



Online Zero-Shot Classification with CLIP

Qi Qian¹* and Juhua Hu²

¹ Alibaba Group, Bellevue, WA 98004, USA

² School of Engineering and Technology,
University of Washington, Tacoma, WA 98402, USA
qi.qian@alibaba-inc.com, juhuah@uw.edu

Abstract. Vision-language pre-training such as CLIP enables zero-shot transfer that can classify images according to the candidate class names. While CLIP demonstrates an impressive zero-shot performance on diverse downstream tasks, the distribution from the target data has not been leveraged sufficiently. In this work, we study a novel online zero-shot transfer scenario, where each image arrives in a random order for classification and is visited only once to obtain prediction immediately without storing its representation. Compared with the vanilla zero-shot classification, the proposed framework preserves its flexibility for online service while considering the statistics of the arrived images as the side information to capture the distribution of target data, which can help improve the performance of real-world applications. To tackle the challenge of effective online optimization, we first develop online label learning to model the target data distribution. Then, the proxy of each class in the vision space is further optimized with the proposed online proxy learning method to mitigate the modality gap between images and text. The convergence of both online strategies can be theoretically guaranteed. By combining the predicted label from the online label learning and proxy learning, our online zero-shot transfer method (OnZeta) achieves 78.94% accuracy on ImageNet without accessing the entire data set. Moreover, extensive experiments on other 13 downstream tasks with different vision encoders show a more than 3% improvement on average, which demonstrates the effectiveness of our proposal.

Keywords: Online learning · Zero-shot classification · CLIP

1 Introduction

Vision-language pre-training has attracted much attention recently due to its impressive zero-shot transfer performance on various downstream tasks. The desired property mainly comes from aligning the vision and text space. For example, one of the most prevalent pre-training methods, *i.e.*, CLIP [17], consists of distinct vision and text encoders for learning image and text representations, respectively. These encoders are optimized by minimizing a contrastive loss defined on image-text pairs. The loss aims to pull images and their corresponding text descriptions together while pushing the irrelevant text or images away [15].

* Corresponding author


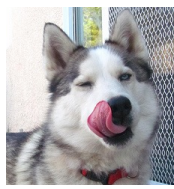


Fig. 1: Illustration of the proposed online zero-shot transfer method (OnZeta). Blue and orange lines denote the inference in the text and vision space, respectively. By incorporating the predictions from online label learning and online proxy learning, OnZeta can leverage the biased prediction from the text space to reduce the variance in the target vision space in an online manner.

After the alignment between vision and text, a novel zero-shot classification paradigm emerges to transfer knowledge in pre-trained models to diverse tasks. Concretely, given all class names from the target task, each class can be represented by a class proxy that applies the text encoder in CLIP for the corresponding class name, which is also termed a text proxy. When an image arrives, its representation is extracted by the vision encoder of CLIP. Then, the label can be assigned by the 1-nearest neighbor (1-NN) classifier with the set of text proxies as the reference. Without any fine-tuning, this simple strategy achieves 77.02% accuracy on ImageNet [18] using the pre-trained ViT [4] as the vision encoder.

Since the success of CLIP, many methods have been developed to explore different side information from the target task to further improve the transfer performance. Firstly, if a few labeled images are available from the target task, the zero-shot transfer can be formulated as a few-shot learning problem and the input text prompt can be learned with the labeled vision data by fine-tuning [26]. Secondly, if the computational resource is too limited to do fine-tuning, the labeled data can be cached as a vision classifier to classify images directly [24]. However, supervised information is often rare, while unsupervised information is more accessible. Therefore, many efforts have been devoted to exploiting the unlabeled images to enhance the zero-shot performance. [19] generates multiple augmented examples from a given image, and then learns the specific text prompt for each image as few-shot learning accordingly. [16] considers reconstructing the class proxy in the vision space, *i.e.*, vision proxy, by leveraging a set of unlabeled images in proxy learning. While [16] shows superior zero-shot performance, it optimizes proxies in an offline manner. Due to the privacy issue, gathering or keeping a set of unlabeled images is still challenging in some real-world applications.

Therefore, in this work, we investigate a novel practical scenario for zero-shot transfer, termed online zero-shot transfer. Concretely, when streaming images arrive as in conventional zero-shot learning, the model has to classify the image immediately without refining. After that, the class proxy can be updated but



Baseline	Alaskan Malamute	Norwegian Elkhound	Husky	Siberian Husky
	32.7%	26.4%	20.0%	9.4%
OnZeta	Husky	Alaskan Malamute	Siberian Husky	Norwegian Elkhound
	23.2%	18.4%	11.7%	10.4%

Fig. 2: Top 4 predicted labels with corresponding probabilities from baseline zero-shot method by CLIP and our proposal.

the representations of arrived images cannot be kept. Unlike [16] that can access the whole unlabeled set, this online setting is way more challenging, where only statistics of seen images can be leveraged for optimization and each image is visited only once. Thereafter, the standard iterative optimization in [16] is not applicable, while the problem (*i.e.*, online service to classify each unlabeled image on the fly without storing data) is ubiquitous in real-world applications (*e.g.*, mobiles and robots).

To address the online service challenge, we first consider exploring the data distribution of the target task to improve the prediction from the text proxy. Note that the vanilla zero-shot prediction focuses on individual images and ignores the distribution over the entire data set. To capture the distribution of all data, we develop an online label learning algorithm to balance the assignment between different classes globally. In addition, the class proxy from the text (*i.e.*, text proxy) can be biased for the vision data due to the modality gap, and thus we propose online proxy learning to obtain the class proxy in the vision space (*i.e.*, vision proxy) directly in an online manner. Finally, the prediction from both the text space and vision space can be combined to help reduce the variance of learning in the vision space as illustrated in the framework of our online zero-shot transfer (OnZeta) shown in Fig. 1. More importantly, with online optimization, our proposal OnZeta can leverage the information from target data effectively and efficiently, which helps correct the prediction compared to the baseline zero-shot method. For example, Fig. 2 shows that an image of “Husky” will be misclassified by vanilla CLIP. By considering the distribution over the whole data set and incorporating the learned vision proxy, our method can obtain the right label in an online manner. The main contributions of this work can be summarized as follows.

- Considering the practical need for online service (*i.e.*, classifying each image on the fly without storage) and improving the performance of vanilla zero-shot classification, we study a novel setting, that is, online zero-shot transfer, using pre-trained vision-language models.
- To address the online service challenge, we propose an online label learning method to trace the distribution of assigned labels over the entire data in an online manner. Moreover, an online vision proxy learning method is proposed

to further improve the class proxy in vision space for classification. The convergence of both methods can be theoretically guaranteed.

- Experiments on 14 downstream tasks demonstrate the effectiveness of our proposal OnZeta. Concretely, when visiting each image only once, OnZeta outperforms the baseline by a margin of more than 3% on average without storing any unlabeled data.

2 Related Work

CLIP consists of two encoders (*i.e.*, vision and text), and the optimization can be conducted on each of them individually. In this section, we briefly review research efforts in these two directions.

2.1 Optimization with Text Encoder

CLIP demonstrates that the zero-shot performance heavily depends on the text prompts for generating appropriate class proxies, and many methods investigate learning the optimal text prompt as input to the text encoder. CoOp [26] marks the input text prompts as learnable variables and optimizes them via few-shot learning with a limited number of labeled images for each class. CoCoOp [25] further includes the vision representation as the conditional context for prompt learning. TPT [19] relaxes the requirement of labeled examples and optimizes the consistency between different augmentations from the same image for prompt learning. However, it has to optimize the prompt for each image with 64 augmentations, which is inefficient for computation and memory, making it infeasible for real-time applications as the online scenario in this work.

2.2 Optimization with Vision Encoder

While prompt learning can improve the performance on the target data, recent research shows that optimizing the vision representation can be more effective for classification [6, 16]. CLIP-Adapter [6] introduces an additional adapter network after the vision encoder and optimizes the vision representation with few-shot learning. Although a similar adapter can be attached for the text encoder, the empirical study shows that optimizing the vision space is sufficient for transfer and can achieve better performance than either the version with only a text adapter or that with both vision and text adapters. [22] shows that fine-tuning the vision classifier directly with few-shot learning can achieve better performance than CoOp. Recently, the analysis in [16] demonstrates that the optimal class proxy for visual classification lives in the vision space. To mitigate the modality gap, [16] proposes proxy learning that aims to learn class proxy for each class in vision space. However, it requires a set of unlabeled data and will learn proxies iteratively. In this work, we also aim to refine the class proxy in the vision space, but with the learned labels in an online manner, which is more applicable for online real applications.

3 Online Zero-shot Transfer

3.1 Zero-shot Transfer in CLIP

Given a data set consisting of image-text pairs as $\{I_i, T_i\}_{i=1}^n$, where I_i denotes an image and T_i is the corresponding text such as class names or captions, CLIP trains vision and text encoders jointly to align the vision subspace and its corresponding text subspace. Let $f(\cdot)$ and $g(\cdot)$ denote the vision encoder and text encoder, respectively. The representations for images and text can be extracted as $\mathbf{x}_i = f(I_i)$ and $\mathbf{z}_i = g(T_i)$.

After pre-training two encoders, zero-shot classification can be implemented by the 1-nearest neighbor (1-NN) classifier. Concretely, given the class names of the target task, the proxy of j -th class \mathbf{z}_j can be obtained by applying the text encoder for the prompt as “a photo of a {class name}”. Then, the image I_i will be classified as

$$y_i = \arg \max_j \mathbf{x}_i^\top \mathbf{z}_j \quad (1)$$

While the conventional zero-shot paradigm demonstrates an impressive transfer performance, the side information from the target task has not been explored sufficiently. Recent work shows that with only a set of unlabeled target data, the zero-shot performance can be improved dramatically [16]. In this work, we further relax the requirement of unlabeled target data and investigate a novel online zero-shot transfer scenario, where each unlabeled image arrives in an online manner without storing.

3.2 Online Label Learning

First, we develop an online label learning method to leverage the information from arrived images to capture the distribution over the entire set.

Given text proxy \mathbf{z}_j , the predicted distribution over classes for I_i can be computed as

$$q_{i,j} = \frac{\exp(\mathbf{x}_i^\top \mathbf{z}_j / \tau_T)}{\sum_k^C \exp(\mathbf{x}_i^\top \mathbf{z}_k / \tau_T)} \quad (2)$$

where τ_T is the temperature optimized by CLIP.

If the whole image set is available, the label can be refined by considering the distribution over classes. Let $\{p_i\}$ denote the learnable labels, and the optimization problem with C classes can be cast as

$$\min_{p_i \in \Delta} \frac{1}{n} \sum_i D_{KL}(p_i || q_i) \quad s.t. \quad \forall j, \quad \frac{1}{n} \sum_i p_{i,j} \geq \frac{\alpha}{C} \quad (3)$$

where $p_i = [p_{i,1}, \dots, p_{i,C}]$, $q_i = [q_{i,1}, \dots, q_{i,C}]$ and Δ denotes the simplex as $\Delta = \{p_i \in \mathbb{R}^C | \sum_j^C p_{i,j} = 1; \forall j, p_{i,j} \geq 0\}$. The constraint bounds the minimal

Algorithm 1 Online Label Learning (OnLab)

-
- 1: **Input:** Unlabeled image set $\{I_i\}$, class names $\{T_j\}$, pre-trained CLIP encoders (f, g) , dual variables $\{\rho_j\}$, ratio α , temperature τ_T , initial learning rate c_ρ
 - 2: Obtain representation of text proxy as $\mathbf{z}_j = g(T_j)$
 - 3: Initialize $\forall j, \rho_j = 0$
 - 4: **while** the i -th image arrives **do**
 - 5: Obtain vision representation as $\mathbf{x}_i = f(I_i)$
 - 6: Observe the probability p_i^* by Eqn. 6
 - 7: Have the prediction as $y_i = \arg \max_j p_{i,j}^*$
 - 8: Update dual variables as in Eqn. 7
 - 9: **end while**
-

size of each class to avoid collapsing and $\alpha \in [0, 1]$ is the ratio for balancing the assignment between different classes. $\alpha = 1$ denotes a well-balanced distribution that each class has the same proportion of images, while $\alpha = 0$ keeps the prediction of the original zero-shot classification from q_i .

The problem can be solved effectively given all examples, but optimizing the objective with streaming images is challenging. To address the online problem, we introduce the Lagrange dual variables [2] $\{\rho_j\}$ for constraints and rewrite the problem as

$$\mathcal{L}(\rho, p_i) = \max_{\rho_j \geq 0} \min_{p_i \in \Delta} \frac{1}{n} \sum_i D_{KL}(p_i || q_i) - \sum_j \rho_j \left(\frac{1}{n} \sum_i p_{i,j} - \frac{\alpha}{C} \right) \quad (4)$$

When the i -th example arrives for prediction, the sub-problem with fixed dual variables becomes

$$\mathcal{L}'(p_i) = \min_{p_i \in \Delta} D_{KL}(p_i || q_i) - \sum_j \rho_j^{i-1} p_{i,j} \quad (5)$$

where ρ_j^{i-1} indicates the dual variable used for the $(i-1)$ -th arrived image.

Fortunately, this problem has the closed-form solution as follows.

Proposition 1 *The optimal solution to the problem in Eqn. 5 is*

$$p_{i,j}^* = \frac{q_{i,j} \exp(\rho_j^{i-1})}{\sum_k q_{i,k} \exp(\rho_k^{i-1})} \quad (6)$$

Proof. It is from the K.K.T. condition [2].

Remark Proposition 1 shows that the dual variable helps balance the assignment between classes. When $\rho = 0$ without optimization, it degenerates to the original prediction from CLIP.

With the prediction for the image, the dual variable can be updated by gradient ascent as

$$\rho_j^i = \max\{0, \rho_j^{i-1} - \eta_\rho^i (p_{i,j}^* - \alpha/C)\} \quad (7)$$

where η_ρ^i is the learning rate for the dual variable. When $\{\rho_j\}$ is initialized as 0 and $\alpha = 0$, dual variables will not be updated, which leads to the original prediction.

Our **online label learning** (OnLab) algorithm is summarized in Alg. 1. The prediction will be obtained by a close-form solution and the updating for dual variables is also efficient, which makes the proposed method applicable for real-time applications.

Besides the efficiency, the effectiveness of OnLab can be theoretically guaranteed in the following theorem. The detailed proof can be found in the supplementary.

Theorem 1 *By running Alg. 1 for solving the problem in Eqn. 4 with the step size $\eta_\rho^i = c_\rho/\sqrt{i}$, where c_ρ is a constant, the convergence can be guaranteed as*

$$\frac{1}{n} \max_{\rho} \sum_i \mathcal{L}(\rho, p_i^*) - \min_{p_i} \sum_i \mathcal{L}(\rho^{i-1}, p_i) \leq \mathcal{O}\left(\frac{1}{\sqrt{n}}\right) \quad (8)$$

3.3 Online Proxy Learning

After optimizing the label by online label learning, we consider reconstructing the class proxy in the vision space to reduce the modality gap from the text proxy as suggested in [16]. Given the whole set of images, vision proxy learning can be cast as an optimization problem

$$\min_{\mathbf{w}} \mathcal{L}(\mathbf{w}) = - \sum_i \sum_j p_{i,j}^* \log(p'_{i,j}) \quad (9)$$

where $p_{i,j}^*$ is the output of online label learning and

$$p'_{i,j} = \frac{\exp(\mathbf{x}_i^\top \mathbf{w}_j / \tau_I)}{\sum_k \exp(\mathbf{x}_i^\top \mathbf{w}_k / \tau_I)} \quad (10)$$

\mathbf{w}_j denotes the learnable vision proxy for the j -th class. τ_I is the temperature that should be larger than τ_T to obtain the class proxy in the vision space [16].

When solving the problem in an online manner, only a single example will be received at each iteration. Therefore, the vision proxy will be updated according to the gradient from the i -th example as

$$\mathbf{w}^i = \Pi_{\|\mathbf{w}\|_2=1}(\mathbf{w}^{i-1} - \eta_w^i \nabla \mathcal{L}(\mathbf{w}^{i-1}, \mathbf{x}_i)) \quad (11)$$

where $\Pi(\cdot)$ projects the updated vision proxy to the unit norm. Note that the loss function is convex in \mathbf{w} , and the convergence can be guaranteed by the standard online learning theory [7].

Theorem 2 *By updating \mathbf{w} as in Eqn. 11 and setting $\eta_w^i = c_w/\sqrt{i}$, where c_w is a constant, the regret can be bounded as*

$$\frac{1}{n} \sum_i \mathcal{L}(\mathbf{w}^{i-1}, \mathbf{x}_i) - \min_{\mathbf{w}} \mathcal{L}(\mathbf{w}) \leq \mathcal{O}\left(\frac{1}{\sqrt{n}}\right) \quad (12)$$

Algorithm 2 Online Zero-shot Transfer (OnZeta)

-
- 1: **Input:** Unlabeled image set $\{I_i\}$, class names $\{T_j\}$, pre-trained CLIP encoders (f, g) , dual variables $\{\rho_j\}$, vision proxy $\{\mathbf{w}_j\}$, ratio α , temperature τ_T and τ_I , initial learning rate c_ρ and c_w , weight parameter β
 - 2: Obtain representation of text proxy as $\mathbf{z}_j = g(T_j)$
 - 3: Initialize $\forall j, \rho_j = 0$ and $\mathbf{w}_j = \mathbf{z}_j$
 - 4: **while** the i -th image arrives **do**
 - 5: Obtain vision representation as $\mathbf{x}_i = f(I_i)$
 - 6: Observe the probability p_i^* by Eqn. 6
 - 7: Observe the probability p_i' by Eqn. 10
 - 8: Have the prediction as $y_i = \arg \max_j \tilde{p}_{i,j}$ where \tilde{p}_i is from Eqn. 13
 - 9: Update dual variables as in Eqn. 7
 - 10: Update vision proxy as in Eqn. 11
 - 11: **end while**
-

Since images arrive in a random order and only a single image will be leveraged for updating \mathbf{w} , the variance in learning can be large, especially for the first few iterations. Unlike offline optimization in [16], where the performance is evaluated after the final iteration, online learning has to accumulate the accuracy of predictions from different timestamps. Therefore, we consider combining the pseudo label p_i^* from online label learning in text space to stabilize the prediction.

For the i -th example, assume that \hat{p}_i is the ground-truth distribution in the vision space. Since p_i^* is mainly from the text proxy, it is a biased estimation from the text space. While p_i' is estimated by the vision proxy, the variance can be large due to the online updating. Therefore, these predictions can be mixed to trade-off bias from text space and variance from vision space. Concretely, let

$$\tilde{p}_i = \lambda_i p_i' + (1 - \lambda_i) p_i^* \quad (13)$$

then we have

$$\begin{aligned} E[\|\tilde{p}_i - \hat{p}_i\|_2^2] \\ = E[\lambda_i^2 \|p_i' - \hat{p}_i\|_2^2 + (1 - \lambda_i)^2 \|p_i^* - \hat{p}_i\|_2^2 + 2\lambda_i(1 - \lambda_i)\langle p_i' - \hat{p}_i, p_i^* - \hat{p}_i \rangle] \end{aligned} \quad (14)$$

For the sake of simplicity, we assume that p_i' is an unbiased estimation for \hat{p}_i such that $E[p_i'] = \hat{p}_i$. Then

$$E_w[\|\tilde{p}_i - \hat{p}_i\|_2^2] = \lambda_i^2 \underbrace{E[\|p_i' - \hat{p}_i\|_2^2]}_{\text{variance}} + (1 - \lambda_i)^2 \underbrace{\|p_i^* - \hat{p}_i\|_2^2}_{\text{bias}} \quad (15)$$

To minimize the approximation error, the trade-off ratio λ_i has the optimal solution as

$$\lambda_i^* = \frac{\|p_i^* - \hat{p}_i\|_2^2}{\|p_i^* - \hat{p}_i\|_2^2 + E[\|p_i' - \hat{p}_i\|_2^2]} \quad (16)$$

We find that the bias can be considered as a constant but the variance will be reduced with the learning of \mathbf{w} , which implies a monotonically increasing function for combining the pseudo labels from different spaces. While various weight functions for λ can be developed according to the analysis, for the i -th iteration, we empirically set

$$\lambda_i = \beta\sqrt{i/n} \quad (17)$$

where $\beta \in [0, 1]$ is a constant. The combination strategy helps online transfer when the vision proxy is not trained sufficiently.

Incorporating online label learning and online proxy learning, the proposed **online zero-shot transfer** (OnZeta) algorithm is summarized in Alg. 2. All operations in the proposed method can be computed efficiently without significantly increasing the cost of conventional zero-shot transfer.

4 Experiments

To evaluate the proposed method, we conduct experiments on 14 diverse downstream tasks. Considering that text prompts are essential for obtaining the appropriate text proxy for zero-shot transfer, we follow the common practice of including an ensemble of 7 text prompts as suggested in [17] to generate the text proxy for baseline and our method. The initial learning rates for dual variables and vision proxy are set as $c_\rho = 20$ and $c_w = 0.5$, respectively. The learning rate will be decayed as $\eta_\rho^i = c_\rho/\sqrt{i}$ and $\eta_w^i = c_w/\sqrt{i}$ where i denotes the order of the incoming image. Unlabeled images will arrive in a random order. For the temperature, we fix $\tau_T = 0.01$ that is the optimized value in CLIP and $\tau_I = 0.04$ as suggested in [16]. All experiments are implemented on a single V100 GPU.

4.1 Ablation Study

The main parameters in our method are the ratio for class assignment α and the ratio for pseudo label combination β . In this subsection, we will investigate the influence of these parameters along with the setting of online zero-shot transfer. The ablation experiments are conducted on ImageNet with ResNet-50 [8] as vision encoder in CLIP.

Effect of α α is the ratio to capture the distribution over the entire data set as in Eqn. 3. While the vanilla zero-shot transfer ignores the side information from the arrived examples, the proposed online label learning can leverage the information of distribution by α . Since $\alpha \in [0, 1]$, we vary α in $\{0, 0.4, 0.6, 0.8, 1\}$ and summarize the results in Tab. 1.

Table 1: Effect of α in Eqn. 3 for online label learning on ImageNet with ResNet-50 as vision encoder. “#Min class” denotes the proportion of examples in the smallest class.

α	Acc(%)	#Min class(%)
0	60.32	13.7
0.4	60.41	18.5
0.6	60.63	23.0
0.8	61.20	33.3
1	61.96	39.0




Fig. 3: Illustration of data distribution over 1,000 classes with different α on ImageNet. Best viewed in color.

First, according to our analysis, $\alpha = 0$ will not update dual variables and implies the accuracy of baseline zero-shot transfer as in CLIP. Then, increasing α can consistently improve the accuracy over the baseline, which shows the efficacy of online label learning. Moreover, the size of the smallest class is also increased as α increases and it shows that our proposed method OnLab with only the online label learning can adjust the class assignment effectively.

To further investigate the data distribution derived by different α , we compare the size of all 1,000 classes in Fig. 3. Evidently, the proposed method is capable of balancing the distribution in an online manner. Note that even with $\alpha = 1$, the obtained distribution is not perfectly balanced, which is from the approximation error in the online scenario as analyzed in Theorem 1. We find that our method is not sensitive to α and fix $\alpha = 1$ in the following experiments if not specified.

Effect of β With the predictions from text and vision space in online zero-shot transfer, a combination strategy with β in Eqn. 17 can trade-off between the bias and variance. Tab. 2 reports the performance with different β .

Table 2: Effect of β in Eqn. 17 for online proxy learning on ImageNet with ResNet-50 as vision encoder.

β	0	0.2	0.4	0.6	0.8	1
Acc(%)	61.96	62.17	62.23	62.38	62.50	61.84

Table 3: Comparison of strategies for obtaining the combination ratio λ in Eqn. 13.

β	Fixed	Dynamic
Acc(%)	62.28	62.50

When $\beta = 0$, only the label from the online label learning (*i.e.*, OnLab) is adopted for prediction. If mixing the label from the vision proxy, the accuracy improves. The mixture prediction outperforms OnLab by 0.54% with $\beta = 0.8$ and is 2.18% better than the baseline CLIP. If excluding the label from text space at the end of data streaming as $\beta = 1$, the performance degenerates due to the large variance from the vision space. The phenomenon illustrates the challenge

of online learning that the vision proxy may not be well learned with a limited number of arrived images and the pseudo label from text proxy can help reduce the variance.

According to our analysis in Section 3.3, a dynamic ratio for the bias-variance trade-off is more effective for the proposed online learning scenario. To verify the claim, we compare the dynamic ratio to the fixed one in Tab. 3. The value of fixed λ is searched in $[0, 1]$ by a step size of 0.1 and the best result is reported.

While the fixed λ still improves the performance from the online label learning, a monotonically increasing λ captures the training dynamics better and shows an additional gain of 0.22%. It confirms our analysis about the bias-variance trade-off.

Effect of iterations in optimization Our method has to learn the dual variables and vision proxy with streaming data and those parameters may not be learned well at the early stage of online learning. To investigate the influence of iterations in online learning, we report the accumulated accuracy of different iterations in one epoch on ImageNet in Tab. 4.

Table 4: Comparison of accuracy (%) with different iterations on ImageNet.

#Iters	5k	10k	20k
Acc(%)	60.96	61.50	61.71
	30k	40k	50k
	62.02	62.41	62.50

Table 5: Effect of epochs for optimization on ImageNet. Prediction from the last epoch is evaluated for OnZeta after running multiple epochs.

#Epoch	1(online learning)	2	3	4	5
Acc(%)	62.50	63.24	63.35	63.39	63.46

First, we can observe that even with 5,000 iterations for learning, the accuracy over 1,000 classes is already better than the baseline. With the arrival of more images, the performance increases steadily. It is because that the convergence rate for dual variables and vision proxy is $\mathcal{O}(1/\sqrt{n})$ as analyzed in Theorems 1-2. It implies that these variables can reach a reasonable performance with a limited number of iterations, while it is challenging to obtain the optimal solution as discussed in the following experiment.

Effect of epochs in optimization In the proposed online zero-shot transfer problem, all examples can be visited only once when they arrive. The setting is prevalent in real-world applications but very challenging for optimization. This ablation experiment aims to investigate the performance gain when each example can be visited multiple times.

Concretely, we run the proposed method on multiple epochs of data, where images arrive in a random order within each epoch. For multiple epochs, the prediction from the last one is reported for evaluation. Tab. 5 compares the performance of multiple epochs to that of a single epoch in the online scenario.

First, we can observe that with an additional epoch, the performance of OnZeta can be boosted by 0.74%, which demonstrates the challenge of online zero-shot transfer. When increasing the number of epochs, the performance can be further improved and it achieves 63.46% with 5-epoch learning. The result approaches the performance of learning with the entire data set as in [16] that shows a 63.74% accuracy when visiting all data by 2,000 epochs. The competitive performance of OnZeta confirms the effectiveness of the proposed online learning algorithm. More experiments can be found in the supplementary.

4.2 Comparison on ImageNet

After the ablation study, we compare our method OnZeta to state-of-the-art methods on ImageNet. Multiple vision encoders from CLIP are adopted for evaluation, which include ResNet-50, ViT-B/32, ViT-B/16, ViT-L/14, and ViT-L/14@336. Considering that images come in a random order in our online setting, the proposed method is repeated by 5 times and the averaged result is reported for the comparison. We note that the standard deviation is quite small and also report the best result among different runs by *. The same parameters are shared by different vision encoders in our method. Tab. 6 summarizes the comparison, where the performance of InMaP [16] is marked by gray due to the application of the whole unlabeled data set. All other methods only access a single image at each iteration. Our OnLab with only the online label learning is denoted as “OnZeta_{0.5}” and the vanilla zero-shot transfer in Eqn. 1 is “Baseline”.

Table 6: Comparison of accuracy (%) on ImageNet with different vision encoders in CLIP. The overall best performance is in bold. “-” denotes that the result is unavailable in their original papers. The averaged performance over 5 random trials is reported for our method, where the best one is denoted by *. InMaP is marked by gray since it leverages the whole unlabeled image set at each iteration.

Vision encoder	Baseline [17]	TPT [19]	OnZeta _{0.5}	OnZeta _{0.5} *	OnZeta	OnZeta*	InMaP [16]
ResNet-50	60.32	60.74	61.89	61.96	62.42	62.50	63.74
ViT-B/32	63.77	-	65.06	65.18	65.71	65.80	67.29
ViT-B/16	68.75	68.98	70.18	70.23	70.87	70.99	72.55
ViT-L/14	75.96	-	76.98	77.07	77.84	77.94	79.29
ViT-L/14 ₃₃₆	77.02	-	77.93	78.05	78.75	78.94	80.21

First, by capturing the overall distribution using only the proposed online label learning component, OnZeta_{0.5} already gains about 1% for all vision encoders over the baseline. It demonstrates that the side information from the arrived images is useful for improving the zero-shot transfer. By learning the vision proxy and combining the predictions from both the text and vision spaces, OnZeta shows about 1% additional improvement over OnZeta_{0.5}, which illustrates that online proxy learning is complementary to online label learning for zero-shot classification. Moreover, OnZeta obtains more gains with a large backbone, showing its potential for large models.

Then, we can observe that the prompt learning method TPT outperforms the baseline with an ensemble of artificial text prompts, which indicates the importance of the appropriate prompt for zero-shot transfer. However, OnZeta surpasses TPT by 1.76% on ResNet-50 and 2.01% on ViT-B/16. It is because that we learn a better class proxy in the vision space for visual classification, which helps reduce the modality gap and is more effective than optimizing the text prompt in the text space. The observation is consistent with the analysis in [16]. Moreover, compared with InMaP that can access the whole unlabeled set, OnZeta is only about 1% worse using different vision encoders. As analyzed in the ablation study, online learning that visits each example once without storing is more challenging than the offline method with multiple iterations over the entire set. Nevertheless, online learning is practical for real-world applications. Finally, all vision encoders share the same parameters in experiments. It implies that our method OnZeta is not sensitive to the selection of vision encoder and is applicable in different configurations.

4.3 Comparison on Other 13 Downstream Tasks

Besides the comparison on ImageNet, we conduct experiments on other 13 downstream tasks to evaluate the performance of our method. The benchmark data sets for zero-shot transfer are adopted in the comparison, which includes Aircraft [12], Caltech101 [5], Stanford Cars [10], CIFAR-10 [11], CIFAR-100 [11], CUB200-2011 [21], Describable Textures Dataset (DTD) [3], EuroSAT [9], Flowers [13], Food101 [1], Oxford-IIIT Pet (Pets) [14], Sun397 [23], and UCF101 [20]. A broad range of downstream tasks are involved by these diverse data sets, *e.g.*, low-resolution image classification, fine-grained visual categorization (FGVC), land cover classification with satellite images, scene categorization, texture recognition, *etc.* These tasks contain fewer images than ImageNet, which increases the challenge of our online transfer with the convergence rate of $\mathcal{O}(1/\sqrt{n})$. Therefore, we reduce β from 0.8 to 0.4 for all tasks to exploit the label from the text space. For data sets that the baseline already achieves satisfied performance, *e.g.*, Caltech101 and CIFAR-10 with ViT, online label learning can be skipped by letting $\alpha = 0$, while $\alpha = 0.4$ can provide a slight improvement. Other parameters on the rest data sets remain the same as those for ImageNet. Two distinct vision encoders, *i.e.*, ResNet-50 and ViT-B/16, are applied for evaluation. The performance averaged over 5 random trials is reported for our method.

The comparison is summarized in Tab. 7. A similar phenomenon as ImageNet can be observed on these downstream tasks. Concretely, online label learning, *i.e.*, OnZeta_{0.5}, already improves the performance over baselines, where the gain is 3.12% on ResNet and 2.58% on ViT, respectively. Compared with the baseline, only streaming images are leveraged by our method, and the representation of each arrived image is not stored, which preserves the flexibility of zero-shot transfer but captures the distribution over the entire data set in an online manner. In addition, learning the vision proxy with the pseudo label from the text space further improves the performance with different vision encoders, which demonstrates the effectiveness of learning the class proxy directly in the target

Table 7: Comparison of accuracy (%) on 13 downstream tasks with ResNet-50 and ViT-B/16 as vision encoders in CLIP. The overall best performance is in bold. “-” denotes that the result is unavailable in their original papers. The averaged performance over 5 random trials is reported for our method, where the best one is denoted by *. InMaP is marked by gray since it leverages the whole unlabeled image set at each iteration.

	Aircraft	Caltech	Cars	Cifar10	Cifar100	CUB	DTD	EuroSAT	Flowers	Food	Pets	SUN	UCF101	Avg.
<i>ResNet-50:</i>														
Baseline	16.62	86.00	56.31	73.15	40.60	41.37	41.13	26.90	62.97	74.10	81.85	59.25	55.56	55.06
TPT	17.58	87.02	58.46	-	-	-	40.84	28.33	62.69	74.88	84.49	61.46	60.82	-
OnZeta _{0.5}	18.29	86.17	58.57	75.62	46.96	45.68	43.24	33.79	62.42	76.41	83.05	61.64	60.12	57.84
OnZeta _{0.5} *	18.51	86.45	58.91	75.75	47.22	45.93	43.74	34.32	63.17	76.53	83.59	61.81	60.38	58.18
OnZeta	18.31	86.26	59.40	75.84	47.55	46.06	43.38	33.99	63.16	76.83	83.77	62.32	61.00	58.30
OnZeta*	18.54	86.41	59.71	75.95	47.71	46.34	43.91	34.35	63.95	76.97	84.66	62.40	61.38	58.64
InMaP	18.96	86.73	63.30	78.84	49.26	49.17	44.86	35.88	66.68	78.36	87.93	63.82	63.36	60.55
<i>ViT-B/16:</i>														
Baseline	23.13	93.51	66.29	91.09	67.29	49.36	45.09	50.17	66.99	84.43	87.00	65.68	65.21	65.79
TPT	24.78	94.16	66.87	-	-	-	47.75	42.44	68.98	84.67	87.79	65.50	68.04	-
OnZeta _{0.5}	27.19	93.65	67.99	91.13	69.85	52.81	47.51	56.20	68.49	85.83	88.44	68.06	68.50	68.13
OnZeta _{0.5} *	27.48	93.79	68.11	91.18	70.02	53.09	47.87	56.42	69.06	85.97	88.72	68.27	68.86	68.37
OnZeta	27.57	93.69	68.74	91.31	70.39	53.66	48.04	56.52	69.37	86.21	89.01	68.87	69.58	68.69
OnZeta*	28.29	93.83	69.03	91.35	70.52	53.80	48.58	56.74	69.63	86.35	89.32	69.01	69.94	68.95
InMaP	28.35	93.59	73.00	93.27	72.51	57.87	50.77	63.98	71.28	87.76	92.94	70.85	72.06	71.40

vision space. Moreover, OnZeta outperforms TPT on 9 out of 10 (only 10 tasks are available in TPT’s original paper) data sets with different vision encoders. Compared with TPT that has to optimize the text prompt with multiple augmentations of the image by back-propagation through the text encoder, OnZeta only contains the forward pass to extract the representation of the image, which is more efficient and is capable of online deployment. By comparing the best performance of the proposed method (*i.e.*, indicated by * in the table) to the averaged one, we can find that the gap averaged over 13 data sets is only about 0.3%. While the images arrive in a random order, the performance of the proposed method can be theoretically guaranteed as indicated in Theorems 1-2, which is robust to the order of input images. Finally, most data sets share the same parameters and it implies that OnZeta is not sensitive to hyper-parameters and is applicable for different tasks.

5 Conclusion

While CLIP demonstrates impressive zero-shot transfer performance, the information from the target data has not been explored sufficiently. In this work, we investigate an online zero-shot scenario, where a random image arrives at each iteration and the model has to predict its label without storing its representation. By capturing the overall label distribution and refining the class proxy in an online manner, our method OnZeta demonstrates a substantial improvement over the vanilla zero-shot transfer on diverse downstream tasks with different vision encoders. In addition, the operations in the proposed method can be implemented efficiently for real-time applications. Exploring more effective algorithms for challenging online zero-shot transfer can be our future work.

References

1. Bossard, L., Guillaumin, M., Gool, L.V.: Food-101 - mining discriminative components with random forests. In: ECCV. Lecture Notes in Computer Science, vol. 8694, pp. 446–461. Springer (2014)
2. Boyd, S.P., Vandenberghe, L.: Convex Optimization. Cambridge University Press (2014)
3. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., Vedaldi, A.: Describing textures in the wild. In: CVPR. pp. 3606–3613 (2014)
4. 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. OpenReview.net (2021)
5. 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: CVPR workshop. pp. 178–178. IEEE (2004)
6. 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. CoRR **abs/2110.04544** (2021)
7. Hazan, E.: Introduction to online convex optimization. CoRR **abs/1909.05207** (2019)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778. IEEE Computer Society (2016)
9. Helber, P., Bischke, B., Dengel, A., Borth, D.: Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. IEEE J. Sel. Top. Appl. Earth Obs. Remote. Sens. **12**(7), 2217–2226 (2019)
10. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: ICCV workshop. pp. 554–561 (2013)
11. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
12. Maji, S., Rahtu, E., Kannala, J., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. arXiv preprint arXiv:1306.5151 (2013)
13. Nilsback, M., Zisserman, A.: Automated flower classification over a large number of classes. In: ICVGIP. pp. 722–729. IEEE Computer Society (2008)
14. Parkhi, O.M., Vedaldi, A., Zisserman, A., Jawahar, C.: Cats and dogs. In: CVPR. pp. 3498–3505. IEEE (2012)
15. Qian, Q., Shang, L., Sun, B., Hu, J., Li, H., Jin, R.: Softtriple loss: Deep metric learning without triplet sampling. In: ICCV. pp. 6449–6457. IEEE (2019)
16. Qian, Q., Xu, Y., Hu, J.: Intra-modal proxy learning for zero-shot visual categorization with clip. In: NeurIPS (2023)
17. 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. In: ICML. pp. 8748–8763. PMLR (2021)
18. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M.S., Berg, A.C., Fei-Fei, L.: Imagenet large scale visual recognition challenge. Int. J. Comput. Vis. **115**(3), 211–252 (2015)
19. Shu, M., Nie, W., Huang, D., Yu, Z., Goldstein, T., Anandkumar, A., Xiao, C.: Test-time prompt tuning for zero-shot generalization in vision-language models. In: NeurIPS (2022)

20. Soomro, K., Zamir, A.R., Shah, M.: UCF101: A dataset of 101 human actions classes from videos in the wild. CoRR **abs/1212.0402** (2012)
21. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset (2011)
22. Wortsman, M., Ilharco, G., Kim, J.W., Li, M., Kornblith, S., Roelofs, R., Lopes, R.G., Hajishirzi, H., Farhadi, A., Namkoong, H., Schmidt, L.: Robust fine-tuning of zero-shot models. In: CVPR. pp. 7949–7961. IEEE (2022)
23. Xiao, J., Hays, J., Ehinger, K.A., Oliva, A., Torralba, A.: SUN database: Large-scale scene recognition from abbey to zoo. In: CVPR. pp. 3485–3492. IEEE Computer Society (2010)
24. Zhang, R., Zhang, W., Fang, R., Gao, P., Li, K., Dai, J., Qiao, Y., Li, H.: Tip-adapter: Training-free adaption of CLIP for few-shot classification. In: ECCV. vol. 13695, pp. 493–510. Springer (2022)
25. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Conditional prompt learning for vision-language models. In: CVPR. pp. 16795–16804. IEEE (2022)
26. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. *Int. J. Comput. Vis.* **130**(9), 2337–2348 (2022)