

KPL: Training-Free Medical Knowledge Mining of Vision-Language Models

Jiaxiang Liu^{1,2}, Tianxiang Hu^{1,2}, Jiawei Du^{3,4}, Ruiyuan Zhang¹, Joey Tianyi Zhou^{3,4}, Zuozhu Liu^{1,2*}

¹ ZJU-Angelalign R&D Center for Intelligence Healthcare, ZJU-UIUC Institute, Zhejiang University, China

² Zhejiang Key Laboratory of Medical Imaging Artificial Intelligence, Zhejiang University, China

³ Centre for Frontier AI Research (CFAR), Agency for Science, Technology and Research (A*STAR), Singapore

⁴ Institute of High Performance Computing (IHPC), Agency for Science, Technology and Research (A*STAR), Singapore
{jiaxiang.21, zuozhuliu}@intl.zju.edu.cn

Abstract

Visual Language Models such as CLIP excel in image recognition due to extensive image-text pre-training. However, applying the CLIP inference in zero-shot classification, particularly for medical image diagnosis, faces challenges due to: 1) the inadequacy of representing image classes solely with single category names; 2) the modal gap between the visual and text spaces generated by CLIP encoders. Despite attempts to enrich disease descriptions with large language models, the lack of class-specific knowledge often leads to poor performance. In addition, empirical evidence suggests that existing proxy learning methods for zero-shot image classification on natural image datasets exhibit instability when applied to medical datasets. To tackle these challenges, we introduce the Knowledge Proxy Learning (KPL) to mine knowledge from CLIP. KPL is designed to leverage CLIP’s multimodal understandings for medical image classification through *Text Proxy Optimization* and *Multimodal Proxy Learning*. Specifically, KPL retrieves image-relevant knowledge descriptions from the constructed knowledge-enhanced base to enrich semantic text proxies. It then harnesses input images and these descriptions, encoded via CLIP, to stably generate multimodal proxies that boost the zero-shot classification performance. Extensive experiments conducted on both medical and natural image datasets demonstrate that KPL enables effective zero-shot image classification, outperforming all baselines. These findings highlight the great potential in this paradigm of mining knowledge from CLIP for medical image classification and broader areas.

Code — <https://github.com/JXLiu-AI/KPL>

Introduction

Zero-shot image classification is a crucial task in various domains (Radford et al. 2021; Wang et al. 2022a; Qian, Xu, and Hu 2024; Ren, Su, and Liu 2024; Zhou et al. 2023; Jiao et al. 2023; Novack et al. 2023; Ali and Khan 2023), especially where comprehensive datasets covering all possible classes are unattainable and annotated images are scarce. An exemplary application is zero-shot medical image diagnosis (Wang et al. 2022b; Liu et al. 2023b; Mahapatra, Borzorgtabar, and Ge 2021; Mahapatra, Ge, and Reyes 2022;

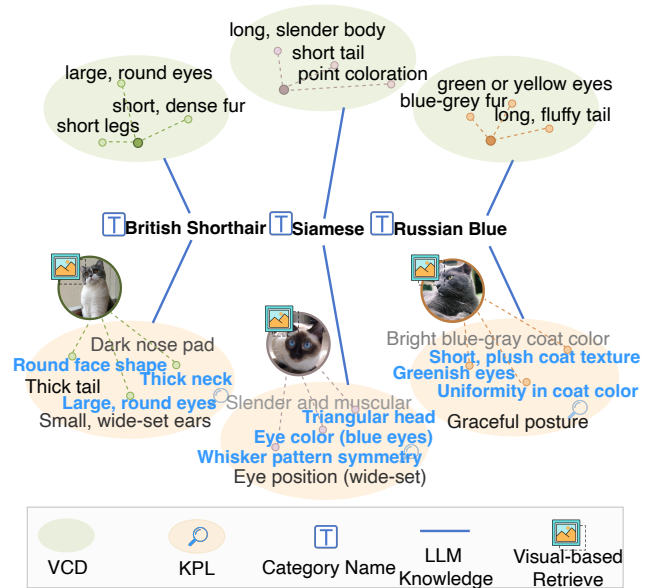


Figure 1: VCD utilizes a limited number of descriptions that may be irrelevant to the images, whereas KPL generates richer descriptions from the knowledge of LLMs. These are visually-based and retrieved to ensure the descriptions’ relevance to the images (Text Proxy Optimization).

Xian et al. 2018; Lampert, Nickisch, and Harmeling 2013), which is essential in real-world healthcare settings. Visual Language Models (VLMs) like CLIP have demonstrated promising performance in these tasks (Radford et al. 2021; Liu et al. 2023a; Eslami, de Melo, and Meinel 2021; Gai et al. 2024; Chen et al. 2024; Wang et al. 2023). After training CLIP on internet-sourced images with their corresponding captions or descriptions to gain image-text alignment (Radford et al. 2021; Wang et al. 2023; Esmailpour et al. 2022), the CLIP inference procedure utilizes class name features as proxies to classify images by comparing the similarity between image features and these text-based proxies (Radford et al. 2021). While CLIP effectively uses unlabeled internet data for large-scale training to foster a rich multimodal understanding, its performance often falls short in specialized classification tasks, such as the medical domain (Liu et al. 2023b; Lai et al. 2024; Liu et al. 2024a,b).

*Corresponding author.




Figure 2: KPL is designed to leverage CLIP’s capabilities for medical image classification through *Text Proxy Optimization* and *Multimodal Proxy Learning*. First, semantic information is enriched for the cataract category through Text Proxy Optimization. Then, Multimodal Proxy Learning generates a multimodal proxy combining textual semantics and image knowledge to guide classification.

This has led us to reconsider the existing CLIP inference paradigm.

We question whether the CLIP scheme fully utilizes its capabilities during inference. We contend that several facets of the vanilla CLIP inference methodology might substantially constrain the effectiveness of zero-shot image classification. Firstly, using single class names as standalone representatives is questionable (Ren, Su, and Liu 2024; Menon and Vondrick 2023), as CLIP models are trained to align images with descriptions, not just single class names (eg, synonyms might show up in the pre-training text). This simplification, while practical, is suboptimal and fails to fully leverage the rich semantic alignments developed during training. Secondly, in zero-shot scenarios like medical image diagnosis, class names are often abstract, complex, and rarely present in CLIP’s pre-training data, posing challenges to the text encoder’s ability to interpret them accurately (Liu et al. 2023b). Thirdly, prior research indicates a persistent disparity between the text and vision feature spaces in CLIP (Qian, Xu, and Hu 2024; Liang et al. 2022), suggesting that direct similarity calculations may not fully capture the nuanced relationships needed for accurate classification.

Recent studies (Menon and Vondrick 2023; Ren, Su, and Liu 2024; Liu et al. 2023b; Novack et al. 2023; Shi, Fan, and Yan 2024) have been exploring how to leverage CLIP’s knowledge. To address the inadequacy of using single class names for inference, Menon et al. and Ren et al. (Menon and Vondrick 2023; Ren, Su, and Liu 2024) proposed description-based classification methods that utilize large language models (LLMs) to enrich class names with descriptions to generate more sophisticated text proxies for conducting similarity calculations with images. In the medical field, Liu et al. (Liu et al. 2023b) enhanced CLIP models with disease details, mimicking human diagnostics. Despite the efforts of unleashing more of CLIP’s textual ability via providing additional text, current description-based methods face the issue that the enriched descriptions are not semantically rich or relevant enough for the image classification tasks. For example, VCD (Menon and Vondrick 2023) often misclassifies semantically similar categories, such as Siamese, British Shorthair, and Russian Blue in the Oxford Pets dataset (Zhang et al. 2022). This misclassification stems from its reliance on descriptions like “short tail” and “long, slender body” which do not accurately reflect the im-

ages, compounded by the absence of relevant descriptors like “round face shape” and “thick neck”. (see Figure 1) These observations underscore the urgent need for developing a more precisely tailored description generation strategy.

To address the modal gap issue, Qian et al. (Qian, Xu, and Hu 2024) proposed learning proxies on the vision space, and employing the Sinkhorn to refine pseudo labels computed by text proxies, generated using the original category names, to guide the learning. However, current proxy learning methods, just like the CLIP inference, lacks text semantic richness, which can lead to inaccurate vision proxies that deviate from the intended optimal in the vision space. Moreover, as illustrated in Figure 4, the Sinkhorn used for pseudo label refinement exhibits significant instability for medical image classification tasks, which highlights the need for a more robust algorithm in medical imaging to minimize information loss during the generation of pseudo labels which carry the knowledge from CLIP to guide the identification of optimal proxies across the modal gap.

In this paper, we propose a **Knowledge Proxy Learning (KPL)** to unlock CLIP’s potential by mining knowledge from CLIP. Specifically, KPL employs LLMs to build a knowledge-enhanced base for each category, ensuring semantically rich text resources. Then, KPL enriches semantic information of category names through multi-level text and image retrieval for category-specific image-relevant knowledge from these comprehensive databases to optimize text proxies. Finally, The Stable GreenHorn (SG) is proposed to refine the text proxies to guide the learning of multimodal proxies, which bridges the modal gap and results in improved decision boundaries. Comprehensive experimental results show that our method consistently outperforms all baseline methods on various datasets for diseases like pneumonia, glaucoma, cataracts, diabetic retinopathy, and malaria. Key contributions are summarized as follows:

- Our findings reveal a discernible gap between current methodologies and the theoretical optimal. This research advocates for a focused and in-depth exploration into the generation of multimodal proxy that exploits the existing capabilities of the CLIP model, with the ultimate goal of establishing a new CLIP-based inference paradigm.
- We propose a two-step proxy optimization method, KPL. KPL first builds semantic databases, optimizes text proxies with relevant image descriptions, and reduces semantic ambiguity between categories and images (Text Proxy Optimization). Afterwards, it optimizes multimodal proxies to address modal disparities and reduce performance degradation in medical contexts (Multimodal Proxy Learning).
- Extensive experiments and ablation studies conducted on nine datasets substantiate the efficacy of KPL. Notably, experiments reveal that KPL can also effectively mine knowledge from specialized domain CLIP such as BioMedCLIP (Zhang et al. 2023) trained specifically on medical data to achieve performance improvements. These findings underscore the adaptability of the proposed paradigm.

Preliminaries and Related Work

CLIP Inference Let F, G denote the vision encoder and text encoder in CLIP (Radford et al. 2021; Wang et al. 2022b; Zhang et al. 2023; Jeong et al. 2023). Denote the image dataset for inference as $\mathcal{D} = \{x_i\}_{i=1}^N$, and the category names as $\mathcal{C} = \{c_j\}_{j=1}^K$, the image and text features can be represented as $\{\mathbf{x}_i\}_{i=1}^N, \{\mathbf{z}_j\}_{j=1}^K$ where $\mathbf{x}_i = F(x_i), \mathbf{z}_j = G(c_j)$. The prediction of CLIP inference involves taking the features of category names as proxies $\mathbf{w}_j = \mathbf{z}_j, j = 1, \dots, K$ and calculating similarities between the image features and these proxies for classification $\hat{y}_i = \arg \max_j (\mathbf{x}_i \cdot \mathbf{w}_j)$.

Description-based Zero-shot Classification Method To leverage CLIP’s comprehension of text, Menon et al. (Menon and Vondrick 2023) and Ren et al. (Ren, Su, and Liu 2024) proposed to use LLM-generated category descriptions and a created multilevel knowledge tree as a supplement to category names. Liu et al. (Liu et al. 2023b) have explored description-based CLIP inference in the medical domain. These methods built enriched category text $\mathcal{C}' = \{C_1, \dots, C_K\}$ for classification based on category names $\mathcal{C} = \{c_1, \dots, c_K\}$, where each C_j contains relevant text about c_j , and then took $\mathbf{w}_j = \frac{1}{|\mathcal{C}'_j|} \sum_{t \in \mathcal{C}'_j} G(t)$ as proxies.

Proxy Learning Recent research (Qian, Xu, and Hu 2024; Liang et al. 2022) have shown there is a gap between the text and vision space learned by CLIP. To reduce the modal gap, Qian et al. (Qian, Xu, and Hu 2024) introduced a proxy learning approach InMaP which obtains visual proxies from text proxies for classification via $\mathbf{w}_1, \dots, \mathbf{w}_K = \arg \min_{\mathbf{w}_1, \dots, \mathbf{w}_K} \mathbb{E}_{x \sim \mathcal{D}} [KL(P'(x) || P(x; \mathbf{w}_1, \dots, \mathbf{w}_K))]$, where $P(x; \mathbf{w}_1, \dots, \mathbf{w}_K) = \left[\frac{\exp(\mathbf{x} \cdot \mathbf{w}_1 / \tau)}{\sum_{j=1}^K \exp(\mathbf{x} \cdot \mathbf{w}_j / \tau)}, \dots, \frac{\exp(\mathbf{x} \cdot \mathbf{w}_K / \tau)}{\sum_{j=1}^K \exp(\mathbf{x} \cdot \mathbf{w}_j / \tau)} \right]$ denotes the class distribution calculated by the proxies to learn $\{\mathbf{w}_j\}_{j=1}^K$, and $P'(x)$ denotes the pseudo ground truth class distribution generated by text proxies $\{\mathbf{z}_j\}_{j=1}^K$. Specifically, InMaP obtains the pseudo ground truth by solving the following optimization problem using Sinkhorn algorithm (Sinkhorn 1964; Altschuler, Niles-Weed, and Rigollet 2017): $\max_{P' \in \mathcal{P}} (\langle M, P' \rangle + \tau H(P'))$, $\mathcal{P} = \{P' | \forall i, \sum_j P'_{i,j} = \frac{1}{n}; \forall j, \sum_i P'_{i,j} = q_j; \forall i, j, P'_{i,j} \geq 0\}$, where $M_{i,j} = \mathbf{x}_i^T \mathbf{z}_j$, H is entropy function of matrices, $q \in \mathbb{R}^K$ is a reference distribution over classes, and $\tau \in \mathbb{R}$ is the temperature.

IPFP Iterative Proportional Fitting Procedure (IPFP), also known as biproportional fitting in various fields such as statistics, economics, and computer science (Stephan 1942; Bacharach 1965; Idel 2016; Chang et al. 2023), has a rich history, being reexplored many times since its first appearing in (Yule 1912; Kruihof 1937). The attribution of the proof for uniqueness and convergence within this process is credited to Sinkhorn’s seminal work (Sinkhorn 1964), which subsequently led to the naming of the algorithm in his honor. Recent efforts to analyze and enhance the Sinkhorn algorithm have led to developments like Greenhorn (Altschuler, Niles-Weed, and Rigollet 2017). Studies (Lin, Ho, and Jordan 2022, 2019; Altschuler, Niles-Weed, and Rigollet 2017) demonstrate that Greenhorn retains similar theoretic

cal guarantees and has better performance in practice.

Methodology

Overview

We focus on zero-shot image classification using CLIP (Radford et al. 2021). Recognizing the limitations of relying solely on category names and the existing text enrichment methods, as well as the issue of modality gaps, we propose a novel CLIP inference paradigm, KPL. Let F, G denote the vision encoder and text encoder in CLIP (Radford et al. 2021). Denote the image dataset for inference as $\mathcal{D} = \{x_i\}_{i=1}^N$, and the class names as $\mathcal{C} = \{c_j\}_{j=1}^K$. The KPL pipeline, illustrated in Figure 3, starts with \mathcal{C} , processed through a Text Proxy Optimization (TPO) which leverages CLIP capability, followed by encoding via text encoder G to produce semantic embedding centers. Visual images \mathcal{D} are processed by the visual encoder F to obtain visual features. Subsequently, semantic embedding centers are employed to guide the Multimodal Proxy Learning to obtain multimodal proxies:

$$\mathbf{W}_{KPL}(\mathcal{D}, \mathcal{C}) = \arg \min_{\mathbf{W}} \mathbb{E}_{x \sim \mathcal{D}} [d(Q_{F,G}(x; \mathcal{D}, \mathcal{C}), P_F(x, \mathbf{W}))], \quad (1)$$

$$P_F(x, \mathbf{W})_j = \frac{\exp(F(x) \cdot \mathbf{w}_j / \tau)}{\sum_{k=1}^K \exp(F(x) \cdot \mathbf{w}_k / \tau)}, \forall j = 1, \dots, |\mathcal{C}|$$

$Q_{F,G}$ leverages CLIP and data to generate pseudo labels, and d is KL divergence by default.

Generated multimodal proxies \mathbf{W}_{KPL} are integrated with visual features $\{F(x)\}_{x \in \mathcal{D}}$ to produce classification results $\hat{y} = f(x) = \arg \max_{j \in \{1, \dots, |\mathcal{C}|\}} F(x) \cdot \mathbf{w}_{KPL_j}, \forall x \in \mathcal{D}$. In the

following sections, we introduce in details the *Text Proxy Optimization* and the *Multimodal Proxy Learning*.

Text Proxy Optimization

Knowledge-Enhanced Base (KEB) To overcome the inadequacies of current text representations for categories (CLIP inference or current description-based CLIP inference), we propose to first utilize LLMs to generate category descriptions $\mathcal{C}_1 = \{C_1, \dots, C_K\}$:

$$C_j = \{t_{C_j}^1, \dots, t_{C_j}^m\} = \text{LLM}(\text{prompt}, c_j), \quad (2)$$

which build our knowledge-enhanced base. These prompts should be designed to elicit a certain amount of category descriptions that aid in classification. In our experiments, the LLM used is ChatGPT, and the prompt for generating these descriptions is: *According to published literature, what are useful medical visual features for distinguishing [Category] in an image? List as many features as possible, at least 50.* This approach aims to generate as many features as possible, while avoiding semantic loss due to the randomness inherent in LLMs. For instance, for the category of [Cataract], we utilize LLMs to generate descriptive terms, such as [“Cloudy or opaque area in the lens”] and [“Vacuoles in the lens”], which we employ to construct a comprehensive knowledge base for [Cataract] (Figure 2 and Appendix Figure 6).

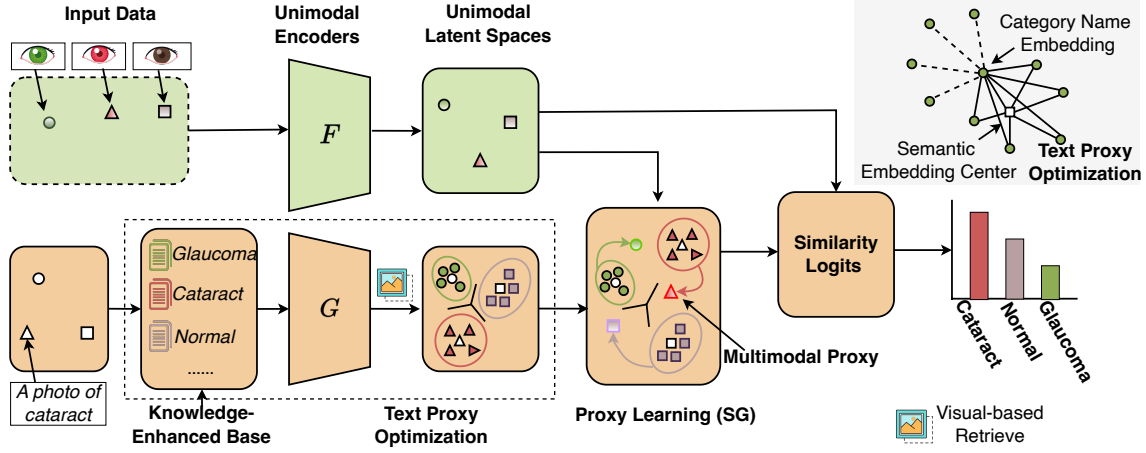


Figure 3: Overview of KPL. 1) Category texts (e.g., cataract) pass through a Knowledge-Enhanced Base and are encoded by a text encoder to produce semantic embedding centers via visual-based retrieval. Visual images are processed through a visual encoder to obtain visual features. 2) CLIP uses the category name embeddings as text proxies to compare with visual features for classification results. 3) KPL uses semantic embedding centers as text proxies to direct Multimodal Proxy Learning. Final classification results are obtained via the generated multimodal proxies and visual features.

Visual-based Retrieval Despite having an extensive knowledge base, not all descriptions within it are directly related to images. Therefore, we further leverage CLIP’s text-image alignment capabilities to retrieve descriptions that is more relevant to images for our classification tasks. Specifically, we select the top k descriptions from the knowledge-enhanced base with the highest similarity to the image x , which finally forms $\mathcal{C}_2 = \{C'_1, \dots, C'_K\}$ ($\mathcal{C}_2 = \text{KEB}(\mathcal{C})$).

$$C'_j = \{t_{C'_j}^{\sigma(1)}, \dots, t_{C'_j}^{\sigma(k)}\}, \quad (3)$$

where $\{s_{\sigma(1)}, \dots, s_{\sigma(n)}\} = \text{sort}_{\downarrow}(\{s_1, \dots, s_n\})$,

$$s_l = \cos\left\langle \frac{1}{|\mathcal{D}|} \sum_{x \in \mathcal{D}} F(x), G(t_{C'_j}^l) \right\rangle, \forall l = 1, \dots, n.$$

For example, for a cataract image, visual-based retrieval can identify from KEB relevant descriptions such as [“Cloudy or opaque area in the lens”] as shown in Figure 2.

As depicted in Figure 3, the core of the Text Proxy Optimization involves constructing an extensive knowledge base (by setting the number requirement of descriptions in the prompts to LLMs) that encompasses detailed descriptions of each class, and leveraging CLIP’s existing alignment to identify the descriptions most helpful for classifying dataset \mathcal{D} . While $\frac{1}{|C'_j|} \sum_{t \in C'_j} G(t)$ can function as text proxies for classification by itself, the presence of a modal gap necessitates a more nuanced approach. Therefore, we propose to utilize it to guide the Multimodal Proxy Learning, leveraging it in a more effective manner.

Multimodal Proxy Learning

Following previous work on proxy learning, we use the solution of the optimization problem in the following to guide the learning of Eq1:

$$Q_{F,G}(\mathcal{D}, \mathcal{C}) = \arg \max_{P: P_{i,j} \geq 0, \sum_j P_{i,j} = \frac{1}{|\mathcal{D}|}, \sum_i P_{i,j} = q_j} (\langle M, P \rangle + \tau H(P)),$$

$$M_{i,j} = F(x_i) \cdot \frac{1}{|C'_j|} \sum_{t \in C'_j} G(t), \quad C'_j \in \text{KEB}(\mathcal{C}),$$

$$\forall i = 1, \dots, |\mathcal{D}|, \quad \forall j = 1, \dots, |\mathcal{C}| \quad (4)$$

H denotes the entropy function, q in $\mathbb{R}^{|\mathcal{C}|}$ is the reference distribution over classes, and τ in \mathbb{R} is the temperature.

Note $Q_{F,G}(x; \mathcal{D}, \mathcal{C})$ indicates the row of $Q_{F,G}(\mathcal{D}, \mathcal{C})$ corresponding to image x .

Proposition 1 The optimization problem in Eq4 has a unique solution of the form

$$P = \text{diag}(u) \text{A} \text{diag}(v), \quad (5)$$

where $\text{diag}(u), \text{diag}(v)$ are two diagonal matrices with diagonals taken from vectors u, v , and $A = e^{\frac{M}{\tau}}$. See proof in Appendix.

Consider jointly the form of matrix P from Proposition 1 and the constraint on its rows and columns as delineated in the optimization problem, the optimization task essentially equates to estimate matrix P given matrix $A = e^{\frac{M}{\tau}}$ such that $P_{i,j} = u_i A_{i,j} v_j$, $\sum_j P_{i,j} = \frac{1}{N}$, $\sum_i P_{i,j} = q_j$. This can be commonly achieved by the Sinkhorn algorithm (η is iteration index of Sinkhorn) (Kruithof 1937; Sinkhorn 1964):

$$\text{Set } P_{i,j}^{(0)} = A_{i,j}, \text{ and for all } \eta \geq 1, \quad (6)$$

$$P_{i,j}^{(2\eta-1)} = \frac{P_{i,j}^{(2\eta-2)} * \frac{1}{N}}{\sum_{k=1}^K P_{i,k}^{(2\eta-2)}}, P_{i,j}^{(2\eta)} = \frac{P_{i,j}^{(2\eta-1)} * q_j}{\sum_{k=1}^N P_{k,j}^{(2\eta-1)}}. \quad (7)$$

However, our empirical experiments indicate that pseudo labels generated by Sinkhorn did not consistently improve medical image classification performance, highlighting potential limitations in its efficacy (see Appendix Figure 5).

Notably, prior research has explored the Greenkhorn algorithm (Altschuler, Niles-Weed, and Rigollet 2017), a greedy variant of Sinkhorn which maintains similar theoretical guarantees to Sinkhorn but has been shown to significantly outperform it in practice (Lin, Ho, and Jordan 2022, 2019; Altschuler, Niles-Weed, and Rigollet 2017). Additionally, the traditional Sinkhorn algorithm faces computational challenges like underflow and overflow issues. Computing in logarithmic space can improve stability and increase error tolerance (see Appendix Proof), which is essential for handling the high dynamic range and potential extremes in medical imaging (Karim et al. 2004; Hung et al. 2013). Therefore, we proposed the SG algorithm for Multimodal Proxy Learning:

$$\text{Set } \log(P_{i,j}^{(0)}) = \frac{M_{i,j}}{\tau}, \text{ and for all } 1 \leq \eta \leq N_{\max},$$

$$\text{if } \max_i \left| \sum_j P_{i,j}^{(\eta-1)} - \frac{1}{N} \right| > \max_j \left| \sum_i P_{i,j}^{(\eta-1)} - q_j \right| :$$

$$r = \arg \max_i \left| \sum_j P_{i,j}^{(\eta-1)} - \frac{1}{N} \right|,$$

$$\log(P_{r,j}^{(\eta)}) = \log(P_{r,j}^{(\eta-1)}) + \log\left(\frac{1}{N}\right) - \log\left(\sum_{k=1}^K P_{r,k}^{(\eta-1)}\right),$$

$$\text{elif } \max_i \left| \sum_j P_{i,j}^{(\eta-1)} - \frac{1}{N} \right| \leq \max_j \left| \sum_i P_{i,j}^{(\eta-1)} - q_j \right| :$$

$$r = \arg \max_j \left| \sum_i P_{i,j}^{(\eta-1)} - q_j \right|,$$

$$\log(P_{i,r}^{(\eta)}) = \log(P_{i,r}^{(\eta-1)}) + \log(q_r) - \log\left(\sum_{k=1}^N P_{k,r}^{(\eta-1)}\right).$$

By selectively updating rows or columns with the greatest deviation from the target distribution, SG can address the inherent non-uniformity and complexity of medical datasets more effectively, thereby ensuring the preservation of knowledge mined from CLIP (see further analysis in Appendix). Consequently, the pseudo labels $Q_{F,G}(\mathcal{D}, \mathcal{C})$, which encapsulate this knowledge, guide the learning in Eq1 to bridge the modal gap and establish multimodal proxies \mathbf{W}_{KPL} . Finally, the classification is conducted on the image dataset \mathcal{D} : $\hat{y} = f(x) = \underset{j \in \{1, \dots, K\}}{\operatorname{argmax}} F(x) \cdot \mathbf{w}_{KPL_j}, \forall x \in \mathcal{D}$.

Experiments

Experimental Setup

To validate KPL in medical zero-shot scenarios, we utilized several datasets including Shenzhen (Jaeger et al. 2014), IDRiD (Porwal et al. 2018), MalariaCell (Hassan et al. 2022), Cataract (Rokhana, Herulambang, and Indraswari 2022), and Montgomery (Jaeger et al. 2014). These datasets cover a diverse range of diseases (Appendix Table 1), which align with previous research (Liu et al. 2023b). Additionally, to test the versatility of KPL, we applied it to a wider array of natural image datasets including CUB (Wah et al. 2011), Places365 (López-Cifuentes et al. 2020), Oxford Pets (Zhang et al. 2022), and ImageNet (Deng et al. 2009). Classification accuracy is used as the evaluation metric in our experiments (Liu et al. 2023b; Guo et al. 2024).

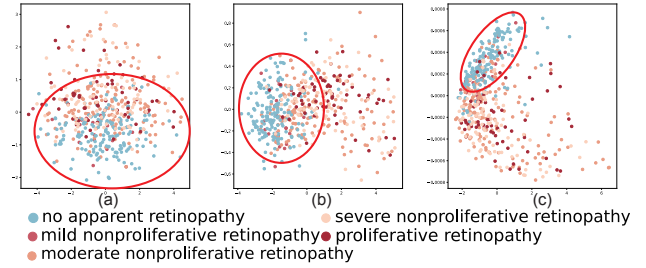


Figure 4: PCA Visualization on the IDRiD with CLIP: (a) CLIP using text proxies for feature visualization shows that the features of "no apparent retinopathy" are very dispersed and difficult to distinguish. (b) Using Sinkhorn for Multimodal Proxy Learning, the features of "no apparent retinopathy" are slightly more clustered. (c) Using the SG for Multimodal Proxy Learning, the features of "no apparent retinopathy" are tightly clustered into a distinct group.

To ensure the applicability across different domains, we employed various pre-trained CLIP vision encoders for different datasets. For the medical image datasets, we used ViT-L/14 and ViT-L/14@336px. For the natural image datasets, we utilized RN-50, ViT-B/16, and ViT-L/14. For the determination of key hyper-parameters, please refer to Appendix Figure 2. We construct the KEB using LLMs (GPT-3.5-turbo) to generate as many descriptions as possible for each category name, ensuring a minimum of $n \geq 50$ descriptions. The parameters k , N_{\max} , and τ are determined through Grid Search (Appendix Figure 1). Pseudocode is provided in Algorithm 1 and 2. We perform experiments in PyTorch framework on NVIDIA GEFORCE RTX 3090 GPU.

Main Results

For medical datasets, we evaluate KPL and compare it to the state-of-the-art (SoTA) zero-shot image classification methods, including CLIP (Radford et al. 2021), VCD (Menon and Vondrick 2023), CMD (Liu et al. 2023b), and InMaP (Qian, Xu, and Hu 2024), with two backbones on five datasets. Results are presented in Table 1. We observe that our method consistently and significantly outperforms the baselines across all datasets and with any backbone used, demonstrating its effectiveness and robustness in medical image classification. This performance extends across various medical disease datasets, including tuberculosis, diabetic retinopathy, malaria cells, glaucoma, and cataracts. Notably, on the MalariaCell dataset, traditional methods struggled to accurately identify malaria cell images. However, KPL manages to effectively mine CLIP's knowledge to increase the accuracy by 50.83% to 80.86%. This significant improvement transformed previously unrecognizable images into highly accurate identifications, suggesting KPL's potential as a new paradigm of employing VLMs for zero-shot medical image classification.

To validate the efficacy of the KPL paradigm, we compared KPL (with Sinkhorn) with SoTA methods such as CLIP (Radford et al. 2021), VCD (Menon and Vondrick 2023), Hierarchical (Ren, Su, and Liu 2024), and InMaP

		Shenzhen	IDRiD	MalariaCell	Cataract	Montgomery
ViT-L/14	CLIP (Radford et al. 2021)	50.76	06.80	48.53	16.64	60.14
	VCD (Menon and Vondrick 2023)	64.65	18.45	49.97	16.64	63.77
	CMD (Liu et al. 2023b)	68.13	20.38	49.88	16.64	59.42
	InMaP (Qian, Xu, and Hu 2024)	69.79	32.75	50.29	34.44	68.12
	KPL (Ours)	75.38	44.77	77.64	37.10	71.74
ViT-L@336px	CLIP (Radford et al. 2021)	50.76	11.65	50.83	16.64	57.97
	VCD (Menon and Vondrick 2023)	69.94	18.45	50.00	16.80	62.32
	CMD (Liu et al. 2023b)	68.88	13.59	49.92	16.80	57.97
	InMaP (Qian, Xu, and Hu 2024)	73.72	38.37	54.66	31.11	60.87
	KPL (Ours)	78.55	47.09	80.86	44.43	70.29

Table 1: Comparison of KPL with baselines across different backbones on medical image datasets.

		CUB	Places365	Pets	ImageNet
RN50	CLIP	49.02	32.48	74.60	54.85
	VCD	49.10	35.79	78.58	60.20
	Hierarchical	49.86	36.83	79.99	60.63
	InMaP	47.38	40.21	86.08	63.17
	KPL ^s (Ours)	50.50	41.60	88.81	63.83
ViT-B/16	CLIP	56.66	38.21	81.68	63.53
	VCD	58.85	40.87	86.56	68.43
	Hierarchical	59.25	41.52	87.19	68.95
	InMaP	57.31	43.79	91.79	71.92
	KPL ^s (Ours)	59.49	44.56	92.67	72.43
ViT-L/14	CLIP	63.46	38.65	87.92	76.60
	VCD	65.02	40.13	92.10	75.36
	Hierarchical	65.45	41.13	92.84	79.64
	InMaP	64.06	44.15	94.80	78.46
	KPL ^s (Ours)	67.82	45.59	95.06	79.26

Table 2: Comparison of KPL with baselines across different backbones on nature image datasets. KPL^s denotes Multi-modal Proxy Learning of KPL with Sinkhorn.

(Qian, Xu, and Hu 2024), using three backbones on four natural image datasets, as shown in Table 2 (The reason for using KPL with Sinkhorn is explained in the Appendix). To ensure a fair comparison, we did not use any designed prompts as templates for the comparative methods. We observed that in most cases, KPL outperforms the baselines across all datasets and backbones, demonstrating its high generalizability in various image domains.

To demonstrate the improvements in representation learning brought by KPL, we perform PCA visualization (Abid et al. 2018) according to the predictions of the model with CLIP (ViT-L/14@336px) on dataset IDRiD and compare the effects in Figure 4. Subfigures (a), (b), (c) denote our models with different learning scheme versions, as described in the caption. We can observe from the visualization that KPL complete model (c) has distinct improvements over the model with naive learning schemes (b). As shown in Figure 4 (c), the individual characteristics of the same type of disease are closer, and the feature distinction of different clusters of disease categories is more obvious, which shows the

	KEB	Ret.	SG	Shenzhen	IDRiD	Mal.	Cataract	Mon.
		✓	✓	76.13	39.15	31.00	31.78	68.12
	✓	✓		72.81	32.36	60.02	32.45	68.84
			✓	76.13	46.51	60.98	31.11	63.77
	✓	✓	✓	78.55	47.09	80.86	44.43	70.29

Table 3: Ablation study on different components with CLIP (ViT-L@336px) backbone on five datasets. Note that when operating without KEB and using visual-based retrieval with VCD descriptions (Menon and Vondrick 2023), $k = 5$. Note that operating without SG denotes using KPL with Sinkhorn.

effectiveness of our designs for feature learning.

Ablation Studies

Effects of KPL in Domain-Specific CLIP We applied KPL to BioMedCLIP (Zhang et al. 2023), trained specifically for the medical domain, with results shown in Table 4. Across all datasets and regardless of the domain knowledge used to train the CLIP (Radford et al. 2021) models, applying KPL with a training-free approach resulted in substantial improvements. This underscores the potential of KPL’s inference paradigm to become a mainstream method for mining knowledge in future CLIP applications. It is noteworthy that the CLIP (ViT-L/14), trained on internet data, surpassed the performance of BioMedCLIP (trained extensively on medical images) when enhanced with KPL. Furthermore, applying KPL to BioMedCLIP still significantly improves its classification performance. Besides, as shown in Figure 6, whether in medical datasets or natural image datasets, VCD experiences a performance leap after the application of Text Proxy Optimization, followed by another leap after Multi-modal Proxy Learning.

Effects of KEB We compared the performance of descriptions from VCD with those from the KEB, as shown in Figure 5. On five medical datasets, the best performance of visual-based retrieval enhanced by VCD descriptions still falls short compared to the performance of the KEB. This underscores the limitations of VCD’s knowledge descriptions and highlights the superior guidance for classification learning provided by the KEB due to its richer knowledge content. Table 3 also shows that KPL outperforms VCD in

Model	Shenzhen	IDRiD	MalariaCell	Cataract	Montgomery
CLIP (Radford et al. 2021)	50.76	06.80	48.53	16.64	60.14
KPL (CLIP)	75.38(+24.62)	44.77(+37.97)	77.64(+29.11)	37.10(+20.46)	71.74(+11.6)
BioMedCLIP (Zhang et al. 2023)	67.98	43.02	72.13	16.47	84.06
KPL (BioMedCLIP)	77.79(+9.81)	48.45(+5.43)	91.61(+19.48)	51.25(+34.78)	89.86(+5.8)

Table 4: KPL on Medical-Domain CLIP.




Figure 5: Ablation studies with the KEB show that Multimodal Proxy Learning with KPL descriptions (KPL_D) outperforms VCD (VCD_D), highlighting KEB’s effectiveness.

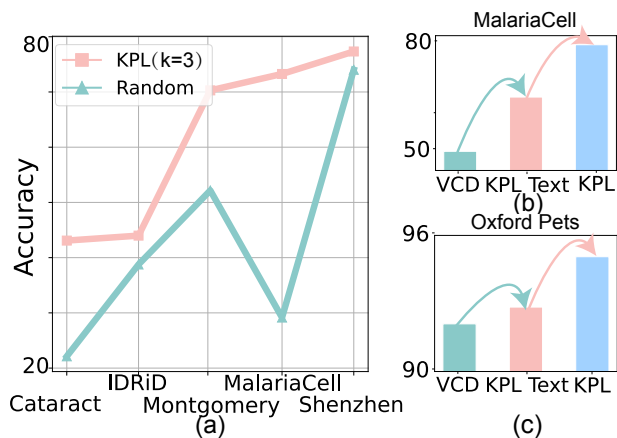


Figure 6: (a) Ablation study for visual-based retrieval at $k = 3$ show that descriptions obtained through visual-based retrieval outperform randomly selected descriptions across five datasets. (b),(c) On the MalariaCell and the Oxford Pets dataset, VCD performance significantly improves after Text Proxy Optimization (KPL Text) and again after Multimodal Proxy Learning (KPL).

retrieval scenarios when $k = 5$. Furthermore, in the Appendix Figure 3, grid search experiments for all k parameters in the Oxford Pets dataset between VCD and KPL further demonstrate the effectiveness of the KEB.

Effects of Visual-based Retrieval Within the KEB, we compared the experimental results of selecting retrieval descriptions with $k = 3$ versus selecting three descriptions at random, as shown in Figure 6. It is evident that on all medical imaging datasets, the performance of enhanced visual-based retrieval outperforms random selection. Notably, on the MalariaCell, Montgomery, and Cataract, the performance of randomly selected descriptions was very low. This could be due to the selection of descriptions irrelevant to the disease category in the image or even related to other diseases, leading to a significant drop in performance.

This demonstrates the effectiveness of visual-based retrieval.

Effects of SG We also compared the Sinkhorn (Sinkhorn 1964) with the proposed SG, as shown in Table 3 and Appendix Table 3. Replacing SG in KPL with the Sinkhorn resulted in performance declines across all datasets. The parameters for the Sinkhorn were set according to the reference (Qian, Xu, and Hu 2024), demonstrating that stable updates to the most imbalanced rows and columns can preserve more knowledge and better guide the learning of multimodal proxies for medical image classification. Furthermore, we also conduct PCA Visualization (Abid et al. 2018) to visualize features under three settings (original CLIP, Sinkhorn-guided proxy learning, and SG-guided proxy learning), as shown in Figure 4. It can be observed that the PCA Visualization features of CLIP’s text proxy remain dispersed and do not effectively distinguish between normal (blue) and diseased retinal features. Using the Sinkhorn-guided multimodal proxy, features of the normal retina show some clustering. Finally, the SG-guided multimodal proxy allows the normal retinal features to cluster into a distinct group, as shown in Figure 4 (c).

Conclusion

In this work, we focus on mining CLIP’s knowledge for zero-shot classification across multiple image domains. Our analysis reveals that the main obstacles to current classification performance are semantically insufficient text descriptions that are unrelated to the images and the modal disparity between visual and textual data. To address these challenges, we introduced a Text Proxy Optimization to generate more accurate descriptions, and then stabilize the pseudo label generation process by utilizing Multimodal Proxy Learning to reduce the modal gap more effectively. With KPL, the zero-shot accuracy of CLIP has been significantly enhanced on both medical image datasets and natural image datasets.

Acknowledgments

This work is supported by the Natural Science Foundation of Zhejiang Province, China (Grant No. LZ23F020008), the National Natural Science Foundation of China (Grant No. 62106222), and the Zhejiang University-Angelalign Inc. R&D Center for Intelligent Healthcare.

References

- Abid, A.; Zhang, M. J.; Bagaria, V. K.; and Zou, J. 2018. Exploring patterns enriched in a dataset with contrastive principal component analysis. *Nature communications*, 9(1): 2134.
- Ali, M.; and Khan, S. 2023. CLIP-Decoder: ZeroShot Multilabel Classification using Multimodal CLIP Aligned Representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4675–4679.
- Altschuler, J.; Niles-Weed, J.; and Rigollet, P. 2017. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. *Advances in neural information processing systems*, 30.
- Bacharach, M. 1965. Estimating nonnegative matrices from marginal data. *International Economic Review*, 6(3): 294–310.
- Chang, S.; Qu, Z.; Leskovec, J.; and Ugander, J. 2023. Inferring Networks from Marginals Using Iterative Proportional Fitting. In *The Second Learning on Graphs Conference*.
- Chen, X.; Lai, Z.; Ruan, K.; Chen, S.; Liu, J.; and Liu, Z. 2024. R-LLaVA: Improving Med-VQA Understanding through Visual Region of Interest. *arXiv preprint arXiv:2410.20327*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Eslami, S.; de Melo, G.; and Meinel, C. 2021. Does CLIP Benefit Visual Question Answering in the Medical Domain as Much as it Does in the General Domain? *arXiv preprint arXiv:2112.13906*.
- Esmailpour, S.; Liu, B.; Robertson, E.; and Shu, L. 2022. Zero-shot out-of-distribution detection based on the pre-trained model clip. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 6568–6576.
- Gai, X.; Zhou, C.; Liu, J.; Feng, Y.; Wu, J.; and Liu, Z. 2024. MedThink: Explaining Medical Visual Question Answering via Multimodal Decision-Making Rationale. *arXiv preprint arXiv:2404.12372*.
- Guo, D.; Li, K.; Hu, B.; Zhang, Y.; and Wang, M. 2024. Benchmarking Micro-action Recognition: Dataset, Method, and Application. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(7): 6238–6252.
- Hassan, E.; Shams, M. Y.; Hikal, N. A.; and Elmougy, S. 2022. A novel convolutional neural network model for malaria cell images classification. *Computers, Materials & Continua*, 72(3): 5889–5907.
- Hung, A. H.; Liang, T.; Sukerkar, P. A.; and Meade, T. J. 2013. High dynamic range processing for magnetic resonance imaging. *PloS one*, 8(11): e77883.
- Idel, M. 2016. A review of matrix scaling and Sinkhorn’s normal form for matrices and positive maps. *arXiv preprint arXiv:1609.06349*.
- Jaeger, S.; Candemir, S.; Antani, S.; Wang, Y.-X. J.; Lu, P.-X.; and Thoma, G. 2014. Two public chest X-ray datasets for computer-aided screening of pulmonary diseases. *Quantitative imaging in medicine and surgery*, 4(6): 475.
- Jeong, J.; Zou, Y.; Kim, T.; Zhang, D.; Ravichandran, A.; and Dabeer, O. 2023. Winclip: Zero-/few-shot anomaly classification and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19606–19616.
- Jiao, S.; Wei, Y.; Wang, Y.; Zhao, Y.; and Shi, H. 2023. Learning mask-aware clip representations for zero-shot segmentation. *Advances in Neural Information Processing Systems*, 36: 35631–35653.
- Karim, K. S.; Yin, S.; Nathan, A.; and Rowlands, J. A. 2004. High-dynamic-range pixel architectures for diagnostic medical imaging. In *Medical Imaging 2004: Physics of Medical Imaging*, volume 5368, 657–667. SPIE.
- Kruithof, J. 1937. Telefoonverkeersrekening. *De Ingenieur*, 52: 15–25.
- Lai, H.; Yao, Q.; Jiang, Z.; Wang, R.; He, Z.; Tao, X.; and Zhou, S. K. 2024. CARZero: Cross-Attention Alignment for Radiology Zero-Shot Classification. *arXiv preprint arXiv:2402.17417*.
- Lampert, C. H.; Nickisch, H.; and Harmeling, S. 2013. Attribute-based classification for zero-shot visual object categorization. *IEEE transactions on pattern analysis and machine intelligence*, 36(3): 453–465.
- Liang, V. W.; Zhang, Y.; Kwon, Y.; Yeung, S.; and Zou, J. Y. 2022. Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. *Advances in Neural Information Processing Systems*, 35: 17612–17625.
- Lin, T.; Ho, N.; and Jordan, M. 2019. On efficient optimal transport: An analysis of greedy and accelerated mirror descent algorithms. In *International Conference on Machine Learning*, 3982–3991. PMLR.
- Lin, T.; Ho, N.; and Jordan, M. I. 2022. On the efficiency of entropic regularized algorithms for optimal transport. *Journal of Machine Learning Research*, 23(137): 1–42.
- Liu, J.; Hu, T.; Xiong, H.; Du, J.; Feng, Y.; Wu, J.; Zhou, J.; and Liu, Z. 2024a. VPL: Visual Proxy Learning Framework for Zero-Shot Medical Image Diagnosis. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, 9978–9992.
- Liu, J.; Hu, T.; Zhang, Y.; Feng, Y.; Hao, J.; Lv, J.; and Liu, Z. 2023a. Parameter-Efficient Transfer Learning for Medical Visual Question Answering. *IEEE Transactions on Emerging Topics in Computational Intelligence*.

- Liu, J.; Hu, T.; Zhang, Y.; Gai, X.; FENG, Y.; and Liu, Z. 2023b. A ChatGPT Aided Explainable Framework for Zero-Shot Medical Image Diagnosis. In *ICML 3rd Workshop on Interpretable Machine Learning in Healthcare (IMLH)*.
- Liu, J.; Wang, Y.; Du, J.; Zhou, J.; and Liu, Z. 2024b. Med-CoT: Medical Chain of Thought via Hierarchical Expert. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 17371–17389.
- López-Cifuentes, A.; Escudero-Vinolo, M.; Bescós, J.; and García-Martín, Á. 2020. Semantic-aware scene recognition. *Pattern Recognition*, 102: 107256.
- Mahapatra, D.; Bozorgtabar, B.; and Ge, Z. 2021. Medical image classification using generalized zero shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3344–3353.
- Mahapatra, D.; Ge, Z.; and Reyes, M. 2022. Self-supervised generalized zero shot learning for medical image classification using novel interpretable saliency maps. *IEEE Transactions on Medical Imaging*, 41(9): 2443–2456.
- Menon, S.; and Vondrick, C. 2023. Visual Classification via Description from Large Language Models. In *The Eleventh International Conference on Learning Representations*.
- Novack, Z.; McAuley, J.; Lipton, Z. C.; and Garg, S. 2023. Chils: Zero-shot image classification with hierarchical label sets. In *International Conference on Machine Learning*, 26342–26362. PMLR.
- Porwal, P.; Pachade, S.; Kamble, R.; Kokare, M.; Deshmukh, G.; Sahasrabudhe, V.; and Meriaudeau, F. 2018. Indian diabetic retinopathy image dataset (IDRiD): a database for diabetic retinopathy screening research. *Data*, 3(3): 25.
- Qian, Q.; Xu, Y.; and Hu, J. 2024. Intra-Modal Proxy Learning for Zero-Shot Visual Categorization with CLIP. *Advances in Neural Information Processing Systems*, 36.
- 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*, 8748–8763. PMLR.
- Ren, Z.; Su, Y.; and Liu, X. 2024. ChatGPT-Powered Hierarchical Comparisons for Image Classification. *Advances in Neural Information Processing Systems*, 36.
- Rokhana, R.; Herulambang, W.; and Indraswari, R. 2022. Classification of ocular diseases on fundus images using weighted MobileNetV2. In *2022 International Electronics Symposium (IES)*, 570–575. IEEE.
- Shi, L.; Fan, J.; and Yan, J. 2024. OT-CLIP: Understanding and Generalizing CLIP via Optimal Transport. In *Forty-first International Conference on Machine Learning*.
- Sinkhorn, R. 1964. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The annals of mathematical statistics*, 35(2): 876–879.
- Stephan, F. F. 1942. An iterative method of adjusting sample frequency tables when expected marginal totals are known. *The Annals of Mathematical Statistics*, 13(2): 166–178.
- Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The caltech-ucsd birds-200-2011 dataset.
- Wang, H.; Li, Y.; Yao, H.; and Li, X. 2023. Clipn for zero-shot ood detection: Teaching clip to say no. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1802–1812.
- Wang, Z.; Codella, N.; Chen, Y.-C.; Zhou, L.; Yang, J.; Dai, X.; Xiao, B.; You, H.; Chang, S.-F.; and Yuan, L. 2022a. Clip-td: Clip targeted distillation for vision-language tasks. *arXiv preprint arXiv:2201.05729*.
- Wang, Z.; Wu, Z.; Agarwal, D.; and Sun, J. 2022b. Med-CLIP: Contrastive Learning from Unpaired Medical Images and Text. In *2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022*.
- Xian, Y.; Lorenz, T.; Schiele, B.; and Akata, Z. 2018. Feature generating networks for zero-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5542–5551.
- Yule, G. U. 1912. On the Methods of Measuring Association Between Two Attributes. *Journal of the Royal Statistical Society*, 75(6): 579–652.
- Zhang, H.; Zhou, S.; Li, G. Y.; and Xiu, N. 2022. 0/1 deep neural networks via block coordinate descent. *arXiv preprint arXiv:2206.09379*.
- Zhang, S.; Xu, Y.; Usuyama, N.; Bagga, J.; Tinn, R.; Preston, S.; Rao, R.; Wei, M.; Valluri, N.; Wong, C.; et al. 2023. Large-scale domain-specific pretraining for biomedical vision-language processing. *arXiv preprint arXiv:2303.00915*, 2(3): 6.
- Zhou, Z.; Lei, Y.; Zhang, B.; Liu, L.; and Liu, Y. 2023. Zegclip: Towards adapting clip for zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11175–11185.