Understanding Transferable Representation Learning and Zero-shot Transfer in CLIP

Anonymous Author(s)
Affiliation
Address
email

Abstract

Multi-modal learning has become increasingly popular due to its ability to leverage information from different data sources. Recently, CLIP has emerged as an effective approach that employs vision-language contrastive pretraining to learn joint image and text representations and exhibits remarkable performance in zero-shot learning and text-guided natural image generation. Despite the huge practical success of CLIP, its theoretical understanding remains elusive. In this paper, we formally study transferrable representation learning underlying CLIP and demonstrate how features from different modalities get aligned. We also analyze its zero-shot transfer performance on the downstream tasks. Inspired by our analysis, we propose a new CLIP-type approach, which achieves better performance than CLIP and other state-of-the-art methods on benchmark datasets.

1 Introduction

Recently, CLIP (Radford et al., 2021) emerged as a milestone work that leverages vision-language contrastive pretraining to jointly learn image and text embeddings, using the vast amounts of image-text data available on the web. This approach has achieved remarkable success in zero-shot transfer (Lei Ba et al., 2015). Inspired by CLIP's groundbreaking zero-shot capabilities, subsequent studies (Yao et al., 2022; Li et al., 2022; Mu et al., 2022; Goel et al., 2022; Zhai et al., 2022; Alayrac et al., 2022) emerged with the primary objective of further enhancing CLIP's zero-shot performance. Despite the empirical success of CLIP in zero-shot transfer, the theoretical understanding of how it works remains elusive.

This paper delves into the mechanisms through which CLIP learns transferable representations and demonstrates how such representations ensure successful zero-shot transfer for downstream tasks. We present our theoretical result for transferable representation learning in CLIP and summarize our contributions as follows.

- We theoretically examine transferable representation learning in CLIP. Our analysis shows that if a near-optimal network is obtained on the training data, features from different modalities become aligned, enabling zero-shot learning if appropriate prompts are issued.
- Building upon our general theoretical findings, we delve deeper into specific cases. We illustrate how multi-modal learning aligns different features and reveal when the learned features obtained by CLIP can outperform those obtained through naive square loss.
- We conduct experiments on real data to confirm our theoretical predictions. Furthermore, inspired by our theoretical findings, we propose a new regularization technique for CLIP that effectively leads to improved zero-shot performance.

2 Problem Setting and Preliminaries

2.1 Data Distribution

In our paper, we focus on the setting where the image $\mathbf{x}$ and the text $\mathbf{y}$ are conditionally independent given the shared feature $\mathbf{z}$.

Submitted to the Mathematics of Modern Machine Learning Workshop at NeurIPS 2023. Do not distribute.

Assumption 2.1. Let $(\mathbf{x}, \mathbf{y})$ be generated from the joint distribution $\mathcal{D}_{\mathbf{x}\times\mathbf{y}}$. We assume $\mathbf{z}$ to be a shared feature of $\mathbf{x}, \mathbf{y}$ satisfying $\mathbf{x} \perp \mathbf{y}|\mathbf{z}$, and further denote $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ that follows the joint distribution $\mathcal{D}_{\mathbf{x}\times\mathbf{y}\times\mathbf{z}}$ with marginal distributions $\mathcal{D}_{\mathbf{x}\times\mathbf{z}}, \mathcal{D}_{\mathbf{y}\times\mathbf{z}}$. We further assume $\mathbf{z}$ to be a discrete and sparse random variable $\mathbf{z} \in \mathcal{V} = \{\mathbf{v}_1, \ldots, \mathbf{v}_K\}$ with $p_k := \mathbb{P}(\mathbf{z} = \mathbf{v}_k)$.

2.2 Learning via Contrastive Loss

The CLIP architecture has three main components: (i) an image encoder network $\mathbf{g}$ that can encode the image $\mathbf{x}$ into the embedding $\mathbf{g}(\mathbf{x}) \in \mathbb{R}^d$; (ii) a text encoder network $\mathbf{h}$ that can encode the text $\mathbf{y}$ into an embedding vector $\mathbf{h}(\mathbf{y}) \in \mathbb{R}^d$; and (iii) a score function $f(\mathbf{x}, \mathbf{y}) = \mathbf{sim}(\mathbf{g}, \mathbf{h})$ that measures the similarity between the image $\mathbf{x}$ and the text $\mathbf{y}$ given their embeddings $\mathbf{g}, \mathbf{h}$ $\big($e.g., $f(\mathbf{x}, \mathbf{y}) = \langle \mathbf{g}(\mathbf{x}), \mathbf{h}(\mathbf{y})\rangle\big)$. During the training, we will sample a batch of image-captions pairs $S' = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^B \subseteq S$. The contrastive loss in CLIP aims to align the image representation $\mathbf{g}(\mathbf{x})$ and text representations $\mathbf{h}(\mathbf{y})$ by minimizing the empirical version of the following population loss,

$$
\begin{aligned}
L_{\mathcal{D}^B}(f, \tau) = \mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]} \exp\big(\big[f(\mathbf{x}_1, \mathbf{y}_t) - f(\mathbf{x}_1, \mathbf{y}_1)\big]/\tau\big)\bigg)\bigg] \\
+ \mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]} \exp\big(\big[f(\mathbf{x}_t, \mathbf{y}_1) - f(\mathbf{x}_1, \mathbf{y}_1)\big]/\tau\big)\bigg)\bigg],
\end{aligned} \tag{2.1}
$$

where $\tau > 0$ is a temperature parameter and the expectation is taken with respect to all B random pairs $(\mathbf{x}_t, \mathbf{y}_t)$ i.i.d. sampled from $\mathcal{D}_{\mathbf{x}\times\mathbf{y}}$. Therefore, CLIP learns the score function f with the corresponding representations $\mathbf{g}$ and $\mathbf{h}$ by minimizing $L_{\mathcal{D}^B}(f, \tau)$. In fact, we can divide the training dataset S into n batches $\cup_{k\in[n]}\mathcal{S}_k$. Further discussion of problem setting is deferred to Appendix B.

3 Zero-shot Transfer

The key idea of CLIP is to pull the embeddings of positive image-text pairs together while pushing the embeddings of negative pairs apart. For the data pair $(\mathbf{x}, \mathbf{y}')$ generated with $\mathbf{x} \sim \mathcal{D}_{\mathbf{x}|\mathbf{z}}, \mathbf{y}' \sim \mathcal{D}_{\mathbf{x}|\mathbf{z}'}$, $(\mathbf{x}, \mathbf{y}')$ is a positive pair if $\mathbf{z} = \mathbf{z}'$ and a negative pair if $\mathbf{z} \neq \mathbf{z}'$.

Assumption 3.1 ((α, β, γ)-Completeness)**.** There exists a score function f^* bounded by 1 (i.e., $|f^*| \le 1$) with $f^* = \mathbf{sim}(\mathbf{g}^*, \mathbf{h}^*)$ satisfying the following properties,

- For any $\mathbf{z} \neq \mathbf{z}'$, let $\mathbf{x} \sim \mathcal{D}_{\mathbf{x}|\mathbf{z}}, \mathbf{y} \sim \mathcal{D}_{\mathbf{y}|\mathbf{z}}, \mathbf{x}' \sim \mathcal{D}_{\mathbf{x}'|\mathbf{z}'}, \mathbf{y}' \sim \mathcal{D}_{\mathbf{y}'|\mathbf{z}'}$. With probability at least $1 - \alpha$, we have $f^*(\mathbf{x}', \mathbf{y}) \le f^*(\mathbf{x}, \mathbf{y}) - \gamma$ and $f^*(\mathbf{x}, \mathbf{y}') \le f^*(\mathbf{x}, \mathbf{y}) - \gamma$.
- Let $(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mathcal{D}_{\mathbf{x}\times\mathbf{y}\times\mathbf{z}}$, assume $\mathbb{E}_{(\mathbf{y},\mathbf{z})}\big[\mathrm{Var}_{\mathbf{x}|\mathbf{z}}(f^*(\mathbf{x}, \mathbf{y}))\big], \mathbb{E}_{(\mathbf{x},\mathbf{z})}\big[\mathrm{Var}_{\mathbf{y}|\mathbf{z}}(f^*(\mathbf{x}, \mathbf{y}))\big] \le \beta$.

Further discussion on Assumption 3.1 can be found in Appendix G. In the zero-shot transfer task, we have K prompts $\{\mathbf{y}_k, k \in [K]\}$ where $\mathbf{y}_k \sim \mathcal{D}_{\mathbf{y}|\mathbf{v}_k}$. For a new image $\mathbf{x}$ generated from $\mathcal{D}_{\mathbf{x}}$, we want to predict the label of the shared feature $\mathbf{z}$ in $\mathbf{x}$. The following theorem provides the guarantee of zero-shot transfer learning for CLIP.

Theorem 3.2 (Informal)**.** Suppose Assumption 3.1 hold and we can find an ϵ approximate minimum $\widehat{f} \in \mathcal{F}$ with respect to the temperature τ such that $\widehat{f}$ is bounded by M and

$$
L_{\mathcal{D}^B}(\widehat{f}, \tau) \le L_{\mathcal{D}^B}(f^*, \tau) + \epsilon. \tag{3.1}
$$

For the zero-shot downstream task, we calculate the similarity score $\widehat{f}(\mathbf{x}, \mathbf{y}_k)$ for all $k \in [K]$ and pick the indices of the top-r scores within the set $\{\widehat{f}(\mathbf{x}, \mathbf{y}_k)\}$ as the predictions of the image $\mathbf{x}$. The top-r error is bounded by $\epsilon'/\log(1 + r)$, where $\epsilon' = (C_B + 2) \cdot \big[\epsilon + C\tau^{-1}MB\alpha + C\tau^{-1}(\beta MB)^{1/3} + 2B\exp(-\gamma/\tau)\big]$ and $C = \widetilde{O}(1), C_B = \widetilde{O}(\max_k p_k^{-1}/B)$.

A formal discussion is presented in Appendix G and H. Next, we will introduce a specific problem to illustrate how CLIP can learn transferable features with distinguishable margins, which is hard to achieve by simple square loss.

Definition 3.3 (A Case Study)**.** Let shared feature $\mathbf{z} \in \mathbb{R}^{K_1}$ be random variable uniformly drawn from the set $\mathcal{V} = \{\mathbf{v}_1, \ldots, \mathbf{v}_K\}$ where $\|\mathbf{v}_k\|_2 = 1$, $\max_{k\neq k'}\langle \mathbf{v}_k, \mathbf{v}_k'\rangle = 1 - \gamma$. Let $\boldsymbol{\xi} \in \mathbb{R}^{K_2}, \boldsymbol{\zeta} \in \mathbb{R}^{K_3}$ be unique random features satisfying $\|\boldsymbol{\xi}\|_2, \|\boldsymbol{\zeta}\|_2 \le R$ and are mutually independent given $\mathbf{z}$. The image-text pair is generated as

$$
\mathbf{x} = \mathbf{G}\begin{bmatrix}\mathbf{z}\\ \boldsymbol{\xi}\end{bmatrix} = \mathbf{G}_1\mathbf{z} + \mathbf{G}_2\boldsymbol{\xi}, \qquad \mathbf{y} = \mathbf{H}\begin{bmatrix}\mathbf{z}\\ \boldsymbol{\zeta}\end{bmatrix} = \mathbf{H}_1\mathbf{z} + \mathbf{H}_2\boldsymbol{\zeta},
$$

where $\mathbf{G} \in \mathbb{R}^{d_1 \times (K_1+K_2)}$ is the image dictionary with full rank $(K_1 + K_2)$, $\mathbf{H} \in \mathbb{R}^{d_2 \times (K_1+K_3)}$ is the text dictionary with full rank $(K_1 + K_3)$.

We verify Assumptions 3.1 for the specified distribution in Appendix H. The following theorem gives convergence guarantees for CLIP and provides the upper bound of its zero-shot transfer error.

Theorem 3.4. For sufficiently large n, set the learning rate $\eta = O(\epsilon\tau^2\|\mathbf{G}\|^{-2}\|\mathbf{H}\|_2^{-2}(1+R)^{-4})$, gradient descent can find $\widehat{\mathbf{W}}$ within $4\|\mathbf{W}^{(0)} - \mathbf{W}^*\|_F^2/(\eta\epsilon)$ iterations such that $L_{\mathcal{D}^B}(\widehat{f}, \tau) \leq L_{\mathcal{D}^B}(f^*, \tau) + \epsilon$ where $\widehat{f} = \langle \widehat{\mathbf{W}}\mathbf{x}, \mathbf{y}\rangle$. In addition, the top-$r$ zero-shot transfer error is bounded by $\epsilon'/\log(1+r)$, where $\epsilon' = (C_B + 2) \cdot \Big[\epsilon + 2B\exp(-\gamma/\tau)\Big]$ and $C_B = \widetilde{O}(K/B)$.

Square Loss Fails Zero-Shot Learning. Suppose square loss $\mathbb{E}[\|\mathbf{g}(\mathbf{x}) - \mathbf{y}\|_2^2]$ is used to learn the embedding $\mathbf{g}$, We find that even if we can train with population risk and get the Bayesian optimal predictor, the learned representation $\mathbf{g}$ is not suitable for the zero-shot transfer. We consider the data introduced in Definition 3.3 for the following.

Theorem 3.5. The Bayesian optimal representation $\mathbf{g}$ is $\mathbf{g}(\mathbf{x}) = \mathbf{H}\begin{bmatrix}\mathbf{z} \\ \mathbb{E}[\boldsymbol{\zeta}|\mathbf{z}]\end{bmatrix}$.

The following corollary formally states the negative result.

Corollary 3.6. For the distribution in Definition 3.3 with $\mathbf{H} = \begin{bmatrix}\mathbf{I} \\ \mathbf{0}\end{bmatrix}$, margin $\gamma < 1/3$, text unique feature $\boldsymbol{\zeta} \in \mathbb{R}^{K_3}$ drawn from $\{\mathbf{e}_1, \mathbf{e}_2\}$ with probability $1/3, 2/3$ respectively. Then, the zero-shot top-1 error is at least $1/(3K)$ under various similarity scores, including cosine similarity.

Remark 3.7. By Theorem 3.4, we can achieve arbitrarily small top-1 error by CLIP as long as ϵ and τ are sufficiently small. However, for the representation learned from the square loss, the top-1 error is at least a constant even if we can achieve the Beyasian optimal predictor.

4 Learn Better Representation via Regularization

In Corollary 3.2, we know that CLIP can achieve a small error for zero-shot transfer tasks. In this section, we investigate how large the margin can be achieved between different features $\mathbf{z}$'s. Under the same condition of Corollary 3.2, we present the following corollary.

Corollary 4.1. Suppose the result of Theorem G.1 holds for the learned similarity function $\widehat{f}$. We calculate the similarity score $\widehat{f}(\mathbf{x}, \mathbf{y}_k)$ for all $k \in [K]$. Then with probability at least $1 - 4\epsilon'$, the top-1 result gives the correct answer with a margin τ.

Here, the margin depends on the temperature parameter τ. Note that we only achieve the margin with τ instead of γ guaranteed in the Assumption 3.1. This indicates a theoretical gap in the learned margin.

Theorem 4.2. Under the same condition as Theorem 3.4, there exists a special case with initialization $\mathbf{W}^{(0)}$, such that when we train the model with polynomial iterations $T = \mathrm{poly}(\eta^{-1}, \epsilon, d_1, d_2)$, with probability at least 0.99, the top-1 result can only give the correct answer with a margin $\widetilde{O}(\tau)$.

Such a phenomenon also exists in real data: the margin will decrease when temperature τ decreases (see Figure 1). To obtain a larger margin, we propose to use the following regularization,

$$R(f) = -\frac{1}{|S|}\sum_{(\mathbf{x},\mathbf{y})\in S} f(\mathbf{x}, \mathbf{y}).$$

The following theorem shows that the regularization can improve the margin.

Theorem 4.3. Under the same condition as Theorem 4.2, with sufficiently small τ and appropriately chosen λ, within polynomial iterations $T = \mathrm{poly}(\eta^{-1}, \epsilon, d_1, d_2)$, we can find a score function $\widehat{f}$ with large margin. In particular, with a probability of at least 0.99, the top-1 result gives the correct label with a margin $\widetilde{\Omega}(\gamma)$.

5 Experiments

Datasets. For performance evaluation, we consider Conceptual Captions 3M (CC3M) (Sharma et al., 2018) and MSCOCO (Chen et al., 2015) as the pretraining datasets, in alignment with prior literature (Li et al., 2022; Goel et al., 2022).

Architectures. We consider the same setting for experiments on all baseline CLIP-objectives. Following the original CLIP paper, we employ ResNet (He et al., 2016) as the image encoder and the Transformer architecture (Vaswani et al., 2017) as the text encoder. We use pre-trained weights for both encoders. Detailed hyperparameters and additional experiments are presented in Appendix E.

5.1 Effect of Temperature on Margin

In support of our theoretical discussions in Corollary 4.1 and Theorem 4.2 that find the positive correlation between the margin and the temperature parameter, we conduct real data experiments. In Figure 1, we examine the margin distribution of CLIP models trained at varying temperatures. The margin is considered as the difference between a diagonal value and an off-diagonal value within a batch: $f(\mathbf{x}_i, \mathbf{y}_i) - f(\mathbf{x}_j, \mathbf{y}_i)$ and $f(\mathbf{x}_i, \mathbf{y}_i) - f(\mathbf{x}_i, \mathbf{y}_j)$. We collect the margins of untrained and trained CLIP models on all batches within the MSCOCO training dataset.

As depicted in Figure 1, a CLIP model with random initialization at the projection layers has margins normally distributed near zero, whereas trained models exhibit positive margins, signifying successful training. Furthermore, we consider CLIP models trained at fixed temperature values of 0.07 and 0.01. As observed in the figure, the margin distribution shifts to the left as temperature τ decreases, suggesting that an extremely small τ leads to small margins, aligning with the results in Corollary 4.1.

5.2 Zero-shot Transfer

To confirm Theorem 4.3, we investigate the advantages of incorporating our regularization term during training by evaluating zero-shot transfer accuracy and linear probing on various datasets. We consider the following training objectives when adding our regularization: (1) the original CLIP (Radford et al., 2021), and (2) CyCLIP (Goel et al., 2022) with cross-modal and in-modal consistency regularizations, adopting the same hyperparameters for the regularizations as outlined in Goel et al. (2022). All models are trained on CC3M using the same model architecture, batch size, and optimizer settings. Further experimental details are provided in Appendix E.

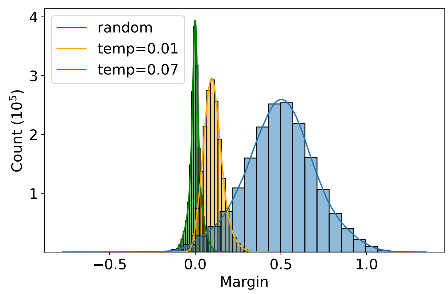

Figure 1: The distribution of the margins with regard to CLIP models trained at different temperature values. Margin is computed within each batch of the data.

In Table 1, we present the zero-shot test accuracy of CLIP models trained with the original CLIP objective and the CyCLIP objective. Firstly, we demonstrate the model's performance when training solely on the regularization term (L2) and compare to that of the CLIP objective. In alignment with our Corollary 3.6, we can observe on real data that training exclusively on the L2 objective leads to a large error and even random guessing on the zero-shot datasets. Combining with our theoretical analysis, we show that a naive square loss fails to learn transferable representations. In the context of multi-modal learning, contrastive loss is important. Moreover, confirming our result from Theorem 4.3, incorporating the regularization term into the contrastive objective effectively enhances performance across the majority of zero-shot transfer tasks. It improves over the baseline on 5 out of 6 datasets by a good margin. The best performance achieved by adding regularization to the CLIP objective outperforms its original objective by 3.62% on CIFAR10 and by 2.06% on average of all datasets.

Table 1: Zero-shot top-1 accuracy ($\%$). Notably, adding the regularization term successfully improves the baselines on 5 out of the 6 datasets.

	CIFAR10	CIFAR100	STL10	Food101	ImageNetV2	DTD	Average
Reg	10.04	1.05	9.95	1.08	0.11	2.07	3.47
CLIP	63.85	31.17	90.35	8.39	20.24	**21.22**	39.20
CyCLIP	60.71	28.87	89.98	9.72	19.66	20.21	38.19
CLIP+Reg	**67.47**	**33.33**	**92.64**	**12.14**	**22.36**	19.63	**41.26**

6 Conclusion

In this paper, we rigorously investigated the theoretical underpinnings of transferable representation learning in CLIP We provided insights through specific cases and corroborated our theory with empirical evidence. Lastly, we proposed a regularization term grounded in our theoretical findings to enhance CLIP's performance in zero-shot transfer.

References

Alayrac, J.-B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M. et al. (2022). Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems* **35** 23716–23736.

Bartlett, P. L. and Mendelson, S. (2002). Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research* **3** 463–482.

Bossard, L., Guillaumin, M. and Van Gool, L. (2014). Food-101–mining discriminative components with random forests. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part VI 13*. Springer.

Chen, X., Fang, H., Lin, T.-Y., Vedantam, R., Gupta, S., Dollár, P. and Zitnick, C. L. (2015). Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325* .

Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S. and Vedaldi, A. (2014). Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.

Coates, A., Ng, A. and Lee, H. (2011). An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings.

Desai, K. and Johnson, J. (2021). Virtex: Learning visual representations from textual annotations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.

Fukumizu, K., Bach, F. R. and Jordan, M. I. (2004). Dimensionality reduction for supervised learning with reproducing kernel hilbert spaces. *Journal of Machine Learning Research* **5** 73–99.

Fukumizu, K., Bach, F. R. and Jordan, M. I. (2009). Kernel dimension reduction in regression .

Gao, Y., Liu, J., Xu, Z., Zhang, J., Li, K., Ji, R. and Shen, C. (2022). Pyramidclip: Hierarchical feature alignment for vision-language model pretraining. *Advances in Neural Information Processing Systems* **35** 35959–35970.

Goel, S., Bansal, H., Bhatia, S., Rossi, R., Vinay, V. and Grover, A. (2022). Cyclip: Cyclic contrastive language-image pretraining. *Advances in Neural Information Processing Systems* **35** 6704–6719.

Gomez, L., Patel, Y., Rusinol, M., Karatzas, D. and Jawahar, C. (2017). Self-supervised learning of visual features through embedding images into text topic spaces. In *Proceedings of the ieee conference on computer vision and pattern recognition*.

HaoChen, J. Z., Wei, C., Gaidon, A. and Ma, T. (2021). Provable guarantees for self-supervised deep learning with spectral contrastive loss. *Advances in Neural Information Processing Systems* **34**.

He, K., Zhang, X., Ren, S. and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.

Huang, Y., Du, C., Xue, Z., Chen, X., Zhao, H. and Huang, L. (2021). What makes multi-modal learning better than single (provably). *Advances in Neural Information Processing Systems* **34** 10944–10956.

Jia, C., Yang, Y., Xia, Y., Chen, Y.-T., Parekh, Z., Pham, H., Le, Q., Sung, Y.-H., Li, Z. and Duerig, T. (2021). Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*. PMLR.

Krizhevsky, A. (2009). Learning multiple layers of features from tiny images. Tech. rep.

Lee, J. D., Lei, Q., Saunshi, N. and Zhuo, J. (2020). Predicting what you already know helps: Provable self-supervised learning. *arXiv preprint arXiv:2008.01064* .

LEE, J. D., LEI, Q., SAUNSHI, N. and ZHUO, J. (2021). Predicting what you already know helps: Provable self-supervised learning. *Advances in Neural Information Processing Systems* **34** 309–323.

LEI BA, J., SWERSKY, K., FIDLER, S. ET AL. (2015). Predicting deep zero-shot convolutional neural networks using textual descriptions. In *Proceedings of the IEEE international conference on computer vision*.

LI, Y., LIANG, F., ZHAO, L., CUI, Y., OUYANG, W., SHAO, J., YU, F. and YAN, J. (2022). Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. In *International Conference on Learning Representations*.

LIANG, P. P., DENG, Z., MA, M., ZOU, J., MORENCY, L.-P. and SALAKHUTDINOV, R. (2023). Factorized contrastive learning: Going beyond multi-view redundancy. *arXiv preprint arXiv:2306.05268* .

MITROVIC, J., MCWILLIAMS, B., WALKER, J., BUESING, L. and BLUNDELL, C. (2020). Representation learning via invariant causal mechanisms. *arXiv preprint arXiv:2010.07922* .

MU, N., KIRILLOV, A., WAGNER, D. and XIE, S. (2022). Slip: Self-supervision meets language-image pre-training. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVI*. Springer.

NAKADA, R., GULLUK, H. I., DENG, Z., JI, W., ZOU, J. and ZHANG, L. (2023). Understanding multimodal contrastive learning and incorporating unpaired data. In *International Conference on Artificial Intelligence and Statistics*. PMLR.

NILSBACK, M.-E. and ZISSERMAN, A. (2008). Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*. IEEE.

PARKHI, O. M., VEDALDI, A., ZISSERMAN, A. and JAWAHAR, C. (2012). Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*. IEEE.

PHAM, H., DAI, Z., GHIASI, G., KAWAGUCHI, K., LIU, H., YU, A. W., YU, J., CHEN, Y.-T., LUONG, M.-T., WU, Y. ET AL. (2021). Combined scaling for zero-shot transfer learning. *arXiv preprint arXiv:2111.10050* .

RADFORD, A., KIM, J. W., HALLACY, C., RAMESH, A., GOH, G., AGARWAL, S., SASTRY, G., ASKELL, A., MISHKIN, P., CLARK, J. ET AL. (2021). Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*. PMLR.

RECHT, B., ROELOFS, R., SCHMIDT, L. and SHANKAR, V. (2019). Do imagenet classifiers generalize to imagenet? In *International conference on machine learning*. PMLR.

SAITO, K., SOHN, K., ZHANG, X., LI, C.-L., LEE, C.-Y., SAENKO, K. and PFISTER, T. (2022). Prefix conditioning unifies language and label supervision. *arXiv preprint arXiv:2206.01125* .

SARIYILDIZ, M. B., PEREZ, J. and LARLUS, D. (2020). Learning visual representations with caption annotations. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VIII 16*. Springer.

SAUNSHI, N., ASH, J., GOEL, S., MISRA, D., ZHANG, C., ARORA, S., KAKADE, S. and KRISHNAMURTHY, A. (2022). Understanding contrastive learning requires incorporating inductive biases. *arXiv preprint arXiv:2202.14037* .

SAUNSHI, N., PLEVRAKIS, O., ARORA, S., KHODAK, M. and KHANDEPARKAR, H. (2019). A theoretical analysis of contrastive unsupervised representation learning. In *International Conference on Machine Learning*. PMLR.

SHARIATNIA, M. M. (2021). Simple CLIP.

SHARMA, P., DING, N., GOODMAN, S. and SORICUT, R. (2018). Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

THOMEE, B., SHAMMA, D. A., FRIEDLAND, G., ELIZALDE, B., NI, K., POLAND, D., BORTH, D. and LI, L.-J. (2016). Yfcc100m: The new data in multimedia research. *Communications of the ACM* **59** 64–73.

TIAN, Y., YU, L., CHEN, X. and GANGULI, S. (2020). Understanding self-supervised learning with dual deep networks. *arXiv preprint arXiv:2010.00578* .

TOSH, C., KRISHNAMURTHY, A. and HSU, D. (2021a). Contrastive estimation reveals topic posterior information to linear models. *Journal of Machine Learning Research* **22** 1–31.

TOSH, C., KRISHNAMURTHY, A. and HSU, D. (2021b). Contrastive estimation reveals topic posterior information to linear models. *Journal of Machine Learning Research* **22** 1–31.

TSAI, Y.-H. H., WU, Y., SALAKHUTDINOV, R. and MORENCY, L.-P. (2020). Demystifying self-supervised learning: An information-theoretical framework. *arXiv preprint arXiv:2006.05576* .

VASWANI, A., SHAZEER, N., PARMAR, N., USZKOREIT, J., JONES, L., GOMEZ, A. N., KAISER, Ł. and POLOSUKHIN, I. (2017). Attention is all you need. In *Advances in neural information processing systems*.

WANG, T. and ISOLA, P. (2020). Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*. PMLR.

WEN, Z. and LI, Y. (2021). Toward understanding the feature learning process of self-supervised contrastive learning. In *International Conference on Machine Learning*. PMLR.

WIGHTMAN, R. (2019). Pytorch image models. https://github.com/rwightman/pytorch-image-models.

WOLF, T., DEBUT, L., SANH, V., CHAUMOND, J., DELANGUE, C., MOI, A., CISTAC, P., RAULT, T., LOUF, R., FUNTOWICZ, M., DAVISON, J., SHLEIFER, S., VON PLATEN, P., MA, C., JERNITE, Y., PLU, J., XU, C., SCAO, T. L., GUGGER, S., DRAME, M., LHOEST, Q. and RUSH, A. M. (2020). Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Online.

WU, K., ZHANG, J., PENG, H., LIU, M., XIAO, B., FU, J. and YUAN, L. (2022). Tinyvit: Fast pretraining distillation for small vision transformers. In *European Conference on Computer Vision*. Springer.

YAO, L., HUANG, R., HOU, L., LU, G., NIU, M., XU, H., LIANG, X., LI, Z., JIANG, X. and XU, C. (2022). FILIP: Fine-grained interactive language-image pre-training. In *International Conference on Learning Representations*.

ZADEH, A., LIANG, P. P. and MORENCY, L.-P. (2020). Foundations of multimodal co-learning. *Information Fusion* **64** 188–193.

ZHAI, X., WANG, X., MUSTAFA, B., STEINER, A., KEYSERS, D., KOLESNIKOV, A. and BEYER, L. (2022). Lit: Zero-shot transfer with locked-image text tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

ZHANG, T. (2002). Covering number bounds of certain regularized linear function classes. *Journal of Machine Learning Research* **2** 527–550.

ZHANG, Y., JIANG, H., MIURA, Y., MANNING, C. D. and LANGLOTZ, C. P. (2022). Contrastive learning of medical visual representations from paired images and text. In *Machine Learning for Healthcare Conference*. PMLR.

A Related Work

Vision-Language Pre-Training. While labeled data are expensive and relatively scarce, images paired with text descriptions are available in much larger volumes (Thomee et al., 2016). Consequently, numerous studies (Gomez et al., 2017; Sariyildiz et al., 2020; Desai and Johnson, 2021; Zhang et al., 2022; Liang et al., 2023) have focused on leveraging free-form natural language supervision to learn visual representations. Recently, CLIP (Radford et al., 2021) and ALIGN (Jia et al., 2021) have emerged as prominent works extending contrastive learning to the vision-language pre-training framework. Built upon CLIP's success, several studies (Pham et al., 2021; Gao et al., 2022; Saito et al., 2022) have refined CLIP's contrastive methodology to better learn from web-scale image-text data. Notably, FILIP (Yao et al., 2022) introduces a fine-grained contrastive loss tailored for transformer architectures. DeCLIP (Li et al., 2022) and SLIP (Mu et al., 2022) additionally incorporate single-modality self-supervised learning. CyCLIP (Goel et al., 2022) introduces two regularizing terms enforcing cross-modal and in-modal consistency. LiT (Zhai et al., 2022) and Flamingo (Alayrac et al., 2022) consider training from pre-trained single-modality models. In our empirical validation of theoretical findings, we employ the same setting and train from pre-trained image and text encoders.

Theory of self-supervised learning. To understand self-supervised learning, numerous studies have been conducted, particularly focusing on *unimodal* contrastive learning, a widely used self-supervised learning approach rooted in data augmentation (Saunshi et al., 2019; Tsai et al., 2020; Mitrovic et al., 2020; Tian et al., 2020; Wang and Isola, 2020; Tosh et al., 2021a,b; HaoChen et al., 2021; Wen and Li, 2021; Saunshi et al., 2022). In multimodal learning, theoretical explanation has been explored in several studies (Zadeh et al., 2020; Huang et al., 2021; Lee et al., 2020; Nakada et al., 2023). These works have established that multimodal learning can surpass unimodal learning in terms of performance. For instance, Lee et al. (2020) employed square loss prediction to learn image representations under certain conditional independence assumptions, offering generalization performance guarantees. Meanwhile, Nakada et al. (2023) examined CLIP within specific linear representation settings and emphasized its correlation with singular value decomposition (SVD). We note that, these related works have not considered the zero-shot transfer mechanism and thus can't adequately explain the zero-shot transfer capability of CLIP.

B Preliminaries

Notation. We use lowercase letters, lowercase boldface letters, and uppercase boldface letters to denote scalars, vectors, and matrices, respectively. For a vector $\mathbf{x}$, we use $\|\mathbf{x}\|_2$ to denote its Euclidean norm. For a matrix $\mathbf{W}$, we use $\|\mathbf{W}\|_F$ to denote its Frobenius norm. Given two sequences $\{x_n\}$ and $\{y_n\}$, we denote $x_n = \mathcal{O}(y_n)$ if $|x_n| \le C_1|y_n|$ for some absolute positive constant C_1, $x_n = \Omega(y_n)$ if $|x_n| \ge C_2|y_n|$ for some absolute positive constant C_2, and $x_n = \Theta(y_n)$ if $C_3|y_n| \le |x_n| \le C_4|y_n|$ for some absolute constants $C_3, C_4 > 0$. We also use $\widetilde{\mathcal{O}}(\cdot)$ to hide logarithmic factors of d in $\mathcal{O}(\cdot)$. Additionally, we denote $x_n = \mathrm{poly}(y_n)$ if $x_n = \mathcal{O}(y_n^D)$ for some positive constant D, and $x_n = \mathrm{polylog}(y_n)$ if $x_n = \mathrm{poly}(\log(y_n))$. We also denote by $x_n = o(y_n)$ if $\lim_{n\to\infty} x_n/y_n = 0$. Finally we use $[N]$ to denote the index set $\{1, \dots, N\}$. In the function space, let $B_r(f)$ denote the ball of radius r centered at f, with the metrics $\|\cdot\|_\infty$. A set C is the covering of function class $\mathcal{F}$ with radius r, if and only if $\mathcal{F} \subseteq \cup_{f\in C} B_r(f)$. The covering number of $\mathcal{F}$ with radius r is the minimum cardinality of any covering of $\mathcal{F}$, denoted as $\mathcal{N}(\mathcal{F}, r)$.

Loss function of CLIP.

$$
\begin{aligned}
L_{S'}(f,\tau) &= \frac{1}{B}\sum_{i\in S'} -\log\left(\frac{\exp\big(f(\mathbf{x}_i,\mathbf{y}_i)/\tau\big)}{\sum_{j\in S'}\exp\big(f(\mathbf{x}_j,\mathbf{y}_i)/\tau\big)}\right) + \frac{1}{B}\sum_{i\in S'} -\log\left(\frac{\exp\big(f(\mathbf{x}_i,\mathbf{y}_i)/\tau\big)}{\sum_{j\in S'}\exp\big(f(\mathbf{x}_i,\mathbf{y}_j)/\tau\big)}\right)\\
&= \frac{1}{B}\sum_{i\in S'}\log\left(\sum_{j\in S'}\exp\big(\big[f(\mathbf{x}_j,\mathbf{y}_i) - f(\mathbf{x}_i,\mathbf{y}_i)\big]/\tau\big)\right)\\
&\quad + \frac{1}{B}\sum_{i\in S'}\log\left(\sum_{j\in S'}\exp\big(\big[f(\mathbf{x}_i,\mathbf{y}_j) - f(\mathbf{x}_i,\mathbf{y}_i)\big]/\tau\big)\right), \qquad \text{(B.1)}
\end{aligned}
$$

where $\tau > 0$ is a temperature parameter. The training loss $L_{S'}$ over a single epoch can be viewed as the empirical version of (2.1).

Remark B.1. In Assumption 2.1, the assumption of conditional independence is frequently made in the analysis of self-supervised learning (Saunshi et al., 2019; Lee et al., 2021) and dimension reduction algorithms (Fukumizu et al., 2004, 2009). Under the premise that $\mathbf{x}, \mathbf{y}$ are conditionally independent (CI) given $\mathbf{z}$, it can be posited that any additional patterns found within $\mathbf{x}|\mathbf{z}$ and $\mathbf{y}|\mathbf{z}$ should be interpreted as unique features. Notably, in the absence of discrete and sparse constraints, a suitable $\mathbf{z}$ can always be found, given that one could simply assign $\mathbf{z} = \mathbf{x}$ or $\mathbf{z} = \mathbf{y}$. From the generative model's point of view, Assumption 2.1 naively holds when the data are from some generator with $\mathbf{x} = T_1(\mathbf{z}, \boldsymbol{\xi})$ and $\mathbf{y} = T_2(\mathbf{z}, \boldsymbol{\zeta})$ where $\boldsymbol{\xi} \perp \boldsymbol{\zeta}|\mathbf{z}$.

Given the population loss in (2.1), the following theorem shows that the empirical loss $\widehat{\mathbb{E}}_S(f, \tau) := (1/n)\sum_{k\in[n]} L_{S_k}(f, \tau)$ concentrates on the population loss when n is large enough.

Theorem B.2. Suppose $\delta \in (0, 1)$ and $n \geq (8\tau^{-1}\epsilon^{-2}M \log B) \log(2\mathcal{N}(\mathcal{F}, \epsilon/8M)/\delta)$, then with probability at least $1 - \delta$, we have

$$|\widehat{L}_S(f, \tau) - L_{\mathcal{D}^B}(f, \tau)| \leq \epsilon$$

for all function $f \in \mathcal{F}$ and $|f| \leq M$, where $\mathcal{N}(\mathcal{F}, \epsilon)$ is the covering number of $\mathcal{F}$.

Theorem B.2 shows that the generalization gap $|\widehat{L}_S(f, \tau) - L_{\mathcal{D}^B}(f, \tau)|$ approaches zero as the number of batches n increase. In practice, the batch size is limited by the GPU's memory and is smaller than the number of batches (or the number of training examples). Therefore, instead of letting the batch size B go to infinity like in prior studies (Wang and Isola, 2020; Pham et al., 2021), we keep the batch size B as a constant in (2.1) and Theorem B.2 to enable the analysis of CLIP even for small batches. Pham et al. (2021) also provided the generalization gap for CLIP. However, their result is for $B \to \infty$ and a loss function without the log term, i.e., $\exp\big(f(\mathbf{x}_i, \mathbf{y}_i)/\tau\big)/\Big(\sum_{j\in S'} \exp\big(f(\mathbf{x}_j, \mathbf{y}_i)/\tau\big)\Big)$.

C Discussion on the Margin

In Assumption 3.1, we introduce the (α, β, γ) completeness. In this section, we will discuss how to verify the assumption and formally measure the quality of the learned function.

Sample two independent tuple $(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mathcal{D}_{\mathbf{x}\times\mathbf{y}\times\mathbf{z}}$ and $(\mathbf{x}', \mathbf{y}', \mathbf{z}') \sim \mathcal{D}_{\mathbf{x}\times\mathbf{y}\times\mathbf{z}}$, we introduce a measure as follows.

$$\alpha_\gamma = \mathbb{P}\Big(\mathbf{z} \neq \mathbf{z}', f(\mathbf{x}, \mathbf{y}) - f(\mathbf{x}, \mathbf{y}') \leq \gamma\Big) + \mathbb{P}\Big(\mathbf{z} \neq \mathbf{z}', f(\mathbf{x}, \mathbf{y}) - f(\mathbf{x}', \mathbf{y}) \leq \gamma\Big)$$

By Assumption 3.1, we know that we want to find a large γ with small α_γ.

However, in real applications, we can access $\mathcal{D}_{\mathbf{x}\times\mathbf{y}}$ but have little knowledge of the model $\mathcal{V}$ and the latent variable $\mathbf{z}$. Thus, we introduce another measure $\widehat{\alpha}_\gamma$ as follows,

$$\widehat{\alpha}_\gamma = \mathbb{P}\Big(f(\mathbf{x}, \mathbf{y}) - f(\mathbf{x}, \mathbf{y}') \leq \gamma\Big) + \mathbb{P}\Big(f(\mathbf{x}, \mathbf{y}) - f(\mathbf{x}', \mathbf{y}) \leq \gamma\Big)$$

$\widehat{\alpha}_\gamma$ differs from the α_γ since we didn't extinguish different classes in the probability. Therefore we can easily calculate $\widehat{\alpha}_\gamma$ without observe $\mathbf{z}$. Besides, we have the following upper and low bounds, which show that $\widehat{\alpha}_\gamma$ can approximate α_γ.

Theorem C.1. Let $\gamma \geq 0$, then we have that

$$\widehat{\alpha}_\gamma \geq \alpha_\gamma \geq \widehat{\alpha}_\gamma - \sum_{k\in[K]} p_k^2.$$

where p_k is the probability of the classes in Assumption 2.1. Besides the second inequality become exact equality for $\gamma = 0$.

Proof.

$$\begin{aligned}
\widehat{\alpha}_\gamma &= \mathbb{P}\Big(f(\mathbf{x}, \mathbf{y}) - f(\mathbf{x}, \mathbf{y}') \leq \gamma\Big) + \mathbb{P}\Big(f(\mathbf{x}, \mathbf{y}) - f(\mathbf{x}', \mathbf{y}) \leq \gamma\Big)\\
&= \underbrace{\mathbb{P}\Big(\mathbf{z} \neq \mathbf{z}', f(\mathbf{x}, \mathbf{y}) - f(\mathbf{x}, \mathbf{y}') \leq \gamma\Big) + \mathbb{P}\Big(\mathbf{z} \neq \mathbf{z}', f(\mathbf{x}, \mathbf{y}) - f(\mathbf{x}', \mathbf{y}) \leq \gamma\Big)}_{=\alpha_\gamma}\\
&\quad + \underbrace{\mathbb{P}\Big(\mathbf{z} = \mathbf{z}', f(\mathbf{x}, \mathbf{y}) - f(\mathbf{x}, \mathbf{y}') \leq \gamma\Big) + \mathbb{P}\Big(\mathbf{z} = \mathbf{z}', f(\mathbf{x}, \mathbf{y}) - f(\mathbf{x}', \mathbf{y}) \leq \gamma\Big)}_{\mathrm{ApproximateError}}.
\end{aligned}$$

The Approximate Error has a naive lower bound of 0 and we can upper bound it as follows

$$\begin{aligned}\mathbb{P}\Big(\mathbf{z}=\mathbf{z}', f(\mathbf{x},\mathbf{y})-f(\mathbf{x},\mathbf{y}')\leq\gamma\Big) &= \mathbb{P}\Big(f(\mathbf{x},\mathbf{y})-f(\mathbf{x},\mathbf{y}')\leq\gamma|\mathbf{z}=\mathbf{z}'\Big)\cdot\mathbb{P}(\mathbf{z}=\mathbf{z}')\\ &\leq \mathbb{P}\Big(f(\mathbf{x},\mathbf{y})-f(\mathbf{x},\mathbf{y}')\leq 0|\mathbf{z}=\mathbf{z}'\Big)\cdot\mathbb{P}(\mathbf{z}=\mathbf{z}')\\ &= 1/2\sum_{k\in[K]}p_k^2.\end{aligned}$$

were the the inequality is due to fact that $\gamma\geq 0$ and the last equality is because $\mathbf{y}'$ and $\mathbf{y}$ are symmetric give $\mathbf{z}=\mathbf{z}'$. Finally, the inequality is an exact equality for $\gamma=0$. □

By Theorem C.1, α_γ and $\widehat{\alpha}_\gamma$ are close to each other if $\max_{k\in[K]}p_k$ is small, since

$$\sum_{k\in[K]}p_k^2\leq\sum_{k\in[K]}p_k\cdot\max_{k\in[K]}p_k=\max_{k\in[K]}p_k\cdot\Big(\sum_{k\in[K]}p_k\Big)=\max_{k\in[K]}p_k.$$

Relation with the Figure 2: $\widehat{\alpha}_\gamma$ has a strong relationship with Figure 2, where we have plot the distribution of $f(\mathbf{x},\mathbf{y})-f(\mathbf{x},\mathbf{y}')$ and $f(\mathbf{x},\mathbf{y})-f(\mathbf{x}',\mathbf{y})$. The figure can be viewed as the figure of the probability density function and $\widehat{\alpha}_\gamma$ can be viewed as the cumulative probability function which is the integral of probability mass smaller than γ. From Figure 2, we can deduce that the CLIP learned with regularization has consistently smaller $\widehat{\alpha}_\gamma$ for all $\gamma\geq 0$.

D Discussion on the Trainable Temperature Parameter τ

This section considers the setting where the temperature τ is also trainable with the following loss.

$$\begin{aligned}L_{\mathcal{D}^B}(f,\tau)=\mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]}\exp\big(\big[f(\mathbf{x}_1,\mathbf{y}_t)-f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)\bigg)\bigg]\\ +\mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]}\exp\big(\big[f(\mathbf{x}_t,\mathbf{y}_1)-f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)\bigg)\bigg].\end{aligned}$$

Suppose τ is clipped to be within the range $[\tau_{\min},\tau_{\max}]$, it is natural to assume that we can obtain function $\widehat{f}$ with temperature $\widehat{\tau}\in[\tau_{\min},\tau_{\max}]$ such that

$$\begin{aligned}L_{\mathcal{D}^B}(\widehat{f},\widehat{\tau}) &\leq \min_{\tau\in[\tau_{\min},\tau_{\max}]}L_{\mathcal{D}^B}(f^*,\tau)+\epsilon &\text{(D.1)}\\ &= L_{\mathcal{D}^B}(f^*,\widehat{\tau})+\epsilon-\Big(L_{\mathcal{D}^B}(f^*,\widehat{\tau})-\min_{\tau\in[\tau_{\min},\tau_{\max}]}L_{\mathcal{D}^B}(f^*,\tau)\Big) &\text{(D.2)}\\ &= L_{\mathcal{D}^B}(f^*,\widehat{\tau})+\widetilde{\epsilon} &\text{(D.3)}\end{aligned}$$

where $\widetilde{\epsilon}=\epsilon-\Big(L_{\mathcal{D}^B}(f^*,\widehat{\tau})-\min_{\tau\in[\tau_{\min},\tau_{\max}]}L_{\mathcal{D}^B}(f^*,\tau)\Big)\leq\epsilon$. Since $\widetilde{\epsilon}$ is smaller than ϵ, we can get smaller ϵ' in Theorem G.1, and thus get smaller top-r error in zero-shot transfer task by Corollary 3.2. This observation implies that the representation $(\widehat{f},\widehat{\tau})$ found by trainable temperature can be better than the representation $(\widehat{f}',\widehat{\tau})$ found with fixed temperature $\widehat{\tau}$.

E Additional Experiment Results

We consider the same model architecture as CLIP (Radford et al., 2021) and consider ResNet-50 (He et al., 2016) as the image encoder and transformer (Vaswani et al., 2017) achitecture as the text encoder. Specifically, we use pre-trained weights for the encoders for faster convergence in training. We follow the code framework in Shariatnia (2021) and use pre-trained ResNet-50 from the PyTorch Image Models library (Wightman, 2019) and pre-trained DistilBERT from the Huggingface Transformers library (Wolf et al., 2020). We further have linear projection layers on both image and text encoder as the same as in CLIP and consider embedding dimension of 512. As we are training at small-scale data with pre-trained encoders, we follow Shariatnia (2021) and use AdamW optimizer with learning rate 1e-4 on the image encoder, 1e-5 on the text encoder, and 1e-3 on the projection layers, with weight decay coefficient 1e-3. Our code is provided anonymously on Github[1].

[1] https://anonymous.4open.science/r/CLIP_theory-BC8F/README.md

E.1 Effect of Temperature on Margin

Setup. For lightweight exploration in section 5.1, we use the training dataset from MSCOCO (Chen et al., 2015) Image Captioning Task as the data for vision-language contrastive pre-training. Specifically, the dataset contains $82,783$ images where each image is coupled with 5 captions. We consider each image-caption pair as a data example in pre-training and therefore arrive at $413,915$ pre-training data pairs. We further randomly split the data to keep 20% of the data as validation set and stops training as the contrastive loss on validation data no longer decreases to avoid overfitting on the small dataset.

Margin. Given a training data batch, the margin is consider as the difference between a diagonal value and an off-diagonal value: $f(\mathbf{x}_i, \mathbf{y}_i) - f(\mathbf{x}_j, \mathbf{y}_i)$ and $f(\mathbf{x}_i, \mathbf{y}_i) - f(\mathbf{x}_i, \mathbf{y}_j)$. We consider CLIP models trained at fixed temperature $\tau = 0.07$ and $\tau = 0.01$. We note that 0.07 is the default value for τ to start training in CLIP and 0.01 is the clamping value (equivalently as the maximum logit scale of 4.6052.) In Figure 1, we collected the margins from all batches of size 64 in the MSCOCO training data, where the data is randomly shuffled.

Additional Experiments. Here, we additionally compare the margin distribution of CLIP trained at temperature $\tau = 0.01$, without or with our regularization term. We could observe that the margin distribution shifts to the right with the regularization term, which alleviates the negative influence of an extremely small temperature value.

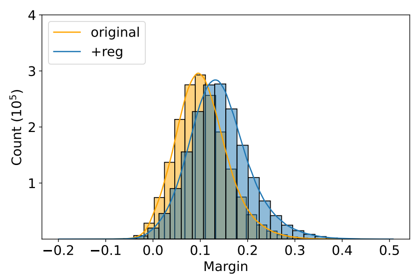

Figure 2: The distribution of the margins with regard to CLIP models trained $\tau = 0.01$ with or withour regularization. Margin is computed within each batch of the data.

E.2 Zero-shot Transfer and Linear Probing

Setup. In the evaluation of zero-shot transfer and linear probing, we use CC3M (Sharma et al., 2018) as the pre-training dataset, which contains around $3,318,332$ image-caption pairs gathered from the web. While some URLs are broken so that we cannot download the images, we eventually reached a pre-training dataset of $2,786,288$ data pairs. When training CLIP models, we use the default coefficients of CyCLIP regularization terms of $\lambda_1 = 0.25$ and $\lambda_2 = 0.25$. For our regularization term, we use a coefficient of $\lambda = 0.1$. As in CLIP, we set the temperature τ from 0.07, equivalently having maximum logit scale at 2.6593. Lastly, we use a training batch size of 32 and trained for 8 epochs in the results reported in section 5.2.

Table 2: Summary of datasets used for zero-shot transfer and linear probing.

Dataset	Classes	Class Description
CIFAR10	10	Categories of animals and vehicles
CIFAR100	100	Categories of objects including animals, foods, vehicles and people
STL10	10	Categories of animals and vehicles
Food101	101	Categories of foods/dishes
ImageNetV2	1000	Categories of objects including animals, foods, vehicles and people
DTD	47	Categories of textures
Flowers102	102	Categories of flower species
Oxford-IIIT Pet	37	Categories of cats and dogs

Evaluations. As similar in previous works (Radford et al., 2021; Yao et al., 2022; Mu et al., 2022; Goel et al., 2022), we consider the following image classification tasks for zero-shot transfer and linear probing: CIFAR10/100 (Krizhevsky, 2009), STL10 (Coates et al., 2011), Food101 (Bossard et al., 2014), ImageNetV2 (Recht et al., 2019), DTD (Describable Textures,Cimpoi et al. (2014)), Flowers102 (Nilsback and Zisserman, 2008) and Oxford-IIIT Pet (Parkhi et al., 2012). The dataset statistics are reported in Table 2. For zero-shot transfer, we use the same prompt engineering and ensembling as the original CLIP and report the top-1 accuracy. For linear probing, as the same in CLIP, we train a logistic regression classifier on the image embeddings generated by the image encoder of pre-trained CLIP models on the training data from the considered datasets. The classifiers are all trained to convergence and we report the test accuracy on each of the test dataset of the tasks. We note that, due to the limitation of the training data CC3M, the zero-shot test accuracy of all CLIP-objectives on Flowers102 and Oxford-IIIT Pet are near random guesses. Therefore, we omit these datasets for zero-shot transfer.

Additional Experiments. In Table 3, we report the results of linear probing, where logistic regression classifiers are fitted to the embeddings learned by the image encoders of our compared models. This table offers an assessment of the visual representation learning for each training objective. Similarly supporting Corollary 3.6, training on the regularization term only results in learning bad representations that yield unsatisfactory performances on linear probing. Moreover, in alignment with Theorem 4.3, we observe that adding the regularization term consistently improves CLIP's performance across various datasets by an average of 1.54%.

Table 3: Linear probing accuracy ($\%$). All logistic regression models are trained till convergence. Adding our regularization term to CLIP provides decent improvements across all datasets. On CyCLIP, we also makes improvements on the majority of datasets.

	CIFAR10	CIFAR100	STL10	Food101	DTD	Flowers	OxfordPets	Average
Reg	14.09	2.17	17.86	1.73	3.40	2.18	4.12	6.51
CLIP	87.30	66.03	93.26	62.8	56.70	70.24	72.91	72.75
CyCLIP	86.31	63.93	93.69	61.57	56.86	70.56	70.46	71.91
CLIP+Reg	**88.49**	**66.16**	**94.98**	**63.39**	**57.66**	**72.21**	**77.13**	**74.29**

We additionally report the zero-shot transfer results of the original CLIP objective and adding our regularziation term, on a different visual encoder architecture of TinyViT (Wu et al., 2022) with pre-trained weights from Huggingface.

Table 4: Zero-shot top-1 accuracy ($\%$). Notably, adding the regularization term successfully improves the baselines on 5 out of the 6 datasets.

	CIFAR10	CIFAR100	STL10	Food101	ImageNetV2	DTD	Average
CLIP	52.02	15.57	81.89	7.92	**16.91**	**11.80**	31.02
CLIP+Reg	**53.30**	**19.67**	**83.76**	**7.99**	16.06	11.53	**32.05**

F Proof of Results in Section 2

Proof of Theorem B.2. We first prove that $L_{S'}(f,\tau)$ is upper bounded by $4M\log B/\tau$.

$$\begin{aligned} L_{S'}(f,\tau) &= \frac{1}{B}\sum_{i\in S'}\log\bigg(\sum_{j\in S'}\exp\big([f(\mathbf{x}_j,\mathbf{y}_i)-f(\mathbf{x}_i,\mathbf{y}_i)]/\tau\big)\bigg) \\ &\quad + \frac{1}{B}\sum_{i\in S'}\log\bigg(\sum_{j\in S'}\exp\big([f(\mathbf{x}_i,\mathbf{y}_j)-f(\mathbf{x}_i,\mathbf{y}_i)]/\tau\big)\bigg) \\ &\leq \frac{1}{B}\sum_{i\in S'}\log\bigg(\sum_{j\in S'}\exp\big(2M/\tau\big)\bigg) + \frac{1}{B}\sum_{i\in S'}\log\bigg(\sum_{j\in S'}\exp\big(2M/\tau\big)\bigg) \\ &= 4M\log B/\tau. \end{aligned} \tag{F.1}$$

where the inequality is by the fact the $|f|\leq M$. On the other hand, we have that

$$\begin{aligned} L_{S'}(f,\tau) &= \frac{1}{B}\sum_{i\in S'}\log\bigg(\sum_{j\in S'}\exp\big([f(\mathbf{x}_j,\mathbf{y}_i)-f(\mathbf{x}_i,\mathbf{y}_i)]/\tau\big)\bigg) \\ &\quad + \frac{1}{B}\sum_{i\in S'}\log\bigg(\sum_{j\in S'}\exp\big([f(\mathbf{x}_i,\mathbf{y}_j)-f(\mathbf{x}_i,\mathbf{y}_i)]/\tau\big)\bigg) \\ &\geq \frac{2}{B}\sum_{i\in S'}\log\bigg(\exp\big([f(\mathbf{x}_i,\mathbf{y}_i)-f(\mathbf{x}_i,\mathbf{y}_i)]/\tau\big)\bigg) \\ &\geq 0. \end{aligned}$$

where the inequality is because Exp function is greater than 0. Therefore we have proved that $L_{S'}(f, \tau) \in (0, 4M\log(B)/\tau]$. For all $f_1, f_2 \in \mathcal{F}$ and any batch S' with size B, we have that

$$
\begin{aligned}
L_{S'}(f_1, \tau) - L_{S'}(f_2, \tau) &= \frac{1}{B}\sum_{i\in S'}\log\bigg(\sum_{j\in S'}\exp\big([f_1(\mathbf{x}_j, \mathbf{y}_i) - f_1(\mathbf{x}_i, \mathbf{y}_i)]/\tau\big)\bigg) \\
&\quad - \frac{1}{B}\sum_{i\in S'}\log\bigg(\sum_{j\in S'}\exp\big([f_2(\mathbf{x}_j, \mathbf{y}_i) - f_2(\mathbf{x}_i, \mathbf{y}_i)]/\tau\big)\bigg) \\
&\quad + \frac{1}{B}\sum_{i\in S'}\log\bigg(\sum_{j\in S'}\exp\big([f_1(\mathbf{x}_i, \mathbf{y}_j) - f_1(\mathbf{x}_i, \mathbf{y}_i)]/\tau\big)\bigg) \\
&\quad - \frac{1}{B}\sum_{i\in S'}\log\bigg(\sum_{j\in S'}\exp\big([f_2(\mathbf{x}_i, \mathbf{y}_j) - f_2(\mathbf{x}_i, \mathbf{y}_i)]/\tau\big)\bigg) \\
&\le \frac{1}{B}\sum_{i\in S'}\log\bigg(\sum_{j\in S'}\exp\big([f_1(\mathbf{x}_j, \mathbf{y}_i) - f_1(\mathbf{x}_i, \mathbf{y}_i)]/\tau\big)\bigg) \\
&\quad - \frac{1}{B}\sum_{i\in S'}\log\bigg(\sum_{j\in S'}\exp\big([f_1(\mathbf{x}_j, \mathbf{y}_i) - f_1(\mathbf{x}_i, \mathbf{y}_i) - 2\|f_1 - f_2\|_\infty]/\tau\big)\bigg) \\
&\quad + \frac{1}{B}\sum_{i\in S'}\log\bigg(\sum_{j\in S'}\exp\big([f_1(\mathbf{x}_i, \mathbf{y}_j) - f_1(\mathbf{x}_i, \mathbf{y}_i)]/\tau\big)\bigg) \\
&\quad - \frac{1}{B}\sum_{i\in S'}\log\bigg(\sum_{j\in S'}\exp\big([f_1(\mathbf{x}_i, \mathbf{y}_j) - f_1(\mathbf{x}_i, \mathbf{y}_i) - 2\|f_1 - f_2\|_\infty]/\tau\big)\bigg) \\
&= 4\|f_1 - f_2\|_\infty/\tau.
\end{aligned}
$$

Similarly, we can get another direction $L_{S'}(f_2, \tau) - L_{S'}(f_1, \tau) \le 4\|f_1 - f_2\|_\infty/\tau$, which yields to $|L_{S'}(f_2, \tau) - L_{S'}(f_1, \tau)| \le 4\|f_1 - f_2\|_\infty/\tau$. Taking the expectation gives that $|L_{\mathcal{D}^B}(f_2, \tau) - L_{\mathcal{D}^B}(f_1, \tau)| \le 4\|f_1 - f_2\|_\infty/\tau$. By the definition of the covering set, the function class $\mathcal{F}$ can be covered by K subsets $\mathcal{B}_1, \ldots, \mathcal{B}_K$, that is $\mathcal{F} = \mathcal{B}_1 \cup \ldots \cup \mathcal{B}_K$, where $K = \mathcal{N}(\mathcal{F}, \tau\epsilon/16)$ and $\mathcal{B}_1, \ldots \mathcal{B}_K$ are the balls of the radius $\tau \cdot \epsilon/16$ centered at $f_1, \ldots, f_K$. Then we have that

$$
\begin{aligned}
&\mathbb{P}_{S\sim\mathcal{D}^n}\Big[\sup_{f\in\mathcal{F}}\big|L_{\mathcal{D}^B}(f, \tau) - \widehat{L}_S(f, \tau)\big| \ge \epsilon\Big] \\
&\le \sum_{k\in[K]}\mathbb{P}_{S\sim\mathcal{D}^n}\Big[\sup_{f\in\mathcal{B}_k}\big|L_{\mathcal{D}^B}(f, \tau) - \widehat{L}_S(f, \tau)\big| \ge \epsilon\Big] \\
&\le \sum_{k\in[K]}\mathbb{P}_{S\sim\mathcal{D}^n}\Big[\big|L_{\mathcal{D}^B}(f_k, \tau) - \widehat{L}_S(f_k, \tau)\big| \ge \epsilon/2\Big] \\
&= \sum_{k\in[K]}\mathbb{P}_{S\sim\mathcal{D}^n}\Big[\big|L_{\mathcal{D}^B}(f_k, \tau) - (1/n)\sum_{i\in[n]}L_{S_i}(f_k, \tau)\big| \ge \epsilon/2\Big] \\
&\le 2K\exp\Big(-\frac{n\epsilon^2\tau}{8M\log B}\Big) \\
&= 2\mathcal{N}(\mathcal{F}, \tau\epsilon/16)\exp\Big(-\frac{n\epsilon^2\tau}{8M\log B}\Big),
\end{aligned} \tag{F.2}
$$

the first inequality is by union bound, the second is by triangle inequality, and the third is by Hoeffding's inequality and (F.1). Finally, plugging the condition $n \ge (8\tau^{-1}\epsilon^{-2}M\log B)\log(2\mathcal{N}(\mathcal{F}, \epsilon/8M)/\delta)$ into (F.2) we have that

$$
\mathbb{P}_{S\sim\mathcal{D}^n}\Big[\sup_{f\in\mathcal{F}}\big|L_{\mathcal{D}^B}(f, \tau) - \widehat{L}_S(f, \tau)\big| \ge \epsilon\Big] \le \delta,
$$

which completes the proof. $\square$

G Transferrable Representation Learning

Discussion on Assumption 3.1. In simple terms, Assumption 3.1 is made on the data distribution to allow the *existence* of good encoding functions $\mathbf{g}^*$ and $\mathbf{h}^*$. Specifically, the first bullet guarantees that the data with different $\mathbf{z}$, the underlying shared feature, is well distinguishable with margin γ. If the data from different $\mathbf{z}$ does not satisfy this condition, the majority of the diagonal term $f(\mathbf{x}_i, \mathbf{y}_i)$ in (B.1) can be smaller than the off-diagonal term $f(\mathbf{x}_j, \mathbf{y}_i)$, which is not favored by the mechanism of CLIP. In other words, all encoding functions may yield higher similarity score for negative pairs than positive pairs The second bullet requires the similarity score within each underlying shared feature not vary too much, which is naturally satisfied if the learned embeddings $\mathbf{g}(\mathbf{x}), \mathbf{h}(\mathbf{y})$ are consistent and do not vary too much given the same $\mathbf{z}$.

Theorem G.1 (Formal)**.** Suppose Assumption 3.1 hold and we can find an ϵ approximate minimum $\widehat{f} \in \mathcal{F}$ with respect to the temperature τ such that $\widehat{f}$ is bounded by M and

$$L_{\mathcal{D}^B}(\widehat{f}, \tau) \leq L_{\mathcal{D}^B}(f^*, \tau) + \epsilon. \tag{G.1}$$

Then the following results hold:

1. For $(\mathbf{x}, \mathbf{z}) \sim \mathcal{D}_{\mathbf{x}\times\mathbf{z}}$, $\{\mathbf{y}_k \sim \mathcal{D}_{\mathbf{y}|\mathbf{v}_k}, k \in [K]\}$, let $\mathbf{y}^* = \sum_{k\in[K]} \mathbb{1}(\mathbf{z} = \mathbf{v}_k)\mathbf{y}_k$, we have

$$\mathbb{E}\bigg[\log\bigg(\sum_{k\in[K]} \exp\big(\big[\widehat{f}(\mathbf{x}, \mathbf{y}_k) - \widehat{f}(\mathbf{x}, \mathbf{y}^*)\big]/\tau\big)\bigg)\bigg] \leq \epsilon'. \tag{G.2}$$

2. For $(\mathbf{y}, \mathbf{z}) \sim \mathcal{D}_{\mathbf{y}\times\mathbf{z}}$,$\{\mathbf{x}_k \sim \mathcal{D}_{\mathbf{x}|\mathbf{v}_k}, k \in [K]\}$, let $\mathbf{x}^* = \sum_{k\in[K]} \mathbb{1}(\mathbf{z} = \mathbf{v}_k)\mathbf{x}_k$, we have

$$\mathbb{E}\bigg[\log\bigg(\sum_{k\in[K]} \exp\big(\big[\widehat{f}(\mathbf{x}_k, \mathbf{y}) - \widehat{f}(\mathbf{x}^*, \mathbf{y})\big]/\tau\big)\bigg)\bigg] \leq \epsilon'. \tag{G.3}$$

3. For $(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mathcal{D}_{\mathbf{x}\times\mathbf{y}\times\mathbf{z}}$, variance $\mathbb{E}_{(\mathbf{y},\mathbf{z})}\big[\mathrm{Var}_{\mathbf{x}|\mathbf{z}}(\widehat{f}(\mathbf{x}, \mathbf{y}))\big] + \mathbb{E}_{(\mathbf{x},\mathbf{z})}\big[\mathrm{Var}_{\mathbf{y}|\mathbf{z}}(\widehat{f}(\mathbf{x}, \mathbf{y}))\big] \leq 16M^2\epsilon'$.

where $\epsilon' = (C_B+2)\cdot\big[\epsilon + C\tau^{-1}MB\alpha + C\tau^{-1}(\beta MB)^{1/3} + 2B\exp(-\gamma/\tau)\big]$ and $C = \widetilde{O}(1), C_B = \widetilde{O}(\max_k p_k^{-1}/B)$.

Remark G.2. Theorem G.1 establishes a soft margin between CLIP's learned embeddings on data of different $\mathbf{z}$'s. For instance, if an image $\mathbf{x}$ has a shared feature $\mathbf{z} = \mathbf{v}_1$, we have its accurate description $\mathbf{y}^* = \sum_{k\in[K]} \mathbb{1}(\mathbf{z} = \mathbf{v}_k)\mathbf{y}_k = \mathbf{y}_1$. From (G.2), it follows that $\log\Big(\sum_{k\in[K]} \exp\big(\big[\widehat{f}(\mathbf{x}, \mathbf{y}_k) - \widehat{f}(\mathbf{x}, \mathbf{y}_1)\big]/\tau\big)\Big)$ is small. This can only occur when $\widehat{f}(\mathbf{x}, \mathbf{y}_k) < \widehat{f}(\mathbf{x}, \mathbf{y}_1)$ for all $k \geq 2$, i.e., the trained model always yield higher similarity score for this image-text pair as compared to all other texts generated on different topics. This outcome aligns with the expectation that image-text pairs with the same shared feature will yield the highest similarity score.

Remark G.3 (Choice of temperature parameter)**.** In Theorem G.1, when the data is well separated (i.e., $\alpha, \beta = 0$), a smaller temperature will invariably lead to a smaller ϵ' and, consequently, better performance. In practice, τ is typically set to be 0.01, a sufficiently small value that ensures the term $\exp(-\gamma/\tau)$ is less than 0.0000454 for $\gamma = 0.1$. However, when the data is nonseparable (i.e., α and β exceed 0), a balance must be struck between the terms related to τ. As a consequence, τ should not be too small. A reasonable choice would be $\tau = O(\gamma/\log(B/\epsilon))$.

Remark G.4 (Batch size)**.** In Theorem G.1, while we do not demand an increasing batch size B, our analysis does suggest a preference for larger batch sizes, as they can reduce the constant C_B and consequently ϵ'.

Lemma G.5. For $b_j \geq 0, j \in [m]$, we have that

$$\log\bigg(1 + \sum_{j\in[m]} b_j\bigg) \leq \sum_{j\in[m]} \log(1 + b_j).$$

Proof. Notice that

$$\Pi_{j\in[J]}(1 + b_j) \geq 1 + \sum_{j\in[J]} b_j.$$

Taking the logarithm over both sides completes the proof. ∎

Lemma G.6. Suppose that $a_1, \dots a_m$ are i.i.d random variable sample lies in $[-R, R]$ where $R \geq 1$, with mean $\mu := \mathbb{E}[a_1]$ and variance $\sigma^2 := \mathbb{E}[(a_1 - \mathbb{E}[a_1])^2]$. Then we have that

$$\mathbb{E}[\log\Big(\sum_{i=1}^{m} \exp(a)\Big) \geq \log(m) + \mu + \frac{m-1}{4mR^2}\sigma^2.$$

Proof. Let $\bar{a} = \big[\sum_{i=1}^{m} a_i\big]/m$

$$\begin{aligned}
\log\Big(\sum_{i=1}^{m} \exp(a_i)\Big) &= \log(m) + \frac{1}{m}\sum_{i=1}^{m} a_i + \log\Big(\frac{1}{m}\sum_{i=1}^{m} \exp(a - \bar{a})\Big) \\
&\geq \log(m) + \frac{1}{m}\sum_{i=1}^{m} a_i + \log\Big(1 + \frac{1}{3mR^2}\sum_{i=1}^{m}[a - \bar{a}]^2\Big) \\
&\geq \log(m) + \frac{1}{m}\sum_{i=1}^{m} a_i + \frac{1}{4mR^2}\sum_{i=1}^{m}[a - \bar{a}]^2.
\end{aligned}$$

where the first inequality is by $\exp(t) \geq 1 + t + t^2/(3R^2), \forall t \in [-R, R]$, the second inequality is due to $\log(1+t) \geq 3t/4, \forall t \in [0, 1/3]$.

□

Lemma G.7. Suppose f^* is the function that satisfies Assumption 3.1, then we have that

$$L_{\mathcal{D}^B}(f^*, \tau) \leq 2\mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t = \mathbf{z}_1)\bigg)\bigg] + 6MB\alpha/\tau + 3\sqrt[3]{6MB\beta}/\tau + 2B\exp(-\gamma/\tau)$$

Proof. Let the event $\mathcal{E}_t$ be the case that either i) $\mathbf{z}_t = \mathbf{z}_1$ and $|f^*(\mathbf{x}_t, \mathbf{y}_1) - f^*(\mathbf{x}_1, \mathbf{y}_1)| \leq \rho$ or ii) $\mathbf{z}_t \neq \mathbf{z}_1$ and $f^*(\mathbf{x}_t, \mathbf{y}_1) - f^*(\mathbf{x}_1, \mathbf{y}_1) \leq -\gamma$. We also denote the complementary set of $\mathcal{E}_t$ to be $\mathcal{E}_t^c$. By Assumption 3.1, we have that

$$\begin{aligned}
&\mathbb{P}(\mathcal{E}_t, \mathbf{z}_t = \mathbf{z}_1) \leq \beta/\rho^2 \\
&\mathbb{P}(\mathcal{E}_t, \mathbf{z}_t \neq \mathbf{z}_1) \leq \alpha.
\end{aligned}$$

the first inequality is by Chebyshev's inequality, and the second is by margin assumption. Therefore, we have that $\mathbb{P}(\mathcal{E}_t^c) \leq \alpha + \beta/\rho^2$. Next, let us decompose $L_{\mathcal{D}^B}(f^*, \tau)$ into three parts,

$$\begin{aligned}
L_{\mathcal{D}^B}(f^*, \tau) = \mathbb{E}\bigg[\log\bigg(&\sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t \neq \mathbf{z}_1)\,\mathbb{1}(\mathcal{E}_t)\exp\big(\big[f^*(\mathbf{x}_1, \mathbf{y}_t) - f^*(\mathbf{x}_1, \mathbf{y}_1)\big]/\tau\big) \\
&+ \sum_{t\in[B]} \mathbb{1}(\mathcal{E}_t^c)\exp\big(\big[f^*(\mathbf{x}_1, \mathbf{y}_t) - f^*(\mathbf{x}_1, \mathbf{y}_1)\big]/\tau\big) \\
&+ \sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t = \mathbf{z}_1)\,\mathbb{1}(\mathcal{E}_t)\exp\big(\big[f^*(\mathbf{x}_1, \mathbf{y}_t) - f^*(\mathbf{x}_1, \mathbf{y}_1)\big]/\tau\big)\bigg)\bigg] \\
&+ \mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t \neq \mathbf{z}_1)\,\mathbb{1}(\mathcal{E}_t)\exp\big(\big[f^*(\mathbf{x}_t, \mathbf{y}_1) - f^*(\mathbf{x}_1, \mathbf{y}_1)\big]/\tau\big) \\
&+ \sum_{t\in[B]} \mathbb{1}(\mathcal{E}_t^c)\exp\big(\big[f^*(\mathbf{x}_t, \mathbf{y}_1) - f^*(\mathbf{x}_1, \mathbf{y}_1)\big]/\tau\big) \\
&+ \sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t = \mathbf{z}_1)\,\mathbb{1}(\mathcal{E}_t)\exp\big(\big[f^*(\mathbf{x}_t, \mathbf{y}_1) - f^*(\mathbf{x}_1, \mathbf{y}_1)\big]/\tau\big)\bigg)\bigg] \\
\leq 2\mathbb{E}\bigg[\log\bigg(&1 + B\exp\big(-\gamma/\tau\big) + \sum_{t\geq 2} \mathbb{1}(\mathcal{E}_t^c)\exp\big(2M/\tau\big) + \sum_{t\geq 2} \mathbb{1}(\mathbf{z}_t = \mathbf{z}_1)\exp\big(\rho/\tau\big)\bigg)\bigg] \\
\leq \underbrace{2\,\mathbb{E}\bigg[\log\bigg(1 + B\exp\big(-\gamma/\tau\big)\bigg)\bigg]}_{I_1} &+ \sum_{t\geq 2} \underbrace{2\,\mathbb{E}\bigg[\log\bigg(1 + \mathbb{1}(\mathcal{E}_t^c)\exp\big(2M/\tau\big)\bigg)\bigg]}_{I_2}
\end{aligned}$$

$$+ \underbrace{2\,\mathbb{E}\Big[\log\Big(1+\sum_{t\geq 2}\mathbb{1}(\mathbf{z}_t=\mathbf{z}_1)\exp\big(\rho/\tau\big)\Big)\Big]}_{I_3} \tag{G.4}$$

where the first inequality is by Assumption 3.1, the second inequality is due to Lemma G.5. Next, we will bound I_1, I_2, I_3 separately.

$$I_1 \leq B\exp(-\gamma/\tau), \tag{G.5}$$

where the inequality is due to the fact that $\log(1+x)\leq x$.

$$I_2 = \mathbb{E}\Big[\mathbb{1}(\mathcal{E}_t^c)\log\Big(1+\exp\big(2M/\tau\big)\Big)\Big] \leq \mathbb{P}(\mathcal{E}_t^c)\frac{3M}{\tau} = (\alpha+\beta/\rho^2)\cdot\frac{3M}{\tau}. \tag{G.6}$$

where the first equality is due to $\log\Big(1+\mathbb{1}(\mathcal{E}_t^c)\exp\big(2M/\tau\big)\Big)=0)$ when $\mathbb{1}(\mathcal{E}_t^c)=0$, the first inequality is due to $\log\Big(1+\exp\big(2M/\tau\big)\Big)\leq 3M/\tau$. The last inequality is due to $\mathbb{P}(\mathcal{E}_t^c)\leq\alpha+\beta/\rho^2$.

$$\begin{aligned} I_3 &\leq \mathbb{E}\Big[\log\Big(\exp\big(\rho/\tau\big)+\sum_{t\geq 2}\mathbb{1}(\mathbf{z}_t=\mathbf{z}_1)\exp\big(\rho/\tau\big)\Big)\Big] \\ &= \rho/\tau + \mathbb{E}\Big[\log\Big(\sum_{t\in[B]}\mathbb{1}(\mathbf{z}_t=\mathbf{z}_1)\Big)\Big]. \end{aligned} \tag{G.7}$$

where the inequality is because $1\leq\exp(\rho/\tau)$.
Plugging (G.5), (G.6) and (G.7) into (G.4) gives that,

$$\begin{aligned} L_{\mathcal{D}^B}(f^*,\tau) &\leq 2B\exp(-\gamma/\tau)+6MB\alpha/\tau+6MB\beta/(\tau\rho^2)+2\rho/\tau+2\mathbb{E}\Big[\log\Big(\sum_{t\in[B]}\mathbb{1}(\mathbf{z}_t=\mathbf{z}_1)\Big)\Big] \\ &\leq 2\mathbb{E}\Big[\log\Big(\sum_{t\in[B]}\mathbb{1}(\mathbf{z}_t=\mathbf{z}_1)\Big)\Big]+6MB\alpha/\tau+3\sqrt[3]{6MB\beta}/\tau+2B\exp(-\gamma/\tau), \end{aligned}$$

where the second inequality is by choosing $\rho=\sqrt[3]{6MB\beta}$. ☐

Proof of Theorem G.1. First by Lemma G.7, we have that

$$L_{\mathcal{D}^B}(\widehat{f},\tau)\leq L_{\mathcal{D}^B}(f^*,\tau)+\epsilon\leq 2\mathbb{E}\Big[\log\Big(\sum_{t\in[B]}\mathbb{1}(\mathbf{z}_t=\mathbf{z}_1)\Big)\Big]+\epsilon' \tag{G.8}$$

where $\epsilon'=\epsilon+6MB\alpha/\tau+3\sqrt[3]{6MB\beta}/\tau+2B\exp(-\gamma/\tau)$. Notice that

$$\begin{aligned} L_{\mathcal{D}^B}(\widehat{f},\tau) = &\underbrace{\mathbb{E}\Big[\log\Big(\sum_{t\in[B]}\exp\big(\big[\widehat{f}(\mathbf{x}_1,\mathbf{y}_t)-\widehat{f}(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)\Big)\Big]}_{\mathrm{I}_1} \\ &+\underbrace{\mathbb{E}\Big[\log\Big(\sum_{t\in[B]}\exp\big(\big[\widehat{f}(\mathbf{x}_t,\mathbf{y}_1)-\widehat{f}(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)\Big)\Big]}_{I_2} \end{aligned} \tag{G.9}$$

Next, we prove the bullets in Theorem G.1 one by one.

First and Second Bullet in Theorem G.1: Denote the event $\mathcal{E}$ as the case that for all $t\geq 1$, $\mathbf{z}_t\neq\mathbf{z}_1$, which is the event that CLIP favored. We first lower bound I_1.

$$\begin{aligned} I_1 = \mathbb{E}\Big[\log\Big(&\sum_{t\in[B]}\mathbb{1}(\mathbf{z}_t\neq\mathbf{z}_1)\exp\big(\big[\widehat{f}(\mathbf{x}_t,\mathbf{y}_1)-\widehat{f}(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big) \\ &+\sum_{t\in[B]}\mathbb{1}(\mathbf{z}_t=\mathbf{z}_1)\exp\big(\big[\widehat{f}(\mathbf{x}_t,\mathbf{y}_1)-\widehat{f}(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)\Big)\Big] \end{aligned}$$

$$
\begin{aligned}
&= \mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t \neq \mathbf{z}_1)\exp\big([\widehat{f}(\mathbf{x}_t,\mathbf{y}_1) - \widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big) \\
&\qquad + \sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t = \mathbf{z}_1)\exp\big([\widehat{f}(\mathbf{x}_t,\mathbf{y}_1) - \widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big)\bigg)\bigg] \\
&\geq \mathbb{E}\bigg[\mathbb{1}(\mathcal{E})\log\bigg(\sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t \neq \mathbf{z}_1)\exp\big([\widehat{f}(\mathbf{x}_t,\mathbf{y}_1) - \widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big) + 1\bigg)\bigg] \\
&\qquad + \mathbb{E}\bigg[\mathbb{1}(\mathcal{E}^c)\log\bigg(\sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t = \mathbf{z}_1)\exp\big([\widehat{f}(\mathbf{x}_t,\mathbf{y}_1) - \widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big)\bigg)\bigg] \\
&= \mathbb{E}\bigg[\mathbb{1}(\mathcal{E})\log\bigg(\sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t \neq \mathbf{z}_1)\exp\big([\widehat{f}(\mathbf{x}_t,\mathbf{y}_1) - \widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big) + 1\bigg)\bigg] \\
&\qquad + \mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t = \mathbf{z}_1)\exp\big([\widehat{f}(\mathbf{x}_t,\mathbf{y}_1) - \widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big)\bigg)\bigg] \\
&\geq \mathbb{E}\bigg[\mathbb{1}(\mathcal{E})\log\bigg(\sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t \neq \mathbf{z}_1)\exp\big([\widehat{f}(\mathbf{x}_t,\mathbf{y}_1) - \widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big) + 1\bigg)\bigg] \\
&\qquad + \mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t = \mathbf{z}_1)\exp\big(\mathbb{E}[\widehat{f}(\mathbf{x}_t,\mathbf{y}_1) - \widehat{f}(\mathbf{x}_1,\mathbf{y}_1)\big|\mathbf{z}_t,\mathbf{z}_1]/\tau\big)\bigg)\bigg] \\
&= \mathbb{E}\bigg[\mathbb{1}(\mathcal{E})\log\bigg(\sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t \neq \mathbf{z}_1)\exp\big([\widehat{f}(\mathbf{x}_t,\mathbf{y}_1) - \widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big) + 1\bigg)\bigg] \\
&\qquad + \mathbb{E}\bigg[\log\bigg(\Big|\Big\{t\in[B]\Big|\mathbf{z}_t = \mathbf{z}_1\Big\}\Big|\bigg)\bigg]. \qquad \text{(G.10)}
\end{aligned}
$$

where the first inequality is because when $\mathcal{E}$ holds $\sum_{t\in[B]}\mathbb{1}(\mathbf{z}_t = \mathbf{z}_1)\exp\big([\widehat{f}(\mathbf{x}_t,\mathbf{y}_1) - \widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big) = 1$ when $\mathcal{E}^c$ holds $\sum_{t\in[B]}\mathbb{1}(\mathbf{z}_t \neq \mathbf{z}_1)\exp\big([\widehat{f}(\mathbf{x}_t,\mathbf{y}_1)-\widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big) \geq 0$, the last second equality is because when $\mathcal{E}$ holds $\sum_{t\in[B]}\mathbb{1}(\mathbf{z}_t = \mathbf{z}_1)\exp\big([\widehat{f}(\mathbf{x}_t,\mathbf{y}_1) - \widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big) = 1$, the second inequality is because LogSumExp function is convex, and the last equality is due to $\mathbb{E}[[\widehat{f}(\mathbf{x}_t,\mathbf{y}_1) - \widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]|\mathbf{z}_t,\mathbf{z}_1] = 0$ when $\mathbf{z}_t = \mathbf{z}_1$. Similarly, we can prove

$$
\begin{aligned}
I_2 &\geq \mathbb{E}\bigg[\mathbb{1}(\mathcal{E})\log\bigg(\sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t \neq \mathbf{z}_1)\exp\big([\widehat{f}(\mathbf{x}_1,\mathbf{y}_t) - \widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big) + 1\bigg)\bigg] \\
&\qquad + \mathbb{E}\bigg[\log\bigg(\Big|\Big\{t\in[B]\Big|\mathbf{z}_t = \mathbf{z}_1\Big\}\Big|\bigg)\bigg]. \qquad \text{(G.11)}
\end{aligned}
$$

Notice that when event $\mathcal{E}$ holds, $\mathbf{z}_t \neq \mathbf{z}_1$ holds for all $t \geq 2$. Therefore, plugging the (G.10) and (G.11) into (G.9) gives,

$$
\mathbb{E}\bigg[\mathbb{1}(\mathcal{E})\log\bigg(\sum_{t\geq 2}\exp\big([\widehat{f}(\mathbf{x}_t,\mathbf{y}_1) - \widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big) + 1\bigg)\bigg] \leq \epsilon' \qquad \text{(G.12)}
$$

$$
\mathbb{E}\bigg[\mathbb{1}(\mathcal{E})\log\bigg(\sum_{t\geq 2}\exp\big([\widehat{f}(\mathbf{x}_1,\mathbf{y}_t) - \widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big) + 1\bigg)\bigg] \leq \epsilon'. \qquad \text{(G.13)}
$$

(G.14)

Let us compute the probability of $\mathcal{E}$ given $\mathbf{z}_1$. Let $\mathbf{z}_1 = \mathbf{v}_1$ without loss of generality, we have that

$$
\mathbb{P}(\mathcal{E}|\mathbf{z} = \mathbf{v}_1) = (1 - p_1)^{B-1}.
$$

Therefore $\mathbb{P}(\mathcal{E}|\mathbf{z} = \mathbf{v}_1)$ is always positive and is greater than $1/2$ as long as $B \leq 1/p_1$.
Next, consider the following situation. Given $\mathbf{z}_1 = \mathbf{v}_1$, we generate sequence $\mathbf{z}_1', \ldots, \mathbf{z}_L'$ with length $L = \lceil \log(2K)/(B-1)\min p_k \rceil (B-1)$, such that each $\mathbf{z}_1', \ldots, \mathbf{z}_L'$ are generated from $\mathcal{D}_{\mathbf{z}|\mathbf{z}\neq\mathbf{v}_1}$.

The probability that the sequence includes $\mathbf{v}_k$ is

$$1-(1-p_k/(1-p_k))^L \geq 1-(1-p_k)^L \geq 1-\exp(-Lp_k) \geq 1-\exp(-L\min p_k).$$

Therefore the probability that the sequence can cover all the other $K-1$ classes is at least

$$1-K\exp(-L\min p_k) \geq 1/2.$$

Then we look deeper into

$$\mathbb{E}\bigg[\log\bigg(\sum_{t\geq 2}\exp\big([\widehat{f}(\mathbf{x}_t,\mathbf{y}_1)-\widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big)+1\bigg)\bigg|\mathbf{z}_1=\mathbf{v}_1,\mathbf{z}_2\neq\mathbf{v}_1,\ldots,\mathbf{z}_K\neq\mathbf{v}_1\bigg].$$

We can introduce $L/(B-1)$ copies $\mathbf{x}_t^{(l)}$ with $l\in[L/(B-1)]$ for $t\geq 2$, then we have that

$$\begin{aligned}
&\Big(L/(B-1)\Big)\cdot\mathbb{E}\bigg[\log\bigg(\sum_{t\geq 2}\exp\big([\widehat{f}(\mathbf{x}_t,\mathbf{y}_1)-\widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big)+1\bigg)\bigg|\mathbf{z}_1=\mathbf{v}_1,\mathbf{z}_2\neq\mathbf{v}_1,\ldots,\mathbf{z}_K\neq\mathbf{v}_1\bigg]\\
&=\mathbb{E}\bigg[\sum_l\log\bigg(\sum_{t\geq 2}\exp\big([\widehat{f}(\mathbf{x}_t^{(l)},\mathbf{y}_1)-\widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big)+1\bigg)\bigg|\mathbf{z}_1=\mathbf{v}_1,\mathbf{z}_2^{(l)},\ldots,\mathbf{z}_K^{(l)}\neq\mathbf{v}_1\bigg]\\
&\geq\mathbb{E}\bigg[\log\bigg(\sum_l\sum_{t\geq 2}\exp\big([\widehat{f}(\mathbf{x}_t^{(l)},\mathbf{y}_1)-\widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big)+1\bigg)\bigg|\mathbf{z}_1=\mathbf{v}_1,\mathbf{z}_2^{(l)},\ldots,\mathbf{z}_K^{(l)}\neq\mathbf{v}_1\bigg]\\
&\geq\mathbb{E}\bigg[\log\bigg(\sum_{k\in[K]}\exp\big([\widehat{f}(\mathbf{x}_k,\mathbf{y})-\widehat{f}(\mathbf{x}^*,\mathbf{y})]/\tau\big)\bigg)\bigg|\mathbf{z}=\mathbf{v}_1\bigg].
\end{aligned}\tag{G.15}$$

where the first inequality is by Lemma G.5, the second inequality is by the fact that the Exp function is greater than 0, and the $\mathbf{x}_k,\mathbf{x}^*$ in the last line are the ones that defined in Theorem G.1. Plugging (G.15) into (G.12) and applying total expectation completes the proof for the second bullet. The proof for the first bullet is the same.

Third Bullet in Theorem G.1: By the third equality in (G.10), we have that

$$\begin{aligned}
I_1&\geq\mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]}\mathbb{1}(\mathbf{z}_t=\mathbf{z}_1)\exp\big([\widehat{f}(\mathbf{x}_t,\mathbf{y}_1)-\widehat{f}(\mathbf{x}_1,\mathbf{y}_1)]/\tau\big)\bigg)\bigg]\\
&=\mathbb{E}\bigg[\mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]}\mathbb{1}(\mathbf{z}_t=\mathbf{z}_1)\exp\big(\widehat{f}(\mathbf{x}_t,\mathbf{y}_1)/\tau\big)\bigg)\bigg|\mathbf{z}_1,\ldots,\mathbf{z}_B\bigg]\bigg]-\mathbb{E}[\widehat{f}(\mathbf{x}_1,\mathbf{y}_1)/\tau]\\
&\geq\mathbb{E}\bigg[\log\bigg(\Big|\Big\{t\in[B]\Big|\mathbf{z}_t=\mathbf{z}_1\Big\}\Big|\bigg)\bigg]+\mathbb{E}\Bigg[\frac{\Big|\Big\{t\in[B]\Big|\mathbf{z}_t=\mathbf{z}_1\Big\}\Big|-1}{4M^2\Big|\Big\{t\in[B]\Big|\mathbf{z}_t=\mathbf{z}_1\Big\}\Big|}\mathrm{Var}_{\mathbf{x}_1|\mathbf{z}_1}(\widehat{f}(\mathbf{x}_1,\mathbf{y}_1))\Bigg].
\end{aligned}\tag{G.16}$$

where the inequality is by Lemma G.6. Next we will We analyze the distribution of $\Big\{t\in[B]\Big|\mathbf{z}_t=\mathbf{z}_1\Big\}$. Without loss of generality, fix $\mathbf{z}_1=\mathbf{v}_1$. We know that the probability that $\Big\{t\in[B]\Big|\mathbf{z}_t=\mathbf{z}_1\Big\}\geq 2$ is

$$1-\mathbb{P}(\mathbf{z}_2\neq\mathbf{z}_1)\cdot\ldots\cdot\mathbb{P}(\mathbf{z}_B\neq\mathbf{z}_1)\geq 1-(1-\min p_k)^{B-1}\geq\min\{0.25*\min p_k\cdot(B-1),0.25\},$$

the last inequality holds since the strictly increasing function $F(s)=1-(1-\min p_k)^s$ is 0 at $s=0$ and have derivative lower bounded by 0.25 when $s\leq 1/\min p_k$. Therefore we can further lower bound (G.16) as follows,

$$I_1\geq\mathbb{E}\bigg[\log\bigg(\Big|\Big\{t\in[B]\Big|\mathbf{z}_t=\mathbf{z}_1\Big\}\Big|\bigg)\bigg]+\mathbb{E}\bigg[\frac{\min\{0.25*\min p_k\cdot(B-1),0.25\}}{8M^2}\mathrm{Var}_{\mathbf{x}_1|\mathbf{z}_1}(\widehat{f}(\mathbf{x}_1,\mathbf{y}_1))\bigg]$$

Similarly, we can prove that

$$I_2\geq\mathbb{E}\bigg[\log\bigg(\Big|\Big\{t\in[B]\Big|\mathbf{z}_t=\mathbf{z}_1\Big\}\Big|\bigg)\bigg]+\mathbb{E}\bigg[\frac{\min\{0.25*\min p_k\cdot(B-1),0.25\}}{8M^2}\mathrm{Var}_{\mathbf{y}_1|\mathbf{z}_1}(\widehat{f}(\mathbf{x}_1,\mathbf{y}_1))\bigg].$$

Plugging the bound of I_1,I_2 into (G.9) completes the proof for the third bullet of Theorem G.1. $\square$

H Proof of the Results in Section 3

Corollary H.1. Suppose the result of Theorem G.1 holds for the learned similarity function $\widehat{f}$. Then we calculate the similarity score $\widehat{f}(\mathbf{x}, \mathbf{y}_k)$ for all $k \in [K]$ and pick the indices of the top-r scores within the set $\{\widehat{f}(\mathbf{x}, \mathbf{y}_k)\}$ as the predictions of the image $\mathbf{x}$. Then the top-r error is bounded by $\epsilon'/\log(1+r)$.

Proof of Corollary H.1. For $(\mathbf{x}, \mathbf{z}) \sim \mathcal{D}_{\mathbf{x}\times\mathbf{z}}$, $\{\mathbf{y}_k \sim \mathcal{D}_{\mathbf{y}|\mathbf{v}_k}, k \in [K]\}$, let $\mathbf{y}^* = \sum_{k\in[K]} \mathbb{1}(\mathbf{z} = \mathbf{v}_k)\mathbf{y}_k$. Denote $\mathcal{E}$ to be the event that the top-r choice gives the wrong prediction. Then we have that,

$$
\begin{aligned}
\epsilon' &\geq \mathbb{E}\bigg[\log\bigg(\sum_{k\in[K]} \exp\big(\big[\widehat{f}(\mathbf{x}, \mathbf{y}_k) - \widehat{f}(\mathbf{x}, \mathbf{y}^*)\big]/\tau\big)\bigg)\bigg] \\
&\geq \mathbb{E}\bigg[\mathbb{1}(\mathcal{E})\log\bigg(\sum_{k\in[K]} \exp\big(\big[\widehat{f}(\mathbf{x}, \mathbf{y}_k) - \widehat{f}(\mathbf{x}, \mathbf{y}^*)\big]/\tau\big)\bigg)\bigg] \\
&\geq \mathbb{E}\bigg[\mathbb{1}(\mathcal{E})\log(1+r)\bigg] \\
&= \mathbb{P}(\mathcal{E})\log(1+r),
\end{aligned}
$$

where the first inequality is by the first bullet of Theorem G.1, the second inequality is due to the fact that $\log\Big(\sum_{k\in[K]} \exp\big(\big[\widehat{f}(\mathbf{x}, \mathbf{y}_k) - \widehat{f}(\mathbf{x}, \mathbf{y}^*)\big]/\tau\big)\Big) > 0$, the last inequality is due to $\log\Big(\sum_{k\in[K]} \exp\big(\big[\widehat{f}(\mathbf{x}, \mathbf{y}_k) - \widehat{f}(\mathbf{x}, \mathbf{y}^*)\big]/\tau\big)\Big) \geq \log(1+r)$ since there are at least $r+1$ number of $\widehat{f}(\mathbf{x}, \mathbf{y}_k)$ are greater than $\widehat{f}(\mathbf{x}, \mathbf{y}^*)$ if the prediction is wrong. Therefore, we have that $\mathbb{P}(\mathcal{E}) \leq \epsilon'/\log(1+r)$ which completes the proof. $\square$

Remark H.2. The result in Corollary H.1 can be generalized to out-of-distribution zero-shot transfer. For example, we can deal with the case where the distribution of the prompts $\mathcal{D}_{\mathbf{y}|\mathbf{v}_k}$ and the image distribution $\mathcal{D}_{\mathbf{x}}$ are shifted. As long as the χ^2 distance between the shifted distributions is bounded, we can provide a top-r error guarantee.

Discussion for out-of-distribution zero shot learning. The result in Corollary H.1 can be generalized to out-of-distribution zero-shot transfer learning. For example, we can deal with the case where the distribution of the prompts $\mathcal{D}_{\mathbf{y}|\mathbf{v}_k}$ and the image distribution $\mathcal{D}_{\mathbf{x}}$ are shifted. In particular, let us consider the case that the distribution of the prompts is shifted to $\mathcal{D}'_{\mathbf{y}|\mathbf{v}_k}$ and the image distribution $\mathcal{D}_{\mathbf{x}}$ is shifted to $\mathcal{D}'_{\mathbf{x}}$. Then the original joint cumulative distribution function function $P(\mathbf{x}, \mathbf{z}, \mathbf{y}_1, \ldots, \mathbf{y}_K)$ is shifted to $Q(\mathbf{x}, \mathbf{z}, \mathbf{y}_1, \ldots, \mathbf{y}_K)$. Suppose Q is absolutely continuous with respect to P, and the Pearson χ^2 distance is bounded

$$
\int \bigg(\frac{dQ}{dP} - 1\bigg)^2 dP \leq C.
$$

Then we have that

$$
\begin{aligned}
&\int \sqrt{\log\bigg(\sum_{k\in[K]} \exp\big(\big[\widehat{f}(\mathbf{x}, \mathbf{y}_k) - \widehat{f}(\mathbf{x}, \mathbf{y}^*)\big]/\tau\big)\bigg)} dQ \\
&= \int \sqrt{\log\bigg(\sum_{k\in[K]} \exp\big(\big[\widehat{f}(\mathbf{x}, \mathbf{y}_k) - \widehat{f}(\mathbf{x}, \mathbf{y}^*)\big]/\tau\big)\bigg)} \bigg(\frac{dQ}{dP}\bigg) dP \\
&\leq \sqrt{\int \log\bigg(\sum_{k\in[K]} \exp\big(\big[\widehat{f}(\mathbf{x}, \mathbf{y}_k) - \widehat{f}(\mathbf{x}, \mathbf{y}^*)\big]/\tau\big)\bigg) dP} \cdot \sqrt{\int \bigg(\frac{dQ}{dP}\bigg)^2 dP} \\
&= \sqrt{(C+1)\epsilon'},
\end{aligned}
$$

where the first inequality is by Cauchy Schwartz inequality and the last equality is due to $\int \left(\frac{dQ}{dP}\right)^2 dP = \int \left(\frac{dQ}{dP} - 1\right)^2 dP + 1 = C + 1$. Then we can follow a similar analysis in the proof of Corollary H.1 and have that top-r test error is smaller than $\sqrt{(C+1)\epsilon'/\log(1+r)}$. Therefore, if the χ^2 distance between the shifted distributions is bounded, we can still provide a top-r error guarantee. It is worth noting the bound for out-of-distribution zero-shot learning is looser. If we want to do a more general zero shot analysis, we may need to add more data structure in Assumption 3.1.

Lemma H.3 (Completeness)**.** There exist a score function $f^*(\mathbf{x}, \mathbf{y}) = \langle \mathbf{W}^*\mathbf{x}, \mathbf{y}\rangle$ with $\mathbf{W}^* \in \mathbb{R}^{d_2 \times d_1}$ satisfying

- $|f^*| \leq 1$,
- For $(\mathbf{x}, \mathbf{y}, \mathbf{z}) \sim \mathcal{D}_{\mathbf{x}\times\mathbf{y}\times\mathbf{z}}$, variance $\mathbb{E}_{(\mathbf{y},\mathbf{z})}\big[\mathrm{Var}_{\mathbf{x}|\mathbf{z}}(f^*(\mathbf{x}, \mathbf{y}))\big] = \mathbb{E}_{(\mathbf{x},\mathbf{z})}\big[\mathrm{Var}_{\mathbf{y}|\mathbf{z}}(f^*(\mathbf{x}, \mathbf{y}))\big] = 0$,
- Let $\mathbf{x} \sim \mathcal{D}_{\mathbf{x}|\mathbf{z}}, \mathbf{y} \sim \mathcal{D}_{\mathbf{y}|\mathbf{z}}, \mathbf{x}' \sim \mathcal{D}_{\mathbf{x}'|\mathbf{z}'}, \mathbf{y}' \sim \mathcal{D}_{\mathbf{y}'|\mathbf{z}'}$ where $\mathbf{z} \neq \mathbf{z}'$. With probability 1, we have that $f^*(\mathbf{x}', \mathbf{y}) \leq f^*(\mathbf{x}, \mathbf{y}) - \gamma$ and $f^*(\mathbf{x}, \mathbf{y}') \leq f^*(\mathbf{x}, \mathbf{y}) - \gamma$.

Proof of Lemma H.3. We can construct $\mathbf{W}^* = \mathbf{H}(\mathbf{H}^\top\mathbf{H})^{-1}\mathbf{P}(\mathbf{G}^\top\mathbf{G})^{-1}\mathbf{G}^\top$, where $\mathbf{P} \in \mathbb{R}^{(K_1+K_2)\times(K_1+K_3)}$ is the projection matrix $\begin{bmatrix}\mathbf{I} & \mathbf{0}\\ \mathbf{0} & \mathbf{0}\end{bmatrix}$ with rank K_1.

It is easy to verify that $\mathbf{H}^\top\mathbf{W}^*\mathbf{G} = \mathbf{P}$. Therefore we have that

$$\langle \mathbf{W}^*\mathbf{x}, \mathbf{y}'\rangle = \langle \mathbf{z}, \mathbf{z}'\rangle.$$

Then applying $\|\mathbf{v}_k\|_2 = 1$, $\langle \mathbf{v}_k, \mathbf{v}'_k\rangle \leq 1 - \gamma, \forall k \neq k'$ completes the proof . □

Lemma H.4. $\|\nabla L_S(f_{\mathbf{W}}, \tau)\|_F \leq L$ where $L = 2\tau^{-1}\|\mathbf{G}\|_2\|\mathbf{H}\|_2(R^2+1)$.

Proof. First, we have that

$$\|\nabla_{\mathbf{W}}\langle \mathbf{W}\mathbf{x}, \mathbf{y}\rangle\|_F = \|\mathbf{x}\mathbf{y}^\top\|_F \leq \|\mathbf{x}\|_2\|\mathbf{y}\|_2 \leq \|\mathbf{G}\|_2\|\mathbf{H}\|_2(R^2+1).$$

Therefore we have that $\|\nabla L_S(f_{\mathbf{W}}, \tau)\|_F \leq 2\tau^{-1}\|\mathbf{G}\|_2\|\mathbf{H}\|_2(R^2+1)$ since LogSumExp function is an 1-Lipschitz function. □

Proof of Theorem 3.4. By the gradient update rule, we have that

$$\begin{aligned}
&\|\mathbf{W}^{(t)} - \mathbf{W}^*\|_F^2 - \|\mathbf{W}^{(t+1)} - \mathbf{W}^*\|_F^2\\
&= 2\eta\langle \nabla\widehat{L}_S(\mathbf{W}^{(t)}, \tau), \mathbf{W}^{(t)} - \mathbf{W}^*\rangle - \eta^2\|\nabla\widehat{L}_S(\mathbf{W}^{(t)}, \tau)\|_F^2\\
&\geq 2\eta\widehat{L}_S(\mathbf{W}^{(t)}, \tau) - 2\eta\widehat{L}_S(\mathbf{W}^*, \tau) - \eta^2L^2.
\end{aligned} \tag{H.1}$$

Take the telescope sum of (H.1) from 0 to $T-1$ we have that

$$\begin{aligned}
\frac{\sum_{t=0}^{T-1}\widehat{L}_S(\mathbf{W}^{(t)}, \tau)}{T} &\leq \widehat{L}_S(\mathbf{W}^*, \tau) + \eta L^2 + \frac{\|\mathbf{W}^{(0)} - \mathbf{W}^*\|_F^2 - \|\mathbf{W}^{(T)} - \mathbf{W}^*\|_F^2}{2\eta T}\\
&\leq \widehat{L}_S(\mathbf{W}^*, \tau) + \epsilon/4 + \epsilon/4\\
&= \widehat{L}_S(\mathbf{W}^*, \tau) + \epsilon/2,
\end{aligned}$$

where the second inequality is by $\eta \leq \epsilon/(4L^2)$ and $T = 4\|\mathbf{W}^{(0)} - \mathbf{W}^*\|_F^2/(\eta\epsilon)$. Therefore, there exist $t' \leq T-1$ such that $\widehat{L}_S(\mathbf{W}^{(t')}, \tau) \leq \widehat{L}_S(\mathbf{W}^*, \tau) + \epsilon/2$. Let $\widehat{T}$ to be the first time that $\widehat{L}_S(\mathbf{W}^{(\widehat{T})}, \tau) \leq \widehat{L}_S(\mathbf{W}^*, \tau) + \epsilon/2$. Again take telescope sum of (H.1) from 0 to $\widehat{T}-1$, we have that

$$\begin{aligned}
\|\mathbf{W}^{(\widehat{T})} - \mathbf{W}^*\|_F^2 &\leq 2\eta\widehat{T}\widehat{L}_S(\mathbf{W}^*, \tau) - 2\eta\widehat{T}\sum_{t=0}^{\widehat{T}-1}\widehat{L}_S(\mathbf{W}^{(t)}, \tau) + 2\eta^2L^2\widehat{T} + \|\mathbf{W}^{(0)} - \mathbf{W}^*\|_F^2\\
&\leq -\eta\widehat{T}\epsilon + 0.5\eta\widehat{T}\epsilon + \|\mathbf{W}^{(0)} - \mathbf{W}^*\|_F^2\\
&\leq \|\mathbf{W}^{(0)} - \mathbf{W}^*\|_F^2,
\end{aligned}$$

where the second inequality is due to the definition of $\widehat{T}$, the last inequality is due to $-0.5\eta\widehat{T}\epsilon \leq 0$. Therefore, within $T = 4\|\mathbf{W}^{(0)} - \mathbf{W}^*\|_F^2/(\eta\epsilon)$ we can find $\widehat{\mathbf{W}} = \mathbf{W}^{(\widehat{T})}$ such that $\widehat{L}_S(\widehat{\mathbf{W}}, \tau) \leq \widehat{L}_S(\mathbf{W}^*, \tau) + \epsilon/2$ and

$$\|\mathbf{W}^{(\widehat{T})}\|_F^2 \leq 2\|\mathbf{W}^*\|_F + \|\mathbf{W}^{(0)}\|_F^2$$

where the inequality is by triangle inequality. Therefore, for any $\mathbf{x}, \mathbf{y}$

$$\begin{aligned}\widehat{f}(\mathbf{x}, \mathbf{y}) &= \langle \mathbf{W}^* \mathbf{x}, \mathbf{y} \rangle + \langle \widehat{\mathbf{W}} - \mathbf{W}^* \mathbf{x}, \mathbf{y} \rangle \\ &\leq 1 + \|\widehat{\mathbf{W}} - \mathbf{W}^*\|_F \|\mathbf{x}\mathbf{y}^\top\|_F \\ &\leq 1 + \|\widehat{\mathbf{W}} - \mathbf{W}^*\|_F \|\mathbf{G}\|_2 \|\mathbf{H}\|_2 (R^2 + 1) \\ &\leq 1 + \|\mathbf{W}^* - \mathbf{W}^{(0)}\|_F \|\mathbf{G}\|_2 \|\mathbf{H}\|_2 (R^2 + 1).\end{aligned}$$

Therefore the function $\widehat{f}$ is bonded by $M = 1 + \|\mathbf{W}^* - \mathbf{W}^{(0)}\|_F \|\mathbf{G}\|_2 \|\mathbf{H}\|_2 (R^2 + 1)$. Moreover, the function $\widehat{f}$ must belong to the class $\mathcal{F} = \{\langle \mathbf{W}\mathbf{x}, \mathbf{y} \rangle | \|\mathbf{W}\|_F \leq 2\|\mathbf{W}^*\|_F + \|\mathbf{W}^{(0)}\|_F^2\}$. Since the linear function class $\mathcal{F}$ has finite covering the set $\mathcal{N}(\mathcal{F}, \epsilon)$ (Bartlett and Mendelson, 2002; Zhang, 2002), by Theorem B.2 we know that when $n \geq (8\tau^{-1}\epsilon^{-2} M \log B) \log(2\mathcal{N}(\mathcal{F}, \epsilon/32M)/\delta)$, with probability at least $1 - \delta$ we have that

$$\begin{aligned}|\widehat{L}_S(\widehat{f}, \tau) - L_{\mathcal{D}^B}(\widehat{f}, \tau)| &\leq \epsilon/4 \\ |\widehat{L}_S(f^*, \tau) - L_{\mathcal{D}^B}(f^*, \tau)| &\leq \epsilon/4.\end{aligned}$$

Thus, we can conclude that

$$\begin{aligned}\widehat{L}_{\mathcal{D}^B}(\widehat{f}, \tau) - \widehat{L}_{\mathcal{D}^B}(f^*, \tau) &\leq \widehat{L}_S(\widehat{f}, \tau) - \widehat{L}_S(f^*, \tau) + |\widehat{L}_S(\widehat{f}, \tau) - L_{\mathcal{D}^B}(\widehat{f}, \tau)| \\ &\quad + |\widehat{L}_S(f^*, \tau) - L_{\mathcal{D}^B}(f^*, \tau)| \\ &\leq \epsilon/2 + \epsilon/4 + \epsilon/4 \\ &= \epsilon.\end{aligned}$$

where the first inequality is by the triangle inequality, the second inequality is by the bounded gap between empirical and population loss. □

Proof of Theorem 3.5.

$$\begin{aligned}\mathbb{E}\Big[\|\mathbf{g}(\mathbf{x}) - \mathbf{y}\|_2^2 \Big| \mathbf{z}\Big] &= \mathbb{E}\Big[\|\mathbf{g}(\mathbf{x}) - \mathbb{E}[\mathbf{y}|\mathbf{z}] + \mathbb{E}[\mathbf{y}|\mathbf{z}] - \mathbf{y}\|_2^2 \Big| \mathbf{z}\Big] \\ &= \mathbb{E}\Big[\|\mathbf{g}(\mathbf{x}) - \mathbb{E}[\mathbf{y}|\mathbf{z}]\|_2^2 \Big| \mathbf{z}\Big] + \mathbb{E}\Big[\|\mathbb{E}[\mathbf{y}|\mathbf{z}] - \mathbf{y}\|_2^2 \Big| \mathbf{z}\Big]\end{aligned}$$

where the second equality is due to $\mathbf{x} \perp \mathbf{y} | \mathbf{z}$ and $\mathbb{E}\Big[\mathbb{E}[y|z] - y \Big| \mathbf{z}\Big] = \mathbf{0}$. Then taking a total expectation over both sides over $\mathbf{z}$ gives that

$$\mathbb{E}\big[\|\mathbf{g}(\mathbf{x}) - \mathbf{y}\|_2^2\big] = \mathbb{E}\big[\|\mathbf{g}(\mathbf{x}) - \mathbb{E}[\mathbf{y}|\mathbf{z}]\|_2^2\big] + \mathbb{E}\big[\|\mathbf{y} - \mathbb{E}[\mathbf{y}|\mathbf{z}]\|_2^2\big] \geq \mathbb{E}\big[\|\mathbf{y} - \mathbb{E}[\mathbf{y}|\mathbf{z}]\|_2^2\big].$$

Obviously, $\mathbb{E}\big[\|\mathbf{g}(\mathbf{x}) - \mathbf{y}\|_2^2\big]$ achieves global minima when

$$\mathbf{g}(\mathbf{x}) = \mathbb{E}[\mathbf{y}|\mathbf{z}] = \mathbf{H} \begin{bmatrix} \mathbf{z} \\ \mathbb{E}[\boldsymbol{\zeta}|\mathbf{z}] \end{bmatrix}.$$

This function $\mathbf{g}$ is also achievable. We can construct function $\mathbf{g}_2(\mathbf{z}) = \mathbf{H} \begin{bmatrix} \mathbf{z} \\ \mathbb{E}[\boldsymbol{\zeta}|\mathbf{z}] \end{bmatrix}$, and projection function $\mathbf{g}_1(\mathbf{x}) = \mathbf{z}$ that is linear. Then we can define $\mathbf{g} = \mathbf{g}_2 \circ \mathbf{g}_1$. □

Proof of Corollary 3.6. Since $\boldsymbol{\zeta}$ is independent with $\mathbf{z}$, we have that

$$\mathbf{g}(\mathbf{x}) = \mathbf{H} \begin{bmatrix} \mathbf{z} \\ \mathbb{E}[\boldsymbol{\zeta}|\mathbf{z}] \end{bmatrix} = 1/3 \cdot \begin{bmatrix} \mathbf{z} \\ \mathbf{e}_1 \\ \mathbf{0} \end{bmatrix} + 2/3 \cdot \begin{bmatrix} \mathbf{z} \\ \mathbf{e}_2 \\ \mathbf{0} \end{bmatrix}.$$

Besides, we have that

$$\mathbf{y}' = \mathbf{H} \begin{bmatrix} \mathbf{z}' \\ \boldsymbol{\zeta}' \end{bmatrix} = \begin{bmatrix} \mathbf{z}' \\ \boldsymbol{\zeta}' \\ \mathbf{0}. \end{bmatrix}$$

Inner product similarity. We have that $f(\mathbf{x}, \mathbf{y}') = \langle \mathbf{z}, \mathbf{z}' \rangle + 1/3 + 1/3 \cdot \mathbb{1}(\boldsymbol{\zeta}' = \mathbf{e}_2)$. Since margin $\gamma < 1/3$. There exist j, k such that $\langle \mathbf{v}_j, \mathbf{v}_k \rangle > 2/3$. Then for $\mathbf{z} = \mathbf{v}_j$, we will sample K prompt $\mathbf{y}_1, \ldots, \mathbf{y}_K$. When $\mathbf{y}_j = \begin{bmatrix} \mathbf{v}_j \\ \mathbf{e}_1 \\ \mathbf{0}. \end{bmatrix}$ and $\mathbf{y}_k = \begin{bmatrix} \mathbf{v}_k \\ \mathbf{e}_2 \\ \mathbf{0}. \end{bmatrix}$, we have that

$$f(\mathbf{x}, \mathbf{y}_j) = 4/3 < \langle \mathbf{v}_j, \mathbf{v}_k \rangle + 2/3 = f(\mathbf{x}, \mathbf{y}_k),$$

which leads to the wrong top-1 prediction. The key insight behind this consequence is that $f(\mathbf{x}, \mathbf{y}') = \langle \mathbf{z}, \mathbf{z}' \rangle + 1/3 + 1/3 \cdot \mathbb{1}(\boldsymbol{\zeta}' = \mathbf{e}_2)$ is greatly influenced by the unique feature $\boldsymbol{\zeta}$. A similar case also exists for $\mathbf{z} = \mathbf{v}_k$ with $\mathbf{y}_j = \begin{bmatrix} \mathbf{v}_j \\ \mathbf{e}_2 \\ \mathbf{0}. \end{bmatrix}$ and $\mathbf{y}_k = \begin{bmatrix} \mathbf{v}_k \\ \mathbf{e}_1 \\ \mathbf{0}. \end{bmatrix}$. The probability that the above event occurs is at least $2/K \cdot 1/3 \cdot 2/3 = 4/(9K) \geq 1/(3K)$. Therefore, the test error is at least $1/(3K)$.

Cosine similarity. Notice that $\|\mathbf{g}(\mathbf{x})\|_2 = \sqrt{1 + 1/9 + 4/9} = \sqrt{14}/3$, and $\|\mathbf{y}\|_2 = 1$, therefore the cosine similarity is proportional to inner product similarity with factor $\sqrt{14}/3$. Thus, the test error is still at least $1/(3K)$.

L_2 similarity. We have that $f(\mathbf{x}, \mathbf{y}') = -\|\mathbf{z} - \mathbf{z}'\|_2^2 - 8/9 + 2/3 \cdot \mathbb{1}(\boldsymbol{\zeta}' = \mathbf{e}_2)$. Since margin $\gamma < 1/3$. There exist j, k such that $\|\mathbf{v}_j - \mathbf{v}_k\|_2^2 < 2/3$. Then for $\mathbf{z} = \mathbf{v}_j$, we will sample K prompt $\mathbf{y}_1, \ldots, \mathbf{y}_K$. When $\mathbf{y}_j = \begin{bmatrix} \mathbf{v}_j \\ \mathbf{e}_1 \\ \mathbf{0}. \end{bmatrix}$ and $\mathbf{y}_k = \begin{bmatrix} \mathbf{v}_k \\ \mathbf{e}_2 \\ \mathbf{0}. \end{bmatrix}$, we have that

$$f(\mathbf{x}, \mathbf{y}_j) = -8/9 < -\|\mathbf{v}_j, \mathbf{v}_k\|_2^2 + 2/3 = f(\mathbf{x}, \mathbf{y}_k),$$

which leads to the wrong top-1 prediction. The key insight behind this consequence is that $f(\mathbf{x}, \mathbf{y}') = -\|\mathbf{z} - \mathbf{z}'\|_2^2 - 8/9 + 2/3 \cdot \mathbb{1}(\boldsymbol{\zeta}' = \mathbf{e}_2)$ is greatly influenced by the unique feature $\boldsymbol{\zeta}$. A similar case also exists for $\mathbf{z} = \mathbf{v}_k$ with $\mathbf{y}_j = \begin{bmatrix} \mathbf{v}_j \\ \mathbf{e}_2 \\ \mathbf{0}. \end{bmatrix}$ and $\mathbf{y}_k = \begin{bmatrix} \mathbf{v}_k \\ \mathbf{e}_1 \\ \mathbf{0}. \end{bmatrix}$. The probability that the above event occurs is at least $2/K \cdot 1/3 \cdot 2/3 = 4/(9K) \geq 1/(3K)$. Therefore, the test error is at least $1/(3K)$.

$\square$

I Proof of Results in Section 4

Proof of Corollary 4.1. For $(\mathbf{x}, \mathbf{z}) \sim \mathcal{D}_{\mathbf{x} \times \mathbf{z}}$, $\{\mathbf{y}_k \sim \mathcal{D}_{\mathbf{y}|\mathbf{v}_k}, k \in [K]\}$, let $\mathbf{y}^* = \sum_{k \in [K]} \mathbb{1}(\mathbf{z} = \mathbf{v}_k)\mathbf{y}_k$. Denote $\mathcal{E}$ to be the event that the top-1 choice gives the wrong prediction or the margin is smaller than τ. Then we have that,

$$\begin{aligned}
\epsilon' &\geq \mathbb{E}\bigg[\log\bigg(\sum_{k \in [K]} \exp\big(\big[\widehat{f}(\mathbf{x}, \mathbf{y}_k) - \widehat{f}(\mathbf{x}, \mathbf{y}^*)\big]/\tau\big)\bigg)\bigg] \\
&\geq \mathbb{E}\bigg[\mathbb{1}(\mathcal{E}) \log\bigg(\sum_{k \in [K]} \exp\big(\big[\widehat{f}(\mathbf{x}, \mathbf{y}_k) - \widehat{f}(\mathbf{x}, \mathbf{y}^*)\big]/\tau\big)\bigg)\bigg] \\
&\geq \mathbb{E}\bigg[\mathbb{1}(\mathcal{E}) \log(1 + \exp(-1))\bigg] \\
&= \mathbb{P}(\mathcal{E}) \log(1 + e^{-1}),
\end{aligned}$$

where the first inequality is by the first bullet of Theorem G.1, the second inequality is due to the fact that $\log\Big(\sum_{k \in [K]} \exp\big(\big[\widehat{f}(\mathbf{x}, \mathbf{y}_k) - \widehat{f}(\mathbf{x}, \mathbf{y}^*)\big]/\tau\big)\Big) > 0$, the last inequality is due to $\log\Big(\sum_{k \in [K]} \exp\big(\big[\widehat{f}(\mathbf{x}, \mathbf{y}_k) - \widehat{f}(\mathbf{x}, \mathbf{y}^*)\big]/\tau\big)\Big) \geq \log(1 + e^{-1})$ since there exists at least one similarity score $\widehat{f}(\mathbf{x}, \mathbf{y}_k)$ greater than $\widehat{f}(\mathbf{x}, \mathbf{y}^*) - \tau$ with $\mathbf{y}_k \neq \mathbf{y}^*$. Therefore, we have that $\mathbb{P}(\mathcal{E}) \leq \epsilon'/\log(1 + e^{-1}) \leq 4\epsilon'$ which completes the proof. $\square$

Discussion of Theorem 4.2. The reason is that softmax function $L(\mathbf{a}) = \log(\sum_i \exp(a_i))$ is convex but not strongly convex and has an exponential-decaying tail. Once the score function f with the features $\mathbf{g}$ and $\mathbf{h}$ achieves the margin of order $\Omega(\tau)$, the gradient will exponentially decrease. Therefore, the weight will not be updated effectively.

Proof of Theorem 4.2. Consider the simplest setting where $\boldsymbol{\xi}$ and $\boldsymbol{\zeta}$ are all zero vectors, and we can access to the population loss and its gradient (notice that we are constructing the negative example). We will show that even under this ideal setting, the learned score function with corresponding representations may not achieve a margin greater than $\widetilde{O}(\tau)$. Notice that

$$
\begin{aligned}
\nabla_{\mathbf{W}} \mathbb{E}_{\mathcal{D}^B} L(f,\tau) &= \nabla_{\mathbf{W}} \mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]} \exp\big(\big[f(\mathbf{x}_1,\mathbf{y}_t) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)\bigg)\bigg] \\
&\quad + \nabla_{\mathbf{W}} \mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]} \exp\big(\big[f(\mathbf{x}_t,\mathbf{y}_1) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)\bigg)\bigg] \\
&= \mathbb{E}\bigg[\nabla_{\mathbf{W}} \log\bigg(\sum_{t\in[B]} \exp\big(\big[f(\mathbf{x}_1,\mathbf{y}_t) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)\bigg)\bigg] \\
&\quad + \mathbb{E}\bigg[\nabla_{\mathbf{W}} \log\bigg(\sum_{t\in[B]} \exp\big(\big[f(\mathbf{x}_t,\mathbf{y}_1) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)\bigg)\bigg] \\
&= \mathbb{E}\bigg[\sum_{t\in[B]} \frac{\exp\big(\big[f(\mathbf{x}_1,\mathbf{y}_t) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)}{\sum_s \exp\big(\big[f(\mathbf{x}_1,\mathbf{y}_s) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)} (\mathbf{y}_t - \mathbf{y}_1)\mathbf{x}_1^\top\bigg] \\
&\quad + \mathbb{E}\bigg[\sum_{t\in[B]} \frac{\exp\big(\big[f(\mathbf{x}_t,\mathbf{y}_1) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)}{\sum_s \exp\big(\big[f(\mathbf{x}_s,\mathbf{y}_1) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)} \mathbf{y}_1(\mathbf{x}_t - \mathbf{x}_1)^\top\bigg] \\
&= \mathbb{E}\bigg[\sum_{t\in[B]} \frac{\mathbb{1}(\mathbf{z}_t \neq \mathbf{z}_1)\exp\big(\big[f(\mathbf{x}_1,\mathbf{y}_t) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)}{\sum_s \exp\big(\big[f(\mathbf{x}_1,\mathbf{y}_s) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)} (\mathbf{y}_t - \mathbf{y}_1)\mathbf{x}_1^\top\bigg] \\
&\quad + \mathbb{E}\bigg[\sum_{t\in[B]} \frac{\mathbb{1}(\mathbf{z}_t \neq \mathbf{z}_1)\exp\big(\big[f(\mathbf{x}_t,\mathbf{y}_1) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)}{\sum_s \exp\big(\big[f(\mathbf{x}_s,\mathbf{y}_1) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)} \mathbf{y}_1(\mathbf{x}_t - \mathbf{x}_1)^\top\bigg]
\end{aligned}
$$

where the last inequality is by $\mathbf{x}_t = \mathbf{x}_1$ and $\mathbf{y}_t = \mathbf{y}_1$ when $\mathbf{z}_t = \mathbf{z}_1$. Therefore suppose function f can achieve a margin greater than $\log\Big(16\|\mathbf{G}\|_2^2\|\mathbf{H}\|_2^2(R^2+1)^2 B\tau^{-1}\eta T\Big)\tau$, we have that the gradient

$$
\begin{aligned}
&\Big\|\nabla_{\mathbf{W}} \mathbb{E}_{\mathcal{D}^B} L(f,\tau)\Big\|_F \\
&\leq 2\|\mathbf{G}\|_2\|\mathbf{H}\|_2(R^2+1)\cdot \mathbb{E}\bigg[\sum_{t\in[B]} \frac{\mathbb{1}(\mathbf{z}_t \neq \mathbf{z}_1)\exp\big(\big[f(\mathbf{x}_1,\mathbf{y}_t) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)}{\sum_s \exp\big(\big[f(\mathbf{x}_1,\mathbf{y}_s) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)}\bigg] \\
&\quad + 2\|\mathbf{G}\|_2\|\mathbf{H}\|_2(R^2+1)\cdot \mathbb{E}\bigg[\sum_{t\in[B]} \frac{\mathbb{1}(\mathbf{z}_t \neq \mathbf{z}_1)\exp\big(\big[f(\mathbf{x}_t,\mathbf{y}_1) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)}{\sum_s \exp\big(\big[f(\mathbf{x}_s,\mathbf{y}_1) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)}\bigg] \\
&\leq 2\|\mathbf{G}\|_2\|\mathbf{H}\|_2(R^2+1)\cdot \mathbb{E}\bigg[\mathbb{1}(\mathbf{z}_t \neq \mathbf{z}_1)\sum_{t\in[B]} \exp\big(\big[f(\mathbf{x}_1,\mathbf{y}_t) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big) \\
&\quad + 2\|\mathbf{G}\|_2\|\mathbf{H}\|_2(R^2+1)\cdot \mathbb{E}\bigg[\sum_{t\in[B]} \mathbb{1}(\mathbf{z}_t \neq \mathbf{z}_1)\exp\big(\big[f(\mathbf{x}_t,\mathbf{y}_1) - f(\mathbf{x}_1,\mathbf{y}_1)\big]/\tau\big)\bigg] \\
&\leq 0.25\tau\|\mathbf{G}\|_2^{-1}\|\mathbf{H}\|_2^{-1}(R^2+1)^{-1}\eta^{-1}T^{-1},
\end{aligned} \tag{I.1}
$$

is very small. Now suppose the SGD trajectory start at $\mathbf{W}^{(0)} = 2\log\Big(16\|\mathbf{G}\|_2^2\|\mathbf{H}\|_2^2(R^2+1)^2 B\tau^{-1}\eta T\Big)\cdot(\tau/\gamma)\mathbf{W}^*$. Obviously the score function with weight $\mathbf{W}^{(0)}$ achieve a margin $2\log\Big(16\|\mathbf{G}\|_2^2\|\mathbf{H}\|_2^2(R^2+1)^2 B\tau^{-1}\eta T\Big)\tau$. Suppose there exists a time $t \leq T$ such that $\langle \mathbf{W}^{(t)}\mathbf{x},\mathbf{y}\rangle$ can achieve margin larger than $3\log\Big(16\|\mathbf{G}\|_2^2\|\mathbf{H}\|_2^2(R^2+1)^2 B\tau^{-1}\eta T\Big)\tau$ or can achieve margin larger than $\log\Big(16\|\mathbf{G}\|_2^2\|\mathbf{H}\|_2^2(R^2+1)^2 B\tau^{-1}\eta T\Big)\tau$. Then there must exist a first time $t < t'$ such that the margin at time t lies outsize the range between $\log\Big(16\|\mathbf{G}\|_2^2\|\mathbf{H}\|_2^2(R^2+1)^2B\tau^{-1}\eta T\Big)\tau$ and $3\log\Big(16\|\mathbf{G}\|_2^2\|\mathbf{H}\|_2^2(R^2+1)^2B\tau^{-1}\eta T\Big)\tau$. By definition of t (margin gap), we know that there exist $\mathbf{x},\mathbf{y}$ such that $|\langle\mathbf{W}^{(t)}\mathbf{x},\mathbf{y}\rangle-\langle\mathbf{W}^{(0)}\mathbf{x},\mathbf{y}\rangle|>\tau$. On the other hand, we have that

$$\begin{aligned}\big|\langle\mathbf{W}^{(t)}\mathbf{x},\mathbf{y}\rangle-\langle\mathbf{W}^{(0)}\mathbf{x},\mathbf{y}\rangle\big| &\le \|\mathbf{W}^{(t)}-\mathbf{W}^{(0)}\|_F\|\mathbf{x}\mathbf{y}^\top\|_F\\ &\le 2\|\mathbf{G}\|_2\|\mathbf{H}\|_2(R^2+1)\|\mathbf{W}^{(t)}-\mathbf{W}^{(0)}\|_F\\ &\le 2\|\mathbf{G}\|_2\|\mathbf{H}\|_2(R^2+1)\cdot\eta T\cdot 0.25\tau\|\mathbf{G}\|_2^{-1}\|\mathbf{H}\|_2^{-1}(R^2+1)^{-1}\eta^{-1}T^{-1}\\ &\le 0.5\tau,\end{aligned}$$

a contradiction! Therefore, such a t doesn't exist. The score function learned by SGD within T iterations can't achieve a margin greater than $3\log\Big(16\|\mathbf{G}\|_2^2\|\mathbf{H}\|_2^2(R^2+1)^2B\tau^{-1}\eta T\Big)\tau$. $\square$

Theorem I.1 (Formal statement of Theorem 4.3)**.** Under the same condition as Theorem 3.4, with $\boldsymbol{\zeta}=\mathbf{0}$. (This problem setting includes the special case considered in Theorem 4.2.) Let $\epsilon\le\lambda\gamma^2\min p_k/(3200\|\mathbf{H}\|_2^2)$ and $\tau\le\gamma/\log(\gamma^2\min p_k/(6400B\|\mathbf{H}\|_2^2))$, within polynomial iterations, we can find a score function $\widehat{f}$ with large margin. In particular, with a probability of at least 0.99, the top-1 result gives the correct label with a margin of at least 0.5γ.

Proof. For simplicity, consider the case that we can access the population loss and its gradient, i.e., $n\to\infty$. The regularized loss then becomes,

$$L^{new}=L_{\mathcal{D}^B}(f,\tau)+\lambda\mathbb{E}[\|\mathbf{g}(\mathbf{x})-\mathbf{h}(\mathbf{y})\|_2^2].$$

Since the new loss is still convex and even strongly convex. By applying the same technique in the proof of the Theorem 3.4, within polynomial iterations, we can find $L^{new}(f,\tau,\lambda)\le L^{new}(f^*,\tau,\lambda)+\epsilon$. Besides,

$$L^{new}(f^*,\tau,\lambda)=L_{\mathcal{D}^B}(f^*,\tau)\le 2\mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]}\mathbb{1}(\mathbf{z}_t=\mathbf{z}_1)\bigg)\bigg]+2B\exp(-\gamma/\tau)$$

where the first equality is by plugging in $\mathbf{W}^*=\mathbf{H}(\mathbf{H}^\top\mathbf{H})^{-1}\mathbf{P}(\mathbf{G}^\top\mathbf{G})^{-1}\mathbf{G}^\top,\mathbf{g}(\mathbf{x})=\mathbf{W}\mathbf{x},\mathbf{h}(\mathbf{y})=\mathbf{y}$, the inequality is by Lemma G.7. Thus we have that

$$L_{\mathcal{D}^B}(f,\tau)+\lambda\mathbb{E}[\|\mathbf{g}(\mathbf{x})-\mathbf{h}(\mathbf{y})\|_2^2]\le 2\mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]}\mathbb{1}(\mathbf{z}_t=\mathbf{z}_1)\bigg)\bigg]+\epsilon',$$

where $\epsilon'=\epsilon+2B\exp(-\gamma/\tau)$. By (G.10) and (G.11), we know that $L_{\mathcal{D}^B}(f,\tau)\ge 2\mathbb{E}\bigg[\log\bigg(\sum_{t\in[B]}\mathbb{1}(\mathbf{z}_t=\mathbf{z}_1)\bigg)\bigg]$. Therefore, we can conclude that

$$\mathbb{E}[\|\mathbf{g}(\mathbf{x})-\mathbf{h}(\mathbf{y})\|_2^2]\le\epsilon'/\lambda\le\gamma^2\min p_k/(1600\|\mathbf{H}\|_2^2),$$

where the last inequality is by choose $\epsilon\le\lambda\gamma^2\min p_k/(3200\|\mathbf{H}\|_2^2)$ and $\tau\le\gamma/\log(\gamma^2\min p_k/(6400B\|\mathbf{H}\|_2^2))$. Then by Chebyshev's inequality, for any $\mathbf{z}$, with probability $1-0.01$ we have $\|\mathbf{g}(\mathbf{x})-\mathbf{h}(\mathbf{y})\|_2\le\sqrt{100\max p_k^{-1}\mathbb{E}[\|\mathbf{g}(\mathbf{x})-\mathbf{h}(\mathbf{y})\|_2^2]}\le\gamma/(4\|\mathbf{H}\|_2)$. Then for any $\mathbf{y}'$ that has the different shared feature from $\mathbf{y}$ (i.e., $\mathbf{z}'\neq\mathbf{z}$) we have that

$$\begin{aligned}&\langle\mathbf{g}(\mathbf{x}),\mathbf{h}(\mathbf{y}')\rangle-\langle\mathbf{g}(\mathbf{x}),\mathbf{h}(\mathbf{y})\rangle\\ &\le\langle\mathbf{h}(\mathbf{y}),\mathbf{h}(\mathbf{y}')\rangle-\langle\mathbf{h}(\mathbf{y}),\mathbf{h}(\mathbf{y})\rangle+\|\mathbf{g}(\mathbf{x})-\mathbf{h}(\mathbf{y})\|_2\cdot(\|\mathbf{h}(\mathbf{y}')\|_2+\|\mathbf{h}(\mathbf{y})\|_2)\\ &\le-\gamma+\gamma/2\\ &\le-\gamma/2,\end{aligned}$$

where the first inequality is by triangle inequality, the second inequality is by $\|\mathbf{g}(\mathbf{x})-\mathbf{h}(\mathbf{y})\|_2\le\gamma/(4\|\mathbf{H}\|_2)$ and $\|\mathbf{h}(\mathbf{y}')\|_2=\|\mathbf{h}(\mathbf{y})\|_2\le\|\mathbf{H}\|_2$ since $\boldsymbol{\zeta}=\mathbf{0}$. $\square$