

Unified Language Representation for Question Answering over Text, Tables, and Images

Bowen Yu, Cheng Fu, Haiyang Yu, Fei Huang, Yongbin Li*

DAMO Academy, Alibaba Group

{yubowen.ybw, fucheng.fuc, yifei.yhy, f.huang, shuide.lyb}@alibaba-inc.com

Abstract

When trying to answer complex questions, people often rely on multiple sources of information, such as visual, textual, and tabular data. Previous approaches to this problem have focused on designing input features or model structure in the multi-modal space, which is inflexible for cross-modal reasoning or data-efficient training. In this paper, we call for an alternative paradigm, which transforms the images and tables into unified language representations, so that we can simplify the task into a simpler textual QA problem that can be solved using three steps: retrieval, ranking, and generation, all within a language space. This idea takes advantage of the power of pre-trained language models and is implemented in a framework called Solar. Our experimental results show that Solar outperforms all existing methods by 10.6-32.3 pts on two datasets, Multi-modalQA and MMCQA, across ten different metrics. Additionally, Solar achieves the best performance on the WebQA leaderboard¹.

1 Introduction

Information overload is a major problem in today's society due to the vast amount of information available. Question answering (QA) systems can help alleviate this problem by providing users with concise and accurate answers² (Rajpurkar et al., 2016). However, traditional QA models are limited to text or structured data and are unable to utilize the vast amount of multimodal knowledge available on the internet (Hannan et al., 2020), much of which is in non-text formats, such as images (Talmor et al., 2021), data tables (Zhu et al., 2021), etc. For example, on the Wikipedia page for the United States, a user's questions about the US census, the color

of the Statue of Liberty, and the US capital city could be answered using data tables, images, and paragraphs on the page, respectively.

To address this challenge, one strategy is to train separate QA models for each modality and use a classifier to determine which modality to obtain information from (Talmor et al., 2021; Li et al., 2022b). While this method is straightforward, it lacks the ability of cross-modal reasoning as models of different modalities are integrated without interaction. However, both human thought and the current benchmarks require reasoning across modalities (Talmor et al., 2021). An alternative approach is to train a multi-modal model that can accept input from multiple sources simultaneously and produce the answer directly (Chen et al., 2022; Yang et al., 2022a). This method allows for reasoning between modalities, but it can be challenging to train due to the diversified inputs (e.g. text, image, and table can be combined in seven modal combinations). For example, MuRAG (Chen et al., 2022) requires 3 million image-text pairs for pre-training to achieve decent results.

In this paper, we introduce Solar, a unified language representation for multi-modal QA. Unlike existing QA methods performing in the multi-modal space, Solar transforms images and tables into natural languages at the input stage, simplifying the multi-modal QA task into a simpler textual QA problem. This transformation has at least three advantages as follows: (1) it helps to overcome the diversified modal combinations and enable cross-modal reasoning by putting together the textual representation of relevant clues belonging to multiple modalities; (2) language models have learned an enormous amount of world knowledge and patterns by reading billion-level text corpora (Hao et al., 2022), compared with the much smaller multi-modal pre-training datasets used in MuRAG. The memorized information can serve as reusable background knowledge and basic skills for com-

* Corresponding author.

¹The code of this work is available at <https://github.com/AlibabaResearch/DAMO-ConvAI/tree/main/solar>.

²In this paper, we present a uniform solution to question answering and conventional question answering. To avoid confusion, we will refer to these two tasks as QA.

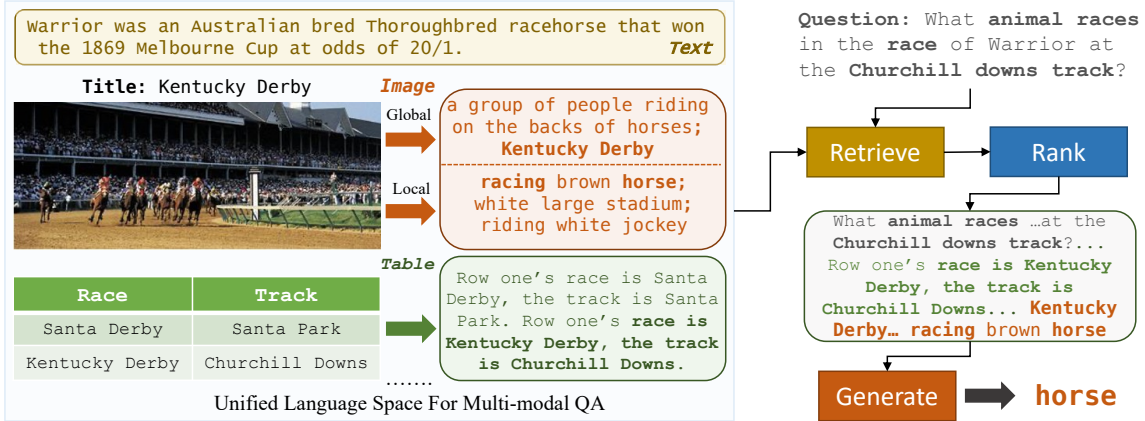


Figure 1: A schematic overview of our approach.

plex QA (Dai et al., 2022); (3) Language has a higher information density compared to images (He et al., 2022) and tables, the transformation from low-density to high-density can be done losslessly for tables and with minimal loss for images using global and local textualization strategies. Then we retrieve relevant clues in the unified language space for each question, rank clues to get the most related top- N ones, and input them into the decoding model to generate the desired answer.

Though simple, our experiments show that Solar achieves state-of-the-art performance on three different multi-modal QA datasets. On Multi-ModalQA and MMCoQA, we outperform sophisticated baselines by 10.6-32.3 pts across ten metrics. Without the need of pre-training on external vision-language data or additional tricks, Solar also achieves the first position on the official WebQA leaderboard at the time of writing³. At the same time, Solar has reduced 99.7% of storage space.

2 Method

In this section, we will introduce Solar in detail. As shown in Figure 1, it consists of two parts: a unified language representation and a unified QA model. The former is responsible for mapping images, tables and texts in different modal spaces into a unified language space. The latter generates an answer through three steps of retrieval, ranking and decoding in the language space. We discuss the related work in Appendix A.

2.1 Unified Language Representation

Multi-modal QA often contains heterogeneous information sources of text, table and image. There-

fore, if we want to implement QA in a unified language space, we need to first convert tables and images into text representations in the input layer.

Table. We use simple natural language templates to transform tables into sentences that sound natural to humans. As an example, consider the table in Figure 1. We can turn the table into a sentence by arranging the cells in a linear fashion, like this: "Row one's race is Santa Derby, the track is Santa Park. Row two is..." and so on. This process doesn't lose any information - the original structure of the table can be reconstructed from the resulting natural language sequence (Chen et al., 2020).

Image. Compared with tables, image has no natural language structure, so transforming image into natural language must be information damaging; In order to alleviate information loss, we propose two transformation strategies: global and local. The global strategy involves using a trained image caption model to generate text descriptions of the macro scenes of the image (Changpinyo et al., 2022; Yang et al., 2022b), including the image's title if it has one. The local strategy leverages a trained object-attribute detection model (Gao et al., 2022) to describe the individual objects and their attributes in the image such as "racing brown horse" in Figure 1, to compensate for missing detail semantics in the overall description. Then the global and local sequences are spliced through punctuations as the language representation of the image.

Question. We hope that Solar will be able to handle both standard QA and more complex conventional QA tasks that involve taking into account the conversation history. To accomplish this, we concatenate the current question with the previous questions and answers in the conversation to create a contextual question q for the system to process.

³<https://eval.ai/web/challenges/challenge-page/1255/leaderboard/3168>

2.2 Unified Multi-Modal QA Model

After building the unified space \mathbb{Z} containing the language representation of all heterogeneous sources, we can convert the multi-modal QA task into a simple textual QA problem. To find the answer to a question q from \mathbb{Z} , Solar follows these steps: first, it retrieves the K clues in \mathbb{Z} most relevant to the question; next, it accurately sorts these clues for the question to get the top- N valuable clues; finally, it sends these clues to the decoding component to generate the answer. The training details of these steps are presented in **Appendix 2.3**

Retrieve. The first step is to retrieve several question-relevant clues in the language space \mathbb{Z} to ground the following generation. We use DPR (Karpukhin et al., 2020) as our retriever, which projects questions and textual clues to a shared space using BERT (Devlin et al., 2019). Formally, we retrieve K most relevant clues $\mathbb{Z}_K = \{z_{[1, \dots, K]} \in \mathbb{Z} \text{ for question } q \text{ as:}$

$$\mathbb{Z}_K = \{z_i \in \mathbb{Z} | \text{topK} \{ \text{BERT}(q)^\top \text{BERT}(z_i) \} \}. \quad (1)$$

Rank. The ranker we use is based on the sequence-pair classification. The question q and each candidate clue $z_i \in \mathbb{Z}_K$ are input together to a BERT followed by a projection layer and Sigmoid function to calculate the ranking score of z_i :

$$s_i = \text{Sigmoid}(\text{Linear}(\text{BERT}(z_i \oplus q))). \quad (2)$$

Cross attention is applied over the tokens of both sequences jointly. Therefore, compared with the independent coding of DPR, it can more accurately model the correlation between q and z_i . In the overall architecture, the ranker plays two roles: first, to find the most relevant clue, which is one of the evaluation metrics; Then, delete irrelevant clues of error recall in the DPR stage to reduce noise in the input of decoding component.

Generate. The top- N clues with the highest ranking scores in \mathbb{Z}_K are then joined with q and passed to an encoder-decoder network such as T5 (Raffel et al., 2020) to generate the answer a : Note that the top- N clues may contain related messages belonging to different modalities. Through the cross attention in the T5 transformer decoding, reasoning between and within modalities can be naturally realized. For instance, consider the text sequence shown in the lower right corner of Figure 1. This sequence includes language representations of tables (highlighted in green) and images

(highlighted in red), which are combined with the question to create a reasoning chain of *Churchill downs track -> Kentucky Derby -> racing -> horse*.

2.3 Model Training

Our training is carried out in three phases: retrieval training, ranking training, and generation training.

2.3.1 Retrieval Training

The goal of retrieval training is to develop an encoder that maps a given question q and all relevant clues in the language space \mathbb{Z} into an embedding space such that the question is close in proximity to its corresponding ground-truth evidence z^+ . The objective is as follows:

$$P_{retr}(z^+ | q, \mathbb{Z}) = \frac{\exp(\text{sim}(q, z^+))}{\sum_{z \in \mathbb{Z}} \exp(\text{sim}(q, z))} \quad (3)$$

where $\text{sim}(q, z)$ is the cosine similarity between the normalized embeddings of the question and evidence, generated by the BERT encoder. In order to perform contrastive learning, a set of negative evidences must be sampled as it is not feasible to enumerate all other evidences. This is done by using the BM25 algorithm to retrieve the most difficult negative clue for each positive clue and then placing them into batches of 32 instances. The training loss is then calculated as the negative log-likelihood for the positive evidence.

2.3.2 Ranking Training

The next phase of training, ranking, begins by gathering the initial retrieval results on the training set. The top 30 samples (excluding the ground-truth evidence z^+) returned by the retrieval module are used as negative examples, and the ranker model is trained to distinguish positive cases from negative cases. In multimodal QA, where multiple positive evidences are required for reasoning, the negative of the summed log-likelihood for the positive evidences is used as the loss function. The logits \mathbf{z} generated by the ranker are used as the clues and the indices for the correct evidences are taken from the ground-truth provenance Pos .

$$\mathcal{L}_{ranking} = \sum_{i \in Pos} \log(\text{Softmax}(\mathbf{z})_i) \quad (4)$$

2.3.3 Generation Training

Current generative models are trained to maximize the probability of generating the correct tokens at each decoding step. Given a question q , a set

of collected clues c , and the ground-truth answer sequence $\mathbf{y}_{1:T} = \{y_1, \dots, y_T\}$, the generation objective is to minimize the loss calculated by the following equation:

$$\mathcal{L}_{generation} = \sum_{t=1}^T \log p(y_t | y_1, \dots, t_t, q, c) \quad (5)$$

3 Experiments

3.1 Datasets

WebQA (Chang et al., 2022) includes QA pairs that require one or two images and text snippets to answer. Each question has a set of distractors that the model must consider along with the correct clues to provide a answer. WebQA uses BARTScore to assess the fluency, and keyword accuracy to assess the correctness of the answer. These two scores are then multiplied to obtain the QA score. The clue retrieval is easily evaluated via F1 score.

MultimodalQA (Talmor et al., 2021) involves combining information from text, tables, and images to answer questions. Each question also includes visual and text distractors. The performance is measured by F1 score at the word level and the Exact Match (EM) of the predicted answer.

MMCoQA (Li et al., 2022b) is the first attempt at addressing the multi-modal conversational QA problem. 57.7% of the conversations in the dataset involve two different modalities, while 24.4% involve three modalities. The statistics all three datasets are demonstrated in Table 6.

3.2 Implementation Details

We conduct experiments on three datasets: WebQA, MMQA, and MMCoQA. The information source for WebQA includes both text and image modalities, while MMCoQA and MMQA focus on text, images, and tables. Our method includes retrieval, ranking, and generation, though not all datasets require all three steps. For WebQA, a candidate clue list is given, and the model needs to find the most relevant clue to evaluate the accuracy of the clue retrieval, so a ranking module is necessary to sort the relevance. Similarly, MMQA also provides a clue list, so retrieval is not necessary. Following previous work (Li et al., 2022b), in MMCoQA, the positive clues are manually included in the retrieval results. However, the number of candidate clues is not small (more than 20), which makes it difficult for the generation model to handle the long sequence, so a ranking model is

Model	QA-FL	QA-Acc	QA	Retr
VLP	42.6	36.7	22.6	68.9
VLP + VinVL	44.2	38.9	24.1	70.9
MuRAG	55.7	54.6	36.1	74.6
Solar	60.9	58.9	40.9	89.4

Table 1: WebQA official test-set results.

Model	QA-FL	QA-Acc	QA	Retr
VitaminC	59	57	39	84
CMU ITL	60	58	39	81
HIT TMG	57	58	39	89
Solar	61	58	41	89

Table 2: WebQA results indicated on leaderboard.

needed to screen the candidate clues. In all datasets, the value of K in retrieval is 30 and the value of N in ranking is 10. The backbone for retrieval and ranking is BERT (Devlin et al., 2019), and the generation model is based on T5 (Raffel et al., 2020). We utilize the Transformers library and pre-trained parameters from HuggingFace⁴ and conduct experiments using 80G GPU cards. Further, AdamW (Loshchilov and Hutter, 2018) is used as the optimization algorithm with a learning rate of $1e-4$. The batch sizes for retrieval, ranking, and generation are 32, 3, and 1, respectively. The image captioning is generated by BLIP (Li et al., 2022a). And the image-attribute features are obtained with VinVL (Zhang et al., 2021).

3.3 Results

WebQA. We conduct experiments on the WebQA dataset, which includes both visual and textual data. Our results, presented in Table 1, demonstrate that Solar significantly outperforms the current SOTA method, MuRAG (Chen et al., 2022), across four metrics. MuRAG encodes image patches and text tokens as sequences of vectors, which are then concatenated and fed to a decoder for answer generation. However, this approach requires pre-training the backbone with 300 million (image, text) pairs. In contrast, Solar converts images into text and performs QA in the language space, leveraging the world knowledge from public pre-trained models. Comparing to the language pre-training corpus such as C4 (34 billion words) (Raffel et al., 2020), multi-modal pre-training datasets are much smaller, which leads to less knowledge.

⁴<https://huggingface.co/>

Model	Single-Modal		Mutli-Modal		All	
	EM	F1	EM	F1	EM	F1
AR	51.7	58.5	34.2	40.2	44.7	51.1
ID	51.6	58.4	44.6	51.2	48.8	55.5
Solar	69.7	74.8	55.5	65.4	59.8	66.1

Table 3: MultimodalQA results.

Model	Dev		Test	
	EM	F1	EM	F1
ORConvQA	1.0	3.0	1.0	1.9
ManyModelQA	0.7	2.3	1.0	1.8
MAE	21.5	30.2	24.9	32.3
Solar	56.8	62.5	57.3	64.6

Table 4: MMCoQA test-dev-set results.

Furthermore, Solar ranks first on the WebQA leaderboard even without using model ensemble or additional post-processing strategies, as shown in Table 2. Note that the other methods we compared all employ external tricks to improve performance (Yang et al., 2022a).

MultimodalQA. As demonstrated in Table 3, Solar also achieves significant performance improvements on the more complex MultimodalQA dataset, which contain text, images, and tables. There are 11.0 pts and 20.6 pts improvements on EM and F1, respectively. The SOTA approach, ID (Talmor et al., 2021), requires an additional classifier to determine the appropriate modality for generating answers. This additional step can hinder the performance, as it relies on accurate classification results. In contrast, Solar blurs the boundaries between different modes at the input, allowing the model to generate answers from a unified text input without considering the modality type.

MMCoQA. Solar demonstrates the highest performance improvement on the most challenging MMCoQA dataset (Table 4). The EM and F1 scores have both increased by 32.4 and 32.3 pts, respectively. In comparison to MultimodalQA, MMCoQA requires the model to correctly incorporate dialog history. The previous SOTA method MAE (Li et al., 2022b), relies on ensembling three conventional QA models and utilizing a modality classifier to determine the final output. However, MAE lacks the ability to reason between modalities, which is crucial for MMCoQA, as 35.7% of the questions require cross-modal reasoning. Our proposed Solar is able to effectively retrieve multiple relevant clues in the unified language space (potentially belonging to different modalities), and

utilize cross attention in transformer decoding to perform cross-modal reasoning.

In addition, Solar also significantly reduces the feature storage. Text storage is more efficient than image feature, For example, AR requires 78G of space to store Faster R-CNN (Ren et al., 2015) extracted image features, while Solar only requires 186M of space to store text representation of three modalities, thus saving 99.7% of storage space. This is particularly important for online services, as feature extraction, such as converting images to embeddings or text, can be done offline, but feature storage is always present. We also give an ablation study in **Appendix B**. Our analysis shows that both global and local textualization of images have a positive effect on the results, with global features being more significant.

4 Conclusion

In this paper, we approach the multi-modal QA task from a new perspective, where all the information sources are aligned into the language space to take advantage of the power of pre-trained language models and facilitate cross-modal reasoning. Our experiments show the promise of this approach, as it outperforms baselines by a large margin and achieves the top position at public leaderboard with a single model. We hope this work can lead to a discussion: can language become a unified interface for people and models to understand the world?

5 Limitations

While our Solar has demonstrated its superior performance on three benchmarks, it still has several limitations. Firstly, Solar relies on accurate recognition of image caption and object-attribute detection models. If the features of these two parts are not correctly recognized, it will cause subsequent cascading errors. Secondly, it only demonstrates that language can serve as a unified representation in multi-modal QA, but has not been tested in other more multi-modal tasks, which we will leave as future work. Lastly, the experimental results do not delve deeper into which cases a unified language representation is better and in which cases a multi-modal model performs better. We speculate that an integration of language models and multi-modal models will yield better results.

6 Ethics Statement

Our submission adheres to the ACL Code of Ethics as it utilizes open public datasets and models, and does not raise any ethical concerns regarding human subjects, harmful insights, conflicts of interest, discrimination, privacy, security, legal compliance, or research integrity. Additionally, we have thoroughly referenced related works and have compared our results accordingly.

References

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. [VQA: visual question answering](#). In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 2425–2433. IEEE Computer Society.
- Jiangxia Cao, Shaoshuai Li, Bowen Yu, Xiaobo Guo, Tingwen Liu, and Bin Wang. 2023. Towards universal cross-domain recommendation. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, pages 78–86.
- Yingshan Chang, Mridu Narang, Hisami Suzuki, Guihong Cao, Jianfeng Gao, and Yonatan Bisk. 2022. Webqa: Multihop and multimodal qa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16495–16504.
- Soravit Changpinyo, Doron Kukliansy, Idan Szpektor, Xi Chen, Nan Ding, and Radu Soricut. 2022. [All you may need for VQA are image captions](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1947–1963, Seattle, United States. Association for Computational Linguistics.
- Wenhu Chen, Hexiang Hu, Xi Chen, Pat Verga, and William W Cohen. 2022. [Murag: Multimodal retrieval-augmented generator for open question answering over images and text](#). *ArXiv preprint*, abs/2210.02928.
- Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyu Zhou, and William Yang Wang. 2020. [Tabfact: A large-scale dataset for table-based fact verification](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. 2022. [Knowledge neurons in pretrained transformers](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8493–8502, Dublin, Ireland. Association for Computational Linguistics.
- 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Feng Gao, Qing Ping, Govind Thattai, Aishwarya Rganti, Ying Nian Wu, and Prem Natarajan. 2022. Transform-retrieve-generate: Natural language-centric outside-knowledge visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5067–5077.
- Michael Glass, Gaetano Rossiello, Md Faisal Mahbub Chowdhury, Ankita Naik, Pengshan Cai, and Alfio Gliozzo. 2022. [Re2G: Retrieve, rerank, generate](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2701–2715, Seattle, United States. Association for Computational Linguistics.
- Darryl Hannan, Akshay Jain, and Mohit Bansal. 2020. [Manymodalqa: Modality disambiguation and QA over diverse inputs](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 7879–7886. AAAI Press.
- Yaru Hao, Haoyu Song, Li Dong, Shaohan Huang, Zewen Chi, Wenhui Wang, Shuming Ma, and Furu Wei. 2022. [Language models are general-purpose interfaces](#). *ArXiv preprint*, abs/2206.06336.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009.
- Sebastian Hofstätter, Jiecao Chen, Karthik Raman, and Hamed Zamani. 2022. [Fid-light: Efficient and effective retrieval-augmented text generation](#). *ArXiv preprint*, abs/2209.14290.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe

- Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022a. [Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation](#). *ArXiv preprint*, abs/2201.12086.
- Yongqi Li, Wenjie Li, and Liqiang Nie. 2022b. [MM-CoQA: Conversational question answering over text, tables, and images](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4220–4231, Dublin, Ireland. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2018. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Jiasen Lu, Christopher Clark, Rowan Zellers, Roozbeh Mottaghi, and Aniruddha Kembhavi. 2022. [Unified-io: A unified model for vision, language, and multi-modal tasks](#). *ArXiv preprint*, abs/2206.08916.
- Yuning Mao, Pengcheng He, Xiaodong Liu, Yelong Shen, Jianfeng Gao, Jiawei Han, and Weizhu Chen. 2021. [Generation-augmented retrieval for open-domain question answering](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4089–4100, Online. Association for Computational Linguistics.
- Jia Ning, Chen Li, Zheng Zhang, Zigang Geng, Qi Dai, Kun He, and Han Hu. 2023. [All in tokens: Unifying output space of visual tasks via soft token](#). *ArXiv preprint*, abs/2301.02229.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Revant Gangi Reddy, Xilin Rui, Manling Li, Xudong Lin, Haoyang Wen, Jaemin Cho, Lifu Huang, Mohit Bansal, Avirup Sil, Shih-Fu Chang, et al. 2022. [Mumuqa: Multimedia multi-hop news question answering via cross-media knowledge extraction and grounding](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, 10, pages 11200–11208.
- Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. 2015. [Faster R-CNN: towards real-time object detection with region proposal networks](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 91–99.
- Alon Talmor, Ori Yoran, Amnon Catav, Dan Lahav, Yizhong Wang, Akari Asai, Gabriel Ilharco, Hananeh Hajishirzi, and Jonathan Berant. 2021. [Multimodalqa: complex question answering over text, tables and images](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Qian Yang, Qian Chen, Wen Wang, Baotian Hu, and Min Zhang. 2022a. [Enhancing multi-modal and multi-hop question answering via structured knowledge and unified retrieval-generation](#). *ArXiv preprint*, abs/2212.08632.
- Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Yumao Lu, Zicheng Liu, and Lijuan Wang. 2022b. 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.
- Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. 2021. [Vinvl: Revisiting visual representations in vision-language models](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5579–5588.
- Yingxiu Zhao, Bowen Yu, Haiyang Yu, Bowen Li, Chao Wang, Fei Huang, Yongbin Li, and Nevin L Zhang. 2023. [Causal document-grounded dialogue pre-training](#). *arXiv preprint arXiv:2305.10927*.
- Fengbin Zhu, Wenqiang Lei, Youcheng Huang, Chao Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng, and Tat-Seng Chua. 2021. [TAT-QA: A question answering benchmark on a hybrid of tabular and textual content in finance](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3277–3287, Online. Association for Computational Linguistics.

A Related Work

A.1 Multimodal Question Answering

The problem of multimodal question answering has been extensively studied. VQA (Antol et al., 2015) is firstly proposed to answer questions from visual-only inputs. Later, WebQA (Chang et al., 2022), MuMuQA (Reddy et al., 2022), Many-ModalQA (Hannan et al., 2020) provide questions

which require reasoning over images and explicitly provided text snippets. More recently, Multi-modalQA (Talmor et al., 2021) and MMCoQA (Li et al., 2022b) require integrating information across free text, semi-structured tables, and images. To address the challenge of finding answers from multiple sources of information, MuRAG (Chen et al., 2022) designs a multi-modal transformer architecture to accept both text and image feature inputs, and builds a million-scale dataset for pre-training the model. AutoRouting (Talmor et al., 2021) and MAE (Li et al., 2022b) train separate models for each modality, fully utilizing existing pre-trained models and then using classifiers to determine which modality to answer a question from, but unable to perform cross-modal reasoning.

In contrast to these two methods, which perform multi-modal QA within the multi-modal space, we propose Solar, which transforms tables and images to the language space and uses a unified pre-trained language model to achieve multi-modal QA. This brings two advantages: first, it fully utilizes the language model pre-trained on tens of billions of tokens, having more world knowledge than multi-modal models with only a few million pre-training data; second, it can naturally achieve cross-modal reasoning. Experiments on three datasets have demonstrated the advantages of Solar in terms of performance, parameter amount, and storage space.

A.2 Unified Model

One of the most ambitious goals in the field of artificial intelligence is the development of a unified model that can perform a wide range of tasks across different modalities (Lu et al., 2022; Ning et al., 2023; Hao et al., 2022; Cao et al., 2023). However, this is a challenging task due to the diversity and complexity of real-world scenarios. Despite these difficulties, recent advances in natural language processing, specifically with large-scale language models such as GPT-3, have demonstrated impressive capabilities as a general-purpose solution for language-based tasks. This success has encouraged researchers to investigate the potential for universal models in other areas, such as computer vision, through works like knowledge-based VQA (Yang et al., 2022b; Changpinyo et al., 2022; Gao et al., 2022). In this paper, we explore an even more complex multi-modal task: finding answers from three different data sources, including text, images, and tables. We demonstrate that language can be used

Datasets	Text?	Table?	Image?	Retrieve?	Rank?	Generate?
WebQA	✓	✗	✓	✗	✓	✓
MMQA	✓	✓	✓	✗	✓	✓
MMCoQA	✓	✓	✓	✓	✗	✓

Table 5: Comparison of Multi-modal QA datasets.

Datasets	Train	Dev	Test
WebQA	34.2k	5k	7.5k
MultiModalQA	23.8K	2.4K	-
MMCoQA	45.8k	0.6k	0.6k

Table 6: Overall Statistics of downstream datasets.

as a unified interface to solve such a complex task.

A.3 Retrieval-augmented Generation

Solar is also closely related to retrieval-augmented generation. This approach combines the benefits of retrieval and generation methods to provide more comprehensive and accurate responses to complex questions (Lewis et al., 2020; Mao et al., 2021). Retrieval-augmented generation models have been shown to outperform standalone language models by leveraging a knowledge base for additional context and provenance (Hofstätter et al., 2022; Glass et al., 2022; Zhao et al., 2023). The results of Solar demonstrate the effectiveness of this approach in multi-modal QA.

B Ablation Study

Model	EM	F1
Solar	56.8	62.5
– Global textualization	48.9	54.1
– Local textualization	56.3	60.9

Table 7: Ablation study of Solar on MMCoQA dev set.

In this experiment, we ablate the Global and Local textualization strategies for images. Note that we cannot ablate retrieval, ranking, and generation as all datasets require generation, and each dataset only contains one of the steps of retrieval or ranking. Linearization of tables also cannot be ablated otherwise tables cannot be mapped to the language space. Results of the ablation experiments are shown in Table 7. It can be seen that both global and local features have a positive impact on the results, with global features being particularly significant. We hypothesize that this may be partly due to the importance of image captions for understanding the overall semantic meaning of an image, and partly due to the fact that image titles are often relevant answers to many visual-related questions.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section 5
- A2. Did you discuss any potential risks of your work?
Section 6
- A3. Do the abstract and introduction summarize the paper’s main claims?
Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
Not applicable. Left blank.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Not applicable. Left blank.
- 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)?
Not applicable. Left blank.
- 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.?
Not applicable. Left blank.
- 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.
Left blank.

C Did you run computational experiments?

Section 3

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

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?

Appendix C

- 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?

Section 3

- 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.)?

No response.

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.?

No response.

- 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)?

No response.

- 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?

No response.

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

No response.

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

No response.