

CAPSFUSION: Rethinking Image-Text Data at Scale

Qiyang Yu^{1,2*} Quan Sun^{2*} Xiaosong Zhang² Yufeng Cui² Fan Zhang²
 Yue Cao³ Xinlong Wang^{2†} Jingjing Liu¹

¹ Institute for AI Industry Research (AIR), Tsinghua University ² Beijing Academy of Artificial Intelligence
³ Independent Researcher

Code & Dataset: <https://github.com/baaivision/CapsFusion>

Abstract

Large multimodal models demonstrate remarkable generalist ability to perform diverse multimodal tasks in a zero-shot manner. Large-scale web-based image-text pairs contribute fundamentally to this success, but suffer from excessive noise. Recent studies use alternative captions synthesized by captioning models and have achieved notable benchmark performance. However, our experiments reveal significant Scalability Deficiency and World Knowledge Loss issues in models trained with synthetic captions, which have been largely obscured by their initial benchmark success. Upon closer examination, we identify the root cause as the overly-simplified language structure and lack of knowledge details in existing synthetic captions. To provide higher-quality and more scalable multimodal pretraining data, we propose CAPSFUSION, an advanced framework that leverages large language models to consolidate and refine information from both web-based image-text pairs and synthetic captions. Extensive experiments show that CAPSFUSION captions exhibit remarkable all-round superiority over existing captions in terms of model performance (e.g., 18.8 and 18.3 improvements in CIDEr score on COCO and NoCaps), sample efficiency (requiring 11-16 times less computation than baselines), world knowledge depth, and scalability. These effectiveness, efficiency and scalability advantages position CAPSFUSION as a promising candidate for future scaling of LMM training.

1. Introduction

Large Multimodal Models [3, 36, 51] (LMMs), which as versatile multimodal generalists bridge powerful pretrained large language models [52, 53] and vision encoders [43, 50], have garnered significant success in zero-shot multimodal tasks. Although image-text pairs harvested directly from the web [45] contribute instrumentally to the success of current LMMs, such web-scale data tend to be noisy

*Equal contribution. † Correspondence to wangxinlong@baai.ac.cn.




Figure 1. Training process of models trained on different captions.



Raw Caption:

A male Northern Cardinal is feeding a fledgling on the top of a tree branch.

Synthetic Caption:

Two birds sitting on a tree branch.

(a) Raw caption and synthetic caption (from caption model)



Raw Caption:

Justin Timberlake performs at the 2017 Pilgrimage Music Festival on September 23, 2017 in Franklin, Tennessee.

Conceptual Caption:

pop artist performs at the festival in a city

(b) Data Collection Process of Conceptual Captions

Figure 2. (a) Comparison of raw and synthetic captions for training. (b) Data processing of Conceptual Captions [47], where real-world information is substituted with generic concepts.

and sub-optimal for model training [27, 34]. Thus, strategies have been devised to harness synthetic captions generated by image captioning model [34], which has augmented model performance notably [5, 14, 15, 22, 35, 51] by adopting large-scale synthetic caption datasets such as LAION-COCO [1] and BLIP-LAION [34].

Although achieving promising performance on classic benchmarks such as COCO Caption [12], our further evaluations on recent benchmarks such as SEED-Bench [32] reveal that training LMMs with large-scale synthetic captions

Input Image:				
Input Prompt:	Describe the image in detail. The image shows	Introduce some books by him. answer:	where is this photo taken? short answer:	where is it?
M1: (trained on Raw)	President Bill Clinton and First Lady Hillary Rodham Clinton cutting the ribbon.....	George Orwell's 1984, Animal Farm, and the collected essays	Melbourne, Australia, long answer: melbourne, australia, novak Djokovic australian open trophy	the 12 apostles, victoria, australia
M2: (trained on Synthetic)	two men and one woman cutting a red ribbon.	I'm not going to tell you what I'm doing	The tennis player is kissing his trophy	I don't know.
M3: (trained on CapsFusion)	President Bill Clinton and First Lady Hillary Rodham Clinton cutting the ribbon.....	George Orwell's 1984, Animal Farm, Down and Out in Paris and London	Melbourne, Australia, long answer: melbourne, australia, where novak djokovic kisses the australian open trophy	The Twelve Apostles are a group of cliffs located along the Great Ocean Road in Victoria, Australia.

Figure 3. Outputs of models trained with different caption datasets. Models trained on raw and CAPSFUSION captions (M1 and 3) possess strong world knowledge (in blue), while the model trained on synthetic captions (M2) can only generate generic concepts (in red).

alone is problematic. We conduct a closer examination of the large-scale training process of LMMs and observe that model training on synthetic captions rapidly reaches a saturation point, beyond which the model performance may even degrade (as illustrated by the green lines in Fig. 1). While this severe *Scalability Deficiency* may not be readily apparent on traditional benchmarks such as COCO caption (Fig. 1-a), it becomes notably pronounced (Fig. 1-b) on the new benchmark SEED-Bench, which supports a much more comprehensive assessment of LMMs than COCO. We conduct further analysis on the generated outputs from different models trained with captions of varying quality. Fig. 3 illustrates system responses trained on Raw captions (M1), Synthetic captions (M2), and our captions (M3). These examples demonstrate that the outputs from M2, in particular, suffer from severe *World Knowledge Loss*, constituting only high-level concepts while missing all the details about well-known people, locations, events, etc. The generated sentences by M3 (trained on our captions) are more natural and semantically richer than those from M1 and M2.

Through examining the differences between raw caption data and synthetic data used in training, we observe that the simplistic syntactic and semantic structures in synthetic captions (Fig. 2-a) may have potentially attributed to the *Scalability Deficiency* and *World Knowledge Loss* issues, which so far have been obscured by their initial benchmark success. The root cause is that currently used captioning models (e.g. BLIP [34] used in LAION-COCO [1]) for generating synthetic captions heavily rely on academic datasets such as COCO and Conceptual Captions [47] for training. These datasets replace specific details (e.g. people’s names, locations, landmarks) with more generic *conceptual* placeholders (e.g. ‘person’, ‘city’) in the data collection process (Fig. 2-b). Although this eases the training of captioning models, it inevitably results in the loss of a substantial reser-

voir of valuable real-world information in the trained model, which learns an overly-simplified language structure with basic semantics. Consequently, LMMs trained on the synthetically simplified datasets generated by these captioning models suffer from a deficiency in language complexity and knowledge depth.

Therefore, to train a scalable LMM with abundant real-world knowledge, it is crucial to develop an effective strategy to better synthesize caption data while distilling real-world knowledge from web-based image-text pairs. There have been some recent attempts to leverage both raw and synthetic captions straightforwardly, by simply mixing them with a fixed hand-tuned ratio [16, 19, 39]. In this work, we propose CAPSFUSION, a more advanced pipeline that leverages large language models (LLMs) to enhance the quality of large-scale image-text data. CAPSFUSION first uses a captioning model [34] (following [1, 34]) to generate synthetic captions for images. Then, it utilizes ChatGPT [46] to organically integrate raw and synthetic captions, by extracting real-world knowledge from the structure-flawed raw captions while merging with structured but syntactically simplified synthetic captions. Our evaluations show that ChatGPT excels in this task, but is non-scalable due to its restrictive access. To overcome this limitation, we use the outputs generated by ChatGPT as training data to finetune a LLaMA [53]. Evaluation of this finetuned, task-specific LLM demonstrates that it performs *on par with* ChatGPT and consistently produces high-quality captions, while easy to scale up. The trained model is then employed for large-scale caption fusion (examples are presented in Fig. 4, which clearly demonstrate the advantages of CAPSFUSION).

Extensive experiments show that CAPSFUSION captions demonstrate remarkable all-around superiority, as a better substitute for both synthetic and raw captions in the train-

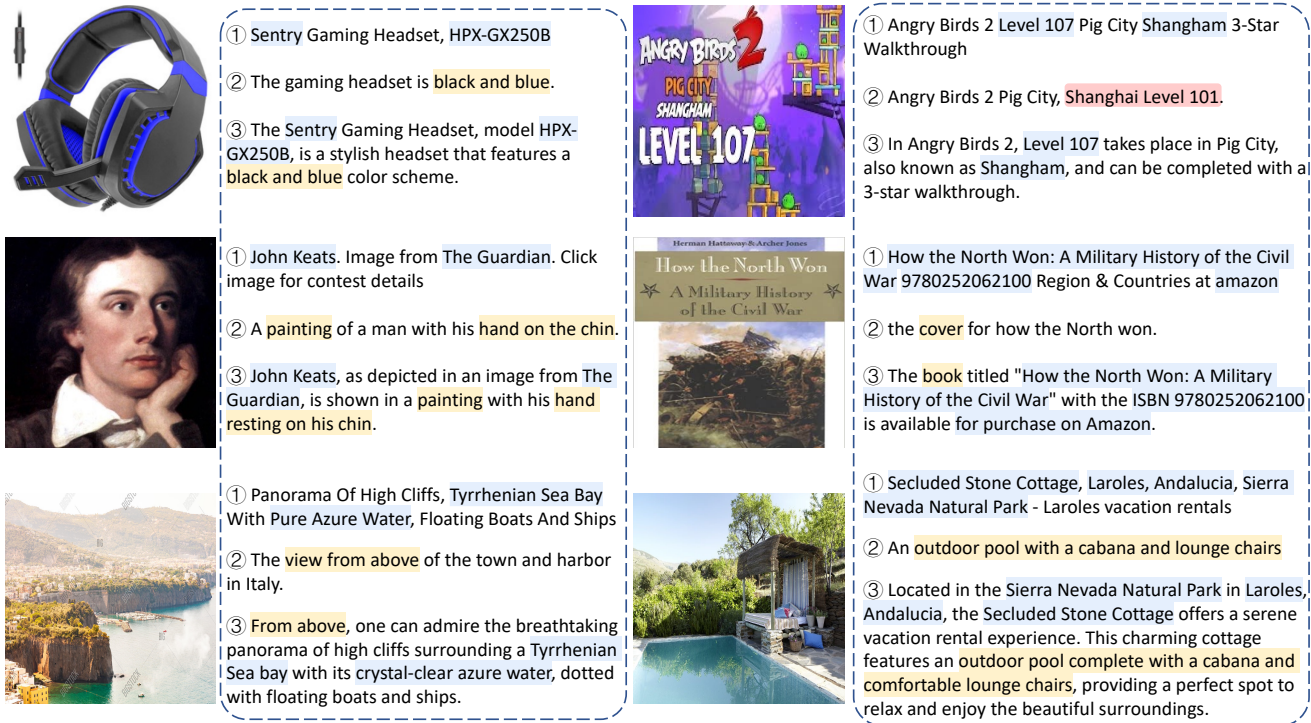


Figure 4. Examples of ① raw captions (from LAION-2B), ② synthetic captions (from LAION-COCO, generated by BLIP), and their corresponding ③ CAPSFUSION captions. Knowledge from raw captions (in blue) and information from synthetic captions (in yellow) are organically fused into integral CAPSFUSION captions. CAPSFUSION captions can also correct false information in synthetic captions (in red). More examples can be found in Fig. 8.

ing of LMMs. In terms of *model performance*, CAPSFUSION captions clearly outperform synthetic captions by substantial margins, with an improvement of 18.8, 18.3, 19.7, and 15.6 in CIDEr score on COCO, NoCaps, TextCaps, and Flickr30K datasets, respectively. This compelling advantage extends to *sample efficiency* as well. Refined captions from CAPSFUSION require 11-16 times less computation to achieve high performance similar to synthetic captions. Furthermore, our investigation unveils that CAPSFUSION captions surpass raw captions by a considerable margin when evaluated on *world knowledge*. Also importantly, CAPSFUSION captions demonstrate greater *scalability*, meaning that model performance continues to improve with an increased volume of training samples. This scalability advantage, critical for the training of large-scale models, positions CAPSFUSION as a promising candidate for further scaling efforts in LMM training.

2. Related Work

Image-text Data Enhancement LaCLIP [16] utilizes LLM to rewrite raw captions, whose performance can be limited due to severe hallucination, because of limited visual information and low-quality raw captions. [19, 39] in-

vestigate how to filter and then mix raw and synthetic captions to induce a better CLIP model [43]. FuseCap [44] uses visual experts like object detector to improve captions, while we use large-scale data from the web to enhance. Our concurrent work VeCLIP [31] proposes to use LLM to combine information from raw and synthetic captions. The difference is that they directly use an existing LLM for inference, while we finetune a state-of-the-art open-source LLM with training data generated by ChatGPT. In addition, they have no explicit instructions such as extracting world knowledge present in raw captions and referring sentence structure of synthetic captions, which we use to help LLMs make informed decisions during the caption fusion process.

All recent studies focus on training CLIP models. We instead investigate LMMs and derive insights from a new perspective, such as mixing raw and synthetic captions [16, 31, 39] induces no improvement than separate captions.

Large Multimodal Models With the success of large language models [8, 53] (LLMs), recent studies explore building large multimodal models [9–11, 13, 20, 23, 26, 33, 40, 56, 60–62, 65, 67–69] (LMMs) on LLMs with pretrained vision encoders [43, 50, 55]. Most existing works commonly use the prediction of the next text token as the ob-

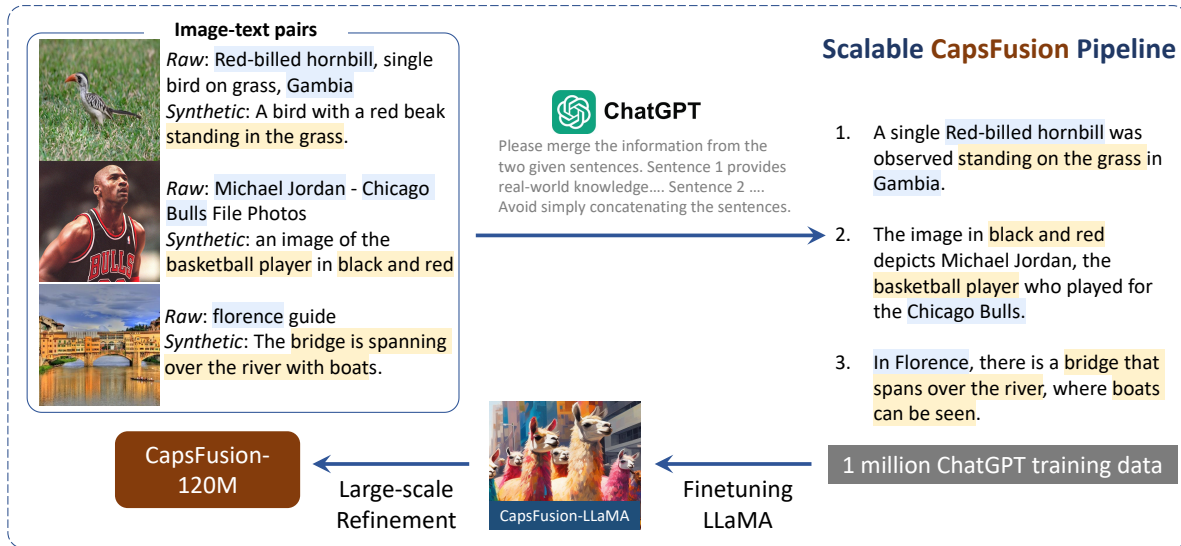


Figure 5. Illustration of the scalable CAPSFUSION pipeline for generating high-quality large-scale image-text data.

jective [14, 25, 35, 36]. Another type of LMMs learns to predict both image and text tokens [15, 21, 51, 63], endowing models with more versatile abilities of processing both text and image generation tasks, while maintaining image-to-text performance comparable to LMMs trained with only text token supervision.

3. CAPSFUSION

Large Multimodal Models [3, 28, 35, 49, 57, 66, 71] serve as a powerful generalist for diverse multimodal tasks. Typical LMM generalist unifies image-to-text tasks only (e.g. image captioning and visual question answering). Recent studies such as Emu [51] further enhance the capabilities of multimodal generalist by enabling it to perform both image-to-text and text-to-image tasks in a zero-shot manner [22, 29, 30, 58, 70].

Learning Objective of LMM. The LMM generalist ability originates from a GPT-style auto-regressive training objective [42], wherein the model learns to predict the next token in a sequence. As a result of this training paradigm, during inference, the model exhibits a remarkable capability to generate appropriate completions for a wide range of tasks.

Image-text pairs are the most commonly used multimodal pretraining data for learning vision-language alignment. Specifically, given a dataset \mathcal{D} consisting of image-text pairs (I, T) , where I represents the image and T represents the text represented by a sequence of text tokens $T = \{t_1, t_2, \dots, t_n\}$. The typical training objective is maximizing the conditional likelihood of text tokens T given I

in an auto-regressive manner:


$$\max_{\theta} \frac{1}{|\mathcal{D}|} \sum_{(I, T) \in \mathcal{D}} \sum_{i=1}^n \log P(t_i | t_1, \dots, t_{i-1}, I; \theta) \quad (1)$$

Under this training objective, the presence of noisy captions can lead the model to generate extraneous words. Conversely, if the captions are overly simplistic in nature, the model may learn a simplified output style, resulting in a loss of language complexity. Therefore, high-quality image-text pairs are in urgent need to power new-generation LMMs.

Caption Generation. Given raw image-text pairs, CAPSFUSION first generates synthetic captions using image captioning models following [1, 34]. In previous analysis (Figs. 1 to 3), we find that raw captions contain a wealth of real-world knowledge but are noisy, while synthetic captions have clean structures but lack in-depth real-world knowledge, which exhibits severe scalability issues. Thus, our objective is to develop a scalable framework to organically integrate information from both raw and synthetic captions, to create a comprehensive refined image-text dataset.

Caption Fusion via ChatGPT. In CAPSFUSION, we use ChatGPT to fuse raw and synthetic captions given a prompt. The task instruction is structured in three key elements: the task description, caption property, and the desired output specifications. Specifically, we first include a task description that conveys the following objective to ChatGPT: *Please merge the information from two provided sentences.* Furthermore, we provide the distinct properties of the two captions involved, with the following contextual guidance:

Raw captions offer detailed real-world information, yet it suffers from flaws in sentence structure and grammar.



Raw	Wedding Day Blackboard	Minecraft Gets an 'Honest Game Trailer'	""""2016 Bloomer 4H w/ 19'6"""" Short Wall LQ and 6' Hyd. Slide Horse Trailer""""
Synthetic	A sign that says please leave your advice and wishes.	an image of a man holding a creeper	A large horse trailer parked in the parking lot
CapsFusion	On the wedding day, there is a blackboard where guests are encouraged to leave their advice and wishes.	An 'Honest Game Trailer' for Minecraft has been released, featuring an image of a man holding a creeper.	In the parking lot, there is a 2016 Bloomer 4H horse trailer with a 19'6" Short Wall LQ and a 6' Hyd. Slide.
ChatGPT	On the wedding day, there is a blackboard with a sign inviting guests to leave their advice and well wishes.	Minecraft receives an 'Honest Game Trailer,' featuring an image of a man holding a creeper.	In the parking lot, there is a 2016 Bloomer 4H horse trailer with a 19'6" short wall living quarters and a 6' hydraulic slide.
LaCLIP	Romantic vintage style photo of two people on their wedding day, with a chalkboard message in the background	It's a wonder Minecraft even has a story, seeing as how you spend half your time digging endlessly downwards.	"2016 Bloomer 4H w/ 19'6"""" Short Wall LQ and 6' Hyd. Slide Horse Trailer""""

Figure 6. Comparison among CAPSFUSION-LLaMA, ChatGPT, and LaCLIP. CAPSFUSION-LLaMA performs on par with ChatGPT on the caption fusion task, while LaCLIP suffers severe hallucination because only raw text is considered (hallucinations are highlighted in red in image 1 and 2). LaCLIP also fails when the raw caption is too noisy, while CAPSFUSION-LLaMA and ChatGPT can extract useful information from noise (image 3).

Synthetic captions exhibit impeccable sentence structure but often lack in-depth real-world details and may contain false information.

This nuanced description helps ChatGPT make informed decisions during fusion. Finally, we outline our expectations for the output captions with the following directive:

Ensure a well-structured sentence while retaining the detailed real-world information provided in the raw caption.

This guideline succinctly encapsulates the desired characteristics of the generated captions.

In our experimentation, we observe that in a portion of samples, ChatGPT resorts to a straightforward concatenation of the raw and synthetic captions for fusion. To address this, we explicitly instruct ChatGPT to *avoid simply concatenating two sentences*, a directive we have found highly effective in mitigating this issue. The full instruction template is presented in Sec. 7.

During human evaluation, ChatGPT is shown to be exceptionally effective at this caption fusion task. Examples are provided in the fourth row of Fig. 6. We acquired 1 million fused captions using the `gpt-3.5-turbo` API.

Refinement Model with Fused Caption. Although ChatGPT is effective, time and computational costs are pro-

hibitive. For scaling, we opt to employ LLaMA-2 [53], a state-of-the-art open-source LLM. We finetune the 13B version of LLaMA-2 specifically for the task of caption fusion, using triplets obtained from ChatGPT. These triplets consist of raw and synthetic captions as inputs, with CAPSFUSION captions as the target outputs. Training hyperparameters can be found in Sec. 8. The finetuned model, referred to as CAPSFUS-LLaMA, is rigorously evaluated through human evaluation on 100 validation cases. The evaluation results are presented in Tab. 2, revealing that the performance of the finetuned CAPSFUS-LLaMA performs on par with ChatGPT, with 80 out of 100 samples performing equally or better. LaCLIP [16] also leverages LLM for enhancing image-text captions, but simply asks LLM to rewrite raw captions. Qualitative comparisons among LaCLIP, CAPSFUSION, and ChatGPT are illustrated in Figure 6. Notably, LaCLIP tends to hallucinate information not present in the associated image, due to the absence of detailed visual information represented in the raw captions. On the other hand, CAPSFUS-LLaMA exhibits outputs similar to ChatGPT and delivers exceptional performance.

Large-scale Caption Fusion. The trained CAPSFUS-LLaMA, being as effective as ChatGPT, now possesses the ability to organically fuse and harness raw and synthetic captions in a manner that is both scalable and highly effec-

Scale	Captions	COCO		NoCaps		TextCaps		Flickr30K	
		SPICE	CIDEr	SPICE	CIDEr	SPICE	CIDEr	SPICE	CIDEr
10M	Raw [45]	15.5	75.1	9.0	64.0	10.5	46.4	13.6	54.4
	Synthetic [1]	19.8	102.5	11.7	84.2	12.7	42.3	15.0	63.9
	Language Rewrites [16]	14.6	71.6	8.6	59.0	9.3	38.3	11.6	49.0
	Mixing (Raw & Syn.) [16, 39]	17.9	90.5	10.6	76.7	12.4	51.7	15.1	64.0
	Mixing (Raw & LR) [16]	15.0	72.6	9.0	61.1	10.3	44.6	12.2	51.7
	CAPSFUSION	20.7_(+0.9)	107.7_(+5.2)	12.6_(+0.9)	92.4_(+8.2)	13.9_(+1.2)	56.3_(+4.6)	15.9_(+0.8)	68.4_(+4.4)
50M	Raw [45]	16.4	81.0	9.7	68.4	11.7	55.2	14.3	60.3
	Synthetic [1]	19.2	100.9	11.5	82.5	13.2	46.7	14.3	60.2
	Mixing (Raw & Syn.) [16, 39]	18.5	93.3	10.9	79.7	12.7	55.5	15.1	64.6
	CAPSFUSION	21.3_(+2.1)	112.4_(+11.5)	13.6_(+2.1)	99.2_(+16.7)	14.9_(+1.7)	62.7_(+7.2)	16.9_(+1.8)	74.5_(+9.9)
100M	Raw [45]	17.1	85.5	10.1	72.8	12.3	59.6	14.6	62.2
	Synthetic [1]	18.5	96.9	11.0	81.6	13.1	46.5	13.7	57.4
	Mixing (Raw & Syn.) [16, 39]	18.0	95.0	10.5	77.9	12.3	55.1	15.0	66.5
	CAPSFUSION	21.7_(+3.2)	115.7_(+18.8)	13.5_(+2.5)	99.9_(+18.3)	15.2_(+2.1)	66.2_(+11.0)	16.8_(+1.8)	73.0_(+6.4)

Table 1. Zero-shot evaluation of models trained with different caption datasets on a broad range of image captioning benchmarks.

	ChatGPT win	CAPSFUS-LLaMA win	Similar quality	(Nearly) Identical
Number	20	15	46	19

Table 2. Human evaluation on CAPSFUS-LLaMA vs. ChatGPT over 100 validation samples.

tive. We randomly select a subset containing 127,897,754 image-text pairs from LAION-COCO [1], which contains both raw captions from the web and synthetic captions generated by BLIP [34]. Subsequently, we apply CAPSFUS-LLaMA to organically integrate the captions of these image-text pairs. This process costs about 12 days using 128 A100-40G GPUs. After filtering with heuristic rules, we retain a total of 120,724,312 image-text pairs, which we term as the CAPSFUS-120M dataset.

CAPSFUS-120M Dataset. Tab. 3 provides a comparison of CAPSFUS-120M with existing image-text datasets. We compute the number of unique trigrams and the average length of these captions (word instead of token as unit) in each dataset. Notably, CAPSFUS-120M exhibits the highest count of unique trigrams and the longest average sentence length, underscoring superb diversity within its captions. In contrast, synthetic captions (LAION-COCO) exhibit a considerably lower number of trigrams, signifying a notable lack of language complexity.

4. Experiments

We present a comprehensive analysis of different caption datasets. Extensive experiments show that CAPSFUSION exhibits all-around superiority over existing image-text pair datasets, in terms of effectiveness, efficiency, world knowledge depth, and scalability.

Datasets	# Unique Trigrams	Avg. Length
LAION-2B	5.51×10^7	10.95
LAION-COCO	1.00×10^7	8.99
La-CLIP	5.46×10^7	14.63
CAPSFUS-120M	7.13×10^7	22.74

Table 3. Statistics of different caption datasets (on a randomly selected 10 million subset of CAPSFUS-120M images).

4.1. Setup

For a fair comparison, we compare CAPSFUSION with other caption datasets under the same set of images from LAION-COCO [1], isolating caption quality as the only varying factor. Experiments are conducted across three scales: 10, 50 and 100 million image-text pairs.

Model Architecture. We adopt the most prevalent LMM architecture, consisting of three components: an LLM, a vision encoder, and a vision-language bridging module. We use LLaMA-2-7B [53] and EVA-01-CLIP-g [17, 50] to initialize the LLM and vision encoder modules, respectively. For the bridging module, we follow Emu [51] to use a randomly initialized Causal Transformer to bridge the vision and language modalities. This module transforms bi-directional image representations from the vision encoder into a causal sequence that aligns better to the nature of LLMs, which excel at modeling causal sequences in an autoregressive fashion. The LLM and vision encoder are frozen during training to save computation cost following [35], and only the bridging module is tuned.

Training Schedule. The training schedule is set as the same for all compared captions. For each evaluation scale, we train the model for 1 epoch. This practice follows Datacomp [19], a benchmark for evaluating image-text pair




Figure 7. Comparison of scalability and sample efficiency across different datasets.

datasets on CLIP training. The peak learning rate is $3e-4$, with the initial 2,000 (100M) / 1,000 (50M) / 500 (10M) steps as warm-up, after which the learning rate decreases to $3e-5$ with a cosine learning rate decay schedule. Batch size is set to 8192 for all scales. Detailed training hyperparameters can be found in Sec. 8. The 100M scale training costs 40 hours with 16 A800-80G GPUs.

Baselines. We establish two baselines using raw captions from LAION-2B [45] and synthetic captions from LAION-COCO [1]. Additionally, two state-of-the-art methods for improving image-text pairs in CLIP training are evaluated: language rewrites (LaCLIP [16]) and random mixing [16, 39]. For [16], we adopt their in-context strategy and employ LLaMA-2-7B to rewrite 10M captions for comparison, taking 30 hours with 8 A100-40G GPUs. For random mixing [16, 39], we set the mixing ratio of two types of captions as 1:1 [16] and do not tune this ratio as in [39].

Evaluation. We comprehensively assess the performance of LMMs across a wide range of evaluation benchmarks. These benchmarks encompass both traditional benchmarks and recently introduced assessments, including COCO [12], NoCaps [2], TextCaps [48], Flickr30k [41], and SEED-Bench [32]. For image captioning tasks, we employ SPICE [4] and CIDEr [54] metrics. For the comprehensive SEED-Bench in the form of multiple-choice questions, we evaluate LMMs using accuracy.

4.2. Model Performance

The performances of models trained with different captions on COCO, NoCaps, TextCaps, and Flickr30K benchmarks are presented in Tab. 1. We observe that CAPSFUSION outperforms all baseline captions in all settings by a large margin, across all datasets evaluated. For example, on the 100M scale, CAPSFUSION surpasses the best baseline by a substantial margin, achieving 18.8 and 18.3 CIDEr score improvements on COCO and NoCaps, respectively.

Rewriting Captions Fails at Image Captioning. On the 10M scale, our examination reveals that Language Rewrites captions [16], generated through the process of rewriting raw captions, fail to achieve decent performance. This can be attributed to the severe hallucination issue we observed in the rewrites captions (Fig. 6), which introduces extrane-

ous text that is irrelevant to the content depicted in the accompanying images. The underlying cause of the hallucination phenomenon can be traced back to the input data, which consists solely of noisy raw captions, providing a suboptimal starting point for the rewriting process.

Mixing Captions does not Bring Consistent Gains. Another notable observation is that mixing captions cannot yield better performance. For instance, on the 10M-scale over COCO benchmark, mixing raw and LR captions (72.62 CIDEr and 15.01 SPICE scores) achieves a median performance between Raw (75.13 CIDEr and 15.48 SPICE) and LR (71.61 CIDEr, 14.6 SPICE) captions. This finding is contrarian to the observation in CLIP training [16, 31, 39], where mixing raw and generated captions has proven to be a strong strategy for enhancing CLIP performance, with raw captions being an indispensable component [16, 31]. In contrast, our experiments show that in LMMs training, the exclusive use of a single caption type (CAPSFUSION) can outperform both raw and synthetic captions.

Synthetic Captions Shout at Small Scale. A noteworthy observation is that synthetic caption demonstrates exceptional results on the 10M dataset (102.5 COCO CIDEr), while exhibiting inferior performance (96.93 COCO CIDEr) on the larger-scale 100M dataset. This aligns with our earlier observation of the *Scalability Deficiency* issue in synthetic captions, a potential threat to the effective training of LMMs. But even at small scales, the effectiveness of synthetic captions consistently falls behind that of CAPSFUSION across all datasets.

4.3. Sample Efficiency

In addition to comparing performance across different dataset scales, we probe deeper into training sample efficiency. In Tab. 1, we find that with only 10M image-text pairs, CAPSFUSION captions outperform other captions with much larger scale (50M and 100M), demonstrating exceptional sample efficiency. We visualize the updates of evaluation metrics on NoCaps, TextCaps, Flickr30K, and COCO benchmarks when the number of seen training samples increases from 0 to 100 million image-text pairs, presented in Fig. 7. The horizontal grey dashed lines approximately represent the best-saturated performance of baseline

Captions	Scene U	Inst Iden	Inst Loc	Inst Attr	Inst Cnt	Spatial Rel	Inst Inter	Vis Reason	Text Rec	Total
Raw [45]	57.9	51.2	39.8	47.7	44.6	35.3	47.4	48.6	34.1	48.7
Synthetic [1]	52.7	48.9	36.7	42.2	35.7	34.5	48.4	35.0	12.9	43.2
CAPSFUSION	58.8	52.7	41.0	48.0	46.3	35.9	57.7	47.1	20.0	49.8

Table 4. Zero-shot evaluation of models trained with different caption datasets on SEED-Bench.

captions when trained with 100 million image-text pairs. The vertical dashed line reveals the number of samples employed by CAPSFUSION to achieve a similar level of performance as the best-performing baseline captions. It is worth noting that CAPSFUSION attains the same level of performance as the best baseline captions with only 6M, 9M, 6M, and 8M samples for NoCaps, TextCaps, Flickr30K, and COCO captions, respectively. This achievement underscores CAPSFUSION’s ability of 11-16 times speedup and demonstrates its superior sample efficiency.

4.4. Scalability Analysis

Scalability stands as a crucial attribute in large model training. Our investigation reveals that synthetic captions, among all the caption types considered, exhibit the worst scalability. This can be observed from Fig. 7 (a), (b), and (d), wherein the blue lines exhibit early saturation with a mere 30 million image-text pairs. Subsequently, their performance gradually deteriorates. In contrast, raw caption (orange lines) displays commendable scalability, with its performance showing a consistent upward trajectory. However, the inherent high noise level in raw caption hampers its ability to achieve strong performance. CAPSFUSION caption (red lines) exhibits remarkable scalability on all datasets, outperforming both synthetic and raw captions by a substantial margin throughout the entire scale.

Note: Our investigation reveals that synthetic captions have severe scalability limitations and typically saturate with only 30 million pairs, after which *more computation imposes an adverse impact on model performance*. However, current synthetic caption datasets used are typically much larger in scale (*e.g.* 600M in LAION-COCO). We hope our findings raise concerns about the efficiency issue in training LMMs with such massive synthetic caption datasets.

4.5. Further Evaluation on SEED-Bench

Recently, new comprehensive benchmarks are proposed for thorough evaluations of LMMs on granular functionalities [6, 7, 18, 37, 59, 64]. We evaluate our proposed model on a representative benchmark, SEED-Bench [32], over its 9 image-text tasks (dataset details can be found in Sec. 9.). Results are presented in Tab. 4. We find CAPSFUSION outperforms raw and synthetic captions in 7 out of 9 evaluated tasks, which underscores the remarkable capabilities of CAPSFUSION in instance counting, instance interaction, scene understanding and other multimodal functionalities.

Method	COCO [12]	SEED-Bench [32]	MMLU [24]
LLaMA-2-7B	-	-	45.78
Raw	74.9	48.5	43.7
Synthetic	55.5	40.2	42.9
CAPSFUSION	111.3	51.5	<u>44.1</u>

Table 5. Model performance with LLM tuned over different caption datasets.

4.6. Qualitative Evaluation on World Knowledge

In Fig. 3 and Fig. 9 (Appendix), we provide a qualitative evaluation on the outputs generated by models trained with different datasets. The first row is the input image with text prompt, the lower three rows show the outputs from models trained on raw, synthetic, and CAPSFUSION captions.

We observe that models trained on raw and CAPSFUSION captions exhibit rich real-world knowledge, able to identify celebrities (Fig. 3 image 1), recognize famous artworks (Fig. 9 image 2), attribute literature works to their authors (Fig. 3 image 2), and pinpoint the location where the specific event occurred (Fig. 3 image 3). Models trained on synthetic captions totally lost such capabilities.

4.7. Effects when Firing LLM

We investigate the impact of different training captions when firing the LLM. We conduct experiments at the 10M scale on COCO [12], SEED [32], MMLU [24]. Results are summarized in Tab. 5. Notably, we observe a significant decline in the performance of synthetic captions. This indicates a deterioration in the LLM’s capabilities when it is trained on the simplified language of synthetic captions.

5. Conclusion

In this work, we identify severe *Scalability Deficiency* and *World Knowledge Loss* issues in LMMs trained with synthetic captions. On the other hand, web-based image-text pairs possess rich world knowledge but are too noisy to achieve decent performance. We thus propose CAPSFUSION, an advanced framework to generate high-quality captions in a scalable and effective manner. The resulting CAPSFUSION-120M dataset exhibits all-around superiority over existing image-text datasets, which poses CAPSFUSION as a promising framework to generate large-scale high-quality image-text data for scalable LMM training.

References

- [1] Laion coco: 600m synthetic captions from laion2b-en. <https://laion.ai/blog/laion-coco/>. 1, 2, 4, 6, 7, 8
- [2] Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. Nocaps: Novel object captioning at scale. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8948–8957, 2019. 7
- [3] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. 1, 4
- [4] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part V 14*, pages 382–398. Springer, 2016. 7
- [5] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023. 1
- [6] Shuai Bai, Shusheng Yang, Jinze Bai, Peng Wang, Xingxuan Zhang, Junyang Lin, Xinggang Wang, Chang Zhou, and Jingren Zhou. Touchstone: Evaluating vision-language models by language models. *arXiv preprint arXiv:2308.16890*, 2023. 8
- [7] Yonatan Bitton, Hritik Bansal, Jack Hessel, Rulin Shao, Wanrong Zhu, Anas Awadalla, Josh Gardner, Rohan Taori, and Ludwig Schimdt. Visit-bench: A benchmark for vision-language instruction following inspired by real-world use. *arXiv preprint arXiv:2308.06595*, 2023. 8
- [8] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 3
- [9] Chi Chen, Ruoyu Qin, Fuwen Luo, Xiaoyue Mi, Peng Li, Maosong Sun, and Yang Liu. Position-enhanced visual instruction tuning for multimodal large language models. *arXiv preprint arXiv:2308.13437*, 2023. 3
- [10] Delong Chen, Jianfeng Liu, Wenliang Dai, and Baoyuan Wang. Visual instruction tuning with polite flamingo. *arXiv preprint arXiv:2307.01003*, 2023.
- [11] Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing multimodal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023. 3
- [12] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 1, 7, 8
- [13] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointly-scaled multilingual language-image model. *arXiv preprint arXiv:2209.06794*, 2022. 3
- [14] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven C. H. Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. *CoRR*, abs/2305.06500, 2023. 1, 4
- [15] Runpei Dong, Chunrui Han, Yuang Peng, Zekun Qi, Zheng Ge, Jinrong Yang, Liang Zhao, Jianjian Sun, Hongyu Zhou, Haoran Wei, et al. Dreamllm: Synergistic multimodal comprehension and creation. *arXiv preprint arXiv:2309.11499*, 2023. 1, 4
- [16] Lijie Fan, Dilip Krishnan, Phillip Isola, Dina Katabi, and Yonglong Tian. Improving clip training with language rewrites. *arXiv preprint arXiv:2305.20088*, 2023. 2, 3, 5, 6, 7
- [17] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. Eva: Exploring the limits of masked visual representation learning at scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19358–19369, 2023. 6
- [18] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin, Jinrui Yang, Xiawu Zheng, et al. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*, 2023. 8
- [19] 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 arXiv:2304.14108*, 2023. 2, 3, 6
- [20] Peng Gao, Jiaming Han, Renrui Zhang, Ziyi Lin, Shijie Geng, Aojun Zhou, Wei Zhang, Pan Lu, Conghui He, Xianguyue Yue, et al. Llama-adapter v2: Parameter-efficient visual instruction model. *arXiv preprint arXiv:2304.15010*, 2023. 3
- [21] Yuying Ge, Yixiao Ge, Ziyun Zeng, Xintao Wang, and Ying Shan. Planting a seed of vision in large language model. *arXiv preprint arXiv:2307.08041*, 2023. 4
- [22] Yuying Ge, Sijie Zhao, Ziyun Zeng, Yixiao Ge, Chen Li, Xintao Wang, and Ying Shan. Making llama see and draw with seed tokenizer. *arXiv preprint arXiv:2310.01218*, 2023. 1, 4
- [23] Tao Gong, Chengqi Lyu, Shilong Zhang, Yudong Wang, Miao Zheng, Qian Zhao, Kuikun Liu, Wenwei Zhang, Ping Luo, and Kai Chen. Multimodal-gpt: A vision and language model for dialogue with humans. *arXiv preprint arXiv:2305.04790*, 2023. 3
- [24] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2020. 8
- [25] Shaohan Huang, Li Dong, Wenhui Wang, Yaru Hao, Saksham Singhal, Shuming Ma, Tengchao Lv, Lei Cui, Owais Khan Mohammed, Qiang Liu, et al. Language is

- not all you need: Aligning perception with language models. *arXiv preprint arXiv:2302.14045*, 2023. 4
- [26] Yupan Huang, Zaiqiao Meng, Fangyu Liu, Yixuan Su, Nigel Collier, and Yutong Lu. Sparkles: Unlocking chats across multiple images for multimodal instruction-following models. *arXiv preprint arXiv:2308.16463*, 2023. 3
- [27] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021. 1
- [28] Yang Jin, Kun Xu, Liwei Chen, Chao Liao, Jianchao Tan, Bin Chen, Chenyi Lei, An Liu, Chengru Song, Xiaoqiang Lei, et al. Unified language-vision pretraining with dynamic discrete visual tokenization. *arXiv preprint arXiv:2309.04669*, 2023. 4
- [29] Jing Yu Koh, Daniel Fried, and Ruslan Salakhutdinov. Generating images with multimodal language models. *arXiv preprint arXiv:2305.17216*, 2023. 4
- [30] Jing Yu Koh, Ruslan Salakhutdinov, and Daniel Fried. Grounding language models to images for multimodal generation. *arXiv preprint arXiv:2301.13823*, 2023. 4
- [31] Zhengfeng Lai, Haotian Zhang, Wentao Wu, Haoping Bai, Aleksei Timofeev, Xianzhi Du, Zhe Gan, Jiulong Shan, Chen-Nee Chuah, Yinfei Yang, et al. From scarcity to efficiency: Improving clip training via visual-enriched captions. *arXiv preprint arXiv:2310.07699*, 2023. 3, 7
- [32] Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. Seed-bench: Benchmarking multimodal llms with generative comprehension. *arXiv preprint arXiv:2307.16125*, 2023. 1, 7, 8
- [33] Bo Li, Yuanhan Zhang, Liangyu Chen, Jinghao Wang, Jingkang Yang, and Ziwei Liu. Otter: A multi-modal model with in-context instruction tuning. *arXiv preprint arXiv:2305.03726*, 2023. 3
- [34] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. 1, 2, 4, 6
- [35] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 1, 4, 6
- [36] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. 1, 4
- [37] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*, 2023. 8
- [38] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 1
- [39] Thao Nguyen, Samir Yitzhak Gadre, Gabriel Ilharco, Sewoong Oh, and Ludwig Schmidt. Improving multi-modal datasets with image captioning. *arXiv preprint arXiv:2307.10350*, 2023. 2, 3, 6, 7
- [40] Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint arXiv:2306.14824*, 2023. 3
- [41] 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 *Proceedings of the IEEE international conference on computer vision*, pages 2641–2649, 2015. 7
- [42] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 4
- [43] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 1, 3
- [44] Noam Rotstein, David Bensaid, Shaked Brody, Roy Ganz, and Ron Kimmel. Fusecap: Leveraging large language models for enriched fused image captions. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5689–5700, 2024. 3
- [45] 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. 1, 6, 7, 8
- [46] John Schulman, Barret Zoph, Christina Kim, Jacob Hilton, Jacob Menick, Jiayi Weng, Juan Felipe Ceron Uribe, Liam Fedus, Luke Metz, Michael Pokorny, et al. Chatgpt: Optimizing language models for dialogue. *OpenAI blog*, 2022. 2
- [47] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018. 1, 2
- [48] Oleksii Sidorov, Ronghang Hu, Marcus Rohrbach, and Amanpreet Singh. Textcaps: a dataset for image captioning with reading comprehension. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 742–758. Springer, 2020. 7
- [49] Yixuan Su, Tian Lan, Huayang Li, Jialu Xu, Yan Wang, and Deng Cai. Pandagpt: One model to instruction-follow them all. *arXiv preprint arXiv:2305.16355*, 2023. 4
- [50] Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. Eva-clip: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*, 2023. 1, 3, 6
- [51] Quan Sun, Qiyang Yu, Yufeng Cui, Fan Zhang, Xiaosong Zhang, Yueze Wang, Hongcheng Gao, Jingjing Liu, Tiejun

- Huang, and Xinlong Wang. Generative pretraining in multimodality. *arXiv preprint arXiv:2307.05222*, 2023. 1, 4, 6
- [52] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 1
- [53] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 1, 2, 3, 5, 6
- [54] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575, 2015. 7
- [55] Guangzhi Wang, Yixiao Ge, Xiaohan Ding, Mohan Kankanhalli, and Ying Shan. What makes for good visual tokenizers for large language models? *arXiv preprint arXiv:2305.12223*, 2023. 3
- [56] Weiyun Wang, Min Shi, Qingyun Li, Wenhai Wang, Zhenhang Huang, Linjie Xing, Zhe Chen, Hao Li, Xizhou Zhu, Zhiguo Cao, et al. The all-seeing project: Towards panoptic visual recognition and understanding of the open world. *arXiv preprint arXiv:2308.01907*, 2023. 3
- [57] Xinyu Wang, Bohan Zhuang, and Qi Wu. Switchgpt: Adapting large language models for non-text outputs. *arXiv preprint arXiv:2309.07623*, 2023. 4
- [58] Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. Next-gpt: Any-to-any multimodal llm. *arXiv preprint arXiv:2309.05519*, 2023. 4
- [59] Peng Xu, Wenqi Shao, Kaipeng Zhang, Peng Gao, Shuo Liu, Meng Lei, Fanqing Meng, Siyuan Huang, Yu Qiao, and Ping Luo. Lvlm-ehub: A comprehensive evaluation benchmark for large vision-language models. *arXiv preprint arXiv:2306.09265*, 2023. 8
- [60] Zhewei Yao, Xiaoxia Wu, Conglong Li, Minjia Zhang, Heyang Qi, Olatunji Ruwase, Ammar Ahmad Awan, Samyam Rajbhandari, and Yuxiong He. Deepspeed-visualchat: Multi-round multi-image interleave chat via multi-modal causal attention. *arXiv preprint arXiv:2309.14327*, 2023. 3
- [61] Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023.
- [62] Haoxuan You, Haotian Zhang, Zhe Gan, Xianzhi Du, Bowen Zhang, Zirui Wang, Liangliang Cao, Shih-Fu Chang, and Yinfei Yang. Ferret: Refer and ground anything anywhere at any granularity. *arXiv preprint arXiv:2310.07704*, 2023. 3
- [63] Lili Yu, Bowen Shi, Ramakanth Pasunuru, Benjamin Muller, Olga Golovneva, Tianlu Wang, Arun Babu, Binh Tang, Brian Karrer, Shelly Sheynin, et al. Scaling autoregressive multimodal models: Pretraining and instruction tuning. *arXiv preprint arXiv:2309.02591*, 2023. 4
- [64] Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. Mm-vet: Evaluating large multimodal models for integrated capabilities. *arXiv preprint arXiv:2308.02490*, 2023. 8
- [65] Yan Zeng, Hanbo Zhang, Jiani Zheng, Jiangnan Xia, Guoqiang Wei, Yang Wei, Yuchen Zhang, and Tao Kong. What matters in training a gpt4-style language model with multimodal inputs? *arXiv preprint arXiv:2307.02469*, 2023. 3
- [66] Ao Zhang, Hao Fei, Yuan Yao, Wei Ji, Li Li, Zhiyuan Liu, and Tat-Seng Chua. Transfer visual prompt generator across llms. *arXiv preprint arXiv:2305.01278*, 2023. 4
- [67] Shilong Zhang, Peize Sun, Shoufa Chen, Min Xiao, Wenqi Shao, Wenwei Zhang, Kai Chen, and Ping Luo. Gpt4roi: Instruction tuning large language model on region-of-interest. *arXiv preprint arXiv:2307.03601*, 2023. 3
- [68] Bo Zhao, Boya Wu, and Tiejun Huang. Svit: Scaling up visual instruction tuning. *arXiv preprint arXiv:2307.04087*, 2023.
- [69] Yang Zhao, Zhijie Lin, Daquan Zhou, Zilong Huang, Jiashi Feng, and Bingyi Kang. Bubogpt: Enabling visual grounding in multi-modal llms. *arXiv preprint arXiv:2307.08581*, 2023. 3
- [70] Kaizhi Zheng, Xuehai He, and Xin Eric Wang. Minigpt-5: Interleaved vision-and-language generation via generative vokens. *arXiv preprint arXiv:2310.02239*, 2023. 4
- [71] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023. 4