

VeCLIP: Improving CLIP Training via Visual-enriched Captions

Zhengfeng Lai^{1,2}, Haotian Zhang^{2*}, Bowen Zhang², Wentao Wu²,
Haoping Bai², Aleksei Timofeev², Xianzhi Du², Zhe Gan², Jiulong Shan²,
Chen-Nee Chuah¹, Yinfei Yang^{2**}, and Meng Cao²

¹ University of California, Davis

² Apple AI/ML

{lzhengfeng, chuah}@ucdavis.edu

{haotian_zhang2, yinfeiy, mengcao}@apple.com

Abstract. Large-scale web-crawled datasets are fundamental for the success of pre-training vision-language models, such as CLIP. However, the inherent noise and potential irrelevance of web-crawled AltTexts pose challenges in achieving precise image-text alignment. Existing methods utilizing large language models (LLMs) for caption rewriting have shown promise on small, curated datasets like CC3M and CC12M. This study introduces a scalable pipeline for noisy caption rewriting. Unlike recent LLM rewriting techniques, we emphasize the incorporation of visual concepts into captions, termed as **Visual-enriched Captions (VeCap)**. To ensure data diversity, we propose a novel mixed training scheme that optimizes the utilization of AltTexts alongside newly generated VeCap. We showcase the adaptation of this method for training CLIP on large-scale web-crawled datasets, termed VeCLIP. Employing this cost-effective pipeline, we effortlessly scale our dataset up to 300 million samples named VeCap dataset. Our results show significant advantages in image-text alignment and overall model performance. For example, VeCLIP achieves up to **+25.2%** gain in COCO and Flickr30k retrieval tasks under the 12M setting. For data efficiency, VeCLIP achieves **+3%** gain while only using **14%** of the data employed in the vanilla CLIP and **11%** in ALIGN. We also note the VeCap data is complementary with other well curated datasets good for zero-shot classification tasks. When combining VeCap and DFN, our model can achieve strong performance on both of image-text retrieval and zero-shot classification tasks, *e.g.*, **83.1%** accuracy@1 on ImageNet zero-shot for a H/14 model. We release our codes and model weights at <https://github.com/apple/ml-veclip>.

1 Introduction

Large-scale vision-language representation learning, exemplified by CLIP [33], has gained wide attention due to the transferability of knowledge learned from image-text pairs to diverse downstream tasks such as zero-shot image classification

* Equal contribution.

** Corresponding author.



Fig. 1: Noisy web-crawled data and the limitation of LLM rewrite. AltTexts can be noisy and uninformative; it may not describe all visual objects present in the image. Simple LLM rewrite [13] on such raw and noisy captions cannot introduce new image-relevant information. After applying our proposed VeCap, new captions are enriched with more image-specific concepts. We keep all image-text pairs for pre-training rather than filtering out those with noisy AltTexts, as images of rich visual objects still contribute effectively to the training process.

and image-text retrieval [18, 20, 21]. CLIP training is straightforward via the image-text contrastive loss, but involves a large-scale dataset of 400 million image-text pairs crawled from the Web. Consequently, CLIP embeddings lead to consistent improvement across various downstream tasks compared to other vision pre-training methods such as SimCLR [7] and MAE [16]. CLIP achieves success via two scalable paradigms: data and computational resources. First, the massive web-crawled data [36, 37] enable the training to be scalable and meet the requirements of data-hungry backbones (*e.g.*, ViT [12]). Second, the simple contrastive loss grants favorable scaling properties to the computational resources.

Despite the availability of large-scale web-crawled data, their quality can be low or noisy. For example, AltTexts suffer from two major issues: 1) they can be noisy, uninformative, or irrelevant to the images; 2) they may not describe all visual contents in the image. For example, as shown in Figure 1, in the first image, we observe a house with a white roof and a porch. However, the corresponding caption only describes the address, which proves overly abstract for effective vision-language alignment in training. Our observations demonstrate that caption quality plays a pivotal role in CLIP’s performance, as detailed in Table 6b and the Appendix (*e.g.*, CC3M vs. our web-crawled 3M). It is worth noting that the captions in CC3M are derived from human annotations, which may require heavy resources when further scaling up. This motivates the main open research question addressed in this work: Can we devise a **scalable and cost-effective**

pipeline to improve captions within these noisy datasets at scale (e.g., up to million or billion level)?

One natural direction is to deploy Large Language Models (LLMs) to rephrase captions [13]. However, the major limitation of such methods lies in the inability of LLMs to generate and introduce new image-specific details. LLMs can only modify sentence syntax in this scenario. For example, we follow a recent work [13] and use LLM to rewrite the raw captions from the Web: as shown in Fig. 1, LLM rewrite cannot introduce any new information and thus the new caption remains similar to AltText. In other words, if the original AltTexts are noisy, the benefits brought by LLM rewrite might yield only trivial improvements. In essence, the reliance on high-quality captions within pre-training datasets limits the effectiveness of simple LLM rewrite. However, sourcing such high-quality datasets like manually curated CC3M and CC12M [5] remains challenging, and further scaling up to larger datasets becomes both time-consuming and labor-intensive to meet the prerequisites for CLIP pre-training. Therefore, in this work, we focus on building a scalable and cost-effective pipeline tailored to raw and noisy web-crawled data to improve CLIP.

In addition to data quality, the diversity of data significantly impacts VLM pre-training [2, 28]. Methods relying on LLM-based rewriting may diminish data variety, given that LLMs tend to apply a uniform style in their sentence rephrasing. Moreover, existing works mainly focus on image augmentations, while texts are disregarded and unaltered during training without augmentation [13]. This may also incur overfitting issues as the text encoders repeatedly encounter the same texts in each epoch. Since these techniques have exclusively undergone assessment on meticulously curated datasets like CC3M and CC12M [5], their suitability for extensive, uncensored web-crawled data remains uncertain. Consequently, there is a pressing need to build a scalable approach to enhance data quality, diversity, and training methodologies to improve pre-training for VLMs on both model performance and data efficiency.

Concurrently, alongside the evolution of CLIP, there has been substantial progress in the development of instruction fine-tuned LLMs. These models and their multimodal extensions have demonstrated outstanding performance, surpassing human capabilities in various natural language and vision tasks. Inspired by these models, we investigate the potential of utilizing them to improve the noisy captions gathered from the Internet. Specifically, we initially employ LLaVA [25], a Language-Vision Assistant, to leverage visual concepts extracted from the images. Given that AltTexts may lack informativeness, our objective is to integrate the newly derived visual concepts into the caption. However, it is worth noting that LLaVA [25] fine-tuned its language decoder on its own generated dataset, potentially losing its ability to accommodate comprehensive instructions. Consequently, we further propose to utilize an LLM to refine the sentence by fusing the generated caption from LLaVA and the original AltText. This process aims to maximize image-specific information for optimal vision-language alignment. We denote the caption generated from LLM as LLM Visual-enriched **Captions** (VeCap), or **VeCap** for short. For data variety,

we propose VeCLIP and introduce a mixed training scheme, alternating between VeCap and the original AltText. This strategy ensures that the model captures all pertinent information without oversight. We generalize this scalable pipeline to curate five pre-training datasets ranging from small-scale to large-scale up to 300M. Overall, our contributions are summarized below:

- We present a visual-enriched re-captioning technique for CLIP training. This marks the initial endeavor to leverage visual concepts extracted from images and inject them into the captioning process.
- Our pipeline is cost-effective and capable of processing data at a scale exceeding 300M, named VeCap. Then, we propose VeCLIP with a mixed training scheme that uses VeCap to improve CLIP training on model performance.
- VeCLIP can achieve up to **25.2%** improvement over CLIP in retrieval tasks. For training data efficiency, *e.g.*, we use only **5%** data in training but achieve competitive results in image-text retrieval tasks.
- VeCap data is also complementary with other well curated datasets. A CLIP-H/14 model trained on the combination of VeCap and DFN achieves strong performance on both of image-text retrieval and zero-shot classification tasks, with an impressive **83.1%** zero-shot accuracy@1 on ImageNet.

2 Related Work

Contrastive language-image pre-training. CLIP [33] has shown its effectiveness in acquiring transferable image representations via text supervision after large-scale pre-training. Similar models such as ALIGN [18], Florence [45], BASIC [31] and OpenCLIP [9] have shown impressive zero-shot image classification and image-text retrieval capabilities. SLIP [27] and DeCLIP [22] incorporate self-supervised training techniques to improve performance. CoCa [44] introduces an additional decoder alongside the contrastive loss. LiT [48] proposes to keep a pre-trained image encoder frozen and fine-tune text encoders to improve the zero-shot transferability. Nevertheless, the majority of these subsequent studies incorporate supplementary training inputs and losses, potentially exerting adverse effects on both training efficiency and memory usage.

Improving image-text datasets. Given the importance of the pre-training data, many works focus on improving the datasets, such as filtering less informative image-text pairs [1, 4, 13, 26]. However, these methods may disregard a large amount of data even though some images have rich visual concepts. An alternative approach is to rewrite the caption to enhance the alignment between texts and images. For example, LaCLIP [13] employs LLMs to perform rewriting. Nevertheless, their evaluation was conducted on small-scale and meticulously curated datasets like CC3M and CC12M [5], where the initial captions were already of high quality. ShareGPT4V [6] uses GPT4-Vision for obtaining high-quality captions, which may not be scalable considering the cost of GPT4-Vision.

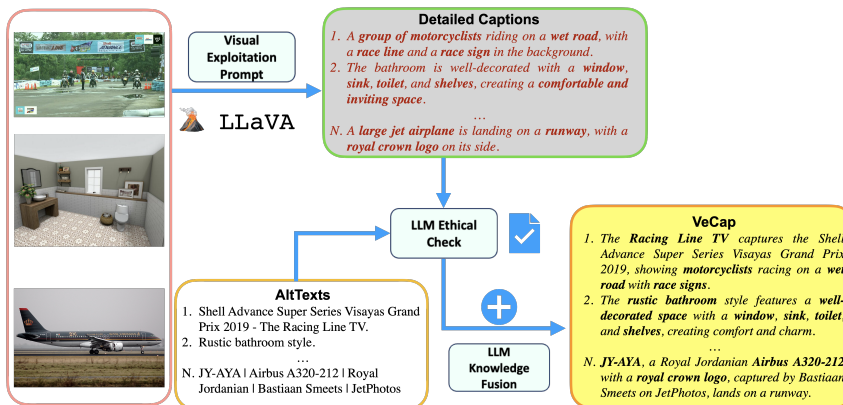


Fig. 2: An overview of the scalable VeCap recaptioning pipeline. First, we focus on exploiting visual concepts in images via leveraging a multimodal LLM (LLaVA) to describe the image with a designed prompt independent of AltText to generate Visual-enriched Captions (VeC). Second, we leverage an LLM to fuse the concepts from both AltText and VeC to generate the final caption, denoted as VeCap.

3 Methodology

3.1 Preliminary

CLIP. The Contrastive Language-Image Pre-training (CLIP) method has shown its effectiveness in training vision models via language supervision. Specifically, a batch of N image-text pairs $\{x_I, x_T\}$ is sampled from the massive training data during each training iteration. We apply data augmentations to the images before inputting them into the vision encoder. We denote f_I and f_T as the normalized features extracted by the vision and text encoders, respectively. We use the contrastive loss to train the model, where the paired images and texts are treated as positive pairs and the remaining as negative samples. The training loss iterating over images can be formulated as follows:

$$L_I = - \sum_{i=1}^N \log \frac{\exp(\text{sim}(f_I(\text{aug}(x_I^i)), f_T(x_T^i))/\tau)}{\sum_{k=1}^N \exp(\text{sim}(f_I(\text{aug}(x_I^i)), f_T(x_T^k))/\tau)}, \quad (1)$$

where (x_I^i, x_T^i) is the i^{th} image-text pair in the batch, and $\text{aug}(\cdot)$ refers to image augmentations. $\text{sim}(\cdot, \cdot)$ is the similarity measurement function. We set τ as a learnable temperature parameter that scales the logits in experiments. The loss iterating over texts is symmetrical and denoted as L_T . Finally, the training loss is $L = (L_I + L_T)/2$.

3.2 Recaptioning with Visual Concept Exploitation

Web-crawled captions (AltTexts) can be noisy and uninformative to the images. LaCLIP [13] used LLM to rewrite the caption. As shown in Fig. 1, this may not

be applicable if the captions are noisy as LLM can only reconstruct the sentence but cannot introduce new information without any information provided by the image. Given the inherent noise in AltTexts, we advocate for the utilization of pre-trained multimodal models to generate augmented captions with richer visual concepts derived from the images. In this subsection, we use LLaVA [25] as one example and present a scalable and cost-effective pipeline for scaling up.

LLaVA and image captioning for Visual-enriched Captions (VeCap).

As a multimodal model, LLaVA connects the open-set visual encoder of CLIP [33] with an LLM, such as LLaMA [39], then fine-tune them on a visual instruction-tuning dataset. LLaVA shows its effectiveness in leveraging the capabilities of pre-trained LLM and vision foundation models. Given an input image x_I , we get f_I from CLIP’s vision encoder. Then, LLaVA applies a trainable projection matrix \mathbf{W} to convert f_I into language embedding tokens to achieve the image-language alignment. To mitigate the influence of AltText, we have devised AltText-independent prompts tailored for LLaVA, ensuring the full exploitation of visual concepts. We refrain from incorporating AltText information into LLaVA, while acknowledging the potential loss of pre-trained knowledge during fine-tuning of the LLM component on the generated dataset. This trade-off, however, may limit its capacity to comprehend more intricate instructions. Thus, we adopt a straightforward yet potent prompt, “Describe the image concisely, less than 20 words”, allowing LLaVA to generate visual concepts directly from the image autonomously. We denote this captions generated by LLaVA as x_{T_v} . Subsequently, the image-text pair is converted as (x_I, x_{T_v}) .

3.3 Scalable LLM Rewrite for Concept Fusion

Given the limited language capacity of LLaVA, we only use LLaVA to extract all possible visual clues. Then, we employ LLMs to refine the caption by fusing both the knowledge from AltText x_T and the novel visual concepts from x_{T_v} . This step has three main advantages: 1) It ensures the retention of information delineated in AltText, thereby amplifying the informativeness of the caption; 2) It can serve as a form of “strong augmentation” in textual data, characterized by a profound restructuring of sentence syntax instead of focusing on word-level modifications used in existing language augmentation techniques [38, 41]; 3) It can mitigate the “hallucination” issue arising from large vision-language models (e.g., LLaVA) to ensure that the entity described in the ultimate caption is present in the image.

Generating rewrites for a vast corpus of texts using closed-source models like ChatGPT or Bard is impractical, considering the substantial financial costs and time incurred through API utilization. Therefore, to facilitate the rewriting tasks on a large-scale dataset, we turn to open-source state-of-the-art LLMs. Due to the license issue, we select Vicuna-1.1 [49], renowned for its robust performance in text completion tasks, as one example of LLM rewriting in this study. We formulate a context input as the following three components. First, we include a sentence designed to apprise the LLM of the task, specifically, rewriting and fusing two attached sentences. This serves as an initial contextual cue to orient the LLM towards comprehending the overarching objective. Second, we impose

several constraints on the ultimate output. For instance, our goal is to position attributes prior to noun entities, all while refraining from introducing any novel semantic interpretations. Furthermore, it is essential that the sentence refrains from commencing with the phrase “The image” and instead directly expounds upon all-encompassed concepts. Finally, the last part of the context includes two sentences (x_v and x_{Tv}) that require fusing and rewriting, followed by the separation symbol. This ensures that the LLM is furnished with the specific texts to be fused and rewritten as part of its context input. By integrating these three components, we establish an all-encompassing context that steers the LLM towards proficiently crafting diverse and knowledge-fused text rewrites.

We conduct this process in a batch-inference process instead of a single-item inference as shown in Fig 2: we group our data into batches and implement a batch-inference process to achieve up to 64 times faster on Nvidia A100. Specifically, we use Vicuna-1.1-13B model to generate the final output as x_{TI} . The final prompt is as follows: [*Rephrase the following two sentences into one short sentence while adhering to the provided instructions: Place attributes before noun entities without introducing new meaning. Do not start with “The image”. + 1. AltText; 2. model generated caption.*] We denote the caption from LLM as LLM-VeCap, or VeCap for short.

Potential ethics of LLM and failure cases processing. While upscaling the LLM rewriting process, we identify two scenarios in which LLM encounters difficulties in executing the designated task: 1) *Ethical Concerns*. If the AltText contains content either illegal or violent, LLM may reply, “I am sorry that I cannot...”; 2) *Length Constraint*. In cases where the AltText exceeds an optimal length, the processing time of the LLM may be significantly prolonged, thus impeding large-scale rewriting. To address the first scenario, we use the model generation captions as the only input to be rewritten via LLM to form VeCap, thereby preemptively excluding potentially unlawful or aggressive content. In the second scenario, we mitigate this issue by preserving the generated caption but truncating the AltText to conform to the maximum allowable length, thus we have more visual concepts aligned with the image.

3.4 VeCLIP: Mixed Training Scheme with Visual-enriched Captions for CLIP

As LLM rewriting may introduce a consistent style, there could be a decline in data diversity for large-scale pre-training, even if data quality is enhanced. To enhance data diversity, we propose a mixed training scheme to serve as additional text augmentations applied in pre-training: $\text{mix}(x_t) \sim \text{Uniform}([x_T, x_{TI}])$. Then, the training loss iterating over the images becomes:

$$L_I = - \sum_{i=1}^N \log \frac{\exp(\text{sim}(f_I(\text{aug}(x_I^i)), f_T(\text{mix}(x_t^i)))/\tau)}{\sum_{k=1}^N \exp(\text{sim}(f_I(\text{aug}(x_I^i)), f_T(\text{mix}(x_t^k)))/\tau)} \quad (2)$$

The only difference with the original CLIP training is that we alternate the AltTexts with our rephrased sentences, with all other components remaining

unaltered. This modification does not incur additional computational complexity or parameter overheads compared to the standard CLIP training process. Through the strategic alternation of AltTexts and our captions, we improve both the quality and diversity of the pre-training dataset without filtering any data points. This approach empowers the model to glean insights from both AltText and VeCap. This simple yet effective strategy elevates the training regimen for CLIP, offering a scalable framework for optimizing other vision-language pre-training efforts utilizing extensive web-crawled data.

4 Experiments

4.1 Pre-training Datasets and Downstream Tasks

Pre-training datasets and training setup. We conduct pre-training experiments on four scales of our datasets (named VeCap) to show the efficiency and scalability of our method. Specifically, we set 3M as small scale, 12M as medium scale, and 100M+ as large scale. We use ViT-B/16 [12] as the vision encoder of CLIP training. Our batch size is 8,192 for small/medium scales (3M/12M), and 32,768 for large scales (100M+). For efficiency purposes, we employ **JAX** [3] and train models on 64 TPUs for the 3M/12M settings, whereas we utilize 512 TPUs for the 100M/200M pre-training configurations. More details can be found in the Appendix A. Although our main focus is to propose a deployable pipeline for scaling up noisy caption data, we also evaluate **CC3M** [5] to show its generalizability and effectiveness on well-curated datasets in our ablation studies and Appendix C.2. We evaluate all pre-trained models on the following three tasks.

Zero-shot image classification. We evaluate all the models on ImageNet [11], ImageNetV2 [34], and VTAB [47]. We select 9 tasks (6 from natural sets and 3 from specialized sets) that are suitable for zero-shot classification tasks such as Flowers102 [29] and Caltech-101 [15] as zero-shot classification tasks. We list the details in the Appendix.

Zero-shot image-text retrieval. We evaluate the pre-trained models on COCO [24] and Flickr30k [32] cross-modal retrieval tasks: Image-to-Text (denoted as I2T) and Text-to-Image (T2I) retrieval. For Flickr30k, we evaluate them on the standard 1K test set. We report the results in terms of Recall@ k as R@1, R@5, and R@10.

Zero-shot image-to-image retrieval. We select GPR1200 [35] for image-to-image retrieval. GPR1200 [35] serves as a general-purpose benchmark for content-based image retrieval, encompassing subsets drawn from six different domains. It includes 1200 categories (10 images per category). Following [35], we do not split images as query and index sets for evaluation. Instead, we perform retrieval of the nearest neighbor for each image and utilize the remaining images as the index set. We report the mean Average Precision (mAP).

4.2 Results on Retrieval Tasks

I2T and T2I retrieval. We summarize the main results in Table 1. We show consistent improvements across Recall@ k metrics in both COCO and Flickr30k

Table 1: Results (Recall@ k) on zero-shot image-to-text and text-to-image retrieval tasks on COCO and Flickr30k. 1.4B-CLIP denotes the in-house CLIP pre-trained on 1.4B web-crawled image-text pairs. We use ViT-B/16 as the vision encoder of CLIP. (*) Denote FLIP uses ViT-L/14.

Data	Model	COCO						Flickr30k					
		Image-to-Text			Text-to-Image			Image-to-Text			Text-to-Image		
		R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
1.8B	ALIGN [18]	58.60	83.00	89.70	45.60	69.80	78.60	88.60	98.70	99.70	75.70	93.80	96.80
400M	FLIP* [23]	60.20	82.60	89.90	44.20	69.20	78.40	89.10	98.50	99.60	75.40	92.50	95.90
400M	OpenAI CLIP	53.76	77.92	85.53	33.09	58.42	68.90	88.00	98.70	99.40	68.70	90.60	95.20
1.4B	In-house CLIP	61.38	82.80	89.78	44.48	69.19	78.28	87.60	97.90	98.80	71.70	91.30	95.24
3M	CLIP	5.46	15.34	22.42	3.28	10.44	15.96	12.20	27.80	37.50	6.36	19.16	27.58
	VeCLIP	22.30	45.00	56.16	13.01	31.61	42.42	40.60	67.30	76.70	27.58	52.44	63.20
	Performance Gain	+16.84	+29.66	+33.74	+9.73	+21.17	+26.46	+28.40	+39.50	+39.20	+21.22	+33.28	+35.62
12M	CLIP	24.52	48.28	59.82	14.28	34.52	46.29	44.70	71.80	80.40	29.06	58.62	70.00
	VeCLIP	47.78	72.54	81.56	31.62	57.19	68.47	73.90	92.30	95.90	55.68	80.78	87.64
	Performance Gain	+23.26	+24.26	+21.74	+17.34	+22.67	+22.18	+29.20	+20.50	+15.50	+26.62	+22.16	+17.64
100M	CLIP	47.24	72.34	81.56	30.61	56.49	67.91	74.40	93.20	96.70	57.16	88.12	88.98
	VeCLIP	64.82	85.56	91.98	46.12	71.19	80.23	89.30	97.70	99.20	73.10	89.12	93.14
	Performance Gain	+17.58	+13.22	+10.42	+15.51	+14.70	+12.32	+14.90	+4.50	+2.50	+15.94	+1.00	+4.16
200M	CLIP	52.20	76.22	85.04	34.97	60.42	71.08	80.90	94.90	97.60	63.26	86.58	92.26
	VeCLIP	67.20	87.28	92.70	48.40	73.26	81.79	91.10	98.50	99.70	76.32	93.50	96.40
	Performance Gain	+15.00	+11.06	+7.66	+13.43	+12.84	+10.71	+10.20	+3.60	+2.10	+13.06	+6.92	+4.14
300M	CLIP	54.24	78.14	86.48	36.98	62.32	72.70	81.30	95.80	97.80	65.80	88.28	93.16
	VeCLIP	67.80	87.94	92.84	48.91	73.54	82.11	91.20	99.10	99.80	76.30	93.00	96.44
	Performance Gain	+13.56	+9.80	+6.36	+11.93	+11.22	+9.41	+9.90	+3.30	+2.00	+10.50	+4.72	+3.28

datasets for both I2T and T2I retrieval tasks. Specifically, for small and medium scales (3M/12M), we attain an improvement of +16.84%/+23.26% in Recall@1 for COCO image-to-text retrieval, respectively. Notably, the strides made in Flickr30k are particularly noteworthy, with a remarkable +28.40%/+29.20% improvement in Recall@1. Subsequently, we scale our approach to 100M and 200M, where we observe sustained and substantial improvements. Notably, we achieve a noteworthy +17.58%/+15.00% enhancement in COCO image-to-text retrieval performance using 100M and 200M, respectively. Furthermore, we observe a diminishing improvement margin as we scale up the dataset. Initially, we achieve a substantial 28.40% improvement in image-to-text retrieval for Flickr30k with the 3M dataset, which subsequently decreases to 10.20% when employing the 200M dataset. These findings show the advantages of our proposed pipeline for enhancing CLIP pre-training. By demonstrating its scalability from 3M to 300M, we provide compelling evidence of its applicability in real-world scenarios, particularly for training CLIP from scratch using WIT datasets.

Image-to-image retrieval. We use GPR1200 [35] with 6 domains for this setting: Google Landmarks V2 (natural and architectural landmarks) denoted as Land, IMDB Faces denoted as Faces, iNat (plants, animals, insects and fungi), INSTRE (planar images and photographs of logos/toys) denoted as INST, ImageNet Sketch denoted as Sketch, and SOP (products and objects, partly isolated). The results (mAP) are summarized in Table 2. We attain a performance gain of 5.22%/3.92% under small/medium scales (3M/12M). Even upon upscaling the dataset to 200M, we observe a notable 1.84% increase in average score across

Table 2: Image-to-image retrieval results (mAP) on 6-domain GPR1200 [35].

Data	Model	Domain Name						
		Land	Faces	iNat	INST	Sketch	SOP	All
3M	CLIP	57.98	20.76	17.61	31.14	18.23	74.29	36.67
	VeCLIP	66.55	23.51	20.43	38.63	24.59	77.65	41.89
12M	CLIP	74.47	30.65	23.60	52.15	30.68	84.25	49.30
	VeCLIP	79.30	31.72	25.53	56.65	41.42	84.69	53.22
100M	CLIP	85.64	51.68	29.66	68.19	42.45	90.38	61.33
	VeCLIP	85.59	42.83	30.72	71.96	52.59	90.54	62.37
200M	CLIP	86.96	56.54	30.95	71.51	46.03	90.95	63.83
	VeCLIP	86.40	48.48	31.72	73.74	56.52	91.16	65.67

Table 3: Zero-shot classification results (Top- k Accuracy) on ImageNet and ImageNetV2 [34].

Data	Model	ImageNet			ImageNetV2		
		Top-1	Top-5	Top-10	Top-1	Top-5	Top-10
3M	CLIP	5.46	21.05	28.70	7.09	18.52	25.83
	VeCLIP	15.98	34.11	43.23	13.51	30.03	38.93
12M	CLIP	31.60	58.80	69.49	27.03	52.68	63.37
	VeCLIP	38.11	66.74	76.36	32.53	60.16	70.50
100M	CLIP	58.64	85.82	91.79	50.96	79.77	86.91
	VeCLIP	60.77	87.77	93.16	54.17	82.51	89.24
200M	CLIP	63.72	89.26	94.11	56.84	83.50	89.79
	VeCLIP	64.62	90.27	94.90	57.67	85.24	91.62

Table 4: Zero-shot classification accuracy. Top-1 accuracies (%) of VTAB [47] across 9 tasks (6 from natural and 3 from specialized sets) are reported. Full table is in Appendix Table A6.

Data	Model	Natural Sets						Specialized Sets			Average
		Caltech101	CIFAR100	SVHN	DTD	OxPet	Flowers102	EuroSAT	RESISC45	Camelyon	
3M	CLIP	39.50	9.83	20.89	7.42	7.44	10.40	11.94	7.93	50.65	18.45
	VeCLIP	54.30	17.74	18.74	11.23	10.09	22.75	7.35	16.54	52.52	23.48
12M	CLIP	70.43	30.06	30.11	30.69	34.51	33.67	8.87	30.05	53.46	35.76
	VeCLIP	70.58	45.10	23.61	30.90	36.22	43.94	27.46	38.09	55.54	41.27
100M	CLIP	81.44	54.75	38.70	57.28	70.51	51.71	34.45	48.56	53.87	54.59
	VeCLIP	81.64	64.62	46.49	57.51	64.81	66.41	46.23	51.75	58.51	59.78
200M	CLIP	82.30	61.87	42.83	64.29	75.60	58.67	46.73	55.59	59.30	60.79
	VeCLIP	83.14	68.14	44.93	61.95	72.61	68.51	47.36	55.10	62.59	62.70

six domains. Notably, our primary performance boost is derived from the Sketch domain, underlining the crucial role of visual concepts in zero-shot transferability. Consequently, our visually-enriched captions play a pivotal role in learning such transferability towards downstream tasks.

Data efficiency for pre-training. To show the data efficiency of VeCLIP, we include ALIGN [18], pre-trained on 1.8B data (denoted as 1.8B-ALIGN), and our in-house CLIP [33] model trained on 1.4B data (denoted as 1.4B-CLIP) as baselines trained at a significantly larger scale. We use these models utilizing over tenfold more data compared to our setting to show the data efficiency of VeCLIP training. VeCLIP can outperform 1.4B-CLIP model when scaling up to 100M, representing approximately 7% of its size, across nearly all downstream tasks. Specifically, in COCO, we achieve +3.44%/+1.64% gain in Recall@1 for both retrieval tasks. Upon further scaling to 200M, the improvement becomes even more pronounced, reaching +5.82%/+3.92%. Furthermore, we achieve a notable +8.60%/+2.80% gain in COCO retrieval, as well as a +2.50%/+0.62% improvement in Flickr30k, when compared to the 1.8B-ALIGN model. Remarkably, these improvements are achieved with only 11.1% of the data utilized in the pre-training process. These results show the data efficacy of VeCLIP. When

we scale it to 300M, the results are similar to 200M. The results on 300M can be found in Appendix. Therefore, we stop further scaling up the dataset.

4.3 Results on Image Classification

ImageNet. We use the same prompt as CLIP (“A photo of a [classname].”) for zero-shot evaluation on both ImageNet [11] and ImageNetV2 [34]. The main results are summarized in Table 3. We report Top-1, Top-5, and Top-10 accuracies. In small and medium-scale settings, we observe a substantial improvement: +10.52%/+6.42% gains in Top-1 accuracy on ImageNet/ImageNetV2 under the 3M setting, and +6.51%/5.50% gains under the 12M setting. While the improvement becomes marginal upon scaling to 100M/200M, we still achieve noteworthy +2.07%/+3.21% and +0.90%/+0.83% gains on 100M and 200M across ImageNet and ImageNetV2, respectively. This shows the data efficiency of our pre-training approach.

Visual Task Adaptation Benchmark (VTAB). Besides ImageNet and ImageNetV2, we also select VTAB [47] for evaluation. Table 4 summarizes zero-shot image classification results for both the original CLIP models and our models, utilizing the identical prompt set from CLIP. Every class label is expanded using a collection of prompt templates, as defined by CLIP, including examples like “A photo of a [classname].” The class embedding is then computed by taking the average of the embeddings of all such templates, followed by L2-normalization. Our approach consistently achieves comparable or superior performance to CLIP across the majority of datasets. For instance, we observe an average accuracy gain of over 5% under settings of 3M, 12M, and 100M. Even upon scaling up to 200M, we maintain a notable gain of +1.91%. These results show great robustness on zero-shot classification tasks across different data distributions.

Performance trend across scales. Besides the performance gain, we also visualize the performance trend across data scales in pre-training. As shown in Figure A1 in Appendix, the performance of CLIP utilizing original AltTexts exhibits a marked surge with the increased data size: while its starting point is poor at 3M, it demonstrates swift progression up to 12M and 100M. However, once scaled beyond 100 million, the performance trend exhibits a gradual and eventually saturated growth. On the other hand, commencing with a higher baseline, VeCLIP employing VeCap demonstrates substantial improvement in comparison to CLIP within small to medium scales (3M and 12M). As we progress beyond 300M, the performance gains of VeCLIP become relatively incremental but still noticeable in retrieval tasks. Both CLIP and VeCLIP reach a saturation point when scaled up to 100M: once over 100M, the performance gain becomes less impressive as a relatively small but high-quality dataset may reach the upper bound of the learning ability constrained by the model’s architecture.

4.4 Complementary to other datasets to achieve SOTA performance

Our VeCap datasets with visual-enriched captions can also be complementary to other well-curated dataset. For example, DFN [14] has shown benefits on CLIP.

Table 5: CLIP training with VeCap and DFN [14], and its comparison with the state-of-the-art models.

Model Resolution Data			COCO (R@1)		Flickr30k (R@1)		ImageNet
			I2T	T2I	I2T	T2I	
B/16	224	DFN [14]	63.0	43.2	87.1	70.4	76.2
		VeCap+DFN	66.3	45.1	88.8	73.6	76.2
<i>Comparison to other state-of-the-art models</i>							
L/14	224	DFN [14]	68.5	48.5	89.2	75.1	81.4
		FLIP [23]	60.2	44.2	89.1	75.4	74.6
		VeCap+DFN	70.8	49.5	92.4	78.4	81.1
	224	MetaCLIP [43]	67.2	49.5	92.1	78.5	80.5
H/14	378	DFN [14]	71.8	55.6	94.0	82.1	84.4
	336	VeCap+DFN	72.8	52.3	93.6	82.6	83.1

To demonstrate that, we train CLIP models with VeCap and DFN separately and also a combination with them. All the models are trained under same configuration for learning rate, maximum steps, and so on. We summarize the results in Table 5. The high-quality descriptive captions from VeCap can achieve superior results compared to DFN in retrieval tasks. However, the performance on classification tasks are inferior. After we combine DFN and VeCap for training, CLIP can achieve the most improvements for all model sizes.

We also train a H/14 model with resolution 336x336, and compare it with the state-of-the-art models like MetaCLIP [43] and DFN [14]. The results are summarized in row 6 to 8 of table 5. Albeit trained on different resolutions and recipes, the CLIP model with VeCap+DFN is compatible with other models and provide yet another option for downstream tasks³. Our VeCLIP with DFN [14] can outperform FLIP [23] and OpenAI CLIP with different backbones (as shown in Table A9 in Appendix). Specifically, our ViT-H/14 model achieves impressive 83.1% of accuracy on ImageNet. We leave the further study of combing the synthetic data (VeCap) with other data curation approaches as a future work.

4.5 Ablation Study

Importance of visual-enriched concepts. Different from previous rewriting methods, our primary emphasis lies in fusing visual-enriched concepts extracted from images. The ablation findings are summarized in Table 6a. We use 3M/12M as examples to show the performance gain in small/medium scales. AltTexts shows their limitation in retrieval tasks due to its noise and limited image-specific information. VeC generated from LLaVA can boost the performance on retrieval tasks but may hurt the performance on ImageNet zero-shot task. Introducing VeCap can further improve in all settings. Intriguingly, the zero-shot ImageNet results still lag behind the original AltText. In summary, our

³ Note we took the DFN-H/14 model from its original paper, which is trained 7 epochs, our model is only trained roughly around 2 epochs.

Data	Caption	Prompt Constraint	COCO (R@1)		Flickr30k (R@1)		ImageNet	ImageNetV2
			I2T	T2I	I2T	T2I		
WIT-3M	AltText	-	5.18	3.40	10.50	6.88	8.02	<u>6.88</u>
	VeC	-	16.76	9.57	32.60	20.06	7.31	6.58
	VeCap	✗	<u>17.34</u>	<u>9.52</u>	<u>37.30</u>	<u>21.62</u>	<u>8.12</u>	6.83
	VeCap	✓	18.10	9.51	40.00	21.94	8.20	7.39
WIT-12M	AltText	-	22.58	14.23	44.40	30.90	31.14	25.91
	VeC	-	40.06	24.59	64.10	43.46	7.29	14.74
	VeCap	✗	<u>44.52</u>	27.46	<u>70.90</u>	<u>50.46</u>	<u>21.05</u>	18.11
	VeCap	✓	46.82	<u>26.61</u>	72.60	50.94	20.99	<u>18.41</u>

(a) Importance of visual-enriched concepts for data quality. We use the AltText with “Highest CLIP Score” (HCS) if multiple AltTexts exist on the same image in all settings.

Data	AltText	VeCap	Training Sampling	COCO (R@1)		Flickr30k (R@1)		ImageNet	ImageNetV2
				I2T	T2I	I2T	T2I		
WIT-3M	✓	✗	HCS	5.18	3.40	10.50	6.88	8.02	6.88
	✓	✗	random	5.46	3.28	12.20	6.36	8.26	7.09
	✗	✓	HCS	18.10	9.51	<u>40.00</u>	21.94	8.20	7.39
	✓	✓	HCS&mixed	<u>19.70</u>	<u>12.14</u>	39.30	<u>25.60</u>	<u>14.83</u>	<u>12.36</u>
	✓	✓	random&mixed	22.30	13.01	40.60	27.58	15.98	13.51
WIT-12M	✓	✗	HCS	22.58	14.23	44.40	30.90	31.14	25.91
	✓	✗	random	23.32	14.28	44.70	29.06	31.60	27.03
	✗	✓	HCS	<u>46.82</u>	26.61	<u>72.60</u>	50.94	20.99	18.41
	✓	✓	HCS&mixed	46.00	<u>31.10</u>	72.50	56.82	<u>37.45</u>	<u>32.41</u>
	✓	✓	random&mixed	47.78	31.62	73.90	<u>55.68</u>	38.11	32.51
CC3M	✓	✗	-	13.88	9.64	26.30	18.04	14.59	12.52
	✓	✓	random&mixed	32.04	22.07	57.20	36.54	20.73	17.90

(b) Importance of the mixed training scheme for data variety. “HCS” refers to using the AltText with “Highest CLIP Score” while “random” refers to randomly selecting one if multiple AltTexts exist.

Table 6: Ablation study of VeCLIP. The highest score is bold, and the second is underlined. “mixed” is our proposed mixed training scheme to alternate among captions.

VeCap exerts a profound influence on retrieval prowess yet exerts a negative effect on classification tasks. There are two possible reasons: 1) there can be a distributional shift in prompts from pre-training to zero-shot inference in ImageNet, particularly noteworthy given the extended length and augmented visual content of VeCap; 2) the data diversity is hurt by LLM rewriting as LLM uses the same writing/paraphrasing style to fuse VeCap and AltText.

Importance of mixed training strategies. To mitigate the aforementioned issues, we propose a mixed training scheme to alternate between AltTexts and VeCap to provide more data variety during pre-training. We summarize the ablation results of VeCLIP in Table 6b. First, we observe a slight performance improvement by randomly selecting one AltText in cases where multiple AltTexts are associated with an image. This practice augments data diversity during pre-training. Second, interchanging between AltText and VeCap proves to be advantageous, not only in retaining substantial performance gains in retrieval tasks but also in markedly elevating zero-shot results on ImageNet. Lastly, leveraging all AltTexts and VeCap within the mixed training approach in VeCLIP achieves superior results across nearly all settings.

Larger backbone architecture. We also investigate a larger backbone architecture, e.g., ViT-L/14 and ViT-H/14. The detailed results can be found in both Table 5 and Appendix C.1. VeCLIP scaled up in backbone size can consistently outperform the original CLIP in all downstream tasks. Besides, a larger backbone (ViT-L/14) can also achieve up to 5.87% improvement compared to ViT-B/16. These findings support the effectiveness of VeCLIP in improving CLIP pre-training, regardless of the specific underlying backbone architecture.

Generalizability of VeCap on well-curated datasets. Besides our WIT datasets, we evaluate VeCap on well-curated CC3M/CC12M. Table 6b shows CLIP achieves better performance when pre-trained on CC3M compared to pre-trained on WIT-3M, indicating the importance of high-quality captions for pre-training. With VeCap to further improve the quality of CC3M’s captions, CLIP can achieve significant improvement, since the captions of CC3M are of higher quality than our noisy WIT dataset. CC3M with its original captions can outperform the performance of our WIT-3M with AltTexts, indicating CC3M is of higher quality. VeCap can significantly improve CLIP under CC3M settings, e.g., +18.16% on the I2T task of COCO and +6.14% on ImageNet, showing its generalizability on well-curated datasets. More results are in Appendix C.2.

5 Discussion

Conclusion. We present a simple yet effective approach to improve CLIP pre-training with leveraging LLaVA and LLMs to rewrite the captions with more visual-enriched concepts. Then, we propose a mixed training scheme to improve data diversity. VeCLIP is intentionally designed to be scalable and adaptable for handling extensive image-text datasets obtained from web crawling. We conduct a thorough evaluation of VeCLIP on a diverse range of raw and noisy datasets, spanning small, medium, and large scales. The results reveal a substantial performance boost, providing compelling evidence for the effectiveness of our strategy in enhancing large-scale VLM pre-training. VeCLIP can significantly reduce the computational cost and the size of training data for large models to reach competitive results as vanilla CLIP.

Future work. We employ CLIP as an illustrative instance to highlight the importance of aligning text and images within the training dataset. For future work, we plan to use the collected large-scale dataset to improve the pre-training of other types of VLMs. Further, LLM can generate outputs that encompass factual inaccuracies and hallucinations. Thus, we also plan to delve into more sophisticated filtering techniques to remove such descriptions.

Limitation. We only leverage LLaVA to exploit the visual concepts. However, the quality measurement metric for such generative AI is still under study. Also, the generated captions may potentially inherit bias from the foundation models we use (LLaVA and Vicuna).

References

1. Abbas, A., Tirumala, K., Simig, D., Ganguli, S., Morcos, A.S.: Semdedup: Data-efficient learning at web-scale through semantic deduplication. arXiv preprint arXiv:2303.09540 (2023)
2. Betker, J., Goh, G., Jing, L., Brooks, T., Wang, J., Li, L., Ouyang, L., Zhuang, J., Lee, J., Guo, Y., Manassra, W., Dhariwal, P., Chu, C., Jiao, Y., Ramesh, A.: Improving image generation with better captions. OpenAI (2023)
3. Bradbury, J., Frostig, R., Hawkins, P., Johnson, M.J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., Zhang, Q.: JAX: composable transformations of Python+NumPy programs. Github (2018), <http://github.com/google/jax>
4. Cao, L., Zhang, B., Chen, C., Yang, Y., Du, X., Zhang, W., Lu, Z., Zheng, Y.: Less is more: Removing text-regions improves clip training efficiency and robustness. arXiv preprint arXiv:2305.05095 (2023)
5. Changpinyo, S., Sharma, P., Ding, N., Soricut, R.: Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3558–3568 (2021)
6. Chen, L., Li, J., Dong, X., Zhang, P., He, C., Wang, J., Zhao, F., Lin, D.: Sharegpt4v: Improving large multi-modal models with better captions. arXiv preprint arXiv:2311.12793 (2023)
7. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
8. Cheng, G., Han, J., Lu, X.: Remote sensing image scene classification: Benchmark and state of the art. Proceedings of the IEEE **105**(10), 1865–1883 (Oct 2017). <https://doi.org/10.1109/jproc.2017.2675998>, <http://dx.doi.org/10.1109/JPROC.2017.2675998>
9. Cherti, M., Beaumont, R., Wightman, R., Wortsman, M., Ilharco, G., Gordon, C., Schuhmann, C., Schmidt, L., Jitsev, J.: Reproducible scaling laws for contrastive language-image learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2818–2829 (2023)
10. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., Vedaldi, A.: Describing textures in the wild. In: Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (2014)
11. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
12. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2020)
13. Fan, L., Krishnan, D., Isola, P., Katabi, D., Tian, Y.: Improving clip training with language rewrites. arXiv preprint arXiv:2305.20088 (2023)
14. Fang, A., Jose, A.M., Jain, A., Schmidt, L., Toshev, A.T., Shankar, V.: Data filtering networks. In: NeurIPS 2023 Workshop on Distribution Shifts: New Frontiers with Foundation Models (2023)
15. Fei-Fei, L., Fergus, R., Perona, P.: Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories.

- In: 2004 conference on computer vision and pattern recognition workshop. pp. 178–178. IEEE (2004)
16. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16000–16009 (2022)
 17. Helber, P., Bischke, B., Dengel, A., Borth, D.: Introducing eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. In: IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium. pp. 204–207. IEEE (2018)
 18. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: International conference on machine learning. pp. 4904–4916. PMLR (2021)
 19. Krizhevsky, A.: Learning multiple layers of features from tiny images. Canadian Institute for Advanced Research (2009)
 20. Kwon, G., Cai, Z., Ravichandran, A., Bas, E., Bhotika, R., Soatto, S.: Masked vision and language modeling for multi-modal representation learning. In: The Eleventh International Conference on Learning Representations (2023)
 21. Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: International Conference on Machine Learning. pp. 12888–12900. PMLR (2022)
 22. Li, Y., Fan, H., Hu, R., Feichtenhofer, C., He, K.: Scaling language-image pre-training via masking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 23390–23400 (2023)
 23. Li, Y., Fan, H., Hu, R., Feichtenhofer, C., He, K.: Scaling language-image pre-training via masking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 23390–23400 (2023)
 24. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13. pp. 740–755. Springer (2014)
 25. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. arXiv preprint arXiv:2304.08485 (2023)
 26. Maini, P., Goyal, S., Lipton, Z.C., Kolter, J.Z., Raghunathan, A.: T-mars: Improving visual representations by circumventing text feature learning. arXiv preprint arXiv:2307.03132 (2023)
 27. Mu, N., Kirillov, A., Wagner, D., Xie, S.: Slip: Self-supervision meets language-image pre-training. In: European Conference on Computer Vision. pp. 529–544. Springer (2022)
 28. Nguyen, T., Gadre, S.Y., Ilharco, G., Oh, S., Schmidt, L.: Improving multimodal datasets with image captioning. arXiv preprint arXiv:2307.10350 (2023)
 29. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: 2008 Sixth Indian conference on computer vision, graphics & image processing. pp. 722–729. IEEE (2008)
 30. Parkhi, O.M., Vedaldi, A., Zisserman, A., Jawahar, C.: Cats and dogs. In: 2012 IEEE conference on computer vision and pattern recognition. pp. 3498–3505. IEEE (2012)
 31. Pham, H., Dai, Z., Ghiasi, G., Liu, H., Yu, A.W., Luong, M.T., Tan, M., Le, Q.V.: Combined scaling for zero-shot transfer learning. arXiv preprint arXiv:2111.10050 (2021)

32. Plummer, B.A., Wang, L., Cervantes, C.M., Caicedo, J.C., Hockenmaier, J., Lazebnik, S.: Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In: Proceedings of the IEEE international conference on computer vision. pp. 2641–2649 (2015)
33. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML. pp. 8748–8763 (2021)
34. Recht, B., Roelofs, R., Schmidt, L., Shankar, V.: Do imagenet classifiers generalize to imagenet? In: International conference on machine learning. pp. 5389–5400. PMLR (2019)
35. Schall, K., Barthel, K.U., Hezel, N., Jung, K.: Gpr1200: a benchmark for general-purpose content-based image retrieval. In: International Conference on Multimedia Modeling. pp. 205–216. Springer (2022)
36. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., 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)
37. Schuhmann, C., Vencu, R., Beaumont, R., Kaczmarczyk, R., Mullis, C., Katta, A., Coombes, T., Jitsev, J., Komatsuzaki, A.: Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. arXiv preprint arXiv:2111.02114 (2021)
38. Sennrich, R., Haddow, B., Birch, A.: Improving neural machine translation models with monolingual data. arXiv preprint arXiv:1511.06709 (2015)
39. Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al.: Llama: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971 (2023)
40. Veeling, B.S., Linmans, J., Winkens, J., Cohen, T., Welling, M.: Rotation equivariant cnns for digital pathology. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part II 11. pp. 210–218. Springer (2018)
41. Wei, J., Zou, K.: Eda: Easy data augmentation techniques for boosting performance on text classification tasks. arXiv preprint arXiv:1901.11196 (2019)
42. Wu, W., Timofeev, A., Chen, C., Zhang, B., Duan, K., Liu, S., Zheng, Y., Shlens, J., Du, X., Gan, Z., et al.: Mofi: Learning image representations from noisy entity annotated images. arXiv preprint arXiv:2306.07952 (2023)
43. Xu, H., Xie, S., Tan, X.E., Huang, P.Y., Howes, R., Sharma, V., Li, S.W., Ghosh, G., Zettlemoyer, L., Feichtenhofer, C.: Demystifying clip data. arXiv preprint arXiv:2309.16671 (2023)
44. Yu, J., Wang, Z., Vasudevan, V., Yeung, L., Seyedhosseini, M., Wu, Y.: Coca: Contrastive captioners are image-text foundation models. arXiv preprint arXiv:2205.01917 (2022)
45. Yuan, L., Chen, D., Chen, Y.L., Codella, N., Dai, X., Gao, J., Hu, H., Huang, X., Li, B., Li, C., et al.: Florence: A new foundation model for computer vision. arXiv preprint arXiv:2111.11432 (2021)
46. Yuval, N.: Reading digits in natural images with unsupervised feature learning. In: Proceedings of the NIPS Workshop on Deep Learning and Unsupervised Feature Learning (2011)
47. Zhai, X., Puigcerver, J., Kolesnikov, A., Ruysen, P., Riquelme, C., Lucic, M., Djolonga, J., Pinto, A.S., Neumann, M., Dosovitskiy, A., et al.: A large-scale study of representation learning with the visual task adaptation benchmark. arXiv preprint arXiv:1910.04867 (2019)

48. Zhai, X., Wang, X., Mustafa, B., Steiner, A., Keysers, D., Kolesnikov, A., Beyer, L.: Lit: Zero-shot transfer with locked-image text tuning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18123–18133 (2022)
49. Zheng, L., Chiang, W.L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E.P., Zhang, H., Gonzalez, J.E., Stoica, I.: Judging llm-as-a-judge with mt-bench and chatbot arena (2023)