A Closer Look at the Few-Shot Adaptation of Large Vision-Language Models

Julio Silva-Rodríguez✉ Sina Hajimiri Ismail Ben Ayed Jose Dolz

ÉTS Montreal
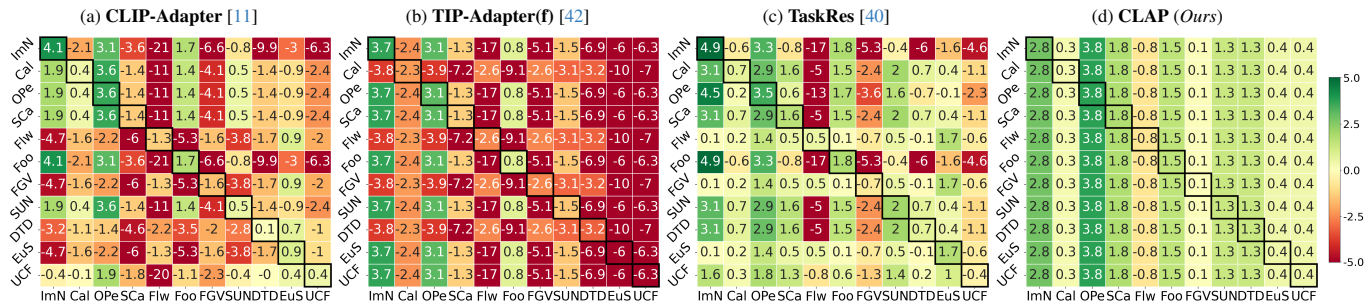
✉ julio-jose.silva-rodriguez@etsmtl.ca

Figure 1. Pitfalls of few-shot adapters due to the absence of a *model selection* strategy. The cross-shift model selection matrices (i, j) depict the relative improvement w.r.t. a zero-shot initialized Linear Probing when using the optimal hyperparameters for the dataset i (*rows*), for adapting in another task j (*columns*), for each SoTA method (*first three plots*) and our approach (*last plot*).

Abstract

Efficient transfer learning (ETL) is receiving increasing attention to adapt large pre-trained language-vision models on downstream tasks with a few labeled samples. While significant progress has been made, we reveal that state-of-the-art ETL approaches exhibit strong performance only in narrowly-defined experimental setups, and with a careful adjustment of hyperparameters based on a large corpus of labeled samples. In particular, we make two interesting, and surprising empirical observations. First, to outperform a simple Linear Probing baseline, these methods require to optimize their hyper-parameters on each target task. And second, they typically underperform –sometimes dramatically– standard zero-shot predictions in the presence of distributional drifts. Motivated by the unrealistic assumptions made in the existing literature, i.e., access to a large validation set and case-specific grid-search for optimal hyperparameters, we propose a novel approach that meets the requirements of real-world scenarios. More concretely, we introduce a CLass-Adaptive linear Probe (CLAP) objective, whose balancing term is optimized via an adaptation of the general Augmented Lagrangian method tailored to this context. We comprehensively evaluate CLAP on a broad span of datasets and scenarios, demonstrating that it consistently outperforms SoTA approaches, while yet being a much more efficient alternative. Code available at

https://github.com/jusiro/CLAP .

1. Introduction

Large vision-language models (VLMs), such as CLIP [30], are reshaping the research landscape with their unprecedented performance. These models undergo training on an extensive dataset consisting of hundreds of millions of image-text pairs, which are leveraged via contrastive learning [30]. Once trained, VLMs offer a remarkable zero-shot performance on a wide span of visual recognition problems thanks to the rich learned representations [27, 30]. Nevertheless, the extensive hardware and data-driven resources that such training demands [3] suggest that these models can only be trained on singular occasions. Furthermore, the large scale of these networks poses important challenges when it comes to adjusting their parameters on small downstream tasks that involve only a few labeled samples, making the full fine-tuning of the entire model impractical.

An emerging alternative to alleviate this issue consists in fine-tuning VLMs by adding a small set of learnable parameters, whose values are optimized during the adaptation step [11, 19, 42, 45, 46]. These tunable weights can be introduced in the input space as visual [19] or text prompts [45, 46], or added in the form of adapters across the network [11, 40, 42]. While both families of approaches fit within the Efficient Transfer Learning (ETL) literature,

prompt learning still requires backpropagating the gradients through the entire network. Thus, besides introducing a burden on resource reuse, these methods preclude *black-box* adaptation, introducing a potential concern about leaking the source data, which is paramount in privacy-oriented applications. In contrast, strategies based on adapters only need gradients on the extra set of parameters, typically in the last layer, avoiding costly fine-tuning processes and data leakage, yet yielding state-of-the-art performance [24, 40].

Despite the progress observed in adapter-based methods for fine-tuning VLMs under the few-shot learning paradigm, improving the performance on the target task while preserving their generalization capabilities remains still a challenge [46]. We argue that this is likely due to the severe overfitting to the support set samples employed during few-shot adaptation, which significantly deviates the updated class prototypes from the zero-shot prototypes initially provided by the pre-trained model. In fact, popular adapter-based ETL strategies, such as CLIP-Adapter [11] and TIP-Adapter [42], carefully adjust the model-specific hyperparameters, in conjunction with other key hyperparameters related to the learning scheduler, to control the trade-off between initial zero-shot inference and the integration of new information from the support set. Furthermore, recent evidence [24] suggests that these works apparently use the large-scale test set to adjust their hyperparameters.

A significant limitation becomes evident in that these hyperparameters, when optimized for one specific task, do not exhibit strong generalizability to other tasks, as illustrated in Fig. 1. Indeed, state-of-the-art (SoTA) methods **struggle to find a homogeneous configuration that outperforms a simple well-initialized Linear Probing (LP) adaptation**. Notably, in a realistic adaptation scenario (Fig. 1), we can observe dramatic performance degradations, up to 21%, compared to this simple baseline. These practices virtually bias the model selection process, as assuming access to a significantly larger set of labeled samples, and adjusting the model hyperparameters in a case-specific manner, is not only *unrealistic* but also *impractical* (grid-search must be done for each case). Thus, we argue that if an ETL method's model selection strategy is not solely based on the support samples, the method is *incomplete*, and impractical for real-world few-shot adaptation problems.

In this work, we seek to redirect the efforts on few-shot ETL to a more strict, but realistic scenario, in which only the support samples are accessible during training. The absence of an evaluation subset urges novel adapters to include a model selection strategy, robust across a large spectrum of tasks. Interestingly, we empirically observed that a carefully designed Linear Probing (ZS-LP), whose weights are initialized with the zero-shot prototypes from CLIP, is a strong baseline that outperforms more convoluted ETL solutions. To further improve the baseline ZS-LP and optimize the trade-off between initial zero-shot representations and updated class prototypes on novel tasks, we propose penalizing large deviations from the original zero-shot prototypes during adaptation. The resulting learning objective, however, presents two major issues. First, the penalty included to control the deviation between original and updated prototypes is a scalar value, uniform across all classes, which can detrimentally affect the model's performance in the presence of harder-to-learn classes. Second, the penalty balancing weight must be set using a validation set, which juxtaposes with our *validation-free* scenario. To address these limitations, we propose CLass-Adaptive linear Probe (CLAP), which is based on an Augmented Lagrangian Multiplier approach. We can summarize our contributions as:

- We empirically observe that SoTA few-shot ETL adapters require careful adjustment of a set of key hyperparameters for each task, which is unrealistic and impractical in real-world settings. Surprisingly, if a fixed configuration is adopted across tasks, these methods are likely to substantially underperform a simple Linear Probing strategy initialized with the zero-shot prototypes from CLIP.

- We propose a principled solution to tackle the trade-off between original and updated class prototypes in Linear Probing, which integrates a penalty term to penalize large deviations from zero-shot prototypes. To address the underlying challenges from the resulting constrained optimization problem, we present a modified Augmented Lagrangian Multiplier (ALM) method. This alleviates the need of having to fine-tune the penalty balancing weight, which is learned in the outer iteration of the optimization process. In order to adapt ALM to the presented scenario, two critical choices were made: *i)* Leveraging class prototypes, as well as data augmentation, motivate the use of class-wise multipliers, instead of sample and class-wise multipliers as in the original ALM; *ii)* In the presented scenario, there is no access to a validation set, and the only feedback available is from the support samples. Hence, we only perform one outer-step update, which can avoid potential overfitting on the support set.

- We provide extensive experiments to assess the performance of CLAP in the proposed scenario, including few-shot adaptation on 11 popular classification benchmarks, domain generalization, comparison to full fine-tuning methods, and ablation studies to validate our choices. As shown in Fig. 1 and in the experimental section, CLAP delivers consistent performance across different tasks with a homogeneous configuration, and largely outperforms SoTA ETL approaches in all scenarios.

2. Related work

Vision-language pre-trained models. The field of machine learning is in the midst of a paradigm shift with the emerging rise of vision-language models (VLMs). These

networks have gained increasing popularity, especially fueled by the significant improvements achieved in computer vision and natural language processing tasks [5, 18, 30, 41]. The prevailing learning paradigm consists of a dual stream of data, which separately encodes images and their text counterparts, leveraging contrastive learning at a large scale to bridge image and text representations in the latent space. Particularly, models such as CLIP [30] and ALIGN [18] have successfully mitigated the distribution discrepancy between text and images, and have shown tremendous zero-shot capabilities on visual recognition tasks, primarily in the context of classification.

Full fine-tuning. A body of work proposes fine-tuning the entire VLMs to adapt to a specific task [12, 22, 36]. This strategy, however, presents several drawbacks. Concretely, fine-tuning increases the complexity of the model being optimized, makes the optimization process more time-consuming compared to ETL methods, and requires access to the backbone weights, which does not allow a black-box adaptation. Furthermore, full fine-tuning methods typically tend to overfit when trained on small datasets, requiring a large corpus of labeled data for the target task, which may be impractical in many real-world scenarios.

Efficient transfer leaning attempts to address these issues by updating a small set of learnable parameters and leveraging a limited amount of annotated samples. Current ETL literature can be categorized into *Prompt Learning* [20, 38, 39, 45–47] and *Adapter-based* [11, 40, 42] approaches. *Prompt Learning* represents a recent advancement in the realm of natural language processing [23, 43], which has been recently adopted with success in VLMs. In these methods, only the text tokens provided to the model are optimized. Nevertheless, these techniques require long training steps due to backpropagating the gradient over the entire network, which juxtaposes with the spirit of *efficient* adaptation. Furthermore, black-box adaptation is also not possible in prompt learning. *Adapter-based* methods, in contrast, offer a much lighter alternative as only a small subset of parameters, typically at the latest layers, are adjusted. For example, CLIP-Adapter [11] integrates a two-layer MLP to modify the visual embedding generated by CLIP. In TIP-Adapter [42], the visual prototypes obtained from the few-shot support samples are leveraged to compute the similarity with the visual embedding of the test image, which is later used to modify the CLIP visual embedding.

3. Preliminaries

3.1. Contrastive vision-language pre-training

Large-scale VLMs, such as CLIP [30], are trained on large heterogeneous datasets, encouraging image and text representations to correlate in a joint embedding space. Formally, CLIP comprises a vision encoder, $f_\theta(\cdot)$, and a text encoder,

$f_\phi(\cdot)$, each aiming at learning a rich representation of their data points. These points are projected in an ℓ_2-normalized shared embedding space, yielding the corresponding visual $\mathbf{v}$ and text $\mathbf{t}$ embeddings. The whole network is optimized to maximize the similarity between the projected embeddings of paired images and texts, using a contrastive loss.

3.2. Transferability

Zero-shot inference. For a particular downstream image classification task, CLIP-based models are able to provide predictions based on the similarity between category prompts, *i.e.*, text descriptions of target classes, and testing images. Given a set of C categories, and an ensemble of N text prompts for each one, $\{\{T_{n,c}\}_{n=1}^N\}_{c=1}^C$, a common practice is to obtain a zero-shot prototype for each target category by computing the center of the ℓ_2-normalized text embeddings for each class, $\mathbf{t}_c = \frac{1}{N}\sum_{n=1}^N f_\phi(T_{n,c})$. Thus, for a given query image $\mathbf{x}$, the zero-shot prediction is obtained from the softmax cosine similarity between the vision embedding $\mathbf{v} = f_\theta(\mathbf{x})$, and category prototypes $\mathbf{t}_c$:

$$\hat{y}_c = \frac{\exp(\mathbf{v} \cdot \mathbf{t}_c^\top / \tau)}{\sum_{i=1}^C \exp(\mathbf{v} \cdot \mathbf{t}_i^\top / \tau)}, \tag{1}$$

where τ is a temperature parameter learned during the pre-training stage, and $\mathbf{v} \cdot \mathbf{t}^\top$ the dot product operator, which is equivalent to cosine similarity, as vectors are ℓ_2-normalized.

Few-shot learning. This scenario assumes access to limited supervisory information on the downstream tasks, in the form of a few examples for each target category, so-called shots. Formally, we denote a support set, $\mathcal{S} = \{(\mathbf{x}^{(m)}, \mathbf{y}^{(m)})\}_{m=1}^{M=K\times C}$, composed of K images for each target category, such that K takes a small value, *e.g.*, $K \in \{1, 2, 4, 8, 16\}$, and where $\mathbf{y} \in \{0, 1\}^C$ is the corresponding one-hot label for a given image $\mathbf{x}$. The objective is to adapt the pre-trained model using this limited support set.

3.3. Efficient transfer learning with adapters

In their general form, ETL methods based on adapters learn a set of transformations over the pre-trained features $(\mathbf{v}', \mathbf{t}' = f_\psi(\mathbf{v}, \mathbf{t}))$, parameterized by the so-called adapter ψ, which produces softmax scores for the new tasks following Eq. (1). The adapter ψ can be optimized by minimizing the popular cross-entropy (CE) loss, $\mathcal{H}(\mathbf{y}, \hat{\mathbf{y}}) = -\sum_{c=1}^C y_c \log \hat{y}_c$, over the support set samples:

$$\min_\psi \frac{1}{M} \sum_{m=1}^M \mathcal{H}(\mathbf{y}^{(m)}, \hat{\mathbf{y}}^{(m)}). \tag{2}$$

3.4. Pitfalls of existing few-shot ETL methods

Recent ETL methods tailored to VLMs focus on enhancing the supervision provided by the support samples with

priors learned by the VLMs at the task at hand. The pretrained model gathers robust knowledge and is able to align visual and textual concepts. Retaining this prior knowledge can therefore produce more robust adapters, able to generalize beyond the specific bias introduced in the few support samples, to more general concepts. In this context, the zero-shot prototypes from CLIP act as a proxy to initialize the learning procedure into a reliable region. For instance, CLIP-Adapter [11] maintains the zero-shot prototypes based inference as in Eq. (1), but includes a residual multi-layered perceptron to modify visual features, such as $\mathbf{v}' = \mathbf{v} + \alpha_\mathrm{r} f_\psi(\mathbf{v})$. TIP-Adapter [42] includes an additional complexity layer, by combining the similarity of the zero-shot prototypes with a weighted similarity to the support samples, $f_\psi(\cdot, \beta)$, controlled by the hyperparameter β, such that the predicted logits are $\mathbf{l}_c = \alpha_\mathrm{tipA} f_\psi(\mathbf{v}, \beta) + \mathbf{v} \cdot \mathbf{t}_c^\top / \tau$. Finally, TaskRes [40] learns a modification of the initial zero-shot prototypes, $\mathbf{w}_{TR}$, using the support samples. The divergence between the initial and final prototypes is controlled by a residual ratio: $\mathbf{t}' = \mathbf{t} + \alpha_\mathrm{TR} \mathbf{w}_{TR}$. Nevertheless, these methods lack a *model selection* strategy to set these hyperparameters (See Supp. Sec. A for details).

4. Proposed approach

4.1. Revisiting Linear Probing

The most straightforward approach used to adapt VLMs is Linear Probing [30], which refers to fitting a multiclass logistic regression linear classifier on top of the pre-trained features. Formally, the objective is to learn a set of classwise prototypes, $\mathbf{w}_c$, to provide softmax class scores for a given visual embedding $\mathbf{v}$:

$$\hat{y}_c = \frac{\exp(\mathbf{v} \cdot \mathbf{w}_c^\top / \tau)}{\sum_{i=1}^C \exp(\mathbf{v} \cdot \mathbf{w}_i^\top / \tau)}. \tag{3}$$

The $\mathbf{w_c}$ prototypes can be trained to minimize the cross-entropy loss on the support samples, as in Eq. (2), using standard SGD. Besides, a common practice in ETL is to regularize the trained weights [24, 30, 40] by minimizing its ℓ_2-norm with an additional term, weighted by an empirically-optimized non-negative balancing term λ_{wd}. Despite its limited performance shown for few-shot adaptation [11, 30], we believe that this requires further exploration, as LP is a lightweight adaptation strategy, especially convenient due to its convexity during optimization. In this work, we present an updated view of Linear Probing. First, the class weights are initialized using the CLIP zero-shot prototypes, as SoTA ETL methods do [11, 40, 42]. Second, we replace the weight decay in the loss function and explicitly perform an ℓ_2-normalization of the prototypes after each update, to exactly meet the pre-training scenario during adaptation, inspired by [12]. Similarly, cosine similarity is also scaled with CLIP's pre-trained temperature

τ. Last, we incorporate data augmentation, usually not included in LP. We refer to this updated Linear Probing version for vision-language models as ZS-LP[1]. Interestingly, ZS-LP serves as a strong baseline (see Tab. 1), which does not require adjusting specific hyperparameters per task.

4.2. Constrained Linear Probing

Albeit a well-initialized Linear Probing offers a strong baseline for efficient transfer learning, the updated prototypes might deviate from the initial regions offering strong generalization. This is especially the case in the few-shot setting, where the few provided support samples might be underrepresentative and contain specific biases that produce spurious correlations, hence harming the generalization after adaptation [34, 44]. Thus, to retain the strong basis provided by the VLM model, and avoid prototype degradation, we resort to a constrained formulation of the loss in Eq. (2).

Retaining prior knowledge. A direct form to avoid prototype degradation from zero-shot points is to constrain the cross-entropy minimization to enforce the resulting prototypes to remain close to the initial solution (*i.e.*, initial set of prototypes $\mathcal{T} = [\mathbf{t}_1, \ldots, \mathbf{t}_c]$). Specifically, this constrained optimization problem can be defined as follows:

$$\min_{\mathcal{W}} \quad \frac{1}{M} \sum_{m=1}^M \mathcal{H}(\mathbf{y}^{(m)}, \hat{\mathbf{y}}^{(m)}) \tag{4}$$
$$\text{s.t.} \quad \mathbf{w}_c = \mathbf{t}_c \quad \forall c \in \{1, \ldots, C\},$$

with $\mathcal{W} = [\mathbf{w}_1, ..., \mathbf{w}_C]$ the set of learnable class prototypes. We can approximate the minimum of the constrained problem in Eq. (4) by a penalty-based optimization approach, transforming the above formulation into an unconstrained problem, and using an ℓ_2-penalty between the class prototypes and the set of zero-shot anchors:

$$\min_{\mathcal{W}} \quad \sum_{m=1}^M \mathcal{H}(\mathbf{y}^{(m)}, \hat{\mathbf{y}}^{(m)}) + \lambda \sum_{m=1}^M \sum_{c=1}^C ||\mathbf{t}_c - \mathbf{w}_c^{(m)}||_2^2, \tag{5}$$

where $\lambda \in \mathbb{R}_+$ is a scalar weight controlling the contribution of the corresponding penalty. Note that $\mathbf{w}_c^{(m)}$ is the optimal class prototype for the support sample m that minimizes the left term. For clarity in the presentation, we have omitted the normalization by the cardinality of each set.

Sample and class-specific constraints. The associated constrained problem in Eq. (4) is approximated by an unconstrained formulation, which uses a single uniform penalty without considering individual data samples or

---

[1]Although the recent work in [24] explores some of these LP improvements, they still resort to a weight-decay regularization of the LP parameters, whose optimum relative weight is found in a few-shot validation set.

classes. Certainly, all samples and categories within a given dataset may indeed present different intrinsic learning challenges. Thus, the problem in Eq. (5) is not solved accurately. A better alternative would consist in integrating multiple penalty weights λ, one for each sample and class, producing a set of penalty weights $\boldsymbol{\Lambda} \in \mathbb{R}_+^{M \times C}$. The resulting optimization problem can then be defined as:

$$\min_{\mathcal{W}} \quad \sum_{m=1}^{M} \mathcal{H}(\mathbf{y}^{(m)}, \hat{\mathbf{y}}^{(m)}) + \sum_{m=1}^{M} \sum_{c=1}^{C} \boldsymbol{\Lambda}_{mc} \, ||\mathbf{t}_c - \mathbf{w}_c^{(m)}||_2^2.$$
(6)

Now, from an optimization standpoint, if we suppose that there exists an optimal set of class-prototypes $\mathcal{W}^*$ for the problem presented in Eq. (4), there also exists $\boldsymbol{\Lambda}^* \in \mathbb{R}_+^{M \times C}$ such that $(\mathcal{W}^*, \boldsymbol{\Lambda}^*)$ represents a saddle point of the Lagrangian associated to Eq. (4). In this scenario, $\boldsymbol{\Lambda}^*$ are the Lagrange multipliers of the presented problem, and is intuitive to consider $\boldsymbol{\Lambda} = \boldsymbol{\Lambda}^*$ as the best choice to solve Eq. (6).

Nevertheless, using the Lagrange multipliers $\boldsymbol{\Lambda}^*$ as the weights for the penalties in Eq. (6) may not be feasible in practice. In particular, a number of conventional strategies employed to train deep neural networks hinder straightforward minimization. First, the use of mini-batch gradient descent averages the updated prototypes for every single observation into a mean prototype per class, making a sample-wise constraint hard to achieve. Furthermore, performing data augmentation over the support samples may yield distinct penalty weights for the augmented versions, which could be harder or easier to classify than their original counterparts.

To alleviate the aforementioned challenges, we propose to relax the sample-wise penalties, which results in solving:

$$\min_{\mathcal{W}} \quad \sum_{m=1}^{M} \mathcal{H}(\mathbf{y}^{(m)}, \hat{\mathbf{y}}^{(m)}) + \sum_{c=1}^{C} \lambda_c \, ||\mathbf{t}_c - \mathbf{w}_c||_2^2, \quad (7)$$

where $\boldsymbol{\lambda} \in \mathbb{R}_+^C$ is a set of C class-wise penalty weights. While the problem complexity has been reduced by removing sample-wise penalty weights, we still need to choose C weights for the class-wise penalties. This poses a challenge in the optimization, particularly for datasets that contain a large number of categories, such as ImageNet [8] ($C = 1000$), where properly selecting the penalty weights $\boldsymbol{\lambda} \in \mathbb{R}_+^C$ can be a laborious process. Furthermore, choosing these values "by hand" juxtaposes with our goal of providing a *validation-free* solution for ETL.

4.3. Class Adaptive Constraint for Linear Probing

General Augmented Lagrangian. Augmented Lagrangian Multiplier (ALM) methods present an appealing alternative for learning the penalty weights. These popular methods

in optimization, which solve a constrained problem by the interplay of penalties and primal-dual steps, present well-known advantages [1, 32]. Formally, we can define a general constrained optimization problem as:

$$\min_{x} \quad g(x) \quad \text{s.t.} \quad h_i(x) \leq 0, \quad i = 1, \ldots, n \quad (8)$$

with $g : \mathbb{R}^d \to \mathbb{R}$ the *objective function* and $h_i : \mathbb{R}^d \to \mathbb{R}, i = 1, \ldots, n$ the *set of constraint functions*. This problem is generally tackled by solving a succession of $j \in \mathbb{N}$ unconstrained problems, each solved approximately w.r.t x:

$$\min_{x,\lambda} \quad \mathcal{L}^{(j)}(x) = g(x) + \sum_{i=1}^{n} P(h_i(x), \rho_i^{(j)}, \lambda_i^{(j)}), \quad (9)$$

with $P : \mathbb{R} \times \mathbb{R}_{++} \times \mathbb{R}_{++} \to \mathbb{R}$ a *penalty-Lagrangian function*, whose derivative w.r.t. its first variable $P'(z, \rho, \lambda) \equiv \frac{\partial}{\partial z} P(z, \rho, \lambda)$ exists, is positive and continuous for all $z \in \mathbb{R}$ and $(\rho, \lambda) \in (\mathbb{R}_{++})^2$. The set of axioms that any penalty function P must satisfy [2] are detailed in Supp. Sec. B. Furthermore, $\boldsymbol{\rho}^{(j)} = (\rho_i^{(j)})_{1 \leq i \leq n} \in \mathbb{R}_{++}^n$ and $\boldsymbol{\lambda}^{(j)} = (\lambda_i^{(j)})_{1 \leq i \leq n} \in \mathbb{R}_{++}^n$ denote the penalty parameters and multipliers associated to the penalty P at the iteration j.

The ALM can be split into two iterations: *outer* iterations (indexed by j), where the *penalty multipliers* $\boldsymbol{\lambda}$ and the *penalty parameters* $\boldsymbol{\rho}$ are updated, and the *inner* iterations, where $\mathcal{L}^{(j)}$ (Eq. (9)) is minimized using the previous solution as initialization. In particular, the penalty multipliers $\boldsymbol{\lambda}^{(j)}$ are updated to the derivative of P w.r.t. to the solution obtained during the last *inner* step:

$$\lambda_i^{(j+1)} = P'(h_i(x), \rho_i^{(j)}, \lambda_i^{(j)}). \quad (10)$$

By doing this, the penalty multipliers increase when the constraint is violated, and decrease otherwise. Thus, this strategy enables an *adaptive* and *learnable* way for determining the penalty weights.

Our solution. We propose to use an ALM approach to solve the problem in Eq. (7). In particular, we reformulate this problem integrating a penalty function P parameterized by $(\boldsymbol{\rho}, \boldsymbol{\lambda}) \in \mathbb{R}_{++}^C \times \mathbb{R}_{++}^C$, formally defined as:

$$\min_{\mathcal{W}, \boldsymbol{\lambda}} \quad \sum_{m=1}^{M} \mathcal{H}(\mathbf{y}^{(m)}, \hat{\mathbf{y}}^{(m)}) + \sum_{c=1}^{C} P(\mathbf{t}_c - \mathbf{w}_c, \rho_c, \lambda_c).$$
(11)

Following our realistic *validation-free* scenario, the only data from which we can obtain feedback during adaptation is the support set $\mathcal{S}$. Thus, the penalty multiplier for class c at epoch $j + 1$ can be defined as:

$$\lambda_c^{(j+1)} = \frac{1}{|\mathcal{S}|} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{S}} P'(\mathbf{t}_c - \mathbf{w}_c, \rho_c^{(j)}, \lambda_c^{(j)}). \quad (12)$$

As suggested by prior work [2, 25], we employ the PHR function as penalty P, defined as:

$$\text{PHR}(z, \rho, \lambda) = \begin{cases} \lambda z + \frac{1}{2}\rho z^2 & \text{if} \quad \lambda + \rho z \geq 0; \\ -\frac{\lambda^2}{2\rho} & \text{otherwise.} \end{cases} \quad (13)$$

Nevertheless, as we empirically found in our experiments (Supp. Sec. C.3), estimating Lagrange multipliers from the support samples might overfit the training data. As we do not have access to additional data points, we follow a simple strategy, consisting in performing only one iteration of the λ update. For a given target task, we rely on text embeddings as an anchor that offers a generalizable representation of concrete concepts along different visual domains. Thus, we consider the zero-shot prototypes $\mathbf{t}_c$ as the initial approximation of the problem in Eq. (12) (first *inner* step). Instead of initializing λ randomly, which might hamper the convergence, we compute the penalty weight for a given class as the average of the zero-shot softmax scores for all support samples belonging to that class, such that $\lambda_c^\star = \frac{1}{|\mathcal{B}_c^+|} \sum_{i \in \mathcal{B}_c^+} \hat{y}_c^{(i)}$, with $\mathcal{B}_c^+ = \{i | i \in M, y_c^{(i)} = 1\}$. Note that these values are obtained by replacing $\mathbf{w}_c$ with the solution found in the *inner* step ($\mathbf{t}_c$) in Eq. (3), which indeed satisfies the constraint $\mathbf{w}_c = \mathbf{t}_c$, resulting in a zero penalty. Taking now the derivative w.r.t. z of PHR, it is straightforward to see that the *learned* value of λ after one iteration is indeed $\lambda_c^\star$.

5. Experiments

5.1. Setup

Datasets: Few-shot adaptation. We follow prior ETL literature [11, 40, 42] and benchmark all the methods on 11 datasets: Imagenet [8], Caltech101 [10], OxfordPets [29], StanfordCars [21], Flowers102 [28], Food101 [4], FGVCAircraft [26], SUN397 [37], DTD [7], EuroSAT [15], and UCF101 [33]. These cover a diverse set of computer vision classification tasks, from general objects to actions or fine-grained categories in specialized applications. To train the few-shot adapters, we randomly retrieve K shots ($K \in \{1, 2, 4, 8, 16\}$) for each class. Last, for evaluation, we used the test sets provided in each dataset, with the same data splits as [40, 46]. **Domain generalization capabilities.** We further assess the model's *robustness* to domain shifts by following existing ETL works. We used ImageNet as a source domain for adaptation, and its variants as target tasks, which include: ImageNetV2 [31], ImageNet-Sketch [35], ImageNet-A [16], and ImageNet-R [17]. In this scenario, the model only sees a few labeled samples from the source domain, and target data are used exclusively for testing. In addition, we also employ this setting to motivate the use of efficient adapters *vs* fine-tuning the entire VLM [12, 22, 40].

Implementation details. All experiments are based on CLIP [30] pre-trained features, using different backbones: ResNet-50 [14] and ViT-B/16 [9] (results for other backbones in Supp. Sec. C.2). We resort to ResNet-50 as backbone in the ablation studies. For each downstream task we first extract all pre-trained features of the support shots and then run adaptation experiments over those. Data augmentation is applied during the feature extraction stage using random zoom, crops, and flips, following [40, 45]. The number of augmentations per support sample is set to 20. We used the same text prompts per dataset as in [40, 46]. Following our claim that using a validation set on few-shot adaptation is unrealistic, we trained ZS-LP and CLAP using the same configuration for all datasets, number of shots, and visual backbones. Concretely, we optimize the adapter for 300 epochs, using SGD optimizer with Momentum of 0.9. We use a relatively large initial learning rate of 0.1 to avoid underfitting on the support set, whose value decreases during training following a cosine decay scheduler. We ran all experiments with three different random seeds, and the results were averaged across runs.

Baselines and adaptation protocol. We selected adapter-based methods as our main competitors based on the similarity to our approach, including Clip-Adapter [11], TIP-Adapter [42], TaskRes [40], and Cross-Modal [24]. It is important to highlight that prior works [11, 40, 42] apparently leverage either the extensive test set, or an independent additional validation subset, to adjust important hyperparameters for few-shot adaptation, such as the learning rate, training epochs, and particular parameters that control each method [24]. Nevertheless, as we exposed in Fig. 1, their performance dramatically decreases when the set of hyperparameters is not adjusted for the testing scenario. To adhere to real-world requirements, we define a strict few-shot adaptation protocol, in which no validation or test samples are available to find the best case-specific configuration for each method, and hyperparameters remain fixed across tasks (details in Supp. Sec. A.4).

5.2. Results

Efficient transfer learning. We report in Tab. 1 the performance of adapter-based approaches averaged across 11 datasets, in the more realistic and practical *validation-free* experimental setting. Furthermore, for prompt-learning-based approaches, we include the results reported in prior literature, for a more comprehensive comparison. From these values, we can make interesting observations. First, a well-initialized Linear Probe, *i.e.*, using the CLIP zero-shot weights, does not show the performance degradation discussed in prior works, and it is indeed a competitive alternative to SoTA approaches. Second, and more surprisingly, more complex approaches such as CLIP-Adapter, or TIP-Adapter, show a significant decline in performance com-

pared to their original results when no validation set is available for model selection. Interestingly, TaskRes(e), which is some sort of two-stage zero-shot initialization Linear Probing with an updated text projection, also offers robust performance. Nevertheless, the absence of a detailed explanation of how the enhanced version is obtained in the original work hampers fair comparisons. Third, constraining the weights update to remain close to the zero-shot knowledge (CLAP) shows consistent improvements across different shots, especially in the very low data regime. This suggests that retaining the previous base knowledge from VLMs is important to avoid diverging because of unrepresentative shots during adaptation. Results per dataset are detailed in Supp. Fig. 8 and Supp. Tab. 9.

Table 1. **Comparison to state-of-the-art methods** for few-shot adaptation of CLIP-based models, using ResNet-50 backbone. ETL methods are trained under the same protocol, *i.e.*, absence of a validation set and using a fixed configuration across datasets, and results are averaged across 11 datasets. Prompt-learning methods results are directly extracted from [6, 13]. Best results in bold.

Method	K=1	K=2	K=4	K=8	K=16
Prompt-learning methods					
CoOp IJCV'22[46]	59.56	61.78	66.47	69.85	73.33
ProGrad ICCV'23[13]	62.61	64.90	68.45	71.41	74.28
PLOT ICLR'23[6]	62.59	65.23	68.60	71.23	73.94
Efficient transfer learning - a.k.a Adapters					
Zero-Shot ICML'21[30]	57.71	57.71	57.71	57.71	57.71
Rand. Init LP ICML'21[30]	30.42	41.86	51.69	60.84	67.54
CLIP-Adapter IJCV'23[11]	58.43	62.46	66.18	69.87	73.35
TIP-Adapter ECCV'22[42]	58.86	60.33	61.49	63.15	64.61
TIP-Adapter(f) ECCV'22[42]	60.29	62.26	65.32	68.35	71.40
CrossModal-LP CVPR'23[24]	62.24	64.48	66.67	70.36	73.65
TaskRes(e) CVPR'23[40]	61.44	65.26	68.35	71.66	74.42
ZS-LP	61.28	64.88	67.98	71.43	74.37
CLAP	**62.79**	**66.07**	**69.13**	**72.08**	**74.57**

Domain generalization. If adaptation is not carefully conducted, the resulting model might distort the pre-trained knowledge and underperform when new data with domain drifts is involved [22], even below the zero-shot (no adaptation) performance. Thus, evaluating the robustness of novel adapters under this scenario of domain generalization is of special interest. To do so, adapters are optimized on ImageNet using 16 shots per class, and directly evaluated on ImageNet variants. In this setting, we also assume the absence of a validation dataset, and hence all adapters are trained until convergence, using the same configuration across backbones. A summary of the results is reported in Tab. 2, while specific numbers across datasets and additional backbones are included in Supp. Tab. 10. From these experiments, we make two striking observations. First, ZS-LP is a strong baseline compared to other more complex adapters on the source domain. Even more remarkably, prior SoTA adapters, such as CLIP-Adapter or TIP-Adapter, fail to generalize to unseen domains. In-

deed, when using recent vision transformers, which are overtaking convolutional neural networks, **none of existing adapters-based approaches outperform standard zero-shot prediction in the presence of distributional drifts.** In contrast, CLAP yields the best in-distribution performance and also shows consistent improvements under domain shifts across all backbones.

Table 2. **Robustness to domain shifts.** Adapters are adjusted on ImageNet and evaluated at out-of-distribution generalization on 4 ImageNet shifts. Bold indicates best performance. Differences with respect to no adaptation (*a.k.a* zero-shot) are highlighted.

	Method	Source (Imagenet)	Target (Average)
ResNet-50	Zero-Shot ICML'21[30]	60.35	40.61
	Rand. Init LP ICML'21[30]	$52.24_{(-8.11)}\downarrow$	$24.61_{(-16.00)}\downarrow$
	CLIP-Adapter IJCV'23[11]	$59.02_{(-1.33)}\downarrow$	$31.21_{(-9.40)}\downarrow$
	TIP-Adapter ECCV'22[42]	$57.81_{(-2.54)}\downarrow$	$40.69_{(+0.08)}\uparrow$
	TIP-Adapter(f) ECCV'22[42]	$62.27_{(+1.92)}\uparrow$	$41.36_{(+0.75)}\uparrow$
	TaskRes(e) CVPR'23[40]	$60.85_{(+0.50)}\uparrow$	$41.28_{(+0.67)}\uparrow$
	ZS-LP	$61.00_{(+0.65)}\uparrow$	$36.58_{(-4.03)}\downarrow$
	CLAP	$\mathbf{65.02}_{(+4.67)}\uparrow$	$\mathbf{42.91}_{(+2.30)}\uparrow$
ViT-B/16	Zero-Shot ICML'21[30]	68.71	57.17
	Rand. Init LP ICML'21[30]	$62.95_{(-5.76)}\downarrow$	$40.41_{(-16.76)}\downarrow$
	CLIP-Adapter IJCV'23[11]	$68.46_{(-0.25)}\downarrow$	$50.72_{(-6.45)}\downarrow$
	TIP-Adapter ECCV'22[42]	$53.81_{(-14.90)}\downarrow$	$41.55_{(-15.62)}\downarrow$
	TIP-Adapter(f) ECCV'22[42]	$51.71_{(-17.00)}\downarrow$	$35.58_{(-21.6)}\downarrow$
	TaskRes(e) CVPR'23[40]	$70.84_{(+2.13)}\uparrow$	$55.35_{(-1.82)}\downarrow$
	ZS-LP	$69.73_{(+1.02)}\uparrow$	$53.65_{(-3.52)}\downarrow$
	CLAP	$\mathbf{73.38}_{(+4.67)}\uparrow$	$\mathbf{60.04}_{(+2.87)}\uparrow$

Table 3. **Fine-tuning (FT) vs. efficient transfer learning (ETL).** A benchmark for the low data regime, *i.e.*, 8 shots for each class. For the sake of fairness, FT methods (above the dashed line) are trained with 4 shots and early-stopped using a validation set containing 4 shots. On the other hand, ETL methods (below the dashed line) are trained using 8 shots and rely solely on the support set. All methods use ViT-B/16 as CLIP backbone.

Method	Source Imagenet	Target -V2	-Sketch	-A	-R	Avg.
Fine-tuning (FT)	69.88	62.44	47.07	47.52	76.08	58.28
LP-FT ICLR'23 [22]	71.29	64.04	48.50	49.49	77.63	59.92
WiSE CVPR'22 [36]	71.17	63.81	49.38	50.59	78.56	60.59
FLYP CVPR'23 [12]	71.51	64.59	49.50	51.32	78.52	**60.98**
Zero-Shot	68.71	60.76	46.18	47.76	73.98	57.17
Rand. Init LP	56.58	47.17	25.82	27.03	47.05	36.77
ZS-LP	68.49	60.07	42.77	42.39	71.73	54.24
CLAP	**71.75**	64.06	47.66	48.40	76.70	59.21

*Specific numbers for FT, LP-FT, WiSE-FT, and FLYP are retrieved from [12].

Is it worth optimizing the entire model? We now compare CLAP to end-to-end full fine-tuning (FT) approaches: LP-FT [22], WiSE-FT [36], and FLYP [12]. The former two methods require a validation set for early stopping, and the latter two use it for both early stopping and tuning the *mixing coefficient* hyperparameter α. Therefore, for a K-shot problem, these methods actually require $2K$ shots for each class, K for training, and K for validation. As the balancing penalty term in CLAP is optimized with the support set, and does not require a validation set, a fair comparison would be to evaluate the K-shot performance of fine-tuning methods against our method's $2K$-shot results. Thus, Tab. 3 in-

cludes the performance of all the models when 8 labeled images are available for each class overall. Analyzing the results, we can conclude that in the low data regime, full fine-tuning is not necessarily superior to ETL when compared properly. More specifically, our approach outperforms fine-tuning methods in in-distribution performance and performs reasonably well on OOD datasets, while having a fraction of the optimizable parameters of fine-tuning methods.

5.3. Ablation experiments

On the need for model selection strategies. Relevant methods (*e.g.*, CLIP-Adapter [11], TIP-Adapter [42], or TaskRes [40]) include different hyperparameters that directly control their performance. Nevertheless, these methods are *incomplete*, since they do not include any strategy for adjusting these parameters, typically referred to as *model selection*. In contrast, and as previously stressed, there is evidence that these works use a large evaluation subset to adapt their settings to each scenario [24]. To investigate this observation, we evaluate these methods in cross-dataset model selection experiments. The best hyperparameters values for a task (*i.e.*, dataset), which are found in an Oracle scenario using the entire test subset, are used during adaptation to another dataset. The matrices showing the relative improvements over a zero-shot initialized Linear Probing (ZS-LP) are depicted in Fig. 1. These results show empirically that the hyperparameters values are highly task-dependent, and **that SoTA methods must adjust their hyperparameters on the target task to outperform this simple baseline,** which is *unrealistic* in practice. In contrast, the proposed CLAP is more robust, showing consistent results across all datasets, even in the worst degradation case, as it does not require particular modifications per task.

Table 4. **Improving Linear Probing.** Using as baseline the proposed ZS-LP configuration detailed in Sec. 4.1, we isolate the effect of removing different parts of the model, while keeping the rest static. Results are averaged across 11 datasets.

Method	$K=1$	$K=2$	$K=4$
ZS-LP	61.28	64.88	67.98
w/o DA	$57.72_{(-3.5)}\downarrow$	$61.94_{(-2.9)}\downarrow$	$65.41_{(-2.5)}\downarrow$
w/o Temp. Scaling (τ)	$58.33_{(-2.9)}\downarrow$	$59.85_{(-5.0)}\downarrow$	$59.91_{(-8.0)}\downarrow$
w/o L^2-norm	$48.67_{(-12.6)}\downarrow$	$55.29_{(-9.6)}\downarrow$	$61.16_{(-6.8)}\downarrow$
Rand. Init.	$30.42_{(-30.8)}\downarrow$	$41.86_{(-23.0)}\downarrow$	$51.69_{(-16.2)}\downarrow$

Details in Linear Probing matter. As described earlier in Sec. 4.1, LP has been discouraged in the prior literature due to its limited performance in few-shot adaptation [11, 30]. Nevertheless, we argue that this behavior stems from the original way in which LP was introduced in [30], inspired by prior self-supervised learning methods. Indeed, a strategy tailored to contrastive VLMs alleviates the performance drop of LP observed in prior works. In particular, using zero-shot initialization, the same temperature scaling as

pre-training, and explicit ℓ_2-normalization of the class prototypes, considerably improves the generalization of few-shot adaptation (Tab. 4). This aligns with relevant literature on other topics such as FT [12], which suggests that the adaptation conditions should match the pre-training setting. Also, including other heuristics such as data augmentation (DA), usually omitted in LP [40, 42], is of special relevance.

Using a few-shot validation set. Cross-Modal adapter [24] uses a validation set composed of $(min(K, 4))$ samples to adjust the experimental setting and early stopping. Even though this setting is more appropriate, it still requires an additional number of shots for model selection. Nevertheless, for the sake of fairness, the performance comparison to methods that do not require a validation set should be carried out by training the latter methods using $K + min(K, 4)$ shots. When this fair benchmark is established (see Tab. 5), simple ZS-LP excels again as a strong baseline, outperforming more complex methods on the low-shot regime. Only when using a large number of shots ($K > 8$) partial fine-tuning and ETL methods marginally benefit from validation samples. However, model selection using a validation set increases the computational workload and processing times during adaptation due to its grid search nature.

Table 5. **Using a few-shot validation set.** Results for priors works on this setting are obtained from [24]. Average across 11 datasets.

Method	$K=1$	$K=2$	$K=4$	$K=8$	$K=16$
Protocol in [24]: K-shots for train + $min(K, 4)$ for validation					
TIP-Adapter [42]	63.3	65.9	69.0	72.2	75.1
CrossModal LP [24]	64.1	67.0	70.3	73.0	76.0
CrossModal Adapter [24]	64.4	67.6	70.8	73.4	75.9
CrossModal PartialFT [24]	64.7	67.2	70.5	**73.6**	**77.1**
Ours: using $K + min(K, 4)$ shots for training					
ZS-LP	64.9	68.0	71.4	73.1	75.0
CLAP	**66.1**	**69.1**	**72.1**	73.5	75.1

6. Limitations

In this work, we have introduced a CLass-Adaptive linear Probe (CLAP) objective, based on an adaptation of the general Augmented Lagrangian method, for efficient adaptation of large vision-language models in realistic scenarios. Despite its superiority, our empirical validation suggests that the benefits of our approach diminish as the number of shots increases, indicating that other strategies might be privileged if the number of adaptation samples is large.

Acknowledgments

References

[1] Dimitri P. Bertsekas. *Constrained Optimization and Lagrange Multiplier Methods (Optimization and Neural Computation Series)*. Athena Scientific, 1 edition, 1996. 5

[2] Ernesto G Birgin, Romulo A Castillo, and José Mario Martínez. Numerical comparison of augmented lagrangian algorithms for nonconvex problems. *Computational Optimization and Applications*, 31(1):31–55, 2005. 5, 6, 15

[3] Rishi Bommasani et al. On the opportunities and risks of foundation models. *ArXiv*, 2021. 1

[4] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – mining discriminative components with random forests. In *European Conference on Computer Vision (ECCV)*, 2014. 6, 16

[5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:1877–1901, 2020. 3

[6] Guangyi Chen, Weiran Yao, Xiangchen Song, Xinyue Li, Yongming Rao, and Kun Zhang. Prompt learning with optimal transport for vision-language models. In *International Conference on Learning Representations (ICLR)*, 2023. 7, 19

[7] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3606–3613, 2014. 6, 16

[8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009. 5, 6, 16

[9] 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. *International Conference on Learning Representations (ICLR)*, 2021. 6

[10] Li Fei-Fei, R. Fergus, and P. Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 178–178, 2004. 6, 16

[11] 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. *International Journal of Computer Vision (IJCV)*, 2023. 1, 2, 3, 4, 6, 7, 8, 11, 12, 13, 14, 15, 16, 19, 21

[12] Sachin Goyal, Ananya Kumar, Sankalp Garg, Zico Kolter, and Aditi Raghunathan. Finetune like you pretrain: Improved finetuning of zero-shot vision models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19338–19347, 2023. 3, 4, 6, 7, 8, 16, 17

[13] Changsheng Xu Hantao Yao, Rui Zhang. Visual-language prompt tuning with knowledge-guided context optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 7

[14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6

[15] 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 *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 3606–3613, 2018. 6, 16

[16] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15262–15271, 2019. 6, 16

[17] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, page 8340–8349, 2021. 6, 16

[18] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning (ICML)*, pages 4904–4916, 2021. 3

[19] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *European Conference on Computer Vision (ECCV)*, pages 709–727, 2022. 1

[20] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19113–19122, 2023. 3

[21] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, page 3498–3505, 2012. 6, 16

[22] Ananya Kumar, Aditi Raghunathan, Robbie Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained features and underperform out-of-distribution. In *International Conference on Learning Representations (ICLR)*, pages 1–42, 2022. 3, 6, 7, 16, 17

[23] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3045–3059, 2021. 3

[24] Zhiqiu Lin, Samuel Yu, Zhiyi Kuang, Deepak Pathak, and Deva Ramanan. Multimodality helps unimodality: Cross-modal few-shot learning with multimodal models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 4, 6, 7, 8, 15, 16, 19

[25] Bingyuan Liu, Jérôme Rony, Adrian Galdran, Jose Dolz, and Ismail Ben Ayed. Class adaptive network calibration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16070–16079, 2023. 6

[26] S. Maji, J. Kannala, E. Rahtu, M. Blaschko, and A. Vedaldi. Fine-grained visual classification of aircraft. In *ArXiv Preprint*, 2013. 6, 16

[27] Sachit Menon and Carl Vondrick. Visual classification via description from large language models. In *International Conference on Learning Representations (ICLR)*, pages 1–17, 2023. 1

[28] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics and Image Processing*, 2008. 6, 13, 16

[29] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, page 3498–3505, 2012. 6, 13, 16

[30] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*, pages 8748–8763, 2021. 1, 3, 4, 6, 7, 8, 19, 21

[31] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning (ICML)*, pages 5389–5400, 2019. 6, 16

[32] Sara Sangalli, Ertunc Erdil, Andeas Hötker, Olivio F Donati, and Ender Konukoglu. Constrained optimization to train neural networks on critical and under-represented classes. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 5

[33] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. In *ArXiv Preprint*, 2012. 6, 16

[34] Rohan Taori, Achal Dave, Vaishaal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig Schmidt. Measuring robustness to natural distribution shifts in image classification. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 4

[35] 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)*, 2019. 6, 16

[36] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo-Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, and Ludwig Schmidt. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7959–7971, 2022. 3, 7, 11, 16, 17

[37] Jianxiong Xiao, James Hays, Krista A. Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3485–3492, 2010. 6, 16

[38] Yinghui Xing, Qirui Wu, De Cheng, Shizhou Zhang, Guoqiang Liang, Peng Wang, and Yanning Zhang. Dual modality prompt tuning for vision-language pre-trained model. *IEEE Transactions on Multimedia*, 2023. 3

[39] Hantao Yao, Rui Zhang, and Changsheng Xu. Visual-language prompt tuning with knowledge-guided context optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6757–6767, 2023. 3

[40] Tao Yu, Zhihe Lu, Xin Jin, Zhibo Chen, and Xinchao Wang. Task residual for tuning vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10899–10909, 2023. 1, 2, 3, 4, 6, 7, 8, 11, 12, 14, 15, 16, 19, 21

[41] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18123–18133, 2022. 3

[42] Renrui Zhang, Rongyao Fang, Wei Zhang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free clip-adapter for better vision-language modeling. In *European Conference on Computer Vision (ECCV)*, pages 1–19, 2022. 1, 2, 3, 4, 6, 7, 8, 11, 12, 14, 15, 16, 19, 21

[43] Zexuan Zhong, Dan Friedman, and Danqi Chen. Factual probing is [mask]: Learning vs. learning to recall. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5017–5033, 2021. 3

[44] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 45:4396–4415, 2022. 4

[45] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 3, 6

[46] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision (IJCV)*, 2022. 1, 2, 6, 7, 15, 16, 19

[47] Beier Zhu, Yulei Niu, Yucheng Han, Yue Wu, and Hanwang Zhang. Prompt-aligned gradient for prompt tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15659–15669, 2023. 3

A Closer Look at the Few-Shot Adaptation of Large Vision-Language Models

Supplementary Material

A. A closer look to the pitfalls of previous vision-language adapters

In this work, we provide a closer view of the pitfalls of current literature on few-shot vision-language adapters of large vision-language models. In particular, we observe that recently proposed adapters rely on a large test subset to adjust important hyperparameters per dataset, and thus become *impractical* in real-world few-shot scenarios. This limitation becomes evident when fixing the hyperparameters on a given scenario and testing the model on other tasks, SoTA methods typically see their performance degrade compared to aa well-initialized Linear Probing (see Supp. Fig. 2). In the following section, we aim to depict a detailed view of these methods and the reasons that underlay their promising reported performance.

A.1. What are SoTA adapters doing?

We observe two concurrent phenomena on current SoTA vision-language adapters: (i) they use a good initialization, based on the zero-shot prototypes; and (ii) they introduce a set of empirically-fixed hyperparameters that control the divergence from the initial set of initial zero-shot prototypes.
CLIP-Adapter [11]. The inference relies on the zero-shot inference as in Eq. (1), using the same prototypes, which remain static. During training, CLIP-Adapter trains a residual MLP block to refine the visual features, such that $\mathbf{v}' = \mathbf{v} + \alpha_{\mathrm{r}} f_\psi(\mathbf{v})$. This method explicitly keeps the class prototypes close to the zero-shot initialization, while modifies the input visual features. This modification can be controlled with the residual ratio, α_{r}, together with the used learning rate and early stopping at specific epochs.
TIP-Adapter [42]. Training-free CLIP proposes a multimodal combination of logits, using two terms: (i) a weighted similarity to the support sample, and (ii), the similarity of the zero-sot prototypes. This dual formulation can be expressed in the following formula:

$$\mathbf{l}_c = \underbrace{\alpha_{\mathrm{tipA}} f_\psi(\mathbf{v}, \beta)}_{\text{vision logits}} + \underbrace{\tau\, \mathbf{v} \cdot \mathbf{t}_c^\top}_{\text{zero-shot logits}} \qquad (14)$$

where α_{tipA} and β are control hyperparameters, which are empirically fixed.

There are two versions of TIP-Adapter. First, a training-free version, in which the vision logits are obtained by the post-processed version of the average cosine similarity between the vision embedding of the target and the support samples per class. Second, a trainable version in which, additionally, the vision embeddings from the support set are tunned, which dramatically increases the number of trainable parameters with the number of shots.

Since details matter, it is worth mentioning that in the combined logits depicted in Supp. Eq. (14), temperature scaling is only applied on the cosine similarity of text prototypes. The τ value is learned during training and usually converges to large values, which are clipped at a maximum value of 100. This scaling makes the logits obtained from the zero-shot weights *dominate* in the combined formulation if α is not properly fixed to large values. As we previously stated, this results in an initialization close to the zero-shot prototypes, and the deviation from this solution is carefully controlled with an α scaling per dataset.
Task Residual Learning [40]. TaskRes uses a linear classifier to obtain class prototypes, following Eq. (3). In particular, the authors propose to train a *"prior-independent task residual"*, which follows a re-parametrization of the learned prototypes $\mathbf{w}$, such that $\mathbf{w} = t + \alpha \mathbf{w}_r$, where t is the language prototypes for the target classes (zero-shot weights), and w_r is a learnable matrix that modifies them. This modification is controlled by the hyperparameter α, which is empirically fixed for each dataset. Since $\mathbf{w}_r$ is initialized to a matrix filled with zeros, the re-parametrization is equivalent to using a zero-shot initialization at the first iteration. In addition, given a feature vector $\mathbf{v}$ of a support sample, and optimizing $\mathbf{w}_r$ via gradient descent using Eq. (2), it is straightforward to derive that this term simply introduces a scaling factor on a given learning rate η, and no additional information is introduced to a simple Linear Probing:

$$\mathbf{w}_r^t = \mathbf{w}_r^{t-1} - \eta \frac{\partial \mathcal{H}(\mathbf{y}, \hat{\mathbf{y}})}{\partial \mathbf{w}_r} = \mathbf{w}_r^{t-1} - \eta \frac{\partial \mathcal{H}(\mathbf{y}, \hat{\mathbf{y}})}{\partial \mathbf{w}} \frac{\partial \mathbf{w}}{\partial \mathbf{w}_r} = \\ \mathbf{w}_r^{t-1} - \underbrace{(\eta\,\alpha)}_{\text{learning rate}} \mathbf{v}\,(\hat{\mathbf{y}} - \mathbf{y}) .$$

$$(15)$$

Other related literature. Albeit proposed in the context of full fine-tuning of vision-language models in the large data regime, WiSE [36] approach also introduces some insights on efficient adaptation using a simple linear classifier. In particular, the authors study the benefit of linear interpolation between fine-tuned and zero-shot (initial) weights. In the case of adjusting uniquely a linear classifier, this method would be equivalent to balancing the text embeddings and the trained prototypes, such that $w = \alpha\, w_{\mathrm{LP}} + (1 - \alpha)\, t$. Concretely, the ratio α is fixed using a validation subset after training, The idea that underlays this method aligns with the observations derived from the few-shot adapters.
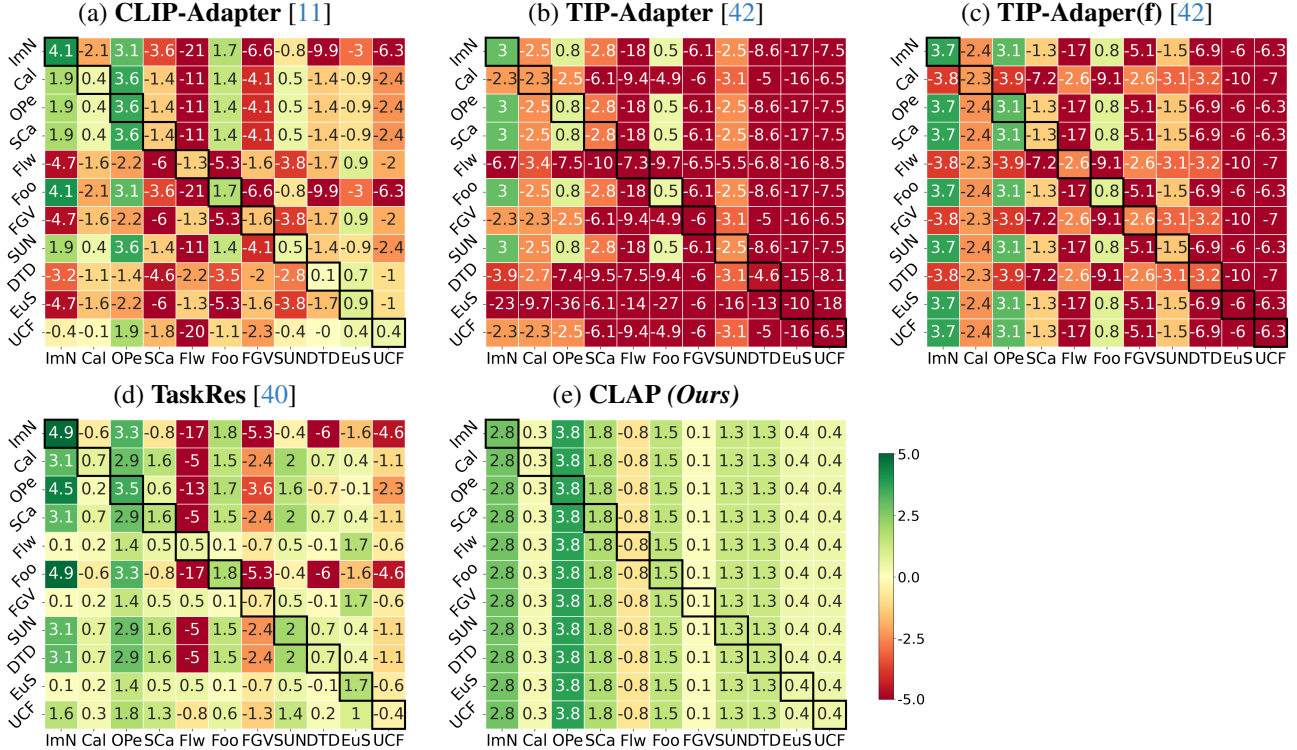
(a) **CLIP-Adapter** [11]

	ImN	Cal	OPe	SCa	Flw	Foo	FGV	SUN	DTD	EuS	UCF
ImN	4.1	-2.1	3.1	-3.6	-21	1.7	-6.6	-0.8	-9.9	-3	-6.3
Cal	1.9	0.4	3.6	-1.4	-11	1.4	-4.1	0.5	-1.4	-0.9	-2.4
OPe	1.9	0.4	3.6	-1.4	-11	1.4	-4.1	0.5	-1.4	-0.9	-2.4
SCa	1.9	0.4	3.6	-1.4	-11	1.4	-4.1	0.5	-1.4	-0.9	-2.4
Flw	-4.7	-1.6	-2.2	-6	-1.3	-5.3	-1.6	-3.8	-1.7	0.9	-2
Foo	4.1	-2.1	3.1	-3.6	-21	1.7	-6.6	-0.8	-9.9	-3	-6.3
FGV	-4.7	-1.6	-2.2	-6	-1.3	-5.3	-1.6	-3.8	-1.7	0.9	-2
SUN	1.9	0.4	3.6	-1.4	-11	1.4	-4.1	0.5	-1.4	-0.9	-2.4
DTD	-3.2	-1.1	-1.4	-4.6	-2.2	-3.5	-2	-2.8	0.1	0.7	-1
EuS	-4.7	-1.6	-2.2	-6	-1.3	-5.3	-1.6	-3.8	-1.7	0.9	-1
UCF	-0.4	-0.1	1.9	-1.8	-20	-1.1	-2.3	-0.4	-0	0.4	0.4

(b) **TIP-Adapter** [42]

	ImN	Cal	OPe	SCa	Flw	Foo	FGV	SUN	DTD	EuS	UCF
ImN	3	-2.5	0.8	-2.8	-18	0.5	-6.1	-2.5	-8.6	-17	-7.5
Cal	-2.3	-2.3	-2.5	-6.1	-9.4	-4.9	-6	-3.1	-5	-16	-6.5
OPe	3	-2.5	0.8	-2.8	-18	0.5	-6.1	-2.5	-8.6	-17	-7.5
SCa	3	-2.5	0.8	-2.8	-18	0.5	-6.1	-2.5	-8.6	-17	-7.5
Flw	-6.7	-3.4	-7.5	-10	-7.3	-9.7	-6.5	-5.5	-6.8	-16	-8.5
Foo	3	-2.5	0.8	-2.8	-18	0.5	-6.1	-2.5	-8.6	-17	-7.5
FGV	-2.3	-2.3	-2.5	-6.1	-9.4	-4.9	-6	-3.1	-5	-16	-6.5
SUN	3	-2.5	0.8	-2.8	-18	0.5	-6.1	-2.5	-8.6	-17	-7.5
DTD	-3.9	-2.7	-7.4	-9.5	-7.5	-9.4	-6	-3.1	-4.6	-15	-8.1
EuS	-23	-9.7	-36	-6.1	-14	-27	-6	-16	-13	-10	-18
UCF	-2.3	-2.3	-2.5	-6.1	-9.4	-4.9	-6	-3.1	-5	-16	-6.5

(c) **TIP-Adaper(f)** [42]

	ImN	Cal	OPe	SCa	Flw	Foo	FGV	SUN	DTD	EuS	UCF
ImN	3.7	-2.4	3.1	-1.3	-17	0.8	-5.1	-1.5	-6.9	-6	-6.3
Cal	-3.8	-2.3	-3.9	-7.2	-2.6	-9.1	-2.6	-3.1	-3.2	-10	-7
OPe	3.7	-2.4	3.1	-1.3	-17	0.8	-5.1	-1.5	-6.9	-6	-6.3
SCa	3.7	-2.4	3.1	-1.3	-17	0.8	-5.1	-1.5	-6.9	-6	-6.3
Flw	-3.8	-2.3	-3.9	-7.2	-2.6	-9.1	-2.6	-3.1	-3.2	-10	-7
Foo	3.7	-2.4	3.1	-1.3	-17	0.8	-5.1	-1.5	-6.9	-6	-6.3
FGV	-3.8	-2.3	-3.9	-7.2	-2.6	-9.1	-2.6	-3.1	-3.2	-10	-7
SUN	3.7	-2.4	3.1	-1.3	-17	0.8	-5.1	-1.5	-6.9	-6	-6.3
DTD	-3.8	-2.3	-3.9	-7.2	-2.6	-9.1	-2.6	-3.1	-3.2	-10	-7
EuS	3.7	-2.4	3.1	-1.3	-17	0.8	-5.1	-1.5	-6.9	-6	-6.3
UCF	3.7	-2.4	3.1	-1.3	-17	0.8	-5.1	-1.5	-6.9	-6	-6.3

(d) **TaskRes** [40]

	ImN	Cal	OPe	SCa	Flw	Foo	FGV	SUN	DTD	EuS	UCF
ImN	4.9	-0.6	3.3	-0.8	-17	1.8	-5.3	-0.4	-6	-1.6	-4.6
Cal	3.1	0.7	2.9	1.6	-5	1.5	-2.4	2	0.7	0.4	-1.1
OPe	4.5	0.2	3.5	0.6	-13	1.7	-3.6	1.6	-0.7	-0.1	-2.3
SCa	3.1	0.7	2.9	1.6	-5	1.5	-2.4	2	0.7	0.4	-1.1
Flw	0.1	0.2	1.4	0.5	0.5	0.1	-0.7	0.5	-0.1	1.7	-0.6
Foo	4.9	-0.6	3.3	-0.8	-17	1.8	-5.3	-0.4	-6	-1.6	-4.6
FGV	0.1	0.2	1.4	0.5	0.5	0.1	-0.7	0.5	-0.1	1.7	-0.6
SUN	3.1	0.7	2.9	1.6	-5	1.5	-2.4	2	0.7	0.4	-1.1
DTD	3.1	0.7	2.9	1.6	-5	1.5	-2.4	2	0.7	0.4	-1.1
EuS	0.1	0.2	1.4	0.5	0.5	0.1	-0.7	0.5	-0.1	1.7	-0.6
UCF	1.6	0.3	1.8	1.3	-0.8	0.6	-1.3	1.4	0.2	1	-0.4

(e) **CLAP** *(Ours)*

	ImN	Cal	OPe	SCa	Flw	Foo	FGV	SUN	DTD	EuS	UCF
ImN	2.8	0.3	3.8	1.8	-0.8	1.5	0.1	1.3	1.3	0.4	0.4
Cal	2.8	0.3	3.8	1.8	-0.8	1.5	0.1	1.3	1.3	0.4	0.4
OPe	2.8	0.3	3.8	1.8	-0.8	1.5	0.1	1.3	1.3	0.4	0.4
SCa	2.8	0.3	3.8	1.8	-0.8	1.5	0.1	1.3	1.3	0.4	0.4
Flw	2.8	0.3	3.8	1.8	-0.8	1.5	0.1	1.3	1.3	0.4	0.4
Foo	2.8	0.3	3.8	1.8	-0.8	1.5	0.1	1.3	1.3	0.4	0.4
FGV	2.8	0.3	3.8	1.8	-0.8	1.5	0.1	1.3	1.3	0.4	0.4
SUN	2.8	0.3	3.8	1.8	-0.8	1.5	0.1	1.3	1.3	0.4	0.4
DTD	2.8	0.3	3.8	1.8	-0.8	1.5	0.1	1.3	1.3	0.4	0.4
EuS	2.8	0.3	3.8	1.8	-0.8	1.5	0.1	1.3	1.3	0.4	0.4
UCF	2.8	0.3	3.8	1.8	-0.8	1.5	0.1	1.3	1.3	0.4	0.4

Figure 2. **Pitfalls of few-shot adapters due to the absence of a *model selection* strategy - Additional methods.** The cross-shift model selection matrices (i, j) depict the relative improvement w.r.t. a zero-shot initialized Linear Probing when using the optimal hyperparameters for the dataset i, for adapting in another task j, for each SoTA method (*first four plots*) and our approach (*last plot*). This is an extended version of Fig. 1 in the main manuscript.

A.2. Remaining close to the initial zero-shot prototypes

As we stated previously, we observed that recently proposed few-shot adapters benefit from a good initialization, which is obtained from robust class-wise text embeddings. Also, in the main manuscript, we introduce a revisited Linear Probing baseline tailored for vision-language models (ZS-LP in Sec. 4.1). This method benefits from this good initialization, together with other training heuristics. Indeed, we demonstrate empirically in the experimental section that it serves as a strong baseline for VLMs adaptation. We now study the convergence of this method during adaptation (curves in Supp. Fig. 3), to shed light on the benefits of using a good set of initial prototypes. Furthermore, our goal is to expose that in different datasets, deviating much from initial prototypes may, or may not, be beneficial. We stress that, as a reminder, the more iterations are performed during adaptation, the more the model predictions deviate from initial zero-shot representations, which can also be controlled with the step size, *a.k.a.*, learning rate. First, we can observe that, the zero-shot CLIP initialized Linear Probe (*orange line*) achieves a maximum in performance over test samples at different epochs, which do not correspond to the convergence on the support set. Indeed, letting the adaptation converge typically yields performance degradation in ZS-LP. Even though this solution (i.e., maximum performance on test samples) could be reached using a large validation subset, which can be used for tuning the hyperparameters and early stopping, its presence is unrealistic on a strict few-shot protocol. In contrast, it is worth mentioning that the proposed learnable class-adaptative Linear Probing (CLAP, see Sec. 4.3) prevents this degradation, and does not require access to any additional data. Last, we would like to highlight an interesting observation from the convergence points seen in these curves. In particular, and interestingly, the range of values for searching the corresponding hyperparameters in methods such as TIP-Adapter, varies with the convergence scenario for the best test performance (more details in the next section).

To provide further empirical evidence, we now study in Supp. Fig. 4 the performance obtained by a zero-shot initialized Linear Probing (ZS-LP) with a fixed scheduler (see Sec. 5.1 for details), just varying the initial learning rate. Larger learning rates might produce solutions farther from the initial data points, and vice-versa. In particular, we focus on two popular datasets used for adaptation: OxfordPets
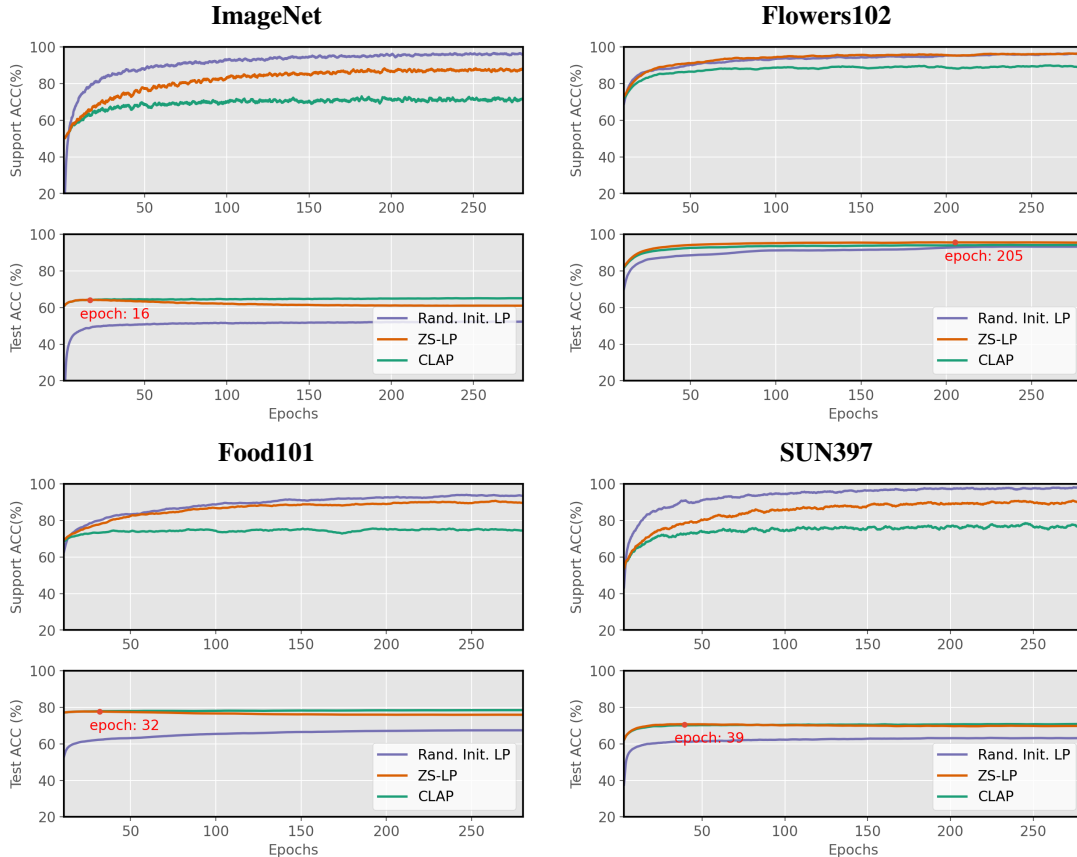
Figure 3. **Linear Probing learning curves**. Results of Linear Probing-based methods when adapted to ImageNet using ResNet-50 as a backbone, 16 shots per class as a support set, and a training scheduler using SGD. During training, both support set accuracy (*top*) and the performance on the test subset (*bottom*) are monitored, and the maximum test accuracy is highlighted in the curves. The training scheduler is described in Sec. 5.1.

[29] and Flowers102 [28]. The experimental results show that even for Linear Probing, adjusting the training specifications per task leads to better test generalization. For some datasets, such as OxfordPets, it is beneficial to underfit on the support set (see Supp. Fig. 4 top-left), and thus using smaller learning rates is beneficial. In other cases, such as FLowers102 (see Supp. Fig. 4 top-right), the degradation from fitting to the support set is not observed. Thus, each adaptation task presents its specific behavior. In a few-shot setting, however, only the support set information is available and model selection for a given adapter should rely only on this data. It is worth mentioning that the proposed class-adaptive solution (CLAP) is able to keep robust performance in both cases, using the same training setting across datasets.

A.3. SoTA methods: is it all about playing with hyperparameters?

We previously introduced the methodological basis of SoTA adapters in Supp. Sec. A.1, and the different hyperparam-

eters they use for model selection. Also, we have introduced in Supp. Sec. A.2 that each adaptation dataset might present different characteristics, and thus the optimum solution might be closer or farther to the zero-shot CLIP initialization. For instance, the Flowers101 dataset presents a particular behavior, which differs from other datasets (see Supp. Fig. 4). Interestingly, for this dataset, our cross-shift dataset experiments unveil that large performance drops are experienced in SoTA methods when using the optimum hyperparameters found for other tasks (see Supp. Fig. 2). In the following, we provide observational evidence that these methods adjust specific hyperparameter values per dataset, using prior knowledge from the test subset, which is unrealistic in practice.

CLIP-Adapter [11]. While an official implementation of the training code is not available, authors explicitly claim in the paper that: "*We perform hyperparameter searching over different value selections of α for each dataset and report the best performance among all searching spaces.*" In addition, we only could replicate their re-

13

Figure 4. **The trade-off between convergence on support set and generalization for zero-shot initialized adapters**. We depict the performance on the support and test subsets (after training) of zero-shot initialized Linear Probing adapters. Red numbers indicate the initial learning rate used, on the fixed scheduler described in Sec. 5.1. Two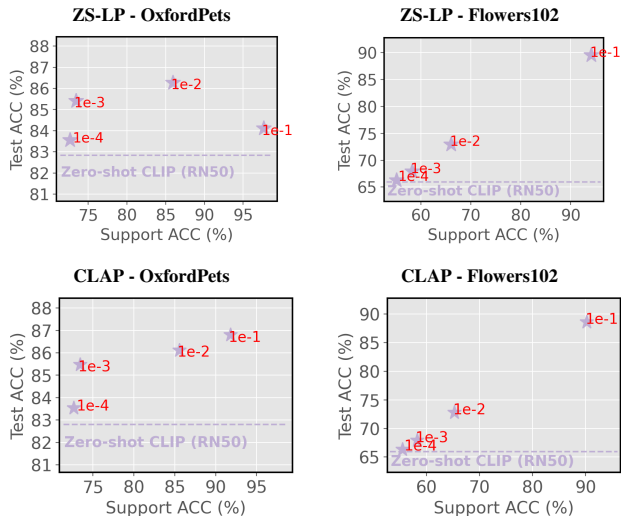 methods are presented: zero-shot initialized Linear Probe (ZS-LP, *top*, see Sec. 4.1), and class adaptive Linear Probe (CLAP, *bottom*, see Sec. 4.3).

sults when directly adjusting the learning rate (swept among $\{10^{-1}, 10^{-2}, 10^{-3}\}$) and residual ratio (searching values are $\alpha_r = \{0.2, 0.4, 0.6, 0.8, 1\}$) on a grid search at the test subset.

TIP-Adapter [42]. The absence of any details in the original publication regarding model selection strategies or the use of validation subsets leaves the GitHub repository as the only available documentation of the official implementation. In this repository, the authors claim (see Issue #13)[2]: "*The alpha and beta are both set to 1 as the tuning baseline. The alpha weighs the importance of CLIP-pre-trained and few-shot knowledge. If the few-shot domain has a large gap to pre-trained data (general images, just like ImageNet), alpha is better to be larger than 1.*" This suggests no explicit strategy for model selection exists. In addition, the official implementation contains a hyperparameter search function that takes as input the test subset in the case of ImageNet and a large validation subset for other tasks. It is worth mentioning that the grid search boundaries per hyperparameter also depend on each specific task. For instance, the α_{tipA} parameter for ImageNet is searched between $[1.17, 7]$, and for Flowers102 dataset the target range is $[10, 50]$, not presenting an overlapping at all. Interestingly, α_{tipA} controls the relative importance of the vision logits, and larger weight values are searched on Flowers102,

a dataset which, as we show in Supp. Fig. 4, benefits from diverging from the zero-shot initialization. These details suggest that the hyperparameters for TIP-Adapter methods are fixed assuming prior knowledge of the test subset for each particular task.

Task Residual Learning [42]. As previously described, TaskRes is equivalent to a zero-shot initialized Linear Probing, and contains an α parameter that regulates the learning rate per dataset. It is worth mentioning that the implementation details describe the use of different learning rates for ImageNet adaptation ($\eta = 2 \cdot 10^{-4}$), and for other tasks ($\eta = 2 \cdot 10^{-3}$), as well as different epochs depending on the number of shots. In addition, it is stated that "*By default, the scaling factor α is set to 0.5 for all datasets except for Flowers102 using 1*". This detail is especially relevant, as it suggests the access to prior knowledge to test performance. Again, a larger adaptation to the support samples is used for the Flowers102 task, which aligns with the low transferability of the hyperparameters set on this task to other datasets in Fig. 2, as well as with the longer convergence on test performance observed in this dataset (Supp. Fig. 3).

A.4. Choosing hyperparameters for a validation-free benchmark

In this work, we seek to provide a realistic protocol for comparing few-shot vision-language adapters. In this setting, we assume access to only the available support samples, and no additional validation examples are used. Next, we describe the implementation details of the different baselines and the motivation for the use of particular hyperparameter values.

For CLIP-Adpater [11], we set the hyperparameter α to 0.2 for all datasets, as it is the best value found on ImageNet evaluation in the original paper. The TIP-Adapter [42] umbrella gathers two methods: training-free, and a trainable version, *i.e.*, TIP-Adapter(f), in which the support samples embeddings are updated. For both methods, we set β and α to 1, as recommended in the official repository (see Supp. Footnote 2). For TaskRes [40], we only used as baseline its enhanced version, referred to as TaskRes(e)[3], which updates the projection layer of the text encoder. The reason for not using the base version of TaskRes is motivated by our findings that suggest that this method is equivalent to a Linear Probe tuning with zero-shot initialization, and a specific learning rate scaling for each dataset (see

---

[3]The training code for TaskRes(e) base is not provided in the official implementation (https://github.com/geekyutao/TaskRes) and might contain specific tuning that indirectly resorts to the test set. Authors uniquely share the enhanced weights, and the lack of specific implementation details might produce unfair comparisons. The only information available in the manuscript is: "*... enhanced base classifier obtained by tuning the text projection layer of CLIP on the target task before starting our task residual tuning ... The aforementioned enhanced base classifier is tuned for 50 epochs*".

Supp. Sec. A.1). We set α to 0.5 in TaskRes(e) since this is the value used in the majority of the tasks in the original publication. Finally, we included Cross-Modal adapters [24], in particular the Linear Probing version, which does not require special hyperparameter tuning. To avoid using an empirical grid search for weight decay, we implicitly applied an ℓ_2-normalization over the weights during training, which provided a better performance on our ablation experiments (see Sec. 5.3). All methods are trained using the same general optimizer and scheduler as our proposed methods, which showed proper convergence on the support set, and all baselines employ the same text prompts for each dataset.

A.5. Trainable parameters

Efficient transfer learning ought to exploit limited supervision during adaptation while being efficient in the number of trainable parameters. We depict in Supp. Fig. 5, a visualization of the trade-off between the number of trainable parameters and test performance of relevant prior methods, and the proposed class adaptive Linear Probing (CLAP, see Sec. 4.3). All results are obtained in the validation-free protocol, using the implementation details described in Sec. 5.1. While CLIP-Adapter is a specially lightweight solution, the obtained performance is limited with respect to the proposed method (CLAP), and even a well-initialized Linear Probing (ZS-LP). On the other hand, TIP-Adapter largely increases the number of tunable weights with the number of shots, which questions its transferability to other tasks, such as dense image segmentation, in which each pixel prototype would constitute an individual parameter. In contrast, CLAP just introduces a negligible set of additional trainable multipliers - one per class - over a Linear Probing solution, which considerably enhances its performance.

B. Penalty functions for ALM: axioms

In this section, we provide the requirements for a penalty function in the Augmented Lagrangian Multiplier (ALM) method, detailed in Sec. 4.3.

A function $P : \mathbb{R} \times \mathbb{R}_{++} \times \mathbb{R}_{++} \to \mathbb{R}$ is a Penalty-Lagrangian function such that $P'(z, \rho, \lambda) \equiv \frac{\partial}{\partial z} P(z, \rho, \lambda)$ exists and is continuous for all $z \in \mathbb{R}$, $\rho \in \mathbb{R}_{++}$ and $\lambda \in \mathbb{R}_{++}$. In addition, a penalty function P should satisfy the following four axioms [2]:

Axiom 1: $P'(z, \rho, \lambda) \geq 0 \quad \forall z \in \mathbb{R}, \rho \in \mathbb{R}_{++}, \lambda \in \mathbb{R}_{++}$

Axiom 2: $P'(0, \rho, \lambda) = \lambda \quad \forall \rho \in \mathbb{R}_{++}, \lambda \in \mathbb{R}_{++}$

Axiom 3: If, for all $j \in \mathbb{N}$, $\lambda^{(j)} \in [\lambda_{\min}, \lambda_{\max}]$, where $0 < \lambda_{\min} \leq \lambda_{\max} < \infty$, then:
$$\lim_{j \to \infty} \rho^{(j)} = \infty \text{ and } \lim_{j \to \infty} y^{(j)} > 0 \text{ imply that}$$
$$\lim_{j \to \infty} P'(y^{(j)}, \rho^{(j)}, \lambda^{(j)}) = \infty$$

Axiom 4: If, for all $j \in \mathbb{N}$, $\lambda^{(j)} \in [\lambda_{\min}, \lambda_{\max}]$, where $0 < \lambda_{\min} \leq \lambda_{\max} < \infty$, then:
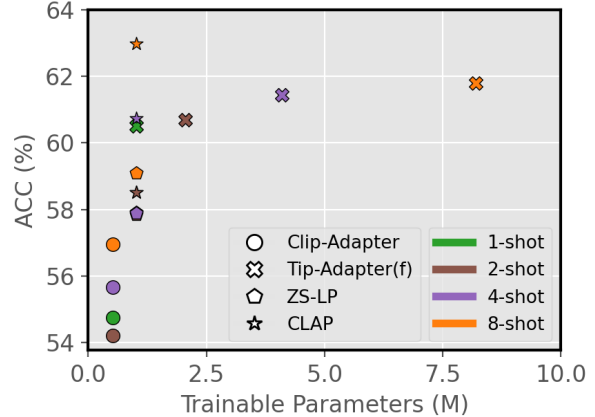
Figure 5. **Trade-off between number of shots, trainable parameters, and adaptation performance.** The test accuracy is presented with respect to the number of trainable parameters for CLIP-Adapter [11], TIP-Adapter(f) [42], and the two proposed solutions in this work: a revisited Linear Probing (ZS-LP, see Sec. 4.1), and a class-adaptive Linear Probing (CLAP, see Sec. 4.3). Results were obtained for 1 to 8 shots in the ImageNet dataset.

$$\lim_{j \to \infty} \rho^{(j)} = \infty \text{ and } \lim_{j \to \infty} y^{(j)} < 0 \text{ imply that}$$
$$\lim_{j \to \infty} P'(y^{(j)}, \rho^{(j)}, \lambda^{(j)}) = 0.$$

The first two axioms guarantee that the derivative of the Penalty-Lagrangian function P w.r.t. z is positive and equals to λ when $z = 0$. The last two axioms guarantee that the derivative tends to infinity when the constraint is not satisfied, and zero otherwise.

C. Supplementary experimental details

C.1. Additional setup information

Datasets details. In our main text, we introduce the datasets employed to evaluate the proposed methods and establish comparisons with relevant literature on the few-shot adaptation of CLIP-based models. In Supp. Tab. 6, we introduce the specific details of each dataset, including the number of categories, test partition size, and particular tasks.

Text prompt templates. We followed the same hand-crafted templates as relevant prior literature of efficient transfer learning for the 11 datasets. Concretely, we followed CoOp [46], TIP-Adapter [42] TaskRes [40], and CrossModal hand-crafted version [24]. These prompts are composed of an ensemble of 8 different templates for Imagenet-like datasets, and 1 template for the others, which are depicted in Supp. Tab. 6. It is worth mentioning that

Table 6. **Summary of datasets details.** Detailed description of the 11 datasets used to validate the SoTA few-shot adapters of VLMs, and 4 ImageNet shifts employed to evaluate the generalization capabilities of those. Also, handcrafted prompts used to obtain the zero-shot predictions and prototypes are detailed. These are the same ones used in relevant prior literature on this topic [24, 40, 46].

Dataset	Classes	Splits Train / Val / Test	Task	Prompt Templates
ImageNet [8]	1000	1.28M / - / 50,000	Natural objects recognition	["itap of a [CLS]" , "a bad photo of a [CLS]" ,
ImageNet-V2 [31]	1000	- / - / 10,000	Natural objects recognition	"a origami of [CLS]" , "a photo of the large [CLS]" ,
ImageNet-Sketch [35]	1000	- / - / 50,889	Sketch-style image classification	"a [CLS] in a video game" , "art of the [CLS]" ,
ImageNet-A [16]	200	- / - / 7,500	Natural objects recognition	"a photo of the small [CLS]" , "a photo of a [CLS]"]
ImageNet-R [17]	200	- / - / 30,000	Natural objects recognition	
Caltech101 [10]	100	4,128 / 1,649 / 2,465	Natural objects classification	["a photo of a [CLS]"]
OxfordPets [29]	37	2,944 / 736 / 3,669	Pets classification (fine-grained)	["a photo of a [CLS], a type of a pet"]
StanfordCars [21]	196	6,509 / 1,635 / 8,041	Cars classification (fine-grained)	["a photo of a [CLS]"]
Flowers102 [28]	102	4,093 / 1,633 / 2,463	Flowers classification (fine-grained)	["a photo of a [CLS], a type of flower"]
Food101 [4]	101	50,500 / 20,200 / 30,300	Foods classification (fine-grained)	["a photo of a [CLS], a type of food"]
FGVCAircraft [26]	100	3,334 / 3,333 / 3,333	Aircrafts classification (fine-grained)	["a photo of a [CLS], a type of aircraft"]
SUN397 [37]	397	15,880 / 3,970 / 19,850	Scenes classification	["a photo of a [CLS]"]
DTD [7]	47	2,820 / 1,128 / 1,692	Textures classification	["[CLS] texture"]
EuroSAT [15]	10	13,500 / 5,400 / 8,100	Satellite image classification	["a centered satellite photo of [CLS]"]
UCF101 [33]	101	7,639 / 1,898 / 3,783	Recognition of actions	["a photo of a person doing [CLS]"]

fine-tuning methods used for the benchmark in Sec. 5.2 (LP-FT [22], FLYP [12]) use a larger set of 80 prompt templates for ImageNet-like datasets, although these are not usually used in the efficient transfer learning literature.

C.2. Results: supplementary details

Efficient transfer learning. We provide detailed numerical results for the few-shot adaptation experiments using relevant baselines and the proposed methods in Supp. Tab. 9, which extend the values reported in Tab. 1. Furthermore, we also depict visual curves of the performance with respect to the number of shots employed by each method in Supp. Fig. 8.

Domain generalization. We show in the main manuscript the domain generalization results using adapters adjusted to ImageNet and evaluated on out-of-distribution shifts (*i.e.*, ImageNet variants). These results are obtained using ResNet-50 and ViT-B/16 CLIP backbones. In the following, we introduce detailed results per dataset, and two additional backbones: ResNet-101 and ViT-B/32, whose results are reported in Tab. 10. We can observe that the results using these additional backbones hold the conclusions elucidated in the main manuscript. In particular, relevant prior methods such as CLIP-Adapter [11] and TIP-Adapter [42] struggle to generalize properly when their hyperparameter setting is held on different backbones than the one used for development, ResNet-50. This is especially the case for Transformer backbones, such as ViT-B/32, which suggests again that existing adapter methods need special care for model selection across each dataset.

Finetuning (FT) vs. efficient transfer learning (ETL), beyond few-shots. Fine-tuning a whole vision encoder to downstream tasks using a few-shot training subset has been historically less favored compared to efficient transfer learning strategies, due to the tendency of FT methods to overfit to the new data, and thus generalizing poorly. Nevertheless, a relevant core of recent literature for VLMs adaptation [12, 22, 36] is showing promising results on this task. As stated in the main body of the paper, this is due to (i) using a few-shot validation dataset, with which they early-stop the training, and (ii) employing small learning rates to not deviate from a good initialization. Nevertheless, if compared properly in the low data regime, *i.e.*, using 4 shots for training and 4 samples per class for validation, and allowing ETL methods that do not require a validation set to use all samples for training, then ETL still seems to yield competitive performance, being a much more computationally-efficient solution. The results previously presented in the main body of the manuscript (see Tab. 3) support these observations.

We now extend this comparison to a scenario in which more data is available. Concretely, a 32-shot scenario, where FT methods use half of it for validation. It is worth mentioning that this experimental setting on ImageNet required 32,000 images (16,000 for validation), which might be hardly considered a few-shot learning protocol. We introduce specific results using 32-shot for ImageNet and its distributional shifts for relevant baselines and the proposed methods in Supp. Tab. 7. It is worth mentioning that FT methods use 16-shots for training and another 16-shots for validation. In addition, we present in Supp. Fig. 6 a study of the performance evolution with respect to the number of shots of relevant FT methods with an increasing number of parameters. More concretely, we include LP-FT [22], which completely fine-tunes the CLIP's vision backbone, and FLYP [12], which trains both vision and text encoders. We compare these results to the proposed class adaptive Linear Probing (CLAP), which only adjusts the

classifier head, and thus brings a negligible computational overhead compared to LP-FT and FLYP. In the 32-shot setting, CLAP shows competitive performance compared to methods that adjust the vision encoder entirely, such as FT, LP-FT [22], and WiSE-FT [36], for both in-distribution and out-of-distribution datasets, while adjusting only the linear classification head. Only FLYP [12], which requires fine-tuning both vision and text encoders, outperforms CLAP using 32,000 images, and by a small margin: 1.2% in ID, and 1.6% in OOD. Nevertheless, this comes at the cost of adjusting the entire CLIP model, which entails a non-negligible computational overhead, making this method an inefficient approach in low-resource scenarios. Note that CLAP is two orders of magnitude lighter than FLYP. In addition, CLAP does not exhibit signs of performance saturation (as LP-FT, for example) with an increase in the number of shots (see Supp. Fig. 6).
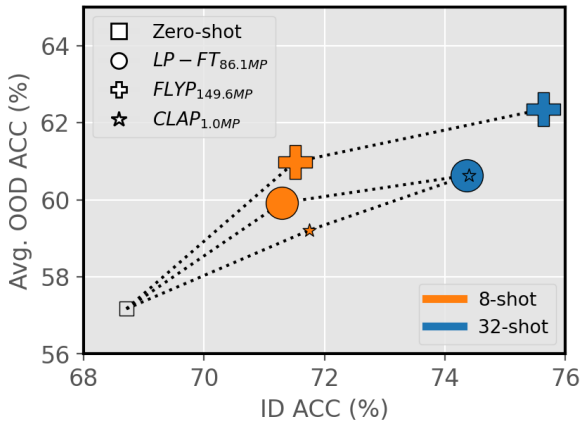
Figure 6. **Finetuning (FT) vs. efficient transfer learning (ETL), performance and trainable parameters**. We compare the generalization performance of relevant full fine-tuning methods, *i.e.*, LP-FT [22] and FLYP [12], and the proposed efficient transfer learning method CLAP (see Sec. 4.3), trained on ImageNet and evaluated on OOD shifts. Relative point size is illustrated as the log ratio of the number of tunable parameters of each method with respect to CLAP. MP: millions of parameters.

C.3. Supplementary ablation experiments

Distilling reliable knowledge. We now study the effect of resorting to a class-adaptive constrained formulation in the proposed CLAP. In particular, we further assess the benefits of using our class-dependent adaptive scaling ($\boldsymbol{\lambda}^\star$) of the imposed constraint, which is initialized on the performance of the zero-shot CLIP prototypes on the support set. To do so, we take as baseline different empirically-set multipliers baselines. First, we explore a homogeneous weight for all classes, such that $\boldsymbol{\lambda} = \mathbf{1}$. Furthermore, we aim to disentangle two different effects that $\boldsymbol{\lambda}^\star$ might have: (i) changing

Table 7. **Finetuning (FT) vs. efficient transfer learning (ETL), beyond few-shots.** Benchmark for the not-so-low data regime, *i.e.*, 32 shots for each class. FT methods (above the dashed line) are trained with 16 shots and early-stopped using a validation set containing 16 shots. WiSE-FT and FLYP use weight ensembling as proposed in [36], and therefore, find the best *mixing coefficient* α using the validation set. On the other hand, ETL methods (below the dashed line) are trained using all the 32 shots given. All methods use ViT-B/16 as CLIP backbone.

Method	Source Imagenet	Target				
		-V2	-Sketch	-A	-R	Avg.
FT	71.86	64.15	47.97	48.23	75.96	59.08
LP-FT [22]	74.36	66.43	49.35	49.84	76.89	60.63
WiSE-FT [36]	73.06	65.70	50.03	51.04	78.22	61.25
FLYP [12]	**75.63**	68.17	50.66	52.09	78.49	**62.35**
Zero-Shot	68.71	60.76	46.18	47.76	73.98	57.17
LP	67.40	56.43	31.71	31.92	51.04	42.71
ZS-LP	71.53	56.59	40.84	41.41	67.98	51.71
CLAP	74.40	66.05	49.16	49.82	77.52	60.64

*Specific numbers for FT, LP-FT, WiSE-FT, and FLYP are retrieved from [12].

the overall relative importance of the constraint term with respect to the cross-entropy term in Eq. (7); and, (ii) providing the capability of capturing class dependent prior knowledge from the pre-trained model. Thus, we further include two alternative ways of computing $\boldsymbol{\lambda}$ in our ablation study: (i) a constant version of the constraint formulation, in which all multipliers are set to a constant $\lambda^{avg} = \frac{1}{C}\sum_c \lambda_c^\star$, *i.e.*, the average importance of the constraint; and, (ii) an importance corrected version of the constrained formulation, $\boldsymbol{\lambda}^{corr} = \boldsymbol{\lambda}^\star/\lambda^{avg}$, such that $\frac{1}{C}\sum_i \lambda_i^{corr} = 1$. The average performance over 11 datasets for the few-shot data regime is shown in Supp. Fig. 7, whereas the full numerical results per dataset are presented in Supp. Tab. 11.
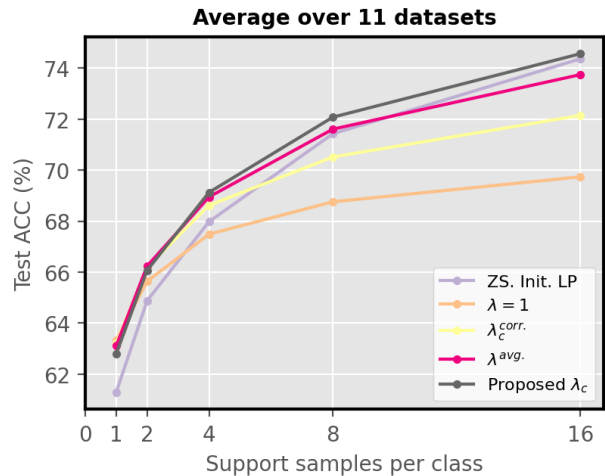
Figure 7. **Ablation study on different alternatives to compute initial $\boldsymbol{\lambda}$ in Eq. (7).** The average performance over 11 datasets is reported.

In the special case of 1-shot or 2-shot, adaptive computation of $\boldsymbol{\lambda}$ may lead to a slight overfitting to the information provided by the few samples, even though the performance gap compared to other strategies is minimal. The more shots provided, the less noisy the information will be and computing class-dependent adaptive importance will prove more useful. Looking at the plots, except for the 1-shot setting, using $\boldsymbol{\lambda} = \mathbf{1}$ is always suboptimal. In the 2-shot scenario, using $\lambda^{avg.}$ for all classes yields the best results, although close to the performance of the proposed $\boldsymbol{\lambda}^\star$. Generally, when the number of labeled samples per class is particularly low, the proposed $\boldsymbol{\lambda}^\star$ is not the best-performing option. However, it proves superior to other alternatives when 4 or more shots are given, and the gap widens as the number of shots increases. It is worth noting that the non-constrained (ZS-LP) version's performance approaches the constrained version in the 16-shot setting, but it underperforms in the lower data regimes, which demonstrates the effectiveness of our formulation across the different regimes.

Table 8. **Augmented Lagrangian multiplier optimization.** We present ablation experiments that motivate the use of only one outer iteration in Eq. (9), which prevents overfitting on the support samples, due to the absence of a validation subset in the realistic few-shot scenario. K denotes the number of shots.

Method	$K=1$	$K=2$	$K=4$
ZS-LP	61.28	64.88	67.98
CLAP - Full outer loop	$61.97_{(+0.7)}\uparrow$	$63.25_{(-1.6)}\downarrow$	$63.83_{(-4.2)}\downarrow$
CLAP - 1 outer loop (*Ours*)	$62.79_{(+1.5)}\uparrow$	$66.07_{(+1,2)}\uparrow$	$69.13_{(+1.2)}\uparrow$

Updating the Lagrangian multipliers beyond the first iteration. The proposed class adaptive Linear Probing (CLAP) is based on an adaptation of the general Augmented Lagrangian Multiplier method, which learns the multiplier weights per class accounting for the particular difficulty of each category. Given the strict few-shot setting used, we propose to use the support samples to validate the satisfaction of the constrained problem. We hypothesized, however, that doing this could increase the risk of overfiting and proposed to stop after one single iteration of the outer optimization iteration in Eq. (9) (i.e., the step where the penalty multiplies $\boldsymbol{\lambda}$ are updated). In this section, we provide the empirical evidence for this hypothesis, which validates our choice. In particular, we also perform the adaptation using both inner and outer iterations during the whole adaptation process. To do so, at each epoch, the Lagrangian multipliers are updated following Eq. (10). Regarding the penalty multipliers ρ, these are initialized per class to the initial penalty value, and for each epoch, their value is fixed after updating the penalty multipliers to the resulting penalty after the inner iteration. Results are reported in Tab. 8, which show that updating these parameters continuously, based on the support set, results in overfitting and thus provides worse generalization.

Table 9. **Efficient transfer learning performance.** Full numerical performance comparison on the few-shot setting, using ResNet-50 as the backbone. All experiments are run with a fixed configuration, and training is done until full convergence on the support set. Results are averaged across 3 random seeds. Results for CoOp and PLOT are directly extracted from [6].

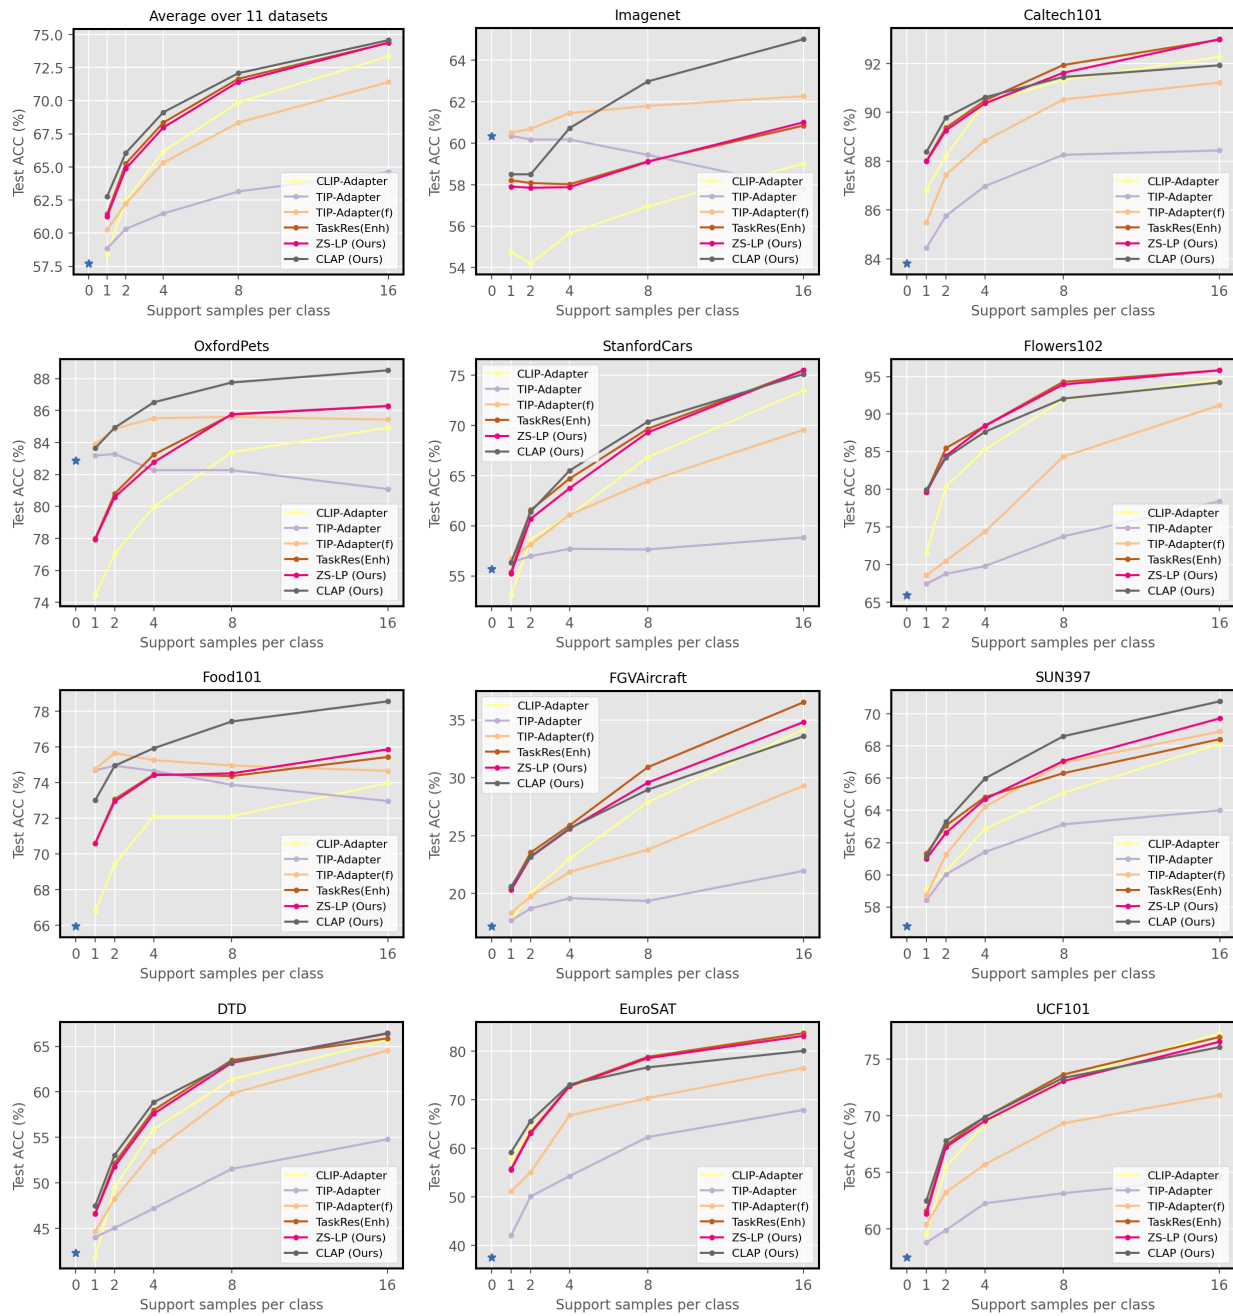| Method	Setting	ImageNet	Caltech101	OxfordPets	StanfordCars	Flowers102	Food101	FGVCAircraft	SUN397	DTD	EuroSAT	UCF101	Average
CoOp IJCV'22[46]		$56.99_{\pm1.03}$	$87.51_{\pm1.02}$	$85.99_{\pm0.28}$	$55.81_{\pm1.67}$	$67.98_{\pm1.98}$	$74.25_{\pm1.52}$	$8.59_{\pm5.79}$	$60.12_{\pm0.82}$	$43.62_{\pm1.96}$	$52.12_{\pm5.46}$	$62.13_{\pm1.14}$	$59.56_{\pm2.06}$
PLOT ICLR'23[6]		$59.54_{\pm0.16}$	$89.83_{\pm0.33}$	$87.49_{\pm0.57}$	$56.60_{\pm0.36}$	$71.72_{\pm0.97}$	$77.74_{\pm0.47}$	$17.90_{\pm0.09}$	$62.47_{\pm0.43}$	$46.55_{\pm2.62}$	$54.05_{\pm5.95}$	$64.53_{\pm0.70}$	$62.59_{\pm1.13}$
Zero-Shot ICML'21[30]		$60.35_{\pm0.00}$	$83.81_{\pm0.00}$	$82.86_{\pm0.00}$	$55.69_{\pm0.00}$	$65.94_{\pm0.00}$	$74.85_{\pm0.00}$	$17.16_{\pm0.00}$	$56.80_{\pm0.00}$	$42.32_{\pm0.00}$	$37.53_{\pm0.00}$	$57.47_{\pm0.00}$	$57.71_{\pm0.00}$
Rand. Init LP ICML'21[30]		$17.62_{\pm0.01}$	$56.82_{\pm1.65}$	$26.61_{\pm0.50}$	$18.41_{\pm0.95}$	$53.42_{\pm2.84}$	$23.92_{\pm0.86}$	$12.14_{\pm0.35}$	$26.66_{\pm0.49}$	$26.69_{\pm1.22}$	$40.69_{\pm4.58}$	$31.69_{\pm0.39}$	$30.42_{\pm1.26}$
CLIP-Adapter IJCV'23[11]		$54.74_{\pm0.19}$	$86.80_{\pm0.10}$	$74.45_{\pm1.57}$	$53.07_{\pm0.37}$	$71.57_{\pm0.78}$	$66.78_{\pm0.76}$	$17.01_{\pm0.06}$	$59.10_{\pm0.17}$	$41.80_{\pm1.84}$	$57.92_{\pm1.68}$	$59.49_{\pm0.33}$	$58.43_{\pm0.71}$
TIP-Adapter ECCV'22[42]		$60.35_{\pm0.08}$	$84.44_{\pm0.35}$	$83.18_{\pm1.01}$	$56.32_{\pm0.47}$	$67.45_{\pm0.16}$	$74.69_{\pm0.15}$	$17.69_{\pm0.26}$	$58.41_{\pm0.04}$	$44.03_{\pm0.30}$	$42.08_{\pm3.08}$	$58.79_{\pm1.06}$	$58.86_{\pm0.54}$
TIP-Adapter(f) ECCV'22[42]	1-shot	$60.51_{\pm0.06}$	$85.50_{\pm0.46}$	$83.90_{\pm0.92}$	$56.71_{\pm0.50}$	$68.60_{\pm0.70}$	$74.76_{\pm0.15}$	$18.33_{\pm0.57}$	$58.73_{\pm0.03}$	$44.64_{\pm0.29}$	$51.11_{\pm2.24}$	$60.38_{\pm0.33}$	$60.29_{\pm0.57}$
TaskRes(r) CVPR'23[40]		$57.91_{\pm0.25}$	$87.99_{\pm0.11}$	$77.94_{\pm2.58}$	$55.25_{\pm0.49}$	$79.62_{\pm0.45}$	$70.60_{\pm0.35}$	$20.30_{\pm0.81}$	$60.99_{\pm0.19}$	$46.59_{\pm1.51}$	$55.60_{\pm1.70}$	$61.32_{\pm0.75}$	$61.28_{\pm0.84}$
TaskRes(e) CVPR'23[40]		$58.21_{\pm0.19}$	$88.01_{\pm0.18}$	$78.01_{\pm2.55}$	$55.36_{\pm0.48}$	$79.83_{\pm0.54}$	$70.60_{\pm0.37}$	$20.50_{\pm0.82}$	$61.33_{\pm0.09}$	$46.63_{\pm1.49}$	$55.75_{\pm1.67}$	$61.59_{\pm0.80}$	$61.44_{\pm0.83}$
CrossModal-LP CVPR'23[24]		$57.40_{\pm0.11}$	$88.06_{\pm0.54}$	$80.09_{\pm1.41}$	$57.43_{\pm0.45}$	$78.51_{\pm0.54}$	$73.00_{\pm0.18}$	$20.72_{\pm0.29}$	$61.32_{\pm0.11}$	$47.81_{\pm1.52}$	$56.85_{\pm3.36}$	$63.49_{\pm0.23}$	$62.24_{\pm0.79}$
ZS-LP		$57.91_{\pm0.25}$	$87.98_{\pm0.13}$	$77.96_{\pm2.57}$	$55.24_{\pm0.46}$	$79.62_{\pm0.47}$	$70.60_{\pm0.35}$	$20.30_{\pm0.81}$	$61.00_{\pm0.19}$	$46.59_{\pm1.51}$	$55.57_{\pm1.71}$	$61.34_{\pm0.73}$	$61.28_{\pm0.83}$
CLAP		$58.50_{\pm0.24}$	$88.38_{\pm0.25}$	$83.64_{\pm1.18}$	$55.69_{\pm0.40}$	$79.90_{\pm0.46}$	$73.00_{\pm0.14}$	$20.62_{\pm0.55}$	$61.15_{\pm0.18}$	$47.46_{\pm1.15}$	$59.21_{\pm0.82}$	$62.48_{\pm0.99}$	$62.79_{\pm0.58}$
CoOp IJCV'22[46]		$56.40_{\pm0.87}$	$87.84_{\pm1.10}$	$82.22_{\pm2.15}$	$58.41_{\pm0.43}$	$77.58_{\pm1.46}$	$72.61_{\pm1.33}$	$16.52_{\pm2.38}$	$59.60_{\pm0.76}$	$45.35_{\pm0.31}$	$59.00_{\pm3.48}$	$64.05_{\pm0.99}$	$61.78_{\pm1.39}$
PLOT ICLR'23[6]		$60.64_{\pm0.06}$	$90.67_{\pm0.21}$	$86.64_{\pm0.63}$	$57.52_{\pm0.71}$	$81.19_{\pm0.79}$	$77.70_{\pm0.02}$	$18.94_{\pm0.44}$	$61.71_{\pm0.65}$	$51.24_{\pm1.95}$	$64.21_{\pm1.90}$	$66.83_{\pm0.43}$	$65.23_{\pm0.72}$
Zero-Shot ICML'21[30]		$60.35_{\pm0.00}$	$83.81_{\pm0.00}$	$82.86_{\pm0.00}$	$55.69_{\pm0.00}$	$65.94_{\pm0.00}$	$74.85_{\pm0.00}$	$17.16_{\pm0.00}$	$56.80_{\pm0.00}$	$42.32_{\pm0.00}$	$37.53_{\pm0.00}$	$57.47_{\pm0.00}$	$57.71_{\pm0.00}$
Rand. Init LP ICML'21[30]		$26.91_{\pm0.36}$	$69.29_{\pm3.92}$	$38.88_{\pm2.99}$	$31.62_{\pm1.20}$	$66.38_{\pm0.52}$	$37.99_{\pm1.24}$	$16.61_{\pm0.75}$	$38.97_{\pm0.75}$	$36.98_{\pm0.19}$	$51.73_{\pm1.22}$	$45.15_{\pm0.39}$	$41.86_{\pm1.26}$
CLIP-Adapter IJCV'23[11]		$54.20_{\pm0.31}$	$88.22_{\pm0.67}$	$77.03_{\pm2.53}$	$58.62_{\pm0.37}$	$80.39_{\pm0.66}$	$69.43_{\pm0.57}$	$20.07_{\pm0.65}$	$60.22_{\pm0.65}$	$49.51_{\pm0.21}$	$63.95_{\pm1.84}$	$65.43_{\pm0.33}$	$62.46_{\pm0.71}$
TIP-Adapter ECCV'22[42]		$60.18_{\pm0.15}$	$85.76_{\pm0.64}$	$83.28_{\pm0.70}$	$56.97_{\pm0.06}$	$68.78_{\pm0.14}$	$74.94_{\pm0.04}$	$18.70_{\pm0.18}$	$60.03_{\pm0.18}$	$45.04_{\pm0.10}$	$50.07_{\pm0.30}$	$59.88_{\pm0.06}$	$60.33_{\pm0.54}$
TIP-Adapter(f) ECCV'22[42]	2-shot	$60.69_{\pm0.14}$	$87.45_{\pm0.36}$	$84.86_{\pm0.47}$	$58.14_{\pm0.05}$	$70.51_{\pm0.05}$	$75.65_{\pm0.25}$	$19.77_{\pm0.26}$	$61.26_{\pm0.26}$	$48.25_{\pm0.13}$	$55.08_{\pm0.29}$	$63.24_{\pm0.33}$	$62.26_{\pm0.57}$
TaskRes(r) CVPR'23[40]		$57.86_{\pm0.05}$	$89.26_{\pm0.21}$	$80.59_{\pm1.57}$	$60.69_{\pm0.41}$	$84.48_{\pm0.20}$	$72.94_{\pm0.36}$	$23.16_{\pm0.36}$	$62.61_{\pm0.35}$	$51.79_{\pm0.20}$	$63.06_{\pm1.51}$	$67.26_{\pm0.75}$	$64.88_{\pm0.84}$
TaskRes(e) CVPR'23[40]		$58.08_{\pm0.12}$	$89.37_{\pm0.33}$	$80.80_{\pm1.54}$	$61.56_{\pm0.84}$	$85.49_{\pm0.68}$	$73.06_{\pm0.09}$	$23.55_{\pm0.49}$	$63.06_{\pm0.65}$	$52.17_{\pm0.39}$	$63.33_{\pm1.49}$	$67.39_{\pm0.80}$	$65.26_{\pm0.83}$
CrossModal-LP CVPR'23[24]		$49.13_{\pm0.20}$	$89.55_{\pm0.36}$	$81.72_{\pm0.72}$	$61.76_{\pm0.27}$	$82.30_{\pm0.55}$	$74.31_{\pm0.32}$	$22.46_{\pm0.30}$	$63.91_{\pm0.24}$	$53.09_{\pm1.44}$	$62.91_{\pm1.41}$	$68.13_{\pm0.67}$	$64.48_{\pm0.59}$
ZS-LP		$57.85_{\pm0.04}$	$89.26_{\pm0.21}$	$80.56_{\pm1.58}$	$60.69_{\pm0.42}$	$84.46_{\pm0.19}$	$72.94_{\pm0.35}$	$23.18_{\pm0.35}$	$62.61_{\pm0.36}$	$51.79_{\pm0.20}$	$63.06_{\pm1.51}$	$67.23_{\pm0.73}$	$64.88_{\pm0.83}$
CLAP		$58.50_{\pm0.24}$	$89.79_{\pm0.15}$	$84.93_{\pm0.66}$	$61.40_{\pm0.38}$	$84.22_{\pm0.35}$	$74.94_{\pm0.24}$	$23.31_{\pm0.32}$	$63.31_{\pm0.03}$	$53.05_{\pm0.13}$	$65.63_{\pm1.15}$	$67.77_{\pm0.99}$	$66.07_{\pm0.58}$
CoOp IJCV'22[46]		$58.48_{\pm0.47}$	$89.52_{\pm0.80}$	$86.65_{\pm0.97}$	$62.74_{\pm0.16}$	$86.10_{\pm1.05}$	$73.49_{\pm2.03}$	$20.63_{\pm2.46}$	$63.24_{\pm0.63}$	$53.94_{\pm1.37}$	$68.61_{\pm3.54}$	$67.79_{\pm0.71}$	$66.47_{\pm1.29}$
PLOT ICLR'23[6]		$61.49_{\pm0.23}$	$90.80_{\pm0.20}$	$88.63_{\pm0.26}$	$63.41_{\pm0.29}$	$87.82_{\pm0.20}$	$77.21_{\pm0.43}$	$22.36_{\pm0.42}$	$65.09_{\pm0.43}$	$56.03_{\pm0.43}$	$72.36_{\pm2.29}$	$69.60_{\pm0.67}$	$68.60_{\pm0.52}$
Zero-Shot ICML'21[30]		$60.35_{\pm0.00}$	$83.81_{\pm0.00}$	$82.86_{\pm0.00}$	$55.69_{\pm0.00}$	$65.94_{\pm0.00}$	$74.85_{\pm0.00}$	$17.16_{\pm0.00}$	$56.80_{\pm0.00}$	$42.32_{\pm0.00}$	$37.53_{\pm0.00}$	$57.47_{\pm0.00}$	$57.71_{\pm0.00}$
Rand. Init LP ICML'21[30]		$36.98_{\pm0.57}$	$78.11_{\pm2.90}$	$50.00_{\pm2.03}$	$44.75_{\pm0.28}$	$77.21_{\pm1.15}$	$50.10_{\pm0.96}$	$21.08_{\pm0.79}$	$49.63_{\pm0.69}$	$47.97_{\pm0.40}$	$58.51_{\pm4.06}$	$54.26_{\pm0.46}$	$51.69_{\pm1.30}$
CLIP-Adapter IJCV'23[11]		$55.66_{\pm0.31}$	$90.39_{\pm0.26}$	$79.99_{\pm1.69}$	$61.04_{\pm0.59}$	$85.28_{\pm0.57}$	$72.10_{\pm0.12}$	$23.03_{\pm0.08}$	$62.85_{\pm0.32}$	$55.89_{\pm0.88}$	$72.49_{\pm3.30}$	$69.28_{\pm0.16}$	$66.18_{\pm0.75}$
TIP-Adapter ECCV'22[42]		$60.18_{\pm0.08}$	$86.98_{\pm0.40}$	$82.27_{\pm1.21}$	$57.70_{\pm0.76}$	$69.81_{\pm0.62}$	$74.65_{\pm0.24}$	$19.60_{\pm0.46}$	$61.42_{\pm0.41}$	$47.18_{\pm0.50}$	$54.29_{\pm3.66}$	$62.26_{\pm0.26}$	$61.49_{\pm0.78}$
TIP-Adapter(f) ECCV'22[42]	4-shot	$61.45_{\pm0.05}$	$88.84_{\pm0.78}$	$85.51_{\pm0.57}$	$61.09_{\pm0.50}$	$74.39_{\pm0.08}$	$75.25_{\pm0.20}$	$21.87_{\pm0.69}$	$64.23_{\pm0.16}$	$53.45_{\pm0.27}$	$66.77_{\pm3.38}$	$65.70_{\pm0.29}$	$65.32_{\pm0.63}$
TaskRes(r) CVPR'23[40]		$57.88_{\pm0.18}$	$90.37_{\pm0.34}$	$82.76_{\pm1.02}$	$63.73_{\pm0.38}$	$88.44_{\pm0.65}$	$74.41_{\pm0.33}$	$25.59_{\pm0.43}$	$64.70_{\pm0.38}$	$57.58_{\pm0.20}$	$72.77_{\pm3.42}$	$69.58_{\pm0.25}$	$67.98_{\pm0.69}$
TaskRes(e) CVPR'23[40]		$58.02_{\pm0.26}$	$90.49_{\pm0.41}$	$83.24_{\pm1.08}$	$64.69_{\pm0.53}$	$89.38_{\pm0.81}$	$74.46_{\pm0.15}$	$25.89_{\pm0.47}$	$64.83_{\pm0.10}$	$57.98_{\pm0.22}$	$72.95_{\pm3.35}$	$69.88_{\pm0.43}$	$68.35_{\pm0.71}$
CrossModal-LP CVPR'23[24]		$42.15_{\pm0.21}$	$90.40_{\pm0.25}$	$84.56_{\pm0.94}$	$63.73_{\pm0.38}$	$85.00_{\pm0.59}$	$75.66_{\pm0.36}$	$24.44_{\pm0.41}$	$66.09_{\pm0.42}$	$58.41_{\pm0.20}$	$71.72_{\pm2.42}$	$69.74_{\pm0.55}$	$66.67_{\pm0.59}$
ZS-LP		$57.88_{\pm0.18}$	$90.36_{\pm0.35}$	$82.76_{\pm1.02}$	$63.73_{\pm0.37}$	$88.47_{\pm0.65}$	$74.41_{\pm0.33}$	$25.57_{\pm0.45}$	$64.70_{\pm0.38}$	$57.58_{\pm0.20}$	$72.78_{\pm3.44}$	$69.55_{\pm0.26}$	$67.98_{\pm0.69}$
CLAP		$60.73_{\pm0.20}$	$90.62_{\pm0.46}$	$86.51_{\pm0.32}$	$65.50_{\pm0.26}$	$87.66_{\pm0.85}$	$75.92_{\pm0.16}$	$25.65_{\pm0.67}$	$65.99_{\pm0.31}$	$58.85_{\pm0.06}$	$73.15_{\pm2.34}$	$69.88_{\pm0.26}$	$69.13_{\pm0.54}$
CoOp IJCV'22[46]		$60.39_{\pm0.47}$	$90.28_{\pm0.42}$	$85.36_{\pm1.00}$	$67.64_{\pm0.06}$	$91.27_{\pm0.83}$	$71.58_{\pm0.79}$	$26.63_{\pm0.86}$	$65.77_{\pm0.02}$	$59.69_{\pm0.13}$	$77.08_{\pm2.42}$	$72.71_{\pm0.50}$	$69.85_{\pm0.69}$
PLOT ICLR'23[6]		$61.92_{\pm0.09}$	$91.54_{\pm0.33}$	$87.39_{\pm0.74}$	$67.03_{\pm0.50}$	$92.43_{\pm0.25}$	$75.31_{\pm0.30}$	$26.17_{\pm0.29}$	$67.48_{\pm0.04}$	$61.70_{\pm0.35}$	$78.15_{\pm0.00}$	$74.45_{\pm0.50}$	$71.23_{\pm0.51}$
Zero-Shot ICML'21[30]		$60.35_{\pm0.00}$	$83.81_{\pm0.00}$	$82.86_{\pm0.00}$	$55.69_{\pm0.00}$	$65.94_{\pm0.00}$	$74.85_{\pm0.00}$	$17.16_{\pm0.00}$	$56.80_{\pm0.00}$	$42.32_{\pm0.00}$	$37.53_{\pm0.00}$	$57.47_{\pm0.00}$	$57.71_{\pm0.00}$
Rand. Init LP ICML'21[30]		$45.06_{\pm0.42}$	$84.00_{\pm2.38}$	$65.65_{\pm1.37}$	$58.06_{\pm0.23}$	$87.47_{\pm0.55}$	$59.65_{\pm0.14}$	$27.99_{\pm0.49}$	$57.18_{\pm0.38}$	$55.24_{\pm1.13}$	$66.58_{\pm0.53}$	$60.84_{\pm1.13}$	
CLIP-Adapter IJCV'23[11]		$56.95_{\pm0.24}$	$91.33_{\pm0.24}$	$83.39_{\pm0.51}$	$66.83_{\pm0.80}$	$91.93_{\pm0.40}$	$72.11_{\pm0.19}$	$27.89_{\pm0.65}$	$65.09_{\pm0.21}$	$61.37_{\pm1.25}$	$78.49_{\pm1.47}$	$73.23_{\pm1.46}$	$69.87_{\pm0.69}$
TIP-Adapter ECCV'22[42]		$59.44_{\pm0.14}$	$88.26_{\pm0.33}$	$82.27_{\pm1.21}$	$57.63_{\pm0.51}$	$73.76_{\pm0.31}$	$73.87_{\pm0.34}$	$19.36_{\pm0.41}$	$63.13_{\pm0.25}$	$51.52_{\pm0.27}$	$62.30_{\pm1.19}$	$63.15_{\pm1.02}$	$63.15_{\pm0.54}$
TIP-Adapter(f) ECCV'22[42]	8-shot	$61.80_{\pm0.05}$	$90.53_{\pm0.28}$	$85.60_{\pm0.35}$	$64.42_{\pm0.06}$	$84.33_{\pm0.23}$	$74.95_{\pm0.66}$	$23.79_{\pm0.48}$	$66.97_{\pm0.09}$	$59.81_{\pm0.46}$	$70.34_{\pm4.31}$	$69.33_{\pm1.04}$	$68.35_{\pm0.73}$
TaskRes(r) CVPR'23[40]		$59.10_{\pm0.16}$	$91.62_{\pm0.29}$	$85.77_{\pm0.39}$	$69.29_{\pm0.10}$	$93.94_{\pm0.31}$	$74.52_{\pm0.29}$	$29.58_{\pm0.81}$	$67.07_{\pm0.04}$	$63.18_{\pm0.99}$	$78.55_{\pm3.05}$	$73.06_{\pm0.85}$	$71.43_{\pm0.66}$
TaskRes(e) CVPR'23[40]		$59.12_{\pm0.15}$	$91.94_{\pm0.24}$	$85.74_{\pm0.35}$	$69.65_{\pm0.47}$	$94.29_{\pm0.37}$	$74.36_{\pm0.26}$	$30.91_{\pm0.60}$	$66.31_{\pm0.25}$	$63.48_{\pm0.51}$	$78.83_{\pm2.89}$	$73.64_{\pm0.35}$	$71.66_{\pm0.59}$
CrossModal-LP CVPR'23[24]		$46.81_{\pm0.11}$	$91.76_{\pm0.06}$	$84.56_{\pm0.41}$	$69.29_{\pm0.14}$	$92.87_{\pm0.24}$	$76.12_{\pm0.22}$	$28.27_{\pm0.79}$	$68.20_{\pm0.29}$	$62.61_{\pm0.82}$	$77.73_{\pm2.72}$	$73.55_{\pm0.52}$	$70.36_{\pm0.61}$
ZS-LP		$59.10_{\pm0.19}$	$91.62_{\pm0.29}$	$85.80_{\pm0.40}$	$69.29_{\pm0.12}$	$93.94_{\pm0.29}$	$74.51_{\pm0.28}$	$29.59_{\pm0.82}$	$67.08_{\pm0.04}$	$63.18_{\pm0.99}$	$78.55_{\pm3.04}$	$73.05_{\pm0.88}$	$71.43_{\pm0.67}$
CLAP		$62.98_{\pm0.13}$	$91.45_{\pm0.05}$	$87.75_{\pm0.40}$	$70.35_{\pm0.30}$	$92.06_{\pm0.43}$	$77.42_{\pm0.31}$	$28.97_{\pm0.89}$	$68.61_{\pm0.20}$	$63.24_{\pm0.65}$	$76.66_{\pm2.78}$	$73.34_{\pm0.49}$	$72.08_{\pm0.60}$
CoOp IJCV'22[46]		$61.91_{\pm0.17}$	$91.99_{\pm0.31}$	$87.02_{\pm0.89}$	$73.60_{\pm0.19}$	$94.49_{\pm0.40}$	$74.48_{\pm0.15}$	$31.43_{\pm0.96}$	$68.36_{\pm0.66}$	$62.51_{\pm0.25}$	$83.69_{\pm0.47}$	$76.90_{\pm0.50}$	$73.33_{\pm0.42}$
PLOT ICLR'23[6]		$63.01_{\pm0.13}$	$92.24_{\pm0.38}$	$87.21_{\pm0.40}$	$72.80_{\pm0.75}$	$94.76_{\pm0.34}$	$77.09_{\pm0.18}$	$31.49_{\pm0.89}$	$69.96_{\pm0.24}$	$65.60_{\pm0.82}$	$82.23_{\pm0.91}$	$77.26_{\pm0.64}$	$73.94_{\pm0.54}$
Zero-Shot ICML'21[30]		$60.35_{\pm0.00}$	$83.81_{\pm0.00}$	$82.86_{\pm0.00}$	$55.69_{\pm0.00}$	$65.94_{\pm0.00}$	$74.85_{\pm0.00}$	$17.16_{\pm0.00}$	$56.80_{\pm0.00}$	$42.32_{\pm0.00}$	$37.53_{\pm0.00}$	$57.47_{\pm0.00}$	$57.71_{\pm0.00}$
Rand. Init LP ICML'21[30]		$52.24_{\pm0.10}$	$87.55_{\pm2.32}$	$71.63_{\pm0.95}$	$69.20_{\pm0.62}$	$92.73_{\pm0.58}$	$66.92_{\pm0.47}$	$34.63_{\pm0.89}$	$63.07_{\pm0.10}$	$60.60_{\pm1.08}$	$73.38_{\pm1.38}$	$70.94_{\pm0.44}$	$67.54_{\pm0.81}$
CLIP-Adapter IJCV'23[11]		$59.02_{\pm0.15}$	$92.28_{\pm0.21}$	$84.92_{\pm0.74}$	$73.49_{\pm0.57}$	$94.56_{\pm0.30}$	$73.96_{\pm0.18}$	$34.19_{\pm0.65}$	$68.14_{\pm0.20}$	$65.70_{\pm0.24}$	$83.24_{\pm0.56}$	$77.30_{\pm0.37}$	$73.35_{\pm0.38}$
TIP-Adapter ECCV'22[42]		$57.81_{\pm0.18}$	$88.44_{\pm0.37}$	$81.09_{\pm1.89}$	$58.83_{\pm0.38}$	$78.41_{\pm0.53}$	$72.96_{\pm0.42}$	$21.96_{\pm0.56}$	$64.00_{\pm0.26}$	$54.79_{\pm0.66}$	$67.90_{\pm2.32}$	$64.52_{\pm0.97}$	$64.61_{\pm0.78}$
TIP-Adapter(f) ECCV'22[42]	16-shot	$62.27_{\pm0.13}$	$91.22_{\pm0.43}$	$85.43_{\pm0.59}$	$69.56_{\pm0.36}$	$91.18_{\pm0.27}$	$74.65_{\pm0.34}$	$29.32_{\pm0.64}$	$68.90_{\pm0.05}$	$64.56_{\pm0.32}$	$76.55_{\pm1.44}$	$71.81_{\pm0.70}$	$71.40_{\pm0.48}$
TaskRes(r) CVPR'23[40]		$61.01_{\pm0.11}$	$93.00_{\pm0.07}$	$86.28_{\pm0.35}$	$75.48_{\pm0.10}$	$95.82_{\pm0.29}$	$75.86_{\pm0.04}$	$34.82_{\pm0.89}$	$69.72_{\pm0.03}$	$66.45_{\pm0.33}$	$83.15_{\pm0.14}$	$76.54_{\pm0.42}$	$74.38_{\pm0.25}$
TaskRes(e) CVPR'23[40]		$60.85_{\pm0.12}$	$93.09_{\pm0.17}$	$86.28_{\pm0.21}$	$75.38_{\pm0.03}$	$96.14_{\pm0.37}$	$75.43_{\pm0.07}$	$36.53_{\pm0.31}$	$68.43_{\pm0.11}$	$65.88_{\pm0.28}$	$83.70_{\pm0.38}$	$76.96_{\pm0.07}$	$74.42_{\pm0.19}$
CrossModal-LP CVPR'23[24]		$52.90_{\pm0.10}$	$92.77_{\pm0.04}$	$87.48_{\pm0.14}$	$75.44_{\pm0.33}$	$95.20_{\pm0.23}$	$77.14_{\pm0.16}$	$33.30_{\pm0.26}$	$70.56_{\pm0.12}$	$66.92_{\pm0.56}$	$82.03_{\pm0.99}$	$76.40_{\pm0.26}$	$73.65_{\pm0.29}$
ZS-LP		$61.00_{\pm0.11}$	$92.98_{\pm0.09}$	$86.27_{\pm0.33}$	$75.49_{\pm0.09}$	$95.82_{\pm0.29}$	$75.86_{\pm0.04}$	$34.82_{\pm0.86}$	$69.72_{\pm0.04}$	$66.43_{\pm0.37}$	$83.16_{\pm0.13}$	$76.54_{\pm0.42}$	$74.37_{\pm0.25}$
CLAP		$65.02_{\pm0.06}$	$91.93_{\pm0.18}$	$88.51_{\pm0.16}$	$75.12_{\pm0.21}$	$94.21_{\pm0.13}$	$78.55_{\pm0.07}$	$33.59_{\pm0.86}$	$70.78_{\pm0.05}$	$66.41_{\pm0.74}$	$80.07_{\pm0.39}$	$76.29_{\pm0.21}$	$74.57_{\pm0.28}$

Figure 8. **Efficient transfer learning Results**. Performance comparison of relevant literature, and the proposed methods for few-shot efficient transfer learning from ResNet-50 CLIP to 11 downstream datasets, using from 1 to 16 shots per class. Average results are depicted in te top-left corner. Full numerical results are introduced in Supp. Tab. 9.

Table 10. **Domain generalization results.** Adapters are adjusted on ImageNet using 16 shots per class, and evaluated at out-of-distribution generalization on 4 ImageNet shifts with multiple CLIP visual backbones. Bold indicates best performance. Relative improvements are obtained for each adapter with respect to no adaptation, *i.e.*, zero-shot prediction.

Method	Visual Backbone	Source Imagenet	Target -V2	-Sketch	-A	-R	Avg.
Zero-Shot ICML'21[30]	ResNet-50	60.35	51.49	33.33	**21.67**	55.93	40.61
Rand. Init LP ICML'21[30]		52.24$_{(-8.11)}$↓	41.85	15.93	10.72	29.95	24.61$_{(-16.00)}$↓
CLIP-Adapter IJCV'23[11]		59.02$_{(-1.33)}$↓	48.15	14.63	15.75	46.29	31.21$_{(-9.40)}$↓
TIP-Adapter ECCV'22[42]		57.81$_{(-2.54)}$↓	50.32	33.59	21.88	56.98	40.69$_{(+0.08)}$↑
TIP-Adapter(f) ECCV'22[42]		62.27$_{(+1.92)}$↑	53.99	33.75	20.48	57.22	41.36$_{(+0.75)}$↑
TaskRes(e) CVPR'23[40]		60.85$_{(+0.50)}$↑	**56.47**	32.80	19.90	55.93	41.28$_{(+0.67)}$↑
ZS-LP		61.00$_{(+0.65)}$↑	51.09	27.90	16.95	50.37	36.58$_{(-4.03)}$↓
CLAP		**65.02**$_{(+4.67)}$↑	56.09	**34.55**	21.52	**59.48**	**42.91**$_{(+2.30)}$↑
Zero-Shot ICML'21[30]	ResNet-101	62.66	54.86	38.69	28.01	64.44	46.50
Rand. Init LP ICML'21[30]		57.51$_{(-5.15)}$↓	44.96	22.61	16.00	40.20	30.94$_{(-15.56)}$↓
CLIP-Adapter IJCV'23[11]		61.87$_{(-0.79)}$↓	52.87	32.49	21.74	54.91	40.50$_{(-6.00)}$↓
TIP-Adapter ECCV'22[42]		60.83$_{(-1.83)}$↓	53.24	38.64	28.88	65.07	46.46$_{(+0.04)}$↓
TIP-Adapter(f) ECCV'22[42]		65.13$_{(+2.47)}$↑	56.48	38.64	26.48	64.57	46.54$_{(+0.04)}$↑
TaskRes(e) CVPR'23[40]		66.10$_{(+3.44)}$↑	56.56	36.76	24.75	61.52	44.90$_{(-1.60)}$↓
ZS-LP		63.79$_{(+1.13)}$↑	53.74	33.64	22.89	58.07	42.09$_{(-4.41)}$↓
CLAP		**67.93**$_{(+5.27)}$↑	**58.98**	**40.68**	28.35	**67.10**	**48.78**$_{(+2.28)}$↑
Zero-Shot ICML'21[30]	ViT-B/32	63.74	54.81	40.84	29.64	66.03	47.83
Rand. Init LP ICML'21[30]		56.87$_{(-6.87)}$↓	46.50	23.44	16.64	41.13	31.93$_{(-15.90)}$↓
CLIP-Adapter IJCV'23[11]		62.70$_{(-1.04)}$↓	52.93	34.27	23.58	57.58	42.09$_{(-5.74)}$↓
TIP-Adapter ECCV'22[42]		47.71$_{(-16.03)}$↓	39.99	23.31	20.02	44.47	31.95$_{(-15.88)}$↓
TIP-Adapter(f) ECCV'22[42]		45.65$_{(-18.09)}$↓	38.00	22.47	12.40	27.44	25.08$_{(-22.75)}$↓
TaskRes(e) CVPR'23[40]		65.18$_{(+1.44)}$↑	55.39	36.54	25.97	61.93	44.96$_{(-2.87)}$↓
ZS-LP		64.02$_{(+0.28)}$↑	53.61	34.93	24.06	60.72	43.33$_{(-4.50)}$↓
CLAP		**68.33**$_{(+4.59)}$↑	**58.38**	**41.27**	29.91	**68.61**	**49.54**$_{(+1.71)}$↑
Zero-Shot ICML'21[30]	ViT-B/16	68.71	60.76	46.18	47.76	73.98	57.17
Rand. Init LP ICML'21[30]		62.95$_{(-5.76)}$↓	52.48	29.22	29.40	50.54	40.41$_{(-16.76)}$↓
CLIP-Adapter IJCV'23[11]		68.46$_{(-0.25)}$↓	59.55	39.88	38.83	64.62	50.72$_{(-6.45)}$↓
TIP-Adapter ECCV'22[42]		53.81$_{(-14.90)}$↓	45.69	29.21	36.04	55.26	41.55$_{(-15.62)}$↓
TIP-Adapter(f) ECCV'22[42]		51.71$_{(-17.00)}$↓	43.07	27.13	27.04	45.07	35.58$_{(-21.59)}$↓
TaskRes(e) CVPR'23[40]		70.84$_{(+2.13)}$↑	62.15	43.76	43.91	71.59	55.35$_{(-1.82)}$↓
ZS-LP		69.73$_{(+1.02)}$↑	60.40	41.63	41.94	70.64	53.65$_{(-3.52)}$↓
CLAP		**73.38**$_{(+4.67)}$↑	**65.00**	**48.35**	**49.53**	**77.26**	**60.04**$_{(+2.87)}$↑

Table 11. **Exploring the proper constraint value in CLAP.** Full numerical performance for the ablation experiment regarding the initial configuration of the Lagrangian multipliers in the class-adaptive Linear Probing. Results using ResNet-50 as the backbone averaged across 3 random seeds.

Method	Setting	ImageNet	Caltech101	OxfordPets	StanfordCars	Flowers102	Food101	FGVCAircraft	SUN397	DTD	EuroSAT	UCF101	Average
ZS-LP		57.91±0.25	87.98±0.13	77.96±2.57	55.24±0.46	79.62±0.47	70.60±0.35	20.30±0.81	61.00±0.19	46.59±1.51	55.57±1.71	61.34±0.73	61.28±0.83
CLAP(Constant-w=1)		59.74±0.18	88.68±0.43	84.23±0.69	58.39±0.29	75.77±0.55	74.44±0.36	20.64±0.09	61.47±0.02	49.45±1.49	59.21±1.82	64.64±0.57	63.33±0.59
CLAP(ClassWise - avgCorrected)	1-shot	59.02±0.24	88.44±0.20	84.29±0.67	57.75±0.43	79.36±0.52	73.59±0.08	20.76±0.20	61.19±0.26	48.09±1.26	59.85±2.61	63.02±1.09	63.21±0.69
CLAP(Constant-w=ZS)		58.80±0.21	88.65±0.34	83.65±0.90	56.61±0.28	78.17±0.30	73.82±0.35	20.67±0.70	61.24±0.16	48.96±1.49	59.77±1.48	63.83±0.53	63.11±0.61
CLAP(ClassWise)		58.50±0.24	88.38±0.25	83.64±1.18	56.35±0.40	79.90±0.46	73.00±0.14	20.62±0.55	61.15±0.18	47.46±1.15	59.21±0.82	62.48±0.99	62.79±0.58
ZS-LP		57.85±0.04	89.26±0.21	80.56±1.58	60.69±0.42	84.46±0.19	72.94±0.35	23.18±0.36	62.61±0.20	51.79±1.77	63.06±3.14	67.23±0.54	64.88±0.80
CLAP(Constant-w=1)		61.29±0.07	89.74±0.07	85.26±0.59	62.02±0.40	77.60±0.25	75.77±0.23	22.29±0.49	64.03±0.15	52.36±1.37	63.72±0.88	67.96±0.41	65.64±0.45
CLAP(ClassWise - avgCorrected)	2-shot	60.42±0.16	89.70±0.13	85.18±0.69	61.87±0.21	83.37±0.32	75.46±0.28	22.61±0.39	63.67±0.09	53.45±0.88	64.07±0.58	68.27±0.67	66.19±0.40
CLAP(Constant-w=ZS)		59.94±0.09	89.99±0.19	85.04±0.45	61.58±0.40	81.88±0.13	75.37±0.24	23.09±0.38	63.71±0.17	53.96±1.64	65.85±1.28	68.19±0.54	66.24±0.52
CLAP(ClassWise)		58.50±0.24	89.79±0.15	84.93±0.66	61.40±0.38	84.22±0.35	74.94±0.24	23.21±0.32	63.31±0.13	53.05±1.03	65.63±1.49	67.77±0.53	66.07±0.50
ZS-LP		57.88±0.18	90.36±0.35	82.76±1.02	63.73±0.37	88.47±0.65	74.41±0.33	25.57±0.45	64.70±0.38	57.58±0.20	72.78±3.44	69.55±0.26	67.98±0.69
CLAP(Constant-w=1)		62.51±0.11	90.51±0.29	86.51±0.23	64.84±0.41	79.81±0.35	76.36±0.10	23.30±0.48	65.83±0.28	55.79±0.88	68.37±1.37	68.60±0.49	67.49±0.45
CLAP(ClassWise - avgCorrected)	4-shot	61.96±0.15	90.41±0.33	86.67±0.18	65.67±0.19	86.36±0.86	76.27±0.17	24.67±0.42	66.12±0.32	57.98±0.46	69.10±1.36	69.56±0.44	68.62±0.44
CLAP(Constant-w=ZS)		61.35±0.21	90.62±0.33	86.31±0.25	65.61±0.39	85.05±0.75	76.04±0.17	25.46±0.54	66.38±0.29	58.51±0.51	73.37±2.25	69.71±0.60	68.95±0.57
CLAP(ClassWise)		60.73±0.20	90.62±0.46	86.51±0.32	65.50±0.26	87.66±0.85	75.92±0.16	25.65±0.67	65.99±0.31	58.85±0.06	73.15±2.34	69.88±0.26	69.13±0.54
ZS-LP		59.10±0.19	91.62±0.29	85.80±0.40	69.29±0.12	93.94±0.29	74.51±0.28	29.59±0.82	67.08±0.04	63.18±0.99	78.55±3.04	73.05±0.88	71.43±0.67
CLAP(Constant-w=1)		63.83±0.07	90.87±0.18	87.15±0.39	66.95±0.24	81.65±0.14	77.42±0.22	23.60±0.24	67.24±0.31	58.18±0.36	69.06±1.41	70.42±0.27	68.76±0.35
CLAP(ClassWise - avgCorrected)	8-shot	63.69±0.12	91.01±0.05	87.50±0.44	68.95±0.08	90.26±0.12	77.49±0.29	25.70±0.65	67.93±0.17	61.21±0.35	70.02±1.44	72.01±0.26	70.52±0.35
CLAP(Constant-w=ZS)		63.41±0.11	91.21±0.07	87.49±0.41	69.99±0.16	88.20±0.19	77.46±0.26	29.20±0.56	68.66±0.24	62.79±0.79	76.51±2.80	72.79±0.68	71.61±0.57
CLAP(ClassWise)		62.98±0.13	91.45±0.05	87.75±0.40	70.35±0.30	92.06±0.43	77.42±0.31	28.97±0.89	68.61±0.20	63.24±0.65	76.66±2.78	73.34±0.49	72.08±0.60
ZS-LP		61.00±0.11	92.98±0.09	86.27±0.33	75.49±0.09	95.82±0.29	75.86±0.04	34.82±0.86	69.72±0.04	66.43±0.30	83.16±0.13	76.54±0.42	74.37±0.25
CLAP(Constant-w=1)		64.76±0.03	91.18±0.16	87.64±0.11	68.97±0.31	82.45±0.17	78.06±0.08	24.43±0.40	68.13±0.04	59.06±0.46	71.46±0.68	71.00±0.15	69.74±0.24
CLAP(ClassWise - avgCorrected)	16-shot	65.02±0.04	91.51±0.20	88.11±0.26	71.55±0.19	92.08±0.03	78.28±0.07	27.80±0.18	69.32±0.13	63.20±0.63	72.56±0.16	74.20±0.23	72.15±0.19
CLAP(Constant-w=ZS)		65.29±0.04	91.81±0.20	88.07±0.05	73.99±0.27	89.61±0.15	78.46±0.09	33.50±0.77	70.46±0.11	65.48±0.52	79.95±0.22	74.61±0.39	73.75±0.26
CLAP(ClassWise)		65.02±0.06	91.93±0.18	88.51±0.16	75.12±0.21	94.21±0.13	78.55±0.07	33.59±0.86	70.78±0.05	66.41±0.74	80.07±0.38	76.07±0.21	74.57±0.28