

Tip-Adapter: Training-free Adaption of CLIP for Few-shot Classification

Renrui Zhang^{*1,2}, Wei Zhang^{*1}, Rongyao Fang², Peng Gao^{†1}, Kunchang Li¹,
Jifeng Dai³, Yu Qiao¹, and Hongsheng Li^{2,4}

¹ Shanghai AI Laboratory

² The Chinese University of Hong Kong

³ SenseTime Research

⁴ Centre for Perceptual and Interactive Intelligence (CPII)

{zhangrenrui, gaopeng, qiaoyu}@pjlab.org.cn, hsli@ee.cuhk.edu.hk

Abstract. Contrastive Vision-Language Pre-training, known as CLIP, has provided a new paradigm for learning visual representations using large-scale image-text pairs. It shows impressive performance on downstream tasks by zero-shot knowledge transfer. To further enhance CLIP’s adaption capability, existing methods proposed to fine-tune additional learnable modules, which significantly improves the few-shot performance but introduces extra training time and computational resources. In this paper, we propose a **Training-free** adaption method for CLIP to conduct few-shot classification, termed as **Tip-Adapter**, which not only inherits the training-free advantage of zero-shot CLIP but also performs comparably to those training-required approaches. Tip-Adapter constructs the adapter via a key-value cache model from the few-shot training set, and updates the prior knowledge encoded in CLIP by feature retrieval. On top of that, the performance of Tip-Adapter can be further boosted to be state-of-the-art on ImageNet by fine-tuning the cache model for $10\times$ fewer epochs than existing methods, which is both effective and efficient. We conduct extensive experiments of few-shot classification on 11 datasets to demonstrate the superiority of our proposed methods. Code is released at <https://github.com/gaopengcuhk/Tip-Adapter>.

Keywords: Vision-language learning, few-shot classification, cache model

1 Introduction

Vision and language are two modalities for humans to perceive the surrounding world and perform diverse interactions with the environment. The accuracy of vision tasks, such as classification [35, 22, 26, 13, 42, 17, 70], detection [51, 5, 72, 64, 69, 9] and 3D understanding [47, 68, 63, 67] has been boosted significantly thanks

* Indicates equal contributions, † Indicates corresponding author

Table 1. Comparison of classification accuracy (%) and time efficiency for different methods on 16-shot ImageNet [10], where our proposed Tip-Adapter and Tip-Adapter-F achieve superior accuracy-efficiency trade-off. All experiments are tested with batch size 32 on a single NVIDIA GeForce RTX 3090 GPU. The column in blue records the performance gain relative to Zero-shot CLIP.

Models	Training	Epochs	Time	Accuracy	Gain	Infer. Speed	GPU Mem.
Zero-shot CLIP [48]	Free	0	0	60.33	0	10.22ms	2227MiB
Linear-probe CLIP [48]	Required	-	13min	56.13	-4.20	-	-
CoOp [73]	Required	200	14h 40min	62.95	+2.62	299.64ms	7193MiB
CLIP-Adapter [16]	Required	200	50min	63.59	+3.26	10.59ms	2227MiB
Tip-Adapter	Free	0	0	62.03	+1.70	10.42ms	2227MiB
Tip-Adapter-F	Required	20	5min	65.51	+5.18	10.53ms	2227MiB

to better neural architecture designs [22, 58] and delicately designed frameworks [51, 37, 5, 71]. Language tasks concerning generation and understanding have also been largely improved due to large-scale self-supervised methods, including pre-training by mask prediction [11] and collected web-scale data [49]. As vision and language normally contain complementary information, joint learning of multi-modality representations has been proven to be quite effective on various tasks, such as visual question answering [2, 1, 31], image captioning [65, 27], and referring expression [66]. Different from previous methods that independently learn vision and language representations on separate datasets [1, 40, 56], CLIP [48] proposed to learn transferable visual features from paired natural language supervisions and exerted amazing zero-shot image classification ability. Due to the interplay between language and vision, the encoded visual representations can be used in open-vocabulary recognition without further re-training.

Many follow-up works have proposed to utilize few-shot data to improve CLIP’s adaption capability on downstream tasks. Following the direction of prompt design [4, 38], CoOp [73] fine-tuned the pre-trained CLIP via learnable textual tokens and achieved strong performance on few-shot image classification. Recently, CLIP-Adapter [16] introduced to equip CLIP with a parametric feature adapter, which generates adapted features and combines them with the original CLIP-encoded features via a residual connection. It demonstrated promising performance for few-shot classification without utilizing prompt designs. Although CoOp [73] and CLIP-Adapter [16] have shown powerful capabilities on few-shot classification benchmarks, in comparison with Zero-shot CLIP [48] and Linear-probe CLIP [48], they require more computational resources to fine-tune the newly introduced learnable parameters. Thus, we ask the following question: can we achieve the best of both worlds, which not only takes the advantage of CLIP’s training-free property for zero-shot classification but also enjoys the strong performance of training-required methods for few-shot classification?

To achieve the goal, we propose a **Training-free** adaption method for **CLIP**, named **Tip-Adapter**, which appends the weight-frozen CLIP model with a novel non-parametric adapter. Different from existing methods, ours does not require extra training, but designs the adapter as a query-key cache model [30, 45,

18] from the few-shot dataset. Specifically, Tip-Adapter extracts visual features of few-shot images by CLIP’s visual encoder and transforms their corresponding labels into one-hot encodings. Then, a cache model containing few-shot visual features and one-hot labels is created, which are viewed as paired keys and values.

By the cache model, the training-free construction of Tip-Adapter exhibits great efficiency compared to traditional fine-tuning via Stochastic Gradient Descent (SGD) [32, 39]. During inference, the test image first calculates its feature similarities with cached keys, and then aggregates cached values to form the adapter’s prediction, which can be regarded as retrieving the few-shot knowledge from the cache model. After that, the adapter’s prediction is combined with the original CLIP’s prediction by a residual connection [22]. In this way, Tip-Adapter simultaneously exploits knowledge from both pre-trained CLIP and the few-shot training dataset. Surprisingly, Tip-Adapter without training could perform comparably to the fine-tuned CoOp and CLIP-Adapter. Furthermore, if we unfreeze the cached keys as learnable parameters and further fine-tune them, Tip-Adapter’s performance could be significantly boosted with just a few training epochs. We term this fine-tuned version as **Tip-Adapter-F**, which only requires 20 epochs on ImageNet [10] to be state-of-the-art compared with 200 epochs adopted by CoOp and CLIP-Adapter. In Table 1, we list the comparison between all existing methods of their performance, training time and inference speed for 16-shot classification on ImageNet, which indicates great accuracy-efficiency trade-off of our methods.

The contributions of our paper are summarised below:

1. We propose Tip-Adapter, a training-free adaption method for CLIP, which discards the conventional SGD-based training by directly setting the adapter with a cache model.
2. Unfreezing the keys of cache model as learnable parameters, the fine-tuned Tip-Adapter, named Tip-Adapter-F, achieves state-of-the-art performance with super-fast convergence on ImageNet.
3. We evaluate Tip-Adapter and Tip-Adapter-F on 11 widely-adopted datasets for few-shot classification and conduct extensive ablation studies to demonstrate their characteristics.

2 Related Work

Data-efficient Transfer Learning. The capability of deep neural networks is revealed with the assistance of large-scale and high-quality datasets [35]. However, collecting such data is challenging and expensive due to long-tail distributions, noisy annotations and the increasing labeling cost. Thus, transfer learning is proposed to alleviate this issue, which has become a popular research field. Supervised pre-training on image classification [10] have been widely adopted as a default basis for fine-tuning on downstream tasks (e.g. detection [51] and segmentation [21]). Self-supervised learning, such as MoCo [20] and BYOL [19], further discards the need of supervised signals and builds a contrastive pretext task for robust feature learning. Recently, CLIP [48], DeCLIP [36] and ALIGN [28]

demonstrate that learning from simple contrastive vision-language pairs obtains promising transferable features for zero-shot recognition over diverse datasets. On top of that, CoOp [73], CLIP-Adapter [16] and WiSE-FT [60] significantly improve the CLIP with limited training data by freezing the pre-trained weights and training additive learnable modules. In contrast, our proposed Tip-Adapter aims at directly infusing few-shot supervisions into the pre-trained CLIP model in a training-free manner. By this, the construction of Tip-Adapter is much more efficient for both time and memory, which only requires calculating the features of few-shot training set once and then caches them.

Cache Model. A cache model stores features of training images and their labels as a key-value database. During inference, the feature encoded from a test sample is treated as query to aggregate information from the cache model by similarity-based retrieval [58]. The whole process is non-parametric [33] and involves no parameter update. The cache model has been equipped on various models to boost the performance for vision or language models, including kNN-LMs [30], Unbounded Cache [18], Matching Network [59] and others [43, 53]. Although simple cache model [45] has shown promising results, the huge storage budget for training data is unaffordable to many applications. To reduce such cost, approximate kNN with highly-optimized similarity search system [29] is proposed, which however is slow and error-prone. Different from previous setup with pure vision or language caches, we construct a blended vision-language cache model by CLIP’s contrastive multi-modality pre-training. Importantly, thanks to our few-shot setting with limited training samples, the total cache size is small and the retrieval can be efficiently calculated by two cascaded matrix multiplications. Moreover, the cache model in Tip-Adapter can be learnable and dynamically updated by Stochastic Gradient Descent (SGD), which further improves its performance.

3 Method

In Section 3.1, we first introduce our proposed training-free Tip-Adapter and its fine-tuned variant, Tip-Adapter-F. Then in Section 3.2, we discuss the relations between our approach and previous methods, such as CLIP-Adapter and cache-based networks.

3.1 Training-free Adaption of CLIP

We propose Tip-Adapter, a training-free adaption method to enhance the few-shot classification performance of CLIP. We construct a key-value cache model from the few-shot training set in a non-parametric manner. Surprisingly, with this well-designed cache model, Tip-Adapter without fine-tuning can achieve comparable performance compared to those training-required approaches, including CoOp [73] and CLIP-Adapter [16]. In addition, if training is allowed, Tip-Adapter-F further surpasses state-of-the-art performance by fine-tuning the cached keys with super-fast convergence.

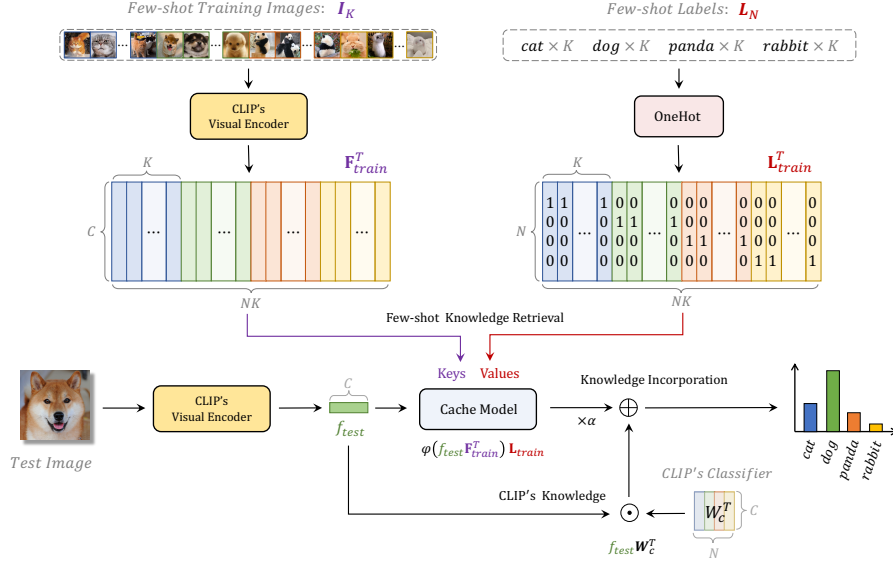


Fig. 1. The Pipeline of Tip-Adapter. Given a K -shot N -class training set, we construct a cache model to adapt CLIP on downstream tasks. It contains few-shot visual features $\mathbf{F}_{\text{train}}^T$ encoded by CLIP and their ground-truth labels $\mathbf{L}_{\text{train}}^T$ under one-hot encodings. After retrieval from the cache model, the few-shot knowledge is incorporated with CLIP’s pre-trained knowledge, achieving the training-free adaption.

Cache Model Construction. Given the pre-trained CLIP [48] model and a new dataset with K -shot N -class training samples for few-shot classification, there are K annotated images in each of the N categories, denoted as I_K with their labels L_N . We aim at creating a key-value cache model as the feature adapter, which contains few-shot knowledge within N classes. For each training image, we utilize the CLIP’s pre-trained visual encoder to extract its C -dimensional L2 normalized feature, and convert its ground-truth label into an N -dimensional one-hot vector. For all NK training samples, we denote their visual features and corresponding label vectors as $\mathbf{F}_{\text{train}} \in \mathbb{R}^{NK \times C}$ and $\mathbf{L}_{\text{train}} \in \mathbb{R}^{NK \times N}$,

$$\mathbf{F}_{\text{train}} = \text{VisualEncoder}(I_K), \quad (1)$$

$$\mathbf{L}_{\text{train}} = \text{OneHot}(L_N). \quad (2)$$

For the key-value cache, the CLIP-encoded representations $\mathbf{F}_{\text{train}}$ are treated as keys, while the one-hot ground-truth vectors $\mathbf{L}_{\text{train}}$ are used as their values. In this way, the key-value cache memorizes all the new knowledge extracted from the few-shot training set, which is for updating the prior knowledge encoded in the pre-trained CLIP.

Tip-Adapter. After constructing the cache model, the adaption of CLIP can be simply achieved by two matrix-vector multiplications. During inference, the L2 normalized feature $f_{\text{test}} \in \mathbb{R}^{1 \times C}$ of the test image is first extracted by the CLIP’s visual encoder and serves as a query for retrieving from the key-value cache. The affinities between the query and keys can be estimated as

$$A = \exp(-\beta(1 - f_{\text{test}}\mathbf{F}_{\text{train}}^T)), \quad (3)$$

where $A \in \mathbb{R}^{1 \times NK}$ and β stands for a modulating hyper-parameter. Since both query and key features are L2 normalized, the term $f_{\text{test}}\mathbf{F}_{\text{train}}^T$ is equivalent to the cosine similarities between test feature f_{test} and all few-shot training features $\mathbf{F}_{\text{train}}^T$. The exponential function is adopted to convert the similarities into non-negative values with β modulating its sharpness. Afterwards, the prediction for cache model can be obtained via linear combination of cached values weighted by the query-key affinities, denoted as $A\mathbf{L}_{\text{train}} \in \mathbb{R}^{1 \times N}$.

Besides the few-shot knowledge retrieved from cache model, the prior knowledge of pre-trained CLIP is calculated by $f_{\text{test}}W_c^T \in \mathbb{R}^{1 \times N}$, where W_c is the weights of CLIP’s classifier generated from its pre-trained textual encoder. By blending both predictions via a residual connection, the output logits of the test image by Tip-Adapter are then calculated as

$$\begin{aligned} \text{logits} &= \alpha A\mathbf{L}_{\text{train}} + f_{\text{test}}W_c^T \\ &= \alpha \varphi(f_{\text{test}}\mathbf{F}_{\text{train}}^T)\mathbf{L}_{\text{train}} + f_{\text{test}}W_c^T, \end{aligned} \quad (4)$$

where α denotes the residual ratio, and we define $\varphi(x) = \exp(-\beta(1 - x))$. Tip-Adapter’s prediction therefore contains two terms, the former term adaptively summarizes information from the few-shot training dataset, and the latter term preserves the prior knowledge from the CLIP’s classifier W_c^T . The two terms are balanced by the weight α . Empirically, α is set to be large if the domain gap between pre-trained and downstream few-shot tasks is large, since more knowledge from the few-shot set is required, and small otherwise.

Tip-Adapter with Fine-tuning. Tip-Adapter can greatly boost CLIP by incorporating new knowledge in the few-shot training set. However, given more shots, Tip-Adapter without training gradually lags behind the training-required CoOp and CLIP-Adapter. To mitigate the gap while preserve the efficiency, we propose Tip-Adapter-F, which treats the keys in the cache model as a good initialization for learnable parameters, and fine-tunes them via SGD. Thanks to the advantageous starting point of cache model, Tip-Adapter-F achieves state-of-the-art performance with only 20-epoch fine-tuning on ImageNet [10], compared with CoOp and CLIP-Adapter’s 200-epoch training.

More specifically, we unfreeze the cached keys $\mathbf{F}_{\text{train}}$, but still freeze the values $\mathbf{L}_{\text{train}}$ and the two encoders of pre-trained CLIP. The intuition is that updating the keys in the cache model can boost the estimation of affinities, which is able to calculate the cosine similarities between the test and training images more accurately. In contrast, values in the cache model are one-hot encodings representing ground-truth annotations and shall be kept frozen to well memorize the category information.

3.2 Relations with Previous Models

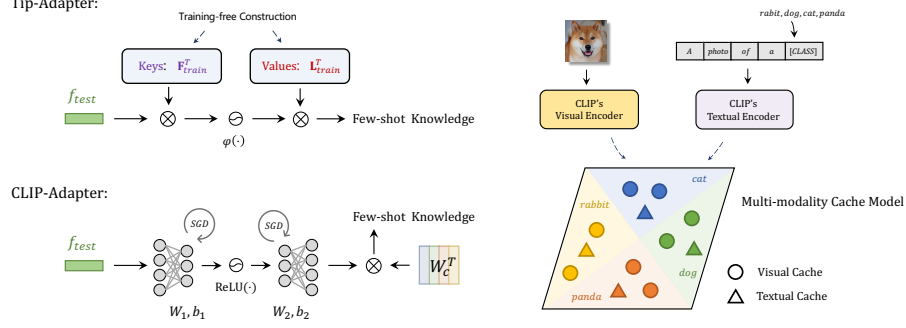


Fig. 2. Comparison of Tip-Adapter and CLIP-Adapter [16] to acquire few-shot knowledge. Tip-Adapter retrieves from the constructed cache model, but CLIP-Adapter encodes the knowledge by the learnable adapter and obtains it aided by CLIP’s classifier W_c .

Fig. 3. The multi-modality cache model of Tip-Adapter. Different from previous networks only with visual cache, Tip-Adapter caches both visual and textual knowledge by CLIP’s encoders.

Relations with CLIP-Adapter. Following the adapter [25] in neural language processing, CLIP-Adapter [16] appends a lightweight two-layer Multi-Layer Perceptron (MLP) to the pre-trained weight-fixed CLIP model and optimizes its parameters via SGD. Specifically, for an input test image, its visual feature f_{test} is first obtained by CLIP’s pre-trained visual encoder. Then, the MLP-based adapter with randomly initialized parameters W_1, b_1, W_2, b_2 , is appended to output the adapted feature,

$$f_{test}^a = \varphi(f_{test}W_1^T + b_1)W_2^T + b_2, \quad (5)$$

where φ denotes the activation function in the MLP. Afterwards, the adapted feature f_{test}^a is linearly combined with the pre-trained CLIP’s feature f_{test} , and output the final classification logits with a hyper-parameter $\alpha \in [0, 1]$,

$$\text{logits} = \alpha f_{test}^a W_c^T + f_{test} W_c^T, \quad (6)$$

where W_c^T is the weights of CLIP’s classifier. The first terms in both Eqs. (4) and (6) represent the ways of Tip-Adapter and CLIP-Adapter to obtain the few-shot knowledge, respectively. As shown in Figure 2, Tip-Adapter acquires the knowledge by retrieval from the cache model, but CLIP-Adapter first utilizes the learnable adapter to predict the adapted feature and then multiplies it with CLIP’s W_c^T to form the final knowledge output.

With further analysis for Eqs. (4) and (6), CLIP-Adapter can be seen as a special form of our proposed Tip-Adapter,

$$W_1 = \mathbf{F}_{train}, \quad W_2 = \mathbf{L}_{train} W_c^{-1}, \quad b_1 = 0, \quad b_2 = 0, \quad (7)$$

$$\varphi(x) = \exp(-\beta(1-x)), \quad \text{where } x \in [0, 1]. \quad (8)$$

They have two key differences. Firstly, CLIP-Adapter randomly initializes both keys and values in the cache model as W_1 and W_2 , and learns them via SGD, while Tip-Adapter directly constructs them with cached training features $\mathbf{F}_{\text{train}}$ and one-hot encodings of the ground-truth labels $\mathbf{L}_{\text{train}}$, which are non-parametric and training-free. Secondly, the bottleneck dimension of Tip-Adapter is equal to NK , while, to prevent over-fitting resulted from training, CLIP-Adapter selects a lower-dimensional bottleneck. This indicates that our cache model could better alleviate the over-fitting problem on few-shot datasets, which further releases the fitting power of large-scale pre-trained models. Thirdly, Tip-Adapter introduces the activation function denoted in Eq. (7). As its inputs are the distances in the normalized feature space, it is naturally bounded between 0 and 1. However, for CLIP-Adapter, the common activation function, $\text{ReLU}(\cdot)$, is chosen to handle unbounded inputs. In short, Tip-Adapter obtains a well-performing adapter without training, which is more efficient on few-shot classification.

Relations with Cache-based Networks. Acquiring a cache model from few-shot training data has been explored by many previous methods, including Matching Network [59], Prototypical Networks [53], MAML [15], Relation Network [55] and others [12, 7, 57, 6]. Our models differ from them in two points for both specific methods and experimental settings.

Firstly, previous works only constructs the cache of visual features, but Tip-Adapter adopts a multi-modality heterogeneous cache model with both visual and textual cached features extracted by CLIP, as shown in Figure 3. In detail, the aforementioned cache model with keys $\mathbf{F}_{\text{train}}$ and values $\mathbf{L}_{\text{train}}$ serves as the visual cache, denoted as \mathbf{F}_{vis} and \mathbf{L}_{vis} here. As CLIP’s classifier W_c is calculated from category texts by the textual encoder, $W_c \in \mathbb{R}^{N \times C}$ can be viewed as language features serving as keys \mathbf{F}_{tex} for textual cache. The values of textual cache is then denoted by an identity matrix $\mathbf{L}_{\text{tex}} \in \mathbb{R}^{N \times N}$, since W_c respectively encodes N category knowledge and each of its row vector corresponds to a certain category. From this perspective, Eq. (4) is reformulated as

$$\text{logits} = \alpha \varphi(f_{\text{test}} \mathbf{F}_{\text{vis}}^T) \mathbf{L}_{\text{vis}} + (f_{\text{test}} \mathbf{F}_{\text{tex}}^T) \mathbf{L}_{\text{tex}}, \quad (9)$$

where the two terms represent knowledge retrieval from both visual and textual cached knowledge.

Secondly, prior works split the same dataset into three sub-sets of different categories, which respectively serve as training, support, and query sets. Although they test on query sets with a new set of categories, it is still within the same semantic domain. In contrast, Tip-Adapter adapts the pre-trained CLIP into a totally new dataset for evaluation, which generalizes to a new domain and thus more challenging. Importantly, we test our models on full test sets, the same as conventional methods [22, 13] trained by the full training set. Compared to existing works [59, 53] on the small query sets, our effectiveness is verified by much more test images of new categories.




Fig. 4. Few-shot classification accuracy of different models on ImageNet [10].

Few-shot Setup	1	2	4	8	16
Zero-shot CLIP [48]: 60.33					
Linear-probe CLIP [48]	22.17	31.90	41.20	49.52	56.13
CoOp [73]	57.15	57.81	59.99	61.56	62.95
CLIP-Adapter [16]	61.20	61.52	61.84	62.68	63.59
Tip-Adapter	60.70	60.96	60.98	61.45	62.03
Tip-Adapter-F	61.32	61.69	62.52	64.00	65.51
	+0.62	+0.73	+1.54	+2.55	+3.48

Table 2. Classification accuracy (%) on ImageNet [10] of different models with quantitative values. The last row in blue records the performance gain of Tip-Adapter-F brought by further fine-tuning over Tip-Adapter.

4 Experiments

4.1 Training Settings

We conduct experiments for Tip-Adapter and Tip-Adapter-F on 11 widely-used image classification datasets: ImageNet [10], StanfordCars [34], UCF101 [54], Caltech101 [14], Flowers102 [44], SUN397 [62], DTD [8], EuroSAT [23], FGV-CAircraft [41], OxfordPets [46], and Food101 [3]. For few-shot learning, we compare the performance of 1, 2, 4, 8, 16 few-shot training sets, and test on the full test sets. For the CLIP backbone, we utilize ResNet-50 [22] as the visual encoder and a transformer [13] as the textual encoder. We obtain the pre-trained weights of both encoders from [48] and freeze them during training. We follow the data preprocessing protocol in CLIP [48], which is composed of random cropping, resizing, and random horizontal flip. Other than the learnable prompts in CoOp, we follow CLIP to adopt prompt ensembling especially on ImageNet and use single handcrafted prompt on other 10 datasets. The Tip-Adapter-F is fine-tuned using batch size 256, learning rate 0.001, and the AdamW [32] optimizer with a cosine scheduler. We set 100-epoch training for EuroSAT dataset and only 20-epoch training for other 10 datasets.

Performance comparison is conducted between Zero-shot CLIP [48], Linear-probe CLIP [48], CoOp [73] and CLIP-Adapter [16]. Therein, Zero-shot CLIP uses no extra training sample and conducts classification purely by pre-trained knowledge. Linear-probe CLIP trains an additional linear classifier after the weight-frozen CLIP on the few-shot training set. CoOp adopts learnable prompts for training, and we select its best-performing variant for comparison, that is, placing the class token at the end of the 16-token prompts without class-specific contexts. CLIP-Adapter appends a feature adapter [25] to narrow the domain gap between the pre-trained features and downstream tasks. We also report the best-performing variant of CLIP-Adapter with only the learnable visual adapter. We report their official scores in the papers for fair comparison.

Table 3. Classification accuracy (%) of different visual encoders on 16-shot ImageNet [10]. ViT-B/32 and ViT-B/16 denote ViT-Base [13] with the patch size 32×32 and 16×16 , respectively, and RN50 \times 16 denotes ResNet-50 [22] with 16 times more computation [48].

Models	ResNet-50	ResNet-101	ViT-B/32	ViT-B/16	RN50 \times 16
Zero-shot CLIP [48]	60.33	62.53	63.80	68.73	70.94
CoOp [73]	62.95	66.60	66.85	71.92	-
CLIP-Adapter [16]	63.59	65.39	66.19	71.13	-
Tip-Adapter	62.03	64.78	65.61	70.75	72.95
Tip-Adapter-F	65.51	68.56	68.65	73.69	75.81

4.2 Comparison on ImageNet

Performance Analysis. As shown in Figure 4 and Table 2, both Tip-Adapter and Tip-Adapter-F show outstanding performance over other methods. Compared to Zero-shot CLIP, Tip-Adapter consistently surpasses it without any training. When the numbers of training samples are limited, Tip-Adapter greatly exceeds the Linear-probe CLIP by +38.53%, +29.06% in 1-shot and 2-shot settings. With further fine-tuning, Tip-Adapter-F updates the keys in the cache model and achieves the best performance over all methods in all few-shot settings. The performance gain over Tip-Adapter becomes larger as the number of training samples increases, from 1-shot’s +0.62% to 16-shot’s +3.44%. This indicates that the fine-tuning with more training samples enables the network to build a more powerful cache model. In Table 3, we also implement different models with various visual encoders over ResNet [22] and ViT [13] backbones, where our Tip-Adapter-F still performs the best.

Efficiency Comparison. In Table 1, we show the comparison of training time and inference speed for different models. CLIP-Adapter, Tip-Adapter and Tip-Adapter-F are able to cache the textual features from CLIP in the beginning and load them during training or inference, but CoOp adopts learnable prompts, which requires to calculate through the whole textual encoder online for every iteration. Linear-probe CLIP utilizes logistic regression [61], so it cannot measure the training time by epochs and the inference speed on GPU. From the comparison, we observe that CoOp takes the most training time for learning prompts and has a +2.26% performance gain over Zero-shot CLIP. CLIP-Adapter significantly reduces the training time with better performance improvement of +3.26%, but still needs 200-epoch training. Aided by the cache model, Tip-Adapter gains +1.70% improvement but requires no extra training time, which makes it a good trade-off between performance and efficiency. Tip-Adapter-F further reaches state-of-the-art accuracy with only 1/10 of CLIP-Adapter and CoOp’s training epochs, achieving the best of both worlds. As for inference speed and GPU memory consumption [52], our Tip-Adapter and Tip-adapter-F only produce marginal extra latency over Zero-shot CLIP and save much GPU memory compared to CoOp, which are quite efficient for applications.

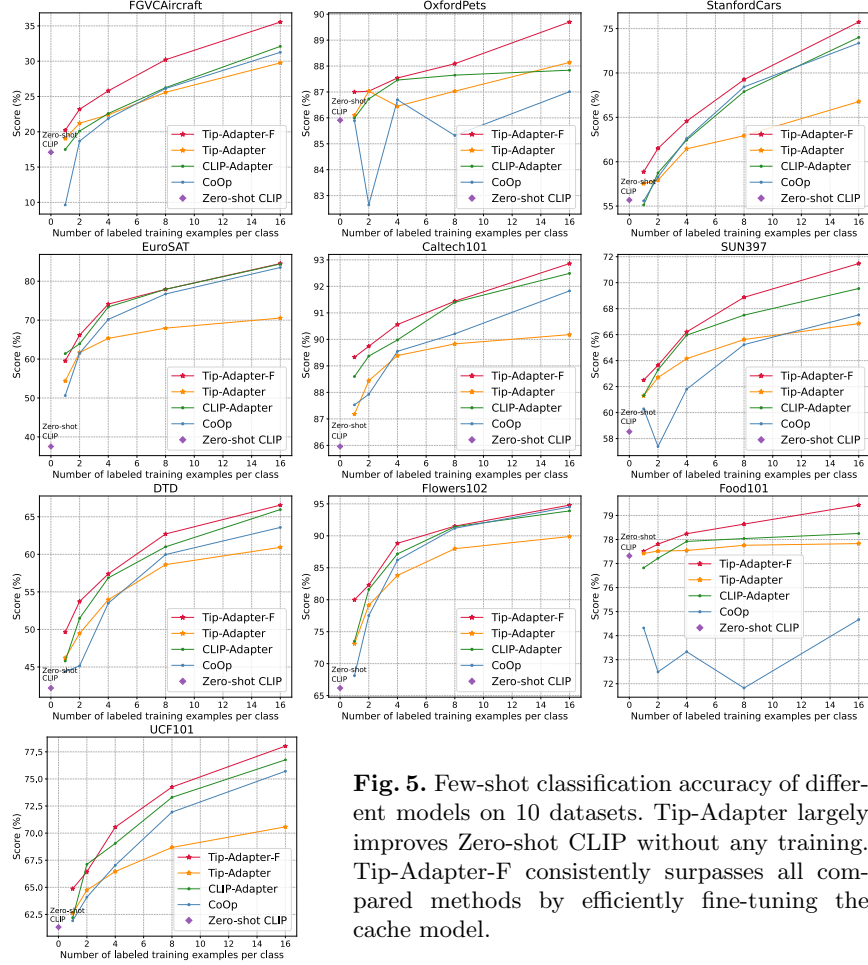


Fig. 5. Few-shot classification accuracy of different models on 10 datasets. Tip-Adapter largely improves Zero-shot CLIP without any training. Tip-Adapter-F consistently surpasses all compared methods by efficiently fine-tuning the cache model.

4.3 Performance on Other Datasets

Figure 5 shows the performance comparison on other 10 datasets listed in Section 4.1. Our training-free Tip-Adapter significantly boosts the classification accuracy over Zero-shot CLIP and surpasses CoOp trained by 1 or 2 shots on most datasets. Although Tip-Adapter underperforms CoOp and CLIP-Adapter trained by more shots, Tip-Adapter-F with a fewer-epoch fine-tuning can eliminate the gap and further surpass all other models, achieving comprehensively leading performance. The consistent superiority of Tip-Adapter-F over 10 datasets fully demonstrates the effectiveness and generality of our proposed cache model.

4.4 Ablation Studies

In this section, we conduct several ablation studies about Tip-Adapter on ImageNet [10]. All experiments adopt the 16-shot setting without training.

Table 4. Four ablation studies (%) of Tip-Adapter on ImageNet [10], from top to bottom: residual ratio α , sharpness ratio β , the size of cache model, and the performance given more shots while fixing cache size 16.

Ablation Studies on Tip-Adapter						
Residual Ratio α	0.0	0.5	1.0	2.0	3.0	4.0
	60.33	61.44	62.03	61.41	60.36	59.14
Sharpness Ratio β	1.5	3.5	5.5	7.5	9.5	11.5
	61.82	61.91	62.03	61.76	61.62	61.40
Cache Size	0	1	2	4	8	16
	60.33	61.45	61.71	61.79	61.83	62.03
More Shots than 16	Shot Setup		16	32	64	128
	Tip-Adapter		62.03	62.51	62.88	63.15
	Tip-Adapter-F		65.47	66.58	67.96	69.74

Residual Ratio α . The hyper-parameter α controls how much to combine newly adapted predictions from the cache model with pre-trained CLIP’s, which can also be interpreted as weighing the visual and textual caches as in Eq. 9. As formulated above, larger α denotes using more knowledge from the few-shot training set and less otherwise. We vary α from 0.0 to 5.0, and set the hyper-parameter β as 5.5. When α equals 0.0, the model is equivalent to Zero-shot CLIP without using few-shot knowledge. From the top part of Table 4, we observe that the classification accuracy is improving as α increases from 0.0 to 1.0, achieving the best 62.03% at 1.0. This indicates that the prior knowledge from CLIP and the few-shot knowledge from cache model are equally important.

Sharpness Ratio β . In Eq. (3), β in the activation function φ controls the sharpness of the affinities. When β is large, only the most similar training samples to the test image in the embedding space have the large influences to the prediction and vice versa. In the second part of Table 4 with the α as 1.0, we observe that the variation of β has a limited impact and a moderate 5.5 for β leads to the best-performing Tip-Adapter.

Size of the Cache Model. We explore the influence of the size for cache model in Tip-Adapter. Given 16-shot training set, rather than caching all 16 samples per category, we construct the cache whose size is more than 0 but less than 16. Taking 8 as an example, we randomly divide 16 samples into 8 uniform groups and obtain 8 prototypes by averaging features of the 2 samples in each group. Considering such random division of samples might influence the performance, we experiment 5 times and report the average scores. The results from the third part of Table 4 illustrate that, the more samples we cache to preserve more few-shot knowledge, the higher accuracy Tip-Adapter can achieve.

Scaling up to More Shots. Given more than 16 shots, we explore a way to still constrain the cache size as 16 and avoid the potential burden for both memory




Fig. 6. Classification performance with different prompt designs: single prompt (Cyan), prompt ensembling (Orange) and learnable prompt (Purple).

Datasets	Source	Target	
	ImageNet [10]	-V2 [50]	-Sketch [24]
Zero-Shot CLIP [48]	60.33	53.27	35.44
Linear Probe CLIP [48]	56.13	45.61	19.13
CoOp [73]	62.95	54.58	31.04
CLIP-Adapter [16]	63.59	55.69	35.68
Tip-Adapter	62.03	54.60	35.90
Tip-Adapter-F	65.51 +3.48	57.11 +2.51	36.00 +0.10

Table 5. The robustness (%) to distribution shift of different methods. The last row in blue records the performance gain of Tip-Adapter-F brought by further fine-tuning over Tip-Adapter..

and computation. Taking 64 shots as an example, following the division strategy in the above paragraph, we obtain 16 prototypes from 4 groups to construct the cache model. The final part of Table 4 indicates that even if the cache size is restrained to 16, Tip-Adapter can well capture the knowledge from 32, 64, and 128 training samples per category. Also, the performance boost gradually slows down when more samples are provided, which implies a possible limit of cache size 16 without training. However, Tip-Adapter-F can break such limit by fine-tuning the keys and achieve better performance by more shots for training.

Prompt Design. We utilize prompt ensembling of 7 templates from [48] for Zero-shot CLIP, CLIP-Adapter, and Tip-Adapter as default. In Figure 6, we test them only using a single prompt, “a photo of a [CLASS].”, and observe slightly worse performance. The accuracy drops are smaller for Tip-Adapter-F and CLIP-Adapter, but larger for Tip-Adapter and Zero-shot CLIP, which indicates the better-performing models are less affected by the prompt variations.

4.5 Distribution Shift

We evaluate the out-of-distribution ability of our proposed Tip-Adapter and Tip-Adapter-F by learning from one dataset but testing on another. We set ImageNet [10] as the source dataset providing 16-shot training set, and adopt two target datasets for testing: ImageNetV2 [50] and ImageNet-Sketch [24], which contain compatible categories to ImageNet but with semantic gaps. As shown in Table 5, Tip-Adapter without training exerts superior robustness to distribution shift, which surpasses CoOp [73] on ImageNet-V2 and CLIP-Adapter [16] on ImageNet-Sketch. This indicates the cache model is more advantageous to out-of-distribution evaluation, whose training-free construction alleviates the risk of over-fitting on the source dataset. Further, Tip-Adapter-F achieves the best of both worlds: the strong out-of-distribution performance brought by cache model and the leading in-distribution ability by fine-tuning.




Fig. 7. t-SNE visualization of $\mathbf{F}_{\text{train}}$ in Tip-Adapter. Dots in different colors stand for embeddings of different categories. From left to right, three distributions indicate the variation of keys in cache model during fine-tuning.

5 Visualization

To better show the variation of cache model during fine-tuning, we utilize t-SNE [48] to visualize the keys $\mathbf{F}_{\text{train}}$ in Figure 7. The dots in different colors denote 10 categories of 16-shot ImageNet [10], and their relative distances reflect the high-dimensional distributions of category embeddings. From left to right, the three sub-figures represent the training-free Tip-Adapter, Tip-Adapter during fine-tuning and the final Tip-Adapter-F, respectively. It could be observed that before training, the distribution has shown good discrimination thanks to the properly designed cache model construction. During fine-tuning, embeddings of the same category gradually converges together and different clusters become more contrastive and separate, contributing to stronger classification capability.

6 Conclusions

We propose Tip-Adapter, a non-parametric adaption method of CLIP, which acquires the adapter by a cache model constructed from the few-shot training set. In this way, the few-shot knowledge is retrieved from the cache model and incorporated with CLIP’s pre-trained knowledge in a training-free manner. On top of that, Tip-Adapter can be further enhanced by fine-tuning the cached keys for just a few epochs, named Tip-Adapter-F, which achieves state-of-the-art performance among existing methods. Considering limitations, although it is marginal, Tip-Adapter-F still requires 20-epoch fine-tuning on ImageNet to learn the best-performing cache model. Our future work will focus on exploring new training-free methods for CLIP to fully unleash its power for visual representation.

Acknowledgement. This work is supported in part by Centre for Perceptual and Interactive Intelligence Limited, in part by the General Research Fund through the Research Grants Council of Hong Kong under Grants (Nos. 14204021, 14207319), in part by CUHK Strategic Fund, and in part by the Shanghai Committee of Science and Technology (Grant No. 21DZ1100100).

References

1. Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., Zhang, L.: Bottom-up and top-down attention for image captioning and visual question answering. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6077–6086 (2018)
2. Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C.L., Parikh, D.: Vqa: Visual question answering. In: Proceedings of the IEEE international conference on computer vision. pp. 2425–2433 (2015)
3. Bossard, L., Guillaumin, M., Van Gool, L.: Food-101—mining discriminative components with random forests. In: European conference on computer vision. pp. 446–461. Springer (2014)
4. Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. arXiv preprint arXiv:2005.14165 (2020)
5. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European Conference on Computer Vision. pp. 213–229. Springer (2020)
6. Chen, W.Y., Liu, Y.C., Kira, Z., Wang, Y.C.F., Huang, J.B.: A closer look at few-shot classification. arXiv preprint arXiv:1904.04232 (2019)
7. Chen, Y., Wang, X., Liu, Z., Xu, H., Darrell, T.: A new meta-baseline for few-shot learning. arXiv preprint arXiv:2003.04390 (2020)
8. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., Vedaldi, A.: Describing textures in the wild. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3606–3613 (2014)
9. Cui, Z., Qi, G.J., Gu, L., You, S., Zhang, Z., Harada, T.: Multitask aet with orthogonal tangent regularity for dark object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2553–2562 (October 2021)
10. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
11. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
12. Dhillon, G.S., Chaudhari, P., Ravichandran, A., Soatto, S.: A baseline for few-shot image classification. arXiv preprint arXiv:1909.02729 (2019)
13. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2021)
14. Fei-Fei, L., Fergus, R., Perona, P.: Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In: 2004 conference on computer vision and pattern recognition workshop. pp. 178–178. IEEE (2004)
15. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: International Conference on Machine Learning. pp. 1126–1135. PMLR (2017)
16. Gao, P., Geng, S., Zhang, R., Ma, T., Fang, R., Zhang, Y., Li, H., Qiao, Y.: Clip-adapter: Better vision-language models with feature adapters. arXiv preprint arXiv:2110.04544 (2021)

17. Gao, P., Ma, T., Li, H., Dai, J., Qiao, Y.: Convmae: Masked convolution meets masked autoencoders. arXiv preprint arXiv:2205.03892 (2022)
18. Grave, E., Cissé, M., Joulin, A.: Unbounded cache model for online language modeling with open vocabulary. arXiv preprint arXiv:1711.02604 (2017)
19. Grill, J.B., Strub, F., Althé, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Doersch, C., Pires, B.A., Guo, Z.D., Azar, M.G., et al.: Bootstrap your own latent: A new approach to self-supervised learning. arXiv preprint arXiv:2006.07733 (2020)
20. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9729–9738 (2020)
21. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 2961–2969 (2017)
22. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
23. Helber, P., Bischke, B., Dengel, A., Borth, D.: Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **12**(7), 2217–2226 (2019)
24. Hendrycks, D., Zhao, K., Basart, S., Steinhardt, J., Song, D.: Natural adversarial examples. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15262–15271 (2021)
25. Houlisby, N., Giurigu, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., Gelly, S.: Parameter-efficient transfer learning for nlp. In: ICML (2019)
26. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861 (2017)
27. Huang, L., Wang, W., Chen, J., Wei, X.Y.: Attention on attention for image captioning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4634–4643 (2019)
28. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q.V., Sung, Y., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: ICML (2021)
29. Johnson, J., Douze, M., Jégou, H.: Billion-scale similarity search with gpus. *IEEE Transactions on Big Data* (2019)
30. Khandelwal, U., Levy, O., Jurafsky, D., Zettlemoyer, L., Lewis, M.: Generalization through memorization: Nearest neighbor language models. arXiv preprint arXiv:1911.00172 (2019)
31. Kim, J.H., Jun, J., Zhang, B.T.: Bilinear attention networks. arXiv preprint arXiv:1805.07932 (2018)
32. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
33. Kossen, J., Band, N., Lyle, C., Gomez, A.N., Rainforth, T., Gal, Y.: Self-attention between datapoints: Going beyond individual input-output pairs in deep learning. arXiv preprint arXiv:2106.02584 (2021)
34. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: Proceedings of the IEEE international conference on computer vision workshops. pp. 554–561 (2013)
35. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS (2012)

36. Li, Y., Liang, F., Zhao, L., Cui, Y., Ouyang, W., Shao, J., Yu, F., Yan, J.: Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. arXiv preprint arXiv:2110.05208 (2021)
37. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
38. Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., Neubig, G.: Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. arXiv preprint arXiv:2107.13586 (2021)
39. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019. OpenReview.net (2019), <https://openreview.net/forum?id=Bkg6RiCqY7>
40. Lu, J., Batra, D., Parikh, D., Lee, S.: Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. arXiv preprint arXiv:1908.02265 (2019)
41. Maji, S., Rahtu, E., Kannala, J., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. arXiv preprint arXiv:1306.5151 (2013)
42. Mao, M., Zhang, R., Zheng, H., Gao, P., Ma, T., Peng, Y., Ding, E., Han, S.: Dual-stream network for visual recognition. arXiv preprint arXiv:2105.14734 (2021)
43. Merity, S., Xiong, C., Bradbury, J., Socher, R.: Pointer sentinel mixture models. arXiv preprint arXiv:1609.07843 (2016)
44. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: 2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing. pp. 722–729. IEEE (2008)
45. Orhan, A.E.: A simple cache model for image recognition. arXiv preprint arXiv:1805.08709 (2018)
46. Parkhi, O.M., Vedaldi, A., Zisserman, A., Jawahar, C.: Cats and dogs. In: 2012 IEEE conference on computer vision and pattern recognition. pp. 3498–3505. IEEE (2012)
47. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 652–660 (2017)
48. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. arXiv preprint arXiv:2103.00020 (2021)
49. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I.: Improving language understanding by generative pre-training (2018)
50. Recht, B., Roelofs, R., Schmidt, L., Shankar, V.: Do imagenet classifiers generalize to imagenet? In: International Conference on Machine Learning. pp. 5389–5400. PMLR (2019)
51. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* **28**, 91–99 (2015)
52. Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning Internal Representations by Error Propagation, p. 318–362. MIT Press, Cambridge, MA, USA (1986)
53. Snell, J., Swersky, K., Zemel, R.S.: Prototypical networks for few-shot learning. arXiv preprint arXiv:1703.05175 (2017)
54. Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402 (2012)

55. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1199–1208 (2018)
56. Tan, H., Bansal, M.: Lxmert: Learning cross-modality encoder representations from transformers. arXiv preprint arXiv:1908.07490 (2019)
57. Tian, Y., Wang, Y., Krishnan, D., Tenenbaum, J.B., Isola, P.: Rethinking few-shot image classification: a good embedding is all you need? In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16. pp. 266–282. Springer (2020)
58. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in neural information processing systems. pp. 5998–6008 (2017)
59. Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. Advances in neural information processing systems **29**, 3630–3638 (2016)
60. Wortsman, M., Ilharco, G., Li, M., Kim, J.W., Hajishirzi, H., Farhadi, A., Namkoong, H., Schmidt, L.: Robust fine-tuning of zero-shot models. arXiv preprint arXiv:2109.01903 (2021)
61. Wright, R.E.: Logistic regression. (1995)
62. Xiao, J., Hays, J., Ehinger, K.A., Oliva, A., Torralba, A.: Sun database: Large-scale scene recognition from abbey to zoo. In: 2010 IEEE computer society conference on computer vision and pattern recognition. pp. 3485–3492. IEEE (2010)
63. Xu, S., Li, Y., Zhao, J., Zhang, B., Guo, G.: Poem: 1-bit point-wise operations based on expectation-maximization for efficient point cloud processing. arXiv preprint arXiv:2111.13386 (2021)
64. Xu, S., Zhao, J., Lu, J., Zhang, B., Han, S., Doermann, D.: Layer-wise searching for 1-bit detectors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5682–5691 (2021)
65. You, Q., Jin, H., Wang, Z., Fang, C., Luo, J.: Image captioning with semantic attention. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4651–4659 (2016)
66. Yu, L., Lin, Z., Shen, X., Yang, J., Lu, X., Bansal, M., Berg, T.L.: Mattnet: Modular attention network for referring expression comprehension. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1307–1315 (2018)
67. Zhang, R., Guo, Z., Gao, P., Fang, R., Zhao, B., Wang, D., Qiao, Y., Li, H.: Pointm2ae: Multi-scale masked autoencoders for hierarchical point cloud pre-training. arXiv preprint arXiv:2205.14401 (2022)
68. Zhang, R., Guo, Z., Zhang, W., Li, K., Miao, X., Cui, B., Qiao, Y., Gao, P., Li, H.: Pointclip: Point cloud understanding by clip. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8552–8562 (2022)
69. Zhang, R., Qiu, H., Wang, T., Xu, X., Guo, Z., Qiao, Y., Gao, P., Li, H.: Monodetr: Depth-aware transformer for monocular 3d object detection. arXiv preprint arXiv:2203.13310 (2022)
70. Zhao, J., Xu, S., Zhang, B., Gu, J., Doermann, D., Guo, G.: Towards compact 1-bit cnns via bayesian learning. International Journal of Computer Vision **130**(2), 201–225 (2022)
71. Zhao, Z., Wu, Z., Zhang, Y., Li, B., Jia, J.: Tracking objects as pixel-wise distributions (2022)

- 72. Zheng, M., Gao, P., Zhang, R., Li, K., Wang, X., Li, H., Dong, H.: End-to-end object detection with adaptive clustering transformer. arXiv preprint arXiv:2011.09315 (2020)
- 73. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. arXiv preprint arXiv:2109.01134 (2021)