

Zero-shot Visual Question Answering with Language Model Feedback

Yifan Du^{1,4}, Junyi Li^{1,3}, Tianyi Tang¹, Wayne Xin Zhao^{1,4}✉ and Ji-Rong Wen^{1,2,4}

¹Gaoling School of Artificial Intelligence, Renmin University of China

²School of Information, Renmin University of China

³DIRO, Université de Montréal

⁴Beijing Key Laboratory of Big Data Management and Analysis Methods

{yifandu1999, batmanfly}@gmail.com

lijunyi@ruc.edu.cn, steventianytang@outlook.com

Abstract

In this paper, we propose a novel language model guided captioning approach, **LAMOC**, for knowledge-based visual question answering (VQA). Our approach employs the generated captions by a captioning model as the context of an answer prediction model, which is a Pre-trained Language model (PLM). As the major contribution, we leverage the guidance and feedback of the prediction model to improve the capability of the captioning model. In this way, the captioning model can become aware of the task goal and information need from the PLM. To develop our approach, we design two specific training stages, where the first stage adapts the captioning model to the prediction model (selecting more suitable caption propositions for training) and the second stage tunes the captioning model according to the task goal (learning from feedback of the PLM). Extensive experiments demonstrate the effectiveness of the proposed approach on the knowledge-based VQA task. Specifically, on the challenging A-OKVQA dataset, LAMOC outperforms several competitive zero-shot methods and even achieves comparable results to a fine-tuned VLP model. Our code is publicly available at <https://github.com/RUCAIBox/LAMOC>.

1 Introduction

Recently, pre-trained language models (PLMs) (Devlin et al., 2019; Brown et al., 2020), especially large language models (Zhao et al., 2023) have demonstrated excellent capabilities in solving tasks that require background knowledge or complex reasoning, such as commonsense reasoning (Sap et al., 2019; Rajani et al., 2019) and logical reasoning (Wei et al., 2022; Kojima et al., 2022). Inspired by these successes, recent studies have proposed utilizing PLMs¹ to solve complex vision-language

✉ Corresponding author.

¹In this paper, PLMs refer to the models trained on text-only corpus, instead of the text encoder/decoder in vision-language pre-trained (VLP) models, which typically have a weaker reasoning capacity in linguistic content.

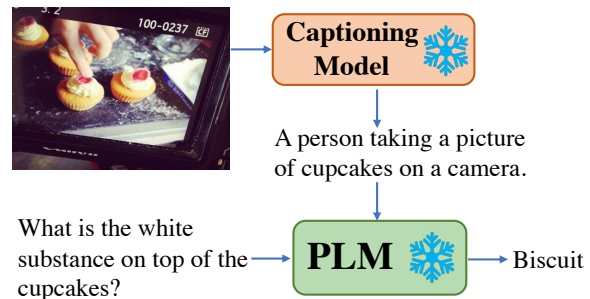


Figure 1: An example that a captioning model (BLIP) fails to provide suitable descriptions for a prediction model (FLAN-T5) of a question in A-OKVQA dataset.

tasks, exemplified by the task of knowledge-based *visual question answering* (VQA) that aims to answer open-ended questions given an image based on outside knowledge (Schwenk et al., 2022). It has been shown that PLM-enhanced approaches (Gui et al., 2022; Lin et al., 2022) typically lead to an improved performance on the knowledge-based VQA task than pure vision-language pre-trained (VLP) models (Schwenk et al., 2022).

In the literature, existing PLM-enhanced VQA approaches can be roughly categorized into two lines. The first line of research focuses on adapting PLMs to the vision modality by introducing specific modular networks or training objectives (Tsimpoukelli et al., 2021; Liang et al., 2022; Alayrac et al., 2022). However, they usually incur a high computational cost during pre-training in order to effectively integrate a vision encoder into the PLM. As another line of research, several studies aim to reduce the cost of tuning PLMs in vision-language tasks by utilizing PLMs in a zero-shot or few-shot manner. They typically generate a caption for an image using a captioning model (e.g., a fine-tuned VLP model), and employ the generated caption as the context (e.g., prompt) to assist PLMs in question answering (Yang et al., 2022; Tiong et al., 2022; Guo et al., 2022). Such an approach is training-free and can be generally applied with

various PLMs.

However, in these existing zero-shot or few-shot methods, the captioning model is unaware of both *task goal* and *information need* for the integrated PLM. They directly reuse the captioning model fine-tuned on caption datasets. As a result, the generated captions tend to be less informative for the VQA task, even irrelevant to the question. Figure 1 presents an example that an inappropriate caption leads to an incorrect answer generated by the PLM. As we can see, the question is highly related to keywords “icing” or “frosting”, while the captioning model misses these information and generates a general description.

To address this issue, we propose **LAMOC**: a novel **L**anguage **M**odel guided **C**aptioning approach for the VQA task. The key idea is to leverage the guidance and feedback of the prediction model (*i.e.*, the PLM) to improve the capability of the captioning model, so that it can be aware of the task goal and information need, and assist the prediction model in answer prediction. Our approach is specially designed with two gradual training stages. At the first stage, the captioning model is trained to align to the prediction model, in which the prediction model selects captions that are more pertinent to a given question from multiple propositions generated by the captioning model. These selected captions are informative and can be used to fine-tune the captioning model to generate informative captions. At the second stage, since the generated caption is used by the PLM as direct evidence for VQA, we employ the feedback from the PLM as reward signals to train the captioning model via reinforcement learning. During training, only the captioning model is tuned while the PLM is fixed, which significantly reduces the computational costs. Meanwhile, since the feedback is from PLM, both training stages do not require any labeled data.

Our contributions can be summarized as follows: (1) We propose **LAMOC**, a novel approach for training captioning models to generate informative captions that can assist PLMs in VQA tasks; (2) Using a small number of randomly sampled unlabeled (image, question) pairs, **LAMOC** consistently outperforms several competitive zero/few-shot baselines without PLM feedback on two knowledge-based VQA datasets: OK-VQA and A-OKVQA; (3) We have demonstrated the effectiveness of our method on PLMs of varying scales, from 223M

to 11B. This not only confirms the robustness of our approach but also demonstrates its potential for generalization to Large Language Models (LLMs).

2 Related Work

PLMs for VQA. After training on large corpora, PLMs exhibit surprising abilities, such as chain-of-thought reasoning (Wei et al., 2022), in-context learning (Brown et al., 2020), and instruction following (Chung et al., 2022), which cannot be obtained by vision-language pre-training. Thus, some works adopt PLM to perform VQA and obtain promising results. One line of research combines a PLM and a vision encoder and trains them end-to-end. Frozen (Tsimpoukelli et al., 2021) and Liang et al. (2022) train a visual encoder or a modular network and keep the PLM frozen to retain its powerful abilities. Flamingo (Alayrac et al., 2022) elaborates the model architecture to combine the vision and language models and scales the model size to 80B. Another line of research tries to deploy PLMs on VQA tasks in a few-shot/zero-shot manner. PICa (Yang et al., 2022) and Img2Prompt (Guo et al., 2022) translate the image to captions or tags and employ GPT-3 to answer a question by in-context learning. PNP-VQA (Tiong et al., 2022) generates question-related captions and utilizes a QA model (Khashabi et al., 2022) for answer prediction. This type of work does not require extra training and can be adapted to new PLMs. Our work follows the second paradigm and is an extension of these works.

Learning from Feedback. A regular paradigm to train a model is defining a loss function and optimizing it. However, certain objectives, such as coherence, diversity, and toxicity in text generation, may not be easily incorporated into the loss function and learned in an end-to-end manner (Paulus et al., 2018; Pang and He, 2021). Thus, explicit feedback on model output is regarded as a learning signal to assist in training. Campos and Shern (2022) utilize a PLM’s refinement and human feedback to fine-tune a summary model. Wang et al. (2022c) leverage compiler feedback to improve the compilability of programs generated by the language model. Ouyang et al. (2022) align a language model with the user’s intention through reinforcement learning from human feedback. We borrow idea from these works, but our feedback comes from a PLM instead of humans, thus saving the annotation cost.

3 Method

In this section, we present the proposed LAMOC: **L**anguage **M**odel guided **C**aptioning method for VQA. The overall architecture of LAMOC is depicted in Figure 2.

3.1 Overview of Our Approach

In this work, we study the task of *visual question answering (VQA)*. Given an image-question pair $x : \langle x_i, x_q \rangle$, the task goal is to predict a correct answer y to the question x_q given the image x_i . Following prior studies (Yang et al., 2022; Tiong et al., 2022), we adopt a captioning-based approach for VQA, in which a captioning model generates auxiliary captions for helping answer prediction. Formally, we represent the above idea in a probabilistic way:

$$p(y|x_i, x_q) \quad (1) \\ = \sum_{z \in \mathcal{Z}} \underbrace{p(z|x_i, x_q; \Theta_C)}_{\text{caption generation}} \cdot \underbrace{p(y|x_q, z; \Theta_P)}_{\text{answer prediction}}$$

where a captioning model Θ_C firstly generates an auxiliary captions z , and then a prediction model Θ_P predicts an answer candidate y based on the caption z and the question x_q . We evaluate this probability by iterating over a set of generated captions. Here, we consider an unsupervised setting: no labeled answer data is available. Although there is no labeled answers, we assume that a small number of image-question pairs can be obtained for training (no overlapping with the task dataset).

To instantiate this probabilistic approach, we adopt a vision-language pre-trained (VLP) model, *i.e.*, BLIP (Li et al., 2022b), as the captioning model, and a pre-trained language model (PLM), *i.e.*, FLAN-T5-XXL (Chung et al., 2022), as the prediction model. The prediction model Θ_P is expected to fulfill the task by accurately predicting the answer, while the captioning model Θ_C plays an assisted role by providing informative evidence for Θ_P . In our approach, the captioning model Θ_C can be tuned while the prediction model Θ_P is fixed during optimization. By leveraging the unlabeled image-question pairs (without the labeled answers), we let the two models cooperate with each other: the captioning model generates informative evidence for helping answer prediction, and the prediction model provides task-specific guidance and feedback to improve the captioning model.

To optimize our approach, we design a gradual training process including two stages: (1) *caption-*

ing adaptation aims to adjust Θ_C to produce informative captions that are suitable for Θ_P (§3.2.1), and (2) *feedback-based learning* aims to optimize Θ_C according to task-specific feedback from Θ_P (§3.2.2). Once the captioning model is well trained, we employ the prediction model for predicting the final answer as in Eq. (1), based on the captions provided by the captioning model (§3.3). Next, we introduce these parts in details.

3.2 Language Model Guided Captioning

The key of our approach (Eq. (1)) is to train an effective captioning model Θ_C for improving the capability of the prediction model Θ_P on VQA. Considering that there are no labeled answers, we employ the prediction model to provide guidance and feedback to optimize the captioning model.

3.2.1 Captioning Adaptation

Since the captioning model is originally intended to describe the given image, it may not be in suited form to assist the prediction model. Thus, we propose a captioning adaptation strategy that tunes the captioning model to fit the prediction model.

Caption Propositions. We first sample n image-question pairs from VQAv2 (Goyal et al., 2017), which is a large VQA dataset containing more than 1M questions and does not overlap with our task dataset. Then we employ the captioning model to propose k captions for each image by nucleus sampling (Holtzman et al., 2019). Among these captions, some may be better suited for the prediction model than the rest. We would like to identify such captions and use them to refine the captioning model.

Instruction-based Captions Selection. Since the prediction model is developed based on the FLAN-T5-XXL, it has encoded a large amount of knowledge in a massive number of parameters. We design the following instruction to prompt FLAN-T5-XXL to identify more informative captions:

“*Question:* [QUESTION] *Caption:* [CAPTION] \backslash \backslash
To what degree does the caption relate to the question: \backslash \backslash A: 0% \backslash \backslash B: 25% \backslash \backslash C: 50% \backslash \backslash D:75%”.

Given the above prompt, FLAN-T5-XXL will generate a corresponding option among the set $\{A, B, C, D\}$. Such an option reflects the correlation between the caption and question, and the captions with the predicted option “D:75%” are more relevant to the question. Since the options are made by the prediction model itself, they tend to be

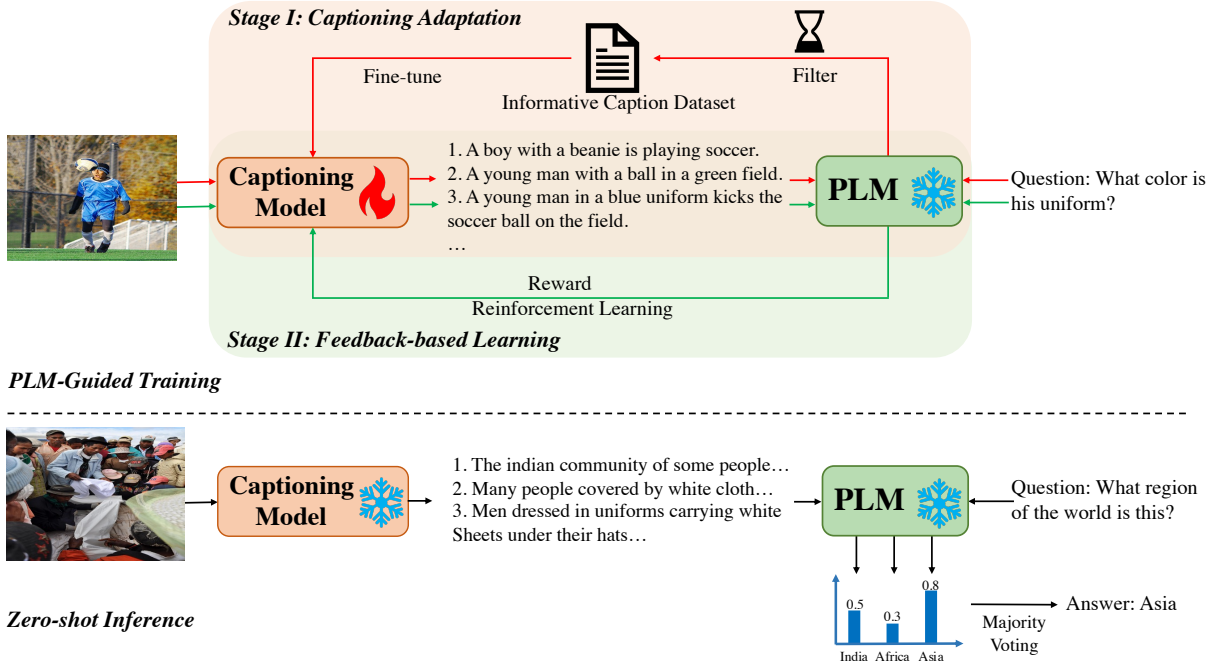


Figure 2: Overview of our proposed approach LAMOC. In captioning adaption, we utilize a PLM to select informative captions and fine-tune the captioning model on them. When learning from PLM feedback, we regard the feedback from the PLM as reward signals and perform reinforcement learning on the captioning model.

more useful for answer prediction. Thus, we keep the captions with the predicted option “ $D:75\%$ ” and discard the rest.

Captioning Model Fine-tuning. Via the above caption selection, we can obtain a set of more informative captions, which are judged by the prediction model. Further, we use them to fine-tune the captioning model by optimizing the following cross-entropy loss:

$$\mathcal{L}_{FT} = -\frac{1}{T} \sum_{t=1}^T \log p(z_t | x_i, z_{<t}), \quad (2)$$

where T is the length of caption, z_t denotes the t -th token of the informative caption selected by FLAN-T5-XXL, $z_{<t}$ represents the generated token up to the $(t-1)$ -th step. After fine-tuning, the captioning model can be better suited for the prediction model.

3.2.2 Feedback-based Learning

Though adapting to the prediction model, the captioning model is still unaware of the answer prediction task for VQA. Thus, we further propose construct pseudo supervision signals based on the PLM feedback from the prediction model. Since the captioning model is only involved as an intermediate component for answer prediction, we design a reinforcement learning method for optimizing it.

Reward From PLM Feedback. A key design consideration of reinforcement learning is the definition of the reward function. In our approach, instead of only generating relevant captions for the images, the effectiveness of the captioning model should be measured by how well it helps find the correct answer. To achieve this goal, we design the following two kinds of reward signals.

- *Prompt-based Reward:* A heuristic method is utilizing the prompt in §3.2.1 to instruct FLAN-T5-XXL to obtain a relevance score, and regard this relevance score as the reward signal:

$$r(x_q, z) = \arg \max_{s \in \{0, 0.25, 0.5, 0.75\}} p(s | x_q, z; \Theta_P), \quad (3)$$

A higher score indicates a more informative caption, which is encouraged.

- *Confidence-based Reward:* Since there is no ground-truth answer during training, following Eq.(1), we employ the probability score of the predicted answer (the most confident candidate) given by the prediction model as the reward:

$$r(x_q, z) = p(\hat{y} | x_q, z; \Theta_P), \quad (4)$$

where z is the generated caption by the captioning model and \hat{y} is the predicted answer from the prediction model. In this way, the PLM (*i.e.*, the prediction model) can inform the captioning model

about the informativeness of the generated caption: the larger probability score, the more informative a caption is, and vice versa. We will verify the reliability of these reward designs in §5.1.

Policy Gradient. In the framework of reinforcement learning, caption generation can be viewed as a sequential decision-making process over the whole vocabulary space. Each generated caption with T tokens is treated as an individual episode of length T in this process. At the t -th time step, the state $(x_i, z_{<t})$ is the combination of the image and caption generated up to the $(t - 1)$ -th token, and the action z_t is the t -th token to be generated. We employ the policy gradient algorithm (Sutton and Barto, 2018) and perform gradient descent to optimize the following objective function:

$$\mathcal{L}_{RL} = - \sum_{t=1}^T r(x_q, z) \log p(z_t | x_i, z_{<t}; \Theta_{cap}), \quad (5)$$

where $z = \langle z_1, \dots, z_t, \dots, z_T \rangle$ is the caption, and $r(x_q, z)$ is the reward given by the PLM. Finally, we jointly optimize the two loss functions:

$$\mathcal{L} = (1 - \alpha) \cdot \mathcal{L}_{FT} + \alpha \cdot \mathcal{L}_{RL}, \quad (6)$$

where α is a weight factor to balance the two parts.

To fully exploit the online feedback provided by FLAN-T5-XXL, we only optimize the captioning adaptation loss function \mathcal{L}_{FT} in the initial epoch, while the reinforcement learning loss function \mathcal{L}_{RL} is optimized throughout the training process.

3.3 Answer Prediction

At inference time, we utilize the updated captioning model to assist the prediction model in answering questions, by calculating the probability $p(y|x_q, z; \Theta_P)$. To increase the diversity of captions and the coverage of answers, we first randomly sample 20% patches from the whole image at each time and apply top- k sampling (Fan et al., 2018) to generate a caption for these patches with the updated captioning model. We repeat this process m times to generate m diverse captions. Then we concatenate each of them with the corresponding question to construct the following prompt:

“Please answer the following question.\n[CAPTION]. [QUESTION]”.

Based on this prompt, the FLAN-T5-XXL is instructed to propose an answer with greedy decoding. We can take the max-voting strategy over all the generated answers.

Different from previous work on learning from feedback (Campos and Shern, 2022; Wang et al., 2022c; Ouyang et al., 2022), our proposed approach explores the guidance and feedback from the prediction model instead of human annotations. As we will see in §5.1, our empirical study shows that there exists a negative correlation between the negative log likelihood assigned by a PLM and the VQA score of a generated answer. This finding suggests that the reward $r(x_q, z)$ given by PLM can potentially serve as a substitute for labeled data to improve the captioning model for the VQA task.

4 Experiment

This section shows the experimental setup and then highlights the main conclusions of our results.

4.1 Experimental Setup

Task Datasets. Since our goal is to improve the performance of PLMs on visual commonsense tasks, we choose two knowledge-based VQA datasets to evaluate our method: (1) **OK-VQA** (Marino et al., 2019) contains 5,046 questions in the test set that require external knowledge resources to answer. (2) **A-OKVQA** (Schwenk et al., 2022) is an augmented dataset based on OK-VQA, which requires additional types of world knowledge compared to OK-VQA. Since the test set of A-OKVQA is not public, we evaluate our method on the validation set. We do not test on VQAv2 (Goyal et al., 2017) because the majority of questions in this dataset are largely focused on recognition and simple visual detection tasks, which can be done without much logical reasoning or external knowledge, and a fine-tuned VLP model could obtain surprising results (Wang et al., 2022b,a). We do not use training data to make a fair comparison with other methods.

Baselines. We divide previous methods into two categories: (1) **Methods without extra large-scale Vision-Language (V-L) pre-training**, which means the models have not been pre-trained on large-scale V-L datasets, including PICa (Yang et al., 2022), PNP-VQA (Tiong et al., 2022), Img2Prompt (Guo et al., 2022). LAMOC also belongs to this category. (2) **Methods with extra large-scale V-L pre-training**, which means that the PLM and the vision encoder are jointly trained on V-L datasets (although the PLM may be fixed, it obtains the ability to understand images), including VL-T5 (Cho et al., 2021), FewVLM (Jin et al.,

Evaluation Setting	Method	Parameters	Use Extra PLM?	With extra V-L Pre-training?	OK-VQA test	A-OKVQA val
Models fine-tuned on training set.						
Supervised learning	BLIP [†]	226M	✗	✗	37.6	38.5
	PromptCap	175B	✓	✗	58.8	58.0
Models without fine-tuning.						
Few-shot	FewVLM _{base}	288M	✗	✓	15.0	-
	FewVLM _{large}	804M	✗	✓	23.1	-
	PICa	175B	✓	✗	48.0	-
Zero-shot	VL-T5 _{no-VQA}	288M	✗	✓	5.8	-
	VLKD _{vIT-B/16}	494M	✓	✓	10.5	-
	VLKD _{vIT-B-L/14}	713M	✓	✓	13.3	-
	FlamingO _{3B}	3B	✓	✓	41.2	-
	FlamingO _{9B}	9B	✓	✓	44.7	-
	FlamingO _{80B}	80B	✓	✓	50.6	-
	Frozen	7B	✓	✓	5.9	-
	PNP-VQA _{3B}	3.9B	✓	✗	34.1	33.4
	PNP-VQA _{11B}	11.9B	✓	✗	35.9	36.0
	Img2Prompt _{6.7B}	8.3B	✓	✗	38.2	33.3
	Img2Prompt _{13B}	14.6B	✓	✗	39.9	33.3
	LAMOC_{11B} (Ours)	11.4B	✓	✗	40.3	37.9

Table 1: Results on OK-VQA and A-OKVQA. The methods are categorized by whether they use extra PLM and whether carry out V-L pre-training. The methods in the upper part have been fine-tuned on the training set, while those in the middle and bottom parts have not. All methods using extra PLM keep it frozen. [†] Instead of first fine-tuning BLIP on VQAv2 and then performing task-specific fine-tuning, we directly fine-tune BLIP on two target datasets for a fair comparison.

2022), VLKD (Dai et al., 2022), Frozen (Tsim-poukelli et al., 2021), and Flamingo (Alayrac et al., 2022). The above methods do not use or use few labeled data (zero-shot/few-shot). Besides, we include two methods, *i.e.*, BLIP (Li et al., 2022b) and PromptCap (Hu et al., 2022), which are fine-tuned on large amounts of labeled data.

Implementation details. For image captioning, we adopt BLIP (Li et al., 2022b) with 446M parameters and load the released checkpoint that has been fine-tuned on the COCO 2014 training set (Lin et al., 2014), which has no overlap with both the OK-VQA and A-OKVQA evaluation datasets. For the PLM, we utilize FLAN-T5-XXL (Wei et al., 2022), which has been fine-tuned on more than 1,800 tasks through instructions and stores considerable world knowledge. We also carry out experiments on PLMs with other sizes, from 223M to 11B parameters, to demonstrate the robustness and generalizability of our approach across PLMs with different sizes. It is noteworthy that the informative caption dataset used in the captioning adaptation stage is selected by FLAN-T5-XXL, because the relevance score given by smaller models is not reliable, as will be illustrated in §5.1. When training the captioning model, we select 1,000 (im-

age, question) pairs without labels from VQAv2 (about 10% of the amount of training data for our target datasets), which has no overlap with the OK-VQA and A-OKVQA. It is worth noting that these 1,000 image-question pairs can be sampled from any datasets or even be generated, we sample from VQAv2 for the sake of reproducibility. The answers are generated by the PLM auto-regressively, without access to the pre-defined answer list. We conduct experiments with 5 random seeds and report the average VQA score according to official evaluation protocols.

4.2 Main Results

Table 1 displays the results of our methods and baselines on OK-VQA and A-OKVQA.

First, LAMOC outperforms all the zero-shot baselines without V-L pre-training on both datasets. Compared to previous state-of-the-art, LAMOC achieves prominent gains on the challenging A-OKVQA dataset (37.9 vs 36.0) and OK-VQA dataset (40.3 vs 39.9). Compared to these baselines, our approach does not require additional image-question matching or question generation modules, thus speeding up the inference speed. Since Flamingo has been trained on a massive V-L dataset, it achieves the best performance among

zero-shot methods. It has been reported that large-scale V-L pre-training can develop a mapping between images and knowledge concepts that can aid in knowledge-based VQA (Tiong et al., 2022).

Second, LAMOC narrows the gap between methods with and without fine-tuning, and even achieves comparable results with the fine-tuned VLP model, *i.e.*, BLIP. For example, the performance gap between PNP-VQA_{11B} and BLIP is 2.5, and has been decreased to 0.6 by LAMOC, which implies the importance of language model feedback.

Finally, we report the results of our methods with different model sizes in Table 2. When increasing the model scale from 223M to 11B, we observe a 1-2 point improvement in VQA scores on the challenging A-OKVQA dataset. This indicates that a larger PLM can not only store more world knowledge to assist with question answering, but also provide more accurate feedback to refine the captioning model. This is further supported by the ablation study in §5.1.

5 Analysis

5.1 The Reliability of Feedback From PLM




Figure 3: The relationship between the caption’s reward and the corresponding answer’s VQA score on A-OKVQA validation set. Figure (a) reflects the reliability of prompt-based reward while figure (b) reflects the reliability of confidence-based reward.

The main idea of our work is leveraging the feedback of a PLM to guide caption generation, so a critical aspect is the reliability of the feedback. LAMOC involves two types of feedback: (1) prompt-based reward and (2) confidence-based reward, which will be evaluated independently.

To evaluate the reliability of the first type of feedback, we analyze the relation between the VQA score and the relevance score provided by the PLM on A-OKVQA validation set (Figure 3(a)). We can observe that as the relevance score provided by FLAN-T5-XXL increases, the VQA score also increases, indicating that FLAN-T5-XXL is a suitable prediction model for providing accurate feed-

back and the relevance scores can be regarded as reward signals. However, this trend is not observed for the other three models, implying that their feedback is unreliable. As a result, we only use FLAN-T5-XXL to select informative captions during captioning adaptation.

To evaluate the reliability of the second type of feedback, we prompt FLAN-T5 to answer the question conditioned on the captions and plot the relationship between the negative log-likelihood (NLL) of the generated answer and its corresponding VQA score. As Figure 3(b) shows, there is a negative correlation between the NLL of the generated answers and their VQA scores, suggesting that captions with lower NLL are more informative and relevant to the questions. Therefore, the probability of the generated answer is a reliable feedback and can be used as the reward signal during reinforcement learning.

5.2 The Effectiveness of Two-stage Training

When training the captioning model, we adopt two gradual training stages: captioning adaptation and feedback-based learning. In this part, we study the effectiveness of this training strategy and explore whether one training stage is more effective than the other. As illustrated in Table 2, different models benefit from different training objectives. For example, the captioning adaptation stage is more beneficial for FLAN-T5-large, leading to an improvement of about 4 points on OK-VQA. On the other hand, FLAN-T5-XXL benefits the most from reinforcement learning with prompt-based rewards and obtains more than 4 points improvement on A-OKVQA. Moreover, the results show that jointly training the two objectives further boosts performance, highlighting the effectiveness of the proposed two-stage training approach.

5.3 Case Study

Figure 4 displays three instances of the captions generated by BLIP and LAMOC, along with the corresponding answers generated by FLAN-T5-XXL. Since LAMOC is trained on the basis of BLIP, the difference can reflect the effect of our method. As can be observed, the captions generated by LAMOC are longer and more comprehensive, containing key information relevant to the question. For example, in Figure 4(a), LAMOC generates captions that include specific details such as “frosting” and “chocolate”, while BLIP only generates general captions about “donuts” and “box”, without sufficient infor-

Method	FLAN-T5-base (223M)		FLAN-T5-large (738M)		FLAN-T5-XL (3B)		FLAN-T5-XXL (11B)	
	OK-VQA	A-OKVQA	OK-VQA	A-OKVQA	OK-VQA	A-OKVQA	OK-VQA	A-OKVQA
BLIP caption	20.42	19.46	23.86	28.86	32.36	31.21	38.48	35.06
+ adaptation	19.72	18.71	27.43	29.19	32.22	31.07	38.35	35.30
+ RL (prompt)	21.24	19.25	27.29	29.73	32.28	30.63	38.74	37.62
+ RL (confidence)	21.14	19.74	25.09	28.98	32.02	32.10	40.31	37.85
+ adaptation + RL	19.72	20.63	24.82	29.84	32.77	32.00	39.72	37.09

Table 2: Results of different model sizes and different training objectives. “BLIP caption” means feeding captions generated by BLIP to PLM without captioning adaptation and feedback-based learning. “adaptation” means captioning adaptation, while “RL (prompt)” means RL with prompt-based reward, and “RL (confidence)” means RL with confidence-based reward.

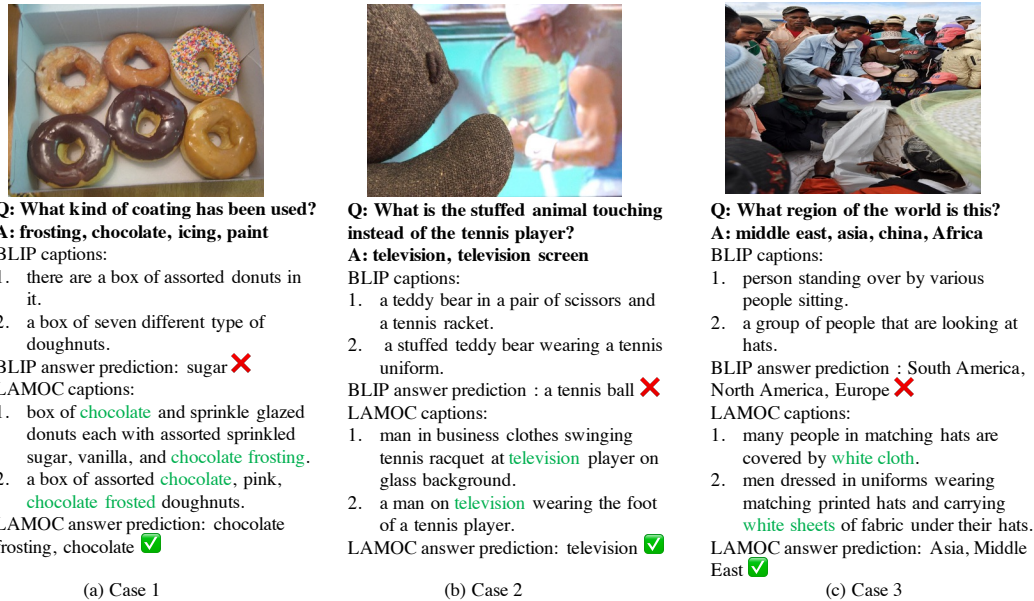


Figure 4: Example captions and predictions generated by BLIP and LAMOC.

mation to help answer the question. These results highlight the importance of training the captioning model under the guidance of PLMs.

One concern is that the PLM may generate correct answers due to the language bias, not attributing to the relevant information contained in the captions. For example, in Figure 4(a), the PLM may generate the answer “chocolate”, even if the captions do not mention chocolate (Li et al., 2023). However, since chocolate often co-occurs with donuts in the training corpora, the PLM may associate chocolate with donuts and generate it as the answer. In order to check how often such a situation happens, we randomly sample 100 questions where the prediction model gives correct answers. For each question, we manually assess whether their answer is derived from the caption. Our analysis reveals that only 6 out of 100 captions are irrelevant to the questions, indicating the reliability of the captions.

Another interesting phenomenon is that the sentences generated by LAMOC can be grammatically incoherent and sometimes incomplete. This indicates that PLM prompting may not always conform to human language patterns, which is consistent with previous studies (Webson and Pavlick, 2022; Deng et al., 2022).

The ablation study of the level of relevance, the number of captions, and the influence of different prompt designs can be found in appendix B.

6 Conclusion

In this paper, we propose LAMOC, a language model guided captioning method that improves a captioning model to generate comprehensive captions for an image to help answer the question. In order to train such a model, we first perform captioning adaptation on a self-generated dataset filtered by FLAN-T5-XXL, and then fine-

tune the updated captioning model through reinforcement learning from PLM feedback. Our method, LAMOC, generates captions that are both informative and able to assist PLMs in VQA tasks, as demonstrated through experiments on two knowledge-based VQA datasets. On the challenging A-OKVQA dataset, LAMOC substantially outperforms previous zero-shot methods and even achieves comparable results to a fine-tuned VLP model. Additionally, we show that LAMOC is generalizable to PLMs of varying sizes, from 223M to 11B parameters, demonstrating its potential to be applied to LLMs, which we leave as future work.

7 Limitations

In our study, we have demonstrated the effectiveness of our proposed method on FLAN-T5 with different sizes. However, we have not yet evaluated its performance on LLMs, which possess an even greater number of parameters and have been pre-trained on larger corpora, thus potentially providing more accurate feedback for both caption adaptation and reinforcement learning. Meanwhile, it is worth noting that PLMs may contain certain biases, and training based on their feedback may amplify these biases. As future work, we aim to investigate the scalability of our method to LLMs, as well as strategies to mitigate the potential negative effects of biases present in PLMs.

Acknowledgement

This work was partially supported by National Natural Science Foundation of China under Grant No. 62222215, Beijing Natural Science Foundation under Grant No. 4222027, and Beijing Outstanding Young Scientist Program under Grant No. BJJWZYJH012019100020098. Xin Zhao is the corresponding author.

References

Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv:2204.14198*.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu,

Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.

Jon Ander Campos and Jun Shern. 2022. Training language models with language feedback. In *ACL Workshop on Learning with Natural Language Supervision*. 2022.

Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. 2021. Unifying vision-and-language tasks via text generation. In *International Conference on Machine Learning*, pages 1931–1942. PMLR.

Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.

Wenliang Dai, Lu Hou, Lifeng Shang, Xin Jiang, Qun Liu, and Pascale Fung. 2022. [Enabling multimodal generation on CLIP via vision-language knowledge distillation](#). In *Findings of the Association for Computational Linguistics: ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 2383–2395. Association for Computational Linguistics.

Mingkai Deng, Jianyu Wang, Cheng-Ping Hsieh, Yihan Wang, Han Guo, Tianmin Shu, Meng Song, Eric P Xing, and Zhiting Hu. 2022. Rlprompt: Optimizing discrete text prompts with reinforcement learning. *arXiv preprint arXiv:2205.12548*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.

Angela Fan, Mike Lewis, and Yann N. Dauphin. 2018. [Hierarchical neural story generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 889–898. Association for Computational Linguistics.

Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913.

- Liangke Gui, Borui Wang, Qiuyuan Huang, Alexander Hauptmann, Yonatan Bisk, and Jianfeng Gao. 2022. **KAT: A knowledge augmented transformer for vision-and-language**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 956–968. Association for Computational Linguistics.
- Jiaxian Guo, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Boyang Li, Dacheng Tao, and Steven CH Hoi. 2022. From images to textual prompts: Zero-shot vqa with frozen large language models. *arXiv preprint arXiv:2212.10846*.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*.
- Yushi Hu, Hang Hua, Zhengyuan Yang, Weijia Shi, Noah A Smith, and Jiebo Luo. 2022. Promptcap: Prompt-guided task-aware image captioning. *arXiv preprint arXiv:2211.09699*.
- Woojeong Jin, Yu Cheng, Yelong Shen, Weizhu Chen, and Xiang Ren. 2022. **A good prompt is worth millions of parameters: Low-resource prompt-based learning for vision-language models**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 2763–2775. Association for Computational Linguistics.
- Daniel Khoshdel, Yeganeh Kordi, and Hannaneh Hajishirzi. 2022. **Unifiedqa-v2: Stronger generalization via broader cross-format training**. *CoRR*, abs/2202.12359.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *arXiv preprint arXiv:2205.11916*.
- Dongxu Li, Junnan Li, Hung Le, Guangsen Wang, Silvio Savarese, and Steven C. H. Hoi. 2022a. **LAVIS: A library for language-vision intelligence**. *CoRR*, abs/2209.09019.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven C. H. Hoi. 2022b. **BLIP: bootstrapping language-image pre-training for unified vision-language understanding and generation**. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 12888–12900. PMLR.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023. **Evaluating object hallucination in large vision-language models**. *CoRR*, abs/2305.10355.
- Sheng Liang, Mengjie Zhao, and Hinrich Schütze. 2022. **Modular and parameter-efficient multimodal fusion with prompting**. In *Findings of the Association for Computational Linguistics: ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 2976–2985. Association for Computational Linguistics.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer.
- Yuanze Lin, Yujia Xie, Dongdong Chen, Yichong Xu, Chenguang Zhu, and Lu Yuan. 2022. **RE-VIVE: regional visual representation matters in knowledge-based visual question answering**. *CoRR*, abs/2206.01201.
- Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*, pages 3195–3204.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.
- Richard Yuanzhe Pang and He He. 2021. **Text generation by learning from demonstrations**. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. **A deep reinforced model for abstractive summarization**. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Nazneen Fatema Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. **Explain yourself! leveraging language models for commonsense reasoning**. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4932–4942. Association for Computational Linguistics.
- Maarten Sap, Ronan Le Bras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A Smith, and Yejin Choi. 2019. Atomic: An atlas of machine commonsense for if-then reasoning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 3027–3035.
- Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. 2022.

- [A-OKVQA: A benchmark for visual question answering using world knowledge](#). In *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part VIII*, volume 13668 of *Lecture Notes in Computer Science*, pages 146–162. Springer.
- Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- Anthony Meng Huat Tiong, Junnan Li, Boyang Li, Silvio Savarese, and Steven CH Hoi. 2022. Plug-and-play vqa: Zero-shot vqa by conjoining large pre-trained models with zero training. *arXiv preprint arXiv:2210.08773*.
- Maria Tsimpoukelli, Jacob Menick, Serkan Cabi, S. M. Ali Eslami, Oriol Vinyals, and Felix Hill. 2021. [Multimodal few-shot learning with frozen language models](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 200–212.
- Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022a. [OFA: unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework](#). In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 23318–23340. PMLR.
- Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, and Furu Wei. 2022b. [Image as a foreign language: Beit pretraining for all vision and vision-language tasks](#). *CoRR*, abs/2208.10442.
- Xin Wang, Yasheng Wang, Yao Wan, Fei Mi, Yitong Li, Pingyi Zhou, Jin Liu, Hao Wu, Xin Jiang, and Qun Liu. 2022c. [Compilable neural code generation with compiler feedback](#). In *Findings of the Association for Computational Linguistics: ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 9–19. Association for Computational Linguistics.
- Albert Webson and Ellie Pavlick. 2022. [Do prompt-based models really understand the meaning of their prompts?](#) In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 2300–2344. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*.
- Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Yumao Lu, Zicheng Liu, and Lijuan Wang. 2022. An empirical study of gpt-3 for few-shot knowledge-based vqa. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 3081–3089.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. [A survey of large language models](#). *CoRR*, abs/2303.18223.

Appendix

A Training Details and Artifacts

For LAMOC training, we adopt the officially released BLIP captioning checkpoint² for model initialization. For both captioning adaptation and reinforcement learning, we adopt the following hyper-parameters: learning rate $2e - 6$, warmup 600 steps, weight decay 0.05, batch size 8. The balance factor α is set to 0.9. We train the model for 10 epochs and choose the one with the highest reward (without labels from the validation set). All the experiments are conducted based on LAVIS (Li et al., 2022a) under BSD 3-Clause License. The A-OKVQA is under the Apache License 2.0.

B Additional Ablation Study

B.1 Level of Relevance

When prompting the PLM to give a correlation score for the caption, the level of relevance is part of the prompt, thus can influence the result. We try different levels for the prompt-based reward and the results are in Table 3. Since four levels gives the highest vqa score, we use four levels in our prompt-based reinforcement learning.

Level	A-OKVQA
A: 0%; B: 100%	27.25
A: 0%; B: 50%; C: 100%	28.29
A: 0%; B: 25%; C: 50%; D: 75%	28.98
A: 0%; B: 25%; C: 50%; D: 75%; E: 100%	27.96

Table 3: VQA score of models trained with different levels of prompt-based rewards.

B.2 Number of Captions

Since the PLM is "blind," all visual information is carried by the captions. Thus, the number of captions is critical for the PLM to answer the question. In Figure 5, we explore the influence of the number of captions. Our results indicate that utilizing a larger number of captions leads to improved performance across various model sizes. Performance gains continue to accumulate even when utilizing 10 captions, leading us to posit that incorporating an even greater number of captions would result in further improvements.

²https://storage.googleapis.com/sfr-vision-language-research/BLIP/models/model_large_caption.pth




Figure 5: VQA score with different number of captions in the A-OKVQA validation set.

B.3 Prompt Design

Another critical design of our method is instructing the FLAN-T5 to provide feedback and answer questions, so we explore the effects of different formats of instruction in Table 4. We can observe that prompt design has a great impact on the results Table 4, which is in line with the conclusion of previous works (Wei et al., 2022).

Prompt	OK-VQA	A-OKVQA
Answer the following question in one word. Q: [caption]. [question]	29.53	29.84
Please answer the following question. [caption]. [question]	28.22	29.73
[caption]. [question]	27.59	27.99
[caption]. [question] Let's think step by step.	18.08	28.72

Table 4: VQA score of the answers generated by FLAN-T5-large conditioned on different prompts.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
7
- A2. Did you discuss any potential risks of your work?
Not applicable. Left blank.
- A3. Do the abstract and introduction summarize the paper's main claims?
1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

4

- B1. Did you cite the creators of artifacts you used?
4
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Appendix
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
4
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
4
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
4

C Did you run computational experiments?

4

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
4

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

4, Appendix

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

4

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Appendix

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

Not applicable. Left blank.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Not applicable. Left blank.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

Not applicable. Left blank.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

Not applicable. Left blank.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Not applicable. Left blank.