

FUSECAP: Leveraging Large Language Models for Enriched Fused Image Captions

Noam Rotstein* David Bensaïd* Shaked Brody Roy Ganz Ron Kimmel
 Technion - Israel Institute of Technology

*Indicates equal contribution.

Abstract

The advent of vision-language pre-training techniques enhanced substantial progress in the development of models for image captioning. However, these models frequently produce generic captions and may omit semantically important image details. This limitation can be traced back to the image-text datasets; while their captions typically offer a general description of image content, they frequently omit salient details. Considering the magnitude of these datasets, manual reannotation is impractical, emphasizing the need for an automated approach. To address this challenge, we leverage existing captions and explore augmenting them with visual details using “frozen” vision experts including an object detector, an attribute recognizer, and an Optical Character Recognizer (OCR). Our proposed method, FUSECAP, fuses the outputs of such vision experts with the original captions using a large language model (LLM), yielding comprehensive image descriptions. We automatically curate a training set of 12M image-enriched caption pairs. These pairs undergo extensive evaluation through both quantitative and qualitative analyses. Subsequently, this data is utilized to train a captioning generation BLIP-based model. This model outperforms current state-of-the-art approaches, producing more precise and detailed descriptions, demonstrating the effectiveness of the proposed data-centric approach. We release this large-scale dataset of enriched image-caption pairs for the community.

1. Introduction

The generation of image captions that effectively capture essential descriptive elements has been a longstanding goal in computer vision [7, 32, 34, 36, 49, 55]. In recent years, image captioning tasks [3, 44] have gained significant research attention and interest due to the success of Vision Language (VL) models. This achievement mainly stems from the ability to efficiently harness the massive amount of image-caption pairs accessible online, using Vi-



Figure 1. **FUSECAP captions.** An illustration comparing our FUSECAP enriched captions with the original ground-truth captions before the fusing process. The examples are from COCO, SBU, CC, and CC12 datasets, displayed from top to bottom.

sion Language Pre-training (VLP) [15, 40, 53], followed by task-specific fine-tuning. However, despite remarkable advancements in image captioning, current state-of-the-art models [14, 28, 37, 38, 68, 70, 73, 81] produce captions that often overlook key semantic elements. As images are rich

sources of information containing intricate and complex content, providing precise descriptions requires highly detailed textual captions.

We hypothesize that the current unsatisfactory captioning results are attributed to the image-caption datasets used for training. Captions in these datasets frequently fail to capture essential elements within images and often omit fine details. For example, consider the original caption of the top image in Figure 1. The caption is missing details such as the laptop and cat. Since these datasets contain a massive number of image-caption pairs, manual re-annotation is unfeasible.

In this work, our primary objective is to develop a framework that produces richer, more accurate captions for images. To achieve this, we introduce FUSECAP, a novel approach designed to automatically augment captions in existing datasets, thereby enhancing training data of the model. This contrasts with methods that primarily focus on improving the caption generator model architecture and resonates with the recently surveyed data-centric artificial intelligence (AI) paradigm [79, 80], which underscores the significance of improving both the quality and quantity of data, rather than merely concentrating on the advancement of model design. We leverage the capabilities of *vision experts* such as object detectors [4], attribute recognizers [82], and Optical Character Recognition (OCR) models [1, 2, 8]. The visual information extracted by these models is intended to provide complementary details to the original simplistic captions. By harnessing the reasoning capabilities of a dedicated LLM, the outputs from the vision experts are fused with the original image caption. This results in a coherent, meaningful natural description of images that is more comprehensive and detailed than the original caption, as illustrated in Figure 1. Specifically, we leverage ChatGPT [11] to generate “fusing” examples which are then used to fine-tune a pre-trained FLAN-T5 model [54]. We apply this method to enrich captions of a human-annotated dataset (COCO [44]) and large-scale datasets collected from the web (CC12 [12], CC [58], and SBU [65]). This process produces an enriched collection consisting of 12M image-text pairs. To confirm the quality of our generated dataset, we first show that humans favor the fused captions, perceiving them as more descriptive and accurate than the original captions. Furthermore, we demonstrate that the fused captions score higher on CLIPScore [26] – a reference-free metric that evaluates text-image alignment without using reference captions — compared to the original ones. To further emphasize their effectiveness, we also assess the captions through image-to-text and text-to-image retrieval tasks.

To illustrate the benefits of the proposed data-centric approach we capitalize on these fused captions to train a caption generator. We use the augmented dataset both for pre-training and for the fine-tuning of an image captioning BLIP model [38]. Despite having fewer parameters and using less

training data, our model surpasses existing state-of-the-art methods [37, 38, 46, 69, 70] in generating comprehensive captions. This superiority is evident both in its higher CLIP-Score and its improved performance in the retrieval tasks. The performance advantage is further illustrated by numerous examples.

Our contributions are as follows:

- Introducing FUSECAP- a novel approach to automatically enrich existing image-captions datasets by fusing outputs from visual experts using an LLM.
- Providing a large dataset of 12M caption-enriched text-image pairs for future research.
- Showcasing that an enriched dataset leads to models capable of generating detailed captions that effectively incorporate previously overlooked key semantic elements.

2. Related Efforts

Image-Caption Generation. Image captioning has been a widely researched topic at the intersection of computer vision and natural language processing. Early strategies for caption generation made use of retrieval-based methods [23, 27] and template-based methods [35], which were limited in their expressiveness [6]. The advent of deep learning marked a shift in this field, as multimodal neural networks enabled the generation of higher quality captions [31]. Subsequently, the encoder-decoder framework, which essentially translates an image into a sentence, became one of the favored approaches [20, 66]. Later advancements incorporated attention mechanisms for focusing on key image aspects [76, 78]. Current best captioning techniques are transformer-based architectures [63], which combine vision [21] and language [63] transformers. The recent advancements in the effectiveness of image captioning can largely be credited to the introduction of *visual language pre-training* (VLP).

VLP uses large-scale image-text pairs to pre-train vision-text models, which are later fine-tuned for downstream tasks, with image captioning as one of the central tasks [25, 37, 38]. VLP can be categorized [13] into single-stream [16, 41] and dual-stream architectures [22, 24, 29, 53], aiming to merge visual and textual modalities into a shared embedding space. Efforts exploring integrating text generative tasks for pre-training [17, 72] have fueled the development of models like BLIP [38] and recent BLIP-2 [37]. OFA [70] proposed unifying multiple unimodal and multimodal pre-training tasks, which led to a significant performance improvement.

Dense captioning can be seen as a task related to ours in generating comprehensive text for images [30, 77]. However, it produces multiple captions for various regions within an image, as opposed to a single descriptive caption for the whole image. Previous approaches to generating more comprehensive captions have consistently identified a common

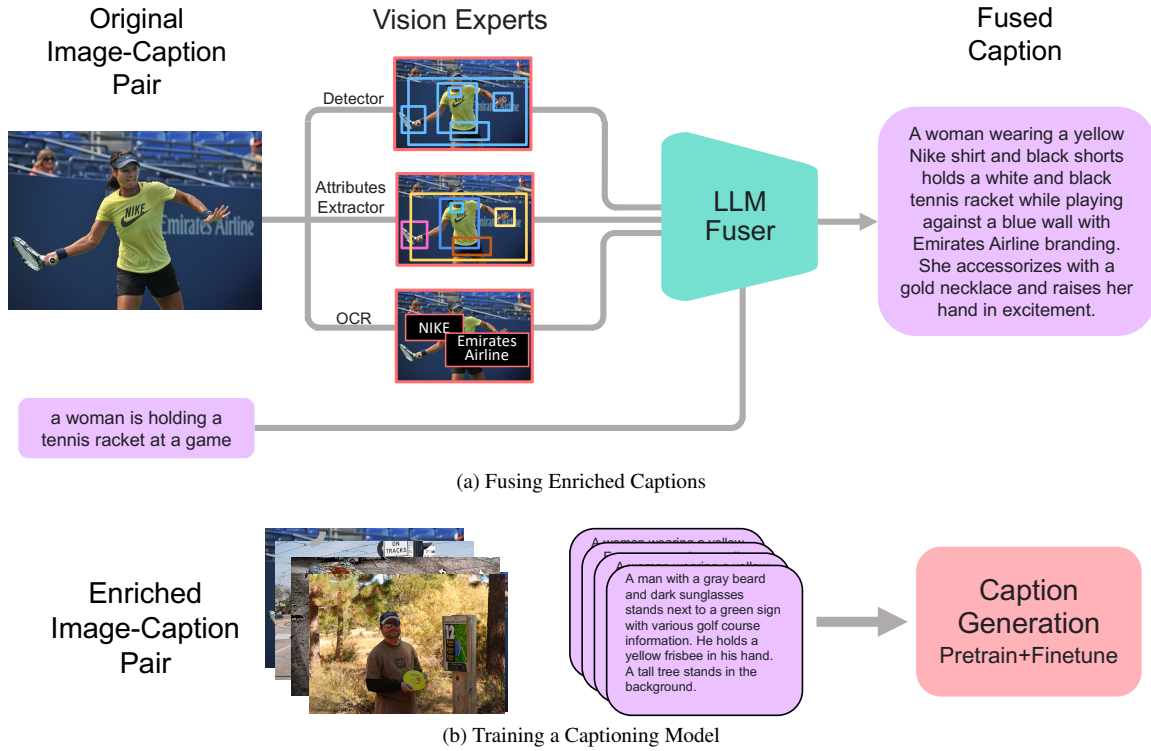


Figure 2. **Our Approach Illustration.** Figure 2a illustrates the automated process of enriching existing image captions using the proposed FUSECAP approach. Visual experts extract meaningful information from images, which is then fused with the original captions by an LLM Fuser, producing rich captions. Following this, Figure 2b illustrates the utilization of image datasets, paired with these augmented captions, in both the pre-training and fine-tuning phases of a comprehensive image-captioning model.

challenge – the unsatisfactory capability of existing image captioners of providing detailed and accurate textual descriptions of images [45, 47, 48]. These approaches focus on improving captions discriminability, whereas our method seeks to enrich captions with additional, meaningful information extracted from the images.

Image-Caption Datasets Current datasets for image-caption pairs fall into two primary categories: specifically human-annotated datasets like COCO [44], and web-crawled datasets such as CC, CC12, and SBU Captions [12, 50, 58]. The conjunction of the datasets serves as the foundation for the VLP, subsequently followed by downstream tasks fine-tuning. In general, the first category of datasets is smaller in size yet exhibits substantially lower noise levels compared to the second category. Both categories, however, are characterized by relatively short and concise captions. Specific examples of these characteristics can be seen in Figure 1.

Data-centric AI underscores the importance of refining data quality to boost model performance rather than fine-tuning model designs. It prioritizes curating, labeling, and cleaning data for superior training datasets [61] as well as automating its processing [79, 80]. This new emphasis suggests that with high-quality data, even basic algorithms can

deliver impressive results. [59] identified the need to enrich image-text datasets. To this end, they harnessed natural language inference (NLI) to fuse multiple existing ground-truth captions into a single one. However, this approach can operate only on datasets with multiple ground-truth captions and cannot be applied to large-scale ones with a single caption (e.g., CC, CC12, and SBU). In contrast, our method can be applied to any image-caption dataset.

Large Language Models Large language models (LLMs) have been shown to be effective in a wide range of tasks, including natural language inference, question answering, and code generation [18, 19, 74]. Further, LLMs such as GPT-3 [11] exhibit impressive zero and few-shot performance on a variety of tasks, including translation, text summarization, and common sense reasoning, without further fine-tuning. Few-shot abilities allow researchers to use LLMs as a tool for data generation. [10] used GPT-3 to generate instructions and edited captions dataset, which is then used to train model for image editing. [57] presented an approach for training LLMs to use external API calls. For example, when the LLM is been asked to solve a mathematical problem, it can use an API of a calculator, instead of generating the output by itself. [57] used GPT-3 few-shot ability to curate

a dataset of external API calls which is then been used to fine-tune another LLM. [52] used GPT-4 outputs to create an instruction-following dataset that can be used later to fine-tune other LLMs in a supervised learning fashion. To fuse together the original caption and the visual expert outputs, in this work, we harness the impressive zero-shot capabilities of OpenAI’s ChatGPT [11]. We employ it to generate a small “fusing” dataset. To establish an open-source framework that scales cost-effectively, we fine-tune Flan-T5 [19], a widely recognized open-source LLM, using this data.

Vision Experts in VLP Several works attempted to improve VLP by incorporating object detectors or other experts [46] as part of their initialization [17, 41], architecture [46], pre-training data [67] or pre-training objectives [42, 60, 62, 82]. In image captioning, such models demonstrate limited capabilities in generating rich captions and have not fully capitalized on the information provided by vision experts. We hypothesize that this limitation stems directly from the succinct captions in existing image-text datasets. In our data-centric approach, we focus on improving such datasets and overcoming this limitation. An example that highlights our approach is the contrast with Prismer [46]. Unlike our method, Prismer augments its model architecture with object detectors and OCR for caption generation but relies on traditional captioning datasets.

A recent concurrent paper [71] proposed an LLM-based model that merges captions of image segments into text. In contrast, we adopt the LLM from a data-centric perspective to enhance existing caption datasets, an approach that is followed by extensive evaluations for quality and consistency.

3. Fusing Enriched Captions

In this section, we introduce FUSECAP, which is illustrated in Figure 2a. This novel strategy is designed to automatically augment existing image captions by integrating important details and inter-object relations within the image. These details are often disregarded in traditional image captioning datasets. First, we elaborate on our use of pre-trained vision models, referred to as *vision experts*, for extracting relevant visual information from images. We then detail how the information gathered from these expert models is subsequently fused with the original caption through a fine-tuned LLM, resulting in the enhanced captions.

3.1. Vision Expert Models

To enrich the information found in the original caption, we employ the following vision experts:

Object Detection A key visual expert we rely on is an object detector model. Following the approach proposed in [82], we utilize a Faster-RCNN [56] with a ResNeXt152 [75] backbone. This model is initially pre-trained on several detection datasets and then fine-tuned on the Visual

Genome (VG) dataset [33]. The data used comprises more than 100K images with 1.6K classes, enabling strong generalization abilities. We regard all objects along with their corresponding bounding boxes as valid detections, provided they exceed a detection confidence threshold. The presence and position of objects in the scene provide essential details to complement the caption.

Attributes Prediction Beyond class identification, we derive a variety of attributes for each object in the image using features generated by our Faster-RCNN within each bounding box. A classification model proposed in [82] was trained on annotations from the VG dataset, which covers a broad array of 400 distinct attributes. These attributes encapsulate various aspects of the objects, including size, condition, and color. For each object, we only consider attributes predicted with a confidence level above the threshold.

Text Detection and Recognition Text within images often contains critical contextual information, leading to a line of different text-image tasks [9, 25]. To incorporate textual information, we utilize robust pre-trained OCR models to detect and extract characters. We first identify text within a scene with CRAFT [5], a robust scene-text detector. We then apply Parseq [8], a state-of-the-art scene text recognizer, to decode the text within the bounding boxes generated by the text detector. To avoid contamination, we do not apply OCR methods on datasets presenting watermarks, such as the CC [58] and CC12 [12].

3.2. LLM Fuser

To generate natural and coherent captions, which are essential for the caption generation task, we utilize a large language model (LLM). We leverage a specifically fine-tuned LLM to fuse insights from various vision experts into the original caption, creating a single coherent and fluent description. The LLM’s advanced reasoning capabilities allow us to articulate the semantic relationships among objects and to seamlessly integrate the diverse knowledge provided by these experts. Consequently, the output captures the essence of the visual content while the impressive generative capacity of LLMs ensures these captions remain coherent and natural. To train the LLM, we create a small “fusing” dataset with ChatGPT [11], and then use it to fine-tune the open-source FlanT5-XL model [19].

ChatGPT Annotation We leverage the zero-shot capabilities of ChatGPT to generate a “fusing” dataset of 20K examples, which is then used to fine-tune an open-source LLM. By employing this approach, we establish a framework based on open-source models that can generate a large-scale dataset at a reasonable cost. To produce enriched captions with ChatGPT, we first extract information from the visual

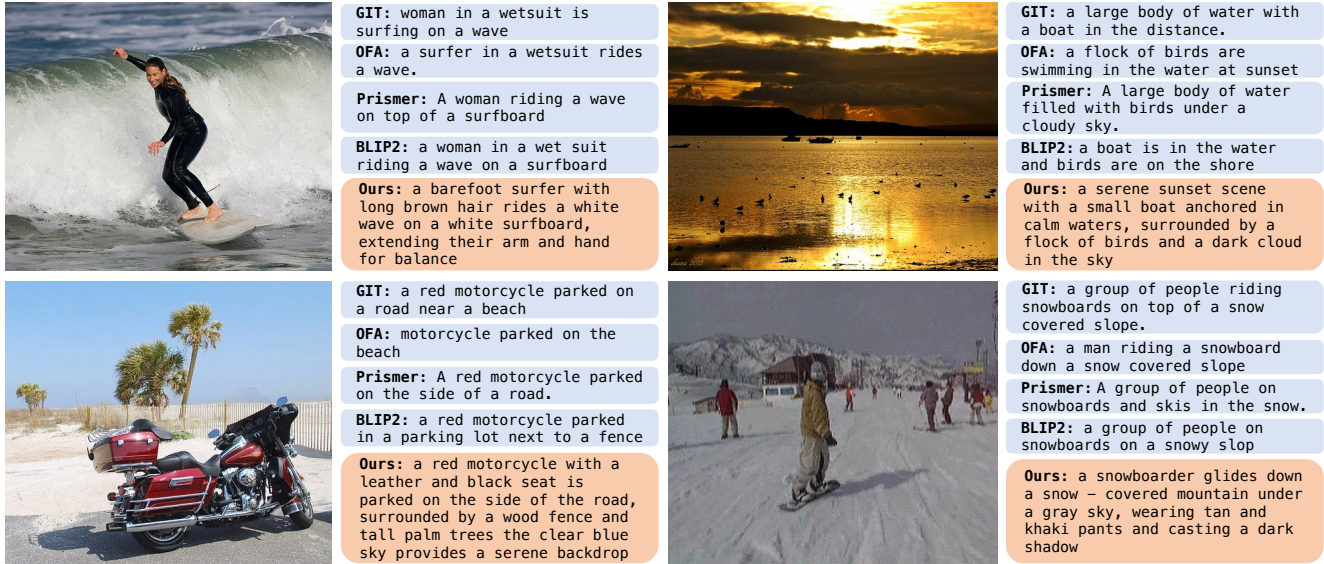


Figure 3. **Image Captioning Results.** While top-performing captioning models tend to provide concise and oversimplified captions, our model outputs rich captions that better describe the images.

experts presented in Section 3.1. The objects and their corresponding attributes and detected texts are then ordered from left to right based on their bounding boxes, providing basic spatial context. The exact prompt can be found in the appendix. We generated 20K examples of such enriched captions from CC, CC12, SBU, and COCO datasets, which served as training data for the FlanT5-XL model.

LLM Fine-tuning Using the “fusing” dataset we created, we fine-tuned Flan-T5 [19], a variant of the T5 encoder-decoder model [54], that has been extensively fine-tuned on numerous tasks to achieve exceptional performance on instruction-based tasks. Specifically, we utilized the Flan-T5-XL checkpoint and fine-tuned it to our curated fuse dataset. During the fine-tuning process, the original caption was concatenated with the output of the visual experts and served as the model input, while the enriched caption was designated as the target. The hyperparameters we used for this section can be found in the appendix. After this phase, the fine-tuned model can be used to generate an enriched captions dataset.

4. Training a Captioning Model

We apply FUSECAP, presented in Section 3, to the COCO, SBU, CC, and CC12 datasets, yielding 12 million augmented image-caption pairs. Following our data-centric approach, and to demonstrate the effectiveness of this augmented dataset, we use it to train a captioning model based on the BLIP architecture [38] (Figure 2b). Adopting the training strategy from the original BLIP paper, we first perform vision-language pre-training for 20 epochs, optimizing three objectives: (a) **Image-Text Contrastive Loss (ITC)**

that aligns visual and textual features by contrasting between matching and non-matching image-text pairs, (b) **Image-Text Matching Loss (ITM)** that classifies image-text pairs as matched or unmatched, and (c) **Generative Language Modeling Loss (LM)** that generates textual captions from images, which are compared to their ground truth captions. Following the pre-training phase, adhering to common practice in VLP [37, 38, 81], we fine-tune our model for a supplementary 5 epochs on the *enriched* COCO dataset, utilizing solely LM loss. To allow the generation of more comprehensive captions, we increase the context length over that used in the original BLIP model from 30 to 60 tokens. The hyperparameters for this section are detailed in the appendix.

5. Experiments

We evaluate two sets of image captions. The first set consists of the fused captions from the FUSECAP dataset, described in Section 3. The second set includes captions generated by the trained captioning model discussed in Section 4. The objective for both sets is to produce captions that are descriptive and accurate. Traditionally, image captioning methods are assessed with n-gram-based metrics like BLEU [51], CIDEr [64], and ROUGE [43]. These metrics compare the tested captions to reference captions, assuming the latter represents the ideal image descriptions that the tested captions aim for. However, as highlighted earlier, captions in existing datasets often fail to provide a comprehensive description of images. This implies that given our emphasis on enhancing descriptiveness, relying on broad and non-specific reference captions for evaluation would be inappropriate. Consequently, n-gram-based metrics do not

effectively measure or promote this quality [84].

We, therefore adopt CLIPScore [26], a reference-free metric that measures the alignment between textual and visual embeddings generated by a pre-trained CLIP model [53]. Since CLIPScore is not dependent on a reference caption, its score is not restricted by the descriptiveness of the original captions. Moreover, in terms of accuracy, it has been shown to exhibit a higher correlation with human judgments than n-gram-based metrics [26].

Besides evaluating enriched captions with CLIPScore, we consider the ability to perform image-text retrieval as a pertinent metric for caption evaluation. Comprehensive captions should inherently serve as distinct image descriptors, thereby improving retrieval precision. To further assess the quality of captions in the FUSECAP dataset, we conducted a human evaluation study.

5.1. FuseCap Dataset

In this subsection, we evaluate the enriched captions produced by the caption augmentation approach introduced in the Section 3. In particular, we carry out a qualitative human-study alongside quantitative evaluations. Examples of enriched captions via FuseCap are showcased in Figure 1.

Qualitative Evaluation We conducted a thorough human evaluation study to assess the ability of our enriched captions to be both descriptive and relevant to the images. Specifically, we randomly sampled 400 pairs from the COCO dataset and provided 40 participants with (1) the original caption and (2) the enriched caption. Our study engaged a pool of random internet users as participants. To minimize biases and ensure an impartial evaluation, they completed the survey unaware of the specific research objectives or goals. The evaluators were asked the following: “Does caption 2 provide an additional meaningful and truthful description of the image compared to caption 1?”. As the enriched captions are much more detailed, as can be seen in Figure 1, the

Dataset	Captions	Mean	Voting
COCO	Original	76.7	31.7%
	FUSECAP	80.3	67.6%
SBU	Original	71.9	32.1%
	FUSECAP	75.5	60.2%
CC	Original	72.6	34.7%
	FUSECAP	75.4	59.7%

Table 1. **FUSECAP data quantitative evaluation.** CLIPScore-based comparison between the original captions of common image-text datasets with our enriched ones. “Mean” indicated the mean CLIPScore and “Voting” expresses CLIP’s preference in a one-vs-one setting. As can be seen, FUSECAP obtains significantly improved results in both metrics.

question focuses on whether they are accurate. Our study results indicate that participants find the enriched captions at least as good as the original ones in 72.9% of the images. This finding highlights the proposed method’s effectiveness in enhancing the captions’ descriptiveness while preserving alignment and relevancy to the images.

Quantitative Evaluation We evaluate the enriched captions in comparison to the original captions for each dataset under consideration. To this end, we randomly selected 5000 images from each dataset and report the CLIPScore obtained with both types of captions. As depicted in Table 1, the enriched captions generated by our proposed method consistently achieve a higher CLIPScore (Mean) on average by 4.6%. In addition, given an image and two captions (original and enriched), we utilize CLIP to measure which caption is preferred. We evaluate this on the different datasets and summarize it under “Voting”, which our captions outperform current ones on average by 29.7%. These results demonstrate the effectiveness of our approach in generating enriched captions that better reflect the content of the images.

Image-Text Retrieval To further demonstrate the descriptiveness and accuracy of the FUSECAP data, we assess its performance in the image-text retrieval task. This task involves matching images to text queries and vice versa. If the enriched data is descriptive and accurate, retrieval performance should improve since the additional details can serve as a discriminative factor in establishing these correspondences. Our training methodology aligns with the original BLIP model. After pre-training the model on a large-scale dataset as in Section 4, the model is fine-tuned for image-text retrieval on the COCO training set using both ITC and ITM losses. For the inference step, we employ the method proposed by [39], previously integrated into BLIP. This method involves the selection of K candidates based on feature similarity, followed by their re-ranking using respective ITM values. We report $R@N$ that corresponds to the accuracy of retrieving the true text/image among the top N retrieved results. The details of the competing models are provided in Section 5.2, with the exception that the versions discussed here have been fine-tuned for the retrieval tasks. As illustrated in Table 2, the use of fused captions contributes significantly to the enhancement of retrieval performance on the COCO test set¹. For example, compared to the model trained on corresponding non-enriched data, the $R@1$ score for image-to-text retrieval increased by 22.1%, and for text-to-image retrieval, it increased by 34.8%.

5.2. Caption Generation

We fine-tune the BLIP model that was pre-trained on the complete FUSECAP dataset for captioning, using the

¹While we used enriched captions for training and testing BLIP_{FUSECAP}, for the other baselines we used the original dataset without enrichment.

Model	COCO Retrieval					
	img → text			text → img		
	R@1	R@5	R@10	R@1	R@5	R@10
BLIP†	75.1	92.7	96.4	58.2	82.4	89.2
BLIP-L	82.4	95.4	97.9	65.2	86.3	91.8
BLIP2	85.4	97.0	98.5	68.3	87.7	92.6
BLIP* _{FUSECAP}	97.2	99.5	99.9	93.0	97.4	98.3

Table 2. **Image-text retrieval results.** Performance on COCO retrieval (test sets). The “*” symbol indicates that the model was trained and tested on our enriched dataset. These results attest that given rich captions, BLIP_{FUSECAP} significantly outperforms existing methods, which utilize standard captions.

enriched COCO dataset. We refer to this fine-tuned version as BLIP_{FUSECAP}. To assess the effectiveness of the captioner, and, by extension, the FUSECAP data it is trained on, we thoroughly compare our results to various state-of-the-art captioning models using CLIPScore. In particular, we consider the following baselines:

- **BLIP†** [38]: An original BLIP model, pre-trained and fine-tuned on the same image set as BLIP_{FUSECAP}, which uses **original** captions, in contrast to BLIP_{FUSECAP} which uses the enriched captions. To guarantee a fair comparison, we set configuration parameters identical to those implemented in BLIP_{FUSECAP}.
- **BLIP-L**: A large version of BLIP, which was pre-trained on a dataset comprising 129M images prior to the captioning fine-tuning. Among all models presented in the original BLIP paper, this model achieves superior captioning results.
- **BLIP2-G-OPT_{2.7}** [37]: A state-of-the-art captioning model that integrates a large frozen vision backbone (ViT G) along with a frozen LLM (OPT [83]), trained on the datasets considered for BLIP-L.
- **OFA** [70]: A state-of-the-art image captioning method trained on unimodal and multimodal pre-training tasks and fine-tuned for captioning.
- **GIT** [69]: A state-of-the-art vision-language model, with an image encoder and text decoder architecture. It is scaled up in terms of both pre-training data and model size, and it is fine-tuned for image captioning.
- **Prismer** [46]: A Vision-Language Model with Multi-Modal Experts such as an object detector and an OCR model. Prismar adopts a strategy that could be seen as dual to ours, integrating vision experts into the model architecture, as opposed to leveraging them to enrich the training data.

Model	Images	Parameters	Val	Test
BLIP†	12M	247M	75.2	75.3
BLIP-L	129M	470M	76.1	76.0
OFA	20M	470M	76.6	76.4
GIT	800M	700M	77.1	77.0
BLIP2-G-OPT _{2.7}	129M	3.8B	77.8	77.5
Prismer	13M	1.6B	76.7	76.7
BLIP _{FUSECAP}	12M	247M	78.3	78.5

Table 3. **Image captioning results.** CLIPScore of leading models and our approach on the COCO captions dataset. Our model outperforms much larger models that have pre-trained with significantly more training data.

Building on the methodology outlined in Section 5, we evaluate the performance of our models using the mean CLIPScore metric. As shown in Table 3, our model not only outperforms the BLIP† model by 4.3% but also surpasses the best performing among other models by 1.3%. This improvement over models with a considerably larger parameter count that have been trained on significantly more image-caption pairs underscores the effectiveness of our data centric approach. Furthermore, as illustrated in Figure 3, our approach exhibits clear superiority in the generation of captions, producing descriptions with greater semantic detail compared to those generated by competing models. Our results surpassing Prismar indicate that for generating comprehensive image captions, leveraging vision experts to enrich data can be more beneficial than incorporating them directly into the model architecture. Additional examples showcasing the performance and capabilities of our model can be found in the supplementary material.

Image-Text Retrieval To further assess the quality of the generated captions, we once again utilize the image-text retrieval task, following the method detailed in Section 5.1. This time, however, we consider captions not from the ground truth COCO or the “ground truth” enriched COCO datasets. Instead, they are produced by the BLIP† and BLIP_{FUSECAP} caption generators. Captions from BLIP† and BLIP_{FUSECAP} are assessed using retrieval models fine-tuned on original and enriched captions, respectively, both of these retrieval models are the ones discussed in Section 5.1. Table 4 shows that by using the captions from the captioners instead of the ground truth captions, BLIP_{FUSECAP} retains its superiority over BLIP† in both image-to-text and text-to-image retrievals. Impressively, using the BLIP_{FUSECAP} captions, the retrieval performance is on par with results achieved using the ground truth enriched captions in Table 2. In contrast, BLIP† exhibits a significant performance drop when compared to its results with the corresponding ground truth captions.

Model	COCO Retrieval					
	img → text			text → img		
	R@1	R@5	R@10	R@1	R@5	R@10
BLIP†	56.3	83.0	90.3	54.5	81.2	88.7
	-18.8%	-9.7%	-6.1%	-3.7%	-1.2%	-0.5%
BLIP* _{FUSECAP}	95.0	98.8	99.2	94.5	98.7	99.3
	-2.2%	-0.7%	-0.7%	+1.5%	+1.3%	+1%

Table 4. **Image-text retrieval results using generated captions.** Performance on COCO retrieval (test sets) using captions generated by the models presented in Section 4 rather than ground truth captions. The models used for evaluation are the ones referenced in Section 5.1. The highlighted percentages indicate the performance difference when compared to the use of corresponding ground truth captions in Table 2. Notably, BLIP_{FUSECAP} captions excel over BLIP†, a trend consistent with ground truth captions. Moreover, BLIP_{FUSECAP} maintains a performance parity with the ground truth datasets, while BLIP† shows a noticeable decline in performance.

5.3. Large-Scale Data Influence

Typical VLP frameworks involve pre-training vision-language models on image-caption pairs before fine-tuning them on downstream tasks. Given this structure, one might argue that the generation of enriched captions could serve as an additional downstream task. This would suggest using the standard captions dataset for pre-training, followed by fine-tuning on a smaller enriched dataset. Yet, as mentioned in Section 5.2, our approach with the BLIP_{FUSECAP} model deviates from this by pre-training on enriched datasets instead of the conventional ones. In this section, we evaluate the impact of this deviation.

We compare the performance of models that are pre-trained on both the standard and enriched datasets. Both models are subsequently fine-tuned using the enriched COCO dataset. Unlike the evaluation method detailed in Section 5, which is reference-free, this section employs conventional, reference-based evaluation metrics, as both models are fine-tuned and tested using the same dataset. Table 5 reveals that the model pre-trained on the enriched captions outperforms its counterpart, which was pre-trained on the original captions, across all metrics when both are fine-tuned on the augmented COCO dataset. This outcome underscores the importance of using large-scale enriched datasets for pre-training and validates our automated dataset creation approach. Interestingly, the model that is pre-trained on the enriched captions and fine-tuned using the standard COCO dataset performs comparably to the model pre-trained on the original captions and fine-tuned on the same dataset.

6. Limitations

In the human evaluation study (Section 5.1), a subset of participants favored the original caption. This may be due to our fusing process occasionally missing inter-element dependencies,

Pre-training Data	Fine-tune +Test Data	B@4	CIDEr	SPICE
Standard	Standard	37.8	126.5	22.9
FUSECAP	Standard	38.4	128.7	23.0
Standard	FUSECAP	35.4	111.4	25.0
FUSECAP	FUSECAP	37.3	123.1	26.8

Table 5. **Influence of Large-Scale Data.** Pre-training with the FUSECAP dataset significantly outperforms pre-training on standard data when fine-tuned for enriched caption generation. This underscores the value of the large-scale FUSECAP dataset and its automated creation approach. Additionally, when fine-tuned for standard captioning, the FUSECAP pre-trained model is comparable to, if not slightly better than, the standard data pre-trained model.

as seen in Figure 1, where the cat’s position between the man and the laptop was not captured. Future research might integrate finer visual details, like segmentation, and refine the LLM fuser accordingly.

Ethical considerations FUSECAP dataset was constructed using FlanT5-XL model [19], which was trained on unfiltered data potentially laden with explicit content or inherent biases. Consequently, the proposed dataset may replicate biases from the original model.

7. Conclusions

In this paper, we address the problem of generating highly descriptive and detailed image captions. We observe that existing state-of-the-art methods produce short and often oversimplified captions that fail to capture the intricate details in images. We hypothesize that this is due to datasets limitation, *i.e.*, existing training data composed of concise captions. Thus, the tendency to provide such captions is distilled into the trained models, regardless of their architecture or training method. Therefore, we introduce a novel and generic data-centric automated strategy to enrich existing image captions, termed FUSECAP. Specifically, this approach harnesses visual experts to extract meaningful information from images and an LLM to fuse such data into the existing captions, yielding enriched ones. We apply FUSECAP to different widespread datasets and generate 12M image-enriched caption pairs. The augmented image captions quality is evaluated qualitatively using a human evaluation survey, and quantitatively, using different evaluation methods. Finally, we demonstrate the effectiveness of the enriched data by utilizing it to train an image captioning model, which outperforms significantly larger state-of-the-art methods. We posit that our research highlights the marked potential of LLMs in enhancing powerful data-centric approaches in computer vision.

References

- [1] Aviad Aberdam, David Bensaïd, Alona Golts, Roy Ganz, Oren Nuriel, Royee Tichauer, Shai Mazor, and Ron Litman. Clipter: Looking at the bigger picture in scene text recognition. *arXiv preprint arXiv:2301.07464*, 2023. [2](#)
- [2] Aviad Aberdam, Roy Ganz, Shai Mazor, and Ron Litman. Multimodal semi-supervised learning for text recognition. *arXiv preprint arXiv:2205.03873*, 2022. [2](#)
- [3] 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. [1](#)
- [4] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086, 2018. [2](#)
- [5] Youngmin Baek, Bado Lee, Dongyoon Han, Sangdoon Yun, and Hwalsuk Lee. Character region awareness for text detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9365–9374, 2019. [4](#)
- [6] Shuang Bai and Shan An. A survey on automatic image caption generation. *Neurocomputing*, 311:291–304, 2018. [2](#)
- [7] K. Barnard and D. Forsyth. Learning the semantics of words and pictures. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, pages 408–415 vol.2, 2001. [1](#)
- [8] Darwin Bautista and Rowel Atienza. Scene text recognition with permuted autoregressive sequence models. In *European Conference on Computer Vision*, pages 178–196, Cham, 10 2022. Springer Nature Switzerland. [2](#), [4](#)
- [9] Ali Furkan Biten, Ruben Tito, Andres Mafla, Lluís Gomez, Marçal Rusinol, Ernest Valveny, CV Jawahar, and Dimosthenis Karatzas. Scene text visual question answering. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4291–4301, 2019. [4](#)
- [10] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. *arXiv preprint arXiv:2211.09800*, 2022. [3](#)
- [11] 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. [2](#), [3](#), [4](#)
- [12] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3558–3568, 2021. [2](#), [3](#), [4](#)
- [13] Fei-Long Chen, Du-Zhen Zhang, Ming-Lun Han, Xiu-Yi Chen, Jing Shi, Shuang Xu, and Bo Xu. Vlp: A survey on vision-language pre-training. *Machine Intelligence Research*, 20(1):38–56, 2023. [2](#)
- [14] Xi Chen, Xiao Wang, Soravit Changpinyo, A. J. Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, Alexander Kolesnikov, Joan Puigcerver, Nan Ding, Keran Rong, Hassan Akbari, Gaurav Mishra, Linting Xue, Ashish Thapliyal, James Bradbury, Weicheng Kuo, Mojtaba Seyedhosseini, Chao Jia, Burcu Karagol Ayan, Carlos Riquelme, Andreas Steiner, Anelia Angelova, Xiaohua Zhai, Neil Houlsby, and Radu Soricut. Pali: A jointly-scaled multilingual language-image model. *CoRR*, 2022. [1](#)
- [15] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. UNITER: universal image-text representation learning. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXX*. Springer, 2020. [1](#)
- [16] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX*, pages 104–120. Springer, 2020. [2](#)
- [17] Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. Unifying vision-and-language tasks via text generation. In *International Conference on Machine Learning*, pages 1931–1942. PMLR, 2021. [2](#), [4](#)
- [18] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022. [3](#)
- [19] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*, 2022. [3](#), [4](#), [5](#), [8](#)
- [20] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2625–2634, 2015. [2](#)
- [21] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [2](#)
- [22] Zi-Yi Dou, Yichong Xu, Zhe Gan, Jianfeng Wang, Shuohang Wang, Lijuan Wang, Chenguang Zhu, Pengchuan Zhang, Lu Yuan, Nanyun Peng, et al. An empirical study of training end-to-end vision-and-language transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18166–18176, 2022. [2](#)
- [23] Ali Farhadi, Mohsen Hejrati, Mohammad Amin Sadeghi, Peter Young, Cyrus Rashtchian, Julia Hockenmaier, and David Forsyth. Every picture tells a story: Generating sentences

- from images. In *Computer Vision–ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5–11, 2010, Proceedings, Part IV 11*, pages 15–29. Springer, 2010. 2
- [24] Roy Ganz and Michael Elad. Clipag: Towards generator-free text-to-image generation. *arXiv preprint arXiv:2306.16805*, 2023. 2
- [25] Roy Ganz, Oren Nuriel, Aviad Aberdam, Yair Kittenplon, Shai Mazor, and Ron Litman. Towards models that can see and read. *arXiv preprint arXiv:2301.07389*, 2023. 2, 4
- [26] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning, 2022. 2, 6
- [27] Micah Hodosh, Peter Young, and Julia Hockenmaier. Framing image description as a ranking task: Data, models and evaluation metrics. *Journal of Artificial Intelligence Research*, 47:853–899, 2013. 2
- [28] Xiaowei Hu, Zhe Gan, Jianfeng Wang, Zhengyuan Yang, Zicheng Liu, Yumao Lu, and Lijuan Wang. Scaling up vision-language pre-training for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17980–17989, 2022. 1
- [29] 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. 2
- [30] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. Denscap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4565–4574, 2016. 2
- [31] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3128–3137, 2015. 2
- [32] Andrej Karpathy, Armand Joulin, and Li Fei-Fei. Deep fragment embeddings for bidirectional image sentence mapping. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8–13 2014, Montreal, Quebec, Canada*, 2014. 1
- [33] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73, 2017. 4
- [34] Girish Kulkarni, Visruth Premraj, Sagnik Dhar, Siming Li, Yejin Choi, Alexander C. Berg, and Tamara L. Berg. Baby talk: Understanding and generating simple image descriptions. In *The 24th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2011, Colorado Springs, CO, USA, 20–25 June 2011*, 2011. 1
- [35] Girish Kulkarni, Visruth Premraj, Vicente Ordonez, Sagnik Dhar, Siming Li, Yejin Choi, Alexander C Berg, and Tamara L Berg. Babytalk: Understanding and generating simple image descriptions. *IEEE transactions on pattern analysis and machine intelligence*, 35(12):2891–2903, 2013. 2
- [36] Polina Kuznetsova, Vicente Ordonez, Alexander C. Berg, Tamara L. Berg, and Yejin Choi. Collective generation of natural image descriptions. In *The 50th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, July 8–14, 2012, Jeju Island, Korea - Volume 1: Long Papers*, 2012. 1
- [37] 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, 2, 5, 7
- [38] 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, 5, 7
- [39] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021. 6
- [40] Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Gotmare, Shafiq R. Joty, Caiming Xiong, and Steven Chu-Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6–14, 2021, virtual*, 2021. 1
- [41] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019. 2, 4
- [42] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX 16*, pages 121–137. Springer, 2020. 4
- [43] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004. 5
- [44] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 1, 2, 3
- [45] Lixin Liu, Jiajun Tang, Xiaojun Wan, and Zongming Guo. Generating diverse and descriptive image captions using visual paraphrases. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4240–4249, 2019. 3
- [46] Shikun Liu, Linxi Fan, Edward Johns, Zhiding Yu, Chaowei Xiao, and Anima Anandkumar. Prismr: A vision-language

- model with an ensemble of experts. *arXiv preprint arXiv:2303.02506*, 2023. 2, 4, 7
- [47] Xihui Liu, Hongsheng Li, Jing Shao, Dapeng Chen, and Xiaogang Wang. Show, tell and discriminate: Image captioning by self-retrieval with partially labeled data. In *Proceedings of the European conference on computer vision (ECCV)*, pages 338–354, 2018. 3
- [48] Ruotian Luo, Brian Price, Scott Cohen, and Gregory Shakhnarovich. Discriminability objective for training descriptive captions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6964–6974, 2018. 3
- [49] Rebecca Mason and Eugene Charniak. Nonparametric method for data-driven image captioning. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22–27, 2014, Baltimore, MD, USA, Volume 2: Short Papers*, 2014. 1
- [50] Vicente Ordonez, Girish Kulkarni, and Tamara Berg. Im2text: Describing images using 1 million captioned photographs. *Advances in neural information processing systems*, 24, 2011. 3
- [51] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002. 5
- [52] Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277*, 2023. 4
- [53] 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, 2, 6
- [54] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020. 2, 5
- [55] Krishnan Ramnath, Simon Baker, Lucy Vanderwende, Motaz Ahmad El-Saban, Sudipta N. Sinha, Anitha Kannan, Noran Hassan, Michel Galley, Yi Yang, Deva Ramanan, Alessandro Bergamo, and Lorenzo Torresani. Autocaption: Automatic caption generation for personal photos. In *IEEE Winter Conference on Applications of Computer Vision, Steamboat Springs, CO, USA, March 24–26, 2014*, 2014. 1
- [56] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 4
- [57] Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. *arXiv preprint arXiv:2302.04761*, 2023. 3
- [58] 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. 2, 3, 4
- [59] Zhan Shi, Hui Liu, and Xiaodan Zhu. Enhancing descriptive image captioning with natural language inference. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 269–277, 2021. 3
- [60] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530*, 2019. 4
- [61] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision*, pages 843–852, 2017. 3
- [62] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490*, 2019. 4
- [63] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2
- [64] 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. 5
- [65] Tomás F Yago Vicente, Le Hou, Chen-Ping Yu, Minh Hoai, and Dimitris Samaras. Large-scale training of shadow detectors with noisily-annotated shadow examples. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VI 14*, pages 816–832. Springer, 2016. 2
- [66] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3156–3164, 2015. 2
- [67] Alex Jinpeng Wang, Pan Zhou, Mike Zheng Shou, and Shuicheng Yan. Position-guided text prompt for vision-language pre-training. *arXiv preprint arXiv:2212.09737*, 2022. 4
- [68] Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. Git: A generative image-to-text transformer for vision and language. *arXiv preprint arXiv:2205.14100*, 2022. 1
- [69] Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. GIT: A generative image-to-text transformer for vision and language. *CoRR*, 2022. 2, 7
- [70] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework, 2022. 1, 2, 7
- [71] Teng Wang, Jinrui Zhang, Junjie Fei, Yixiao Ge, Hao Zheng, Yunlong Tang, Zhe Li, Mingqi Gao, Shanshan Zhao,

- Ying Shan, et al. Caption anything: Interactive image description with diverse multimodal controls. *arXiv preprint arXiv:2305.02677*, 2023. 4
- [72] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. *arXiv preprint arXiv:2108.10904*, 2021. 2
- [73] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*, 2022. 1
- [74] Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*, 2021. 3
- [75] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. 4
- [76] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057. PMLR, 2015. 2
- [77] Linjie Yang, Kevin Tang, Jianchao Yang, and Li-Jia Li. Dense captioning with joint inference and visual context. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2193–2202, 2017. 2
- [78] Zhilin Yang, Ye Yuan, Yuexin Wu, William W Cohen, and Russ R Salakhutdinov. Review networks for caption generation. *Advances in neural information processing systems*, 29, 2016. 2
- [79] Daochen Zha, Zaid Pervaiz Bhat, Kwei-Herng Lai, Fan Yang, and Xia Hu. Data-centric ai: Perspectives and challenges. In *Proceedings of the 2023 SIAM International Conference on Data Mining (SDM)*, pages 945–948. SIAM, 2023. 2, 3
- [80] Daochen Zha, Zaid Pervaiz Bhat, Kwei-Herng Lai, Fan Yang, Zhimeng Jiang, Shaochen Zhong, and Xia Hu. Data-centric artificial intelligence: A survey. *arXiv preprint arXiv:2303.10158*, 2023. 2, 3
- [81] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinvl: Revisiting visual representations in vision-language models. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, 2021. 1, 5
- [82] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5579–5588, 2021. 2, 4
- [83] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022. 7
- [84] Deyao Zhu, Jun Chen, Kilichbek Haydarov, Xiaoqian Shen, Wenxuan Zhang, and Mohamed Elhoseiny. Chatgpt asks, blip-2 answers: Automatic questioning towards enriched visual descriptions. *arXiv preprint arXiv:2303.06594*, 2023. 6