

Fine-grained and Explainable Factuality Evaluation for Multimodal Summarization

Anonymous ACL submission

Abstract

Multimodal summarization aims to generate a concise summary based on the input text and image. However, the existing methods potentially suffer from unfactual output. To evaluate the factuality of multimodal summarization models, we propose two fine-grained and explainable evaluation frameworks (FALLACIOUS) for different application scenarios, i.e. reference-based factuality evaluation framework and reference-free factuality evaluation framework. Notably, the reference-free factuality evaluation framework doesn't need ground truth and hence it has a wider application scenario. To evaluate the effectiveness of the proposed frameworks, we compute the correlation between our frameworks and the other metrics. The experimental results show the effectiveness of our proposed method. We will release our code and dataset via github.

1 Introduction

Sentence summarization, a pivotal task in natural language processing, focuses on creating concise summaries from longer sentences. This area has gained significant attention due to its utility in summarizing events and products, as indicated by various studies. Traditionally, this summarization relied solely on the source sentence. However, with the growing presence of multimedia content combining text and images, recent research has evolved towards multimodal summarization (MMS) (Song et al., 2022; Jing et al., 2023; Lin et al., 2023). This approach integrates visual cues with textual information to enhance the summarization process, making it easier for readers to quickly capture the essence of the information.

In the realm of MMS, notable advancements have been made. Li et al. employed sequence-to-sequence models for better semantic understanding and text generation. They further enhanced MMS by introducing a multimodal selective gate

network, which helped in pinpointing the most relevant parts of a sentence based on its corresponding image. More recently, Song et al. developed a method for generating efficient product summaries using generative pre-trained language models like BART (Lewis et al., 2020). This method begins by transforming product images into attributes, which are then used by the BART model for generating succinct summaries. This shift towards incorporating visual elements marks a significant development in the field of sentence summarization. A significant challenge faced by these models is the issue of hallucination (Wan and Bansal, 2022a). This occurs when the model produces content that is neither present nor implied in the original input text and image.

Recent progress has been made in creating metrics that align closely with human assessments of factual accuracy in summaries (Tang et al., 2022; Zhu et al., 2021). These metrics are designed to evaluate the level of factual consistency between the original document and its generated summary. However, there is only one related work that employs factual accuracy in the multimodal summarization task. Wan and Bansal proposed CLIPBERTSCORE, a simple weighted combination of CLIPScore (Hessel et al., 2021) and BERTScore (Zhang et al., 2020a) to leverage the robustness and strong factuality detection performance between image-summary and document summary, respectively.

However, BERTScore pays more attention to more grammatically correct sentences rather than factual sentences (Hanna and Bojar, 2021). Therefore, the factuality of the text input cannot be measured. In addition, due to CLIPScore's limitations in accurately counting objects (Radford et al., 2021) or conducting compositional reasoning (Ma et al., 2023), the CLIPScore often proves to be unreliable and can yield inaccurate results. Furthermore, the existing evaluation metric only generates an overall

score, which is coarse-grained and less explainable.

To tackle these limitations, we propose two fine-grained and explainable factuality evaluation frameworks for multimodal summarization (FALLA-CIOUS), which can be applied to reference-based and reference-free scenarios. In the reference-based situations, we generate comprehensive questions based on the input textual modality, to provide atomic evaluation. Then, we compare the answers regarding the referenced summary and predicted summary to derive the final factuality score. In the reference-free situation, we generate questions from the model-predicted summary and tested whether the atomic information mentioned in the questions existed in the image and document or not. Based on this, we can also get the final factuality score.

Our contributions can be summarized as: (1) We proposed fine-grained and explainable factuality evaluation frameworks for multimodal summarization under reference-based and reference-free scenarios; (2) We offer an in-depth analysis of our metric along with its components across different benchmarks for evaluating factuality. Through this, we provide substantial empirical proof of its robustness.

2 Related Work

Multimodal Summarization Recently, multimodal summarization, which integrates text with images, has emerged as a significant advancement over traditional text summarization (Li et al., 2020b,a; Zhang et al., 2021; Li et al., 2020c; Palaskar et al., 2019). Early works focused on combining Convolutional Neural Network-based visual models with Recurrent Neural Network-based textual models to improve the multimodal summarization (Zhu et al., 2018; Chen and Zhuge, 2018; Li et al., 2018; Zhu et al., 2020; Zhang et al., 2022). Recent research in the field has shifted towards leveraging pre-trained models for multimodal summarization (Jiang et al., 2023; Lin et al., 2023; Jing et al., 2023; Song et al., 2022). Despite the superior performance of the existing models, these models always suffer from the hallucination where the model generates contents that are not present or entailed by inputs (Wan and Bansal, 2022a).

Faithfulness and Factuality Metrics Evaluating the factuality of generated summaries is crucial for their reliability and accuracy. Factuality metrics fall into two main categories: entailment-based met-

rics and question-answering-based (QGQA) metrics. Entailment-based metrics (Kryscinski et al., 2020; Goyal and Durrett, 2021) assess whether summaries reflect the source document accurately. QGQA approaches (Durmus et al., 2020; Wang et al., 2020; Scialom et al., 2021; Fabbri et al., 2022), evaluate the factuality of the summary by generating and answering input-related questions. Additionally, some other efforts, such as counterfactual estimation (Xie et al., 2021) and embedding-based metrics (Zhang et al., 2020b) have been introduced to this field. Different from them, Wan and Bansal is the first one to consider the textual and visual modalities for multimodal summarization factuality evaluation. Although they have achieved compelling success, these methods do not take into account fine-grained elements and are hard to explain. To address this issue, we introduced a framework involving entity-focused questions to evaluate the model’s fine-grained factuality, which allows for a more detailed assessment of how accurately the model reflects specific entities and relations within the multimodal input.

3 Frameworks

3.1 Problem Definition

Suppose we have an image I and a document D . The existing multimodal summarization models aim to generate a summary based on the multimodal input (i.e. I and D), $S = \mathcal{M}(I, D)$, where \mathcal{M} denotes the multimodal summarization model which takes I and D as inputs and generate a summary. Then we devise the evaluation metric to assess the faithfulness of the generated summary S . We devised the reference-based and reference-free faithfulness evaluation metrics, respectively. The former is based on the ground truth summary, formulated as $S_r = \mathcal{F}(I, D, Y, S)$, where Y is the ground-truth summary. Different from it, the reference-free factuality evaluation metric supposes that there is no available ground truth summary, formulated as $S_s = \mathcal{F}(I, D, S)$. In this setting, the metric has more comprehensive application sceneries.

3.2 Evaluation Framework

We devise two different evaluation frameworks (i.e. reference-based and reference-free frameworks) to facilitate factuality evaluation under various application scenarios.

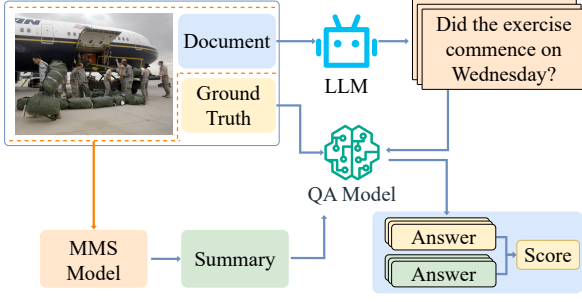


Figure 1: The proposed reference-based framework.

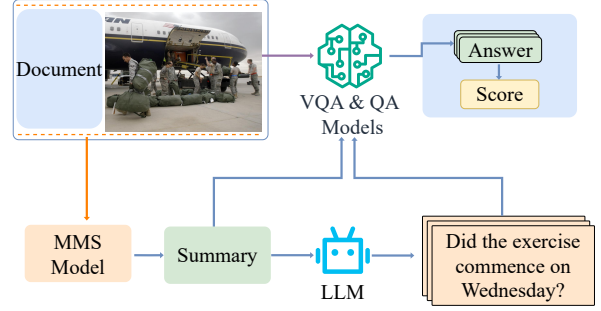


Figure 2: The proposed reference-free framework.

3.2.1 Reference-based

As depicted in Figure 1, the reference-based framework contains three steps: question generation, answer generation, and factuality score aggregation.

Question Generation To get a fine-grained evaluation in terms of factuality, we first generate a question set \mathcal{Q} based on the input document, which can generate comprehensive evaluation questions. All the questions are based on atomic information (Min et al., 2023). We give some examples in Appendix D. The prompt for question generation is detailed in Appendix C. This process is implemented by a Large Language Model (LLM), formalized as $LLM(Y, D) \rightarrow \mathcal{Q}$.

Answer Generation Next, we utilize a Question Answering (QA) model to answer the questions \mathcal{Q} based on the ground-truth summary Y and generated summary S , respectively. We denote the set of answers as \mathcal{A}_r and \mathcal{A}_s corresponding to the ground-truth summary Y and generated summary S . Here, we choose GPT-4 as the model that predicts responses to specific questions, with three possible outcomes: yes, no, or not provided.

Score Aggregation Finally, to evaluate the factuality of the predicted summary S , a factuality score S_r is calculated. When the answers based on the reference summary and generated summary are consistent, the information that corresponds to the question is truthful. Therefore, this score can be calculated by the consistency between the set of answers \mathcal{A}_r and \mathcal{A}_s . The score is computed as the sum of individual match scores for all questions, represented by the following equation,

$$S_r = \frac{\sum_{a_r \in \mathcal{A}_r, a_s \in \mathcal{A}_s} \mathbb{I}(a_r == a_s)}{|\mathcal{Q}|}, \quad (1)$$

where \mathbb{I} is the indicator function and $|\mathcal{Q}|$ is the number of generated questions.

3.2.2 Reference-free

Obviously, manually annotating ground truth summary for multimodal data is labor- and time-extensive. Hence, we further present a reference-free framework illustrated in Figure 2, which consists of three stages: question generation, answer generation, and score aggregation.

Question Generation In this situation, the reference summary is not available for evaluation. Therefore, we resort to the consistency between multimodal input and summary to access the model factuality. Unlike the reference-based framework, we only generate questions \mathcal{Q} based on the summary S generated by the model. It ensures that the questions are directly relevant to the content of the summary, allowing for a more precise assessment of the factuality of the predicted summary, even without the reference. Similarly, we leverage an LLM to generate questions \mathcal{Q} from the generated summary S . This process can be formalized as: $LLM(S) \rightarrow \mathcal{Q}$

Answer Generation To verify the consistency between the multimodal input and summary, we compute the consistency for each modality. Specifically, we use a QA model to answer questions \mathcal{Q} based on the document D and a VQA model to answer questions \mathcal{Q} based on the image I . The QA model also is GPT-4, and the VQA model is implemented by BLIP-2 (Li et al., 2023a). We denote the set of answers as \mathcal{A}_t and \mathcal{A}_i corresponding to the document D and the image I .

Score Aggregation Finally, we aggregate all QA pairs to get the factuality score S_s , which can be represented as:

$$S_s = \frac{\sum_{a_t \in \mathcal{A}_t, a_i \in \mathcal{A}_i} \mathbb{I}(H(a_t) \vee H(a_i))}{|\mathcal{Q}|}, \quad (2)$$

where \mathbb{I} denotes the indicator function, $H(x)$ represents a condition where x is equal to 'Yes', \vee

symbolizes the logical OR operation, and $|Q|$ is the total number of questions in Q .

In this formulation, we consider that if either the a_t or the a_i contains 'Yes', it indicates factuality, signifying that the model has not fabricated information. Otherwise, it indicates a hallucination.

4 Experiments

4.1 Experimental Setting

We employed the BART-MMSS (Lin et al., 2023) model and MMSS dataset (Li et al., 2018) to evaluate the effectiveness of our proposed method. We employed GPT-4 to generate questions and as the QA model. We used BLIP-2 (Li et al., 2023b) as the VQA model.

4.2 Human Evaluation

To verify the effectiveness of the reference-based evaluation framework, we randomly selected 200 samples from the MMSS dataset. In addition, we randomly annotated 200 samples generated by another model V2P (Song et al., 2022) on another dataset CEPsum (Li et al., 2020a). More annotation details are detailed in Appendix A.

4.3 Result Analysis

Evaluating QA Model We evaluate the performance of the QA model to get the reliability of this component. We collected all QA pairs from the annotated dataset and got an accuracy of 95.3%, which shows this component is reliable.

Evaluating VQA Model Similarly, to evaluate the performance of the VQA model, we collected all VQA questions from the annotated dataset. The accuracy is 82.7%, which demonstrates the robustness of this component.

Comparison We selected document-based metrics: 1) ROUGE-1 (Lin, 2004), 2) ROUGE-L (Lin, 2004), 3) BLEU (Papineni et al., 2002), and 4) BertScore (Zhang et al., 2020a); image-based metric: CLIPScore (Hessel et al., 2021); and combined metric: comCLIPBERTScore (Wan and Bansal, 2022b) as baseline metrics. To the best of our knowledge, CLIPBERTScore has been identified as the best metric for evaluating faithfulness in previous assessments of multi-modal summarization. Besides, we also conduct different ablation methods derived from our frameworks based on image input and document input, respectively.

Table 1: Pearson correlation coefficients between automatic metrics and human judgments of factuality on MMSS.

Metric	Document	Image	Combined
BLEU	0.15	-	0.15
ROUGE-1	0.23	-	0.23
ROUGE-L	0.13	-	0.13
BERTScore	0.45	-	0.45
CLIPScore	-	0.13	0.13
CLIPBERTScore	0.45	0.13	0.42
FALLACIOUS (reference-based)	0.48	-	0.48
FALLACIOUS (reference-free)	0.63	0.14	0.51

Table 2: Pearson correlation coefficients between automatic metrics and human judgments of factuality on CEPsum.

Metric	Document	Image	Combined
BLEU	0.52	-	0.52
ROUGE-1	0.49	-	0.49
ROUGE-L	0.53	-	0.53
BERTScore	0.63	-	0.63
CLIPScore	-	0.21	0.21
CLIPBERTScore	0.63	0.21	0.66
FALLACIOUS (reference-based)	0.78	-	0.78
FALLACIOUS (reference-free)	0.78	0.31	0.88

We computed Pearson correlation coefficients between the human evaluation results and all atomic metrics (all baselines and our proposed metrics) in terms of factuality. From Table 1 and 2, we have two observations: (1) Compared with document-based metrics, image-based metrics always perform worse. The potential reason may be that the summary has more overlapped content with the input text rather than the input image. (2) Our frameworks achieved the best performance compared to other baselines, which demonstrates the superiority of our method. In addition, we also use the ChatGPT for evaluation on MMSS and achieve a 0.48/0.47 Pearson correlation coefficient for reference-free/based metric which surpasses all the baselines and demonstrates our framework has the generalization ability.

5 Conclusion

We propose two fine-grained and explainable factuality evaluation frameworks (FALLACIOUS) for multimodal summarization in terms of reference-based and reference-free sceneries. In the human evaluation, we found that our proposed frameworks can achieve good performance compared with other multimodal summarization metrics.

Limitations

Currently, our work relies on Large Language Models (i.e., GPT-4) for evaluation, which is expensive when you evaluate massive datasets. In addition, we used the changing LLMs and the black-box model, which may cause difficulties in reproductivity. The LLMs used in the evaluation metrics may introduce propagated bias.

References

Jingqiang Chen and Hai Zhuge. 2018. [Abstractive text-image summarization using multi-modal attentional hierarchical RNN](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 4046–4056. Association for Computational Linguistics.

Esin Durmus, He He, and Mona T. Diab. 2020. [FEQA: A question answering evaluation framework for faithfulness assessment in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 5055–5070. Association for Computational Linguistics.

Alexander R. Fabbri, Chien-Sheng Wu, Wenhao Liu, and Caiming Xiong. 2022. [Qafacteval: Improved qa-based factual consistency evaluation for summarization](#). 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 2587–2601. Association for Computational Linguistics.

Tanya Goyal and Greg Durrett. 2021. [Annotating and modeling fine-grained factuality in summarization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 1449–1462. Association for Computational Linguistics.

Michael Hanna and Ondrej Bojar. 2021. [A fine-grained analysis of bertscore](#). In *Proceedings of the Sixth Conference on Machine Translation, WMT@EMNLP 2021, Online Event, November 10-11, 2021*, pages 507–517. Association for Computational Linguistics.

Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. 2021. CLIPScore: a reference-free evaluation metric for image captioning. In *EMNLP*.

Chaoya Jiang, Rui Xie, Wei Ye, Jinan Sun, and Shikun Zhang. 2023. [Exploiting pseudo image captions for multimodal summarization](#). In *Findings of the Association for Computational Linguistics: ACL 2023*,

Toronto, Canada, July 9-14, 2023, pages 161–175. Association for Computational Linguistics.

Liqiang Jing, Yiren Li, Junhao Xu, Yongcan Yu, Pei Shen, and Xuemeng Song. 2023. [Vision enhanced generative pre-trained language model for multi-modal sentence summarization](#). *Mach. Intell. Res.*, 20(2):289–298.

Wojciech Kryscinski, Bryan McCann, Caiming Xiong, and Richard Socher. 2020. [Evaluating the factual consistency of abstractive text summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 9332–9346. Association for Computational Linguistics.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7871–7880. Association for Computational Linguistics.

Haoran Li, Peng Yuan, Song Xu, Youzheng Wu, Xiaodong He, and Bowen Zhou. 2020a. [Aspect-aware multimodal summarization for chinese e-commerce products](#). 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 8188–8195. AAAI Press.

Haoran Li, Junnan Zhu, Tianshang Liu, Jiajun Zhang, and Chengqing Zong. 2018. Multi-modal sentence summarization with modality attention and image filtering. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 4152–4158. ijcai.org.

Haoran Li, Junnan Zhu, Jiajun Zhang, Xiaodong He, and Chengqing Zong. 2020b. [Multimodal sentence summarization via multimodal selective encoding](#). In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 5655–5667. International Committee on Computational Linguistics.

Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. 2023a. [BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 19730–19742. PMLR.

Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. 2023b. [BLIP-2: bootstrapping language-image](#)

- pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 19730–19742. PMLR.
- Mingzhe Li, Xiuying Chen, Shen Gao, Zhangming Chan, Dongyan Zhao, and Rui Yan. 2020c. **VMSMO: learning to generate multimodal summary for video-based news articles**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 9360–9369. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. **ROUGE: A package for automatic evaluation of summaries**. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Dengtian Lin, Liqiang Jing, Xuemeng Song, Meng Liu, Teng Sun, and Liqiang Nie. 2023. **Adapting generative pretrained language model for open-domain multimodal sentence summarization**. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, July 23-27, 2023*, pages 195–204. ACM.
- Zixian Ma, Jerry Hong, Mustafa Omer Gul, Mona Gandhi, Irena Gao, and Ranjay Krishna. 2023. **@CREPE: can vision-language foundation models reason compositionally?** In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 10910–10921. IEEE.
- Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. **Factscore: Fine-grained atomic evaluation of factual precision in long form text generation**. *CoRR*, abs/2305.14251.
- Shruti Palaskar, Jindrich Libovick’y, Spandana Gella, and Florian Metze. 2019. **Multimodal abstractive summarization for how2 videos**. 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 6587–6596. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **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, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. **Learning transferable visual models from natural language supervision**. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Thomas Scialom, Paul-Alexis Dray, Sylvain Lamprier, Benjamin Piwowarski, Jacopo Staiano, Alex Wang, and Patrick Gallinari. 2021. **Questeval: Summarization asks for fact-based evaluation**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 6594–6604. Association for Computational Linguistics.
- Xuemeng Song, Liqiang Jing, Dengtian Lin, Zhongzhou Zhao, Haiqing Chen, and Liqiang Nie. 2022. **V2P: vision-to-prompt based multi-modal product summary generation**. In *SIGIR ’22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022*, pages 992–1001. ACM.
- Xiangru Tang, Arjun Nair, Borui Wang, Bingyao Wang, Jai Desai, Aaron Wade, Haoran Li, Asli Celikyilmaz, Yashar Mehdad, and Dragomir R. Radev. 2022. **CONFIT: toward faithful dialogue summarization with linguistically-informed contrastive fine-tuning**. 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 5657–5668. Association for Computational Linguistics.
- David Wan and Mohit Bansal. 2022a. **Evaluating and improving factuality in multimodal abstractive summarization**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 9632–9648. Association for Computational Linguistics.
- David Wan and Mohit Bansal. 2022b. **Evaluating and improving factuality in multimodal abstractive summarization**. In *EMNLP 2022*.
- Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020. **Asking and answering questions to evaluate the factual consistency of summaries**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 5008–5020. Association for Computational Linguistics.
- Yuexiang Xie, Fei Sun, Yang Deng, Yaliang Li, and Bolin Ding. 2021. **Factual consistency evaluation for text summarization via counterfactual estimation**. In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*, pages 100–110. Association for Computational Linguistics.

Litian Zhang, Xiaoming Zhang, and Junshu Pan. 2022. [Hierarchical cross-modality semantic correlation learning model for multimodal summarization](#). In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 11676–11684. AAAI Press.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020a. [Bertscore: Evaluating text generation with bert](#).

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020b. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Zijian Zhang, Chenxi Zhang, Qinpei Zhao, and Jiangfeng Li. 2021. [Abstractive sentence summarization with guidance of selective multimodal reference](#). *CoRR*, abs/2108.05123.

Chenguang Zhu, William Hinthorn, Ruochen Xu, Qingkai Zeng, Michael Zeng, Xuedong Huang, and Meng Jiang. 2021. [Enhancing factual consistency of abstractive summarization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 718–733. Association for Computational Linguistics.

Junnan Zhu, Haoran Li, Tianshang Liu, Yu Zhou, Jiajun Zhang, and Chengqing Zong. 2018. [MSMO: multimodal summarization with multimodal output](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 4154–4164. Association for Computational Linguistics.

Junnan Zhu, Yu Zhou, Jiajun Zhang, Haoran Li, Chengqing Zong, and Changliang Li. 2020. [Multimodal summarization with guidance of multimodal reference](#). 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 9749–9756. AAAI Press.

A Details of Human Evaluation

For the dataset MMSS, each sample has a maximum of 12 QA pairs, and a minimum of 4 QA pairs, with an average of 8.34 QA pairs per text, totaling 1668 QA pairs. On average, the length of the source sentence and the target summary are 22 and 8, respectively. Similarly, for the reference-free evaluation framework, we also utilized these 200 samples. For this framework, the number of generated questions is 3. For the dataset CEPsum, each sample has a maximum of 15 QA pairs, and a minimum of 5 QA pairs, with an average of 10.2 QA pairs per text, totaling 2040 QA pairs. The average token lengths of the input document and output summary are 316 and 79, respectively. Similarly, for the reference-free evaluation framework, we also utilized these 200 samples. For this framework, the number of generated questions is 6.

To explore the reliability of the QA and VQA models, we employ annotators to answer all the generated questions from the two human evaluation datasets. Specifically, the annotator manually annotated answers for each generated question based on the original text and images separately. In addition, to verify the effectiveness of our frameworks, for each sample (original text, image, model summary), the annotator subjectively conducted manual scoring for the faithfulness of the model summaries on a scale of 1 to 5, representing completely incorrect, mostly incorrect, roughly equal parts correct and incorrect, mostly correct, and completely correct, respectively.

B More Analysis

For the reference-based framework, we found the ground truth replaces words or rephrases input documents, which makes the generated questions more difficult compared to the reference-free framework. This may cause the performance of a reference-free framework to be better than a reference-based framework. Compared to text, images show a low correlation with humans, which has also been found in another dataset (Wan and Bansal, 2022a). Therefore, we argue that our method can be applied to more datasets. In addition, results on another model V2P, and another dataset CEPsum show our method is reliable and can be transferred to other datasets.

CLIPBERTScore is a simple weighted combination of CLIPSCORE and BERTScore. However, it’s important to note that while BERTScore

tends to focus more on the grammatical correctness of sentences, it may not necessarily prioritize factual accuracy (Hanna and Bojar, 2021). Furthermore, CLIPScore has its limitations, notably in accurately modeling entities (Radford et al., 2021) and in conducting compositional reasoning (Ma et al., 2023). These limitations can lead to reliability issues, often resulting in inaccurate outcomes. As a result, the effectiveness of CLIPBERTScore in delivering accurate factuality scores is potentially compromised. In contrast, our proposed metric offers a more nuanced approach. It is not only fine-grained, allowing for a detailed analysis, but also interpretable, providing clear insights into its scoring process.

C Prompt

We illustrate our prompt for question generation in Table 3.

D Evaluation Example

We presented reference-free and reference-based evaluation metrics in Figure 3 and Figure 4.

Table 3: Prompt for question generation and question answering.

Question Generation for the Task of With Ground Truth:
<p>User: Now, you will receive a news article. Please generate questions and provide answers based on the information in the news article. Your questions should be proportional to the amount of information and the number of named entities in the article. The question should be a Yes-No Question. The answers should be "yes," "no," or "not provided," and the answer is based on news article. In other words, you do not need to ask open-ended questions. All questions must have answers, either "yes," "no," or "not provided." Please intentionally ask some incorrect questions to verify the accuracy of the answers. For example, you can intentionally create incorrect answers using names or places that have not appeared before to obtain QA pairs with answers as "NO.". Try to maintain a roughly equal ratio of "yes," "no," and "not provided" answers and cover as much content as possible. The format for providing answers should be ["Question": question, "Answer": answer, "Question": question, "Answer": answer, "Question": question, "Answer": answer, "Question": question, "Answer": answer]. The news article is: "u.s. , albanian , croatian and macedonian soldiers began a monthlong joint military exercise tuesday , the u.s. embassy said tuesday ."</p>
<p>GPT: ["Question": "Did the joint military exercise begin on Tuesday?", "Answer": "Yes" , "Question": "Is the joint military exercise only between the U.S. and Albania?", "Answer": "No" , "Question": "Is Croatian military involved in the joint exercise?", "Answer": "Yes" , "Question": "Did the French military participate in the joint exercise?", "Answer": "No" , "Question": "Was the announcement made by the U.S. Embassy?", "Answer": "Yes" , "Question": "Did the joint military exercise last for a year?", "Answer": "No" , "Question": "Did Macedonian soldiers participate in the joint exercise?", "Answer": "Yes" , "Question": "Did the news summary provide information about the purpose of the exercise?", "Answer": "No" , "Question": "Did the exercise commence on Wednesday?", "Answer": "No" , "Question": "Is the joint exercise expected to last for a month?", "Answer": "Yes"]</p>
Question Generation for the Task of Without Ground Truth:
<p>User: I will now provide you with a news summary. Please generate three yes/no questions regarding the entities and relationships within this sentence. The answers to these three questions should all be true. Please present the questions in the format of ["question1", "question2", "question3"]. The summary is: "u.s. albanian croatian and macedonian soldiers joint military exercise.s."</p>
<p>GPT: ["Are U.S. soldiers participating in a joint military exercise with Albanian soldiers?", "Are U.S. soldiers participating in a joint military exercise with Croatian soldiers?", "Are U.S. soldiers participating in a joint military exercise with Macedonian soldiers?"]</p>
Question Answering:
<p>User: I will provide you with a news segment and a question; please provide the answers to the questions in the form of 0 or 1, where 0 represents no and 1 represents yes. news: "u.s. , albanian , croatian and macedonian soldiers began a monthlong joint military exercise Tuesday, the u.s. embassy said Tuesday. question: Are U.S. soldiers participating in a joint military exercise with Albanian soldiers?"</p>
<p>GPT: 1</p>

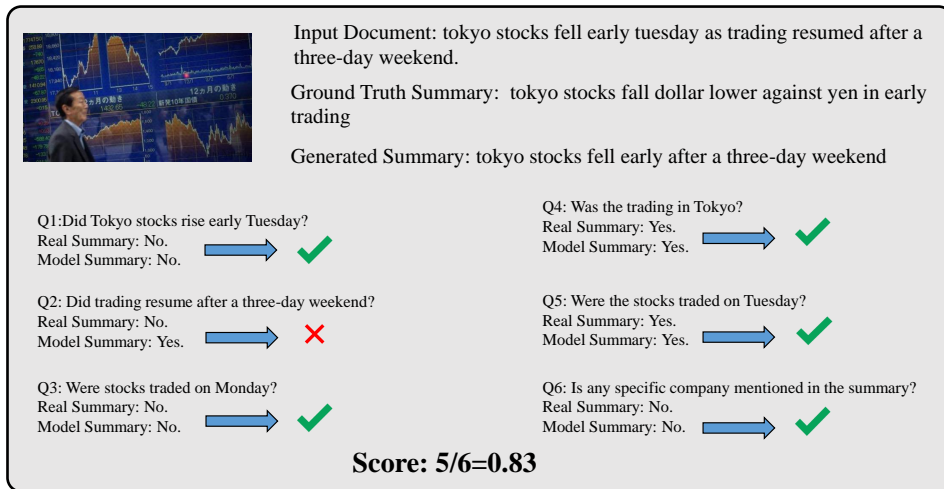


Figure 3: The reference-based evaluation example.

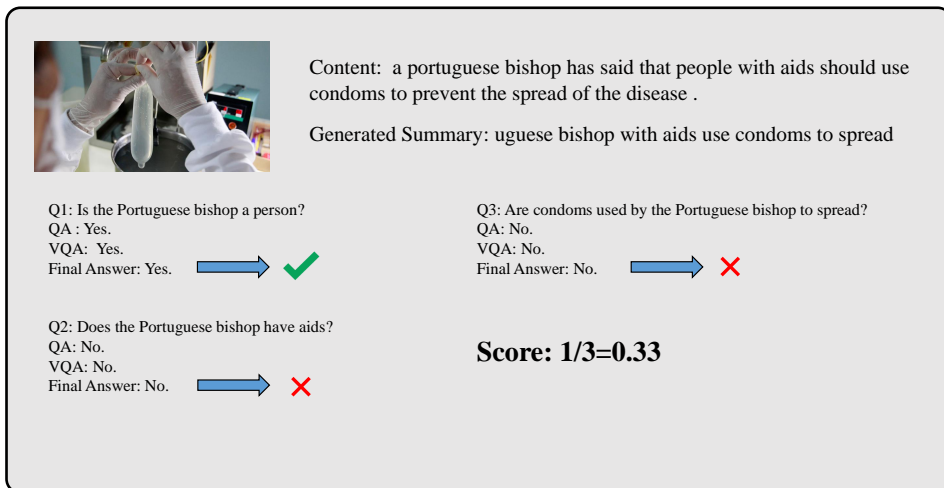


Figure 4: The reference-free evaluation example. QA denotes the answer based on textual modality and VQA denotes the answer based on visual modality. The final answer is the combination of QA and VQA. Score means the score of the proposed metric