard Zemel, and Vincent Dumoulin. Learning a universal template for few-shot dataset generalization. In *International Conference on Machine Learning*, 2021.

Paul A. Viola and Michael J. Jones. 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.

Songtao Wu, Shenghua Zhong, and Yan Liu. Deep residual learning for image steganalysis. *Multimedia Tools and Applications*, 2018.

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

Xiaohua Zhai, J. Puigcerver, Alexander Kolesnikov, P. Ruyssen, Carlos Riquelme, Mario Lucic, Josip Djolonga, André Susano Pinto, Maxim Neumann, A. Dosovitskiy, L. Beyer, Olivier Bachem, M. Tschannen, Marcin Michalski, O. Bousquet, S. Gelly, and N. Houlsby. The visual task adaptation benchmark. *ArXiv*, abs/1910.04867, 2019.

Dinghuai Zhang, Kartik Ahuja, Yilun Xu, Yisen Wang, and Aaron C. Courville. Can subnetwork structure be the key to out-of-distribution generalization? In *ICML*, 2021.

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

Wangchunshu Zhou, Canwen Xu, Tao Ge, Julian McAuley, Ke Xu, and Furu Wei. Bert loses patience: Fast and robust inference with early exit. *ArXiv*, abs/2006.04152, 2020.

Zhi-Hua Zhou, Jianxin Wu, and Wei Tang. Ensembling neural networks: Many could be better than all. *Artificial Intelligence*, 2002. URL `https://www.sciencedirect.com/science/article/pii/S000437020200190X`.

Zhuangdi Zhu, Kaixiang Lin, and Jiayu Zhou. Transfer learning in deep reinforcement learning: A survey. *arXiv preprint arXiv:2009.07888*, 2020.

Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1): 43–76, 2020.

A VALIDATION PROCEDURE FOR DIFFERENT ALGORITHMS

We perform hyper-parameter validation for each VTAB task separately. 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 hyper-parameter 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.

Cost of HEAD2TOE We evaluate different values of F and pick the value with best validation performance. Cost of HEAD2TOE consists of two parts: (1) C_{all}: cost of training the initial head $\boldsymbol{W}_{all}$ (in order to obtain s_i's) (2) $\sum_f C_{F=f}$: 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 ($\boldsymbol{s}$)).

B ADDITIONAL PLOTS FOR EXPERIMENTS USING SINGLE ADDITIONAL LAYER

Test accuracies when using a single additional intermedieate layer from a pretrained ResNet-50 backbone are shown in Fig. 6.

C ADDITIONAL PLOTS FOR LAYER/FEATURE-WISE SELECTION COMPARISON

We compare layer-wise selection strategy discussed in Section 3 to HEAD2TOE in Fig. 7.

D ADDITIONAL PLOTS FOR SCALING BEHAVIOUR OF HEAD2TOE

Improvement over fine-tuning test accuracy for ResNet-50 backbone is shown in Fig. 8-left according to their domain affinity. Similar to earlier plots, we observe a clear correlation between being OOD and benefits of using HEAD2TOE: HEAD2TOE obtains superior few-shot generalization for most of OOD tasks, with the exception of *dSpr-Loc*.

E 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 Fig. 9.

F DETAILED LIST OF INTERMEDIATE REPRESENTATIONS

	dataset	Target Size	F	Learning Rate	Training Steps	$\ell_{2,1}$ Coefficient
7	Caltech101	8192	0.010	0.01	5000	0.00001
0	CIFAR-100	512	0.100	0.01	500	0.00001
8	Clevr-Dist	8192	0.001	0.01	500	0.00100
1	Clevr-Count	512	0.005	0.10	5000	0.00100
9	Retinopathy	8192	0.100	0.01	500	0.00001
10	DMLab	8192	0.020	0.01	500	0.00001
2	dSpr-Orient	512	0.100	0.01	500	0.00001
11	dSpr-Loc	8192	0.005	0.10	500	0.00100
16	DTD	24576	0.005	0.01	5000	0.00001
3	EuroSAT	512	0.100	0.01	500	0.00001
12	KITTI-Dist	8192	0.020	0.01	500	0.00001
4	Flowers102	512	0.100	0.01	5000	0.00001
13	Pets	8192	0.002	0.01	5000	0.00001
5	Camelyon	512	0.020	0.10	500	0.00100
14	Resisc45	8192	0.020	0.01	500	0.00001
17	sNORB-Azim	24576	0.002	0.01	500	0.00001
15	sNORB-Elev	8192	0.050	0.01	500	0.00100
6	Sun397	512	0.100	0.01	5000	0.00100
18	SVHN	24576	0.005	0.01	500	0.00001

	dataset	Target Size	F	Learning Rate	Training Steps	$\ell_{2,1}$ Coefficient
0	Caltech101	768	0.050	0.01	5000	0.00100
1	CIFAR-100	768	0.020	0.01	500	0.00001
12	Clevr-Dist	15360	0.002	0.01	500	0.00100
2	Clevr-Count	768	0.050	0.10	5000	0.00001
3	Retinopathy	768	0.010	0.01	500	0.00100
14	DMLab	32448	0.005	0.01	500	0.00001
4	dSpr-Orient	768	0.100	0.01	5000	0.00001
15	dSpr-Loc	32448	0.002	0.10	500	0.00100
5	DTD	768	0.100	0.01	500	0.00100
6	EuroSAT	768	0.100	0.01	500	0.00100
16	KITTI-Dist	32448	0.050	0.01	5000	0.00100
7	Flowers102	768	0.020	0.01	500	0.00001
8	Pets	768	0.020	0.01	5000	0.00100
9	Camelyon	768	0.100	0.01	500	0.00100
10	Resisc45	768	0.050	0.01	5000	0.00001
17	sNORB-Azim	32448	0.010	0.01	500	0.00001
13	sNORB-Elev	15360	0.100	0.01	500	0.00100
11	Sun397	768	0.050	0.01	5000	0.00100
18	SVHN	32448	0.005	0.01	500	0.00001

Table 2: Hyper parameters selected for the VTAB-1k benchmark tasks when using pretrained (**top**) ResNet-50 or (**bottom**) ViT-B/16 backbones. F refers to fraction of features kept.

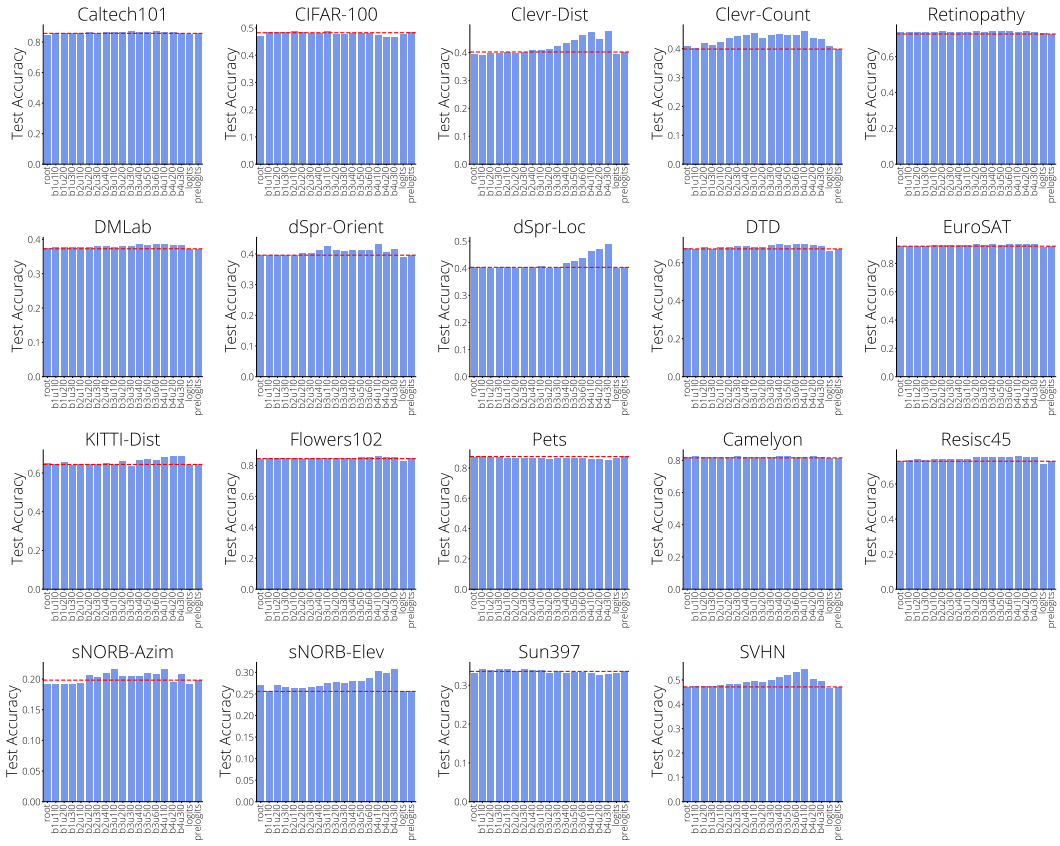

Figure 6: Test accuracies when using a single additional intermedieate layer from a pretrained ResNet-50 backbone.

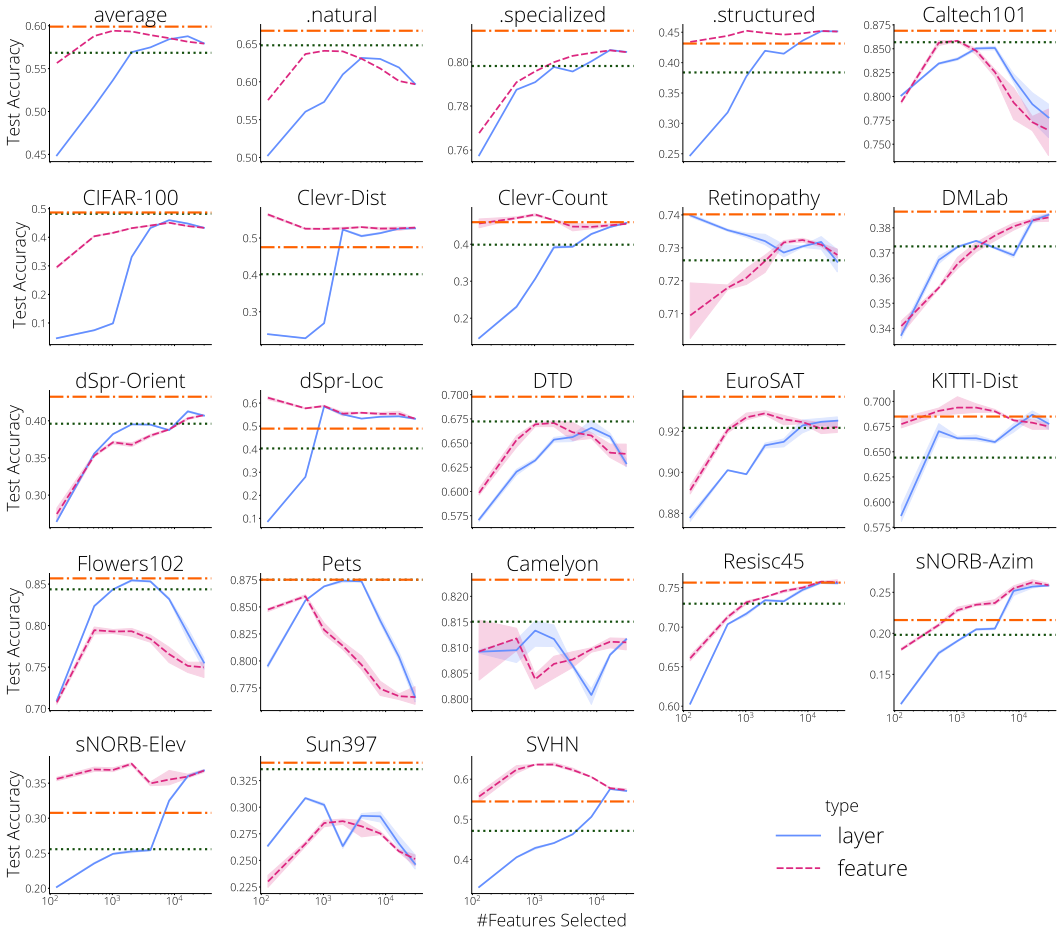

Figure 7: Test accuracies when varying the number of features selected for HEAD2TOE using a pretrained ResNet-50 backbone.

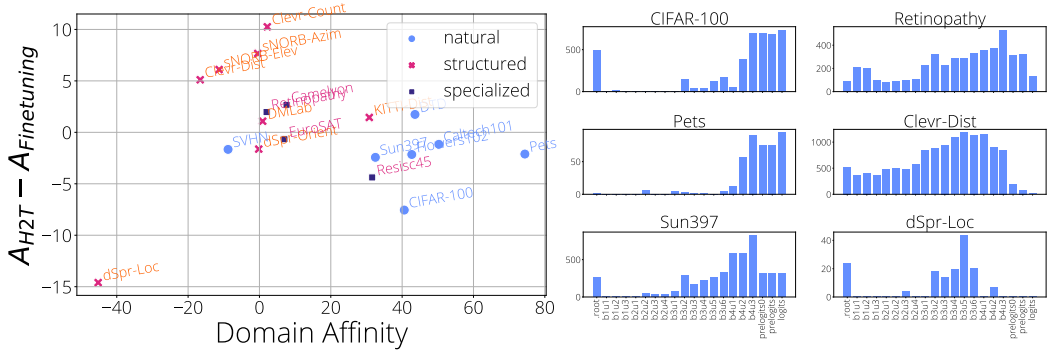

Figure 8: **(left)** Accuracy improvement of HEAD2TOE compared to FINETUNING **(right)** Distribution of selected features over different ResNet-50 layers for 6 different VTAB-1k tasks. We group the layers in each block (group of 3 layers) to reduce numbers of bars.

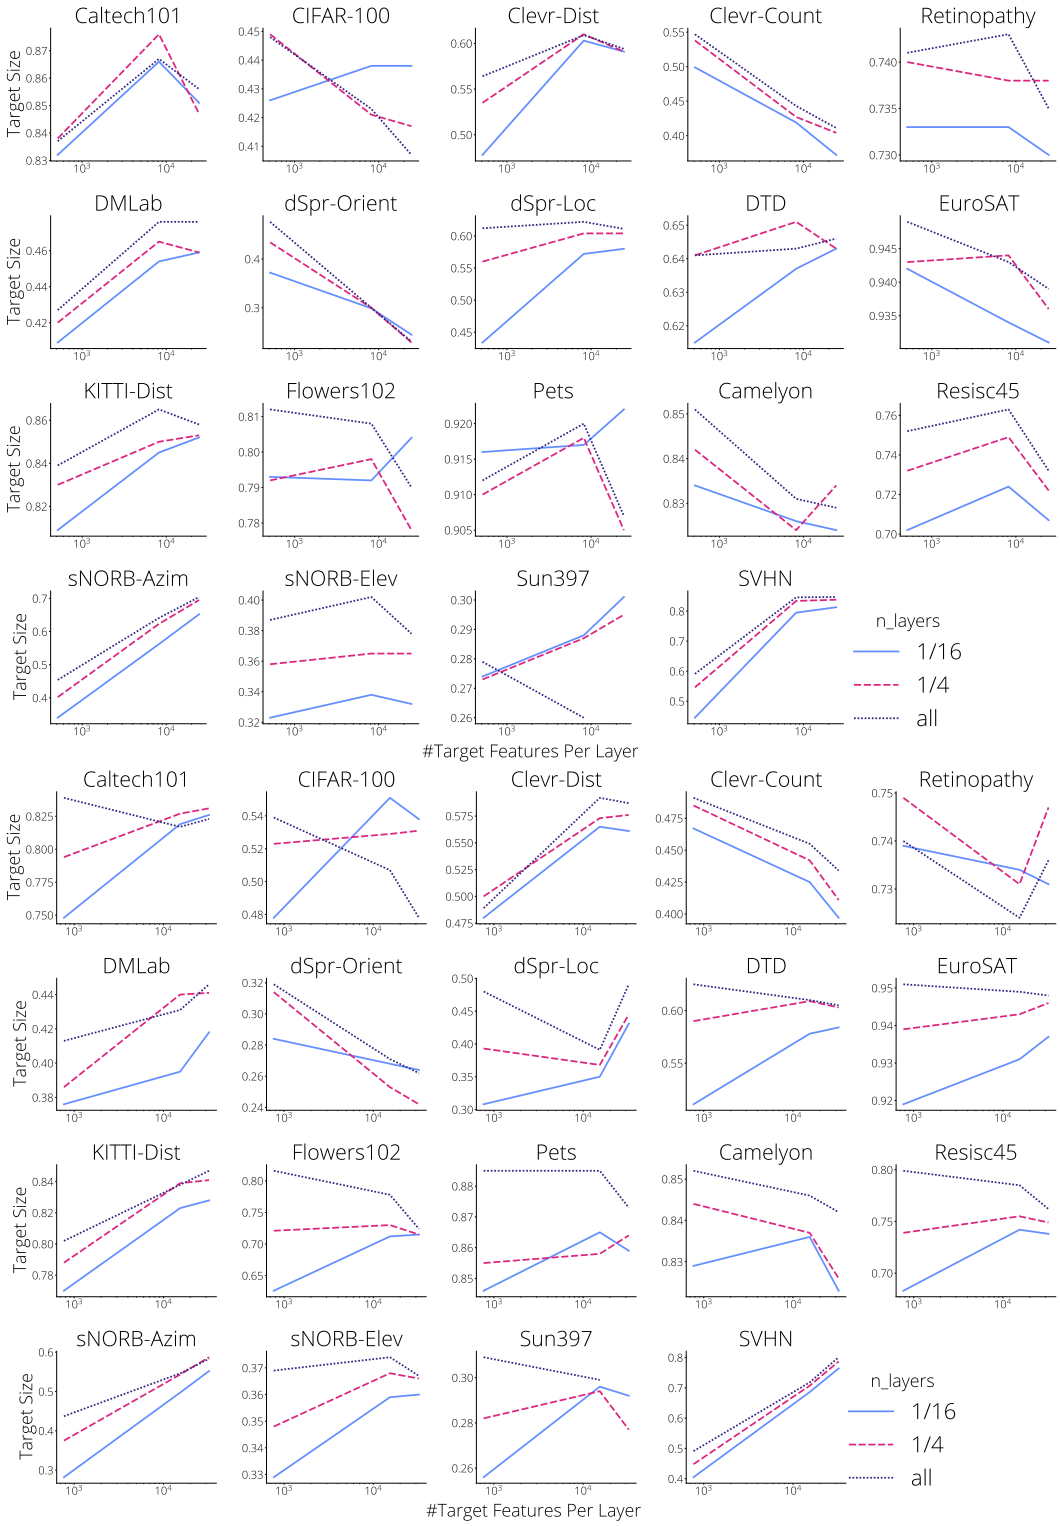

Figure 9: Scaling Behaviour over 19 VTAB-1k tasks when varying feature target size and number of layers utilized for **(top)** ResNet-50 and **(bottom)** ViT-B/16)