

# LP++: A Surprisingly Strong Linear Probe for Few-Shot CLIP

## Supplementary Material

### 5. Proof of Prop. 1

Let us start with  $H_{\mathbf{w}}$ , the Hessian matrix of objective function (3) with respect to block  $\mathbf{w} \in \mathbb{R}^{KD}$ , with block  $\alpha$  being fixed. Let us introduce vector  $\mathbf{F}_{ik} \in \mathbb{R}^{KD}$ , which takes the following form:

$$\mathbf{F}_{ik} = [0, \dots, 0, \mathbf{f}_i^t, 0, \dots, 0]^t \quad (17)$$

where embedding vector  $\mathbf{f}_i \in \mathbb{R}^D$  corresponds to the elements of  $\mathbf{F}_{ik}$  going from the  $((k-1)D+1)^{\text{th}}$  to the  $(kD)^{\text{th}}$  position, with the rest of elements of  $\mathbf{F}_{ik}$  all being equal to zero. Block-coordinatewise  $KD \times KD$  Hessian matrix  $H_{\mathbf{w}}$  reads:

$$\begin{aligned} H_{\mathbf{w}} &= \frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K p_{ik} \mathbf{F}_{ik} \mathbf{F}_{ik}^t \\ &\quad \underbrace{\hspace{10em}}_{\mathbf{D}_i \otimes (\mathbf{f}_i \mathbf{f}_i^t)} \\ &- \frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K \sum_{k'=1}^K p_{ik} p_{ik'} \mathbf{F}_{ik} \mathbf{F}_{ik'}^t \\ &\quad \underbrace{\hspace{10em}}_{(\mathbf{p}_i \mathbf{p}_i^t) \otimes (\mathbf{f}_i \mathbf{f}_i^t)} \\ &= \frac{1}{N} \sum_{i=1}^N (\mathbf{D}_i - \mathbf{p}_i \mathbf{p}_i^t) \otimes (\mathbf{f}_i \mathbf{f}_i^t) \end{aligned} \quad (18)$$

where  $\otimes$  denotes the Kronecker matrix product,  $\mathbf{D}_i$  is the diagonal  $K \times K$  matrix whose elements are given by  $p_{ik}$ :

$$\mathbf{D}_i = \text{Diag}((p_{ik})_{1 \leq k \leq K})$$

and  $\mathbf{p}_i \in \mathbb{R}^K$  is the vector whose elements are  $p_{ik}$ :

$$\mathbf{p}_i = (p_{ik})_{1 \leq k \leq K}$$

Now, let us show the following intermediate result on diagonally dominant, symmetric matrices  $(\mathbf{D}_i - \mathbf{p}_i \mathbf{p}_i^t)$ , which we could establish using the well-known Gershgorin circle theorem.

**Lemma 5.1.** *Let  $\lambda$  denotes an eigen value of matrix  $(\mathbf{D}_i - \mathbf{p}_i \mathbf{p}_i^t)$ , then  $\lambda \in [0, \frac{1}{2}]$ . Therefore,  $(\mathbf{D}_i - \mathbf{p}_i \mathbf{p}_i^t)$  is positive semidefinite and its maximum eigen value verifies:*

$$\lambda_{\max}(\mathbf{D}_i - \mathbf{p}_i \mathbf{p}_i^t) \leq \frac{1}{2}$$

*Proof.* For  $j \in 1, \dots, K$ , let  $R_j$  denotes the sum of the absolute values of the non-diagonal entries in the  $j^{\text{th}}$  row of

a  $K \times K$  matrix. For  $\mathbf{D}_i - \mathbf{p}_i \mathbf{p}_i^t$ , observe the following, due to probability simplex constraint  $\sum_{k=1}^K p_{ik} = 1$ :

$$R_j = p_{ij} \sum_{k \neq j} p_{ik} = p_{ij}(1 - p_{ij}) = p_{ij} - p_{ij}^2$$

Also, notice that  $R_j$  is equal to the diagonal element of the  $j^{\text{th}}$  row, which means matrix  $\mathbf{D}_i - \mathbf{p}_i \mathbf{p}_i^t$  is diagonally dominant. The Gershgorin circle theorem states that, for any eigen value  $\lambda$  of a matrix, there exists at least a row  $j$ , so that  $\lambda$  is within the disc centered at the diagonal element of the row and whose radius is  $R_j$ . Using this fact for matrix  $\mathbf{D}_i - \mathbf{p}_i \mathbf{p}_i^t$  means that, for any eigen  $\lambda$  of the matrix, there exists at least one row  $j$  verifying:

$$|\lambda - (p_{ij} - p_{ij}^2)| \leq R_j$$

This yields:

$$0 \leq \lambda \leq 2(p_{ij} - p_{ij}^2) \leq \frac{1}{2}$$

The last inequality above is due to  $0 \leq p_{ij} - p_{ij}^2 \leq \frac{1}{4}$  for  $p_{ij} \in [0, 1]$ . Since all the eigen values of matrix  $\mathbf{D}_i - \mathbf{p}_i \mathbf{p}_i^t$  are in  $[0, \frac{1}{2}]$ , then the maximum eigen value is also within this interval.  $\square$

From Lemma 5.1, it is straightforward to see that matrix  $\mathbf{D}_i - \mathbf{p}_i \mathbf{p}_i^t - \frac{1}{2} \mathbf{I}_K$  is negative semidefinite and, therefore, its Kronecker multiplication by  $\mathbf{f}_i \mathbf{f}_i^t$  is also negative semidefinite (as  $\mathbf{f}_i \mathbf{f}_i^t$  is positive semidefinite). Hence, the following matrix is negative semi-definite, as it is the sum of negative semidefinite matrices:

$$H_{\mathbf{w}} - \frac{1}{N} \sum_{i=1}^N \frac{1}{2} \mathbf{I}_K \otimes (\mathbf{f}_i \mathbf{f}_i^t) \quad (19)$$

Therefore:

$$\mathbf{x}^t H_{\mathbf{w}} \mathbf{x} \leq \mathbf{x}^t \left( \frac{1}{N} \sum_{i=1}^N \frac{1}{2} \mathbf{I}_K \otimes (\mathbf{f}_i \mathbf{f}_i^t) \right) \mathbf{x} \quad \forall \mathbf{x} \quad (20)$$

This yields:

$$\max_{\|\mathbf{x}\|=1} \mathbf{x}^t H_{\mathbf{w}} \mathbf{x} \leq \max_{\|\mathbf{x}\|=1} \mathbf{x}^t \left( \frac{1}{N} \sum_{i=1}^N \frac{1}{2} \mathbf{I}_K \otimes (\mathbf{f}_i \mathbf{f}_i^t) \right) \mathbf{x} \quad (21)$$

Now, recall the following variational characterization of the maximum eigenvalue of a matrix, following the min-max theorem, also referred to as the variational principle:

$$\lambda_{\max}(H) = \max_{\|\mathbf{x}\|=1} \mathbf{x}^t H \mathbf{x}$$

Using this variational characterization and Eq. (21), we obtain:

$$\lambda_{max}(H_{\mathbf{w}}) \leq \frac{1}{2N} \lambda_{max} \left( \sum_{i=1}^N \mathbf{I}_K \otimes (\mathbf{f}_i \mathbf{f}_i^t) \right) \quad (22)$$

Observe that  $KD \times KD$  matrix  $\sum_{i=1}^N \mathbf{I}_K \otimes (\mathbf{f}_i \mathbf{f}_i^t)$  is block-diagonal with each of its  $D \times D$  diagonal blocks corresponding to matrix  $\sum_{i=1}^N \mathbf{f}_i \mathbf{f}_i^t$ . Therefore, we have:

$$\lambda_{max} \left( \sum_{i=1}^N \mathbf{I}_K \otimes (\mathbf{f}_i \mathbf{f}_i^t) \right) = \lambda_{max} \left( \sum_{i=1}^N \mathbf{f}_i \mathbf{f}_i^t \right)$$

This means that the expression in Eq. (10) in the paper is an upper bound on the maximum eigen value of  $H_{\mathbf{w}}$ , for  $\tau_1 \geq 2$ . Hence, it is a valid block Lipschitz constant for the set of variables in  $\mathbf{w}$ .

It is also possible to have a tighter, but *approximate* Lipschitz constant by omitting the off-diagonal elements in matrices  $\mathbf{D}_i - \mathbf{p}_i \mathbf{p}_i^t$  in the Hessian-matrix expression in Eq. (18). The approximation is motivated by the fact that these matrices are diagonally dominant. This gives the following approximation of the Hessian:

$$\tilde{H}_{\mathbf{w}} = \frac{1}{N} \sum_{i=1}^N \tilde{\mathbf{D}}_i \otimes (\mathbf{f}_i \mathbf{f}_i^t) \quad (23)$$

where  $\tilde{\mathbf{D}}_i$  is the diagonal  $K \times K$  matrix whose elements are given by  $p_{ik} - p_{ik}^2$ ,  $k = 1, \dots, K$ . Since  $p_{ik} - p_{ik}^2 \leq \frac{1}{4}$ , we have:  $\lambda_{max}(\tilde{\mathbf{D}}_i) \leq \frac{1}{4}$ . Therefore, following the same arguments as before, we have:

$$\lambda_{max}(\tilde{H}_{\mathbf{w}}) \leq \frac{1}{4N} \lambda_{max} \left( \sum_{i=1}^N (\mathbf{f}_i \mathbf{f}_i^t) \right) \quad (24)$$

Hence,  $\tau_1 \geq 1$  in Eq. (10) in the paper provides an approximate Lipschitz constant for the block of variables in  $\mathbf{w}$ .

Notice that we still need an eigenvalue decomposition of  $\sum_{i=1}^N \mathbf{f}_i \mathbf{f}_i^t$ . However, for feature embeddings of size 1024, i.e., a  $1024 \times 1024$  matrix, this computation, which has to be performed only once, is manageable; it took 0.3 seconds on a single NVIDIA RTX A6000 GPU.

Now, let us look at the Hessian matrix of objective function (3) w.r.t block  $\boldsymbol{\alpha} \in \mathbf{R}^K$ , with the variables in  $\mathbf{w}$  being fixed. Let vector  $\mathbf{T}_{ik} \in \mathbf{R}^K$  takes the form:

$$\mathbf{T}_{ik} = [0, \dots, 0, \mathbf{f}_i^t \mathbf{t}_k, 0, \dots, 0]^t$$

where the  $k^{\text{th}}$  element is equal to  $\mathbf{f}_i^t \mathbf{t}_k$  and the rest of elements are equal to zero. This block-coordinatewise  $K \times K$

Hessian reads:

$$\begin{aligned} H_{\boldsymbol{\alpha}} &= \frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K p_{ik} \mathbf{T}_{ik} \mathbf{T}_{ik}^t \\ &\quad - \frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K \sum_{k'=1}^K p_{ik} p_{ik'} \mathbf{T}_{ik} \mathbf{T}_{ik'}^t \\ &= \frac{1}{N} \sum_{i=1}^N (\mathbf{D}_i - \mathbf{p}_i \mathbf{p}_i^t) \odot \mathbf{T}_i \end{aligned} \quad (25)$$

where  $\odot$  denotes the Hadamard product and  $\mathbf{T}_i$  is the  $K \times K$  matrix whose element at  $(k, k') \in [1, \dots, K]^2$  is given by:  $\mathbf{f}_i^t \mathbf{t}_k \mathbf{f}_i^t \mathbf{t}_{k'}$ . As before, omitting the off-diagonal elements in Eq. (25) yields an approximate Hessian,  $\tilde{H}_{\boldsymbol{\alpha}}$ , which corresponds to the diagonal  $K \times K$  matrix whose diagonal elements are given by:

$$\frac{1}{N} \sum_{i=1}^N (p_{ik} - p_{ik}^2) (\mathbf{f}_i^t \mathbf{t}_k)^2, k = 1, \dots, K$$

Thus, using  $p_{ik} - p_{ik}^2 \leq \frac{1}{4}$ , we obtain for  $\tau_2 \geq 1$ :

$$\lambda_{max}(\tilde{H}_{\boldsymbol{\alpha}}) \leq \max_k \frac{\tau_2}{4N} \sum_{i=1}^N (\mathbf{f}_i^t \mathbf{t}_k)^2 \quad (26)$$

This gives the approximate Lipschitz constant in Eq. (11).

Finally, it is well-known that, for two blocks of variables with Lipschitz constants  $\gamma_{\mathbf{w}}$  and  $\gamma_{\boldsymbol{\alpha}}$ , the following is a Global Lipschitz constant (See [1]):

$$\gamma = 2 \max(\gamma_{\mathbf{w}}, \gamma_{\boldsymbol{\alpha}}) \quad (27)$$

Therefore, for  $\tau \geq 2$ , the expression in Eq. (12) provides an approximate global Lipschitz constant.

## 6. Proof of Proposition 2

Consider the following expressions of  $g_1$  and  $g_2$ , for some  $\lambda > 0$ :

$$\begin{aligned} g_1 &= -\frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K y_{ik} \mathbf{f}_i^t (\mathbf{w}_k + \alpha_k \mathbf{t}_k) + \frac{\lambda}{2} \sum_{k=1}^K \|\mathbf{w}_k\|^2 \\ g_2 &= \frac{1}{N} \sum_{i=1}^N \ln \left( \sum_{j=1}^K \exp \mathbf{f}_i^t (\mathbf{w}_j + \alpha_j \mathbf{t}_j) \right) - \frac{\lambda}{2} \sum_{k=1}^K \|\mathbf{w}_k\|^2 \end{aligned} \quad (28)$$

It is easy to verify that  $L = g_1 + g_2$ . Let us now write the gradients of  $g_1$  and  $g_2$  w.r.t  $\mathbf{w}_k$ ,  $k \in [1, \dots, K]$ :

$$\begin{aligned} \frac{\partial g_1}{\partial \mathbf{w}_k} &= -\frac{1}{N} \sum_{i=1}^N y_{ik} \mathbf{f}_i + \lambda \mathbf{w}_k \\ \frac{\partial g_2}{\partial \mathbf{w}_k} &= \frac{1}{N} \sum_{i=1}^N p_{ik} \mathbf{f}_i - \lambda \mathbf{w}_k \end{aligned} \quad (29)$$

Independently of the value of  $\lambda > 0$ ,  $g_1$  is convex w.r.t each  $\mathbf{w}_k$ ,  $k \in [1, \dots, K]$ , as it is the sum of linear and convex functions. By setting the gradient of  $g_1$  w.r.t  $\mathbf{w}_k$  to 0, we obtain the minimum given in Eq. (13). As for  $g_2$ , we can ensure it is convex w.r.t each  $\mathbf{w}_k$  by setting  $\lambda$  as a function of positive semi-definite (PSD) matrices  $\mathbf{A}_k$ , as given in Prop. 2. This ensures that the Hessian of  $g_2$  is PSD, as  $\mathbf{A}_k$  is the Hessian of the first term in  $g_2$ . Setting the gradient of  $g_2$  w.r.t  $\mathbf{w}_k$  to 0 yields the minimum in Eq. (14).

## 7. More details on the initialization

The technical observations made in the paper in Eq. (13) and Eq. (15) suggest expressions for initializing variables  $\alpha_k$  and  $\mathbf{w}_k$ :

$$\begin{aligned}\alpha_k^0 &= \frac{1}{\beta N} \sum_{i=1}^N y_{ik} \mathbf{f}_i^t \mathbf{t}_k = \frac{1}{\beta N} \tilde{\alpha}_k \\ \mathbf{w}_k^0 &= \frac{1}{\lambda N} \sum_{i=1}^N y_{ik} \mathbf{f}_i = \frac{1}{\lambda N} \tilde{\mathbf{w}}_k\end{aligned}\quad (30)$$

There are two non-negative multiplicative factors,  $\lambda$  and  $\beta$ , which appear in these expressions. Here, we provide more details on how to set  $\lambda$  and  $\beta$  systematically (without hyper-parameter search).

Consider the logit scores for each sample  $j$ , according to our image-text linear probe model:

$$l_{j,k} = \frac{1}{\lambda N} \mathbf{f}_j^t \tilde{\mathbf{w}}_k + \frac{1}{\beta N} \tilde{\alpha}_k \mathbf{f}_j^t \mathbf{t}_k \propto \mathbf{f}_j^t \tilde{\mathbf{w}}_k + \frac{\lambda}{\beta} \tilde{\alpha}_k \mathbf{f}_j^t \mathbf{t}_k \quad (31)$$

Notice that multiplying  $\lambda$  and  $\beta$  by the same non-negative multiplicative factor does not change the order of the class scores, i.e., we only need to set a value for the ratio  $\frac{\lambda}{\beta}$ , which balances in Eq. (31) the contribution of the text-embedding term vs. the contribution of the visual-prototype term. We set  $\frac{\lambda}{\beta}$  systematically (without additional hyper-parameters) as a monotonically decreasing function of the number of support samples ( $S = \frac{N}{K}$ ):

$$\frac{\lambda}{\beta} = \frac{250}{S} \quad (32)$$

The choice in Eq. (32) is motivated by two intuitive reasons. First, the text-embedding scores are at least one-order of magnitude smaller than the visual-embedding scores. Second, a monotonically decreasing function of  $S$  makes sense: the smaller the number of visual support samples, the bigger the weight given to the text-knowledge. Finally, by setting  $\frac{\lambda}{\beta}$  according to Eq. (32) and  $\lambda = \frac{1}{N}$ , we initialize the vari-

ables as follows:

$$\begin{aligned}\alpha_k^0 &= \frac{250}{S} \sum_{i=1}^N y_{ik} \mathbf{f}_i^t \mathbf{t}_k \\ \mathbf{w}_k^0 &= \sum_{i=1}^N y_{ik} \mathbf{f}_i\end{aligned}\quad (33)$$

## 8. More details on the training-free version of our method

Here we gave more details on the training-free version of our method LP++. To be specific, for a given test sample, we can compute the class scores (logits) as follows:

$$l_{\text{test},k} = \mathbf{f}_{\text{test}}^t (\mathbf{w}_k^0 + \alpha_k^0 \mathbf{t}_k) \quad \forall k$$

where  $\mathbf{f}_{\text{test}}$  is the vision embedding of the test sample and the initial, hyper-parameter free variables  $\alpha_k^0$  and  $\mathbf{w}_k^0$  are set according to Eq. (33). Then, the class of test sample  $\mathbf{f}_{\text{test}}^t$  is predicted as:

$$\hat{k} = \arg \max_k l_{\text{test},k}$$

## 9. Detailed explanation of the flaw in the experimental evaluation of TIP-adapter [28]

By examining the official GitHub repository<sup>6</sup>, we found that the original implementation of TIP-adapter [28] uses large test-data performance feedback for choosing the hyper-parameters and best model. Therefore, we re-evaluated TIP-adapter using the same small validation sets as those used for all the competing methods, for fairness. Indeed, the paper [28] does not provide details as to the hyper-parameter selection and use of validation data. According to the GitHub repository, the authors mentioned that key hyper-parameters ( $\alpha$  and  $\beta$  in [28]), which control the trade-off between the text and vision supervision, are both set to 1 as tuning baseline (see Issue #13 in the GitHub). They also mentioned that  $\alpha$  has to be increased when the domain gap between the CLIP pre-trained model and the downstream task is large. Starting from some initial values of the hyper-parameters, with  $\beta$  set to 1 and  $\alpha \in [1, 10]$  but predetermined specifically for each dataset, the authors deployed an additional search function after the adaptation procedure. **This function finds the best combination of  $\alpha$  and  $\beta$  on either the test set for ImageNet or the entire validation set for the other 10 datasets, which invalidates the few-shot assumption.** Furthermore, the grid search intervals for  $\alpha$  and  $\beta$  vary significantly among the different datasets (and sometimes do not even overlap), e.g.,  $\alpha \in [1.17, 7]$  for ImageNet whereas  $\alpha \in [10, 50]$  for Flowers102.

<sup>6</sup><https://github.com/gaopengcuhk/Tip-Adapter/issues/13>

Number of shots ( $S$ )	1	2	4	8	16
BMM ( $\text{iter}_w=10, \text{iter}_\alpha=1$ )	<b>63.43</b>	<b>66.20</b>	<b>69.16</b>	<b>72.04</b>	<b>74.42</b>
BMM ( $\text{iter}_w=10, \text{iter}_\alpha=10$ )	62.85	66.01	69.12	71.98	74.33
BCGD	62.92	66.04	69.03	71.94	74.24
GD (our Lipschitz cst)	63.04	66.14	69.15	71.99	74.28

Table 5. Comparison of different block-cycling strategies using our initialization. BCGD corresponds to BMM with both  $\text{iter}_w$  and  $\text{iter}_\alpha$  equal to 1. GD (our Lipschitz cst) corresponds to the single-block version of LP++. Results are averaged over 11 datasets.

Through our experiments, we found that the initial value of  $\alpha$  plays a substantial role in the final performance of the Tip-Adapter-F model. Hence, for a fair comparison, we re-evaluated Tip-Adapter-F\* (Table 1) by (i) finding the initial value of  $\alpha_{\text{init}} \in [1, 10]$  on the small validation set deployed for all methods, and (ii) setting  $\beta_{\text{init}} = 1$ . Then, we maintain the implementation of the original publication, and perform the final grid search after adaptation. Instead of employing a data-set specific range for these two hyper-parameters, we use  $[\alpha_{\text{init}}, 50]$  for  $\alpha$  and  $[\beta_{\text{init}}, 50]$  for  $\beta$  during grid search. Although these grid-search steps add significantly to the overall computational overhead of TIP-adapter, they are needed to achieve performances close to those reported in the paper [28] (refer to Fig. 1 for a comparison between Tip-Adapter-F and Tip-Adapter-F\*).

As for the training-free version of Tip-adapter [28], which we evaluated in Table 3, we fixed these hyper-parameters, i.e.,  $\alpha = \beta = 1$ , for a fair comparison with our training-free LP++. Indeed, one could argue that searching for these hyper-parameters on large validation sets would invalidate the claim that the method is “training-free”.

In addition to the mentioned flaws, the original implementation of [28] finds the best model using the performance on the test set. In our re-implementation, we addressed this issue by applying early stopping on the small validation set.

## 10. Ablation on the block-cycling strategies

Table 5 reports the performances for different block-cycling strategies in our Block Majorize-Minimize (BMM) procedure. For example, BMM ( $\text{iter}_w=10, \text{iter}_\alpha=1$ ) corresponds to the LP++ version in Alg. 1, whose performance is reported in the paper. GD corresponds to the single-block MM version of LP++, and BCGD (Block Coordinate Gradient Descent) corresponds the block-wise version of LP++ with both  $\text{iter}_w$  and  $\text{iter}_\alpha$  equal to 1.

## 11. Method ranking

In Fig. 5, we show the ranking of different methods using Autorank [12] (on the 11 datasets); lower value indicates better performance.



Figure 5. The ranking of different methods using Autorank [12]. The results are averaged over 11 datasets.

## 12. Implementation of the L-BFGS optimizer

When assessing the standard LP as in Table 6, we follow the instructions provided in [29], where the L-BFGS optimizer is implemented by NumPy, and used to optimize the standard CE loss jointly with L2 regularizers. As for the ablation study in Table 2, all the optimizers are implemented by Pytorch, including L-BFGS<sup>7</sup>. In this case, we did not use regularizers, for fair comparisons.

## 13. Details of prompt design

We follow the same strategy as in [28] for the text prompts. To be specific, we use the average of 7 templates per class for ImageNet, while for the other datasets, we use a single template. The latter might be different according to the type of the images those datasets include. For example, for Caltech101, we simply use “a photo of a [class<sub>k</sub>]” for each class, while for Food101, we use “a photo of [class<sub>k</sub>], a type of food” for each class.

## 14. Proof Lemma 2.1.

Assume a function  $L(\mathbf{v})$  is  $\gamma$ -smooth, i.e., its gradient is Lipschitz (with a Lipschitz constant  $\gamma > 0$ ):  $\nabla^2 L(\mathbf{v}) \preceq \gamma \mathbf{I}$ , with  $\mathbf{I}$  the identity matrix. From this condition on the Hessian of  $L$ , it is easy to see that the following function is convex:  $G(\mathbf{v}) = \frac{\gamma}{2} \mathbf{v}^t \mathbf{v} - L(\mathbf{v})$ . Using the first-order condition for convexity of  $G$  gives:

$$G(\mathbf{v}) \geq G(\mathbf{v}^j) + \nabla G(\mathbf{v}^j)^t (\mathbf{v} - \mathbf{v}^j) \quad (34)$$

Then using the expression of  $G$  and its gradient in Eq. (34) yield the following quadratic upper bound on  $L$ , which corresponds to majorizing function  $M(\mathbf{v}, \mathbf{v}^j)$  in the paper:

$$L(\mathbf{v}) \leq L(\mathbf{v}^j) + \nabla L(\mathbf{v}^j)^t (\mathbf{v} - \mathbf{v}^j) + \frac{\gamma}{2} \|\mathbf{v} - \mathbf{v}^j\|^2 \quad (35)$$

By setting the derivative of this convex upper bound to zero, we get the following solution, which corresponds to a specific gradient-descent step:

$$\mathbf{v}^{j+1} = \mathbf{v}^j - \frac{1}{\gamma} \nabla L(\mathbf{v}^j) = \arg \min_{\mathbf{v}} M(\mathbf{v}, \mathbf{v}^j) \quad (36)$$

<sup>7</sup><https://github.com/hjmshi/PyTorch-LBFGS>

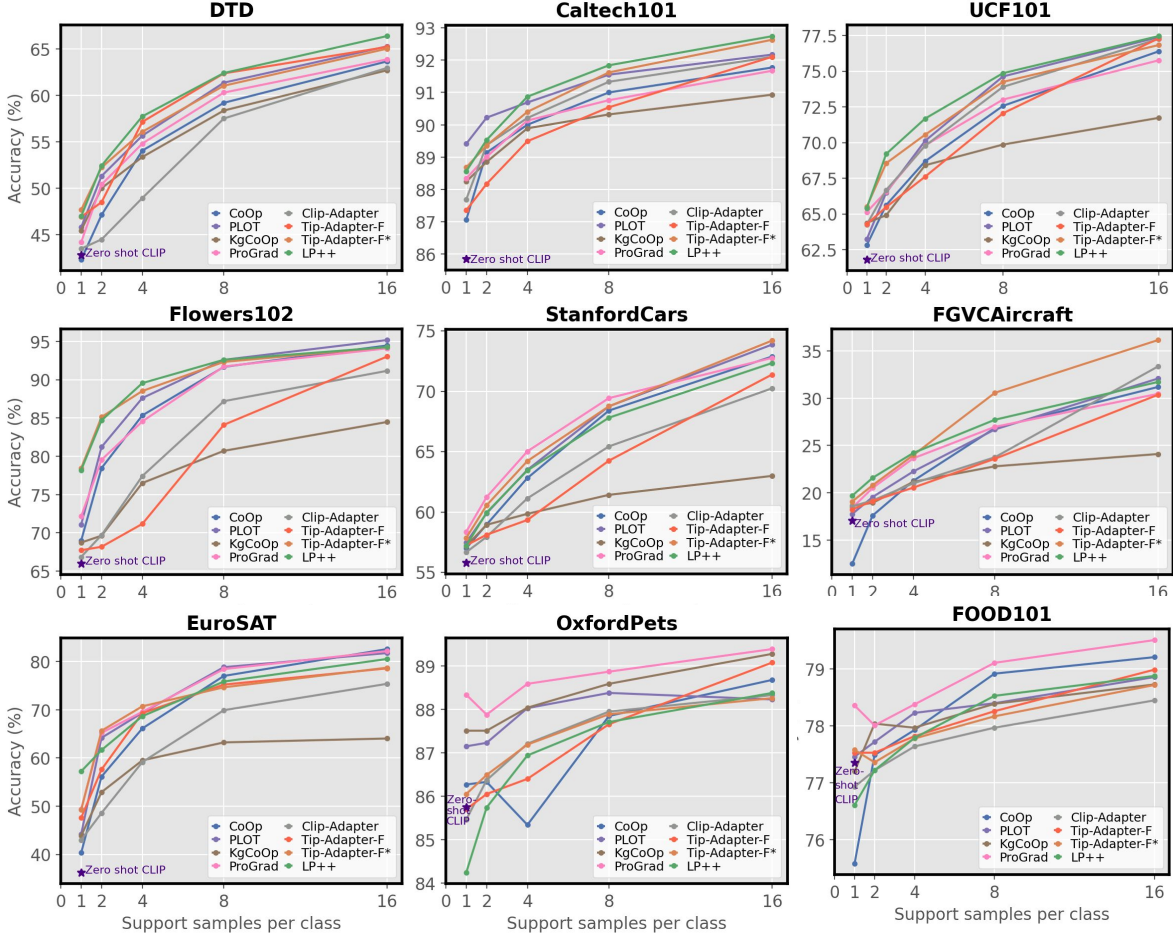


Figure 6. Comparison of LP++ with different adaptation methods on 9 benchmarks, which are averaged over 10 tasks.

Applying Eq. (35) to  $\mathbf{v} = \mathbf{v}^{j+1}$ , and plugging the gradient step in Eq. (36) into Eq. (35), one gets the following decrease of objective  $L$ , by at least  $\frac{1}{2\gamma} \|\nabla L(\mathbf{v}^j)\|^2$ :

$$L(\mathbf{v}^{j+1}) \leq L(\mathbf{v}^j) - \frac{1}{2\gamma} \|\nabla L(\mathbf{v}^j)\|^2 \quad (37)$$

## 15. Proof Theorem 2.2.

The standard result in Theorem 2.2 provides, for a convex,  $\gamma$ -smooth function, the convergence rate of gradient steps with step sizes  $\frac{1}{\gamma}$ . As  $L$  is convex, we have:

$$L(\mathbf{v}^j) \leq L(\mathbf{v}^*) + \nabla L(\mathbf{v}^j)^t (\mathbf{v}^j - \mathbf{v}^*) \quad (38)$$

Plugging Eq. (38) in Eq. (37) yields:

$$L(\mathbf{v}^{j+1}) - L(\mathbf{v}^*) \quad (39)$$

$$\leq \nabla L(\mathbf{v}^j)^t (\mathbf{v}^j - \mathbf{v}^*) - \frac{1}{2\gamma} \|\nabla L(\mathbf{v}^j)\|^2$$

$$= \frac{\gamma}{2} (\|\mathbf{v}^j - \mathbf{v}^*\|^2 - \|\mathbf{v}^j - \frac{1}{\gamma} \nabla L(\mathbf{v}^j) - \mathbf{v}^*\|^2)$$

$$= \frac{\gamma}{2} (\|\mathbf{v}^j - \mathbf{v}^*\|^2 - \|\mathbf{v}^{j+1} - \mathbf{v}^*\|^2) \quad (40)$$

Summing both parts of the inequality in Eq. (39), we get:

$$\sum_{j=0}^{J-1} (L(\mathbf{v}^{j+1}) - L(\mathbf{v}^*)) \quad (41)$$

$$\leq \frac{\gamma}{2} \sum_{j=0}^{J-1} (\|\mathbf{v}^j - \mathbf{v}^*\|^2 - \|\mathbf{v}^{j+1} - \mathbf{v}^*\|^2)$$

$$= \frac{\gamma}{2} (\|\mathbf{v}^0 - \mathbf{v}^*\|^2 - \|\mathbf{v}^J - \mathbf{v}^*\|^2)$$

$$\leq \frac{\gamma}{2} (\|\mathbf{v}^0 - \mathbf{v}^*\|^2) \quad (42)$$

Finally, using the fact that  $L$  is decreasing at each iteration, along with inequality Eq. (41), we have:

$$\begin{aligned}
 L(\mathbf{v}^J) - L(\mathbf{v}^*) &\leq \frac{1}{J} \sum_{j=0}^{J-1} (L(\mathbf{v}^{j+1}) - L(\mathbf{v}^*)) \\
 &\leq \frac{\gamma}{2J} (\|\mathbf{v}^0 - \mathbf{v}^*\|^2) \quad (43)
 \end{aligned}$$

## 16. Additional numerical results on 11 benchmarks

Tab. 6 shows the average accuracy and standard deviation over 10 tasks for LP++ and the baseline methods on 11 benchmarks. The per-dataset performance across all the methods and for different number of shots is depicted in Fig. 6, which complements Fig. 3 in the main paper. In these plots, one may observe that LP++ (*green curve*) yields the best overall performance in 6 out of 11 data sets. Furthermore, LP++ outperforms all adapter-based methods in nearly all the remaining datasets and scenarios (i.e., when varying the number of shots). Only in a few cases, e.g., FGVCAirCRAFT and OxfordPets, Tip-Adapter-F\* yields a performance better than LP++. Indeed, the largest improvement gains are observed in FGVCAirCRAFT, where the accuracy obtained by all methods is considerably small. We believe that, in this particular scenario, the more intensive grid search carried out by Tip-Adapter-F\* has a greater impact on its results. Considering prompt learning strategies, ProGrad is the top-performing approach, which obtains the best overall results in several datasets. Nevertheless, in addition to the limitations already highlighted for these methods (i.e., time complexity and non-suitability for black-box adaptation), the large standard deviations observed across runs (see also Fig. 1), make these approaches highly unreliable, specially in the low-labeled data regime. As discussed, this could be due to the fact that these approaches learn prompts, through the text encoder, which are “too specialized” for a given image support set.

Dataset	Number of shots ( $S$ )	1	2	4	8	16
ImageNet	Zero-shot CLIP [ICML'21] <sup>[23]</sup>	60.35				
	CoOp [ICV'22] <sup>[29]</sup>	61.19 ± 0.17	61.58 ± 0.40	62.22 ± 0.22	62.87 ± 0.21	63.70 ± 0.13
	PLOT [ICLR'23] <sup>[4]</sup>	60.46 ± 0.34	60.73 ± 0.60	61.79 ± 0.39	62.48 ± 0.32	63.08 ± 0.26
	KgCoOp CVPR'23 <sup>[27]</sup>	60.90 ± 0.16	61.44 ± 0.15	62.00 ± 0.11	62.20 ± 0.15	62.43 ± 0.11
	ProGrad [ICCV'23] <sup>[30]</sup>	<b>61.58 ± 0.27</b>	<b>62.14 ± 0.13</b>	<b>62.59 ± 0.09</b>	63.04 ± 0.11	63.54 ± 0.18
	CLIP-Adapter [ICV'23] <sup>[9]</sup>	59.82 ± 0.11	59.94 ± 0.05	59.97 ± 0.04	59.98 ± 0.09	61.31 ± 0.39
	Tip-Adapter-F ECCV'22 <sup>[28]</sup>	60.59 ± 0.14	61.42 ± 0.05	62.12 ± 0.06	63.41 ± 0.07	<b>65.06 ± 0.04</b>
	Tip-Adapter-F* ECCV'22 <sup>[28]</sup>	60.98 ± 0.15	61.23 ± 0.11	61.72 ± 0.25	62.84 ± 0.10	64.03 ± 0.12
	Standard LP [ICML'21] <sup>[23]</sup>	22.21 ± 0.31	31.96 ± 0.25	41.48 ± 0.25	49.49 ± 0.16	56.04 ± 0.13
LP++	61.18 ± 0.08	61.56 ± 0.14	62.55 ± 0.12	<b>63.76 ± 0.07</b>	64.73 ± 0.05	
SUN397	Zero-shot CLIP [ICML'21] <sup>[23]</sup>	58.85				
	CoOp [ICV'22] <sup>[29]</sup>	61.79 ± 0.45	63.32 ± 0.47	65.79 ± 0.44	67.89 ± 0.38	70.15 ± 0.32
	PLOT [ICLR'23] <sup>[4]</sup>	62.53 ± 0.30	63.87 ± 0.26	65.85 ± 0.48	67.83 ± 0.36	69.90 ± 0.31
	KgCoOp CVPR'23 <sup>[27]</sup>	<b>62.91 ± 0.59</b>	64.38 ± 0.30	66.06 ± 0.37	66.66 ± 1.10	67.68 ± 0.78
	ProGrad [ICCV'23] <sup>[30]</sup>	62.79 ± 0.50	64.12 ± 0.55	66.32 ± 0.59	68.33 ± 0.28	70.18 ± 0.27
	CLIP-Adapter [ICV'23] <sup>[9]</sup>	60.78 ± 0.16	61.79 ± 0.23	63.84 ± 0.35	66.26 ± 0.14	67.66 ± 1.05
	Tip-Adapter-F ECCV'22 <sup>[28]</sup>	61.02 ± 0.36	62.15 ± 0.28	63.86 ± 0.19	67.25 ± 0.16	70.94 ± 0.13
	Tip-Adapter-F* ECCV'22 <sup>[28]</sup>	62.58 ± 0.22	63.79 ± 0.13	65.49 ± 0.35	67.43 ± 0.11	69.25 ± 0.16
	Standard LP [ICML'21] <sup>[23]</sup>	32.56 ± 0.40	43.77 ± 0.41	54.49 ± 0.39	61.83 ± 0.30	67.03 ± 0.16
LP++	62.47 ± 0.27	<b>64.65 ± 0.25</b>	<b>67.28 ± 0.27</b>	<b>69.34 ± 0.14</b>	<b>71.23 ± 0.07</b>	
DTD	Zero-shot CLIP [ICML'21] <sup>[23]</sup>	42.69				
	CoOp [ICV'22] <sup>[29]</sup>	42.31 ± 1.86	47.13 ± 1.93	54.06 ± 1.49	59.21 ± 0.91	63.67 ± 0.83
	PLOT [ICLR'23] <sup>[4]</sup>	45.82 ± 1.72	51.32 ± 1.61	55.67 ± 1.14	61.38 ± 1.04	65.29 ± 1.05
	KgCoOp CVPR'23 <sup>[27]</sup>	45.46 ± 2.83	50.01 ± 2.71	53.37 ± 0.71	58.38 ± 1.34	62.71 ± 0.92
	ProGrad [ICCV'23] <sup>[30]</sup>	44.19 ± 2.38	50.41 ± 1.74	54.82 ± 1.28	60.31 ± 0.99	63.89 ± 0.88
	CLIP-Adapter [ICV'23] <sup>[9]</sup>	43.49 ± 0.68	44.49 ± 1.07	48.95 ± 0.85	57.52 ± 0.67	62.97 ± 0.60
	Tip-Adapter-F ECCV'22 <sup>[28]</sup>	46.92 ± 1.01	48.50 ± 1.08	57.16 ± 0.53	62.38 ± 0.47	65.23 ± 0.82
	Tip-Adapter-F* ECCV'22 <sup>[28]</sup>	<b>47.68 ± 1.43</b>	<b>52.24 ± 0.74</b>	56.09 ± 0.99	61.05 ± 0.71	65.04 ± 0.21
	Standard LP [ICML'21] <sup>[23]</sup>	29.63 ± 1.53	41.19 ± 1.45	51.72 ± 0.57	58.78 ± 0.52	64.56 ± 0.69
LP++	46.97 ± 1.37	<b>52.44 ± 0.99</b>	<b>57.75 ± 0.82</b>	<b>62.42 ± 0.53</b>	<b>66.40 ± 0.50</b>	
Caltech101	Zero-shot CLIP [ICML'21] <sup>[23]</sup>	85.84				
	CoOp [ICV'22] <sup>[29]</sup>	87.06 ± 1.24	89.14 ± 0.87	90.00 ± 0.63	91.00 ± 0.66	91.77 ± 0.29
	PLOT [ICLR'23] <sup>[4]</sup>	<b>89.41 ± 0.41</b>	<b>90.22 ± 0.25</b>	90.69 ± 0.37	91.55 ± 0.38	92.17 ± 0.30
	KgCoOp CVPR'23 <sup>[27]</sup>	88.24 ± 0.49	88.85 ± 0.43	89.89 ± 0.31	90.32 ± 0.43	90.93 ± 0.26
	ProGrad [ICCV'23] <sup>[30]</sup>	88.34 ± 1.64	89.01 ± 0.61	90.13 ± 0.45	90.76 ± 0.32	91.67 ± 0.39
	CLIP-Adapter [ICV'23] <sup>[9]</sup>	87.69 ± 0.41	89.37 ± 0.29	90.21 ± 0.32	91.33 ± 0.15	92.10 ± 0.20
	Tip-Adapter-F ECCV'22 <sup>[28]</sup>	87.35 ± 0.64	88.17 ± 0.29	89.49 ± 0.25	90.54 ± 0.34	92.10 ± 0.25
	Tip-Adapter-F* ECCV'22 <sup>[28]</sup>	88.68 ± 0.44	89.36 ± 0.59	90.40 ± 0.26	91.62 ± 0.23	92.63 ± 0.21
	Standard LP [ICML'21] <sup>[23]</sup>	68.88 ± 1.68	78.41 ± 0.54	84.91 ± 0.45	88.70 ± 0.40	91.14 ± 0.19
LP++	88.56 ± 0.43	89.53 ± 0.35	<b>90.87 ± 0.19</b>	<b>91.84 ± 0.24</b>	<b>92.73 ± 0.17</b>	
UCF101	Zero-shot CLIP [ICML'21] <sup>[23]</sup>	61.80				
	CoOp [ICV'22] <sup>[29]</sup>	62.80 ± 1.26	65.62 ± 1.09	68.69 ± 0.76	72.57 ± 0.80	76.39 ± 0.54
	PLOT [ICLR'23] <sup>[4]</sup>	63.22 ± 1.05	66.49 ± 0.92	70.12 ± 0.62	74.63 ± 0.79	77.39 ± 0.53
	KgCoOp CVPR'23 <sup>[27]</sup>	64.37 ± 1.66	64.91 ± 1.01	68.41 ± 0.38	69.86 ± 0.33	71.73 ± 0.78
	ProGrad [ICCV'23] <sup>[30]</sup>	65.13 ± 0.87	66.57 ± 0.62	69.80 ± 0.62	73.01 ± 0.52	75.76 ± 0.47
	CLIP-Adapter [ICV'23] <sup>[9]</sup>	64.25 ± 0.54	66.68 ± 0.31	69.77 ± 0.40	73.90 ± 0.50	77.26 ± 0.39
	Tip-Adapter-F ECCV'22 <sup>[28]</sup>	64.28 ± 0.54	65.48 ± 0.43	67.61 ± 0.28	72.05 ± 0.53	77.30 ± 0.21
	Tip-Adapter-F* ECCV'22 <sup>[28]</sup>	<b>65.50 ± 0.59</b>	68.55 ± 0.45	70.55 ± 0.58	74.25 ± 0.48	76.83 ± 0.24
	Standard LP [ICML'21] <sup>[23]</sup>	40.80 ± 1.05	51.71 ± 0.79	61.64 ± 0.50	68.47 ± 0.44	73.38 ± 0.43
LP++	65.41 ± 0.37	<b>69.20 ± 0.52</b>	<b>71.68 ± 0.41</b>	<b>74.86 ± 0.36</b>	<b>77.46 ± 0.39</b>	
Flowers102	Zero-shot CLIP [ICML'21] <sup>[23]</sup>	65.98				
	CoOp [ICV'22] <sup>[29]</sup>	69.00 ± 2.44	78.47 ± 1.88	85.34 ± 1.69	91.68 ± 0.82	94.47 ± 0.36
	PLOT [ICLR'23] <sup>[4]</sup>	71.09 ± 1.44	81.22 ± 0.92	87.61 ± 0.79	92.60 ± 0.55	<b>95.18 ± 0.40</b>
	KgCoOp CVPR'23 <sup>[27]</sup>	68.73 ± 1.97	69.63 ± 1.25	76.51 ± 0.51	80.71 ± 0.63	84.48 ± 0.70
	ProGrad [ICCV'23] <sup>[30]</sup>	72.16 ± 1.74	79.55 ± 0.88	84.56 ± 1.41	91.73 ± 0.35	94.10 ± 0.41
	CLIP-Adapter [ICV'23] <sup>[9]</sup>	66.86 ± 0.73	69.71 ± 0.46	77.42 ± 0.60	87.20 ± 0.52	91.16 ± 0.23
	Tip-Adapter-F ECCV'22 <sup>[28]</sup>	67.73 ± 0.57	68.18 ± 0.84	71.17 ± 0.67	84.11 ± 0.49	93.02 ± 0.28
	Tip-Adapter-F* ECCV'22 <sup>[28]</sup>	<b>78.46 ± 1.01</b>	<b>85.14 ± 0.81</b>	88.53 ± 0.54	92.33 ± 0.32	94.26 ± 0.38
	Standard LP [ICML'21] <sup>[23]</sup>	56.98 ± 1.56	73.40 ± 0.87	84.38 ± 0.53	91.81 ± 0.34	95.05 ± 0.29
LP++	78.21 ± 1.01	84.69 ± 0.70	<b>89.56 ± 0.45</b>	<b>92.61 ± 0.32</b>	94.26 ± 0.24	
StanfordCars	Zero-shot CLIP [ICML'21] <sup>[23]</sup>	55.78				
	CoOp [ICV'22] <sup>[29]</sup>	57.00 ± 0.93	58.96 ± 0.78	62.81 ± 0.71	68.40 ± 0.61	72.87 ± 0.50
	PLOT [ICLR'23] <sup>[4]</sup>	57.47 ± 0.58	59.89 ± 0.60	63.49 ± 0.80	68.75 ± 0.46	73.86 ± 0.39
	KgCoOp CVPR'23 <sup>[27]</sup>	57.19 ± 0.65	58.94 ± 0.33	59.85 ± 0.51	61.42 ± 0.40	62.99 ± 1.33
	ProGrad [ICCV'23] <sup>[30]</sup>	<b>58.63 ± 0.39</b>	<b>61.23 ± 0.65</b>	<b>65.02 ± 0.78</b>	<b>69.43 ± 0.40</b>	72.76 ± 0.45
	CLIP-Adapter [ICV'23] <sup>[9]</sup>	56.67 ± 0.22	57.94 ± 0.27	61.13 ± 0.30	65.43 ± 0.10	70.24 ± 0.79
	Tip-Adapter-F ECCV'22 <sup>[28]</sup>	57.24 ± 0.23	58.12 ± 0.50	59.34 ± 0.20	64.25 ± 0.19	71.38 ± 0.23
	Tip-Adapter-F* ECCV'22 <sup>[28]</sup>	57.85 ± 0.33	60.55 ± 0.34	64.22 ± 0.52	68.75 ± 0.31	<b>74.19 ± 0.30</b>
	Standard LP [ICML'21] <sup>[23]</sup>	22.94 ± 0.61	35.48 ± 0.51	47.49 ± 0.67	59.34 ± 0.30	69.11 ± 0.18
LP++	57.20 ± 0.65	59.95 ± 0.36	63.44 ± 0.34	67.81 ± 0.24	72.33 ± 0.18	

Table 6. **Comparison to state-of-the-art methods.** Average classification accuracy (%) and standard deviation over 10 tasks for 11 benchmarks, Best values are highlighted in bold.

Dataset	Number of shots ( $S$ )	1	2	4	8	16
FGVCAircraft	Zero-shot CLIP [ICML'21][23]	17.07				
	CoOp [ICV'22][29]	12.50 ± 6.16	17.59 ± 3.70	21.27 ± 2.54	26.85 ± 0.63	31.20 ± 0.40
	PLOT ICLR'23[4]	17.75 ± 1.36	19.55 ± 0.99	22.26 ± 0.89	26.70 ± 0.70	32.09 ± 0.68
	KgCoOp CVPR'23[27]	18.61 ± 0.76	18.93 ± 1.01	21.16 ± 0.82	22.80 ± 0.44	24.10 ± 0.59
	ProGrad [CCV'23][30]	18.41 ± 0.98	20.51 ± 1.11	23.65 ± 0.50	26.98 ± 0.50	30.47 ± 0.76
	CLIP-Adapter [ICV'23][9]	18.56 ± 0.20	19.18 ± 0.28	21.00 ± 0.21	23.76 ± 0.33	33.37 ± 0.23
	Tip-Adapter-F [ECCV'22][28]	18.23 ± 0.19	19.12 ± 0.20	20.55 ± 0.20	23.60 ± 0.29	30.37 ± 0.25
	Tip-Adapter-F <sup>+</sup> [ECCV'22][28]	19.08 ± 0.15	20.79 ± 0.59	23.99 ± 0.57	<b>30.58 ± 0.29</b>	<b>36.16 ± 0.34</b>
	Standard LP [ICML'21][23]	12.66 ± 0.59	16.92 ± 0.56	21.11 ± 0.83	26.53 ± 0.38	32.42 ± 0.54
	LP++	<b>19.69 ± 0.39</b>	<b>21.58 ± 0.46</b>	<b>24.22 ± 0.60</b>	<b>27.73 ± 0.48</b>	<b>31.73 ± 0.44</b>
EuroSAT	Zero-shot CLIP [ICML'21][23]	36.22				
	CoOp [ICV'22][29]	40.36 ± 7.19	56.15 ± 5.82	66.13 ± 3.62	77.02 ± 1.78	<b>82.59 ± 1.00</b>
	PLOT ICLR'23[4]	44.22 ± 9.14	64.19 ± 6.24	69.37 ± 3.26	<b>78.84 ± 1.33</b>	81.76 ± 1.43
	KgCoOp CVPR'23[27]	43.86 ± 9.17	52.92 ± 5.92	59.51 ± 3.46	63.23 ± 3.03	64.04 ± 1.40
	ProGrad [CCV'23][30]	49.37 ± 5.03	65.22 ± 4.01	69.57 ± 2.88	78.44 ± 1.69	82.17 ± 0.98
	CLIP-Adapter [ICV'23][9]	43.00 ± 2.27	48.60 ± 2.76	59.15 ± 2.26	69.92 ± 1.49	75.38 ± 0.78
	Tip-Adapter-F [ECCV'22][28]	47.63 ± 2.64	57.62 ± 1.86	69.30 ± 2.41	75.22 ± 1.32	78.59 ± 1.48
	Tip-Adapter-F <sup>+</sup> [ECCV'22][28]	49.27 ± 2.88	<b>65.66 ± 1.39</b>	<b>70.72 ± 2.73</b>	74.66 ± 3.15	78.73 ± 0.81
	Standard LP [ICML'21][23]	48.29 ± 2.95	56.81 ± 2.93	64.99 ± 3.47	74.56 ± 0.98	80.29 ± 0.90
	LP++	<b>57.23 ± 1.63</b>	61.65 ± 1.66	68.67 ± 2.21	75.86 ± 0.99	80.53 ± 1.05
OxfordPets	Zero-shot CLIP [ICML'21][23]	85.75				
	CoOp [ICV'22][29]	86.27 ± 1.17	86.33 ± 1.13	85.34 ± 1.69	87.85 ± 1.21	88.68 ± 0.71
	PLOT ICLR'23[4]	87.15 ± 0.72	87.23 ± 1.21	88.03 ± 0.49	88.38 ± 0.64	88.23 ± 0.54
	KgCoOp CVPR'23[27]	87.51 ± 0.68	87.51 ± 0.75	88.04 ± 0.46	88.59 ± 0.34	89.28 ± 0.21
	ProGrad [CCV'23][30]	<b>88.34 ± 0.65</b>	<b>87.88 ± 0.69</b>	<b>88.59 ± 0.58</b>	<b>88.87 ± 0.42</b>	<b>89.39 ± 0.47</b>
	CLIP-Adapter [ICV'23][9]	85.46 ± 0.48	86.37 ± 0.25	87.21 ± 0.51	87.95 ± 0.26	88.33 ± 0.33
	Tip-Adapter-F [ECCV'22][28]	85.70 ± 0.16	86.05 ± 0.46	86.40 ± 0.29	87.66 ± 0.28	89.08 ± 0.27
	Tip-Adapter-F <sup>+</sup> [ECCV'22][28]	86.05 ± 0.36	86.49 ± 0.61	87.19 ± 0.36	87.89 ± 0.34	88.26 ± 0.37
	Standard LP [ICML'21][23]	30.62 ± 1.61	42.64 ± 2.03	55.60 ± 0.98	67.32 ± 0.98	76.23 ± 0.38
	LP++	84.24 ± 1.36	85.74 ± 0.56	86.94 ± 0.48	87.71 ± 0.65	88.38 ± 0.61
Food101	Zero-shot CLIP [ICML'21][23]	77.35				
	CoOp [ICV'22][29]	75.58 ± 1.29	77.49 ± 0.41	77.93 ± 0.58	78.92 ± 0.19	79.21 ± 0.36
	PLOT ICLR'23[4]	77.46 ± 0.55	77.72 ± 0.26	78.23 ± 0.25	78.40 ± 0.35	78.86 ± 0.19
	KgCoOp CVPR'23[27]	77.20 ± 0.77	<b>78.04 ± 0.18</b>	77.97 ± 0.28	78.39 ± 0.40	78.73 ± 0.23
	ProGrad [CCV'23][30]	<b>78.36 ± 0.41</b>	78.01 ± 0.70	<b>78.38 ± 0.87</b>	<b>79.11 ± 0.18</b>	<b>79.51 ± 0.23</b>
	CLIP-Adapter [ICV'23][9]	76.93 ± 0.19	77.22 ± 0.15	77.64 ± 0.17	77.97 ± 0.22	78.45 ± 0.14
	Tip-Adapter-F [ECCV'22][28]	77.53 ± 0.14	77.53 ± 0.22	77.82 ± 0.27	78.26 ± 0.22	78.99 ± 0.15
	Tip-Adapter-F <sup>+</sup> [ECCV'22][28]	77.58 ± 0.10	77.36 ± 0.39	77.78 ± 0.15	78.17 ± 0.11	78.72 ± 0.06
	Standard LP [ICML'21][23]	31.59 ± 1.20	44.60 ± 1.03	56.13 ± 0.63	64.45 ± 0.55	70.97 ± 0.19
	LP++	76.61 ± 0.77	77.22 ± 0.55	77.79 ± 0.34	78.53 ± 0.14	78.88 ± 0.19

Table 6. Comparison to state-of-the-art methods (Continued). Average classification accuracy (%) and standard deviation over 10 tasks for 11 benchmarks, Best values are highlighted in bold.