

In Appendix A, we provide training details for each of our fine-tuning run. In Appendix B, we provide details on how to compute the concatenated feature vector  $\Phi_\theta(\mathbf{x})$  from a ViT-B/32 CLIP model for LDIFS. In Appendix C, we provide additional results on **i**) observing how fine-tuning can lead to concept forgetting, **ii**) our investigations on why LP-init-L2SP is a better baseline compared to others in avoiding concept forgetting and **iii**) performance of our proposed LDIFS regularizer.

## A TRAINING DETAILS

**Training datasets and hyper-parameters:** We fine-tune the CLIP ViT-B/32 model on each of the 9 image classification datasets: **a**) Stanford Cars, **b**) CIFAR-10, **c**) CIFAR-100, **d**) DTD, **e**) EuroSAT, **f**) GTSRB, **g**) MNIST, **h**) RESISC45, **i**) SVHN. For each dataset, we have a separate fine-tune run for each of the baselines discussed in §2. For each run, we train using the AdamW Loshchilov & Hutter (2017) optimizer, with an initial learning rate of  $1e - 5$ , a weight decay of 0.1 and a cosine learning rate scheduler with a warmup length of 500. For all the runs, we use a batch-size of 128. Following the code of Ilharco et al. (2022a), we use the following number of epochs to fine-tune each dataset: **a**) Stanford Cars: 35 epochs, **b**) CIFAR-10/100: 10 epochs, **c**) DTD: 76 epochs, **d**) EuroSAT: 12 epochs, **e**) GTSRB: 11 epochs, **f**) MNIST: 10 epochs, **g**) RESISC45: 15 epochs and **h**) SVHN: 10 epochs. For each dataset, we keep a minimum of 10 epochs for fine-tuning.

**Compute and number of experiments:** Each of our fine-tuning run is done on a single NVIDIA A100 GPU. As there are a total of 9 classification datasets and a 8 fine-tuning baselines (including our proposed LDIFS), we perform a total of 72 fine-tune training runs from a pre-trained CLIP ViT-B/32. Furthermore, in order to produce the plots capturing  $\mathcal{A}_{ZS}$ ,  $\mathcal{A}_{LP}$ , as well as  $\ell_2$  distance in parameter and feature space, we store intermediate checkpoints over the course of fine-tuning. Particularly, for each run, we evaluate checkpoints after every 20% of training completion. Finally, in order to obtain the full set of results in Table 6, for the 72 fine-tuned models, we evaluate each of them on 9 datasets (i.e., including the test set of the fine-tuned task), thereby having a total of 648 evaluations. Similar to training, each evaluation is also performed on a single NVIDIA A100 GPU.

**Choosing  $\lambda_{LDIFS}$ :** One hyper-parameter which the LDIFS regularizer introduces is  $\lambda_{LDIFS}$ . A higher value of  $\lambda_{LDIFS}$  encourages the model to preserve features of the original foundation model and vice versa. For each classification task, we performed a grid search over  $\lambda_{LDIFS} \in \{0.01, 0.05, 0.1, 0.5, 1, 10, 100\}$  and cross-validated this hyper-parameter on a held-out validation set, choosing the value which produces the best performance on the validation set. We found  $\lambda_{LDIFS} = 10$  to produce the best performance over datasets in general, so all the results we present in this paper are with  $\lambda_{LDIFS}$  set to 10.

## B COMPUTING $\Phi_\theta(\mathbf{x})$

In this section, we discuss how we compute the concatenated feature vector  $\Phi_\theta(\mathbf{x})$  given input  $\mathbf{x}$  and model parameters  $\theta$ , specifically for a ViT-B/32 model. This is used for computing LDIFS, our proposed regularizer for fine-tuning.

Let the feature output from layer  $l$  in the network for input  $\mathbf{x}$  be  $\Phi_{\theta(l)}(\mathbf{x}) \in \mathbb{R}^l$ . Then the normalized feature vector for layer  $l$  can be represented as  $\frac{\Phi_{\theta(l)}(\mathbf{x})}{\|\Phi_{\theta(l)}(\mathbf{x})\|}$ . In order to form the concatenated feature vector, one can technically take features from every intermediate layer of the network. However, storing all the features is memory intensive. Thus, we follow a similar approach to the LPIPS Zhang et al. (2018) implementation for VGG and AlexNet, and choose 5 intermediate points in the ViT-B/32 architecture to collect features from. For a single input image  $\mathbf{x} \in \mathbb{R}^{3 \times 224 \times 224}$ , each of the intermediate feature representations has a dimension of  $\mathbb{R}^{50 \times 768}$  with 50 tokens and 768 dimensional representation for each token. When normalizing the feature vector, we flatten this vector out to a single 38400 dimensional vector. Thus the full concatenated feature vector  $\Phi_\theta(\mathbf{x})$  has dimensionality  $5 \times 38400 = 192000$ . However, note that there can be other ablations to this design and we leave that for future exploration.

**LDIFS vs LPIPS Zhang et al. (2018):** One can find similarities between our proposed LDIFS regulariser and the LPIPS metric Zhang et al. (2018) used for measuring perceptual similarity between images. However, while LPIPS uses feature space distance on a *pre-trained, frozen* model to find perceptual similarity between pairs or sets of images, LDIFS instead feature space distance

between a pair of pre-trained and fine-tuned model for the same image, to preserve the input-output behaviour of the pre-trained model.

## C ADDITIONAL RESULTS

In this section, we present additional empirical results to supplement our observations and conclusions in the main paper.

### C.1 CRIPPLING EFFECT OF FINE-TUNING

In §3, we observed how most existing fine-tuning methods, while gaining state-of-the-art performance on fine-tuned tasks, can lead to concept forgetting in models. In this section, we provide additional empirical results to further strengthen those observations.

In Figure 7 and Figure 15, we present the  $\mathcal{A}_{ZS}$  and  $\mathcal{A}_{LP}$  accuracy for models fine-tuned on each of the 9 classification tasks, on their respective test sets. This shows expected behaviour as models broadly achieve very high test set accuracy, which increases over the course of fine-tuning, on their respective fine-tuned tasks. Note that LP-init baselines seem to have relatively lower performance on  $\mathcal{A}_{ZS}$  accuracy compared to their ZS-init counter-parts. However, this reduced performance is only limited to  $\mathcal{A}_{ZS}$  as this does not translate to a reduced performance in  $\mathcal{A}_{LP}$ . Moreover, L2SP baselines (both ZS and LP-init) obtain relatively lower fine-tuned accuracy both in case of  $\mathcal{A}_{ZS}$  and  $\mathcal{A}_{LP}$ .

In Figure 8 and Figure 16, we present the  $\mathcal{A}_{ZS}$  and  $\mathcal{A}_{LP}$  for models fine-tuned on EuroSAT (first row), GTSRB (second row) and SVHN (third row) as captured on 8 classification datasets different from their respective fine-tuned set. Broadly, the ZS and LP performance for all models drops on other datasets as they are fine-tuned, thereby capturing concept forgetting in the fine-tuned models. Among the baselines, we consistently observe LP-init-L2SP to perform better than others in avoiding concept forgetting. This is evident through the distinctly higher  $\mathcal{A}_{ZS}$  and  $\mathcal{A}_{LP}$  accuracies over fine-tuning for the LP-init-L2SP baselines.

### C.2 ANALYSING L2SP

In this section, we provide additional results to supplement our observations related to investigating the L2SP baseline in §4 of the main paper. In Figure 19, we present the  $\ell_2$  distance in parameter space  $\|\theta - \theta_0\|_2^2$  between the pre-trained and current models captured over fine-tuning. For all the datasets, we observe that L2SP baselines while initially diverging slightly in the parameter space, converge back to a model having low  $\ell_2$  parameter space distance from the original foundation model. Other baselines on the other hand, completely diverge away from the original model in the parameter space.

To further investigate the change in input-output behaviour of the model over fine-tuning, we measure the distance in feature space (see Equation (3)) over fine-tuning. In Figure 20, we present the  $\ell_2$  distance in feature space captured for models fine-tuned on EuroSAT. Again, consistent with our previous observation, we find that unlike other baselines, L2SP first diverges away from the original foundation model and then converges back to the original input-output behaviour, as is indicative through a decreasing L2 feature space distance in the later stages of fine-tuning. This observation is consistent on the EuroSAT train set, the EuroSAT test set as well as on other datasets, thereby providing our motivation for the LDIFS regularizer.

### C.3 PERFORMANCE OF LDIFS

In this section, we provide additional results to supplement observations related to the performance of the proposed LDIFS regularizer.

**Analysing LDIFS on parameter space and feature space distance:** In Figure 9 and Figure 10, we plot the  $\ell_2$  distance in parameter/weight and feature space respectively (same as Figure 3 and Figure 4 in the main paper), but with the LDIFS baselines added in. In the weight space, LDIFS, while lower than other fine-tuning methods, has a relatively higher  $\ell_2$  distance from the pre-trained model compared to L2SP. On the contrary, in the feature space, LDIFS consistently gets the lowest distance from the pre-trained model. This is expected and was indeed the purpose behind LDIFS’s

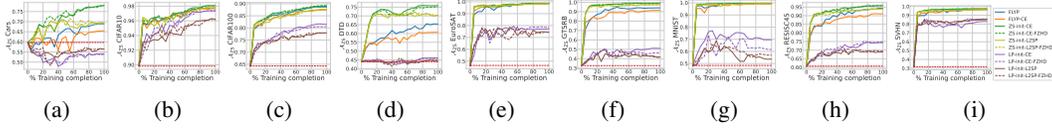


Figure 7: **Test set ZS Accuracy  $\mathcal{A}_{ZS}$  for different fine-tuning methods** on 9 image classification tasks: (a) Stanford Cars, (b) CIFAR-10, (c) CIFAR-100, (d) DTD, (e) EuroSAT, (f) GTSRB, (g) MNIST (h) RESISC45 and (i) SVHN.  $\mathcal{A}_{ZS}$  generally rises over the course of fine-tuning. However, for the ZS-init-L2SP and LP-init-L2SP baselines, the gain in  $\mathcal{A}_{ZS}$  is relatively lower. Furthermore, for LP-init baselines, the performance is consistently lower compared to other baselines

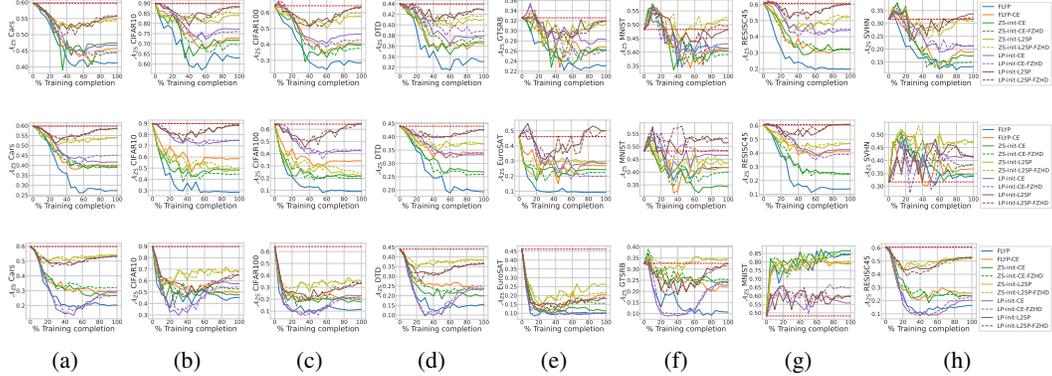


Figure 8: **ZS Accuracy  $\mathcal{A}_{ZS}$  for models fine-tuned on EuroSAT (first row), GTSRB (second row), and SVHN (third row) and evaluated on 8 different datasets different from their fine-tuning dataset.** Most fine-tuning methods show a drop in  $\mathcal{A}_{ZS}$  performance over the course of fine-tuning indicating a reduction in the model’s transferability.

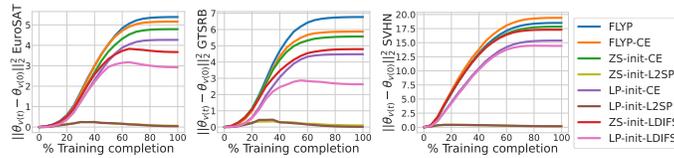


Figure 9:  **$\ell_2$  distance in weight space  $\|\theta_{v(t)} - \theta_{v(0)}\|_2^2$**  between the image encoder fine-tuned to the current time-step  $f_{\theta_{v(t)}}$  and the pre-trained image encoder  $f_{\theta_{v(0)}}$  over the course of fine-tuning.

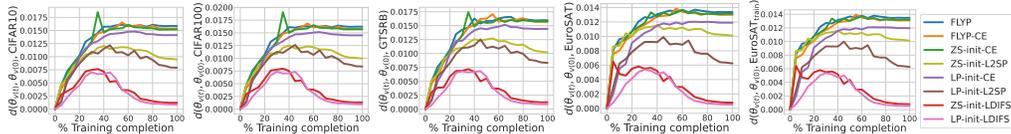


Figure 10:  **$\ell_2$  distance in feature space  $d(\theta_{v(t)}, \theta_{v(0)}, \mathcal{D})$** , between image encoders,  $f_{\theta_{v(t)}}$  and  $f_{\theta_{v(0)}}$  computed over different fine-tuning methods for models fine-tuned on EuroSAT.

design which is to minimize the difference in input-output behaviour, captured through feature space distance between the fine-tuned and pre-trained models and not necessarily constrain the parameter space of the fine-tuned model to lie close to the pre-trained model. Furthermore, note from Figure 10, that although we are applying the LDIFS regularizer only on the EuroSAT samples during fine-tuning, the preservation of features extends even to other datasets like CIFAR10, CIFAR100 and GTSRB, as is indicated by the low feature space distance even on these datasets.

**Results and observations:** Firstly, in Figure 13 and Figure 17, we present the test set accuracy of all fine-tuning methods including LDIFS on all 9 classification tasks. We see that LDIFS performs competitively with all other baselines and improves on L2SP consistently. In Figure 14 and Figure 18, we present the  $\mathcal{A}_{ZS}$  and  $\mathcal{A}_{LP}$  accuracy respectively, for models fine-tuned on EuroSAT (first row),

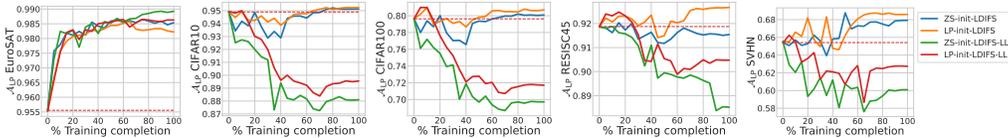


Figure 11:  $\mathcal{A}_{LP}$  for models trained with LDIFS using the full concatenated feature vector vs just the last layer (LL) feature vectors. Full feature vectors with earlier features is crucial for LDIFS’s performance.

GTSRB (second row) and SVHN (third row) over the course of fine-tuning. Furthermore, we also provide the  $\Delta_{ZS}$ ,  $\Delta_{LP}$  and the  $\mathcal{A}_{LP}$  values for fully fine-tuned models on all 9 classification tasks and all baselines in Table 6.

As is crystal clear from our results, LP-init-LDIFS consistently minimizes concept forgetting without compromising on performance on the fine-tuned task. This is evident from its high  $\mathcal{A}_{ZS}$ ,  $\mathcal{A}_{LP}$  (see Figure 14 and Figure 18) and high  $\Delta_{ZS}$  and  $\Delta_{LP}$  (see Table 6) when evaluated on tasks other than the fine-tuning task at hand. Note from Table 6 that even when LP-init-LDIFS is not achieving the best performance on a certain dataset pair, it is often a close second with very little difference compared to the best performing baseline. In addition to this advantage, its  $\mathcal{A}_{LP}$  accuracy on the fine-tuned task itself is very competitive with the top scores obtained by other fine-tuning baselines and provides a consistent improvement on L2SP.

Finally, in Table 7, we present the full results of the 2 task sequence setup for continual fine-tuning. Even in this setting, we find LP-initLDIFS to not only preserve performance of the first fine-tuning task during the second fine-tuning stage but also to preserve performance of all other tasks during the entire sequence of end-to-end fine-tuning.

**Ablation with last layer features:** In Figure 11, we present an ablation where we fine-tune on EuroSAT using LDIFS but using just the last layer features as opposed to the full concatenated feature vector. Clearly, including earlier features when computing the LDIFS regularizer is crucial to its performance as preserving just the last layer features leads to a significant performance drop compared to using the full feature vector.

#### C.4 COMPARISON WITH ADAPTERS AND PROMPT TUNING

In Table 5, we present results of fine-tuning a CLIP ViT-B/32 model using linear probing, CoOp, CLIP-Adapter, Tip-Adapter and full end-to-end fine-tuning using ZS-init-CE loss on 9 downstream image classification tasks. We consistently observe end-to-end fine-tuning to produce the best test set accuracy across all 9 tasks. Furthermore, we also note that in case of tasks where the gap between linear probe performance and end-to-end fine-tuning performance is significant (e.g. SVHN), adapter and prompt tuning approaches don’t cover this performance gap. Hence, if the pre-trained encoder is not performant on a downstream task, the only way to achieve state-of-the-art results on the task requires end-to-end fine-tuning the model on the task itself. This observation thus necessitates study into better end-to-end fine-tuning methods for foundation models.

#### C.5 COMPARISON WITH WISE-FT

In this section, we perform an ablation to see the effect of Wise-FT Wortsman et al. (2022b) on preventing concept forgetting. Wise-FT is a weight interpolation method which forms a linear combination between the pre-trained  $\theta_0$  and fine-tuned  $\theta_f$  model parameters:

$$\theta_{wse} = \alpha\theta_0 + (1 - \alpha)\theta_f \quad (6)$$

Since this can work with any fine-tuned model, we can apply it on all the end-to-end fine-tuning baselines in this work. Specifically, we tune the hyperparameter  $\alpha$  on a held-out validation set of the downstream task at hand to maximum validation accuracy. We evaluate concept forgetting on 3 downstream tasks: CIFAR-10, EuroSAT and SVHN using the  $\Delta_{LP}$  metric and report the results in Table 4. For each dataset, we evaluate the mean  $\Delta_{LP}$  on 5 other downstream tasks: Cars, DTD, GTSRB, RESISC45 and MNIST. Our observations show a consistent reduction in concept forgetting across all fine-tuning baselines. However, the order of performance between the fine-tuning baselines does not change and LDIFS still maintains superior performance over other baselines both pre and

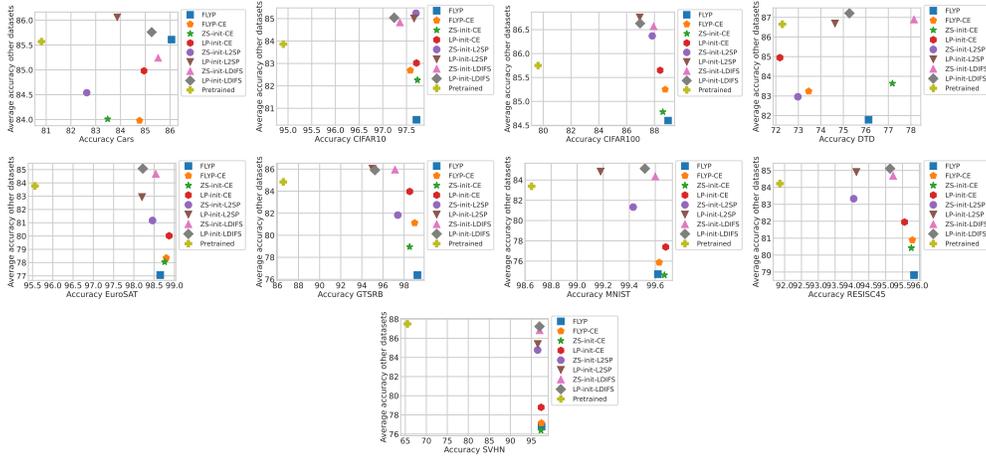


Figure 12: **Visualizing test LP accuracy of fine-tuning baselines on different datasets.** For each dataset, we plot its test set LP accuracy on the x-axis and the average test set LP accuracy of the remaining datasets on the y-axis.

Dataset	FLYP		FLYP-CE		ZS-init-CE		LP-init-CE (LP-FT)				ZS-init-L2SP		LP-init-L2SP		ZS-init-LDIFS (Ours)		LP-init-LDIFS (Ours)	
	$\Delta_{LP}$	$\Delta_{LP} (+wse)$	$\Delta_{LP}$	$\Delta_{LP} (+wse)$	$\Delta_{LP}$	$\Delta_{LP} (+wse)$	$\Delta_{LP}$	$\Delta_{LP} (+wse)$	$\Delta_{LP}$	$\Delta_{LP} (+wse)$	$\Delta_{LP}$	$\Delta_{LP} (+wse)$	$\Delta_{LP}$	$\Delta_{LP} (+wse)$	$\Delta_{LP}$	$\Delta_{LP} (+wse)$	$\Delta_{LP}$	$\Delta_{LP} (+wse)$
CIFAR10	-5.14	-0.11	-1.33	-0.16	-1.67	-0.1	-1.49	-0.01	1.58	1.64	1.49	1.96	1.53	1.99	<b>1.73</b>	<b>2.01</b>		
EuroSAT	-6.5	-1.32	-5.53	-1.04	-5.62	-1.23	-4.08	-0.76	-1.77	-0.12	-0.87	0.08	0.83	1.2	<b>1.24</b>	<b>1.96</b>		
SVHN	-10.82	-6.24	-10.18	-6.37	-10.98	-7.2	-9.2	-4.11	-2.65	-0.66	-2.13	-0.23	-0.57	0.12	<b>-0.18</b>	<b>0.33</b>		

Table 4:  $\Delta_{LP}$  without and with Wise-FT Wortsman et al. (2022b) for all fine-tuning baselines on CIFAR-10, EuroSAT and SVHN.

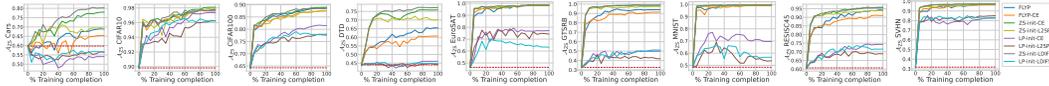


Figure 13: **Comparing test set ZS Accuracy  $\mathcal{A}_{ZS}$  for different fine-tuning methods with LDIFS on 9 image classification tasks.** While LP-init baselines generally underperform on  $\mathcal{A}_{ZS}$ , we find ZS-init-LDIFS to be competitive with the best baselines.

post application of Wise-FT. Thus, the results show that Wise-FT can be combined with any E2E fine-tuning method to further minimize concept forgetting.

### C.6 VISUALIZING TABLE 1

In this section, we present the results of Table 1 using an alternate visualization. For each of the 9 downstream datasets, we plot the LP accuracy on the dataset itself on the x-axis and the average LP accuracy on all other datasets on the y-axis. The best baselines should lie on the top right corner of this plot as that indicates that as the model is fine-tuned on the dataset, it still retains performance on other datasets. Results are in Figure 12. Clearly, we consistently see the LDIFS baselines to lie at the top right, again indicative of their performance in not just minimizing concept forgetting but also obtaining excellent downstream task performance.

### C.7 GENERALITY OF CONCEPT FORGETTING BEYOND CLIP ViT-B/32

To study the relevance of concept forgetting beyond a CLIP ViT model, in Table 8 and Table 9, we apply the end-to-end fine-tuning methods to fine-tune a CLIP RN-50 model and a FLAVA Singh et al. (2022) ViT-B/16 model respectively on EuroSAT. From both these tables, we observe both concept forgetting to exist and LP-init-LDIFS to minimize it, thereby validating our assumptions on models beyond a CLIP ViT model.

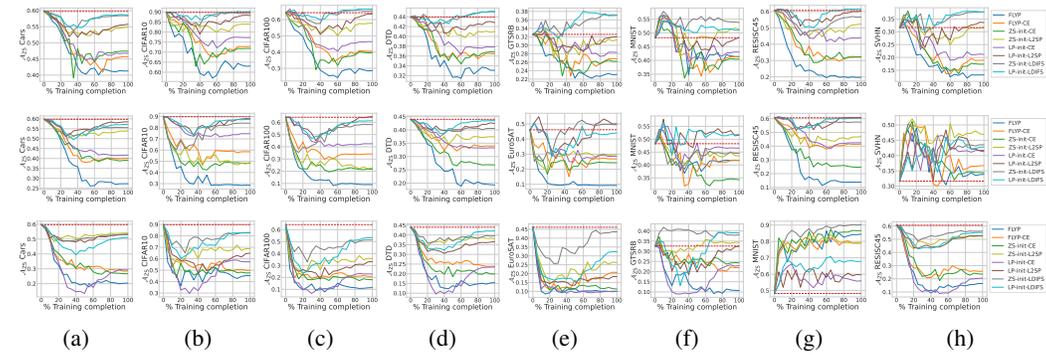


Figure 14: Comparing ZS Accuracy  $\mathcal{A}_{ZS}$  of different fine-tuning methods with LDIFS for models fine-tuned on EuroSAT (first row), GTSRB (second row) and SVHN (third row) and evaluated on 8 datasets different from their fine-tuning dataset. LP-init-LDIFS almost consistently beats all other baselines including LP-init-L2SP in preserving the model’s transferability.

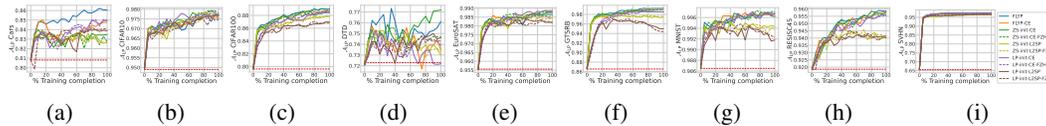


Figure 15: Test set LP Accuracy  $\mathcal{A}_{LP}$  for different fine-tuning methods on 9 image classification tasks: (a) Stanford Cars, (b) CIFAR-10, (c) CIFAR-100, (d) DTD, (e) EuroSAT, (f) GTSRB, (g) MNIST (h) RESISC45 and (i) SVHN.  $\mathcal{A}_{LP}$  generally rises over the course of fine-tuning. However, for the ZS-Init-L2SP and LP-Init-L2SP baselines, the gain in  $\mathcal{A}_{LP}$  is relatively lower. Furthermore, for LP-init baselines, the performance is consistently lower compared to other baselines

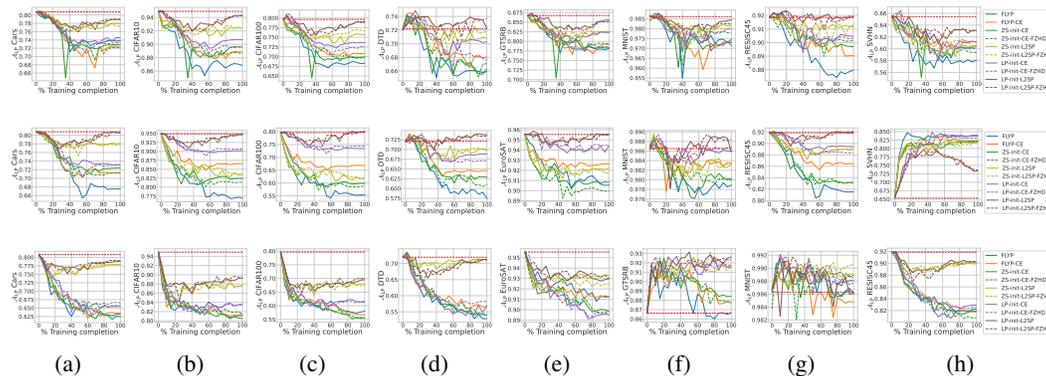


Figure 16: LP Accuracy  $\mathcal{A}_{LP}$  for models fine-tuned on EuroSAT (first row), GTSRB (second row), and SVHN (third row) and evaluated on 8 different datasets different from their fine-tuning dataset. Most fine-tuning methods show a drop in  $\mathcal{A}_{LP}$  performance over the course of fine-tuning indicating a reduction in the model’s transferability.

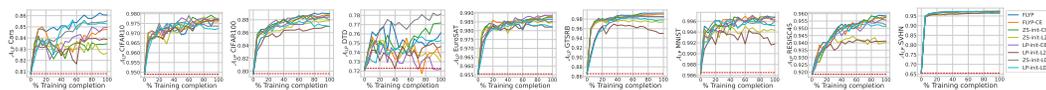


Figure 17: Comparing test set LP Accuracy  $\mathcal{A}_{LP}$  for different fine-tuning methods with LDIFS on 9 image classification tasks.

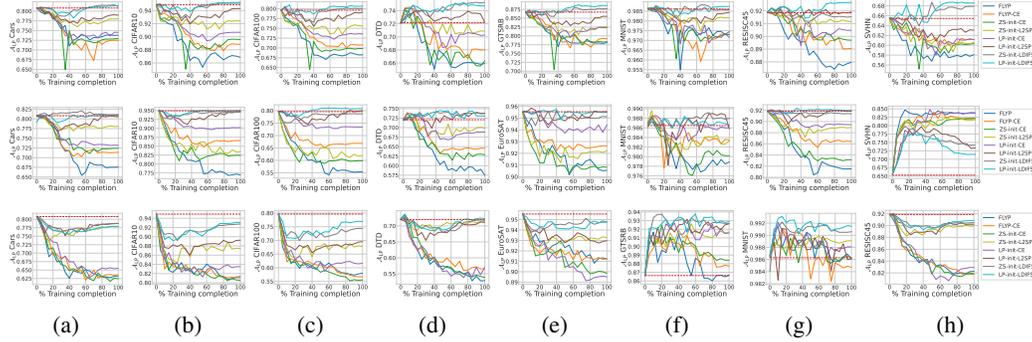


Figure 18: **Comparing LP Accuracy  $A_{LP}$  of different fine-tuning methods with LDIFS for models fine-tuned on EuroSAT (first row), GTSRB (second row) and SVHN (third row) and evaluated on 8 datasets different from their fine-tuning dataset. LP-init-LDIFS almost consistently beats all other baselines including LP-init-L2SP in preserving the model’s transferability.**

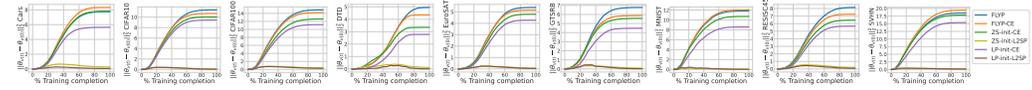


Figure 19:  **$\ell_2$  distance in weight space  $\|\theta - \theta_0\|_2^2$  between pre-trained image encoder  $f_{\theta_0}$  and fine-tuned image encoder  $f_{\theta}$  over the course of fine-tuning. Apart from ZS-init-L2SP and LP-init-L2SP, all fine-tuning baselines diverge in weight space over the course of fine-tuning.**

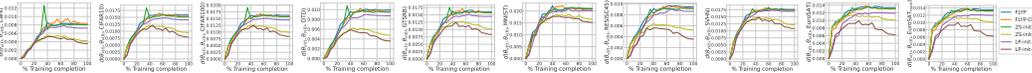


Figure 20:  **$\ell_2$  distance in feature space  $d(\theta, \theta_0, \mathcal{D})$ , between fine-tuned and pre-trained image encoders,  $f_{\theta}$  and  $f_{\theta_0}$  computed over different fine-tuning methods for models fine-tuned on EuroSAT.**

Dataset	LP	CoOp	CLIP-Adapter	Tip-Adapter	E2E Fine-tuning
Cars	80.80	80.88	81.24	81.32	<b>83.48</b>
CIFAR10	94.92	95.02	94.87	95.13	<b>97.73</b>
CIFAR100	79.62	80.12	80.06	80.82	<b>88.6</b>
DTD	72.08	72.01	71.87	72.62	<b>77.18</b>
EuroSAT	95.56	95.14	95.38	96.23	<b>98.76</b>
GTSRB	86.70	86.81	87.45	88.04	<b>98.52</b>
MNIST	98.65	98.84	99.03	99.11	<b>99.67</b>
RESISC45	91.86	91.79	91.82	91.80	<b>96.76</b>
SVHN	65.47	67.28	69.76	69.23	<b>97.3</b>

Table 5: Test set accuracy obtained from linear probing (LP), CoOp, CLIP-Adapter, Tip-Adapter and end-to-end (E2E) fine-tuning using ZS-init-CE loss.

Dataset	Eval	FLYP				FLYP-CE				ZS-init-CE				LP-init-CE (LP-T)				ZS-init-L2SP				LP-init-L2SP				ZS-init-LDIFS (Ours)				LP-init-LDIFS (Ours)										
		$\Delta_{ZS}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{ZS}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{ZS}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{ZS}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{ZS}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{ZS}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{ZS}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{ZS}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )	$\Delta_{LP}$ ( $\uparrow$ )							
Car	CFAR10	-5.05	-0.97	80.62	-3.81	1.96	81.52	-2.39	1.86	81.17	-1.32	1.86	81.17	-1.32	1.86	81.17	-1.32	1.86	81.17	-1.32	1.86	81.17	-1.32	1.86	81.17	-1.32	1.86	81.17	-1.32	1.86	81.17	-1.32	1.86	81.17	-1.32	1.86	81.17			
	CFAR100	-6.91	-0.76	78.89	-7.41	-2.74	76.93	-8.62	-3.58	76.04	-6.84	-2.12	77.51	-3.4	-2.28	77.34	<b>0.23</b>	<b>1.17</b>	<b>79.81</b>	-1.6	-0.37	79.26	-1.76	0.07	79.71	-1.6	-0.37	79.26	-1.76	0.07	79.71	-1.6	-0.37	79.26	-1.76	0.07	79.71			
	EuroSAT	-16.94	-1.26	94.19	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14			
	SVHN	-16.94	-1.26	94.19	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14	-14.36	-1.44	94.14			
	GTSRB	-3.62	0.55	87.09	-4.28	-1.59	85.24	-1.1	-0.66	85.96	-3.58	-1.11	85.62	2.19	-0.4	86.21	<b>1.85</b>	<b>0.64</b>	<b>87.21</b>	<b>2.30</b>	-0.02	86.56	0.74	0.44	87.18	2.30	-0.02	86.56	0.74	0.44	87.18	2.30	-0.02	86.56	0.74	0.44	87.18	2.30	-0.02	86.56
	MNIST	4.28	-0.19	98.56	2.26	-0.52	98.11	<b>4.20</b>	-0.43	98.24	2.21	-0.14	98.48	1.46	-0.34	98.33	2.45	-0.07	98.59	2.33	0.01	98.71	2.15	<b>0.06</b>	<b>98.62</b>	2.15	0.06	98.62	2.15	0.06	98.62	2.15	0.06	98.62	2.15	0.06	98.62	2.15	0.06	98.62
	RESCIS	-6.19	-1.21	90.67	-7.28	-1.57	90.33	-10.7	-2.13	89.75	-4.81	-0.87	91.0	-4.88	-1.4	90.48	-0.21	-0.25	91.62	-1.13	-0.48	91.4	-2.7	-0.13	91.75	-1.13	-0.48	91.4	-2.7	-0.13	91.75	-1.13	-0.48	91.4	-2.7	-0.13	91.75			
	SVHN	0.33	1.14	66.68	-5.4	-3.38	62.15	-1.63	-3.91	61.87	-0.7	-2.09	63.0	<b>2.00</b>	-2.01	63.51	1.54	0.78	66.08	-0.05	-1.13	64.35	1.03	0.53	66.01	1.03	0.53	66.01	1.03	0.53	66.01	1.03	0.53	66.01	1.03	0.53	66.01			
	Can	9.2	<b>6.26</b>	<b>86.06</b>	5.73	3.95	84.77	18.24	2.67	83.48	-5.62	-4.14	84.95	9.44	1.83	82.64	-3.31	3.06	83.87	<b>20.52</b>	4.71	<b>81.52</b>	-3.15	4.45	85.26	4.71	81.52	-3.15	4.45	85.26	4.71	81.52	-3.15	4.45	85.26					
	Can	-58.0	-14.58	66.24	-45.32	-7.04	73.75	-46.08	-5.68	75.12	-36.86	-8.27	72.53	-15.81	0.52	81.32	-4.19	-0.55	80.26	-13.78	<b>0.53</b>	<b>81.33</b>	-3.95	-1.31	79.51	-13.78	<b>0.53</b>	<b>81.33</b>	-3.95	-1.31	79.51	-13.78	<b>0.53</b>	<b>81.33</b>	-3.95	-1.31	79.51			
CFAR10	-19.22	-0.97	80.62	-22.62	-1.96	81.52	-23.9	1.86	81.17	-13.2	1.86	81.17	-13.2	1.86	81.17	-13.2	1.86	81.17	-13.2	1.86	81.17	-13.2	1.86	81.17	-13.2	1.86	81.17	-13.2	1.86	81.17	-13.2	1.86	81.17	-13.2	1.86	81.17				
CFAR100	-14.31	-7.93	64.15	-9.52	-4.52	67.71	-9.73	-6.06	66.01	-9.52	-4.36	67.77	-1.86	<b>1.38</b>	<b>74.46</b>	-2.77	-0.74	71.33	-4.41	-1.14	70.64	-0.60	0.05	72.18	-4.41	-1.14	70.64	-0.60	0.05	72.18	-4.41	-1.14	70.64	-0.60	0.05	72.18				
GTSRB	-2.72	1.24	85.3	-0.8	0.35	87.05	-1.6	-0.66	86.25	-0.33	1.03	87.13	1.63	1.15	87.82	-0.37	1.58	88.23	<b>5.81</b>	2.54	88.23	<b>5.81</b>	2.54	88.23																
MNIST	-0.85	-0.43	98.19	-1.14	-0.54	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38				
RESCIS	-6.32	-4.83	87.05	-10.42	-0.77	85.88	-12.21	-2.41	89.37	-16.32	<b>0.22</b>	<b>89.95</b>	10.29	-0.25	95.62	-1.49	-0.31	91.78	-0.5	-0.17	92.05	1.03	0.13	92.0	-0.5	-0.17	92.05	1.03	0.13	92.0	-0.5	-0.17	92.05							
SVHN	-3.49	1.36	66.84	-7.19	2.9	68.33	-5.45	0.4	65.98	0.62	4.34	69.95	7.15	3.71	69.22	7.08	3.59	69.13	6.17	3.97	69.34	6.13	4.92	70.4	6.17	3.97	69.34	6.13	4.92	70.4	6.17	3.97	69.34							
Can	7.87	2.79	97.71	7.8	2.65	97.51	<b>8.26</b>	<b>2.83</b>	<b>97.78</b>	7.5	2.79	97.71	8.05	2.8	97.7	6.38	2.74	97.66	7.72	2.46	97.36	6.39	2.31	97.24	6.39	2.31	97.24	6.39	2.31	97.24	6.39	2.31	97.24	6.39	2.31	97.24				
Can	-27.76	-3.15	77.66	-12.57	-3.36	77.44	-21.22	-3.58	77.23	-12.42	-1.38	76.42	-6.74	-0.12	80.69	-1.32	0.15	<b>80.96</b>	-12.29	-0.72	80.08	-6.32	-1.07	79.73	-12.29	-0.72	80.08	-6.32	-1.07	79.73	-12.29	-0.72	80.08							
CFAR10	-17.07	1.05	95.95	0.22	1.33	96.25	-0.44	1.5	96.41	1.06	1.52	96.42	2.65	<b>1.96</b>	<b>96.87</b>	<b>3.24</b>	1.62	96.23	0.33	1.34	96.25	0.44	0.88	95.79	0.33	1.34	96.25	0.44	0.88	95.79	0.33	1.34	96.25							
CFAR100	-4.46	-0.17	96.26	-4.83	-2.52	95.54	-9.36	-1.11	95.97	-6.54	-3.1	98.83	-1.52	-0.74	71.33	-2.39	-1.14	<b>72.23</b>	-4.41	-1.14	70.64	-1.97	-1.07	70.11	-4.41	-1.14	70.64	-1.97	-1.07	70.11	-4.41	-1.14	70.64							
GTSRB	-23.19	-0.06	95.48	-12.06	-0.26	95.28	-16.59	-0.48	95.06	-7.11	-0.17	95.37	-0.46	0.19	95.72	-7.13	0.46	96.0	-1.33	0.26	95.8	-2.19	<b>0.56</b>	<b>96.00</b>	-1.33	0.26	95.8	-2.19	<b>0.56</b>	<b>96.00</b>	-1.33	0.26	95.8							
MNIST	-8.99	-0.35	98.36	-7.26	-0.15	98.31	-1.37	-0.3	98.36	-6.36	0.11	98.79	3.09	0.08	98.52	2.85	0.04	98.75	8.1	-0.07	98.68	0.25	0.2	<b>98.85</b>	8.1	-0.07	98.68	0.25	0.2	<b>98.85</b>										
RESCIS	-20.65	-3.24	88.33	-16.31	-3.56	88.32	-17.65	-4.3	88.57	-14.14	-0.32	90.34	0.55	-1.4	90.59	-2.81	-0.57	91.3	-7.68	-0.31	90.78	-6.39	<b>4.44</b>	<b>91.43</b>	-7.68	-0.31	90.78	-6.39	<b>4.44</b>	<b>91.43</b>										
SVHN	2.46	3.06	68.59	6.05	5.04	70.58	2.18	2.57	68.03	5.7	4.64	70.18	7.89	3.3	68.84	6.38	3.7	69.12	6.02	<b>5.32</b>	<b>67.07</b>	8.24	4.68	70.1	6.02	<b>5.32</b>	<b>67.07</b>	8.24	4.68	70.1										
CFAR100	2.44	9.34	88.98	23.16	9.12	88.77	<b>24.63</b>	8.97	88.6	17.26	8.78	88.41	23.33	8.2	87.84	13.77	7.31	86.94	24.28	<b>8.31</b>	<b>87.64</b>	13.77	9.35	88.99	24.28	<b>8.31</b>	<b>87.64</b>	13.77	9.35	88.99										
Can	-18.09	-3.36	77.44	-10.19	-2.87	77.94	-9.85	-2.54	78.25	-5.97	-2.43	78.36	-11.63	-4.15	76.66	-0.06	0.21	81.01	-1.94	<b>0.77</b>	<b>81.58</b>	-0.2	0.36	81.17	-1.94	<b>0.77</b>	<b>81.58</b>	-0.2	0.36	81.17										
CFAR10	-36.84	-10.03	69.61	-18.82	-7.08	72.57	-20.99	-8.23	71.41	-8.12	-3.3	70.34	-17.79	-8.45	71.18	<b>2.39</b>	0.8	80.45	-3.41	-0.72	79.10	1.61	<b>0.84</b>	<b>80.40</b>	-3.41	-0.72	79.10	1.61	<b>0.84</b>	<b>80.40</b>										
CFAR100	-22.07	-0.57	94.96	-20.66	-1.14	94.01	-12.56	-0.41	95.03	-11.02	-0.52	95.05	-22.26	-1.33	94.2	-0.83	94.6	80.09	-7.19	0.43	95.96	2.76	0.1	95.14	-7.19	0.43	95.96	2.76	0.1	95.14										
GTSRB	-18.13	-7.02	79.69	-7.89	-3.52	81.15	-13.89	-2.21	84.51	-4.41	-1.41	85.01	-5.76	-2.62	83.97	0.34	-0.39	86.18	2.11	0.69	87.32	1.43	1.3	<b>87.83</b>	2.11	0.69	87.32	1.43	1.3	<b>87.83</b>										
MNIST	-0.85	-0.43	98.19	-1.14	-0.54	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38	-1.21	-0.61	98.38							
RESCIS	-34.51	-5.0	86.87	-11.89	-2.52	89.37	-13.37	-2.65	89.22	-10.52	-1.48	90.4	-19.27	-4.48																										