

# Boosting Visual-Language Models by Exploiting Hard Pairs

Anonymous authors

Paper under double-blind review

## Abstract

Contrastive Language-Image Pre-training (CLIP) has become the standard for learning cross-modal representations between images and text. Efforts to improve its capabilities typically demand the collection of additional data and retraining with new loss functions. While effective, the added requirements limit their practical use due to the increased resource and time investments needed. In this work, we present HELIP, a cost-effective strategy tailored to enhance the performance of existing CLIP models without the need for training a model from scratch or collecting additional data. Our method allows for effortless integration with existing models’ training pipelines, providing an instant boost by training them with selected challenging text-image pairs from their original training datasets. HELIP treats each text-image pair as a single point in the joint vision-language space, identifying those in close proximity as hard pairs. By incorporating the challenging data, pre-trained CLIP models are refined using both the traditional contrastive loss and the newly introduced hard negative margin loss, ensuring the challenging data is fully utilized. On comprehensive benchmarks, HELIP consistently boosts existing models to achieve leading performance. In particular, it improves the zero-shot classification accuracy on ImageNet for SLIP models pre-trained on CC3M, CC12M and YFCC15M datasets. The improvements are 3.05%, 4.47%, and 10.1% respectively, achieved within two epochs of training. In addition, across fine-grained classification datasets, HELIP improves the zero-shot performance of pre-trained CLIP and SLIP by an average of 8.4% and 18.6%, and their linear probe performance by an average of 9.5% and 3.0%. The code is publicly available at <https://anonymous.4open.science/r/HELIP-7F8E/>.

## 1 Introduction

Contrastive Language-Image Pretraining (CLIP) (Radford et al., 2021) is quickly becoming the standard for foundation models (Awais et al., 2023) due to its effectiveness for a variety of vision-language tasks without task-specific finetuning (Li et al., 2021; Baldrati et al., 2022). However, web-crawled image-text pairs used for the CLIP model pretraining are often loosely connected, resulting in multiple plausible matches beyond the assigned ones (Wu et al., 2022). Several methods have been presented to improve CLIP models by investigating appropriate matches and utilizing widespread supervision among image-text pairs for model training (Li et al., 2022a; 2021; Mu et al., 2022; Radenovic et al., 2023).

Efforts to improve contrastive language-image pretraining models have primarily taken two directions: (1) the addition of objectives to improve the efficacy of supervision (Li et al., 2022a; Mu et al., 2022); and (2) the employment of intra- and inter-modality similarity to select and retrain models using data deemed challenging at the sample level (Li et al., 2021; Radenovic et al., 2023). However, those approaches inevitably require retraining, and those identified as challenging data struggle to bring benefits to model performance. This challenge is partly due to their reliance on finding challenging data within a single batch during training, where truly beneficial challenging data is rare. Additionally, CLIP models’ contrastive loss is not optimally configured to exploit the nuances of difficult data. These limitations significantly restrict the practical application of these enhancements, especially considering the substantial investments already made in pretraining numerous CLIP models (Li et al., 2022a; Mu et al., 2022). This aspect underscores the need for efficient enhancement strategies that do not rely on additional data collection to improve existing pretrained models.

To improve the existing CLIP models, we introduce the HELIP framework, which involves further training the models with challenging data selected from their original training dataset. HELIP defines and identifies the challenging data at the pair level, distinguishing it from traditional methods that compare sample-level similarities between images and texts. Specifically, HELIP treats each text-image pair as a distinct entity within the joint vision-language space, and defines pairs in close proximity as hard pairs. Furthermore, HELIP introduces the **Hard Pair Mining (HPM)** strategy, a novel approach that moves beyond the traditional use of representation spaces learned by CLIP models. Note, the CLIP space is primarily designed for evaluating sample-level similarities—for instance, comparing an image and text (individually, not as a pair)—lacking in evaluating characteristics at the pair level. HPM transforms the task of discovering pairs in close proximity into a solvable proxy task, with the goal of selecting a pair set that optimally supports the target pair’s text-image agreement. HELIP enhances CLIP models not just with the original text-image contrastive loss (Radford et al., 2021), which uniformly pushes all negative samples away from their positive counterpart but also incorporates the **Hard Negative Margin Loss (HNML)** into the loss function. As depicted in Figure 2, HNML imposes an additional geometric structure on the representation space, reflecting the pair-level similarity. Through this approach, HELIP effectively leverages the information within challenging data to boost model performance.

Empirical evidence shows that HELIP improves the performance of existing CLIP models, including pre-trained CLIP, SLIP, and DECLIP, across a variety of benchmarks, such as zero-shot classification, text-image retrieval, and fine-grained linear probing. For zero-shot classification on ImageNet, CIFAR-10, and CIFAR-100, HELIP consistently boosts the performance of all six pre-trained models. Particularly, using HELIP to boost SLIP models pre-trained on CC3M, CC12M, and YFCC15M results in ImageNet zero-shot accuracy gains of 3.05%, 4.47%, and 10.14%, respectively. Further, on seven fine-grained image classification datasets, those pre-trained models achieve better zero-shot and linear probe performance with HELIP. Specifically, the average zero-shot accuracy of CC3M pre-trained CLIP and SLIP are improved by 8.4% and 18.6%. The average linear probe accuracy of CC3M pre-trained CLIP and SLIP are improved by 9.5% and 3.0% respectively. Additionally, the performance gain is also valid in terms of zero-shot retrieval, with 1.1 of R@1 on Flickr30K, and 2.2 of R@1 on COCO for SLIP with HELIP. Our contributions could be summarized as:

- To our best knowledge, our method, HELIP stands out as the first method aimed at improving existing CLIP models in a cost-effective and easily integrable way.
- We introduce the hard pair mining strategy to select challenging data, accompanied by the development of hard negative margin loss. This combination ensures the effective identification and utilization of challenging data, improving the CLIP models.
- Empirical evaluations across zero-shot classification, image-text retrieval, and linear probe benchmarks, consistently show HELIP’s ability to substantially boost the performance of existing CLIP models, underlining its effectiveness and practicality in real-world applications.

## 2 Related work

**Vision-Language pre-training.** Vision Language Pretraining (VLP) is a technique that leverages large-scale image-text datasets to learn a strong joint representation between the two modalities that can be transferred to various downstream vision-language tasks. VLP models can be generally divided into single-stream models and dual-stream models. Dual-stream models (Jia et al., 2021; Li et al., 2022b; Mu et al., 2022; Radford et al., 2021; Yao et al., 2022) typically consist of two separate encoders for image and text respectively and perform cross-modality interactions on the top, are becoming more and more popular because of its flexibility of transferring pre-trained knowledge to downstream tasks. CLIP (Radford et al., 2021), uses a simple contrastive objective to learn visual features from natural language supervision and achieves remarkable zero-shot recognition performance using 400M web-crawled image-text pairs. Recent works boot the performance of CLIP by applying self-supervision within visual modal (Mu et al., 2022), additional nearest neighbor supervision (Li et al., 2022b). These methods are actually doing data augmentations to increase data efficiency and thus bring additional computational costs.

**Contrastive learning with hard negative samples.** Contrastive learning learns a representation of input data that maps semantically comparable examples close together and semantically dissimilar examples far apart (Chen et al., 2020a;b; Wang & Isola, 2020). Recent works include hard negative samples into the loss function and achieve better empirical performance (Cai et al., 2020; Huynh et al., 2022; Kalantidis et al., 2020; Li et al., 2021; Radenovic et al., 2023; Robinson et al., 2021; Shah et al., 2022). For Language-image contrastive learning, current approaches (Li et al., 2021; Radenovic et al., 2023) mine multimodal hard negative examples using intra/inter-modality similarity. Li et al. (2021) choose in-batch hard negative samples with image-text contrastive loss. Hard negative noise contrastive multimodal alignment loss by Radenovic et al. (2023) up-weights the loss term for in-batch hard samples. For previous intra/inter-modality hard sample mining methods, two text-image pairs are considered as hard samples, if the cosine similarity between visual/textual features is high (Li et al., 2021; Radenovic et al., 2023). However, due to the nature of loose assignment for web-crawled image-caption data, a high similarity indicated by intra/inter-modality doesn’t indicate that the two pairs are difficult to tell apart. Contrary to prior works, we design a hard sample mining method to discover similar pairs defined in joint vision-language space and efficiently select samples challenging enough to improve learning.

### 3 Hard pairs for visual-language models

In this section, we first define the notations and revisit CLIP for zero-shot recognition in the preliminary section. Next, we introduce the Hard Pairs Mining strategy (HPM), and the associated Hard Negative Margin Loss (HNML), specifically designed to leverage hard pairs identified by HPM.

#### 3.1 Preliminaries

We consider the task of contrastive image-text pretraining. Given an image-caption dataset  $\mathcal{D} = \{z_i\}_{i=1}^N = \{(x_i^I, x_i^T)\}_{i=1}^N$ ,  $(x_i^I, x_i^T) \in \mathcal{I} \times \mathcal{T}$ , the  $x_i^I$ ,  $x_i^T$  denote the image and its corresponding caption,  $\mathcal{I}$  and  $\mathcal{T}$  indicates visual and textual space respectively, and  $\mathcal{I} \times \mathcal{T}$  indicates the joint Vision-Language space. The goal is to learn a dual encoder model  $\phi = \{\phi_{image}, \phi_{text}\}$ , where  $\phi_{image}$  represents the image encoder and  $\phi_{text}$  denotes the text encoder. We use the shorthand  $I_i = \phi_{image}(x_i^I)$  and  $T_i = \phi_{text}(x_i^T)$  to denote the encoded representation of an image and its caption, respectively. The contrastive objective of CLIP is formulated as,

$$\ell_{CLIP} = -\frac{1}{|B|} \sum_{i \in B} \log \frac{\exp(\text{sim}(I_i, T_i)/\sigma)}{\sum_{j \in B} \exp(\text{sim}(I_i, T_j)/\sigma)}, \quad (1)$$

where  $\text{sim}(\cdot, \cdot)$  is the cosine similarity function,  $B$  is a batch of samples and  $\sigma$  is a trainable parameter controlling the temperature. Intuitively, the above formulation explicitly aligns the representations of image and text from one pair.

#### 3.2 HPM: hard pair mining

In this study, we define *hard pairs* as the pairs that are nearby to a specified target pair within the joint vision-language space,  $\mathcal{I} \times \mathcal{T}$ , which serves as the domain for pair data. Equation 2 depicts the problem of hard pair mining. Here,  $z_i$  represents the target pair,  $\mathcal{H}_i$  denotes a set of pairs chosen from the dataset  $\mathcal{D}_i = \mathcal{D} \setminus z_i$ , and the metric  $\mathbf{S}(\cdot)$  quantifies the similarity between the target pair and a set of pairs,

$$\mathcal{H}_i^* = \arg \max_{\mathcal{H}_i} \mathbf{S}(z_i, \mathcal{H}_i). \quad (2)$$

However, a key challenge arises in defining the similarity metric for pairs,  $\mathbf{S}$ . Existing CLIP methods (Radford et al., 2021; Li et al., 2022b;a) preliminary focus on aligning an image with its caption (Radford et al., 2021; Li et al., 2022a) from a image-text pair. They rarely emphasize on bringing similar pairs closer while distancing the dissimilar ones, which makes current methods fall short in gauging similarity between two pairs. For instance, the cosine similarity between two pairs is ill-defined, within the context of current methods.

**Selecting hard pairs by maximizing pair agreement.** To identify nearby pairs, we introduce the idea of text-image pair agreement maximization. This can be viewed as a proxy task for selecting hard pairs. To

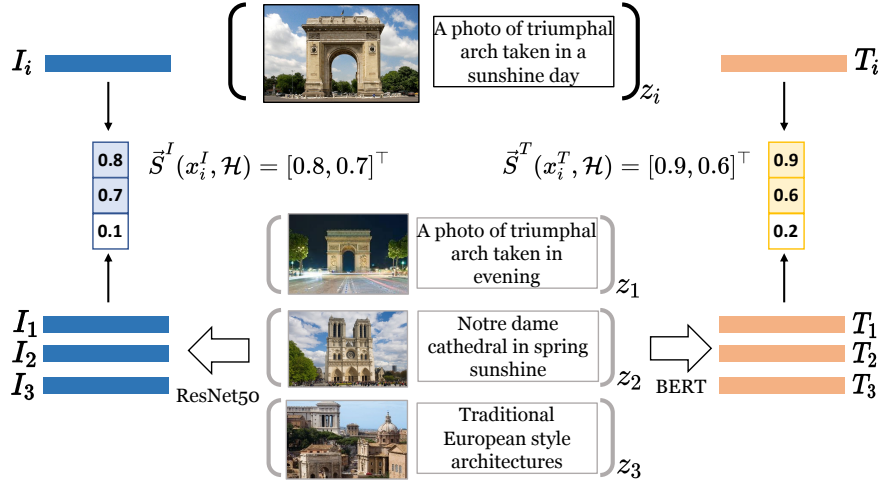


Figure 1: **Hard Pair Mining (HPM)**. Choose hard pairs by optimizing the support set to maximize the agreement prediction of the target pair.

illustrate the rationale for using text-image pair agreement as a proxy for selecting hard pairs, we return to the principle obtained from traditional machine learning methods: the prediction of a model on a test sample is substantially influenced by samples in the training dataset that are similar to the test one. For example, the K-Nearest Neighbors (KNN) algorithm classifies a new instance using the K-closest training examples. The linear regression model predicts the output of a test sample using the weighted sum of the training samples, with higher weights given to samples that are more similar to the test sample. Recent empirical and theoretical studies on model memorization and generalization (Chen et al., 2009; Zhang et al., 2021; Stephenson et al., 2021; Brown et al., 2021) also provide support for this. Intuitively, if a pair agreement prediction model trained on a set of pairs predicts a specific target pair as having a high probability of being a matching pair, the target pair is likely to be similar to the matching pairs on which the model was trained. The challenge of selecting hard pairs is transformed into an optimization task centered on the text-image pair agreement, which is formally represented as:

$$\arg \max_{\mathcal{H}_i} \mathbf{S}(z_i, \mathcal{H}_i) = \arg \max_{\mathcal{H}_i} P_{\mathcal{M}}(z_i | \mathcal{H}_i), \quad (3)$$

where  $P_{\mathcal{M}}(z_i | \mathcal{H}_i)$  denotes the prediction of a pair agreement model,  $\mathcal{M}$ , for the pair  $z_i$  based on a pair set  $\mathcal{H}_i$ . This set is a subset of  $\mathcal{D}_i$ . In this framework, the goal of selecting hard pair is transformed into identifying a training set  $\mathcal{H}_i$  such that the model  $\mathcal{M}$  predicting the target pair as a matching pair.

Designing a suitable pair agreement prediction model for this proxy task is a nontrivial endeavor because the model needs to not only predict the pair matching probability but also allow the optimization of the training set, as indicated in Equation 3. Consequently, a conventional deep neural network design becomes unviable due to the impracticality of retraining across all possible sets  $\mathcal{H}_i$  from  $\mathcal{D}_i$ . Taking inspiration from recent work (Norelli et al., 2022), we propose a data-centric design for the agreement prediction model  $\mathcal{M}$ . As illustrated in Figure 1, the model leverages two pretrained single-modal encoders, i.e.,  $f_{image}$  and  $f_{text}$ , to align representations of images and texts in a unified Vision-Language space. Specifically, the model encodes the target pair  $z_i$  into  $(I_i, T_i)$  using these single-modal encoders. For the visual modality, we determine a similarity vector between the target pair  $z_i$  and the dataset  $\mathcal{D}_i$ . The similarity vector is defined as  $\tilde{S}^I(x_i^I, \mathcal{D}_i) = [\dots, \text{sim}(I_i, I_j), \dots]^T \in \mathbb{R}^{N-1}$ . Here  $I_j = f_{image}(x_j^I)$  with  $(x_j^I, x_j^T)$  being an element of  $\mathcal{D}_i$ , and function  $\text{sim}(\cdot, \cdot)$  denotes the cosine similarity. To counteract noise, values in the vector  $\tilde{S}^I(x_i^I, \mathcal{D}_i)$  are set to zero if  $\text{sim}(I_i, I_j) < \tau$ . This cleaned-up vector is represented as  $\tilde{S}^I$ . The procedure for the textual modality is analogous, producing a vector denoted as  $\tilde{S}^T$ . Note, the representations in this shared space are intuitively interpretable: each dimension corresponding to the visual/textual similarity of the input to a unique pair in the multimodal dataset. This interpretable characteristic enables us to directly optimize the supporting set to maximize the pair matching probability:

$$\mathcal{H}_i^* = \text{Argmax}_{|\mathcal{H}_i|=k} \tilde{S}^I(x_i^I, \mathcal{H}_i)^T \tilde{S}^T(x_i^T, \mathcal{H}_i), \quad (4)$$




Figure 2: **Hard Negative Margin Loss (HNML)**. Hard negative pairs are closer to the positive than the normal negative pairs.




Figure 3: **Further training CLIP with Hard Pairs**. For text-image pairs within a batch, we sample corresponding hard data from the preprocess hard pair set.

where the  $\mathcal{H}_i^*$  is the hard pair set and  $k \in \mathbb{R}^+$  is the number of selected pairs which is much less than  $|\mathcal{D}|$ . The previous problem can be efficiently solved by greedily choosing dimensions that maximize the inner product. Due to the interpretable property, the selected dimensions correspond to the desired pairs.

**Mitigation of noisy data impact.** The prior method assumes the target pair  $z_i$  to be a suitable matching pair. However, in inherently noisy datasets, such as web-crawled ones like LAION (Schuhmann et al., 2022), mismatched pairs might be present. The potential negative effects of hard pairs generated by these mismatched pairs necessitate the development of a strategy for identifying and eliminating them. We create a pair removal strategy based on the availability of hard pairs: A target pair  $z_i$  is deemed as unsuitable and thus removed, if there is a non-empty subset  $\mathcal{H}_i^{sub} \subseteq \mathcal{H}_i^*$  with  $|\mathcal{H}_i^{sub}| > 0$ , such that  $\tilde{S}^I(x_i^I, \mathcal{H}_i^{sub})^\top \tilde{S}^T(x_i^T, \mathcal{H}_i^{sub}) = 0$ . Intuitively, this equation indicates that the number of matching pairs that support  $z_i$  is less than  $k$ . If there is no small subset of size  $k$  to support  $z_i$  as a matching pair in dataset  $\mathcal{D} \setminus z_i$ , it indicates that the target pair is an outlier, possibly due to a mismatch. Such outliers can degrade dataset quality, so they are removed to ensure the reliability of hard data.

**Fast hard pair mining (FastHPM).** It is intuitive to infer that for a dataset collected from a single source, the number of intrinsic hard pairs, which are robust enough to enhance the learned representation, will proportionally increase with the size of the dataset originating from that source. To identify  $k$  (much less than  $|\mathcal{D}|$ ) qualified hard pairs, a portion of the dataset  $\mathcal{D}$  is sufficient. As a result, we present the Fast Hard Pair Mining (FastHPM) approach, which was designed to avoid the time complexity associated with hard pair mining over the entire dataset. FastHPM’s objective can be formalized as follows:

$$\mathcal{H}_i^* \approx \text{Argmax}_{|\mathcal{H}|=k} \tilde{S}^I(x_i^I, \mathcal{H}_i)^\top \tilde{S}^T(x_i^T, \mathcal{H}_i), \quad (5)$$

where  $\mathcal{H}_i \subseteq \overline{\mathcal{D}}_i$  and  $|\overline{\mathcal{D}}_i| = C$  is sampled uniformly from set  $\mathcal{D}_i$ . In this equation, it’s noteworthy that the selection of value  $C$  is solely based on the number of hard pairs  $k$ , instead of the size of  $\mathcal{D}_i$ . Consequently, this optimization reduces the time complexity of FastHPM to  $\mathcal{O}(N)$ . The detailed procedure of the hard pair mining algorithm is consolidated and presented in Appendix A.

### 3.3 HNML: hard negative margin loss

The image-text contrastive loss  $\ell_{CLIP}$ , as illustrated in the preliminary section, aligns the true image-text pairs. But it poses no constraints on the overall geometry among data pairs (Goel et al., 2022). After involving hard data into the finetuning stage, equally maximizing the distance for normal negative pairs and hard negative pairs is an undesired way to utilize the information provided by hard negative pairs. The intuition follows directly from Figure 2. In a desired representation space, the similarity between the

positive and the hard negative,  $\mathbf{S}_1$ , should be greater than the similarity between the positive and those normal negatives,  $\mathbf{S}_2, \mathbf{S}_3$ . Therefore, to impose the additional geometric structure, we introduce the Hard Negative Margin Loss (HNML):

$$\ell_{margin} = \frac{1}{|B|} \sum_{j \in B} \max(0, \text{sim}(I_i, T_j) - \min_{j' \in \mathcal{H}_i^p} \{\text{sim}(I_i, T_{j'})\}), \quad (6)$$

where  $\mathcal{H}_i^p \subseteq \mathcal{H}_i^*$  is the hard negative pairs for the target  $z_i$  involved in one training batch. Note, the HNML is computationally efficient. No extra inner product computation is required. The geometric regularization is applied over the inner product matrix computed in the original CLIP loss, Equation 1. Then, the well-trained model is finetuned with the following loss, where  $\gamma$  is the hyperparameter balancing the two losses,

$$\ell_{finetune} = \ell_{CLIP} + \gamma \ell_{margin}. \quad (7)$$

To boost the performance of well-trained CLIP models without introducing extra data and extra parameters, we introduce the further training strategy which involves the preprocessed hard pairs into the batch composition during training. As shown in Figure 3, for text-image pairs within the batch  $B$ , we randomly sample a subset  $B'$  as seeds. Then, for  $z_i \in B'$ , we randomly select  $|\mathcal{H}_i^p| = p$  pairs from  $\mathcal{H}_i^*$ . The actual

training batch is  $\bar{B} = B \bigcup_{i=0}^{|B'|} \mathcal{H}_i^p$ . We summarize the training pipeline in appendix A.

## 4 Experiments

In the experiments, we conduct a comprehensive empirical investigation to evaluate the efficacy of HELIP in improving zero-shot classification, image-text retrieval, and linear probing performances for existing vision-language models, in Section 4.2. In Sections 4.3 and 4.4, we investigate HELIP’s performance with scaled training data as well as its robustness over noisy datasets. We provide detailed empirical studies on Hard Positive Mining (HPM) and Hard Negative Mining with Margin Loss (HNML) in Sections 4.6 and 4.7.

### 4.1 Experimental setup

**Training datasets.** We used open-source datasets, including Conceptual Captions 3M (CC3M) (Sharma et al., 2018) and Conceptual Captions 12M (CC12M) (Changpinyo et al., 2021), and two 15M subsets of the YFCC100M dataset: v1, collected by Radford et al. (2021), and v2, collected by Li et al. (2022b). The combined datasets of CC3M, CC12M, and YFCC15M v1, which we denote it as Open29M following the term used in prior work (Li et al., 2022b), were not completely obtained due to expired urls. In addition, we independently sampled 7.5M and 8M subsets from the noisier data source, LAION-5B (Schuhmann et al., 2022), labeled as LAION7.5M and LAION8M. These datasets, while smaller than the 400 million pair dataset used in CLIP’s original study (Radford et al., 2021), are well-suited for the data and computational resources we have. Furthermore, they have been widely used in benchmark evaluations for various studies on language-image pretraining, as noted in works by Goel et al. (2022); Li et al. (2022b) and Mu et al. (2022).

**Downstream datasets.** We primarily evaluate the effectiveness of HELIP using zero-shot image classification, linear probing, and zero-shot image-text retrieval. In addition to commonly used ImageNet (Deng et al., 2009), CIFAR10, and CIFAR100 (Krizhevsky et al., 2009), we also verify the performance on 7 fine-grained classification datasets including Caltech101 (Fei-Fei et al., 2004), Food101 (Bossard et al., 2014), Sun397 (Xiao et al., 2010), Flowers102 (Nilsback & Zisserman, 2008), CUB (Wah et al., 2011), Stanford Cars (Krause et al., 2013) and FGVC Aircraft Maji et al. (2013). The zero-shot image-text retrieval task uses MS-COCO (Lin et al., 2014) and Flickr30K (Plummer et al., 2015).

**Implementation details.** Our experiments are conducted across three distinct architectures: ResNet-50, ViT-B/16, and ViT-B/32, tailored to various datasets and pretrained models. Specifically, for loading the pretrained CLIP model on CC3M and CC12M, the ResNet-50 is used as the image encoder. Besides, to align with existing checkpoints established by Mu et al. (2022), we use ViT-B/16 for SLIP model experiments on CC3M and CC12M, respectively. And, we use ViT-B/32 for pretraining on YFCC15M v1, v2, and Open29M

datasets to ensure fair comparison with the results reported in Li et al. (2022b). Furthermore, for the SLIP and DECLIP models, we adapt the pretrained parameters from the publicly available resources\*. The input resolution of the image encoder is  $224 \times 224$  and the maximum context length of the text encoder is 77. All of our experiments are conducted on 8 V100 GPUs with a batch size of 128 for ViT-B/16 models, and a batch size of 512 for ResNet-50 models and ViT-B/32 models. The dimension of the image and text embeddings is 1024 for ResNet-50 models and 512 for ViT-B/16 and ViT-B/32 models. We set  $\tau = 0.5$ ,  $\gamma = 1$  and  $p = 1$  for all the experiments by default. Automatic mixed-precision is used to save GPU memory. To keep the model from overfitting, we use early stopping if there is no performance gain in 5 epochs. Note, to reflect that our method is designed to work with few assumptions on encoder, we used encoders pretrained over a single-modal source rather than multimodally pretrained ones when preparing hard negative pairs. Specifically, we used an unsupervised pre-trained vision transformer, DINO ViTs8 (Caron et al., 2021), and a Sentence Transformer (SentenceT) (Reimers & Gurevych, 2019) to encode text. For DINO ViTs8, the embedding size is 384, while for SentenceT, it is 768.

## 4.2 Main results and discussion

**Zero-shot classification.** We compare the zero-shot performances of the CLIP, SLIP, DECLIP, and those models finetuned by HELIP on CC3M, CC12M, YFCC15M and Open29M. We denote the models finetuned by HELIP as CLIP-HELIP, SLIP-HELIP, and DECLIP-HELIP respectively. Table 1 demonstrates that models fine-tuned by HELIP consistently outperform their counterparts. Specifically, for models pretrained on the CC3M dataset, HELIP boosts the ImageNet zero-shot classification accuracy of the CLIP model from 19.04% to 19.86%. Additionally, on the SLIP model, a performance improvement of over 13% is observed, achieving an accuracy of 26.05%. We additionally include two baseline methods: CYCLIP (Goel et al., 2022) and CLOOB (Fürst et al., 2021) for reference. For CC12M pretraining, we used the SLIP checkpoints released by Mu et al. (2022). On ImageNet, SLIP-HELIP has a 4.47% higher zero-shot accuracy than its counterpart. Due to the lack of openly accessible parameters for DECLIP on the CC3M and CC12M datasets, our analysis focused on comparing DECLIP with DECLIP-HELIP over the YFCC15M v2 dataset. In this context, we present the performance of the SLIP and DECLIP models, as pretrained and released by Li et al. (2022b). The result was obtained by using their evaluation pipeline, denoted with \*. Note, the templates are important for zero-shot tasks. Consequently, for a fair comparison, our analysis and conclusions primarily rely on results obtained from our evaluation pipeline, which is same as the approach used by OpenCLIP. Due to space constraints, we provide more information about the baselines in Appendix B. Notably, both SLIP and DECLIP showed improvements with HELIP, averaging increases of 15.49% and 6.74%, respectively. Further, to demonstrate HELIP’s sustained efficacy across larger datasets, we assessed CLIP and CLIP-HELIP on Open29M. The original CLIP model, upon training with the Open29M dataset, achieves its best performance at the 18th epoch, achieving a zero-shot accuracy of 42.32% on ImageNet. Extending the training duration by two epochs resulted in a marginal decline in accuracy to 42.25%. Notably, our HELIP method instantly boost the performance of the existing CLIP (checkpoint saved at 18th epoch) from 42.32% to 46.33% with just *one additional training epoch*.

**Zero-shot fine-grained classification.** By leveraging hard image-text pairs in contrastive learning, HELIP amplifies the discriminative capability of the CLIP model’s visual embedding. This improvement proves valuable in classification, particularly for fine-grained datasets. Our evaluation on 7 fine-grained classification datasets (Table 2) reveals that SLIP-HELIP boosts the zero-shot accuracy of CC3M and CC12M pretrained SLIP on Caltech101 by 12.88% and 3.95% respectively. Both CLIP and SLIP models witness consistent improvements with their HELIP counterparts.

**Linear probing.** The linear probing task trains a randomly initialized linear classifier on the feature extracted from the frozen image encoder on the downstream dataset. To accomplish this, we train the logistic regression classifier using scikit-learn’s L-BFGS implementation (Pedregosa et al., 2011), with maximum 1,000 iterations on those 7 datasets. For each dataset, we search for the best regularization strength factor on the validation set over 45 logarithmically spaced steps within the range  $1e-6$  to  $1e+5$ . Experimental results in Table 3 demonstrate that both CLIP-HELIP and SLIP-HELIP have consistent improvements over their counterparts on almost all 7 datasets. Note that on CC12M, SLIP-HELIP performs marginally better

\*<https://github.com/facebookresearch/SLIP>, <https://github.com/Sense-GVT/DeCLIP>.

	Method	ImageNet	CIFAR10	CIFAR100
CC3M	CYCLIP (Goel et al., 2022)	22.08	51.45	23.15
	CLOOB (Fürst et al., 2021)	23.97	-	-
	CLIP <sup>†</sup> (Radford et al., 2021)	19.04	33.06	13.77
	CLIP <sup>†</sup> -HELIP	19.86	34.05	14.13
	SLIP (Mu et al., 2022)	23.00	65.61	34.69
	SLIP-HELIP	<b>26.05</b>	<b>68.18</b>	<b>37.77</b>
CC12M	CLIP <sup>†</sup> (Radford et al., 2021)	30.27	51.07	21.94
	CLIP <sup>†</sup> -HELIP	32.05	52.27	24.51
	SLIP (Mu et al., 2022)	41.17	81.30	53.68
	SLIP-HELIP	<b>45.64</b>	<b>82.31</b>	<b>53.79</b>
YFCC15M	SLIP (Mu et al., 2022)	25.29 (34.30*)	60.19	26.80
	SLIP-HELIP	35.43	75.49	47.84
	DECLIP (Li et al., 2022b)	36.05 (43.20*)	78.12	50.60
	DECLIP-HELIP	<b>43.80</b>	<b>84.88</b>	<b>56.31</b>
29M	CLIP <sup>†</sup> (Radford et al., 2021)	42.32	71.98	42.73
	CLIP <sup>†</sup> -HELIP	<b>46.33</b>	<b>77.97</b>	<b>48.33</b>

Table 1: **Zero-shot classification performance on ImageNet, CIFAR10 and CIFAR100.** The <sup>†</sup> indicates baselines pre-trained by us. For all other baselines, publicly available pre-trained parameters were used. Specifically for SLIP and DECLIP on YFCC15M, we report results from two sources: our evaluation using OpenCLIP’s framework with pre-trained parameters released by Li et al. (2022b), and the performance originally reported in Li et al. (2022b), marked with \*.

Dataset	Method	Caltech101	Food101	Sun397	Flowers102	CUB	Stanford Cars	FGVC Aircraft	Average
CC3M	CLIP	42.14	13.02	27.08	13.37	3.45	1.08	1.02	14.45
	CLIP-HELIP	48.08	13.11	28.94	13.61	3.70	1.17	1.11	15.67
	SLIP	54.01	16.03	29.19	12.06	4.70	<b>1.21</b>	<b>1.50</b>	16.96
	SLIP-HELIP	<b>66.89</b>	<b>17.05</b>	<b>33.69</b>	<b>15.16</b>	<b>4.85</b>	1.19	1.29	<b>20.12</b>
CC12M	CLIP	63.78	31.53	37.86	19.56	7.32	14.22	2.49	25.25
	CLIP-HELIP	64.85	36.49	38.22	24.73	8.58	15.59	2.97	27.35
	SLIP	76.33	52.33	44.96	<b>31.81</b>	10.50	22.53	3.06	34.50
	SLIP-HELIP	<b>80.28</b>	<b>54.86</b>	<b>47.53</b>	31.39	<b>10.56</b>	<b>25.67</b>	<b>4.08</b>	<b>36.34</b>

Table 2: **Zero-shot performance on fine-grained image classification.** On a variety of fine-grained classification benchmarks, HELIP consistent boosts the model performance compared to the original versions.

on 5 out of 7 datasets. It’s probably because the self-supervision of SLIP (Mu et al., 2022) within the visual modal can be beneficial for learning fine-grained visual embedding, while SLIP-HELIP doesn’t include image self-supervision during the training. In addition, we did not match the training batch size as SLIP (Mu et al., 2022) because of resource limitations. A combination of HELIP and image self-supervision and larger training batch size may be a potential direction for achieving better linear probe performance.

**Zero-shot retrieval.** We evaluate HELIP on zero-shot image-to-text retrieval tasks on MS-COCO (Lin et al., 2014) and Flickr30K (Plummer et al., 2015). As shown Table 4, both CLIP and SLIP, pre-trained on CC3M and CC12M, consistently improved by HELIP.

### 4.3 Performance of HELIP with Scaled Training Data

To investigate the impact of expanded training dataset sizes on the effectiveness of HELIP, we trained the CLIP model on the YFCC15M dataset. This training yielded a zero-shot classification accuracy of 25.46% on ImageNet. After applying HELIP and one epoch of training, its performance improved to 26.45%. To summarize the zero-shot performance on ImageNet of both the standard CLIP and its enhanced version, CLIP-HELIP, across different data scales, we have illustrated these results in Figure 4. The results show that



Dataset	Method	Caltech101	Food101	Sun397	Flowers102	CUB	Stanford Cars	FGVC Aircraft	Avg.
CC3M	CYCLIP	80.88	54.95	-	83.74	-	22.72	28.02	-
	CLIP	80.11	53.82	56.40	84.07	40.30	22.70	35.61	53.29
	CLIP-HELIP	82.49	59.79	59.56	87.84	46.19	30.01	42.48	58.34
	SLIP	87.96	72.50	66.96	91.91	49.77	39.25	45.87	64.89
	SLIP-HELIP	<b>89.64</b>	<b>73.09</b>	<b>67.67</b>	<b>93.02</b>	<b>53.16</b>	<b>42.44</b>	<b>48.66</b>	<b>66.81</b>
CC12M	CLIP	85.35	68.00	64.45	87.88	48.75	57.80	40.32	64.65
	CLIP-HELIP	85.87	68.89	64.95	88.36	49.41	58.55	40.17	65.17
	SLIP	<b>92.89</b>	83.63	74.34	94.87	<b>60.99</b>	73.43	52.23	76.05
	SLIP-HELIP	92.85	<b>84.25</b>	<b>74.74</b>	<b>95.09</b>	60.53	<b>74.23</b>	<b>52.36</b>	<b>76.29</b>

Table 3: **Linear probe performance on Fine-grained Image Classification.** On average, the linear probe performance of CLIP and SLIP pretrained on CC3M and CC12M are improved.

Pretraining Dataset	Method	COCO		Flickr30K	
		R@1 $\uparrow$	R@5 $\uparrow$	R@1 $\uparrow$	R@5 $\uparrow$
CC3M	CLIP	14.4	34.1	31.7	56.0
	CLIP-HELIP	17.8	39.8	35.4	61.0
	SLIP	22.3	45.6	39.6	68.6
	SLIP-HELIP	<b>23.4</b>	<b>48.3</b>	<b>41.8</b>	<b>69.6</b>
CC12M	CLIP	26.9	52.6	47.2	74.3
	CLIP-HELIP	27.8	54.3	48.2	75.4
	SLIP	39.0	66.0	65.4	<b>90.1</b>
	SLIP-HELIP	<b>39.4</b>	<b>67.2</b>	<b>66.2</b>	89.7

Table 4: **Zero-shot image-text retrieval results on MSCOCO and Flickr.**  $\uparrow$  indicates higher is better. Combining with HELIP, CLIP and SLIP show better performance.

HELIP consistently enhances CLIP’s performance. Most notably, the largest dataset, Open29M, witnessed a remarkable performance increase of 3.06% with HELIP. This result indicates that HELIP can provide immediate performance enhancements for well-trained CLIP models on larger datasets, such as the private 400M dataset mentioned in Radford et al. (2021).

**Top-1 Accuracy comparison between model w/wo HELIP**




Figure 4: **Zero-shot performance on ImageNet for models pre-trained on different dataset sizes.**

#### 4.4 Performance of HELIP on noisy dataset

We expanded our investigation to assess the effectiveness of HELIP on subsets of LAION7.5M and 8M, which are randomly sampled from LAION (Schuhmann et al., 2022). The results are detailed in Table 5. The CLIP model, enhanced with HELIP consistently outperformed its original counterpart on both subsets across a majority of the evaluated datasets, including ImageNet, CIFAR10, CIFAR100, Caltech, and Food. On the 7.5M subset, HELIP enhances performance across all datasets by an average of 3.6%. Although CLIP scores slightly higher on the Sun dataset, HELIP boosts its overall performance with an average improvement of 2.5%

on the 8M subset. These results highlight the enhanced performance achieved through HELIP, demonstrating its robustness and effectiveness in improving existing models that have been pretrained on noisy data.

	ImageNet	CIFAR10	CIFAR100	Caltech	Food	Sun	Avg.
CLIP-7.5M	23.5	34.6	14.5	58.9	28.6	25.3	30.8
CLIP-HELIP-7.5M	<b>25.8</b>	<b>39.9</b>	<b>16.7</b>	<b>61.9</b>	<b>34.1</b>	<b>28.2</b>	<b>34.4</b>
CLIP-8M	25.1	31.1	12.9	60.9	29.5	<b>27.5</b>	31.2
CLIP-HELIP-8M	<b>26.5</b>	<b>38.8</b>	<b>14.6</b>	<b>62.3</b>	<b>33.1</b>	26.6	<b>33.7</b>

Table 5: **Zero-shot performance of CLIP on two LAION subsets.**

#### 4.5 Comparison with other hard data selection method

We evaluate the efficacy of the proposed method in enhancing the discriminative capacity of learned representations by comparing its zero-shot classification performance with that of other hard data mining strategies. As described in the Section 2, a common way to define hard data is through intra-modality similarity. Hence, we introduce the hard data mining methods depending on (sample level) image similarity and text similarity and denote them as IM and TM respectively. For a given target pair, we compute the cosine similarity between its image/text representation and that of the remaining dataset. The image and text representations are encoded using a pretrained Resnet50 and BERT, respectively. As a global preprocessing step, both IM and TM methods mine hard negatives. Subsequently, we integrate the mined hard negative pairs into the training pipeline of the CLIP+IM and CLIP+TM methods and optimize the original contrastive loss to finetune the model. Additionally, we employ the hard negative contrastive loss, HN-NCE, proposed by Radenovic et al. (2023), as a baseline. HN-NCE upsamples the weight of hard-negatives identified by the current model. As shown in Table 6, when the CC3M pretrained CLIP model is combined with HELIP, the performance of our pair-level hard data mining method significantly outperforms other sample-level techniques. In Figure 5, we show the mined hard data that we got from three different preprocessing methods: hard pair mining (HPM), image similarity mining (IM), and text similarity mining (TM). The first row shows the image-text pairs chosen by HPM, while the second and third rows display the pairs chosen by IM and TM, respectively. We observe that the captions of the hard pairs mined with image similarity are only loosely connected with the image of the target pair. For samples mined by TM, their images are even mismatched with the caption of the target pair. The fact that pairs mined by TM are easier than IM is also reflected in Table 6, where the zero-shot performance of the CLIP+IM method consistently outperforms the CLIP+TM method across three datasets.

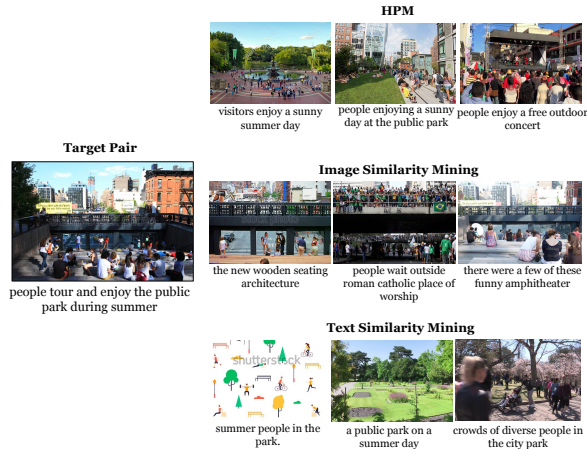


Figure 5: **Hard negative data selected by different methods.** Compared to data mined using the sample-level (image/text modal) similarity, hard pairs mined by HPM are more similar to the target.

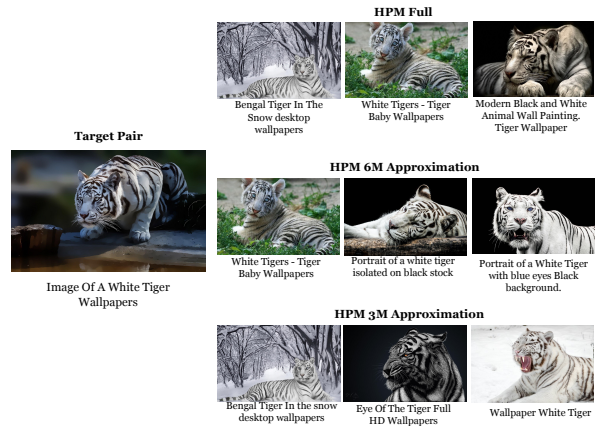


Figure 6: **HPM and fastHPM.** We show the hard pairs mined by HPM and fastHPM. The quality of hard pairs mined by fastHPM is competitive with the pairs mined by HPM.

	Imagenet	CIFAR10	CIFAR100
CLIP + HELIP	<b>19.86</b>	<b>34.05</b>	<b>14.13</b>
CLIP + TM	16.70	28.71	9.67
CLIP + IM	16.93	29.22	10.42
CLIP + HN-NCE	19.47	29.88	11.83

Table 6: **Zero-shot performance of CLIP pre-trained on CC3M boosted by hard data mined by different methods.** HELIP shows superior performance, consistently outperforming local/global hard sample mining techniques by a substantial margin.

#### 4.6 Impact of hard negative margin loss

We investigate the impact of using hard negative margin loss (HNML) on the performance of the SLIP model. In particular, our attention is directed towards an analysis of the SLIP model’s performance, which has been previously pre-trained on the CC3M dataset, when it is both further trained with HPM+HNML and left without HNML. Our approach involves a comparative analysis of the model’s zero-shot classification performance across multiple datasets including ImageNet, CIFAR 100, CIFAR 10, Caltech 101, Food 101, and Sun397. The results of our evaluation are comprehensively detailed in Table 7. These demonstrate that the SLIP model supplemented with HPM and HNML exhibits superior performance, with a performance boost of 4.51 and 3.27 compared to the SLIP and SLIP + HPM models respectively. Interestingly, the model achieved superior performance on the CIFAR 10 dataset without HNML. We postulate that this may be attributed to HNML’s ability to enhance the discriminative power of the learnt representations by employing the class distance as a cost metric. In light of this, our findings suggest that for classification datasets consisting of a larger number of subclasses, employing HNML during the training phase can lead to an increase in classification performance.

	ImageNet	CIFAR10	CIFAR100	Caltech101	Food101	Sun397	Avg.
SLIP	23.00	65.61	34.69	54.01	16.03	29.20	37.09
wo HNML	24.94	<b>69.44</b>	36.35	64.07	16.51	30.91	40.37
w HNML	<b>26.05</b>	68.18	<b>37.77</b>	<b>66.89</b>	<b>17.05</b>	<b>33.68</b>	<b>41.60</b>

Table 7: **SLIP finetuned with and without hard negative margin loss.** When finetuned with hard pairs, the zero-shot performance of CC3M pretrained SLIP can be further enhanced using HNML.

	ImageNet	CIFAR10	CIFAR100	Avg.
CLIP Encoders	19.57	33.28	13.53	22.12
VITs8 + SentenceT	19.86	34.05	14.13	22.68
VITb16 + SentenceT	19.62	35.53	14.67	23.27
VITs8 + T5	19.61	33.99	13.82	22.47

Table 8: **The zero-shot performances of HELIP with different encoders in HPM.** HPM’s performance is insensitive to the selection of encoders.

#### 4.7 Delving into hard pair mining

**Impact of different encoders in HPM.** We explored the effect of different pretrained encoders on HPM’s performance by alternating image and text encoders. Initially, the unsupervised pretrained DINO VITs8 (Caron et al., 2021) was paired with the SentenceT (Reimers & Gurevych, 2019) transformer, trained on over a billion internet-based sentences. This combination was later swapped for the SWAG VITb16 (Singh et al., 2022) and the T5 (Raffel et al., 2020). Additionally, experiments using OpenAI’s CLIP model (Radford et al., 2021) multimodal encoders were conducted. Interestingly, as Table 8 suggests, the encoder choice seemingly has negligible impact on HPM’s performance, likely due to the proficiency of current pretrained models in modeling intra-modal similarities. Moreover, the ability to use single-modal pretrained models and still achieve competitive or superior performance implies that there’s no assumption of having access to a high-quality CLIP model, such as OpenAI’s CLIP-400M.

**Performance Comparison between HPM and FastHPM.** A comparison was made between the zero-shot performances of SLIP models, further trained with hard pairs obtained from both HPM and fastHPM. This comparison, summarized in Table 9, was conducted under three different settings, each maintaining the hyperparameter  $k = 500$ . Additionally, we established subsets  $\tilde{\mathcal{D}}_i$  of sizes 3M and 6M, and accordingly denoted HELIP with these subset sizes as HELIP-3M and HELIP-6M. Table 9 shows that the zero-shot performances of HELIP-3M and HELIP-6M remain competitive with the global HPM hard pair mining approach. These findings suggest that fastHPM offers an efficient strategy for hard pair mining, without compromising performance. Additionally, they hint at fastHPM’s potential to scale up hard pair mining in larger pre-training datasets, a promising direction for future exploration.

	Imagenet	CIFAR10	CIFAR100
SLIP	41.17	81.30	53.68
HELIP- 3M	45.07	<b>82.42</b>	55.22
HELIP- 6M	44.98	81.64	<b>56.62</b>
HELIP- Full	<b>45.64</b>	82.31	53.79

Table 9: **Zero-shot performance for SLIP + HELIP on CC12M with hard samples mined with HPM and fastHPM.** Compared with hard samples mined with HPM, the fast versions are competitive with the full version.

**Visual insights into HPM and FastHPM.** We took the initiative to visualize the hard pairs as identified by the aforementioned three methods. Within Figure 6, the leftmost image-text pairing is earmarked as the target. The pairs in the primary row are those selected via HPM. The subsequent rows, specifically the second and third, present image-text pairings identified by the 6M fastHPM and the 3M fastHPM methods, respectively. Through a comparative visualization, it’s evident that the hard pairs pinpointed by fastHPM bear a significant resemblance to the target pair. For readers keen on delving deeper, we’ve provided an extended set of visualization outcomes in Appendix C.

**Computational time analysis.** Table 10 provides a comparison of the computational time required by HPM and fastHPM. The hard negative pairs preparation times listed were measured on 8 V100 GPUs, with the exception of the \* symbol, which was measured on a single V100 GPU. Given its efficiency and the performance similarities observed in Table 9, fastHPM emerges as a compelling alternative to the full HPM method.

	CC3M	CC12M	YFCC15M
HELIP- 3M	-	2h18min	3h27min
HELIP- 6M	-	5h3min	6h19min
HELIP- Full	1h9min*	9h11min	17h41min

Table 10: **Preparation time for hard pairs.** FastHPM speeds up the hard negative pairs mining process.

## 5 Conclusion

In this study, we delve into boosting pre-trained CLIP models’ performance by more adeptly utilizing their original training dataset. This initiative arose from the recognition of the loosely connected nature of web-crawled image-text pairs, which resulted in suboptimal data utilization due to conventional CLIP loss. Our framework, HELIP, introduces a cost-effective and easily integrable solution for improving existing model performance without extensive retraining or additional datasets. It selects hard pair data from their original training datasets and refines the existing models in a few epochs to immediately boost their performance. Specifically, HELIP treats each text-image pair as a single point in the joint vision-language space, defining those that are close together as hard pairs. The Hard Pair Mining (HPM) strategy effectively identifies challenging hard pairs. The Hard Negative Margin Loss (HNML) was developed to improve existing models by utilizing that hard data. Empirical evaluations across various benchmarks, such as zero-shot classification, image-text retrieval, and linear probing, demonstrate the effectiveness and efficiency of HELIP. We leave the discussion of future work in the appendix. For the discussion of future work, we leave it in the appendix.

## References

- Muhammad Awais, Muzammal Naseer, Salman Khan, Rao Muhammad Anwer, Hisham Cholakkal, Mubarak Shah, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Foundational models defining a new era in vision: A survey and outlook. *arXiv preprint arXiv:2307.13721*, 2023.
- Alberto Baldrati, Marco Bertini, Tiberio Uricchio, and Alberto Del Bimbo. Conditioned and composed image retrieval combining and partially fine-tuning clip-based features. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2022, New Orleans, LA, USA, June 19-20, 2022*, pp. 4955–4964. IEEE, 2022. doi: 10.1109/CVPRW56347.2022.00543. URL <https://doi.org/10.1109/CVPRW56347.2022.00543>.
- Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *European conference on computer vision*, pp. 446–461. Springer, 2014.
- Gavin Brown, Mark Bun, Vitaly Feldman, Adam Smith, and Kunal Talwar. When is memorization of irrelevant training data necessary for high-accuracy learning? In *Proceedings of the 53rd annual ACM SIGACT symposium on theory of computing*, pp. 123–132, 2021.
- Tiffany Tianhui Cai, Jonathan Frankle, David J. Schwab, and Ari S. Morcos. Are all negatives created equal in contrastive instance discrimination? *ArXiv preprint*, 2020.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pp. 9630–9640. IEEE, 2021. doi: 10.1109/ICCV48922.2021.00951. URL <https://doi.org/10.1109/ICCV48922.2021.00951>.
- Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pp. 3558–3568. Computer Vision Foundation / IEEE, 2021. doi: 10.1109/CVPR46437.2021.00356. URL [https://openaccess.thecvf.com/content/CVPR2021/html/Changpinyo\\_Conceptual\\_12M\\_Pushing\\_Web-Scale\\_Image-Text\\_Pre-Training\\_To\\_Recognize\\_Long-Tail\\_Visual\\_CVPR\\_2021\\_paper.html](https://openaccess.thecvf.com/content/CVPR2021/html/Changpinyo_Conceptual_12M_Pushing_Web-Scale_Image-Text_Pre-Training_To_Recognize_Long-Tail_Visual_CVPR_2021_paper.html).
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *Proc. of ICML*, volume 119 of *Proceedings of Machine Learning Research*, pp. 1597–1607. PMLR, 2020a. URL <http://proceedings.mlr.press/v119/chen20j.html>.
- Xinlei Chen, Haoqi Fan, Ross B. Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *ArXiv preprint*, 2020b.
- Yihua Chen, Eric K Garcia, Maya R Gupta, Ali Rahimi, and Luca Cazzanti. Similarity-based classification: Concepts and algorithms. *Journal of Machine Learning Research*, 10(3), 2009.
- Yufeng Cui, Lichen Zhao, Feng Liang, Yangguang Li, and Jing Shao. Democratizing contrastive language-image pre-training: A clip benchmark of data, model, and supervision, 2022.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*, pp. 248–255. IEEE Computer Society, 2009. doi: 10.1109/CVPR.2009.5206848. URL <https://doi.org/10.1109/CVPR.2009.5206848>.
- Li Fei-Fei, Fergus Rob, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *Computer Vision and Pattern Recognition Workshop, 2004. CVPRW’04. Conference on*. IEEE, 2004.
- Andreas Furst, Elisabeth Rumetshofer, Viet Tran, Hubert Ramsauer, Fei Tang, Johannes Lehner, David P. Kreil, Michael Kopp, Günter Klambauer, Angela Bitto-Nemling, and Sepp Hochreiter. CLOOB: modern hopfield networks with infoloob outperform CLIP. *ArXiv preprint*, 2021.

- Samir Yitzhak Gadre, Gabriel Ilharco, Alex Fang, Jonathan Hayase, Georgios Smyrnis, Thao Nguyen, Ryan Marten, Mitchell Wortsman, Dhruva Ghosh, Jieyu Zhang, et al. Datacomp: In search of the next generation of multimodal datasets. *ArXiv preprint*, 2023.
- Shashank Goel, Hritik Bansal, Sumit Bhatia, Ryan A Rossi, Vishwa Vinay, and Aditya Grover. Cyclip: Cyclic contrastive language-image pretraining. *ArXiv preprint*, 2022.
- Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a unified view of parameter-efficient transfer learning. In *Proc. of ICLR*. OpenReview.net, 2022. URL <https://openreview.net/forum?id=0RDcd5Axok>.
- Tri Huynh, Simon Kornblith, Matthew R. Walter, Michael Maire, and Maryam Khademi. Boosting contrastive self-supervised learning with false negative cancellation. In *IEEE/CVF Winter Conference on Applications of Computer Vision, WACV 2022, Waikoloa, HI, USA, January 3-8, 2022*, 2022.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In Marina Meila and Tong Zhang (eds.), *Proc. of ICML*, volume 139 of *Proceedings of Machine Learning Research*, pp. 4904–4916. PMLR, 2021. URL <http://proceedings.mlr.press/v139/jia21b.html>.
- Yannis Kalantidis, Mert Bülent Sariyildiz, Noé Pion, Philippe Weinzaepfel, and Diane Larlus. Hard negative mixing for contrastive learning. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/f7cade80b7cc92b991cf4d2806d6bd78-Abstract.html>.
- Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pp. 554–561, 2013.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Gotmare, Shafiq R. Joty, Caiming Xiong, and Steven Chu-Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 9694–9705, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/505259756244493872b7709a8a01b536-Abstract.html>.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven C. H. Hoi. BLIP: bootstrapping language-image pre-training for unified vision-language understanding and generation. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (eds.), *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pp. 12888–12900. PMLR, 2022a. URL <https://proceedings.mlr.press/v162/li22n.html>.
- Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. In *Proc. of ICLR*. OpenReview.net, 2022b. URL <https://openreview.net/forum?id=zq1iJkNk3uN>.
- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *Proc. of ECCV*, 2014.
- S. Maji, J. Kannala, E. Rahtu, M. Blaschko, and A. Vedaldi. Fine-grained visual classification of aircraft. Technical report, 2013.

- Norman Mu, Alexander Kirillov, David A. Wagner, and Saining Xie. SLIP: self-supervision meets language-image pre-training. In *Proc. of ECCV*, 2022.
- M-E. Nilsback and A. Zisserman. Automated flower classification over a large number of classes. In *Proceedings of the Indian Conference on Computer Vision, Graphics and Image Processing*, Dec 2008.
- Antonio Norelli, Marco Fumero, Valentino Maiorca, Luca Moschella, Emanuele Rodolà, and Francesco Locatello. Asif: Coupled data turns unimodal models to multimodal without training. *ArXiv preprint*, 2022.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 2011.
- Bryan A. Plummer, Liwei Wang, Chris M. Cervantes, Juan C. Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pp. 2641–2649. IEEE Computer Society, 2015. doi: 10.1109/ICCV.2015.303. URL <https://doi.org/10.1109/ICCV.2015.303>.
- Filip Radenovic, Abhimanyu Dubey, Abhishek Kadian, Todor Mihaylov, Simon Vandenhende, Yash Patel, Yi Wen, Vignesh Ramanathan, and Dhruv Mahajan. Filtering, distillation, and hard negatives for vision-language pre-training. *CoRR*, 2023.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In Marina Meila and Tong Zhang (eds.), *Proc. of ICML*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, 2021. URL <http://proceedings.mlr.press/v139/radford21a.html>.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67, 2020. URL <http://jmlr.org/papers/v21/20-074.html>.
- Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proc. of EMNLP*, pp. 3982–3992, Hong Kong, China, 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1410. URL <https://aclanthology.org/D19-1410>.
- Joshua David Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. Contrastive learning with hard negative samples. In *Proc. of ICLR*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=CR1X0QOUTH->.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022.
- Anshul Shah, Suvrit Sra, Rama Chellappa, and Anoop Cherian. Max-margin contrastive learning. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pp. 8220–8230. AAAI Press, 2022. URL <https://ojs.aaai.org/index.php/AAAI/article/view/20796>.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypemymed, image alt-text dataset for automatic image captioning. In *Proc. of ACL*, pp. 2556–2565, Melbourne, Australia, 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1238. URL <https://aclanthology.org/P18-1238>.

- Mannat Singh, Laura Gustafson, Aaron Adcock, Vinicius de Freitas Reis, Bugra Gedik, Raj Prateek Kosaraju, Dhruv Mahajan, Ross Girshick, Piotr Dollár, and Laurens Van Der Maaten. Revisiting weakly supervised pre-training of visual perception models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 804–814, 2022.
- Cory Stephenson, Suchismita Padhy, Abhinav Ganesh, Yue Hui, Hanlin Tang, and SueYeon Chung. On the geometry of generalization and memorization in deep neural networks. In *Proc. of ICLR*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=V8jrrnwGbuc>.
- Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. Technical report, California Institute of Technology, 2011.
- Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *Proc. of ICML*, volume 119 of *Proceedings of Machine Learning Research*, pp. 9929–9939. PMLR, 2020. URL <http://proceedings.mlr.press/v119/wang20k.html>.
- Bichen Wu, Ruizhe Cheng, Peizhao Zhang, Tianren Gao, Joseph E. Gonzalez, and Peter Vajda. Data efficient language-supervised zero-shot recognition with optimal transport distillation. In *Proc. of ICLR*. OpenReview.net, 2022. URL <https://openreview.net/forum?id=G89-1yZLFHk>.
- Jianxiong Xiao, James Hays, Krista A. Ehinger, Aude Oliva, and Antonio Torralba. SUN database: Large-scale scene recognition from abbey to zoo. In *The Twenty-Third IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2010, San Francisco, CA, USA, 13-18 June 2010*, pp. 3485–3492. IEEE Computer Society, 2010. doi: 10.1109/CVPR.2010.5539970. URL <https://doi.org/10.1109/CVPR.2010.5539970>.
- Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. FILIP: fine-grained interactive language-image pre-training. In *Proc. of ICLR*. OpenReview.net, 2022. URL <https://openreview.net/forum?id=cpDhcsEDC2>.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021.



## A Appendix: Algorithm

We summarize the Hard Pair Mining (HPM), the fast Hard Pair Mining (fastHPM) and the training pipeline of HELIP in Algorithm 1, 2 and 3 respectively.

---

### Algorithm 1: Hard Pair Mining (HPM)

---

**Input:** Hard pairs number per sample  $k$   
 Pretrained unimodal vision model:  $f_{text}$   
 Pretrained unimodal vision model:  $f_{image}$   
 Dataset  $\mathcal{D} = \{(x_1^I, x_1^T), (x_2^I, x_2^T), \dots, (x_N^I, x_N^T)\}$   
 Threshold for visual and textual modality  $\tau_I$  and  $\tau_T$   
**Output:** Hard samples  $\mathcal{H} = [\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_N]$

```

for  $i \in [1, N]$  do
   $\mathbf{s} \leftarrow [0, 0, \dots, 0]^\top \in \mathbb{R}^N$ 
   $I_i \leftarrow f_{image}(x_i^I)$ 
   $T_i \leftarrow f_{text}(x_i^T)$ 
  for  $j \in [1, N]$  do
     $I_j \leftarrow f_{image}(x_j^I)$ 
     $T_j \leftarrow f_{text}(x_j^T)$ 
     $\tilde{S}_j^I \leftarrow \frac{I_i \cdot I_j}{\|I_i\|_2 \cdot \|I_j\|_2}$  if  $\frac{I_i \cdot I_j}{\|I_i\|_2 \cdot \|I_j\|_2} > \tau_I$  else 0
     $\tilde{S}_j^T \leftarrow \frac{T_i \cdot T_j}{\|T_i\|_2 \cdot \|T_j\|_2}$  if  $\frac{T_i \cdot T_j}{\|T_i\|_2 \cdot \|T_j\|_2} > \tau_T$  else 0
     $\mathbf{s}_j \leftarrow \tilde{S}_j^I \cdot \tilde{S}_j^T$ 
  end
   $\mathcal{H}_i \leftarrow \text{Argmax}(\mathbf{s}, k)$ 
  if  $\exists j \in \mathcal{H}_i, \mathbf{s}_j = 0$  then
     $\mathcal{H}_i = \emptyset$  # Indicate noise sample
end

```

---

Note, in the inner for loop, shown in Algorithm 1, the image and caption representations will be repeatedly computed. To accelerate the hard pair mining and avoid unnecessary computational overhead, we compute and save the encoded image features and text features. Besides, the outer loop is parallelized in the implementation.

## B Appendix: Discussion about baselines

In our experiments, we utilized CLIP, SLIP, and DECLIP as baseline models on CC3M, CC12M, YFCC15M, and Open29M datasets. To ensure our results are both compelling and reproducible, we primarily employed publicly available checkpoints as our baseline and rigorously tested the effectiveness of HELIP against these checkpoints. On CC3M, the checkpoint of SLIP model is released<sup>†</sup>. We enhanced its performance by applying HELIP which notably improved the zero-shot performance on ImageNet from 23.00 to 26.05. However, we noticed that the CLIP with ResNet50 on CC3M is missing. To address this, we undertook the pretraining ourselves. Our results were encouraging: the performance of our pretrained CLIP with ResNet50 achieved a score of 19.86, surpassing the 17.10 achieved by SLIP’s CLIP with ViT-B/32 as reported in Mu et al. (2022). This outcome suggests the robustness of our implementation. Besides, consistent with several prior studies, we found that on smaller pretraining datasets, CLIP with ResNet50 outperforms CLIP with ViT-B. On the CC12M dataset, a similar situation arose: while the SLIP checkpoint was available, the CLIP model was absent, leading us to undertake its pretraining. On the YFCC15M (v1) collected by Radford et al. (2021), we trained the CLIP model. This resulted in a 25.46 score in the ImageNet zero-shot classification, closely aligning with the 26.10 outcome reported by Cui et al. (2022). Additionally, for the YFCC15M (v2) dataset

<sup>†</sup><https://github.com/facebookresearch/SLIP#results-and-pre-trained-models>

**Algorithm 2:** fast Hard Pair Mining (fastHPM)**Input:** Hard pairs number per sample  $k$ Pretrained unimodal vision model:  $f_{text}$ Pretrained unimodal vision model:  $f_{image}$ Dataset  $\mathcal{D} = \{(x_1^I, x_1^T), (x_2^I, x_2^T), \dots, (x_N^I, x_N^T)\}$ Threshold for visual and textual modality  $\tau_I$  and  $\tau_T$ Candidate pool size  $C$ **Output:** Hard samples  $\mathcal{H} = [\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_N]$ **for**  $i \in [1, N]$  **do**    Uniformly  $C$  samples from Dataset  $\mathcal{D}$ ,  $\bar{\mathcal{D}}_i = \{(x_1^I, x_1^T), (x_2^I, x_2^T), \dots, (x_C^I, x_C^T)\}$      $\mathbf{s} \leftarrow [0, 0, \dots, 0]^\top \in \mathbb{R}^N$      $I_i \leftarrow f_{image}(x_i^I)$      $T_i \leftarrow f_{text}(x_i^T)$     **for**  $j \in [1, C]$  **do**         $I_j \leftarrow f_{image}(x_j^I)$          $T_j \leftarrow f_{text}(x_j^T)$          $\bar{S}_j^I \leftarrow \frac{I_i \cdot I_j}{\|I_i\|_2 \cdot \|I_j\|_2}$  **if**  $\frac{I_i \cdot I_j}{\|I_i\|_2 \cdot \|I_j\|_2} > \tau_I$  **else** 0         $\bar{S}_j^T \leftarrow \frac{T_i \cdot T_j}{\|T_i\|_2 \cdot \|T_j\|_2}$  **if**  $\frac{T_i \cdot T_j}{\|T_i\|_2 \cdot \|T_j\|_2} > \tau_T$  **else** 0         $\mathbf{s}_j \leftarrow \bar{S}_j^I \cdot \bar{S}_j^T$     **end**     $\mathcal{H}_i \leftarrow \text{Argmax}(\mathbf{s}, k)$     **if**  $\exists j \in \mathcal{H}_i, \mathbf{s}_j = 0$  **then**         $\mathcal{H}_i = \emptyset$    # Indicate noise sample**end**

referenced in Li et al. (2022b), both SLIP and DECLIP pretrained parameters were made available by Li et al. (2022b), which we utilized directly as our baselines. On the larger dataset, Open29M, there was a lack of open-source pretrained checkpoints, prompting us to conduct the pretraining ourselves. Notably, the performance of our reimplementation (42.32) closely aligns with the results reported by Li et al. (2022b), indicating the effectiveness of our approach.

## C Appendix: More visualization results

We offer further visualization results pertaining to the hard samples mined by various methods. As depicted in Figure 7, the hard samples sourced by HPM closely resemble the target sample (seen at the top left). Interestingly, for samples with fewer objectives, the image and text mining method can identify a reasonably challenging counterpart, as seen in the case of “the harbor in a small village”. However, for intricate scenes, only the HPM is capable of yielding sufficiently challenging samples, like the scenario “people touring and enjoying the public park during summer”. The dataset acquired from the web encompasses a myriad of such intricate cases. We posit that this is why training with hard samples unearthed by HPM yields more proficient outcomes.

Moreover, we present additional visualization results for hard samples mined via different techniques. Hard samples extracted by HPM exhibit a stronger resemblance to the target sample, as highlighted in Figure 7 (top left). We observed that the image and text mining methods can provide a relatively fitting hard counterpart for simpler samples, like “the harbor in a quaint settlement”. However, for more intricate scenes, only the HPM method produces samples of adequate difficulty, such as “people touring and relishing the public park throughout summer”. The web-based dataset includes a significant proportion of these complex cases. Consequently, we infer that training with hard samples mined by HPM results in enhanced performance.

**Algorithm 3:** Hard sampleE for boosting contrastive Language-Image Pretrained models (HELIP)

---

**Input:**  $\mathcal{D} = \{(x_1^I, x_1^T), (x_1^I, x_1^T), \dots, (x_N^I, x_N^T)\}$   
Hard Pair Mining algorithm,  $\text{HPM}()$  # or the fastHPM()  
Pretrained unimodal vision model:  $f_{\text{text}}$   
Pretrained unimodal vision model:  $f_{\text{image}}$   
Pretrained contrastive language-image model  $\{\phi_{\text{image}}, \phi_{\text{text}}\}$   
hyperparameters:  
Hard pairs number  $k$   
Hard negative margin strength  $\gamma$   
Sampled hard negatives number  $p$   
Learning ratio  $\eta$   
Batch size  $b$   
Training iteration number  $E$   
Visual and textual modality threshold  $\tau_I$  and  $\tau_T$   
**Output:** CLIP model  $\{\phi_{\text{image}}, \phi_{\text{text}}\}$

---

```

 $\mathcal{H} \leftarrow \text{HPM}(\mathcal{D}, f_{\text{text}}, f_{\text{image}}, k, \tau_I, \tau_T)$ 
for  $iter \in [1, E]$  do
     $B \leftarrow \{z_1, \dots, z_b\} \stackrel{\text{i.i.d.}}{\sim} \text{Uniform}(\mathcal{D})$ 
    for  $z_i \in B$  do
         $\mathcal{H}_i^p \leftarrow \{z_i, \dots, z_p\} \stackrel{\text{i.i.d.}}{\sim} \text{Uniform}(\mathcal{H}_i)$ 
         $\bar{B} \leftarrow B \cup \mathcal{H}_i^p$ 
    end
    Compute loss  $\ell_{\text{finetune}}$ , Equation (6), with samples  $\bar{B}$ 
     $\phi_{\text{image}} \leftarrow \phi_{\text{image}} + \eta \cdot \partial_{\phi_{\text{image}}} \ell_{\text{finetune}}$ 
     $\phi_{\text{text}} \leftarrow \phi_{\text{text}} + \eta \cdot \partial_{\phi_{\text{text}}} \ell_{\text{finetune}}$ 
end

```

---

**D Appendix: Future work**

Moving forward, several possibilities for future research emerge. First, we aim to explore composition-aware fine-tuning for VLMs, which could potentially enable more effective utilization of multimodal information. Moreover, we are intrigued by the prospect of combining parameter-efficient tuning (He et al., 2022) with HELIP potentially further enhancing performance. Another area of interest is scaling up the dataset size and examining the applicability of the scaling law to our method. We also intend to investigate how the integration of our boosting algorithm might alter the multimodal dataset curation algorithm (Gadre et al., 2023). Ultimately, we hope our work will serve as a catalyst for additional research in the fine-tuning of pre-trained, large-scale multimodal models.

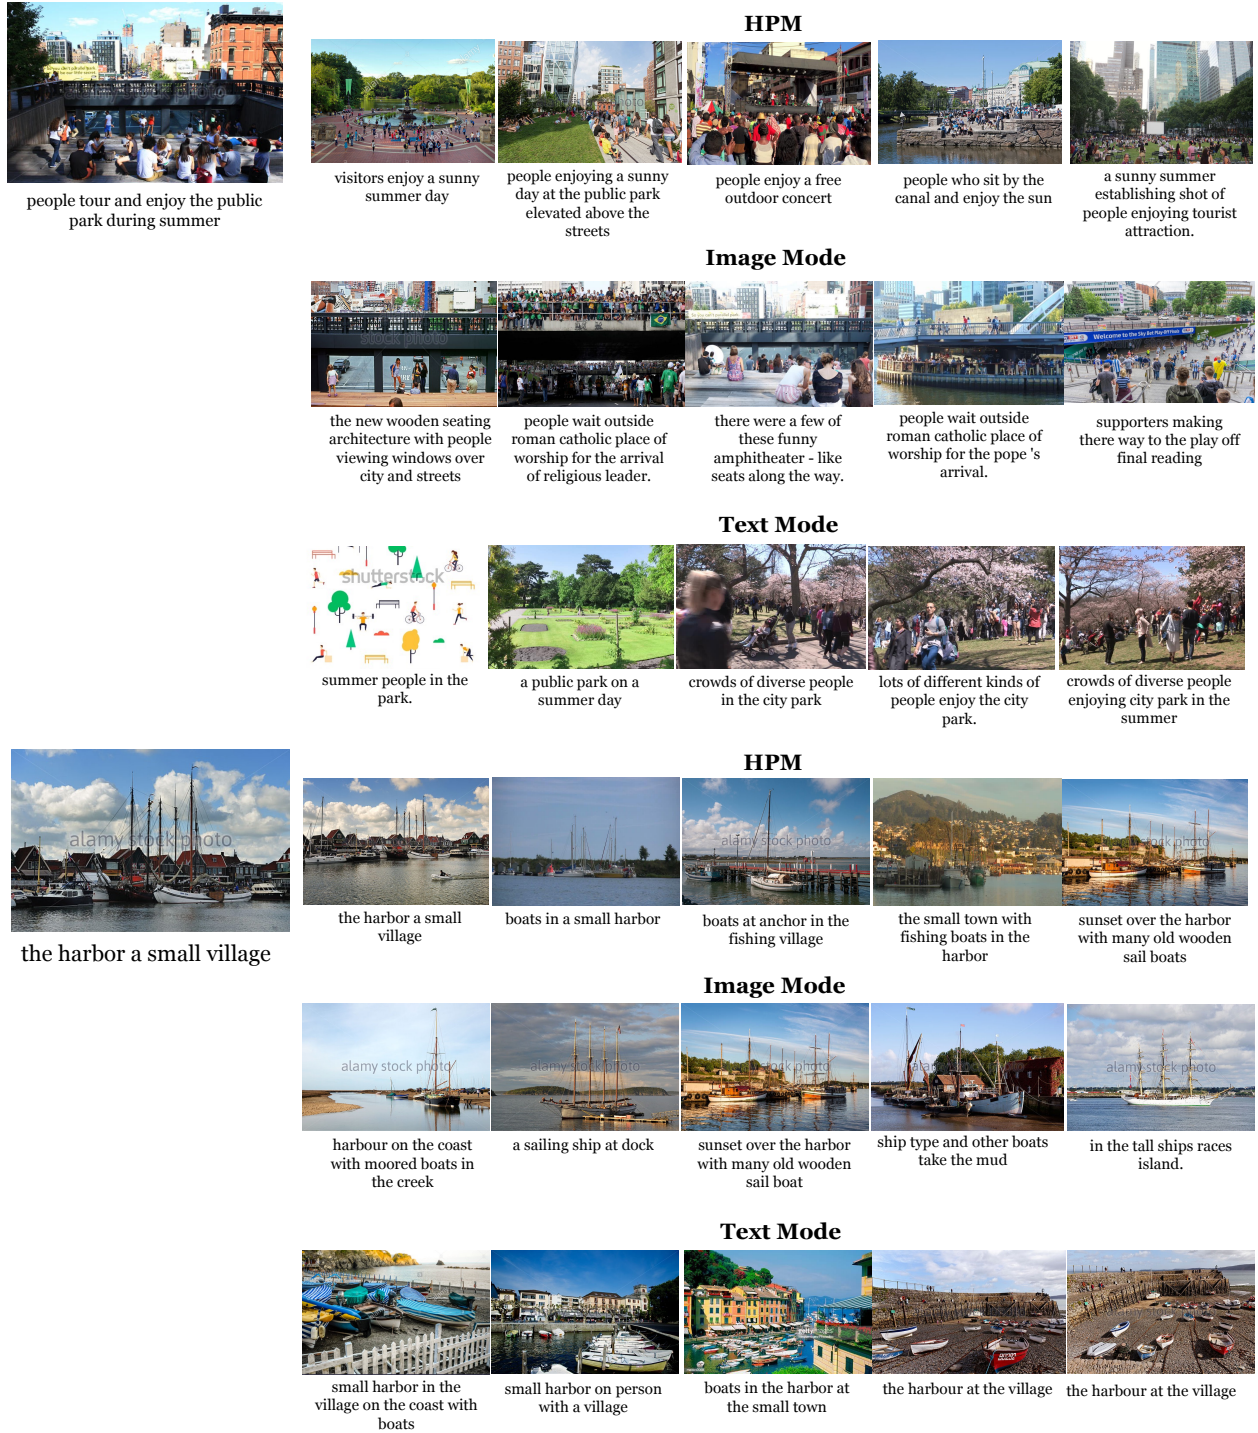


Figure 7: Hard pairs selected by different methods.