

# Noise-Tolerant Few-Shot Unsupervised Adapter for Vision-Language Models

Eman Ali<sup>1, 2</sup>

eman.ali@mbzuai.ac.ae

Muhammad Haris Khan<sup>1</sup>

muhammad.haris@mbzuai.ac.ae

<sup>1</sup> Mohamed Bin Zayed University of  
Artificial Intelligence,  
Abu Dhabi, UAE

<sup>2</sup> Alexandria University  
Alexandria, Egypt

---

## Abstract

Recent advances in large-scale vision-language models have achieved impressive performance in various zero-shot image classification tasks. While prior studies have demonstrated significant improvements by introducing few-shot labelled target samples, they still require labelling of target samples, which greatly degrades their scalability and generalizability while handling various visual recognition tasks. We design NtUA, a Noise-tolerant Unsupervised Adapter that allows the learning of effective target models with few unlabelled target samples. NtUA works as a key-value cache that formulates visual features and predicted pseudo-labels of the few unlabelled target samples as key-value pairs. It consists of two complementary designs. The first is adaptive cache formation that combats pseudo-label noises by weighting the key-value pairs according to their prediction confidence. The second is knowledge-guided cache refinement, which refines pair values (i.e., pseudo-labels) and cache weights by leveraging knowledge distillation from large-scale vision language models. Extensive experiments show that NtUA achieves superior performance consistently across multiple widely adopted benchmarks.

## 1 Introduction

The recent progress in large-scale pretrained vision-language models [1, 2, 3] has significantly advanced image-text relationship modelling. One representative work is CLIP [4], which learns image-text relations by jointly training a visual encoder and a linguistic encoder over web-scale paired image-text data. Thanks to the linguistic diversity of the web data, CLIP can be exploited in various image recognition tasks regardless of the number and nature of image classes. The common procedure is to employ CLIP's linguistic encoder to generate text embeddings of pre-defined class names and then match the text embeddings with the features of test images, which are extracted using CLIP's visual encoder in a zero-shot manner. Although pre-trained CLIP has demonstrated significant effectiveness in image classification tasks, its performance depends heavily on the distribution discrepancy between the image-text pairs in CLIP pretraining and specific classification images in various target domains [5, 6]. Several studies leverage few-shot labelled target samples of each class to adapt the pre-trained CLIP to various target classifi-

cation tasks to mitigate the inter-domain discrepancy [6, 13, 16, 16, 19, 50]. Though these approaches have revealed promising results over several few-shot classification benchmarks, they could end up labelling possibly numerous samples for each class of target domains, which significantly affects their scalability, especially while handling large-scale datasets such as ImageNet [9] with many image classes.

Unsupervised learning can adapt pre-trained CLIP to target classification tasks with unlabelled target samples, which offers a more viable alternative to removing the labelling efforts and significantly improves learning scalability. While recent research has shown some progress in unsupervised adaptation of CLIP across diverse downstream tasks [12, 13, 24, 53], they often need more substantial computational resources [18] or reliance on additional descriptive information for fine-tuning text encoders [24]. Moreover, a prevailing assumption in many of these approaches is the availability of abundantly unlabelled samples in the target domain. However, this assumption may need to consistently align with many realistic settings, particularly in contexts where data collection for the target task presents significant challenges, such as medical applications [8, 43], physiological applications [22], or Image-to-Image translation [21].

Fine-tuning of CLIP model for downstream tasks with a few unlabelled samples is non-trivial and presents significant challenges. First, the model faces the risk of overfitting caused by the scarcity of target domain data [49]. Second, the inherent bias of the CLIP model in generating pseudo-labels can result in noisy labels [13]. This noise tends to increase as the number of samples per shot decreases (Fig. 1), leading to a decline in the adaptation performance (Table. 1 and Table. 4). *To this end, we are the first to explore unsupervised learning for adapting vision-language models towards various downstream classification datasets with a few unlabelled target samples.*

We design NtUA, a Noise-tolerant Unsupervised Adapter that enables robust adaptation of pre-trained vision-language models with a few unlabelled target samples. NtUA achieves noise-tolerant adaptation by generating more accurate pseudo-labels for the few unlabelled target samples. NtUA addresses the challenge posed by the higher noise in pseudo-labels when dealing with limited unlabelled data. Inspired by the adapter idea in supervised methods [4, 14, 23, 44], NtUA introduces a weighted key-value cache, which formulates the CLIP-extracted visual features as keys, the predicted pseudo-labels of few target samples as values, and the corresponding pseudo-label confidence as weights of the key-value pairs (Fig. 2).




Figure 1: A comparison of pseudo-label accuracy across different shots within the DTD dataset illustrates that accuracy tends to increase as the num-

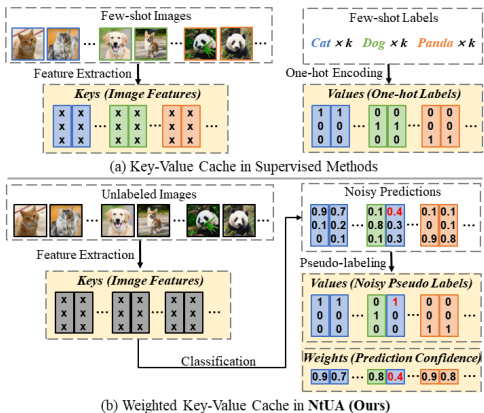


Figure 2: Unlike key-value cache from labelled samples in supervised methods [23, 44], we build weighted key-value cache from unlabelled samples, where the cache weights are determined by the confidence of the pseudo-labels predicted by large-scale vision-language models. The adaptive weighting mechanism makes the unsupervised adaptation more tolerant to noisy pseudo-labels.

Incorporating cache weights significantly enhances NtUA’s tolerance to pseudo-label noises, as the prediction confidence is closely correlated with pseudo-label accuracy [32, 32]. In addition, we design a knowledge-guided cache refinement technique that effectively improves the quality of the predicted pseudo-labels. This technique leverages CLIP-distilled knowledge, updating pair values and cache weights iteratively. To enhance the robustness of fine-tuning, we incorporate a similarity term between each image and its corresponding class prototype into the loss function, giving higher weight to image-prototype pairs with higher similarity. Extensive experiments show that NtUA is simple but effective, with an average accuracy gain of 6.78% across 11 widely studied datasets.

## 2 RELATED WORK

**Cache Model** stores training data features and labels in a key-value format, enabling quick retrieval of relevant information during inference [37]. This technique boosts performance across various models, including language models [2, 25], vision models [28], and vision-language models [46, 47, 51]. For example, Tip-Adapter [48] enhances pre-trained vision-language models with a blended cache model, improving retrieval efficiency. Our novel cache model introduces a weighting mechanism to adapt vision-language models better, prioritizing reliable information for more effective adaptation.

**Knowledge Distillation** transfers knowledge from a larger teacher model to a smaller student model, enhancing the student model’s performance [10]. Early methods primarily focused on mimicking teacher models’ predicted categorical probabilities [10, 48]. Subsequent studies further improved knowledge distillation by transferring teacher knowledge to the student models’ backbone features [19, 29, 51, 38, 42, 44, 45]. Unlike most existing studies; we exploit knowledge distillation to learn from powerful vision-language models in the presence of limited unlabelled target samples.

## 3 METHODOLOGY

This section introduces an innovative approach for unsupervised transfer learning, aiming to augment the capability of the pre-trained CLIP model under limited unlabelled target samples. Section 3.1 presents a concise overview of Tip-Adapter, a closely related study that performs supervised transfer from the pre-trained CLIP model. We then provide details of our proposed Noise-tolerant Unsupervised Adapter (NtUA) in Section 3.2, which is capable of handling a more challenging unsupervised setup that transfers CLIP knowledge with a few unlabelled target samples.

### 3.1 A Revisit of Tip-Adapter

Tip-Adapter [48] is an efficient learning method that adapts the pre-trained CLIP model for supervised few-shot image classification. The task of supervised few-shot image classification involves a labelled dataset  $D_{\text{train}} = \{(x_i, y_i)\}_{i=1}^m$ , with an image  $x_i$  and its corresponding ground-truth label  $y_i$ , and the number of samples  $m$ .  $D_{\text{train}}$  has  $N$ -way- $K$ -shot, where  $N$  represents the number of classes and  $K$  represents the number of training examples per class. For each image  $x_{\text{train}}$  in  $D_{\text{train}}$ , Tip-Adapter utilizes the visual encoder  $E_v$  of the pre-trained CLIP model to extract  $d$ -dimensional L2 normalized image features

$f_{\text{train}} = E_v(x_{\text{train}})$  and converts its ground-truth label into a  $N$ -dimensional one-hot vector. For all  $NK$  training samples, Tip-Adapter extracts  $f_{\text{train}} \in \mathbb{R}^{1 \times d}$  and  $y_{\text{train}} \in \mathbb{R}^{1 \times N}$  from each training image to obtain the image features  $\mathbf{F}_{\text{train}} \in \mathbb{R}^{NK \times d}$  and one-hot vectors  $\mathbf{L}_{\text{train}} \in \mathbb{R}^{NK \times N}$  for the whole training set. To build an efficient feature adapter, Tip-Adapter constructs a key-value cache by storing  $\mathbf{F}_{\text{train}}$  as the keys and  $\mathbf{L}_{\text{train}}$  as the values. During inference, the image features  $f_{\text{test}}$  generated from each test image  $x_{\text{test}} \in D_{\text{test}}$  are utilized as a query to retrieve relevant information from the key-value cache. Finally, the prediction logits from Tip-Adapter  $P_{\text{TA}}$  of a testing image  $f_{\text{test}}$  is obtained as follows:

$$P_{\text{TA}}(f_{\text{test}}) = \alpha \varphi(f_{\text{test}} \cdot \mathbf{F}_{\text{train}}^T) \mathbf{L}_{\text{train}} + f_{\text{test}} \cdot \mathbf{W}^T \quad (1)$$

where  $\varphi(x) = \exp(-\beta(1-x))$  is a mapping function defined in [45],  $\beta$  is a modulating hyperparameter,  $\mathbf{W} \in \mathbb{R}^{N \times d}$  denotes the parameters of the textual encoder  $E_t$  of the pre-trained CLIP, and  $\alpha$  refers to a balancing ratio. The keys in the cache model can be tuned using a loss function defined as follows:

$$\mathcal{L}_{\text{TA}}(f_{\text{train}}, y_{\text{train}}) = \mathcal{L}_{\text{CE}}(P_{\text{TA}}(f_{\text{train}}), \mathbf{L}_{\text{train}}) \quad (2)$$

## 3.2 Noise-tolerant Unsupervised Adapter

Unlike Tip-Adapter, NtUA aims to transfer CLIP knowledge using only a few unlabelled target samples. Although Tip-adapter is specifically designed to adapt the CLIP model to target domains with only a few labelled samples, its effectiveness declines when employing limited unlabelled samples for adaptation (see Table. 2). While CLIP can be utilized to predict pseudo-labels for unlabelled target samples and incorporate these predictions for cache modelling, a significant challenge in this unsupervised scenario arises from inaccuracies in CLIP-generated pseudo-labels. Recent research [46] has shown limitations in CLIP’s visual embedding capabilities, especially for data from less common domains. This weakness often leads to inaccurate pseudo-labels (Table. 3). To address this challenge, NtUA facilitates robust and efficient learning of pseudo-label inaccuracies. As depicted in Fig. 3, NtUA comprises two crucial stages: *adaptive cache formation* and *knowledge-guided cache refinement*, followed by fine-tuning of the weighted cache. The first stage leverages a larger CLIP model to enhance pseudo-labels’ accuracy. The second stage refines the cache model further to improve performance, similar to Tip-Adapter-F [46]. These stages are detailed in the subsequent subsections, providing insights into their design principles and how they address the challenges of unsupervised CLIP adaptation.

**Adaptive Cache Formation:** For unsupervised few-shot learning, we possess a set of unlabelled training images denoted as  $I_{NK} = \{x_i\}_{i=1}^{NK}$ . In order to generate pseudo-labels for the unlabelled target data, we employ the visual encoder  $E_v$  of pre-trained CLIP (ViT-B/32) to extract image features  $f_{\text{train}}$  for all target data. Then, we utilize the CLIP’s textual encoder  $E_t$ , which takes a prompt as input, to generate the classifier’s weights  $\mathbf{W}$  for all the target class names. These weights are applied to the image features, resulting in prediction logits given by  $P(f_{\text{train}}) = f_{\text{train}} \cdot \mathbf{W}^T$ . We then convert these logits into probability distributions using softmax. From these predictions, we derive pseudo-labels as one-hot vectors  $\hat{\mathbf{L}}_{\text{train}}$  for all unlabelled target images, along with confidence scores  $\hat{\mathbf{C}}_{\text{train}}$ .

NtUA incorporates both pseudo-labels  $\hat{\mathbf{L}}_{\text{train}}$  and their corresponding weights  $\hat{\mathbf{C}}_{\text{train}}$  in the weighted cache model. It performs inference of new data by combining the predictions

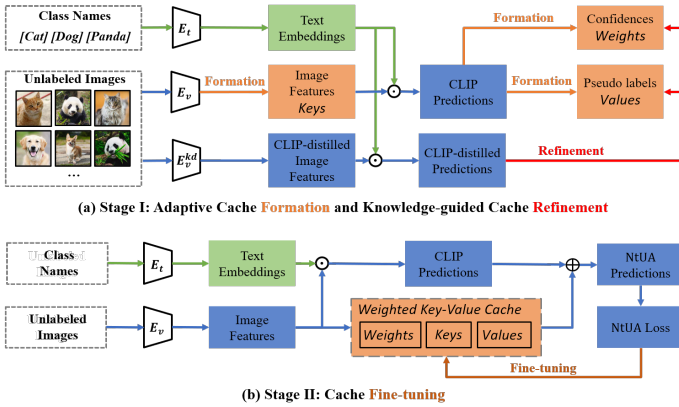


Figure 3: The framework of Noise-Tolerant Unsupervised Adapter (NtUA): (a) At Stage I, NtUA first conducts adaptive cache formation by constructing a weighted key-value cache to store the knowledge of few-shot unlabelled target samples and then applies knowledge-guided cache refinement to rectify both cache values and cache weights. In the cache, the image features extracted with CLIP’s visual encoder  $E_v$  serve as the *keys*, the CLIP-predicted pseudo-labels (generated using  $E_v$  and CLIP’s textual encoder  $E_t$ ) serve as the *values*. The corresponding pseudo-label prediction confidence serves as the *weights* of the key-value pairs. To perform knowledge-guided cache refinement, NtUA generates CLIP-distilled predictions (with CLIP’s visual encoder  $E_v^{kd}$  and textual encoder  $E_t$ ) and leverages such CLIP-distilled knowledge to update both *values* and *weights* in the cache. (b) In Stage II, NtUA updates *the keys* in the constructed *weighted key-value cache* by incorporating knowledge from both the cache and CLIP.

of the weighted cache model and the predictions of the CLIP model as follows:

$$P_{\text{NtUA}}(f_{\text{train}}) = \alpha \hat{\mathbf{C}}_{\text{train}} \varphi(f_{\text{train}} \cdot \mathbf{F}_{\text{train}}^T) \hat{\mathbf{L}}_{\text{train}} + f_{\text{train}} \cdot \mathbf{W}^T \quad (3)$$

Since the accuracy of pseudo-labels is closely related to the prediction confidence [33, 52], as shown in Fig. 4(a), NtUA is more robust to pseudo-label noises by incorporating the prediction confidence  $\hat{\mathbf{C}}_{\text{train}}$  according to the formulation in Eq. 3. To further enhance the quality of both pseudo-labels  $\hat{\mathbf{L}}_{\text{train}}$  and their corresponding cache weights  $\hat{\mathbf{C}}_{\text{train}}$ . We design a novel and effective technique that refines pseudo-labels by leveraging the power of a large CLIP model knowledge. The following subsection provides details about our knowledge-guided cache refinement approach.

**Knowledge-guided Cache Refinement:** To refine the cache model, we leverage the power of another pre-trained CLIP model with a large-scale visual encoder, i.e., ViT-L/14, to effectively enhance the quality of pseudo-labels and their corresponding weights. NtUA initially adopts the visual encoder  $E_v^{kd}$  of the large-scale CLIP model to obtain the CLIP-distilled visual features  $\mathbf{F}_{\text{train}}^{kd} = E_v^{kd}(\mathbf{I}_{NK})$ , then computes CLIP-distilled predictions  $P_{\text{CLIP}}^{kd}$  by multiplying  $\mathbf{F}_{\text{train}}^{kd}$  with the classifier weights  $\mathbf{W}$  of CLIP. Then, the CLIP-distilled pseudo-labels  $\hat{\mathbf{L}}_{\text{train}}^{kd}$  and their corresponding confidence scores  $\hat{\mathbf{C}}_{\text{train}}^{kd}$  can be obtained by converting the CLIP-distilled predictions  $P_{\text{CLIP}}^{kd}$  into one-hot vectors of  $N$  dimensions and determining the maximum probabilities among the CLIP-distilled predictions, respectively.

To perform knowledge-guided cache refinement for the weighted key-value cache, NtUA incorporates the CLIP-distilled pseudo-labels  $\hat{\mathbf{L}}_{\text{train}}^{kd}$  and their corresponding confidence scores  $\hat{\mathbf{C}}_{\text{train}}^{kd}$  to update the cache model’s pair values and cache weights. By integrating the updated weighted key-value cache, NtUA enhances the predictions computed in Eq. 3 using the CLIP-distilled pseudo-labels  $\hat{\mathbf{L}}_{\text{train}}^{kd}$  and the corresponding confidence scores  $\hat{\mathbf{C}}_{\text{train}}^{kd}$




Figure 4: The correct and incorrect pseudo-labels distribution based on (a) the confidence generated from ViT-B/32. (b) The confidence generated from ViT-L/14 (c) the prototype-affinity weights  $\omega$  generated from ViT-L/14. All experiments are done on a 16-shot sample from the Caltech101 training set.

as follows:

$$P_{\text{NtUA}}^{kd}(f_{\text{train}}) = \alpha \hat{\mathbf{C}}_{\text{train}}^{kd} \varphi(f_{\text{train}} \cdot \mathbf{F}_{\text{train}}^T) \hat{\mathbf{f}}_{\text{train}}^{kd} + f_{\text{train}} \cdot \mathbf{W}^T \quad (4)$$

**Weighted Cache Fine-tuning:** During fine-tuning, NtUA updates the keys of the cache model by considering the confidence of pseudo-labels generated from the CLIP model architecture. Specifically, prediction logits from the weighted cache are combined with the CLIP’s prediction logits to fine-tune the cache model based on the confidence score of the pseudo-labels, as defined in Eq. 4. The loss function in NtUA training can be formulated as follows:

$$\mathcal{L}_{\text{NtUA}} = \mathcal{L}_{\text{CE}}(P_{\text{NtUA}}^{kd}(f_{\text{train}}), \hat{\mathbf{L}}_{\text{train}}^{kd}) \quad (5)$$

During inference, NtUA leverages the logits obtained from Eq. 1 to make accurate predictions. Incorporating the confidence score associated with the pseudo-labels when adjusting the cache model’s keys plays a significant role in mitigating the potential adverse effects caused by the presence of inaccurate or noisy pseudo-labels. However, Fig. 4 (a) and (b) illustrate a key challenge associated with pseudo-labels: a somewhattion of the pseudo-labels might be incorrect or assigned with low confidence by the model, even when generated by a large CLIP model like ViT-L/14. While incorporating confidence scores into the adapter helps to some extent, misinterpretations can still occur in specific domains. These misinterpretations can significantly hinder the fine-tuning process.

Inspired by the few-shot learning paradigm proposed by [20, 24, 33], where classification relies on computing the similarity between a data point and class prototypes, we introduce the concept of similarity to prototypes as weights within the loss function. In scenarios involving noisy pseudo-labels with few-shot samples, aggregating image features reduces noise, and the similarity to image prototypes can gauge the reliability of elements belonging to specific classes, as shown in Fig. 4(c). We initially aggregate cache values into  $N$  class prototypes  $\bar{p}^c$  through a simple mean of the embeddings for each class:  $\bar{p}^c = \frac{1}{K} \sum_{i=1}^K f_{\text{train}}^{kd}$ . We then derive weights  $\omega_i^c = f_i^c \cdot \bar{p}^c$  by computing the cosine similarity between each cached value and its corresponding image prototype. These weights are referred to as prototype-affinity weights. Then, we employ  $\omega$  to adjust the loss function accordingly as follows:

$$\mathcal{L}_{\text{NtUA}} = \frac{1}{m} \sum_{i=1}^m \omega_i \cdot \mathcal{L}_{\text{CE}}^i(P_{\text{NtUA}}^{kd}(f_{\text{train}}), \hat{\mathbf{L}}_{\text{train}}^{kd}) \quad (6)$$

By prioritizing pseudo-labels with high confidence scores and similarity to image prototypes, we can ensure that the training relies on accurate pseudo-labels, leading to superior overall performance.

## 4 EXPERIMENT

**Datasets:** All the experiments are conducted on 11 widely-used image classification datasets: ImageNet [9], Caltech101 [9], DTD [9], EuroSAT [9], FGVC Aircraft [23], Food101 [9], Flowers102 [27], OxfordPets [30], SUN397 [40], StanfordCars [17], and UCF101 [32].

**Training Setup:** In our experiments, we utilize CLIP [32] as the pre-trained vision-language model with ViT-B/32 visual encoder. We employ a large-scale CLIP model with ViT-L/14 as the visual encoder to refine the knowledge-guided cache. Unless otherwise stated, all experiments and the ablation study are conducted using 16 samples per class. During fine-tuning, we adopt the Tip-Adapter-F setting proposed by [44]. Specifically, we freeze the pre-trained CLIP encoders and update only the cache model’s keys. The cache model’s keys undergo fine-tuning for 20 epochs using a batch size of 16, a learning rate set to 0.001, and the AdamW optimizer with a cosine scheduler. Our model training is conducted on a single NVIDIA Quadro RTX6000. Additional information regarding the experiments and the selection of the few unlabelled samples can be found in the supplementary material.

### 4.1 Comparison with Baseline Methods

We compare NtUA to four state-of-the-art unsupervised adaptation methods: 1) Zero-shot CLIP [32]; 2) UPL [12]; 3) UPL\* [12]; and 4) LaFTer [26]. In an unsupervised scenario, we also assess the performance of two CLIP-supervised adaptation models, MaPLe [15] and PromptSRC [15]. We employ CLIP-ViT-L/14 to generate pseudo-labels for MaPLe and PromptSRC. The main results of our proposed method obtained from 11 image classification datasets are presented in Table 1. Table 1 shows that NtUA consistently outperforms the state-of-the-art. NtUA notably outperforms zero-shot CLIP by +6.78% with minimal training epochs and samples and remains competitive against UPL and UPL\* with gains of +5.89% and +2.45% respectively. This performance improvement over UPL\* shows that NtUA’s success is not solely reliant on pseudo-label distillation from large models. The confidence and prototype-affinity weights values also play a crucial role. Moreover, NtUA surpasses LaFTer by +5.68% solely through integrating these weights during fine-tuning, despite LaFTer’s use of a pre-trained GPT-3 classifier. Additionally, it outperforms the unsupervised adaptation of two supervised CLIP adaptation methods, MaPLe and PromptSRC, by +5.58% and +4.69%, respectively. For experiments beyond 16 shots, refer to the supplementary material.

While achieving high accuracy is crucial, efficiency is equally essential for practical applications. Fig. 5 demonstrates that NtUA outperforms UPL, UPL\*, LaFTer, MaPLe, and PromptSRC regarding training time. This emphasizes NtUA’s capacity to maintain a delicate equilibrium between performance and computational efficiency.




Figure 5: Comparison of Training Times: NtUA versus Five Baseline Methods using 16 Unlabelled Samples from ImageNet Dataset



Methods	ImgNet	Caltech	DTD	ESAT	FGVCA	Food	Flower	OxPets	SUN	StCars	UCF	Average
CLIP-ViT-B/32	63.77	91.48	44.09	45.27	19.17	80.40	66.59	87.44	62.08	60.12	63.47	62.17
UPL	58.55	90.87	48.29	58.20	18.48	79.17	69.06	84.08	65.99	55.70	65.24	63.06
UPL*	63.17	91.93	52.42	57.40	22.11	79.19	76.65	86.37	<b>66.88</b>	64.58	<b>70.76</b>	66.50
LaFTer	60.41	91.60	48.93	56.06	18.30	78.72	71.49	85.01	62.97	57.09	65.40	63.27
MaPLe	62.04	92.09	46.75	52.26	20.28	80.70	71.42	85.66	63.83	56.75	65.32	63.37
PromptSRC	62.57	91.85	50.83	53.25	<b>22.47</b>	79.79	73.20	85.06	63.73	57.18	66.96	64.26
NtUA (ours)	<b>66.46</b>	<b>94.24</b>	<b>52.96</b>	<b>64.94</b>	21.99	<b>81.42</b>	<b>79.90</b>	<b>90.02</b>	66.67	<b>69.25</b>	70.61	<b>68.95</b>
Supervised	68.52	95.21	69.27	85.25	36.51	82.21	93.79	90.95	74.06	78.22	80.68	77.70

Table 1: Comparison of NtUA with Five SOTA few-shot adaptation methods over 11 widely adopted image classification benchmarks. All experiments are done with CLIP-ViT-B/32 as a backbone and over 16-shot setups. Supervised denotes supervised fine-tuning using labelled target samples.

## 4.2 Ablation Study

**Different Designs:** We conduct a series of ablation experiments to examine the specific contributions of each component in NtUA. Four distinct models are tested (see Table. 2): (1) **NtUA-(KC)**: Baseline CLIP-ViT-B/32 with Key-value Cache (KC) for unsupervised adaptation of Tip-adapter [18]. (2) **NtUA-(KCR)**: Adds Knowledge-guided Cache Refinement (KCR) with pseudo-labels from CLIP-ViT-L/14. (3) **NtUA-(KCR+CKC)**: Integrates Confident weighted Key-value Cache (CKC) on top of **NtUA-(KCR)**. (4) **NtUA-(KCR+CKC+ $\omega$ )**: Incorporates prototype-affinity weights  $\omega$  in the loss function atop **NtUA-(KCR+CKC)**. Compared to **NtUA-(KC)**, **NtUA-(KCR)** gains +3.62% across 11 datasets, showcasing the benefit of using a larger CLIP model for generating pseudo-labels. **NtUA-(KCR+CKC)** consistently outperforms **NtUA-(KCR)**, highlighting the value of incorporating confidence scores into cache refinement. Finally, **NtUA-(KCR+CKC+ $\omega$ )** surpasses all, achieving +6.78% improvement over zero-shot CLIP, emphasizing the efficacy of prototype-affinity weights.

Methods	ImgNet	Caltech	DTD	ESAT	FGVCA	Food	Flower	OxPets	SUN	StCars	UCF	Average
CLIP-ViT-B/32	63.77	91.48	44.09	45.27	19.17	80.40	66.59	87.44	62.08	60.12	63.47	62.17
NtUA-(KC)	64.88	92.74	45.21	50.05	19.23	81.01	69.14	89.40	64.78	62.03	63.89	63.85
NtUA-(KCR)	65.58	92.98	51.24	58.91	<b>22.14</b>	81.46	78.32	89.86	65.60	66.84	69.20	67.47
NtUA-(KCR+CKC)	65.92	92.70	52.01	60.70	22.08	<b>81.52</b>	78.16	89.70	65.49	67.22	69.63	67.74
NtUA-(KCR+CKC+ $\omega$ )	<b>66.46</b>	<b>94.24</b>	<b>52.96</b>	<b>64.94</b>	21.99	81.42	<b>79.90</b>	<b>90.02</b>	<b>66.67</b>	<b>69.25</b>	<b>70.61</b>	<b>68.95</b>

Table 2: Ablation studies are conducted across 11 image classification benchmarks using 16-shot unlabelled target samples. KC, KCR, KCR+CKC, and KCR+CKC+ $\omega$  correspond to the baseline key-value cache [18], our proposed pseudo-label refinement, our proposed weighted key-value cache, and our proposed weighted-based prototypes loss, respectively.

**Pseudo-Labels Accuracy:** We examine the accuracy of pseudo-labels generated using two Vision Transformer (ViT) backbones, ViT-B/32 and ViT-L/14, to assess the effectiveness of knowledge distillation. We focus on understanding how the size of the CLIP model’s image encoder affects the quality of pseudo-labels. Employing both backbones to generate pseudo-labels for unlabelled samples across target training datasets, we observe in Table. 3 that using ViT-L/14 consistently improves pseudo-label quality compared to ViT-B/32. This underscores the importance of model capacity in knowledge distillation tasks, particularly with complex visual data.

**NtUA under Transfer Learning:** We investigate our proposed method, NtUA, in a transfer learning setting, where all unlabelled samples from the target datasets are utilized during training, unlike our default setting, which uses a limited number of images. These



Methods	ImgNet	Caltech	DTD	ESAT	FGVCA	Food	Flower	OxPets	SUN	StCars	UCF	Average
ViT-B/32	64.78	90.84	44.40	45.16	18.69	79.86	64.62	81.08	62.38	49.78	59.80	60.13
ViT-L/14	77.20	95.16	52.94	58.70	32.63	90.64	76.69	89.98	66.93	72.02	74.76	71.60

Table 3: Comparison of Pseudo Label Accuracy for ViT-B/32 and CLIP-ViT-L/14 (Whole Dataset).

labels are then integrated into the construction and fine-tuning of the weighted cache model. Our study utilizes 10 diverse datasets, excluding ImageNet, due to the significant time required to generate pseudo labels and training for its vast size (1.28 million images). By excluding ImageNet, we ensure a more manageable experimental setup while still encompassing a diverse range of datasets for evaluation. Table 4 showcases the detailed results obtained from these experiments. Our proposed NtUA method consistently outperforms state-of-the-art approaches across the majority of the datasets.

Methods	Caltech	DTD	ESAT	FGVCA	Food	Flower	OxPets	SUN	StCars	UCF	Average
CLIP-ViT-B/32	91.48	44.09	45.27	19.17	80.40	66.59	87.44	62.08	60.12	63.47	62.17
UPL	91.12	45.09	51.81	17.85	80.73	68.05	83.78	63.69	50.74	60.98	61.38
UPL*	<b>92.62</b>	52.54	62.96	22.44	82.31	77.02	86.94	64.95	61.45	69.89	67.31
LaFTer	92.13	50.30	<b>72.60</b>	19.11	80.21	71.66	84.93	63.81	57.44	66.32	65.85
MaPLe	91.72	48.11	55.58	12.09	81.37	70.44	85.12	64.62	55.60	66.30	63.10
PromptSRC	91.97	52.07	54.79	<b>23.85</b>	81.69	76.05	86.78	65.00	49.57	68.49	66.03
NtUA (ours)	92.45	<b>57.86</b>	67.91	22.83	<b>83.71</b>	<b>79.13</b>	<b>91.01</b>	<b>67.21</b>	<b>71.12</b>	<b>74.52</b>	<b>70.78</b>
Supervised	94.08	75.06	94.48	39.06	84.75	92.04	90.95	76.39	81.18	83.56	81.16

Table 4: Comparison of NtUA with Unsupervised Adaptation Methods on 10 Image Classification Datasets (Full Training Sets)

**Different Image Encoder Architecture:** We examined NtUA’s effectiveness with different CLIP image encoder architectures (ResNet-50). Even with ResNet-50, NtUA maintains its edge over other SOTA, showing only a slight performance decline compared to ViT-B/32. As depicted in Table 5, NtUA outperforms zero-shot CLIP with ResNet-50, showing a significant +7.79% improvement in the 16-shot setting. Furthermore, NtUA consistently surpasses UPL and UPL\* in this scenario. A direct comparison involving LaFTer, MaPLe, or PromptSRC was impractical due to their incompatibility with ResNet-50.

Methods	ImgNet	Caltech	DTD	ESAT	FGVCA	Food	Flower	OxPets	SUN	StCars	UCF	Average
CLIP-RN50	60.32	85.76	42.79	36.15	17.01	77.38	66.10	85.69	58.81	55.79	61.96	58.89
UPL	58.55	90.43	44.44	56.58	15.48	75.64	67.97	82.09	62.36	53.58	63.76	60.99
UPL*	59.79	88.03	<b>51.60</b>	58.69	<b>21.03</b>	75.30	76.57	84.96	<b>64.32</b>	63.04	68.01	64.67
NtUA (ours)	<b>62.30</b>	<b>92.74</b>	50.47	<b>63.85</b>	20.64	<b>78.28</b>	<b>78.28</b>	<b>88.01</b>	63.64	<b>66.47</b>	<b>68.78</b>	<b>66.68</b>
Supervised	65.26	92.94	66.90	85.07	35.94	79.33	94.36	89.42	71.38	75.45	78.75	75.89

Table 5: Comparison of NtUA with Unsupervised Few-Shot Adaptation (16-Shot, CLIP-RN50).

## 5 CONCLUSION

We introduce NtUA, a novel approach called Noise-tolerant Unsupervised Adapter, which enhances target model learning using limited unlabelled samples. NtUA employs a key-value cache representing visual features and predicted pseudo-labels for unlabelled target samples. NtUA uses an adaptive cache formation technique with confidence-based weights to handle noisy pseudo-labels. NtUA also incorporates similarity to image prototypes into the loss function during cache fine-tuning. Additionally, NtUA refines cache

values and weights through knowledge distillation from large-scale vision-language models. Experimental evaluations across multiple classification datasets consistently demonstrate NtUA's superior performance.

## References

- [1] Bang An, Sicheng Zhu, Michael-Andrei Panaitescu-Liess, Chaithanya Kumar Mummadi, and Furong Huang. PerceptionCLIP: Visual classification by inferring and conditioning on contexts. In *The Twelfth International Conference on Learning Representations*, 2024.
- [2] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *European conference on computer vision*, pages 446–461. Springer, 2014.
- [3] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3606–3613, 2014.
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [5] Li Fei-Fei, Rob Fergus, and Pietro Perona. 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*, pages 178–178. IEEE, 2004.
- [6] 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*, 132(2):581–595, 2024.
- [7] Edouard Grave, Moustapha M Cisse, and Armand Joulin. Unbounded cache model for online language modeling with open vocabulary. *Advances in neural information processing systems*, 30, 2017.
- [8] Mingxuan Gu, Sulaiman Vesal, Ronak Kosti, and Andreas Maier. Few-shot unsupervised domain adaptation for multi-modal cardiac image segmentation. In *Bildverarbeitung für die Medizin 2022*, pages 20–25, Wiesbaden, 2022. Springer Fachmedien Wiesbaden. ISBN 978-3-658-36932-3.
- [9] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. 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.
- [10] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.

- [11] Xuefeng Hu, Ke Zhang, Lu Xia, Albert Chen, Jiajia Luo, Yuyin Sun, Ken Wang, Nan Qiao, Xiao Zeng, Min Sun, et al. Reclip: Refine contrastive language image pre-training with source free domain adaptation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2994–3003, 2024.
- [12] Tony Huang, Jack Chu, and Fangyun Wei. Unsupervised prompt learning for vision-language models. *arXiv preprint arXiv:2204.03649*, 2022.
- [13] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V Le, Yunhsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, 2021.
- [14] Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. Generalization through memorization: Nearest neighbor language models. *arXiv preprint arXiv:1911.00172*, 2019.
- [15] 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*, pages 19113–19122, 2023.
- [16] Muhammad Uzair Khattak, Syed Talal Wasim, Muzammal Naseer, Salman Khan, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Self-regulating prompts: Foundational model adaptation without forgetting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15190–15200, 2023.
- [17] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE International Conference on computer vision workshops*, pages 554–561, 2013.
- [18] Junnan Li, Silvio Savarese, and Steven C. H. Hoi. Masked unsupervised self-training for label-free image classification. In *ICLR*, 2023.
- [19] Zeqi Li, Ruowei Jiang, and Parham Aarabi. Semantic relation preserving knowledge distillation for image-to-image translation. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 648–663, Cham, 2020. Springer International Publishing. ISBN 978-3-030-58574-7.
- [20] Kevin J Liang, Samrudhdi B Rangrej, Vladan Petrovic, and Tal Hassner. Few-shot learning with noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9089–9098, 2022.
- [21] Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, Jaakko Lehtinen, and Jan Kautz. Few-shot unsupervised image-to-image translation. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10550–10559, 2019. doi: 10.1109/ICCV.2019.01065.
- [22] Xin Liu, Ziheng Jiang, Josh Fromm, Xuhai Xu, Shwetak Patel, and Daniel McDuff. Metaphys: few-shot adaptation for non-contact physiological measurement. In *Proceedings of the conference on health, inference, and learning*, pages 154–163, 2021.

- [23] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013.
- [24] Carlos Medina, Arnout Devos, and Matthias Grossglauser. Self-supervised prototypical transfer learning for few-shot classification. *arXiv preprint arXiv:2006.11325*, 2020.
- [25] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*, 2016.
- [26] M. Jehanzeb Mirza, Leonid Karlinsky, Wei Lin, Mateusz Kozinski, Horst Possegger, Rogerio Feris, and Horst Bischof. Lafter: Label-free tuning of zero-shot classifier using language and unlabeled image collections. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2023.
- [27] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729. IEEE, 2008.
- [28] Emin Orhan. A simple cache model for image recognition. *Advances in Neural Information Processing Systems*, 31, 2018.
- [29] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3967–3976, 2019.
- [30] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3498–3505. IEEE, 2012.
- [31] Baoyun Peng, Xiao Jin, Jiaheng Liu, Dongsheng Li, Yichao Wu, Yu Liu, Shunfeng Zhou, and Zhaoning Zhang. Correlation congruence for knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5007–5016, 2019.
- [32] 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*, pages 8748–8763. PMLR, 2021.
- [33] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017.
- [34] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. A dataset of 101 human action classes from videos in the wild. *Center for Research in Computer Vision*, 2(11): 1–7, 2012.
- [35] Korawat Tanwisuth, Shujian Zhang, Huangjie Zheng, Pengcheng He, and Mingyuan Zhou. Pouf: Prompt-oriented unsupervised fine-tuning for large pre-trained models. In *International Conference on Machine Learning*, pages 33816–33832. PMLR, 2023.

- [36] Frederick Tung and Greg Mori. Similarity-preserving knowledge distillation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1365–1374, 2019.
- [37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [38] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. *arXiv preprint arXiv:2006.10726*, 2020.
- [39] Xinyi Wu, Zhenyao Wu, Yuhang Lu, Lili Ju, and Song Wang. Style mixing and patch-wise prototypical matching for one-shot unsupervised domain adaptive semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2740–2749, 2022.
- [40] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE Computer Society Conference on computer vision and Pattern Recognition*, pages 3485–3492. IEEE, 2010.
- [41] Jinyu Yang, Jiali Duan, Son Tran, Yi Xu, Sampath Chanda, Liquan Chen, Belinda Zeng, Trishul Chilimbi, and Junzhou Huang. Vision-language pre-training with triple contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15671–15680, 2022.
- [42] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *ICLR*, 2017.
- [43] Jiajin Zhang, Hanqing Chao, Amit Dhurandhar, Pin-Yu Chen, Ali Tajer, Yangyang Xu, and Pingkun Yan. Spectral adversarial mixup for few-shot unsupervised domain adaptation. In *Medical Image Computing and Computer Assisted Intervention – MIC-CAI 2023*, pages 728–738, Cham, 2023. Springer Nature Switzerland.
- [44] Linfeng Zhang and Kaisheng Ma. Improve object detection with feature-based knowledge distillation: Towards accurate and efficient detectors. In *International Conference on Learning Representations*, 2020.
- [45] Linfeng Zhang, Yukang Shi, Zuoqiang Shi, Kaisheng Ma, and Chenglong Bao. Task-oriented feature distillation. *Advances in Neural Information Processing Systems*, 33: 14759–14771, 2020.
- [46] Renrui Zhang, Wei Zhang, Rongyao Fang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free adaption of clip for few-shot classification. In *European Conference on Computer Vision*, pages 493–510. Springer, 2022.
- [47] Renrui Zhang, Xiangfei Hu, Bohao Li, Siyuan Huang, Hanqiu Deng, Yu Qiao, Peng Gao, and Hongsheng Li. Prompt, generate, then cache: Cascade of foundation models makes strong few-shot learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15211–15222, 2023.

- [48] Ying Zhang, Tao Xiang, Timothy M Hospedales, and Huchuan Lu. Deep mutual learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4320–4328, 2018.
- [49] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022.
- [50] 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*, pages 16816–16825, 2022.
- [51] Xiangyang Zhu, Renrui Zhang, Bowei He, Aojun Zhou, Dong Wang, Bin Zhao, and Peng Gao. Not all features matter: Enhancing few-shot clip with adaptive prior refinement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2605–2615, 2023.
- [52] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5982–5991, 2019.