
Knowledge Composition using Task Vectors with Learned Anisotropic Scaling

Frederic Z. Zhang* Paul Albert* Cristian Rodriguez-Opazo
Anton van den Hengel Ehsan Abbasnejad

Australian Institute for Machine Learning The University of Adelaide

{firstname.lastname}@adelaide.edu.au

<https://github.com/fredzzhang/atlas>

Abstract

Pre-trained models produce strong generic representations that can be adapted via fine-tuning on specialised datasets. The learned weight difference relative to the pre-trained model, known as a task vector, characterises the direction and stride of fine-tuning that enables the model to capture these specialised representations. The significance of task vectors is such that simple arithmetic operations on them can be used to combine diverse representations from different domains. This paper builds on these properties of task vectors and aims to answer (1) whether components of task vectors, particularly parameter blocks, exhibit similar characteristics, and (2) how such blocks can be used to enhance knowledge composition and transfer. To this end, we introduce aTLAS, an algorithm that linearly combines parameter blocks with different learned coefficients, resulting in anisotropic scaling at the task vector level. We show that such linear combinations explicitly exploit the low intrinsic dimensionality of pre-trained models, with only a few coefficients being the learnable parameters. Furthermore, composition of parameter blocks enables modular learning that effectively leverages the already learned representations, thereby reducing the dependency on large amounts of data. We demonstrate the effectiveness of our method in task arithmetic, few-shot recognition and test-time adaptation, with supervised or unsupervised objectives. In particular, we show that (1) learned anisotropic scaling allows task vectors to be more disentangled, causing less interference in composition; (2) task vector composition excels with scarce or no labelled data and is less prone to domain shift, thus leading to better generalisability; (3) mixing the most informative parameter blocks across different task vectors prior to training can reduce the memory footprint and improve the flexibility of knowledge transfer. Moreover, we show the potential of aTLAS as a parameter-efficient fine-tuning method, particularly with less data, and demonstrate that it can be easily scaled up for higher performance.

1 Introduction

One practical advantage of neural networks is the fact that knowledge learned from a previous problem, in the form of network weights, can be transferred to solve other related problems. Commonly referred to as transfer learning [6, 73], this technique is often applied when a model trained on a general-purpose dataset—ImageNet [52] for many years—is fine-tuned on other datasets to improve

*Equal contribution. Listed order was determined by a coin toss. Fred formalised the idea; developed the original codebase; conducted experiments on task arithmetic; and drafted the paper. Paul designed and conducted experiments on few-shot recognition, test-time adaptation, parameter-efficient fine-tuning; investigated numerous properties of task vector compositions; and was crucially involved in every stage of the work.

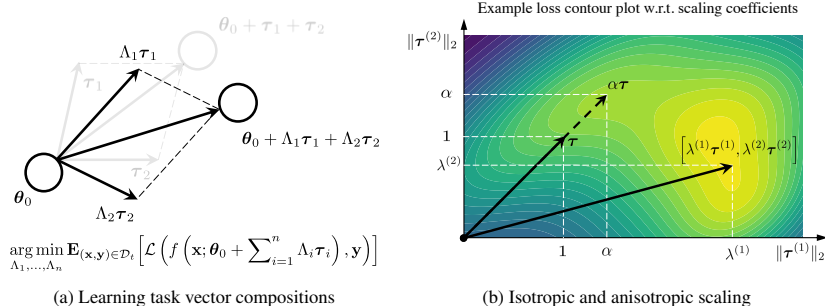


Figure 1: Illustration of (a) learning task vector compositions ($n = 2$, θ_0 denotes the weights of a pre-trained model) and (b) the flexibility of anisotropic scaling. Assume a task vector $\tau = (\tau^{(1)}, \tau^{(2)})$ has two parameter blocks, learning anisotropic scaling grants more flexibility when combining task vectors.

performance on downstream problems. In the past, classification models [18, 53] have been used as the medium for such knowledge transfer, which played a crucial part in the success of detection and segmentation [7, 19, 51, 66–68]. In recent years, foundation models [4] trained on broad data, CLIP [47] particularly, have demonstrated strong performance on a multitude of tasks, even when applied in a zero-shot manner. Besides the conventional way of exploiting the knowledge in these models via fine-tuning, recent works [28, 44, 62] have presented more direct measures to manipulate the network weights. In particular, Ilharco et al. [28] showed that, a task vector, defined as the weight difference between a pre-trained and a fine-tuned model, can be used as a carrier of the task-specific knowledge learned via fine-tuning. As such, multiple task vectors, when combined with simple arithmetic, can form a multi-task model that largely retains its performance across all fine-tuning tasks. Linearisation techniques [44], in addition, have been shown to further enhance this compositionality.

Intrigued by this phenomenon, we investigate the potential of task vectors being knowledge carriers in this paper, by learning linear combinations of them (Figure 1a) for various problems. In particular, parameter blocks, e.g., weights and biases, tend to encode different learned representations in different layers. We thus learn an independent scaling coefficient per block for more precise adjustments tailored to the unique roles of each parameter block. This results in anisotropic scaling of task vectors (Figure 1b), and allows us to exploit their modularity in knowledge composition, granting higher controllability when steering the behaviours of a model for task arithmetic [28].

The potential applications of task vector composition extend beyond model editing. With the coefficients being the only learnable parameters, our method exploits the rich knowledge encapsulated in the task vectors by searching in a low-dimensional coefficient space. As a result, it is a competitive parameter-efficient fine-tuning (PEFT) method, and is particularly effective in cases where labelled data is scarce. This offers new opportunities for few-shot learning [34, 69] and test-time adaptation [35, 57]. Furthermore, for multi-purpose models such as CLIP [47], variants of the model trained with different data sources or fine-tuned on different downstream tasks are often available [26]. These resources constitute a significant knowledge bank, with task vectors being the knowledge carrier. Many learning problems may be simplified to learning a combination of task vectors.

Our primary contribution is a learning algorithm named aTLAS, wherein otherwise complex learning problems can be framed as learning linear combinations of task vectors. The algorithm is broadly applicable to optimising supervised and unsupervised objectives. Its effectiveness is demonstrated in task arithmetic, few-shot recognition, test-time adaptation and parameter-efficient fine-tuning, where we show that (1) learning linear combinations of task vectors directly exploits the low intrinsic dimensionality of pre-trained models [1, 33], resulting in a small number of learnable parameters; (2) standard task vectors, otherwise inferior to linearised variants [44] in task arithmetic, can produce stronger multi-task models with learned anisotropic scaling; (3) aTLAS is effective in low-data regimes, and improves the accuracy of CLIP by 6.5 absolute points averaged over 22 datasets with unlabelled data; (4) aTLAS is complementary to previous few-shot adaptation methods, in that one third of the examples it improves upon are unique; (5) aTLAS as a few-shot learning method is less prone to domain shift, and achieves better generalisation on out-of-domain datasets; (6) the most informative parameter blocks from different task vectors can be mixed prior to training, allowing for flexible and efficient knowledge transfer under memory constraints; (7) aTLAS is a strong PEFT method when data is limited, and existing PEFT methods such as low-rank adaptations (LoRA) [23] can be seamlessly integrated into aTLAS to improve memory efficiency.

2 Models and task vectors

As Ilharco et al. [28] demonstrated, task vectors exhibit many intriguing properties across a wide range of models, such as CLIP [47], GPT-2 [46] and T5-based models [48]. To facilitate more in-depth experimentation and analysis, we focus on the CLIP model in this paper, due to its wide availability and manageable size. In particular, we follow previous practice [28, 44] and acquire task vectors by fine-tuning the image encoder, with the text representations frozen. This ensures that image encoders fine-tuned on different datasets produce features residing in the same representation space, through a common text encoder. The task vectors obtained from these fine-tuned encoders can thus be combined more effectively to form a unified multi-task model.

Formally, denote the CLIP image encoder by $f : \mathcal{X} \times \Theta \rightarrow \mathcal{Z}$, such that for input image $\mathbf{x} \in \mathcal{X}$ and parameters $\theta \in \Theta$, $\mathbf{z} = f(\mathbf{x}; \theta)$ is the learned latent representation for the input image. Denote the weights of a pre-trained model by θ_0 , and the weights of its fine-tuned variant by $\theta_i, i \in \mathbb{N}^+$, where i indexes a dataset \mathcal{D}_i . We follow Ilharco et al. [28] and define a task vector as $\tau_i = \theta_i - \theta_0$. In addition, we investigate task vectors produced by linearised variants of the image encoder using the first-order Taylor expansion,

$$g(\mathbf{x}; \theta) := f(\mathbf{x}; \theta_0) + (\theta - \theta_0)^\top \nabla_{\theta} f(\mathbf{x}; \theta_0). \quad (1)$$

Ortiz-Jiménez et al. [44] showed that, task vectors obtained from fine-tuning the linearised variants have low disentanglement errors, and exhibit strong compositional properties.

3 Learning task vector compositions

Parameters in a neural network, depending on the depth of the layer, often have different significance. For instance, early layers in convolutional neural networks [18, 53] are known for extracting generic, low-level features, such as edges, corners, etc., while deeper layers produce features more specific to the task. We recognise the non-uniform impacts parameters at different layers can have, and do not perform isotropic scaling on task vectors. Instead, weights, biases and any other forms of parameterisation, which we collectively refer to as *parameter blocks*, will be scaled independently.

3.1 Proposed method: aTLAS

Formally, denote a task vector with m parameter blocks by $\tau = (\tau^{(1)}, \dots, \tau^{(m)})$, where each parameter block $\tau^{(j)}$ is vectorised, and round brackets denote column vector concatenation. We learn a block diagonal matrix Λ , parameterised as

$$\Lambda = \begin{bmatrix} \lambda^{(1)} I^{(1)} & \dots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \dots & \lambda^{(m)} I^{(m)} \end{bmatrix}, \quad (2)$$

where $\lambda^{(j)} \in \mathbb{R}$ is a learnable coefficient; $I^{(j)}$ denotes an identity matrix with its number of columns matching the dimension of $\tau^{(j)}$; and the superscript $j \in \mathbb{N}^+$ indexes a parameter block. This results in anisotropic scaling of a task vector, that is,

$$\Lambda_i \tau_i = \left(\lambda_i^{(1)} \tau_i^{(1)}, \dots, \lambda_i^{(m)} \tau_i^{(m)} \right), \quad (3)$$

where the subscript $i \in \mathbb{N}^+$ indexes a task vector. As such, assuming a supervised objective, finding the optimal composition of task vectors can be defined as the following optimisation problem

$$\arg \min_{\Lambda_1, \dots, \Lambda_n} \mathbf{E}_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_t} \left[\mathcal{L} \left(f(\mathbf{x}; \theta_0 + \sum_{i=1}^n \Lambda_i \tau_i), \mathbf{y} \right) \right], \quad (4)$$

where \mathcal{L} is the loss function for a target task; n is the number of task vectors; \mathbf{y} is the labels corresponding to inputs \mathbf{x} ; \mathcal{D}_t denotes a target dataset. The number of learnable parameters, as a result, is precisely mn . Let us denote the solution to the aforementioned optimisation problem by $\{\Lambda_i^*\}_{i=1}^n$. In inference, model $f(\mathbf{x}, \theta_0 + \sum_{i=1}^n \Lambda_i^* \tau_i)$ will be deployed, which incurs no additional computational cost compared to models trained in the conventional way.




Figure 2: Recognition accuracy versus the number of bases when optimising in a low-dimensional subspace. The accuracy is normalised by that of the fully fine-tuned model. Using task vectors to construct the projection matrix performs consistently better than using random bases on (a) MNIST [32], (b) CIFAR100 [31].

In addition, we investigate the task vectors obtained from fine-tuning linearised variants of the model, i.e., $g(x)$ in Eq. 1. Denote such task vectors by $\tilde{\tau}$. The learning objective with linearised task vectors can be derived as follows

$$\arg \min_{\Lambda_1, \dots, \Lambda_n} \mathbf{E}_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_t} \left[\mathcal{L} \left(f(\mathbf{x}; \boldsymbol{\theta}_0) + \left(\sum_{i=1}^n \Lambda_i \tilde{\tau}_i \right)^\top \nabla_{\boldsymbol{\theta}} f(\mathbf{x}; \boldsymbol{\theta}_0), \mathbf{y} \right) \right]. \quad (5)$$

3.2 Relation to intrinsic dimensionality

A notable characteristic of aTLAS is its parameter efficiency. To offer more intuitions, we refer to previous findings [1, 33] that deep neural networks often produce solutions residing in a subspace with much lower intrinsic dimensionality. This is measured by finding a minimum number of d parameters, such that learning these parameters ($\hat{\boldsymbol{\theta}} \in \mathbb{R}^d$) leads to approximately the same performance as optimising in the full parameter space ($\boldsymbol{\theta} \in \mathbb{R}^D$). This can be expressed as follows

$$\boldsymbol{\theta} = \boldsymbol{\theta}_0 + P\hat{\boldsymbol{\theta}}, \quad (6)$$

where $\boldsymbol{\theta}_0 \in \mathbb{R}^D$ denotes the pre-trained weights and $P \in \mathbb{R}^{D \times d}$ is a random projection matrix. We demonstrate that learning task vector compositions leads to the same formulation. For brevity of exposition, let us consider compositions at the block level. For the j -th parameter block, we have

$$\boldsymbol{\theta}^{(j)} = \boldsymbol{\theta}_0^{(j)} + \sum_{i=1}^n \lambda_i^{(j)} \boldsymbol{\tau}_i^{(j)} \quad (7)$$

$$= \boldsymbol{\theta}_0^{(j)} + \underbrace{\left[\boldsymbol{\tau}_1^{(j)}, \dots, \boldsymbol{\tau}_n^{(j)} \right]}_{\text{projection matrix}} \underbrace{\left[\lambda_1^{(j)}, \dots, \lambda_n^{(j)} \right]^\top}_{\text{learnable parameters}}. \quad (8)$$

We draw a parallel between Eqs. 6 and 8 and note that aTLAS explicitly exploits the low intrinsic dimensionality by learning a small set of coefficients. The number of task vectors, i.e., n , is much smaller than the dimension of weight vector $\boldsymbol{\theta}_i^{(j)}$, and is analogous to the intrinsic dimensionality d . However, as opposed to using a random projection matrix P , aTLAS constructs the projection matrix from task vectors, making use of the learned representations. To demonstrate its advantage, we use the same number of bases for task vectors² and random bases³, and show that task vectors consistently achieve higher performance in Figure 2. These results solidify our understanding of task vectors being knowledge carriers. We thus set out to apply aTLAS to various applications.

4 Task arithmetic

Task arithmetic [28] is comprised of a few tasks aimed at editing pre-trained models using task vectors. Following previous practice [28, 44], we conduct experiments under the settings of task negation and task addition on eight image classification datasets (details included in Appendix A).

²A fixed number of task vectors are selected based on the blockwise gradient. Details can be found in Section 5.2 and Appendix D.6.

³Each random basis of the projection is drawn from a Gaussian distribution with the mean and standard deviation to match those of the pre-trained weights in the corresponding parameter block, i.e., $\boldsymbol{\theta}_0^{(j)}$.

Table 1: Performance of task negation averaged across eight datasets. Selected results must maintain at least 95% of the pre-trained accuracy on the control dataset, following previous practice [44]. Best performance in each section is highlighted in bold. Task vector is abbreviated as t.v. Results for each dataset are available in Table 7.

T.V.	Methods	Models	ViT-B/32		ViT-B/16		ViT-L/14	
			Target (↓)	Control (↑)	Target (↓)	Control (↑)	Target (↓)	Control (↑)
n/a	Pre-trained	$f(\mathbf{x}; \theta_0)$	48.14	63.35	55.48	68.33	64.89	75.54
Std.	Search	$f(\mathbf{x}; \theta_0 + \alpha\boldsymbol{\tau})$	23.22	60.71	19.38	64.66	19.15	72.05
	aTLAS (ours)	$f(\mathbf{x}; \theta_0 + \Lambda\boldsymbol{\tau})$	18.76	61.21	17.34	65.84	17.75	73.28
Lin.	Search	$g(\mathbf{x}; \theta_0 + \alpha\tilde{\boldsymbol{\tau}})$	11.54	60.74	10.88	65.54	12.78	72.95
	aTLAS (ours)	$g(\mathbf{x}; \theta_0 + \Lambda\tilde{\boldsymbol{\tau}})$	11.06	61.02	10.16	65.58	12.61	73.14

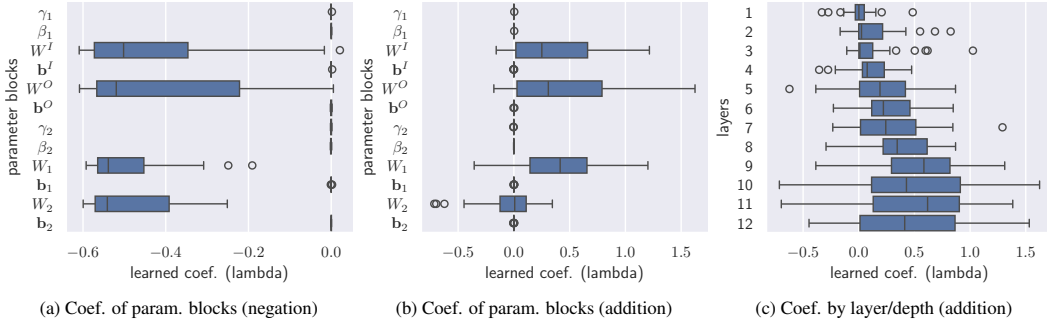


Figure 3: Box-and-whisker plots for the learned coefficients. As each transformer layer consists of a fixed set of parameter blocks, we visualise the distribution of coefficients for these parameter blocks across all layers, for (a) task negation and (b) task addition, as well as (c) distribution of coefficients by layer. We denote the learnable LayerNorm parameters by γ and β . Weights and biases are denoted by W and \mathbf{b} , respective, with attention layer parameters indexed by superscripts and the MLP parameters indexed by subscripts.

Previous works acquire the optimal isotropic scaling factor on task vectors via a hyper-parameter search on validation sets. As such, we learn anisotropic scaling matrices on the same validation sets, and visualise the learned coefficients to shed light on this mechanism.

4.1 Task negation

Task negation aims to reduce undesired biases, characterised by the performance, on a target task, while maintaining performance on a control dataset, ImageNet [52] in this case. Denote the validation sets for the target and control tasks by \mathcal{D}_t and \mathcal{D}_c , respectively. We perform a simultaneous gradient ascent on the target task and gradient descent on the control task, described as follows,

$$\arg \min_{\Lambda_t} \mathbf{E}_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_t} [-\mathcal{L}(f(\mathbf{x}; \theta_0 + \Lambda_t \boldsymbol{\tau}_t), \mathbf{y})] + \mathbf{E}_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_c} [\mathcal{L}(f(\mathbf{x}; \theta_0 + \Lambda_t \boldsymbol{\tau}_t), \mathbf{y})], \quad (9)$$

where $\boldsymbol{\tau}_t$ is the task vector for the target dataset, and cross-entropy loss is used. The learning objectives with linearised task vectors can be derived easily based on Eq. 5, and so are omitted.

We summarise the task negation results in Table 1, and show that our method significantly improves upon standard task vectors, while the improvement upon linear task vectors is less prominent. In particular, we observe that weights matrices tend to have much larger negative coefficients, as shown in Figure 3a. To investigate this, we instead only learn coefficients for the weight matrices, with zero coefficients on other parameter blocks, effectively reducing the number of learnable parameters by two thirds. With ViT-B/32 as the backbone, we observe an average accuracy of 20.14 (vs. 18.76) on target tasks and 61.23 (vs. 61.21) on the control task, which shows that weight matrices carry majority of the knowledge required for task negation.

4.2 Task addition

Task addition aims at producing a multi-task model using task vectors acquired from a range of datasets. We utilise task vectors from the eight image classification datasets, and learn the anisotropic

Table 2: Performance of task addition averaged across eight datasets. We report the absolute accuracy (Abs.) and the relative accuracy (Rel.) with respect to the fine-tuned model. Best performance in each section is highlighted in bold. Task vector is abbreviated as t.v. Results for each dataset are available in Table 8.

T.V.	Methods	Models	ViT-B/32		ViT-B/16		ViT-L/14	
			Abs. (↑)	Rel. (↑)	Abs. (↑)	Rel. (↑)	Abs. (↑)	Rel. (↑)
n/a	Pre-trained	$f(\mathbf{x}; \theta_0)$	48.14	-	55.48	-	64.89	-
Std.	Search	$f(\mathbf{x}; \theta_0 + \alpha \sum_i \tau_i)$	70.12	77.24	73.63	79.85	82.93	87.92
	aTLAS (ours)	$f(\mathbf{x}; \theta_0 + \sum_i \Lambda_i \tau_i)$	84.98	93.79	86.08	93.44	91.36	97.07
Lin.	Search	$g(\mathbf{x}; \theta_0 + \alpha \sum_i \tilde{\tau}_i)$	74.67	85.17	77.51	86.21	84.75	91.86
	aTLAS (ours)	$g(\mathbf{x}; \theta_0 + \sum_i \Lambda_i \tilde{\tau}_i)$	83.42	95.42	85.38	95.10	88.65	96.12




Figure 4: Disentanglement errors between each pair of datasets. Each row reflects the percentage of data in the corresponding dataset that have altered predictions after combining two task vectors. Our method achieves stronger task addition performance as a result of less interference amongst task vectors.

scaling matrices with the objectives described in Eqs. 4, 5 using the cross-entropy loss. The training data is comprised of the validation sets for all eight dataset, i.e., $\mathcal{D}_t = \bigcup_{i=1}^8 \mathcal{D}_i$.

Performance comparison against previous methods is shown in Table 2, where our method yields substantial improvements. Interestingly, we note that with previous methods [28, 44], linear task vectors outperform the standard ones in terms of absolute accuracy, while the converse is true with our method. To investigate this, we compute the pairwise disentanglement error ξ [44], which measures the percentage of data with inconsistent predictions when two task vectors are combined (more details in Appendix C.2). Results in Figure 4 show that standard task vectors with learned anisotropic scaling achieve the lowest average error, indicating less interference in task vector composition. Along with higher fine-tuning accuracy, previously referred to as the non-linear advantage [44], standard task vectors demonstrate stronger performance in task addition.

Furthermore, we again observe that weight matrices have consistently larger coefficients in Figure 3b, and learning coefficients on weight matrices alone results in an accuracy of 84.17 (vs. 84.98) using ViT-B/32. This suggests that weight matrices in transformers are the primary knowledge carrier, which enabled knowledge composition and negation. Note that for better clarity in visualisation, we add L_1 regularisation on the learned coefficients during learning, which causes marginal performance drop (84.23 vs. 84.98) but significantly improves interpretability. In addition, we observe substantially higher coefficients on deeper layers (Figure 3c). This aligns with our understanding that early layers extract generic features that do not vary significantly across datasets [29], while the deeper layers produce task-specific features and require more careful adaptations.

5 Knowledge transfer in low-data regimes

Beyond model editing for task arithmetic, we explore the idea of transferring existing knowledge in task vectors to previously unseen tasks. To this end, we use the CLIP [47] model and a total of 22 image classification datasets, each of which produces a task vector. We defer the details of datasets and the process to acquire task vectors to Appendix A. Denote the set of available task vectors by $T = \{\tau_i\}_{i=1}^n$, and the dataset corresponding to task vector τ_i by \mathcal{D}_i . For each target dataset \mathcal{D}_t , we




Figure 5: Few-shot experiment results averaged across 22 datasets and three seeds, showing (a) comparison against state-of-the-art few-shot methods with ViT-B/32 backbone and (b) percentage of images in the validation sets that become correctly classified after applying few-shot methods. We also show (c) performance difference compared to pre-trained CLIP model on OOD datasets. More detailed results are included in Appendix D.

learn task vector compositions using the subset $T \setminus \{\tau_t\}$, excluding the task vector for the target dataset to avoid information leakage. We test our method in few-shot and test-time adaptation, to demonstrate its effectiveness in low-data regimes. Notably, we observe that task vectors complement existing few-shot methods. Combining aTLAS with them thus leads to significant improvements.

5.1 Few-shot adaptation

Few-shot recognition requires learning new objects or concepts using a limited amount labelled data— k per class for k -shot. Following previous practice [69], we approach this problem by adapting a pre-trained CLIP model [47] to each target dataset \mathcal{D}_t . We use the subset of task vectors $T \setminus \{\tau_t\}$ and $k \in \{1, 2, 4, 8, 16\}$ images from dataset \mathcal{D}_t . During training, we adopt the cross-entropy loss and minimise objectives described in Eqs. 4 and 5 for standard and linear task vectors, respectively.

We compare against Tip-Adapter [69] and LP++ [25] using CLIP with ViT-B/32 backbone, across 22 datasets over three random seeds, and summarise the results in Figure 5a. We show that with $k = 1$, our approach, aTLAS, significantly outperforms previous methods, demonstrating the effectiveness of knowledge transfer with scarce labelled data. More importantly, we note that the idea of task vector composition is highly complementary to those presented in previous methods. As such, combining aTLAS with them results in significant improvements. This is also illustrated in Figure 5b as a Venn diagram, where we show the percentage of examples in the validation set that are incorrectly classified by the pre-trained model but correctly classified with few-shot methods. Out of the examples aTLAS improves upon, around half are unique compared against either Tip-Adapter or LP++, demonstrating its complementarity. We also found that standard task vectors generally perform better than their linearised counterparts, and so defer the results of linear task vectors to Appendix D.2.

In addition, due to the low number of learnable parameters, aTLAS exhibits strong generalisability. To demonstrate this, we learn task vector composition on ImageNet [52], and test it on out-of-domain (OOD) datasets: ImageNet-A [22], ImageNet-R [21], ImageNet-sketch [60] and ImageNetV2 [50]. We summarise the results in Figure 5c, which shows the performance difference against the pre-trained model. Notably, aTLAS is the only method that consistently improves upon the pre-trained model on OOD datasets, and combining aTLAS with other methods can improve their generalisability.

We also test our method and variants integrated with Tip-Adapter and LP++ using other backbones, including ViT-{B/16, L/14} and ResNet-{50, 101}, and find that the results are consistent with those for ViT-B/32. More details can be found in Appendix D.3.

5.2 Task vector budget and selection

In practical applications, there may only be a limited number of task vectors available, or the number of task vectors used in training may be restricted due to memory constraints. To this end, we study the influence of task vector budget b on few-shot recognition performance. We experiment with four selection strategies: (1) random selection; (2) feature-based selection; (3) gradient-based selection; and (4) blockwise gradient-based selection. To elaborate, feature-based selection computes the mean image feature representation of each dataset, and selects b task vectors from datasets most similar

Table 3: Test-time adaptation accuracy averaged over 22 dataset, with $\times 1$ standard error over 3 random seeds. LN refers to tuning the LayerNorm layers. CLIP with the ViT-B/32 backbone is used. Highest performance is highlighted in bold.

Method	Zero-shot	Contrastive (SimCLR)		Entropy (SAR)		Pseudo labelling (UFM)	
		LN	aTLAS	LN	aTLAS	LN	aTLAS
Accuracy	60.4	60.4 \pm 0.0	62.7 \pm 0.1	61.2 \pm 0.1	62.9 \pm 0.0	62.2 \pm 0.1	66.9 \pm 0.1

to the target dataset. Gradient-based selection computes the gradient with respect to each of the learnable coefficients, and either select entire task vectors with the highest L_1 gradient norm, or select task vectors with the highest blockwise gradient for the corresponding parameter block, and repeat the process for all parameter blocks. The blockwise selection therefore allows parameter blocks across different task vectors to be mixed prior to training. More details can be found in Appendix D.6.

For a task vector budget $b \in \{1, 2, 5, 10, 15, 21\}$, we summarise the few-shot recognition performance in Figure 6. First, we note that the accuracy of aTLAS does not plateau with the maximum number of task vectors available (21), indicating that more task vectors could be beneficial. Second, we find that selecting task vectors based on feature similarity is a simple yet effective approach with sufficient budgets ($b > 5$). Selecting whole task vectors with gradient is less effective, generally on par with random selection. Nevertheless, the blockwise variant achieves the best accuracy, particularly for very low budgets ($b \in \{1, 2\}$), as it is able to exploit knowledge from more task vectors than the budget dictates. We thus deduce that parameter blocks can function as knowledge carriers in isolation, independent of the task vectors to which they belong. In fact, a parameter block $\tau^{(1)}$ as part of the task vector $\tau = (\tau^{(1)}, \dots, \tau^{(m)})$ can be considered as a task vector by itself, i.e., $(\tau^{(1)}, \mathbf{0}, \dots, \mathbf{0})$. This modular nature underscores the potential of task vectors for flexible and efficient knowledge transfer.

5.3 Test-time adaptation

Test-time adaptation (TTA) [35, 57, 59] assumes no labelled data is available for the target task, requiring the model to adapt in an unsupervised fashion. We conduct experiments under the offline adaptation setting, which allows access to the target dataset. We consider three categories of self-supervised techniques for TTA: contrastive objectives, entropy objectives and pseudo labelling. Contrastive objectives align representations of the same image under different data augmentations. For this category, we adopt SimCLR [9], a simple yet effective method. Entropy objectives encourage the pre-trained model to produce confident predictions on unseen datasets by minimising the entropy over the predictions. This technique was previously explored by Yang et al. [65] in model merging. While effective in simpler cases, it can lead to catastrophic collapse in TTA. Therefore, we utilise a state-of-the-art sharpness-aware entropy minimisation algorithm named SAR [43]. Last, we experiment with an unsupervised pseudo-labelling algorithm inspired by FixMatch [54], which we refer as unsupervised FixMatch (UFM). UFM selects an equal number of highly confident examples per class as the labelled set, and then employs FixMatch to produce pseudo-labels from rest of the unlabelled examples. Details are available in Appendix E.

We summarise the results in Table 3 and compare our method, i.e., learning task vector compositions, against the conventional approach of tuning the layer normalisation parameters [43, 57, 59]. We show that under all self-supervised objectives, aTLAS achieves higher accuracy than tuning the LayerNorm. In particular, LayerNorm has 30k learnable parameters with ViT-B/32 while our method only has 3.5k learnable parameters. We note that with the UFM objective, aTLAS performs the best and improves the accuracy by an average of 6.5 absolute points over the zero-shot baseline.

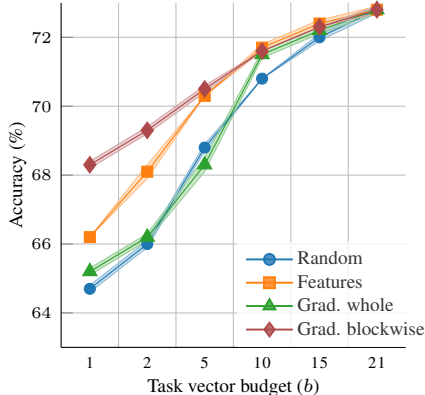


Figure 6: Few-shot performance of aTLAS with various task vector budgets. The accuracy is averaged across 22 datasets and over three random seeds. Standard deviation $\times 1$ is overlaid as the error margin. Performance under the 16-shot setting is visualised, while additional detailed results are included in Table 14.

Table 4: Few-shot recognition performance using standard task vectors or LoRAs as sparse task vectors. Results are averaged across 22 datasets over three seeds, with $\times 1$ standard deviation. The memory consumption for ViT-B/32 backbone is annotated under each variant. For standard task vectors, we learn compositions on all parameter blocks or weight matrices only. For LoRAs as task vectors, we report results with rank 4, 16 and 64.

Shots (k)	Standard task vectors		LoRAs as task vectors		
	All parameter blocks (10.7 GB)	Weight matrices (10.5 GB)	Rank=4 (3.3 GB)	Rank=16 (3.4 GB)	Rank=64 (4.1 GB)
1	66.0 \pm 0.2	66.0 \pm 0.1	64.4 \pm 0.1	64.6 \pm 0.1	65.4 \pm 0.1
2	67.7 \pm 0.1	67.0 \pm 0.2	65.7 \pm 0.0	66.6 \pm 0.2	67.4 \pm 0.1
4	70.0 \pm 0.0	69.4 \pm 0.2	68.2 \pm 0.0	68.7 \pm 0.1	69.5 \pm 0.2
8	71.3 \pm 0.1	70.9 \pm 0.0	70.2 \pm 0.2	70.4 \pm 0.1	70.9 \pm 0.1
16	72.8 \pm 0.1	72.3 \pm 0.0	71.7 \pm 0.1	71.8 \pm 0.1	72.0 \pm 0.1

6 Relation to parameter-efficient fine-tuning

One of the key advantages of aTLAS is its ability to adapt pre-trained models with few learnable parameters, making it suitable for parameter-efficient fine-tuning (PEFT). Similar to popular PEFT methods such as low-rank adaptation (LoRA) [23], our approach does not introduce additional modules, thereby avoiding an increase in inference complexity. In addition, since only the encoded weight matrices in LoRAs have non-zero weight difference, LoRAs are in fact sparse task vectors. They can thus be seamlessly integrated into our method, significantly reducing the memory cost.

6.1 LoRAs as task vectors

Due to the sparsity and rank deficiency, LoRAs as task vectors may have limited representation capacity and carry less knowledge. Therefore, they may be inferior to standard task vectors for knowledge transfer. We investigate this by learning linear combinations of LoRAs⁴ using our method, under the settings of few-shot recognition. Results are summarised in Table 4. We first shed light on the impact of sparsity, and compare two variants of our method that either learns linear combinations of all parameter blocks or just the weight matrices. Results show that sparsity results in an accuracy decrease of around 0.5% on average, except for the one-shot setting. The rank deficiency, on the other hand, causes more substantial accuracy drop. Nevertheless, this can be largely mitigated by increasing the rank. Using a rank of 64 leads to similar performance compared to learning compositions of only weight matrices in standard task vectors. In conclusion, while the sparsity and rank deficiency introduce some performance drops, especially in low-shot settings, LoRAs are competitive alternatives to standard task vectors due to their low memory cost.

6.2 Scalability of aTLAS

Despite the parameter efficiency of aTLAS, its performance is not as competitive when sufficient training data is available. To address this, we devise a strategy to flexibly scale up the number of learnable parameters as needed. Specifically, we randomly divide each parameter block into K partitions, and assign a learnable coefficient to each partition, naturally increasing the number of learnable parameters by K -fold. We denote these variants by aTLAS $\times K$. We conduct experiments with these variants using $\{1, 5, 10, 25, 35, 50, 100\}$ % of the total available training data across the 22 datasets used in Section 5. The results are summarised in Figure 7, showing that our method consistently improves as K increases. Compared to LoRAs, particularly with limited training data,




Figure 7: Scalability of aTLAS. We compare the accuracy of our method against LoRAs, and vary the amount of training data. Results are averaged over 22 datasets. Detailed results are included in Table 17.

⁴Details about the process to acquire LoRAs are included in Appendix D.7.

our method achieves higher performance with fewer learnable parameters. With sufficient training data, the variant aTLAS $\times 1200$ leads to higher performance with a similar number of learnable parameters, as it is able to exploit the knowledge contained in the task vectors that may otherwise be unobtainable from the target dataset.

7 Related work

Task vectors and model compositions. Recent studies have demonstrated the possibility of manipulating the behaviours of neural networks directly in the weight space [27, 62, 64]. In particular, task vectors [28], as a carrier of the domain-specific knowledge learned through fine-tuning, exhibit strong compositional properties. Such compositionality can be enhanced via linearisation using first-order Taylor expansion [44], and improves model editing with simple arithmetic, e.g., addition, negation, etc. Yang et al. [65] also investigated the idea of learning layer-wise coefficients to improve task arithmetic. In addition, low-rank adaptations [23], as special forms of task vectors, were shown to also support such arithmetic operations. A recent study [3] also investigated the idea of learning combinations of LoRAs for few-shot recognition.

Model-based transfer learning. One interpretation of transfer learning [73] is to exploit the knowledge encapsulated in a pre-trained model for a target domain. Amongst various sub-modules of a pre-trained model, transferring the feature extractor is the most extensively studied. This ranges from early convolutional neural networks [18, 53] to modern transformers [58], from vision backbones [14, 37] to language models [13, 46]. For vision applications, classification models trained on ImageNet [52] have been used as the medium for knowledge transfer. In recent years, contrastively pre-trained multi-modal models such as CLIP [47] have emerged as a prevalent choice. Such models are trained on large volumes of data by aligning image and language representations, leading to strong baselines well suited for transfer learning. CLIP representations have since been used for medical imaging [70], semantic segmentation [72], satellite imaging [40], etc.

Model adaptation in low-data regimes. The performance of pre-trained models is often constrained when applied to specific tasks with limited labelled data. To address this limitation, extensive research has been conducted on few-shot adaptation of CLIP [47]. These studies focus on various techniques, including prompt engineering [71], feature adaptation [16], and more recently classifier adaptation [25, 69]. In addition to few-shot adaptation, test-time adaptation represents an even more challenging scenario where no annotated data is available. This typically requires leveraging self-supervised objectives to adapt the model, employing methods such as entropy minimisation [35, 43, 59], contrastive learning [8], pseudo labelling [35] and image rotation prediction [57].

8 Conclusion

In this paper, we introduced aTLAS, a learning algorithm that leverages the rich knowledge encapsulated in task vectors through learned linear combinations with anisotropic scaling. Unlike conventional methods that learn network parameters, our approach focuses on learning coefficients on task vectors, significantly reducing the number of learnable parameters. We conducted experiments across task arithmetic, few-shot recognition, test-time adaptation and parameter-efficient fine-tuning, demonstrating the effectiveness of our method with supervised and unsupervised objectives. In particular, we highlighted several properties of aTLAS, including low disentanglement error, robustness against domain shift, effectiveness in low-data regimes, complementarity with existing few-shot methods, etc. These properties paved the way for efficient knowledge composition and transfer.

Limitations. As a task vector is defined with respect to a specific pre-trained model, knowledge composition and transfer are not yet feasible across different architectures. This may become possible with suitable projections and remains part of the future work. In addition, combining large numbers of task vectors can consume a substantial amount of GPU memory when training larger models. This can be mitigated by selecting a subset of task vectors, using LoRAs as task vectors or by offloading the computation of task vector composition to CPU, at the cost of training speed decrease. It is also possible to perform task vector composition at bit-width lower than floating point precision, e.g., 4-bit. Similar features are being tested with popular deep learning frameworks such as PyTorch, and we expect the memory requirement of larger models to be less of a constraint in the future.

Acknowledgements. This research is funded in part by the Australian Government through the Australian Research Council (Project DP240103278), and the Centre of Augmented Reasoning at the Australian Institute for Machine Learning, established by a grant from the Department of Education. We would like to thank Stephen Gould for his valuable feedback on the paper.

References

- [1] Armen Aghajanyan, Sonal Gupta, and Luke Zettlemoyer. Intrinsic dimensionality explains the effectiveness of language model fine-tuning. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, pages 7319–7328. Association for Computational Linguistics, Aug 2021. URL <https://aclanthology.org/2021.acl-long.568>. 2, 4
- [2] Eric Arazo, Diego Ortego, Paul Albert, Noel E. O’Connor, and Kevin McGuinness. Pseudo-Labeling and Confirmation Bias in Deep Semi-Supervised Learning. In *International Joint Conference on Neural Networks (IJCNN)*, 2020. URL <https://arxiv.org/pdf/1908.02983>. 29
- [3] Nader Asadi, Mahdi Beitollahi, Yasser Khalil, Yinchuan Li, Guojun Zhang, and Xi Chen. Does combining parameter-efficient modules improve few-shot transfer accuracy? *arXiv preprint arXiv:2402.15414*, 2024. URL <https://arxiv.org/pdf/2402.15414>. 10
- [4] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ B. Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri S. Chatterji, Annie S. Chen, Kathleen Creel, Jared Quincy Davis, Dorottya Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah D. Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark S. Krass, Ranjay Krishna, Rohith Kuditipudi, and et al. On the opportunities and risks of foundation models. *arXiv*, abs/2108.07258, 2021. URL <https://arxiv.org/abs/2108.07258>. 2
- [5] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *Proceedings of European Conference on Computer Vision*, pages 446–461. Springer, 2014. URL https://link.springer.com/chapter/10.1007/978-3-319-10599-4_29. 16
- [6] Stevo Bozinovski and Ante Fulgosi. The influence of pattern similarity and transfer learning upon the training of a base perceptron b2. *Proceedings of Symposium Informatica*, 3(125), 1976. 1
- [7] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Proceedings of European Conference on Computer Vision (ECCV)*, pages 213–229, Cham, 2020. Springer International Publishing. URL <https://arxiv.org/pdf/2005.12872>. 2
- [8] Dian Chen, Dequan Wang, Trevor Darrell, and Sayna Ebrahimi. Contrastive test-time adaptation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 295–305, Jun 2022. URL <https://arxiv.org/pdf/2204.10377>. 10
- [9] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A Simple Framework for Contrastive Learning of Visual Representations. In *International Conference on Machine Learning (ICML)*, 2020. URL <https://arxiv.org/pdf/2002.05709>. 8
- [10] Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of IEEE*, 105(10):1865–1883, Oct 2017. URL <https://arxiv.org/abs/1703.00121>. 16
- [11] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3606–3613, Columbus, OH, USA, 24–27 Jun 2014. URL <https://arxiv.org/abs/1311.3618>. 16
- [12] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of International Conference on Artificial Intelligence and Statistics*, pages 215–223. JMLR Workshop and Conference Proceedings, 2011. URL <https://proceedings.mlr.press/v15/coates11a/coates11a.pdf>. 16

- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186, Minneapolis, Minnesota, Jun 2019. URL <https://aclanthology.org/N19-1423.pdf>. 10
- [14] 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 *Proceedings of International Conference on Learning Representations (ICLR)*, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>. 10
- [15] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. The pascal visual objective classes challenges: A retrospective. *International Journal of Computer Vision (IJCV)*, 2015. URL <https://link.springer.com/article/10.1007/s11263-014-0733-5>. 16
- [16] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. CLIP-Adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544*, 2021. URL <https://arxiv.org/pdf/2404.02285>. 10
- [17] Gregory Griffin, Alex Holub, and Pietro Perona. Caltech-256 object category dataset. 2007. URL <https://authors.library.caltech.edu/records/5sv1j-ytw97>. 16
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, Las Vegas, NV, USA, 26 Jun – 1 Jul 2016. URL <https://arxiv.org/pdf/1512.03385>. 2, 3, 10
- [19] Kaiming He, Georgia Gkioxari, Pitor Dollár, and Ross Girshick. Mask R-CNN. In *Proceedings of IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2980–2988, Venice, Italy, 22–29 Oct 2017. URL <https://arxiv.org/pdf/1703.06870>. 2
- [20] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Introducing euroSAT: A novel dataset and deep learning benchmark for land use and land cover classification. In *Proceedings of IEEE International Geoscience and Remote Sensing Symposium*, pages 204–207, Valencia, Spain, 22–27 Jul 2018. URL <https://arxiv.org/abs/1709.00029>. 16
- [21] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8340–8349, 2021. URL <https://arxiv.org/pdf/2006.16241>. 7
- [22] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15262–15271, 2021. URL <https://arxiv.org/pdf/1907.07174>. 7
- [23] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2021. URL <https://arxiv.org/pdf/2106.09685>. 2, 9, 10, 27, 29
- [24] Chengsong Huang, Qian Liu, Bill Yuchen Lin, Tianyu Pang, Chao Du, and Min Lin. LoRAhub: Efficient cross-task generalization via dynamic lora composition. *arXiv preprint arXiv:2307.13269*, 2023. URL <https://arxiv.org/pdf/2307.13269>. 28
- [25] Yunshi Huang, Fereshteh Shakeri, Jose Dolz, Malik Boudiaf, Houda Bahig, and Ismail Ben Ayed. LP++: A surprisingly strong linear probe for few-shot clip. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. URL <https://arxiv.org/pdf/2404.02285>. 7, 10, 23
- [26] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hannaneh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. OpenCLIP, Jul 2021. URL <https://doi.org/10.5281/zenodo.5143773>. 2, 29
- [27] Gabriel Ilharco, Mitchell Wortsman, Samir Yitzhak Gadre, Shuran Song, Hannaneh Hajishirzi, Simon Kornblith, Ali Farhadi, and Ludwig Schmidt. Patching open-vocabulary models by interpolating weights. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, pages 29262–29277. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/bc6cddcd5d325e1c0f826066c1ad0215-Paper-Conference.pdf. 10

- [28] Gabriel Ilharco, Marco Túlio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. In *Proceedings of International Conference on Learning Representations (ICLR)*, Kigali, Rwanda, 1–5 May 2023. OpenReview.net. URL <https://openreview.net/pdf?id=6t0Kwf8-jrj>. 2, 3, 4, 6, 10, 18, 21, 23
- [29] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *Proceedings of International Conference on Machine Learning (ICML)*, volume 97 of *Proceedings of Machine Learning Research*, pages 3519–3529. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/kornblith19a.html>. 6
- [30] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3D object representations for fine-grained categorization. In *IEEE/CVF International Conference on Computer Vision (ICCV) Workshop on 3D Representation and Recognition*, pages 554–561, Sydney, Australia, 1–8 Dec 2013. URL <http://vision.stanford.edu/pdf/3drr13.pdf>. 16
- [31] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 4, 16
- [32] Yann LeCun, Corinna Cortez, and Christopher C. J. Burges. The mnist handwritten digit database, 1998. URL <http://yann.lecun.com/exdb/mnist/>. 4, 16
- [33] Chunyuan Li, Heerad Farkhoor, Rosanne Liu, and Jason Yosinski. Measuring the intrinsic dimension of objective landscapes. In *Proceedings of International Conference on Learning Representations (ICLR)*, Vancouver, Canada, 30 Apr–3 May 2018. URL <https://openreview.net/pdf?id=ryup8-WCW>. 2, 4
- [34] Fei-Fei Li, Robert Fergus, and Pietro Perona. One-shot learning of object categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 28(4):594–611, 2006. URL <http://vision.stanford.edu/documents/Fei-FeiFergusPerona2006.pdf>. 2, 16
- [35] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? Source hypothesis transfer for unsupervised domain adaptation. In *Proceedings of International Conference on Machine Learning (ICML)*, volume 119 of *Proceedings of Machine Learning Research*, pages 6028–6039. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/liang20a.html>. 2, 8, 10
- [36] Baijiong Lin. LoRA-Torch: PyTorch reimplementation of LoRA, 2023. URL <https://github.com/Baijiong-Lin/LoRA-Torch>. 27
- [37] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of IEEE/CVF International Conference on Computer Vision (ICCV)*, 11–17 Oct 2021. 10
- [38] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Proceedings of International Conference on Learning Representations (ICLR)*, New Orleans, LA, USA, 6–9 may 2019. OpenReview.net. URL <https://openreview.net/forum?id=Bkg6RiCqY7>. 16, 23
- [39] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. 16
- [40] Sangwoo Mo, Minkyu Kim, Kyungmin Lee, and Jinwoo Shin. S-clip: Semi-supervised vision-language learning using few specialist captions. *Advances in Neural Information Processing Systems*, 36, 2024. 10
- [41] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In *Neural Information Processing Systems (NeurIPS) Workshop on Deep Learning and Unsupervised Feature Learning*, Granada, Spain, 12–17 Dec 2011. URL http://ufldl.stanford.edu/housenumbers/nips2011_housenumbers.pdf. 16
- [42] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian conference on computer vision, graphics & image processing*, pages 722–729. IEEE, 2008. 16
- [43] Shuaicheng Niu, Jiaxiang Wu, Yifan Zhang, Zhiquan Wen, Yafo Chen, Peilin Zhao, and Minghui Tan. Towards stable test-time adaptation in dynamic wild world. In *International Conference on Learning Representations (ICLR)*, 2023. 8, 10
- [44] Guillermo Ortiz-Jiménez, Alessandro Favero, and Pscal Frossard. Task arithmetic in the tangent space: Improved editing of pre-trained models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, pages 66727–66754, New Orleans, LA, USA, 10–16 Dec 2023. Curran Associates, Inc. 2, 3, 4, 5, 6, 10, 18, 21, 23

- [45] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3498–3505. IEEE, 2012. 16
- [46] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 2019. 3, 10
- [47] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of International Conference on Machine Learning (ICML)*, volume 139, pages 8748–8763. Proceedings of Machine Learning Research (PMLR), 18–24 Jul 2021. 2, 3, 6, 7, 10, 16, 17
- [48] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020. 3
- [49] J. Rapin and O. Teytaud. Nevergrad - A gradient-free optimization platform. <https://GitHub.com/FacebookResearch/Nevergrad>, 2018. 28
- [50] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishal Shankar. Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning (ICML)*, pages 5389–5400. PMLR, 2019. 7
- [51] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 28, pages 91–99, Montréal, Canada, 7–12 Dec 2015. Curran Associates, Inc. 2
- [52] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Fei-Fei Li. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 1, 5, 7, 10, 16, 26
- [53] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *Proceedings of International Conference on Learning Representations (ICLR)*, San Diego, CA, USA, 7–9 May 2015. OpenReview.net. 2, 3, 10
- [54] K. Sohn, D. Berthelot, C.-L. L, Z. Zhang, N. Carlini, E. Cubuk, A Kurakin, H. Zhang, and C. Raffel. FixMatch: Simplifying Semi-Supervised Learning with Consistency and Confidence. *arXiv: 2001.07685*, 2020. 8, 29
- [55] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 16
- [56] Johannes Stallkamp, Marc Schlipsing, Jan Salmen, and Christian Igel. The german traffic sign recognition benchmark: A multi-class classification competition. In *Proceedings of International Joint Conference on Neural Networks (IJCNN)*, pages 1453–1460, San Jose, CA, USA, 31 Jul–5 Aug 2011. URL <https://ieeexplore.ieee.org/document/6033395>. 16
- [57] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *Proceedings of International Conference on Machine Learning (ICML)*, volume 119 of *Proceedings of Machine Learning Research*, pages 9229–9248. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/sun20b.html>. 2, 8, 10
- [58] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, pages 6000–6010. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf. 10
- [59] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2021. URL <https://openreview.net/forum?id=uXl3bZLkr3c>. 8, 10
- [60] Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P Xing. Learning robust global representations by penalizing local predictive power. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 10506–10518, 2019. 7

- [61] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology, 2010. 16
- [62] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *Proceedings of International Conference on Machine Learning (ICML)*, volume 162, pages 23965–23998, 23–29 Jul 2022. 2, 10
- [63] Jianxiong Xiao, Krista A. Ehinger, James Hays, Antonio Torralba, and Aude Oliva. Sun database: Exploring a large collection of scene categories. *International Journal of Computer Vision (IJCV)*, 119(1): 3–22, 2016. ISSN 1573-1405. URL <https://doi.org/10.1007/s11263-014-0748-y>. 16
- [64] Prateek Yadav, Derek Tam, Leshem Choshen, Colin A Raffel, and Mohit Bansal. TIES-Merging: Resolving interference when merging models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, pages 7093–7115. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/1644c9af28ab7916874f6fd6228a9bcf-Paper-Conference.pdf. 10
- [65] Enneng Yang, Zhenyi Wang, Li Shen, Shiwei Liu, Guibing Guo, Xingwei Wang, and Dacheng Tao. Adamerging: Adaptive model merging for multi-task learning. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2024. 8, 10
- [66] Frederic Z. Zhang, Dylan Campbell, and Stephen Gould. Spatially conditioned graphs for detecting human–object interactions. In *Proceedings of IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13319–13327, Oct 2021. URL <https://arxiv.org/pdf/2012.06060>. 2
- [67] Frederic Z. Zhang, Dylan Campbell, and Stephen Gould. Efficient two-stage detection of human–object interactions with a novel unary–pairwise transformer. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20104–20112, New Orleans, LA, USA, Jun 2022. URL <https://arxiv.org/pdf/2112.01838>.
- [68] Frederic Z. Zhang, Yuhui Yuan, Dylan Campbell, Zhuoyao Zhong, and Stephen Gould. Exploring predicate visual context in detecting of human–object interactions. In *Proceedings of IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10411–10421, Oct 2023. 2
- [69] Renrui Zhang, Wei Zhang, Rongyao Fang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-Adapter: Training-free adaption of CLIP for few-shot classification. In *Proceedings of European Conference on Computer Vision (ECCV)*, pages 493–510, Tel Aviv, Israel, 23–27 Oct 2022. Springer Nature Switzerland. ISBN 978-3-031-19833-5. 2, 7, 10, 23, 29
- [70] Zihao Zhao, Yuxiao Liu, Han Wu, Yonghao Li, Sheng Wang, Lin Teng, Disheng Liu, Xiang Li, Zhiming Cui, Qian Wang, et al. Clip in medical imaging: A comprehensive survey. *arXiv preprint arXiv:2312.07353*, 2023. 10
- [71] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision (IJCV)*, 130(9):2337–2348, Sep 2022. ISSN 0920-5691. URL <https://doi.org/10.1007/s11263-022-01653-1>. 10
- [72] Ziqin Zhou, Yinjie Lei, Bowen Zhang, Lingqiao Liu, and Yifan Liu. Zegclip: Towards adapting clip for zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11175–11185, 2023. 10
- [73] 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, 2021. ISSN 0018-9219. 1, 10


A Datasets and task vectors

We acquire task vectors by fine-tuning CLIP [47] on a variety of 22 image recognition datasets: (1) Stanford Cars [30], (2) DTD [11], (3) EuroSAT [20], (4) GTSRB [56], (5) MNIST [32], (6) RESISC45 [10], (7) SUN397 [63], (8) SVHN [41], (9) CIFAR10 [31], (10) CIFAR100 [31], (11) ImageNet [52], (12) STL10 [12], (13) Food101 [5], (14) Caltech101 [34], (15) Caltech256 [17], (16) FGVCAircraft [39], (17) Flowers102 [42], (18) Oxford Pets [45], (19) CUB200 [61], (20) PascalVOC [15], (21) Country211 [47], and (22) UCF101 [55]. Fine-tuning was conducted using AdamW optimiser [38], with a learning rate of 10^{-5} , batch size of 128 and weight decay of 0.1. Details of the datasets, additional dataset-specific hyper-parameters, and the accuracy after fine-tuning for an assortment of backbones are shown in Table 5. We use the same hyper-parameters for the linearised variants of the model.


Table 5: Details of the 22 image classification datasets used in experiments, the number of epochs for fine-tuning and the final accuracy for different backbones of the CLIP model.

#	Datasets	Classes	Splits			Epochs	Fine-tuned accuracy				
			<i>train</i>	<i>val</i>	<i>test</i>		RN50	RN101	ViT-B/32	ViT-B/16	ViT-L/14
(1)	Cars	196	7,330	814	8,041	35	61.92	68.41	78.26	84.14	91.67
(2)	DTD	47	3,384	376	1,880	76	73.14	72.50	78.94	81.91	84.73
(3)	EuroSAT	10	21,600	2,700	2,700	12	98.11	98.07	98.89	98.93	99.81
(4)	GTSRB	43	23,976	2,664	12,630	11	97.33	97.51	99.14	98.84	99.30
(5)	MNIST	10	55,000	5,000	10,000	5	99.62	99.45	99.65	99.69	99.77
(6)	RESISC45	45	17,010	1,890	6,300	15	93.16	93.27	95.94	96.59	97.14
(7)	SUN397	397	17,865	1,985	19,850	14	69.65	72.26	75.40	78.12	81.98
(8)	SVHN	10	68,257	5,000	26,032	4	94.30	94.58	97.38	97.70	97.97
(9)	CIFAR10	10	45,000	5,000	10,000	5	93.55	95.43	98.05	98.54	99.22
(10)	CIFAR100	100	45,000	5,000	10,000	6	77.55	80.15	89.09	89.95	93.01
(11)	ImageNet	1,000	1,276,167	5,000	50,000	10	76.01	78.19	76.41	81.33	85.52
(12)	STL10	10	4,500	500	8,000	4	90.15	91.55	98.55	99.20	99.62
(13)	Food101	101	70,750	5,000	25,250	15	85.14	87.22	88.68	92.85	95.37
(14)	Caltech101	101	6,941	694	1,736	10	87.62	85.89	94.41	95.22	94.82
(15)	Caltech256	257	22,037	2,448	6,122	8	88.29	90.54	92.60	94.58	97.17
(16)	FGVCAircraft	100	3,334	3,333	3,333	60	23.88	26.91	40.65	47.28	68.11
(17)	Flowers102	102	1,020	1,020	6,149	40	60.79	55.47	90.08	94.67	97.84
(18)	OxfordIIITPet	37	3,312	368	3,669	5	75.14	77.49	92.15	93.59	95.91
(19)	CUB200	200	5,395	599	5,794	20	58.11	59.56	73.56	77.37	86.35
(20)	PascalVOC	20	7,844	7,818	14,976	10	74.88	76.87	88.42	90.35	92.05
(21)	Country211	211	31,650	10,550	21,100	15	19.24	20.60	21.99	27.64	38.06
(22)	UCF101	101	7,639	1,898	3,783	20	81.63	83.00	85.01	89.14	92.55

To shed light on the semantic relationships amongst datasets, we extract the features of all images for each dataset, and visualise the distributions as ellipses (Figure 8). Specifically, for each dataset, the mean $\mu_t \in \mathbb{R}^d$ and covariance $\Sigma_t \in \mathbb{R}^{d \times d}$ of image features are computed. Principal component analysis (PCA) is used produce a projection matrix $P \in \mathbb{R}^{d \times 2}$ from the mean features μ_t . Subsequently, the mean and covariance with reduced dimensionality can be expressed as $P^T \mu_t$ and $P^T \Sigma_t P$, respectively.



(a) Distributions of dataset features as ellipses with $1 \times$ standard deviation



(b) Distributions of dataset features as ellipses with $3 \times$ standard deviation

Figure 8: visualisation of dataset image feature distributions as ellipses. The mean image features for all datasets are visualised as the ellipse center, with the dimensionality reduced to 2 using Principal Component Analysis (PCA). The dimensionality of covariance matrices are also reduced using the same principal components. We show visualisations with (a) $\times 1$ and (b) $\times 3$ standard deviations. Pre-trained CLIP [47] with ViT-B/32 is used to extract image features.

Table 6: Learning rates and training epochs for task negation.

	Cars	DTD	EuroSAT	GTSRB	MNIST	RESISC45	SUN397	SVHN
Learning rate	$5 \cdot 10^{-3}$	10^{-2}	$5 \cdot 10^{-3}$	$5 \cdot 10^{-3}$	$3 \cdot 10^{-3}$	$2 \cdot 10^{-3}$	$3 \cdot 10^{-3}$	$5 \cdot 10^{-3}$
Epochs	20	20	3	5	7	10	10	2

Table 7: Accuracy on target and control tasks of task negation for each of the eight datasets. Highest performance in each section is highlighted in bold. The method *search* corresponds to model $f(x; \theta_0 + \alpha \tau_t)$, where α is determined via a hyper-parameter search. Our method *aniso.* corresponds to model $f(x; \theta_0 + \Lambda_t \tau_t)$, where Λ_t is a learnable scaling matrix.

	Cars		DTD		EuroSAT		GTSRB		MNIST		RESISC45		SUN397		SVHN		Average		
	Tgt.	Ctr.	Tgt.	Ctr.	Tgt.	Ctr.	Tgt.	Ctr.	Tgt.	Ctr.	Tgt.	Ctr.	Tgt.	Ctr.	Tgt.	Ctr.	Tgt.	Ctr.	
VT-B/32	Zero-shot	59.73	63.35	43.99	63.35	45.19	63.35	32.56	63.35	48.25	63.35	60.65	63.35	63.18	63.35	31.61	63.35	48.18	63.35
	Std. (search)	35.06	60.72	29.41	60.66	11.89	60.68	7.16	60.39	12.67	60.84	31.27	61.26	51.25	60.48	7.03	60.61	23.22	60.71
	Std. (aniso.)	28.95	60.52	25.21	60.48	10.44	61.62	5.51	60.67	10.76	62.9	20.95	60.72	46.29	60.82	7.28	62.72	18.76	61.21
	Lin. (search)	27.06	60.71	15.27	60.42	0.26	60.63	1.03	61.23	0.06	62.52	6.83	60.68	41.3	59.93	0.54	59.77	11.54	60.74
	Lin. (aniso.)	23.96	60.57	15.05	60.55	0.44	61.86	1.1	61.16	0.06	61.71	4.48	60.26	42.71	60.78	0.76	61.2	11.06	61.02
VT-B/16	Zero-shot	64.61	68.33	45.11	68.33	55.78	68.33	43.34	68.33	51.79	68.33	65.76	68.33	65.5	68.33	51.98	68.33	55.48	68.33
	Std. (search)	24.19	64.41	21.65	63.75	12.41	64.76	7.16	63.95	9.85	65.52	25.48	64.23	47.86	64.16	6.47	66.47	19.38	64.66
	Std. (aniso.)	16.63	63.95	20.69	64.6	15.93	67.65	8.21	66.37	9.51	68.29	21.29	65.17	45.43	65.02	6.84	67.93	17.34	65.84
	Lin. (search)	23.91	64.74	11.01	64.89	0.15	64.12	3.06	66.99	0.21	67.41	5.48	64.52	42.39	65.11	0.88	66.54	10.88	65.54
	Lin. (aniso.)	19.85	64.05	9.68	64.57	0.33	65.91	0.97	65.51	0.01	66.87	7.27	65.39	41.18	65.36	1.97	67.01	10.16	65.58
VT-L/14	Zero-shot	77.75	75.54	55.32	75.54	61.33	75.54	50.55	75.54	76.36	75.54	71.05	75.54	68.28	75.54	58.45	75.54	64.89	75.54
	Std. (search)	24.44	71.34	26.91	71.83	8.63	71.46	6.24	71.78	11.15	72.43	17.98	72.07	51.11	71.99	6.72	73.53	19.15	72.05
	Std. (aniso.)	14.49	71.07	23.94	72.2	12.15	74.81	3.95	71.66	7.29	74.69	25.11	74.29	47.93	72.8	7.16	74.69	17.75	73.28
	Lin. (search)	18.57	71.09	13.03	71.92	0.33	73.15	5.57	74.41	5.31	74.32	3.11	72.03	45.79	72.2	10.54	74.51	12.78	72.95
	Lin. (aniso.)	16.9	71.67	10.48	71.78	1.19	74.49	4.13	74.39	7.6	74.98	6.46	73.38	44.96	72.4	13.23	75.18	12.61	73.14

B Task negation

The evaluation of task negation is conducted on eight classification datasets (1–8 in Table 5), following previous practice [28, 44]. In particular, we learn anisotropic scaling using the validation set of each dataset. We also adjust the learning rates and training epochs on the same validation set. The details are shown in Table 6. We report detailed task negation results for each dataset in Table 7. In addition, for more evidence that weight matrices learn large negative coefficients, we show a detailed visualisation of the learned coefficients in Figure 9 and distribution of the coefficients in Figure 10.

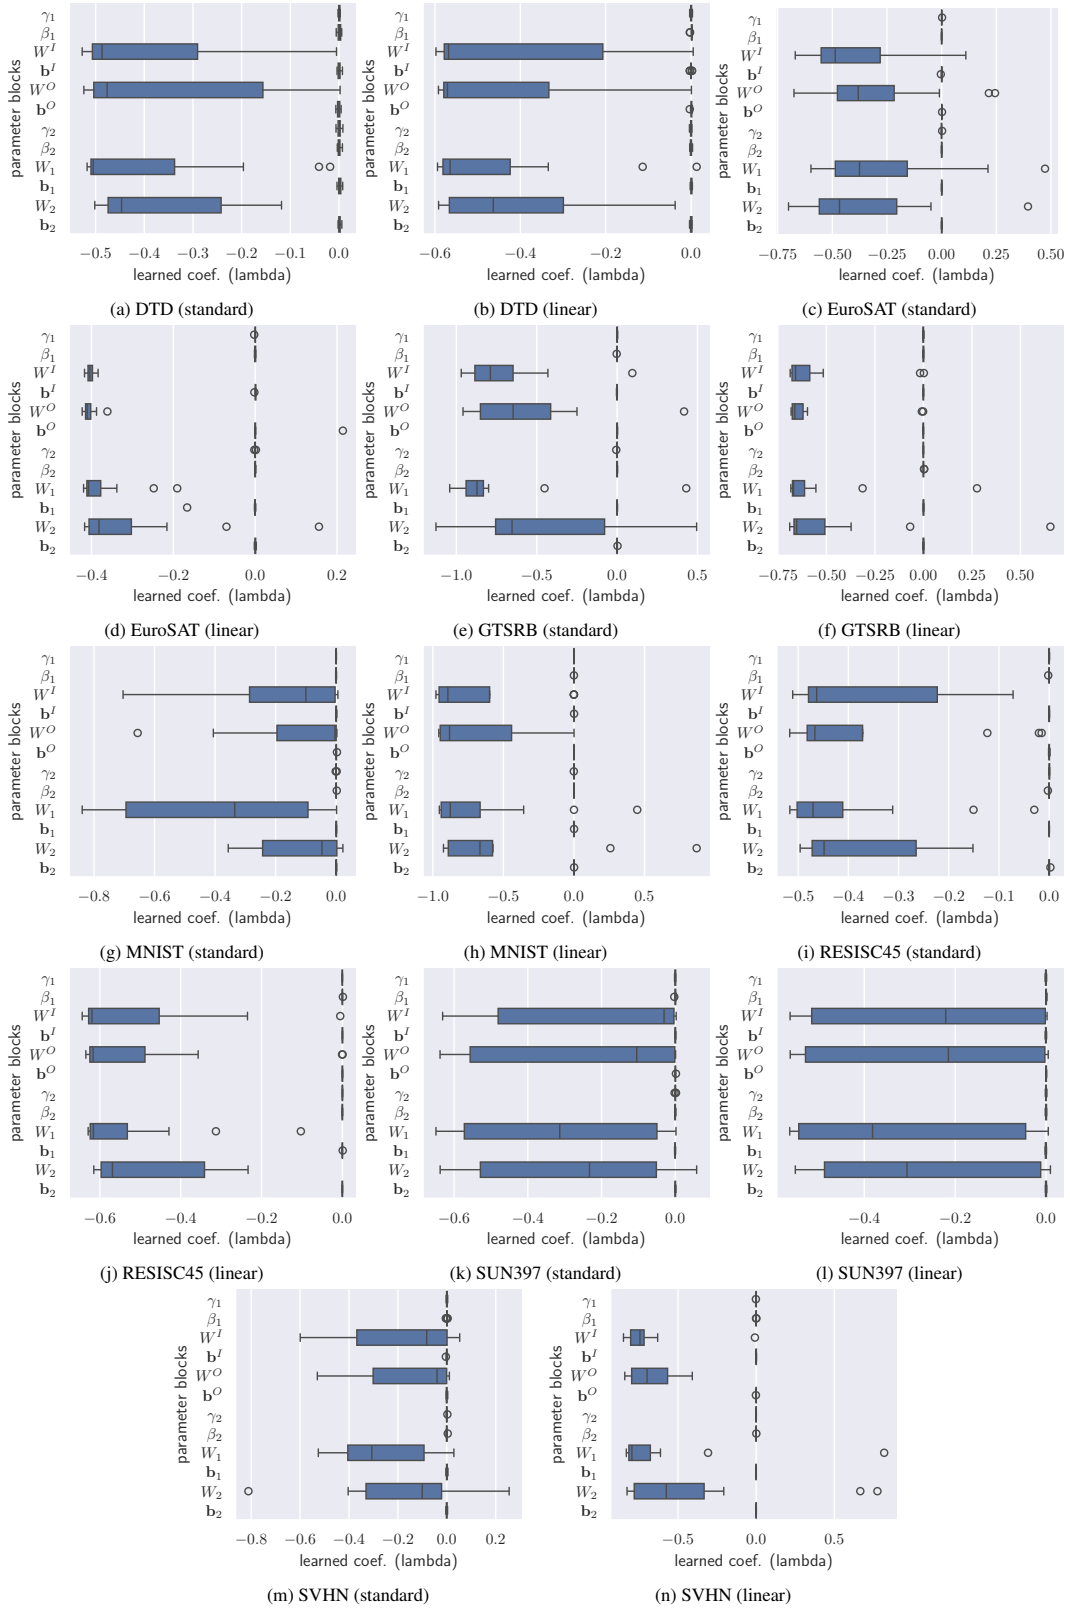


Figure 10: Additional box-and-whisker plots for the learned coefficients in task negation, beside previous visualisation on Cars (Figure 3a), including results on (a, b) DTD, (c, d) EuroSAT, (e, f) GTSRB, (g, h) MNIST, (i, j) RESISC45, (k, l) SUN397 and (m, n) SVHN.

Table 8: Detailed performance for task addition across all eight datasets. We show additional performance with learned isotropic scaling, which have comparable accuracy to the simple hyper-parameter search used in previous methods [28, 44]. Highest performance in each section is highlighted in bold. The method *search* corresponds to model $f(x; \theta_0 + \alpha \sum \tau_i)$, where α is determined via a hyper-parameter search. Methods *iso.* and *aniso.* use learned coefficients and correspond to models $f(x; \theta_0 + \sum \alpha_i \tau_i)$ and $f(x; \theta_0 + \sum \Lambda_i \tau_i)$, respectively.

	Cars		DTD		EuroSAT		GTSRB		MNIST		RESISC45		SUN397		SVHN		Average		
	Abs.	Rel.	Abs.	Rel.	Abs.	Rel.	Abs.	Rel.	Abs.	Rel.	Abs.	Rel.	Abs.	Rel.	Abs.	Rel.	Abs.	Rel.	
VT-B/32	Zero-shot	59.73	-	43.99	-	45.19	-	32.56	-	48.25	-	60.65	-	63.18	-	31.61	-	48.14	-
	Std. (search)	58.97	75.35	52.29	66.24	80.04	80.94	66.74	67.31	95.96	96.3	70.54	73.53	60.74	80.55	75.66	77.69	70.12	77.24
	Std. (iso.)	56.88	72.68	51.97	65.84	87.96	88.95	73.14	73.77	82.23	82.52	61.16	63.75	62.88	83.4	87.98	90.34	70.53	77.66
	Std. (aniso.)	71.56	91.43	72.66	92.05	93.93	94.98	89.8	90.58	96.13	96.47	87.71	91.43	68.9	91.37	91.94	94.41	84.98	93.79
	Lin. (search)	67.14	87.97	58.56	73.2	95.67	97.62	67.58	72.13	94.61	95.25	81.25	85.92	63.99	83.78	68.36	85.5	74.67	85.17
	Lin. (iso.)	63.4	83.07	65.59	81.98	93.78	95.69	82.42	87.97	95.24	95.88	84.78	89.64	58.79	76.96	67.92	84.95	76.49	87.02
Lin. (aniso.)	72.33	94.77	73.03	81.29	95.26	97.2	88.41	94.36	97.27	97.93	88.79	93.89	68.53	89.71	72.53	90.71	83.42	95.42	
VT-B/16	Zero-shot	64.61	-	45.11	-	55.78	-	43.34	-	51.79	-	65.76	-	65.5	-	51.98	-	55.48	-
	Std. (search)	68.47	81.38	52.82	64.48	75.0	75.81	71.03	71.86	96.97	97.27	76.35	79.05	66.57	85.23	81.82	83.75	73.63	79.85
	Std. (iso.)	65.55	77.9	50.05	61.1	81.96	82.85	74.06	74.93	94.96	95.26	80.94	83.8	60.48	77.43	91.65	93.81	74.96	80.88
	Std. (aniso.)	71.79	85.32	67.07	81.88	91.85	92.85	91.9	92.98	97.02	97.32	84.3	87.28	66.6	85.26	92.66	94.85	86.08	93.44
	Lin. (search)	72.7	85.26	60.96	73.65	95.0	96.79	70.39	74.78	95.37	96.17	83.78	87.63	71.47	90.67	70.44	84.71	77.51	86.21
	Lin. (iso.)	72.6	85.14	70.48	85.15	92.19	93.92	82.43	87.58	93.86	94.65	88.14	92.2	61.26	77.72	72.95	87.73	79.24	88.01
Lin. (aniso.)	81.78	95.9	77.13	93.19	96.33	98.15	88.65	94.18	98.6	99.43	90.49	94.65	72.95	92.55	77.11	92.72	85.38	95.10	
VT-L/14	Zero-shot	77.75	-	55.32	-	61.33	-	50.55	-	76.36	-	71.05	-	68.28	-	58.45	-	64.89	-
	Std. (search)	81.69	89.12	64.63	76.27	88.67	88.83	93.88	94.54	98.75	98.98	82.68	85.11	71.3	86.98	81.86	83.56	82.93	87.92
	Std. (iso.)	85.72	93.52	71.38	84.24	83.74	83.9	91.74	92.39	96.5	96.72	91.56	94.25	60.33	73.59	94.94	96.91	84.49	89.44
	Std. (aniso.)	89.58	97.72	80.85	95.42	98.0	98.18	96.75	97.43	98.48	98.71	93.03	95.77	77.96	95.1	96.24	98.23	91.36	97.07
	Lin. (search)	85.13	94.86	74.41	89.28	95.89	97.33	77.82	81.12	98.11	98.75	89.87	93.14	74.29	90.0	82.45	90.38	84.75	91.86
	Lin. (iso.)	84.18	93.81	74.41	89.28	94.89	96.32	82.62	86.13	97.16	97.8	91.33	94.65	73.87	89.48	83.01	90.99	85.18	92.31
Lin. (aniso.)	87.38	97.37	78.51	94.19	95.7	97.14	91.73	95.62	98.39	99.03	93.56	96.96	77.25	93.58	86.7	95.04	88.65	96.12	

C Task addition

Task addition is also evaluated on datasets 1–8 shown in Table 5. The hyper-parameters are identical to fine-tuning, except the learning rate is modified to 10^{-3} . We show detailed performance on each dataset in Table 8, where we compare our method against hyper-parameter search used in previous works [28, 44], and another variant with learned isotropic scaling. We also visualise the learned coefficients with L_1 regularisation in Figure 12. It can be easily observed that weight matrices, particularly those in the deeper layers, have significantly higher learned coefficients, which conforms to our observations in Figures 3b and 3c.

C.1 Comparison against full-parameter optimisation

Since our method involves learning the coefficients, unlike previous methods [28, 44] that only require a hyper-parameter search, we also compare against the direct fine-tuning approach. We fine-tune the pre-trained model on the union of eight datasets, assuming only the validation sets are available. The results are shown in Figure 11. Unsurprisingly, task vector compositions, whether the coefficients are searched or learned, are less susceptible to the lack of data, as the accuracy only starts to drop with less than 35% of the data. The performance of full-parameter fine-tuning, however, drops substantially as the amount of data available decreases.





Figure 11: Task addition accuracy averaged across eight datasets (1–8) versus different percentage of validation data used. Standard task vectors are used.

C.2 Disentanglement error

In addition, we provide more technical details and intuitions on the pairwise disentanglement error [44], which was visualised in Figure 4. Specifically, we make a few changes to the formulation proposed by Ortiz-Jiménez et al. [44], and evaluate the disentanglement error only with the optimal



(a) Learned coefficients on standard task vectors

(b) Learned coefficients on linear task vectors

Figure 12: visualisation of the learned coefficients for (a) standard and (b) linear task vectors in task addition. Note that L_1 regularisation has been applied to the coefficients during training for better clarity. CLIP with ViT-B/32 backbone is used to produce the results.

coefficients. Given two datasets $\mathcal{D}_1, \mathcal{D}_2$ and the respective task vectors τ_1, τ_2 , we overload the definition of function f to denote the mapping from data space \mathcal{X} to the label space \mathcal{Y} , and define the disentanglement error as

$$\xi(\tau_1, \tau_2) = \mathbf{E}_{\mathbf{x} \in \mathcal{D}_1} [\delta(f(\mathbf{x}; \theta_0 + \Lambda_1^* \tau_1), f(\mathbf{x}; \theta_0 + \Lambda_1^* \tau_1 + \Lambda_2^* \tau_2))], \quad (10)$$

$$\xi(\tau_2, \tau_1) = \mathbf{E}_{\mathbf{x} \in \mathcal{D}_2} [\delta(f(\mathbf{x}; \theta_0 + \Lambda_2^* \tau_2), f(\mathbf{x}; \theta_0 + \Lambda_1^* \tau_1 + \Lambda_2^* \tau_2))], \quad (11)$$

where Λ_1^*, Λ_2^* are the learned coefficients in task addition, and δ is defined as

$$\delta(x_1, x_2) = \begin{cases} 0 & x_1 = x_2, \\ 1 & x_1 \neq x_2. \end{cases} \quad (12)$$

The error metric $\xi(\tau_1, \tau_2)$ measures the percentage of data in dataset \mathcal{D}_1 , such that when a second task vector τ_2 is added to the model, the predicted labels differ from when only using task vector τ_1 . As task vector τ_1 is acquired from dataset \mathcal{D}_1 , a low disentanglement error indicates that most predictions made by τ_1 —highly likely to be correct—will be retained, thus resulting in higher performance in task addition.

Table 9: Average accuracy for few-shot recognition over 22 datasets. We report accuracy averaged over 3 random n-shot sample selections, with $1 \times$ standard error. Results are produced using CLIP with ViT-B/32 backbone. For our method, we show results with both standard [28] and linearised [44] task vectors. The best method for each choice of $k \in \{1, 2, 4, 8, 16\}$ is highlighted in bold.

Shots (k)	Tip-Adapter	LP++	aTLAS		aTLAS w/ LP++		aTLAS w/ Tip-Adapter	
			Std.	Lin.	Std.	Lin.	Std.	Lin.
1	64.3 \pm 0.2	64.5 \pm 0.1	66.2 \pm 0.2	64.6 \pm 0.1	68.9 \pm 0.2	69.3 \pm 0.1	68.7 \pm 0.4	66.7 \pm 0.3
2	67.0 \pm 0.1	67.3 \pm 0.1	67.6 \pm 0.1	65.4 \pm 0.1	71.5 \pm 0.1	67.2 \pm 0.2	71.9 \pm 0.2	68.9 \pm 0.1
4	69.7 \pm 0.1	69.9 \pm 0.1	70.0 \pm 0.0	66.6 \pm 0.1	73.7 \pm 0.1	70.8 \pm 0.1	74.3 \pm 0.1	71.6 \pm 0.2
8	71.8 \pm 0.1	72.3 \pm 0.1	71.5 \pm 0.0	68.2 \pm 0.1	75.8 \pm 0.1	73.5 \pm 0.1	76.5 \pm 0.1	74.2 \pm 0.1
16	73.7 \pm 0.1	74.1 \pm 0.1	72.9 \pm 0.1	69.8 \pm 0.1	77.8 \pm 0.0	76.2 \pm 0.1	78.0 \pm 0.1	76.7 \pm 0.0

D Few-shot learning

D.1 Baselines: Tip-Adapter and LP++

Two variants of Tip-Adapter [69] were proposed for few-shot recognition where the weights of the adaptor are either fixed based on features of the few-shot examples or further fine-tuned. We only study the fine-tuned variant due to its higher performance. Tip-Adapter has two hyper-parameters, which in the original paper are optimised through hyper-parameter search on a separate validation set. This practice may not align with the principles of few-shot learning, where access to extensive validation data is typically limited. In addition, Huang et al. [25] note that the performance of Tip-Adapter is very sensitive to these hyper-parameters. We thus opt to learn these two hyper-parameters together with the feature adaptor through gradient descent. The learning rates for the feature adaptor and the hyper-parameters are set to 10^{-3} and 10^{-1} , respectively.

For both Tip-Adapter and LP++ [25], we conduct experiments using the publicly available codebase ⁵. We train both LP++ and Tip-Adapter for 300 epochs on frozen zero-shot features. We apply a cosine annealing decay for Tip-Adapter and maintain fixed learning rates for LP++ as per the official implementation.

D.2 linearised task vectors

We report the average few-shot accuracy over the 22 datasets in Table 9, which corresponds to results in Figure 5a. In particular, we show results with linearised task vectors, as proposed by Ortiz-Jiménez et al. [44]. As highlighted in Section 4, learned anisotropic scaling allows standard task vectors to achieve stronger performance than the linear variants in task addition. For few-shot recognition, we again observe that standard task vectors result in superior performance in most cases. We, however, note the exception that linear task vectors when combined with LP++ achieve higher performance in the 1-shot setting. Nevertheless, the margin over standard task vectors is not very significant, and aTLAS using standard task vectors when integrated with Tip-Adapter is generally a stronger few-shot model.

D.3 Integrating state-of-the-art methods into aTLAS

We use the AdamW [38] optimiser with a learning rate of 10^{-1} and a weight decay of 10^{-1} . Our method by itself is trained for 10 epochs with ViT backbones and 30 epochs with ResNet backbones.

We show that state-of-the-art few-shot methods can be seamlessly integrated into our method, since both Tip-Adapter and LP++ focus on the classifier, while aTLAS improves the feature representations. We experiment with two strategies to combine aTLAS with previous methods, where we either (1) train our method first and use the frozen representations to train a previous method, or (2) train parameters in both methods jointly. Results in Table 10 shows that the joint training strategy results in higher performance, particularly in low-shot settings. We therefore adopt the joint training strategy when combining our method with Tip-Adapter. During training, we adopt different learning rates for different parameter groups, that is, 10^{-1} for learnable coefficients in aTLAS and the hyper-parameters in Tip-Adapter, and 10^{-3} for the adaptor. The joint training takes 20 epochs for ViT backbones and 60 epochs on ResNet backbones, twice the number of epochs when training aTLAS alone.

⁵github.com/fereshreshakeri/fewshot-clip-strong-baseline

Table 10: Comparison of few-shot recognition accuracy between training our method and Tip-Adapter sequentially and jointly over different shots (k). ViT-B/32 is used as the backbone. Results are averaged across three random seeds. Highest performance in each section is highlighted in bold.

k	Strategy	Cars	DTD	EuroSAT	GTSRB	MNIST	RESISC45	SUN397	SVHN	CIFAR10	CIFAR100	ImageNet	STL10	Food101	Caltech256	FGVCAircraft	Flowers102	OxfordIIITPet	CUB200	PascalVOC	Country211	Caltech101	UCF101	Average
1	Seq.	59.7	43.7	64.5	44.6	75.7	63.6	64.0	38.3	89.8	69.2	63.7	97.2	84.0	84.9	19.9	70.2	88.8	53.0	76.2	17.2	89.5	66.0	64.7
	Joint	61.4	49.3	72.5	57.5	84.8	70.8	67.5	57.0	91.0	75.5	64.2	97.2	84.0	84.0	22.6	79.0	88.7	54.4	76.2	17.3	89.3	69.0	68.8
2	Seq.	63.6	48.8	72.6	49.0	85.4	72.8	67.8	42.4	89.8	69.9	64.8	97.2	84.0	86.1	22.8	78.4	88.8	57.5	76.2	17.2	90.3	69.8	68.0
	Joint	65.1	55.9	83.4	67.6	86.7	77.3	69.2	56.3	93.4	75.9	64.5	97.5	84.0	88.0	26.2	85.0	89.6	58.8	76.2	18.1	90.4	72.9	71.9
4	Seq.	68.7	55.5	76.8	58.3	88.4	75.2	70.3	45.5	90.8	71.9	65.8	98.0	84.0	86.6	28.9	86.4	89.0	63.6	76.2	17.2	92.2	73.2	71.0
	Joint	69.0	63.7	79.6	73.1	90.6	78.9	70.8	67.6	93.9	78.3	64.7	97.5	84.0	87.5	29.3	91.6	89.7	62.8	77.3	18.1	90.4	77.5	74.4
8	Seq.	73.7	64.7	82.3	69.5	91.0	79.8	72.3	43.0	91.9	73.4	66.9	97.7	84.0	88.9	34.7	91.2	90.8	69.3	76.2	18.5	91.9	78.8	74.1
	Joint	74.2	69.4	89.5	78.4	94.0	81.1	72.1	60.1	95.0	78.5	65.6	97.8	84.3	88.6	36.6	93.7	90.0	69.0	77.6	19.0	91.7	78.9	76.6
16	Seq.	77.6	69.0	89.3	79.3	92.8	83.7	74.1	64.6	94.0	79.9	66.8	97.5	84.6	89.2	36.1	94.2	91.4	72.8	80.3	20.1	93.6	80.0	77.8
	Joint	79.1	68.8	91.4	82.7	93.9	82.8	73.7	63.7	95.1	80.0	67.0	98.1	84.3	88.6	39.1	94.5	90.1	74.0	77.0	20.4	91.9	80.4	78.0

Table 11: Detailed few-shot accuracy for each dataset over different shots (k) using ViT-B/32 backbone. We report results averaged over 3 random seeds. In the case where the results are worse than the zero-shot accuracy, we report zero-shot accuracy. Highest performance and those within a range of 0.1 in each section are highlighted in bold. Tip-Adapter is abbreviated as Tip.

k	Method	Cars	DTD	EuroSAT	GTSRB	MNIST	RESISC45	SUN397	SVHN	CIFAR10	CIFAR100	ImageNet	STL10	Food101	Caltech256	FGVCAircraft	Flowers102	OxfordIIITPet	CUB200	PascalVOC	Country211	Caltech101	UCF101	Average
0	CLIP	59.7	43.7	42.9	32.7	48.3	60.6	63.2	31.3	89.8	64.2	63.4	97.2	84.0	82.0	19.6	66.6	87.5	53.0	76.2	17.2	84.0	61.6	60.4
1	Tip	62.6	52.3	55.2	37.5	61.3	67.2	66.5	34.7	89.8	66.4	63.9	97.7	84.0	84.8	22.0	80.6	87.6	55.1	76.7	17.3	87.5	68.2	64.5
	LP++	61.5	51.3	60.0	39.5	50.5	68.8	65.8	31.8	89.8	66.3	63.9	97.2	84.1	84.7	23.7	81.9	87.5	55.1	76.2	17.2	88.3	69.7	64.3
	aTLAS	59.7	43.7	74.7	52.4	79.5	64.5	63.2	59.0	89.8	70.2	63.4	97.2	84.0	83.4	19.6	66.6	87.5	53.0	76.2	17.2	89.1	62.5	66.2
	aTLAS w/ LP++	62.2	50.2	72.5	52.8	84.0	69.9	67.2	64.6	92.1	75.6	64.4	97.2	84.0	85.4	24.1	81.6	87.5	56.2	76.2	17.2	89.5	69.4	69.2
	aTLAS w/ Tip	61.4	49.3	72.5	57.5	84.8	70.8	67.5	57.0	91.0	75.5	64.2	97.2	84.0	84.0	22.6	79.0	88.7	54.4	76.2	17.3	89.3	69.0	68.8
2	Tip	64.1	57.4	70.8	43.7	66.0	73.1	68.3	31.8	90.0	66.6	64.5	97.9	84.3	85.6	24.0	87.3	88.0	57.4	76.2	17.5	88.4	72.5	67.1
	LP++	64.2	57.3	69.8	42.3	69.7	74.7	67.4	31.3	89.8	67.6	64.5	97.4	84.0	86.3	25.6	86.8	88.7	59.7	76.2	17.4	90.4	72.3	67.4
	aTLAS	59.7	43.9	79.5	53.8	86.0	68.0	64.0	60.9	91.1	72.5	63.9	97.2	84.0	84.6	21.2	67.4	88.4	53.0	76.2	17.2	89.8	65.0	67.6
	aTLAS w/ LP++	65.8	58.2	80.0	58.4	87.4	74.4	68.5	55.1	93.1	76.8	64.9	97.4	84.0	87.0	26.9	87.0	88.5	60.5	76.2	17.5	89.7	72.7	71.4
	aTLAS w/ Tip	65.1	55.9	83.4	67.6	86.7	77.3	69.2	56.3	93.4	75.9	64.5	97.5	84.0	88.0	26.2	85.0	89.6	58.8	76.2	18.1	90.4	72.9	71.9
4	Tip	65.8	62.3	77.3	48.7	77.9	77.1	70.2	32.8	90.4	68.2	65.0	97.2	84.7	87.1	28.8	91.5	88.4	62.3	76.2	18.3	90.7	73.9	69.8
	LP++	67.5	61.5	74.2	54.0	72.5	79.1	70.5	34.3	90.8	68.1	65.4	98.0	84.3	88.7	27.4	90.1	87.7	62.5	77.3	18.8	91.8	75.5	70.0
	aTLAS	60.9	48.6	84.3	66.3	89.6	72.0	65.1	71.8	91.9	75.1	64.7	97.2	84.0	85.2	22.9	69.9	88.2	53.5	76.2	17.2	90.7	65.2	70.0
	aTLAS w/ LP++	69.9	62.9	85.3	73.6	89.2	76.0	70.8	58.5	93.6	78.0	65.7	97.2	84.4	88.2	28.9	89.2	90.6	64.3	77.0	17.7	91.2	75.8	74.0
	aTLAS w/ Tip	69.0	63.7	79.6	73.1	90.6	78.9	70.8	67.6	93.9	78.3	64.7	97.5	84.0	87.5	29.3	91.6	89.7	62.8	77.3	18.1	90.4	77.5	74.4
8	Tip	71.1	65.6	78.3	58.1	84.9	80.9	71.7	31.3	91.0	68.4	65.0	97.6	85.0	88.0	31.2	93.1	90.4	66.7	76.5	19.4	91.5	76.5	71.9
	LP++	72.2	65.2	79.4	61.2	82.5	81.8	72.2	31.3	91.0	69.6	66.9	97.9	84.7	89.1	31.2	92.2	89.9	66.7	78.7	19.3	92.9	76.8	72.4
	aTLAS	61.6	52.1	90.8	67.2	90.2	74.6	65.9	72.2	93.0	77.3	64.8	97.2	84.0	85.8	24.8	73.0	89.9	55.4	77.0	17.4	91.2	67.8	71.5
	aTLAS w/ LP++	73.5	65.2	86.6	73.2	92.8	80.8	72.8	64.2	94.1	79.6	67.3	97.2	84.7	88.0	32.1	91.4	90.2	66.7	76.2	19.0	92.4	77.3	75.7
	aTLAS w/ Tip	74.2	69.4	89.5	78.4	94.0	81.1	72.1	60.1	95.0	78.5	65.6	97.8	84.3	88.6	36.6	93.7	90.0	69.0	77.6	19.0	91.7	78.9	76.6
16	Tip	73.5	67.5	85.8	66.6	89.1	82.3	73.1	31.3	91.3	69.4	65.9	97.9	84.7	88.8	36.2	94.5	89.6	70.3	76.2	20.3	92.3	79.5	73.9
	LP++	75.2	69.7	86.4	64.8	87.4	83.3	74.4	31.3	91.7	70.8	68.2	98.0	85.5	89.1	35.7	92.2	90.6	69.3	76.9	20.4	93.1	79.5	74.2
	aTLAS	62.9	55.5	92.6	71.3	92.8	77.0	66.5	78.5	93.9	77.8	65.4	97.3	84.3	86.6	24.9	74.1	90.6	56.5	76.9	17.8	92.8	68.4	72.9
	aTLAS w/ LP++	76.9	68.9	89.6	79.3	94.4	84.1	73.5	70.7	94.7	80.6	68.7	97.9	85.1	90.3	34.2	92.0	90.8	69.5	78.4	19.8	93.4	79.9	77.9
	aTLAS w/ Tip	79.1	68.8	91.4	82.7	93.9	82.8	73.7	63.7	95.1	80.0	67.0	98.1	84.3	88.6	39.1	94.5	90.1	74.0	77.0	20.4	91.9	80.4	78.0

On the other hand, The joint training strategy with LP++ is non-trivial, due to LP++’s super-convergence strategy being designed around frozen feature representations, which would have been updated every iteration by aTLAS. We thus use the sequential strategy to combine aTLAS and LP++. We include detailed results for each dataset with ViT-B/32 in Table 11 and additional results with different backbones in Table 12, where we show our method scales well across different datasets and backbones.

Table 12: Detailed few-shot accuracy for each dataset across RN50, RN101, ViT-B/16, and ViT-L/14 backbones. We report results for the same random seed. In the case where the results are worse than the zero-shot accuracy, we report zero-shot accuracy. Highest performance and those within a range of 0.1 in each section are highlighted in bold. Tip-Adapter is abbreviated as Tip.

k	Method	Cars	DTD	EuroSAT	GTSRB	MNIST	RESISC45	SUN397	SVHN	CFP100	CIFAR10	CIFAR100	ImageNet	STL10	Food101	Caltech256	FCVCAircraft	Flowers102	OxfordITPet	CUB200	PascalVOC	Country211	Caltech101	UCF101	Average
RN50	0	CLIP	54.3	41.2	41.5	35.1	58.1	53.1	60.1	28.9	71.5	40.3	59.9	94.2	80.6	77.3	17.0	66.2	85.8	46.5	66.2	15.4	77.6	58.3	55.9
	1	Tip	56.3	50.2	58.4	37.6	66.2	65.6	64.5	29.7	77.0	48.0	60.8	95.0	80.7	79.9	19.4	79.8	86.4	51.3	66.2	16.2	83.3	63.4	60.2
		LP++	57.2	44.8	58.2	35.9	57.4	59.4	63.0	30.9	72.2	45.0	60.7	94.2	80.6	80.9	21.2	80.9	86.0	53.8	68.0	16.2	85.4	64.7	59.9
		aTLAS	54.3	41.2	62.7	44.5	75.6	56.0	61.0	50.7	78.4	54.0	60.7	95.2	80.6	79.8	19.5	66.2	85.8	50.5	66.2	15.4	86.8	60.0	61.1
		aTLAS w/ Tip	56.5	44.2	44.9	37.5	78.3	54.9	62.0	29.0	71.6	40.8	63.5	97.2	84.0	85.2	23.5	75.7	87.4	53.5	76.4	17.4	88.7	68.4	60.9
	2	Tip	58.3	56.4	71.6	37.5	57.4	65.6	64.5	29.7	77.0	48.0	60.8	95.0	80.7	81.2	23.2	84.3	87.2	54.5	67.3	16.2	86.3	66.7	62.1
		LP++	60.0	52.7	66.1	41.7	61.5	64.6	64.8	29.7	72.1	46.9	61.5	94.4	80.6	82.5	22.6	84.2	86.6	56.0	67.5	16.2	88.1	66.1	62.1
		aTLAS	55.2	42.0	76.6	52.3	73.0	59.8	60.6	56.8	79.7	57.0	60.7	94.3	80.6	80.4	20.6	67.0	85.8	51.0	66.2	15.8	86.9	59.6	62.8
		aTLAS w/ Tip	58.9	53.1	52.5	63.3	88.8	62.6	64.3	28.9	72.1	41.0	63.7	97.1	84.0	86.5	24.9	85.9	87.5	57.2	66.5	17.5	89.0	73.0	64.9
	4	Tip	63.3	61.2	74.4	41.4	68.8	71.0	66.3	29.7	76.9	49.1	61.3	94.4	81.4	82.5	24.2	90.2	86.9	58.0	66.2	16.5	86.3	71.5	64.6
		LP++	63.2	59.5	76.6	43.4	70.4	69.8	66.8	29.7	74.9	48.0	62.4	94.2	81.1	83.7	24.6	89.2	86.9	60.1	66.7	16.2	87.0	70.9	64.8
		aTLAS	57.6	43.0	79.7	55.3	82.2	62.0	61.4	53.4	81.6	58.9	61.4	94.2	80.6	80.9	22.6	69.7	86.5	52.3	66.8	15.4	87.1	61.3	64.3
		aTLAS w/ Tip	62.9	58.9	64.2	72.0	90.1	70.0	65.1	29.8	71.7	42.1	64.1	97.4	84.2	87.0	29.2	91.2	88.7	63.1	76.2	17.5	90.1	76.0	67.8
	8	Tip	66.9	63.5	80.8	54.1	77.1	71.8	68.0	29.7	76.7	48.8	61.1	95.1	81.6	83.9	30.4	93.3	85.8	62.7	69.4	17.4	88.6	74.0	67.3
		LP++	68.2	64.0	78.7	50.6	76.7	74.7	70.0	29.7	75.5	50.3	63.8	95.8	81.6	84.3	28.8	92.8	88.0	64.2	68.5	17.0	88.7	74.1	67.5
		aTLAS	58.4	48.5	80.0	57.1	85.4	66.9	62.3	60.8	84.2	61.2	61.9	95.8	81.1	82.2	23.6	71.4	87.8	53.0	69.1	15.6	89.6	63.5	66.3
aTLAS w/ Tip		69.9	61.1	74.4	81.1	91.5	75.9	65.2	28.9	73.6	45.5	64.7	97.1	84.0	87.8	34.0	93.3	89.6	68.0	76.2	18.5	90.7	78.9	70.4	
16	Tip	71.9	67.1	83.2	63.3	86.0	75.1	69.8	32.4	79.3	51.2	62.8	95.0	81.5	85.1	33.1	94.3	87.6	66.9	68.3	18.5	91.9	75.2	70.0	
	LP++	72.9	68.0	84.1	57.2	78.0	75.4	71.3	29.7	76.6	51.8	64.7	96.3	82.1	85.5	31.6	92.5	89.3	67.8	72.3	17.7	92.9	77.5	69.8	
	aTLAS	59.4	51.1	88.8	59.2	87.8	68.5	61.9	67.7	84.9	61.9	62.1	95.6	81.5	83.1	24.4	71.4	88.7	53.4	68.7	16.1	89.9	63.8	67.7	
	aTLAS w/ Tip	72.3	67.2	81.1	83.0	94.3	80.3	68.8	28.9	75.7	47.4	65.7	97.1	84.0	89.0	40.2	94.9	89.8	71.8	77.3	20.2	95.5	81.6	74.4	
RN101	0	CLIP	61.0	43.6	30.7	37.7	51.4	58.5	59.5	31.5	80.8	47.7	62.3	96.8	83.6	80.9	18.5	65.3	86.9	49.6	64.5	16.9	81.8	58.5	57.6
	1	Tip	66.0	50.6	60.1	41.2	54.7	63.5	63.2	39.5	80.9	52.1	63.3	96.7	84.0	84.0	22.5	77.9	87.6	53.2	69.7	17.4	87.2	67.0	62.8
		LP++	64.8	55.4	65.4	43.4	56.4	66.2	62.5	30.8	82.7	54.0	63.0	96.8	83.6	83.0	22.6	79.2	86.8	51.7	65.1	17.4	86.6	68.8	63.0
		aTLAS	62.5	43.6	73.9	55.3	75.9	60.4	61.1	56.3	82.6	61.4	62.9	96.8	83.6	82.7	20.1	65.3	86.9	50.7	66.8	16.9	81.8	58.5	63.9
		aTLAS w/ Tip	61.4	44.7	52.6	49.7	76.0	61.5	63.2	50.9	84.4	58.9	62.9	96.8	83.6	83.2	21.3	66.1	87.1	50.7	65.4	17.0	87.0	59.8	62.9
	2	Tip	67.3	58.4	63.0	37.5	61.9	68.0	65.6	37.0	82.5	52.5	63.8	97.0	84.0	85.0	24.1	87.4	87.7	54.8	64.4	17.6	87.9	70.8	64.5
		LP++	67.5	56.4	66.4	42.5	67.7	69.8	64.8	36.4	81.0	53.7	63.4	97.1	83.7	84.5	23.9	83.8	86.5	57.0	72.1	17.6	86.6	70.0	65.1
		aTLAS	62.9	45.3	79.3	62.2	82.2	66.0	61.4	60.0	82.3	63.2	63.1	96.8	83.6	83.6	21.7	67.8	86.9	51.4	67.7	17.0	83.5	58.5	65.7
		aTLAS w/ Tip	66.8	51.4	65.3	63.9	85.2	59.9	64.3	54.6	83.0	60.4	63.6	96.8	83.6	83.9	22.0	66.3	87.1	52.3	65.9	17.3	88.8	63.3	65.7
	4	Tip	69.9	60.0	74.7	43.8	76.8	74.7	67.9	33.7	80.9	53.3	64.7	97.2	84.3	85.1	27.3	90.0	88.0	59.9	71.8	18.1	91.4	74.3	67.6
		LP++	70.4	61.9	69.7	48.7	71.1	74.4	66.8	32.7	82.5	56.3	64.3	95.9	83.5	85.2	26.3	87.2	89.6	58.2	71.1	18.4	88.6	73.0	67.1
		aTLAS	64.0	47.3	80.8	63.8	80.0	70.0	62.5	59.4	87.4	64.8	64.3	97.0	83.6	83.5	21.7	68.8	87.8	51.4	67.7	17.0	83.5	58.5	66.7
		aTLAS w/ Tip	70.9	61.5	61.8	73.1	89.9	75.2	65.9	59.9	85.2	62.9	64.9	97.1	83.8	84.5	24.3	72.1	89.0	54.5	74.6	18.0	89.9	67.5	69.4
	8	Tip	73.7	66.1	79.8	54.7	77.3	77.9	69.6	34.0	82.3	56.8	65.1	97.4	84.9	86.8	31.2	93.4	88.4	67.3	71.9	18.6	91.0	76.3	70.2
		LP++	71.6	64.5	78.5	54.4	81.0	77.2	69.0	30.8	83.2	57.9	65.6	96.9	84.3	86.0	28.9	88.2	89.4	62.9	73.7	19.4	90.7	76.4	69.6
		aTLAS	64.9	49.5	86.5	66.5	87.6	73.1	63.2	66.6	88.1	67.0	64.0	96.9	83.8	85.3	24.4	72.8	88.2	53.6	71.0	17.1	91.3	66.1	69.4
aTLAS w/ Tip		77.6	63.8	83.6	79.4	93.0	78.6	67.4	76.6	93.3	75.2	64.6	97.2	84.1	88.2	33.5	93.3	89.7	67.8	78.0	17.9	91.3	80.0	76.0	
16	Tip	75.2	68.3	83.1	62.8	82.0	80.5	71.8	33.1	84.9	58.3	65.8	97.8	85.0	88.0	34.5	94.3	88.6	70.5	72.1	19.9	92.3	78.3	72.3	
	LP++	74.8	69.2	81.7	57.5	84.8	80.1	70.1	40.4	84.8	59.6	66.9	97.2	84.8	87.6	30.6	89.5	89.6	65.1	70.7	20.7	91.3	78.0	71.6	
	aTLAS	67.4	55.0	88.5	70.5	90.6	75.2	66.3	69.0	94.5	77.3	65.2	97.8	84.9	86.5	23.1	69.7	91.6	55.5	77.7	17.9	90.3	66.2	71.8	
	aTLAS w/ Tip	80.5	64.9	88.8	84.4	93.9	84.0	72.7	73.8	92.9	78.1	65.8	97.1	84.1	89.0	41.0	94.5	91.2	72.8	76.2	20.0	94.5	80.8	78.2	
ViT-B/16	0	CLIP	64.0	45.0	56.6	42.9	50.6	67.6	65.5	44.6	90.7	68.2	68.8	97.8	85.7	86.9	24.4	71.1	87.2	56.8	77.6	22.9	84.4	66.0	64.8
	1	Tip	67.7	53.8	68.7	47.8	74.2	71.3	68.8	51.9	91.6	70.0	69.1	98.3	88.8	87.6	29.8	87.0	91.8	59.5	78.3	22.9	8		

Table 13: Accuracy of few-shot methods trained on ImageNet [52] and tested on out-of-domain datasets, for $k \in \{4, 16\}$. Results are produced by CLIP with ViT-B/32 backbone and averaged across three random seeds.

Method	$k = 4$					$k = 16$				
	INet	INet-A	INet-R	INet-Sketch	INetV2	INet	INet-A	INet-R	INet-Sketch	INetV2
Zero-shot	63.4	31.5	42.3	69.2	62.7	63.4	31.5	42.3	69.2	62.7
aTLAS	64.7 ± 0.0	30.8 ± 0.0	42.3 ± 0.0	69.9 ± 0.1	64.1 ± 0.0	65.5 ± 0.0	31.7 ± 0.1	42.9 ± 0.0	70.4 ± 0.0	64.7 ± 0.0
Tip-Adapter	64.9 ± 0.1	30.8 ± 0.1	41.3 ± 0.2	68.5 ± 0.1	63.4 ± 0.0	66.1 ± 0.0	29.1 ± 0.2	40.5 ± 0.0	67.7 ± 0.1	63.7 ± 0.1
LP++	65.7 ± 0.0	30.1 ± 0.2	41.0 ± 0.2	67.8 ± 0.1	64.4 ± 0.0	68.0 ± 0.0	29.2 ± 0.1	41.0 ± 0.0	67.0 ± 0.1	66.0 ± 0.0
aTLAS w/ Tip	66.0 ± 0.0	30.2 ± 0.1	41.6 ± 0.1	69.2 ± 0.2	64.5 ± 0.0	68.0 ± 0.0	29.4 ± 0.3	41.4 ± 0.1	68.9 ± 0.2	65.4 ± 0.1
aTLAS w/ LP++	66.0 ± 0.0	29.1 ± 0.2	41.2 ± 0.2	67.9 ± 0.4	64.8 ± 0.0	68.9 ± 0.0	28.7 ± 0.1	41.9 ± 0.0	67.5 ± 0.1	67.0 ± 0.0




Figure 13: Accuracy improvement of aTLAS (16-shot) using one task vector normalised by that of fine-tuning in the full parameter space (all training data). Each column corresponds to a unique task vector, and reflects the relative improvement it leads to on different target datasets. Each row reflects the relative improvement on a dataset, using different task vectors.

D.4 Out-of-domain generalisation

We show detailed results for out-of-domain generalisation over $k \in \{4, 16\}$ shots in Table 13. These results correspond to those presented in Figure 5c. aTLAS is the only method that consistently improves test accuracy over the zero-shot model on out-of-domain images. When combined with LP++ or Tip-Adapter, aTLAS can be observed to improve the out-of-domain generalisation of these methods.

D.5 Relative significance of individual task vectors

In this section, we examine the informativeness of a task vector across different target datasets. To this end, we apply aTLAS to each of the 22 datasets using only one task vector. For each dataset, we compute the relative accuracy improvement, that is, the accuracy improvement of aTLAS normalised by that of fine-tuning in the full parameter space. Note that aTLAS is applied under the 16-shot setting, while standard fine-tuning uses all training data available. Results are shown in Figure 13. We first note that certain datasets are more prone to accuracy improvement, such as EuroSAT, MNIST, etc., as indicated by the high percentage across entire rows. This is most likely due to the low intrinsic dimensionality of the task. In addition, we highlight the average improvement in the last row. Notably, certain task vectors, e.g., ImageNet task vector, are particularly informative while others, such as those from Flowers102 and OxfordPets are much less so. These results illustrate the varying contributions different task vectors can have depending on the target dataset, which also motivated subsequent efforts on careful task vector selection.

Table 14: Few-shot accuracy when only using a budget of b task vectors with different selection strategies. We report results for 4 and 16 shots. The results are averaged over 22 datasets and three random seeds. CLIP with ViT-B/32 backbone is used. Highest performance in each section is highlighted in bold.

Shots (k)	Strategy	$b = 0$	$b = 1$	$b = 2$	$b = 5$	$b = 10$	$b = 15$	$b = 21$
4	Random	60.39	63.5 \pm 0.0	64.8 \pm 0.1	67.1 \pm 0.2	69.3 \pm 0.1	69.7 \pm 0.1	70.0 \pm 0.0
	Features	60.39	65.6 \pm 0.1	67.6 \pm 0.2	68.8 \pm 0.1	69.5 \pm 0.1	69.8 \pm 0.1	70.0 \pm 0.0
	Grad. whole	60.39	64.4 \pm 0.1	65.8 \pm 0.2	67.2 \pm 0.1	69.1 \pm 0.1	69.6 \pm 0.1	70.0 \pm 0.0
	Grad. blockwise	60.39	67.3 \pm 0.2	68.2 \pm 0.0	68.9 \pm 0.1	69.1 \pm 0.2	69.7 \pm 0.2	70.0 \pm 0.0
16	Random	60.39	64.7 \pm 0.1	66.0 \pm 0.1	68.8 \pm 0.1	70.8 \pm 0.0	72.0 \pm 0.1	72.8 \pm 0.1
	Features	60.39	66.2 \pm 0.0	68.1 \pm 0.2	70.3 \pm 0.0	71.7 \pm 0.1	72.4 \pm 0.1	72.8 \pm 0.1
	Grad. whole	60.39	65.2 \pm 0.1	66.2 \pm 0.1	68.3 \pm 0.2	71.5 \pm 0.1	72.2 \pm 0.1	72.8 \pm 0.1
	Grad. blockwise	60.39	68.3 \pm 0.1	69.3 \pm 0.1	70.5 \pm 0.1	71.6 \pm 0.0	72.3 \pm 0.0	72.8 \pm 0.1

D.6 Task vector budget and selection

In this section, we provide details for selecting a budget of b task vectors with feature-based and gradient-based strategies, as introduced in Section 5.2.

Feature based selection. For each dataset \mathcal{D}_i , we compute the average image representation \bar{z}_i of the dataset using the zero-shot model as follows

$$\bar{z}_i = \mathbf{E}_{\mathbf{x} \in \mathcal{D}_i} [f(\mathbf{x}; \theta_0)]. \quad (13)$$

Given a target dataset \mathcal{D}_t , we simply compute the cosine similarity between its feature representation \bar{z}_t and that of each other dataset \bar{z}_i , $i \neq t$. Subsequently, b task vectors corresponding to the datasets with highest similarity will be selected.

Gradient-based selection.

Given a target dataset \mathcal{D}_t , we may directly compute the gradient with respect to the m learnable coefficients for each of the n task vectors. However, as one important motivation behind task vector selection is to reduce memory consumption, using all n task vectors to compute the gradient defeats the purpose. Therefore, we instead only load a group of b task vectors ($b < n$), compute the gradient with respect to their learnable coefficients, and repeat for other groups. With this sequential computation, the gradient across different groups is not calibrated. Nevertheless, we empirically found this strategy to work well. Denote the partial derivative of the loss on dataset \mathcal{D}_t with respect to a learnable coefficient $\lambda_i^{(j)}$ by $\dot{\lambda}_i^{(j)}$, such that

$$\dot{\lambda}_i^{(j)} = \mathbf{E}_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_t} \left[\frac{\partial \mathcal{L} \left(f(\mathbf{x}; \theta_0 + \sum_{i=1}^b \Lambda_i \tau_i), \mathbf{y} \right)}{\partial \lambda_i^{(j)}} \right]. \quad (14)$$

For the i -th task vector, we may compute its L_1 gradient norm, i.e., $\left\| \dot{\lambda}_i^{(1)}, \dots, \dot{\lambda}_i^{(m)} \right\|_1$, and select task vectors with larger gradient. Alternatively, we may select task vectors block by block. Specifically, for the j -th parameter block, we inspect the absolute values of the partial derivatives for the corresponding coefficients, i.e., $\left| \dot{\lambda}_i^{(j)} \right|$, and select task vectors with higher absolute values. This process is repeated for each parameter block, thus allowing different parameter blocks to have different selections. Crucially, for low budgets, particularly $b = 1$, this enables our method to effectively exploit more task vectors than the budget specifies. The impact of this can be observed in Table 14 (corresponding to Figure 6), that blockwise selection significantly outperforms other methods when the budget is low.

D.7 LoRAs as task vectors

We fine-tune LoRAs for ViT-B/32 using the LoRA-Torch [36] library with ranks 4, 16 and 64. We stop at rank 64 as we do not observe improvements beyond it. We train LoRAs on attention and MLP layers and use the same settings as for full finetuning but with a learning rate of 10^{-3} .

Table 15 shows additional results using LoRAs as task vectors. We study learning the effect of fine-tuning the LoRAs task vectors on attention layers only (as done in the original LoRA paper [23]) or on the MLPs. Although the original LoRA paper recommendeds training on the attention layers only [23], we observe that training on MLP layers is important to produce strong LoRA task vectors.

Table 15: Additional few-shot recognition results using LoRAs trained on attention layers, MLP layers or both. Results are averaged across 22 datasets over three seeds, with $\times 1$ standard deviation. Rank 16 is used for LoRAs.

Task vector type	Method	0-shot	1-shot	2-shot	4-shot	8-shot	16-shot
LoRAs (Attn.)	aTLAS	60.4	63.5 \pm 0.1	65.1 \pm 0.1	66.6 \pm 0.1	67.9 \pm 0.1	69.5 \pm 0.0
LoRAs (MLP)	aTLAS	60.4	63.8 \pm 0.1	66.2 \pm 0.1	68.3 \pm 0.1	70.5 \pm 0.1	71.4 \pm 0.0
LoRAs (Attn. & MLP)	aTLAS	60.4	64.6 \pm 0.1	66.6 \pm 0.2	68.7 \pm 0.1	70.4 \pm 0.1	71.8 \pm 0.1
LoRAs (Attn. & MLP)	aTLAS w/ LP++	60.4	67.1 \pm 0.3	70.9 \pm 0.1	73.4 \pm 0.1	75.9 \pm 0.1	78.2 \pm 0.1
LoRAs (Attn. & MLP)	aTLAS w/ Tip-Adapter	60.4	67.5 \pm 0.1	70.0 \pm 0.1	72.4 \pm 0.1	74.9 \pm 0.1	77.0 \pm 0.1
Standard	aTLAS w/ LP++	60.4	68.9 \pm 0.2	71.7 \pm 0.1	74.1 \pm 0.1	75.8 \pm 0.1	77.9 \pm 0.0
Standard	aTLAS w/ Tip-Adapter	60.4	68.6 \pm 0.4	71.6 \pm 0.2	74.3 \pm 0.1	76.4 \pm 0.1	78.2 \pm 0.0

Table 16: Few-shot recognition performance with gradient-free optimisation. Results are averaged accuracy over 22 datasets, with $1\times$ standard error over 3 random seeds.

Scaling	Use gradient	Memory (GB)	0-shot	1-shot	2-shot	4-shot	8-shot	16-shot
Anisotropic	Yes	10	60.4	66.7 \pm 0.23	68.3 \pm 0.28	70.0 \pm 0.01	71.7 \pm 0.11	72.8 \pm 0.08
Isotropic	No	4	60.4	63.1 \pm 0.45	64.2 \pm 0.35	65.0 \pm 0.12	65.7 \pm 0.05	65.4 \pm 0.14
Anisotropic	No	4	60.4	61.3 \pm 0.08	61.5 \pm 0.04	61.5 \pm 0.04	61.6 \pm 0.03	61.6 \pm 0.02

D.8 Gradient-free optimisation

An alternative to save memory during training is to utilise gradient-free methods to learn the coefficients. We follow previous work on the combination of LoRAs [24] and use the nevergrad [49] library. We observe a memory usage reduction of 60% from 10GB to 4GB calculated using a dedicated pytorch function⁶. Results for few-shot recognition are summarised in Table 16. We show that although gradient-free optimisation improves upon the zero-shot model, the performance quickly plateaus as the amount of data increases. In addition, learning anisotropic scaling results in worse performance, most likely due to the relatively high number of parameters.

⁶https://pytorch.org/docs/stable/generated/torch.cuda.memory_allocated.html

Table 17: Accuracy after fine-tuning on different percentage of training data for variants of aTLAS $\times K$ and LoRAs [23]. Results are averaged across 22 datasets. Highest accuracy in each section is highlighted in bold.

Method	# Params	1%	5%	10%	25%	35%	50%	100%
aTLAS	2k	68.5	71.5	72.6	73.6	74.6	75.4	76.4
aTLAS $\times 5$	10k	69.3	72.9	74.7	76.2	76.8	77.5	78.8
aTLAS $\times 20$	40k	69.5	74.0	75.6	77.5	78.2	78.9	80.5
aTLAS $\times 80$	160k	70.2	74.7	76.2	77.9	78.9	80.0	82.0
aTLAS $\times 1200$	2.4M	71.3	75.0	76.6	78.3	80.2	81.5	83.9
LoRA (rank=16)	2.4M	68.8	74.1	75.6	76.8	79.0	80.6	83.6

E Unsupervised FixMatch

We provide more details on the Unsupervised FixMatch (UFM) approach in this section. FixMatch [54] utilises a labelled set to guide training, which is given as part of the semi-supervised learning protocol, while we produce a class-balanced “labelled” set from unlabelled images. Given a target dataset \mathcal{D}_t consisting of N unlabelled images, we first rank the examples by the prediction scores from the zero-shot model across C classes. We then select the top $\min(N/C, 100)$ examples, that is, at most 100 examples per class, as a trusted set in absence of a labelled set. The standard cross-entropy loss is applied to the trusted set. For the rest of the unlabelled images, we use a weakly augmented (Open-CLIP [26] validation augmentations) view of an image to produce pseudo-labels, and incur a loss on the strongly augmented view (Tip-Adapter [69] augmentations). Denote an image with weak augmentation by \mathbf{x} , its strongly augmented view by \mathbf{x}' , and the predictions made by network by $\hat{\mathbf{y}}$ and $\hat{\mathbf{y}}'$, respectively, the unsupervised loss can be expressed as

$$\ell_u(\hat{\mathbf{y}}, \hat{\mathbf{y}}') = -\mathbb{1}(\max(\sigma(\hat{\mathbf{y}})) > \omega) \sigma(\hat{\mathbf{y}})^\top \log(\hat{\mathbf{y}}'), \quad (15)$$

$$\sigma(\hat{\mathbf{y}}) = \frac{\hat{\mathbf{y}}^{0.5}}{\mathbf{1}^\top \hat{\mathbf{y}}^{0.5}}, \quad (16)$$

where $\mathbb{1}(\cdot)$ denotes the indicator function, $\sigma(\cdot)$ performs re-normalisation with adjusted temperature scaling, and ω is a confidence threshold that is linearly adjusted from 0.9 to 1 during training. The trusted set is re-estimated at the beginning of each epoch to account for the improving accuracy of the model. In training, images in the trusted set are over-sampled to constitute one fourth of each batch, as this practice prevents the model from diverging due to confirmation bias [2, 54].

F Details of aTLAS $\times K$ variants

Dividing a parameter block into K random partitions allows us to introduce more learnable coefficients to each block, thus scaling up our method flexibly. One draw back of this approach, however, is that masks for the partitions have to be stored in memory, resulting in a linear memory increase with respect to the size of the parameter block and the value K . To reduce the memory consumption of aTLAS $\times K$ variants, we only apply it to LoRAs task vectors. Nevertheless, these memory requirements could most likely be reduced by exploiting sparse matrices or memory efficient matrix indexing techniques, which we plan to investigate in the future.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: All claims made in the introducing summarise findings in Sections 4, 5 and 6. We clearly enunciate our claims and hypothesis by numbering them by order of appearance in the main body of the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalise to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss limitations of our work in a paragraph found in the conclusion, Section 8. The dominant limitation we observe to our work is possible memory limitations when applied to larger models with LoRA task vectors. We otherwise tested our approach on a larger array of varied image classification datasets, ranging from simple to larger datasets. Some of the datasets we tested on are specialized while others are more generic. This should ensure that our results are generalisable to a large array of tasks.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognise that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: We do not contribute any theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data is provided or not)?

Answer: [Yes]

Justification: The algorithms we propose is fully explained in Section 3 with all the information needed to reproduce our results being available the appendix Sections A, D.1, D.3, E and F. Furthermore, the pre-trained models and every dataset we use are publicly available.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data is provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognise that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility.

In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We have released the entirety of our code base and task vector checkpoint under the link provided in the paper. All the data used in this paper is not owned by us and is publicly available.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimiser, etc.) necessary to understand the results?

Answer: [Yes]

Justification: All hyper-parameters, number of training samples, optimiser used, model architectures and everything else needed to reproduce our results are available the appendix Sections A, D.1, D.3, E and F.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We perform 3 independent runs in Figures 5a, ?? and ?? with standard errors being reported. The standard error is computed as the standard deviation over 3 runs divided by the number of runs (3). This is calculated with the numpy library.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [No]

Justification: Although we do not report compute requirements in the paper, on a single A100 or 4090GPU except for ViT-L/14 experiments that were performed on 2 A100. The typical run time for an experiment is 2 hours with ViT-L/14 experiments going up to 8h. We estimate the total compute needed to 1000 GPU-hours for repeated results over 3 seeds. The research project includes failed experiments and iterations on the method that are not reported.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: We find no direct societal impact for this research paper as research is conducted on controlled open-sourced datasets. This paper did not study applicability to specialized datasets that can be used to impact society.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimising neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognise that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All datasets we use are credited in Section A.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: No participants.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognise that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.