

Generating Faithful and Salient Text from Multimodal Data

Tahsina Hashem¹, Weiqing Wang¹, Derry Tanti Wijaya²,
Mohammed Eunus Ali³, Yuan-Fang Li¹

¹Department of Data Science & AI, Monash University, Australia

²Department of Data Science, Monash University, Indonesia

³Department of CSE, Bangladesh University of Engineering and Technology, Bangladesh

{tahsina.hashem, Teresa.Wang, derry.wijaya, yuanfang.li}@monash.edu;

eunus@cse.buet.ac.bd

Abstract

While large multimodal models (LMMs) have obtained strong performance on many multimodal tasks, they may still hallucinate while generating text. Their performance on detecting salient features from visual data is also unclear. In this paper, we develop a framework to generate faithful and salient text from mixed-modal data, which includes images and structured data (represented in knowledge graphs or tables). Specifically, we train a small *vision critic model* to identify hallucinated and non-salient features from the image modality. The critic model also generates a list of salient image features. This information is used in the *post editing* step to improve the generation quality. Experiments on two datasets show that our framework improves LMMs' generation quality on both faithfulness and saliency, outperforming recent techniques aimed at reducing hallucination. The dataset and code are available at <https://github.com/TahsinaHashem/FaithD2T>.

1 Introduction

In many real-world scenarios, data is presented in mixed modalities, in which complementary information is contained. Examples include product brochures, scientific/technical publications, and news articles. Structured data-to-text generation is the task of generating natural language sentences from the data in a structured format, such as tables, knowledge graphs, or databases. Researchers have proposed several models to make this structured information more accessible to humans, aiming to generate fluent, informative, and faithful text descriptions or summaries from the structured data. This task has a wide range of applications across different industries and domains i.e. house advertising, financial reporting, automated journalism, medical reporting, e-commerce product descriptions, generating biographies, etc.

Significant progress has been made in data-to-text generation tasks. Several well-known models have utilized pre-trained language models (PLMs) such as BART (Lewis et al., 2019), T5 (Raffel et al., 2020) or GPT (Radford et al., 2019) with appropriate structure-aware frameworks (Colas et al., 2022; Han and Shareghi, 2022; Li et al., 2024) to generate text descriptions from the structured data. However, the importance of multimodal input with structured data was not extensively addressed. The problem was explored on a small scale by (Gatti et al., 2022; Yang et al., 2023). Their proposed model aimed to generate a one-line summary sentence from a given table and an associated image. They showed that integrating vision data with structured data would lead to more informative and relevant text. However, the research did not consider their generation task's saliency and faithfulness.

Recently, several open-sourced large multimodal models (LMMs) (Liu et al., 2023d; Zhu et al., 2023; Dai et al., 2023) show promising performance in a variety of multimodal tasks (Bai et al., 2023; Liu et al., 2024; Lu et al., 2022; Yin et al., 2023a; Gupta et al., 2023) i.e. image captioning, visual question answering, multimodal conversation, cross-modal retrieval, etc. In this research, we exploit these powerful LMMs to generate text from structured data (knowledge graph and table) with images. We have examined the performance of two prominent LMMs, LLaVA-1.5 (Liu et al., 2023d) and MiniGPT4 (Zhu et al., 2023) on two advertising multimodal (structured data with images) datasets i.e. the real-estate house dataset (Das et al., 2021) and the e-commerce product dataset (Shao et al., 2019). The models generate good-quality advertising text but have two types of limitations: (1) generate some hallucinated information that is not aligned with the vision input; (2) unable to detect salient image features. These limitations hamper the faithfulness and saliency of the generated text.

Figure 1 shows an example input and output

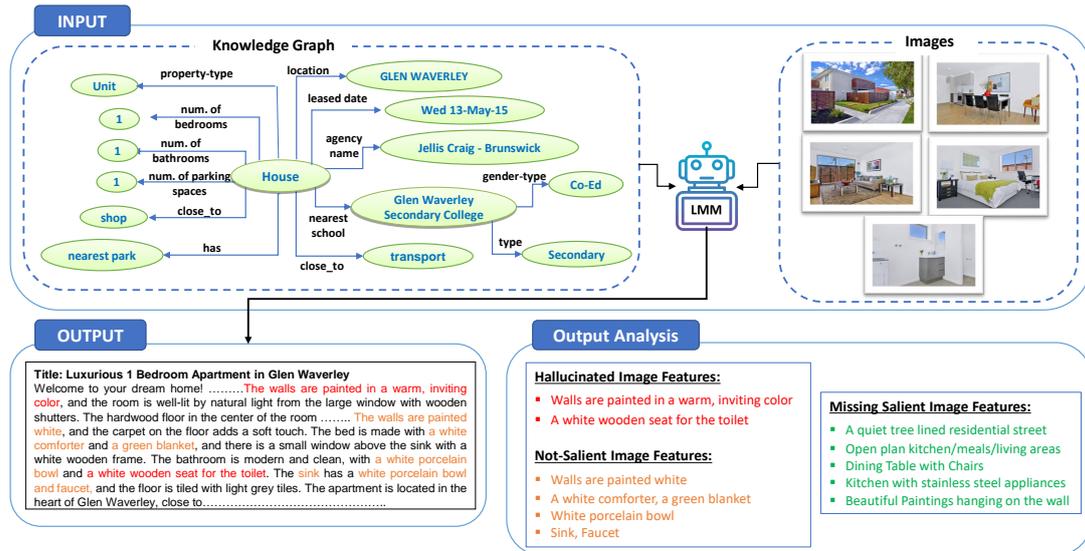


Figure 1: A Sample Input and Output of an LLM: MiniGPT4. The Output Analysis lists the errors.

of an LLM: MiniGPT4. The input consists of a small KG about a house, which contains information on its internal features and neighborhood, and the corresponding images of the house, which gives a detailed outlook of the properties, from a real-world real-estate KG (Das et al., 2021). The output shows the text generated by the LLM. The generated text describes the graph features of the house accurately but struggles to describe the image features accurately. The Output Analysis lists some limitations of LLM: The LLM mentions some features (i.e. hallucination, highlighted in red) that are not aligned with the input images. The LLM also lists some features (i.e. not-salient features, highlighted in orange) that deteriorate the saliency of the generated text while missing some features in the ground-truth text (i.e. salient features, highlighted in green) that are important to make the text attractive for advertising purposes.

Visual hallucination problems of LLMs cause a serious negative impact on visual-to-text generation and reasoning tasks (Liu et al., 2023a; Wang et al., 2023a; Gunjal et al., 2024a; Jing et al., 2023). Researchers have started proposing different strategies (Liu et al., 2023a; Wang et al., 2024; Sun et al., 2023a; Zhou et al., 2024; Yin et al., 2023b) to reduce object hallucinations. Most of the techniques are based on instruction-tuning (Liu et al., 2023a,b) or filtering the hallucination information from the training data (Wang et al., 2024; Yu et al., 2023) and then fine-tuning the models with the revised version of the dataset. This process of preparing such a good number of high-quality instructions or datasets is time-consuming and costly. Some re-

searchers (Sun et al., 2023a) have utilized reinforcement learning from human feedback in training the LLMs using reward models. Another alternate way of mitigating hallucination is post-hoc detection and correction frameworks (Zhou et al., 2024; Yin et al., 2023b). These methods are cost-friendly and showed good performance in mitigating hallucinations in the generated text.

Our proposed framework follows the detection and correction strategy but instead of fine-tuning LLMs, we train a small vision language model (VLM) (Li et al., 2023a) as a transparent vision critic model that can detect the errors of the text generated by LLMs with an explanation and list the missing salient image features of the text generated by LLMs. Finally, we update the generated text using LLM from the feedback of the critic model using an appropriate prompt.

The contributions of our research work are:

- Propose a novel task of generating faithful and salient natural language text from structured data and images.
- Design a framework to train a small vision model to act like an interpretable vision critic model that can verify the faithfulness and saliency of the features as well as list the missing salient image features of the text generated by LLMs.
- Experimental Results demonstrate the effectiveness of our model over existing baselines.

2 Related Work

2.1 Multimodal Data to Text generation

Several structure-enhanced pre-trained language models (Han and Shareghi, 2022; Li et al., 2024; Tang et al., 2023; Liu et al., 2022) showed good performance in structured data-to-text generation tasks. However, very few works (Gatti et al., 2022; Yang et al., 2023) have been done in multimodal data-to-text generation tasks. An initial attempt was made by Gatti et al. (Gatti et al., 2022) to generate a one-line summary sentence from vision-augmented tabular data. They proposed a VT3 multimodal transformer that consists of a BART model (Lewis et al., 2019) and a vision transformer (Liu et al., 2021), that can generate text auto-regressively. A different approach was proposed to overcome a large amount of annotated training data requirement (Yang et al., 2023). They proposed a multimodal prompt learning framework to accurately generate titles for novel products with limited labels. However, both models aim to generate a one-line summary sentence. They cover a small number of vision features without verifying the saliency and faithfulness of their generated text. Whereas, in our problem, we focus on generating a long advertising text that should contain all the salient and faithful features of the vision data.

Recently, the large multimodal models (Zhu et al., 2023; Liu et al., 2023d; Dai et al., 2023; Ye et al., 2023) have shown remarkable success in various multimodal tasks such as image captioning (Lin et al., 2014), visual question-answering (VQA) (Antol et al., 2015) and multimodal conversation (Liu et al.). Hence, our research work exploits these powerful LMMs to generate salient and faithful text from multimodal data.

2.2 Hallucination in LMMs

Although LMMs demonstrate strong performance across multiple benchmark tasks and produce quality results, they struggle with the problem of visual hallucination. This issue occurs when the generated responses do not align with the visual input. Researchers investigated this phenomenon in the realm of object hallucination (Li et al., 2023b; Liu et al., 2023c; Biten et al., 2022), where the generated content features objects that do not match or are not present in the input image. Recently, it has been shown (Zhai et al., 2023) that this multimodal hallucination happens because the vision encoder does not accurately ground images. They tend to

depend more on their built-in knowledge rather than the visual input provided. Furthermore, empirical studies by Wang et al. (Wang et al., 2023b), have shown that these models focus more on previously generated tokens than on the image features.

2.3 Hallucination Mitigation of LMMs

Researchers have already proposed a number of alternative strategies to minimize the visual hallucination problem of LMMs. Some focus on improving the quality of instruction tuning data. LRV-Instruction dataset (Liu et al., 2023c), VIGC (Wang et al., 2024), M-HalDetect (Gunjal et al., 2024b) are examples of such high-quality prepared datasets. Some tried to refine the model training techniques like reinforcement learning from human feedback (RLHF) in LLaVA-RLHF (Sun et al., 2023b), or optimization models in FDPO (Gunjal et al., 2024b). Some researchers apply post hoc detection and correction strategies such as LURE (Zhou et al., 2024) that is based on object co-occurrence, uncertainty, and position in text; and Woodpecker (Yin et al., 2023b) that extracts key concepts and validates the visual knowledge using object detector and VQA model.

For a more cost-effective approach, we adopt the post-hoc detection and correction approach. We train a small pre-trained VLM that can be used in cooperation with LLM to mitigate both the visual hallucinated features and non-salient features.

3 Method

To generate high-quality text, finetuning LMM may not be feasible for proprietary models, and it may not be practical due to the prohibitively high resource and data requirements. Thus, we propose a cost-effective post-hoc detection and correction approach that trains a small VLM to act as a *vision critic model* that identifies errors in the LMM-generated text. With the feedback provided by the critic model, a capable LLM (such as GPT-3.5) is then employed to update the text using this feedback.

Figure 2 depicts the overall architecture of our proposed method. Given the text generated from the mixed-modal data by the LMM, we first prompt a capable LLM (such as GPT-3.5) to extract the list of image features from the text by filtering out features from the structured data. With this list and the images as input, we employ our trained vision critic model to identify hallucinated image features and

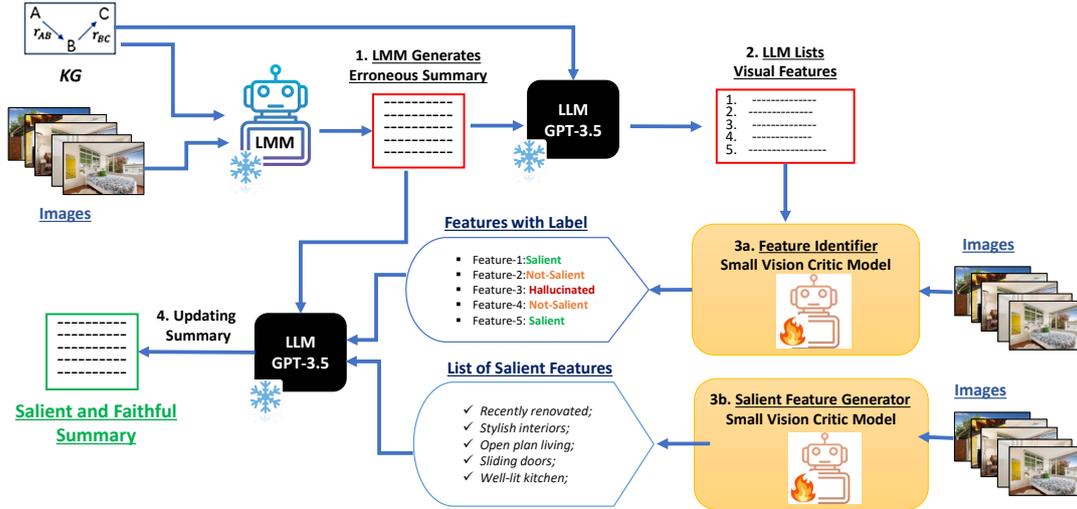


Figure 2: The Pipeline of our Framework for Salient and Faithful Multimodal Data to Text Generation 1) Generating Text using LMM 2) Extracting Image Features from the Text using GPT-3.5 3) Trained Vision Critic Model gives feedback to LMM 4) LLM update the Text by making corrections.

non-salient image features in the text. The critic model also generates salient image features that are missing from the text. Finally, we prompt the LLM to remove the hallucinated and non-salient features from the text and append the missing salient image features to the text.

3.1 Problem Formulation

Given a training dataset $\mathcal{D} = (X, Y)$, in which $X = \{(s_1, i_1), (s_2, i_2), \dots, (s_{|\mathcal{D}|}, i_{|\mathcal{D}|})\}$ is a mixed-modal dataset that consists of pairs of structured data s_i (i.e. knowledge graphs or tables) and (multiple) images i_i , and $Y = \{y_1, \dots, y_{|\mathcal{D}|}\}$ is a set of reference text for each x_i , our aim is to train a model that generates a text passage \hat{y}_j for $x_j = (s_j, i_j)$ that is both faithful to s_j and i_j and contains the salient image features in i_j . Note that Y may contain hallucinated information.

We assume the structured data is either a knowledge graph or a table. Let $KG = (V, E)$ represent a knowledge graph, where $V = \{e_1, e_2, \dots, e_{|V|}\}$ represents the entity set and $E = \{r_{ij}\} \subseteq V \times V$ represents the relations connecting the entities. For the tabular data, let $T = \{(a_1, v_1), (a_2, v_2), \dots, (a_m, v_m)\}$ represents a table with m number of attribute-value pairs. Every type of structured data contains an image set. Let $I = \{i_1, i_2, \dots, i_l\}$ represents the corresponding image set.

3.2 Training a Small Vision Language Model

We choose a small vision language model (VLM) BLIP-2 (Li et al., 2023a) to act as a critic model.

BLIP-2 addresses the modality gap by employing a lightweight Querying Transformer (Q-Former). BLIP-2 utilizes a generic and efficient pretraining strategy that bootstraps vision-language pretraining from off-the-shelf frozen pretrained image encoders and frozen large language models (LLMs). It shows good performance on visual question-answering tasks, image captioning tasks, image-text retrieval tasks and visual commonsense reasoning tasks (Park et al., 2024).

Recently, it has been shown (Kim et al.) that Parameter-efficient fine-tuning (PEFT) (Mangrulkar et al., 2022) maintains competitive performance while requiring much less computational memory. Thus, we apply LoRA (Hu et al., 2021) to the Q-Former and the base LLMs, Flan-T5-XL, of the BLIP-2 model. This allows us to fine-tune the BLIP-2 model in a cost-effective way.

In the following two subsections we discuss our training process in detail for the two critic tasks:

3.2.1 Classifying Image Feature

We observe that LMMs cannot reliably distinguish salient features from non-salient features and hallucinated features, degrading the quality of generated text. Thus, we train the vision critic model to become an expert in detecting whether a feature is salient, non-salient or hallucinated given an input image. We formulate this task as a generation problem, where a set $x = (i, f)$ is given to BLIP-2 vision critic model, with i being an image and f being a feature, and the output is a textual output of the label

$y \in \{salient, non-salient, hallucinated\}$ with an appropriate explanation. We use a standard conditional language modeling loss function:

$$L_{CE} = - \sum_{i=1}^n \log P(y_i | y_{<i}, X) \quad (1)$$

Our training data consists of a set of labeled image-feature pairs along with the corresponding rationales for the three categories.

3.2.2 Listing Salient Image Features

We train our vision critic model to identify the important i.e. salient features of a given image. We formulate this task also as a generation problem, where the vision critic model outputs a list of salient features, $S_i = \{[s_1]; [s_2], \dots; [s_m]\}$ given an image i . We fine-tune the vision critic model by maximizing the log-likelihood:

$$L_{S_i} = -\mathbb{E}_{(I, S_i) \sim \mathcal{D}'} \log P(S_i | i) \quad (2)$$

Here, the training dataset $D' = (i, S_i)$ consists of an image and a list of salient image features. The training data generation process is discussed in detail below.

3.2.3 Training Data Generation

We prepare labeled data (i.e., image features labeled with salient, non-salient, and hallucinated) for training the critic model to classify the image features. To generate this data, we take samples of ground-truth texts and the corresponding LLMs-generated texts. We also prepare image-features pairs where each pair is a list of salient features for the corresponding image. This data is used to train the critic model to generate salient features of an image. The entire training data generation process involves the following steps:

(1) Extracting features from text: Both the ground-truth text and the generated text contain features from the structured data and the images in an aggregated form. We use an LLM, i.e. GPT-3.5, to list the features one by one from every sentence of the texts following some in-context examples. An example prompt can be found in Figure: 5

(2) Listing visible and non-visible features: Both the ground-truth text and the LLM-generated text contain hallucinated information. To prepare labeled visible (i.e., salient or non-salient) and not visible (i.e., potentially hallucinated) features from the images for training the critic model, we prompt GPT-4V with input images and the list of extracted

features from (1) to verify whether the feature is visible or not visible in the input images.

(3) Listing hallucinated features: We input GPT-3.5 the structured data and the list of not visible features that we obtain from (2) and prompt it to list the features not aligned with the structured data. Features that are not aligned with the structured data will be labeled as hallucinated features since they are neither visible in the image nor exists in the structured data.

(4) Listing salient and non-salient features: We ask GPT-3.5 to compare the visible image features in the LMM-generated text with the ground-truth visible image features. All visible features from the ground-truth text are salient image features. Visible features in the generated text that are similar to any of the features mentioned in the ground-truth text or the structured data are also salient features. The remaining visible features in the LMM-generated text are the non-salient features.

(5) Generating rationale for feature labels:

After preparing the labeled image-features pairs for salient, non-salient, and hallucinated categories, we prompt GPT-3.5 to generate a one-sentence explanation for why a feature might be labeled salient or not salient. For the hallucinated feature, we use the default explanation that: "The feature is not visible in the image". These rationales make our vision critic model interpretable and leads to improve the accuracy of the critic model in feature labeling tasks.

All the prompt templates are shown in appendix B.

3.3 Post-hoc Text Editing from the Feedback given by the Critic Model

We design an appropriate prompt to utilize an LLM (GPT-3.5) for updating the LMM-generated text according to the feedback of the trained vision critic model. The update operation is done in two steps. Firstly, non-salient and hallucinated image features are pruned from the text. Secondly, salient image features are appended to the pruned text. Figure 10 shows the prompt template in the Appendix.

4 Experiments

4.1 Dataset

We conduct experiments and evaluation on two multimodal data-to-text generation datasets: the House dataset of real-estate house listings (Das et al., 2021), containing images and knowledge graphs; and the Product dataset of Chinese e-

commerce clothing products (Shao et al., 2019), containing images and attribute-value information in a tabular format. In both datasets, the ground-truth text contains a significant amount of hallucinated information, making the task of generating faithful text especially challenging. Table 1 shows brief statistics of these two datasets.

Table 1: Statistics of the House and Product Datasets.

Dataset	Avg. # triples	Avg. # images	Avg. Text length
House	22.8	3	153
Product	7.4	1	110

The House dataset is a large real-estate and point-of-interests (POI) dataset of Melbourne, Australia (Das et al., 2021). It includes 53,220 records of house sales transactions from 2013 to 2015. It consists of three types of POIs, namely regions, schools, and train stations, along with their corresponding features. Each sample in the dataset includes (1) a ground-truth advertisement text, (2) a KG describing house and POI features, and (3) multiple images of the house. However, the given ground-truth text contains a significant level of hallucinated information. We use 3,100 samples for training the vision critic model and 100 test samples for testing the performance of the critic model. We prepared labeled image-feature pairs according to section 3.2.3. Details of the training data-split ratio are shown in Appendix E.

The Product dataset is from a Chinese e-commerce platform of clothings, consisting of 119K samples of advertising text, a clothing specification table, and a single image of the clothing. Each table is a set of attribute-value pairs describing a piece of clothing. The ground-truth advertising text also contains hallucinated information. For training of the critic model, we have used 4,700 samples and for testing, we used 340 samples. We prepared labeled image-feature pairs according to section 3.2.3. Details of the training data-split ratio are shown in Appendix E.

4.2 Baseline Models

Two prominent LMMs, namely MiniGPT-4 (Zhu et al., 2023) and LLaVA-1.5 (Liu et al., 2023d) are used as the baseline models. We also compare with two recent post-hoc hallucination detection and correction models, LURE (Zhou et al.,

2024) and Woodpecker (Yin et al., 2023b). Due to resource constraints, we only experiment with LURE on MiniGPT-4. The backbone model of LURE is MiniGPT-4. Woodpecker utilizes GPT-3.5-turbo as its corrector, grounding DINO (Liu et al., 2023e) as its object detector and BLIP-2-FlanT5-XXL (Li et al., 2023a) as its visual question answering model.

4.3 Preliminary Analysis

We conducted a preliminary analysis of the performance of MiniGPT-4 and LLaVA-1.5 with the 100 test samples of the House and 340 test samples of the product datasets. For each model, we input the structured data and images and prompt it to generate an advertising text passage. The structured data (KG or table) is given in a linearized format for a better understanding by the LMM (Li et al., 2024). As LMMs are unable to accept as input multiple images simultaneously, for the House dataset, we input images one by one and ask the LMM to list the key features of the input image. Detailed prompt templates are shown inref A in the supplementary files. We observe that the LMMs can accurately list features from the structured data, but struggle to list the image features correctly.

The following common errors are observed in the generated texts by the LMMs:

- **Missing salient image features:** LMMs sometimes miss some important image features in the generated text that are essential for advertising purposes. We consider the image features listed in the ground-truth text as the standard salient image features.
- **Hallucinated image features:** LMM-generated text sometimes contain image features that are not present in input images.
- **Non-salient image features:** LMMs sometimes mention features from the images that are not attractive to customers. These features deteriorate the saliency of the text.

Figure 1 shows the text generated by MiniGPT4 from a sample in the House dataset. The output analysis lists the erroneous features (i.e., hallucinated or not salient) in the text as well as the missing salient image features.

4.4 Experimental Settings

In our framework, we keep the LMM and the LLM frozen. We finetune the small vision language model Blip-2 (Li et al., 2023a). We apply PEFT fine-tuning (Mangrulkar et al., 2022) to the BLIP-

Table 2: Main results on the House dataset. **Bold** font denotes the best results for each backbone model.

Model	Saliency				Faithfulness
	BLEU	METEOR	ROUGE-L	BERTScore	CLIP Score
Baseline Model					
MiniGPT4	8.08	25.28	13.44	83.98	23.92
MiniGPT4-Woodpecker	8.15	27.13	13.34	83.91	23.89
MiniGPT4-LURE	11.01	16.63	14.44	84.33	23.86
Our Model					
MiniGPT4-Pruned	10.13	23.14	15.13	84.71	24.30
MiniGPT4-Appended	10.69	28.08	15.58	85.15	24.26
MiniGPT4-Combined	11.98	26.09	16.50	85.36	24.59
Baseline Model					
LLaVA-1.5	11.34	29.82	16.06	85.10	23.92
LLaVA-1.5-Woodpecker	9.81	29.93	14.86	84.66	23.89
Our Model					
LLaVA-1.5-Pruned	13.74	27.36	16.87	85.67	24.29
LLaVA-1.5-Appended	13.52	31.65	17.35	85.72	24.49
LLaVA-1.5-Combined	15.01	29.29	17.33	86.01	24.63

2-FlanT5-XL model (Li et al., 2023a). We apply LoRA (Hu et al., 2021) to both the Q-Former and the base LLMs, Flan-T5-XL. For the House dataset, we fine-tune both critic models (for feature classification and for generating missing salient features respectively) for 25 epochs. For the Product dataset, we fine-tune both critic models for 50 epochs. The batch size for both datasets is set to 16. The maximum length of the output text sequences is set to 350 tokens for the House dataset and 200 tokens for the Product dataset. We adopt Adam (Kingma and Ba, 2014) as the optimizer and set the learning rate to be $5e-5$. We used one A40 48GB GPU for all the experiments.

4.5 Main Results

Our main experiment aims at the faithfulness and saliency of the text generated by LMMs from the mixed-modal data. For saliency evaluation, we consider the image features contained in the ground-truth text as ground-truth salient features. For faithfulness evaluation, we need to pre-process the text to obtain faithful features as the ground-truth text contains hallucinated information. Specifically, we prompt GPT-3.5 to list features from ground-truth text, and prompt GPT-4V to remove hallucinated features from this list (i.e., features that are neither visible in the image nor exist in the structured data). Finally, we prompt GPT-3.5 to generate a paragraph containing the faithful features. Note that, the faithful features in the ground-truth text are also salient. Thus, in this way, we obtain the salient and faithful ground-truth text. The prompt template can be found in the appendix.

We use automatic metrics to measure both faithfulness and saliency of generated text. We employ standard metrics BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), ROUGE-L (Lin, 2004) and BERTScore (Zhang et al., 2019) to measure saliency of the generated text by comparing it with the pre-processed ground-truth text. To verify the faithfulness of the generated text with respect to the input image(s), we utilize the CLIP score (Hessel et al., 2021), which is widely used (Zhou et al., 2024; Jing et al., 2023) to measure text-image alignment.

Table 2 and Table 3 present the results on the House and Product datasets respectively. We evaluate our method, denoted “**-Combined**”, against the baselines and other post-hoc hallucination detection and correction models (LURE and Woodpecker). From the results of both datasets, we can observe that our method achieves the best performance, outperforming the baseline models and other hallucination-reduction techniques on most settings. Specifically, our method not only outperforms Woodpecker and LURE in reducing hallucination (i.e. improving faithfulness), it also achieves the best result in preserving saliency.

Some qualitative examples of pre-processed ground-truth text and the text generated by different models can be found in Appendix F.

4.6 Ablation Studies

To investigate the effect of our two trained critic models, we experiment on both datasets with two variants of our full method (i.e., **-Combined**): “**-Pruned**”, which only removes hallucinated and

Table 3: Main Results on the Product dataset. **Bold** font denotes the best results for each backbone model.

Model	Saliency				Faithfulness
	BLEU	METEOR	ROUGE-L	BERTScore	CLIP Score
Baseline Model					
MiniGPT4	9.49	23.24	15.83	85.63	22.62
MiniGPT4-Woodpecker	10.42	23.79	16.66	86.20	22.99
MiniGPT4-LURE	10.19	20.39	15.42	85.47	22.71
Our Model					
MiniGPT4-Pruned	10.69	21.06	16.48	86.06	22.81
MiniGPT4-Appended	10.19	24.59	16.23	86.15	22.74
MiniGPT4-Combined	11.17	22.51	16.84	86.34	22.96
Baseline Model					
LLaVA-1.5	13.89	24.79	18.52	87.47	23.14
LLaVA-1.5-Woodpecker	12.47	24.74	18.20	86.99	23.19
Our Model					
LLaVA-1.5-Pruned	13.90	21.55	18.48	87.58	23.34
LLaVA-1.5-Appended	13.71	25.58	18.43	87.47	23.18
LLaVA-1.5-Combined	15.07	22.84	18.58	87.61	23.34

non-salient features identified by our critic model; and “-**Appended**”, which only appends missing salient image features generated by our critic model. As we see in Table 2 and Table 3, both variants positively contribute to improving saliency and faithfulness.

We also assess our trained critic models’ (based on fine-tuning BLIP-2 on our training data) performance with the non-fine-tuned BLIP-2 model at the feature-level. Table 4 shows the feature classification accuracy of our trained critic model-3a and non-fine-tuned BLIP-2 model on three types of image features: hallucinated, salient, and non-salient in the test set of the House data. It is observed that although the non-fine-tuned BLIP-2 model achieves equal accuracy in identifying hallucinated features, its performance is significantly worse in identifying salient and non-salient features compared to our trained critic model.

Table 4: Evaluation of image feature classification accuracy into hallucinated (Hal), salient (Sal) and non-salient (Non-Sal) labels on the House Dataset.

Model	Hal	Sal	Non-Sal
Trained BLIP2 Model-3a	96.12	92.93	71.20
Non-fine-tuned BLIP2 Model	96.12	57.32	41.77

Our critic model-3b generates a list of salient features from the input image. We measure the quality of the generated list of salient features in terms of saliency and faithfulness. Table 5 shows the comparison between the two models. We measure the saliency of the generated features list by comparing this generated features

list with the list of ground-truth salient features using Sentence-BERT (SBERT) similarity score (Reimers and Gurevych, 2019). For faithfulness, considering the images, we use CLIPScore (Hessel et al., 2021). The SBERT score shown in Table 5 shows that our model-generated salient features are more similar to ground-truth salient features compared to the salient features generated by the non-fine-tuned BLIP-2 model. The CLIPscore shown in Table 5 shows the generated features are comparably-aligned with the input images.

Table 5: Evaluation of generated salient features on the House Dataset.

Model	SBERT Score	CLIP Score
Trained BLIP2 Model-3b	54.87	27.05
Non-fine-tuned BLIP2 Model	45.01	27.46

5 Conclusion

In this paper, we propose a novel approach to generating text that is both faithful and salient from mixed-modal data that includes images and structured data. To ensure salient and faithful text generation, we train a small vision critic model to: (i) identify the hallucinated, salient and non-salient features, and (ii) generate a list of salient features from images. This information is used in the *post editing* step to improve generation quality. Experimental results on two mixed-modal datasets demonstrate that our framework outperforms recent large multimodal models as well techniques specifically

designed to reduce hallucination in terms of faithfulness and saliency metrics.

Limitation and Future work Our critic model sometimes prunes subjective features such as "Eye-catching", "Amazing opportunity", "Elegant beauty", "Piece of luxury" etc, which are essential for making the advertising text attractive. In future, we will consider this issue. In addition, we also plan to explore the saliency and hallucination problem in other modalities such as videos and audios.

Ethical Considerations

Our model utilizes existing pre-trained vision language model, thus the ethical concerns associated with these models would also be applicable to our proposed framework.

Acknowledgments

This material is based on research sponsored by Defense Advanced Research Projects Agency (DARPA) under agreement number HR0011-22-2-0047. The U.S. Government is authorised to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government.

References

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond.
- Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.
- Ali Furkan Biten, Lluís Gómez, and Dimosthenis Karatzas. 2022. Let there be a clock on the beach: Reducing object hallucination in image captioning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1381–1390.
- Anthony Colas, Mehrdad Alvandipour, and Daisy Zhe Wang. 2022. Gap: A graph-aware language model framework for knowledge graph-to-text generation. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5755–5769.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning. *Preprint*, arXiv:2305.06500.
- Sarkar Snigdha Sarathi Das, Mohammed Eunus Ali, Yuan-Fang Li, Yong-Bin Kang, and Timos Sellis. 2021. Boosting house price predictions using geospatial network embedding. *Data Mining and Knowledge Discovery*, 35:2221–2250.
- Prajwal Gatti, Anand Mishra, Manish Gupta, and Mithun Das Gupta. 2022. Vistot: vision-augmented table-to-text generation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9936–9949.
- Anisha Gunjal, Jihan Yin, and Erhan Bas. 2024a. Detecting and preventing hallucinations in large vision language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18135–18143.
- Anisha Gunjal, Jihan Yin, and Erhan Bas. 2024b. Detecting and preventing hallucinations in large vision language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18135–18143.
- Devaansh Gupta, Siddhant Kharbanda, Jiawei Zhou, Wanhua Li, Hanspeter Pfister, and Donglai Wei. 2023. Cliptrans: Transferring visual knowledge with pre-trained models for multimodal machine translation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2875–2886.
- Jiuzhou Han and Ehsan Shareghi. 2022. Self-supervised graph masking pre-training for graph-to-text generation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4845–4853.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. 2021. Clipscore: A reference-free evaluation metric for image captioning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7514–7528.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2021. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.

- Liqiang Jing, Ruosen Li, Yunmo Chen, Mengzhao Jia, and Xinya Du. 2023. Faithscore: Evaluating hallucinations in large vision-language models. *arXiv preprint arXiv:2311.01477*.
- Sungkyung Kim, Adam Lee, Junyoung Park, Sounho Chung, Jusang Oh, and Jayyoon Lee. Parameter-efficient fine-tuning of instructblip for visual reasoning tasks.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*.
- Shujie Li, Liang Li, Ruiying Geng, Min Yang, Binhua Li, Guanghu Yuan, Wanwei He, Shao Yuan, Can Ma, Fei Huang, et al. 2024. Unifying structured data as graph for data-to-text pre-training. *Transactions of the Association for Computational Linguistics*, 12:210–228.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023b. Evaluating object hallucination in large vision-language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 292–305.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- 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 *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer.
- Ao Liu, Haoyu Dong, Naoaki Okazaki, Shi Han, and Dongmei Zhang. 2022. Plog: Table-to-logic pretraining for logical table-to-text generation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5531–5546.
- Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. 2023a. Aligning large multi-modal model with robust instruction tuning. *arXiv preprint arXiv:2306.14565*.
- Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. 2023b. Mitigating hallucination in large multi-modal models via robust instruction tuning. In *The Twelfth International Conference on Learning Representations*.
- Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. 2023c. Mitigating hallucination in large multi-modal models via robust instruction tuning. In *The Twelfth International Conference on Learning Representations*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023d. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge (january 2024). URL <https://llava-vl.github.io/blog/2024-01-30-llava-next>.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. 2023e. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022.
- Haoyu Lu, Nanyi Fei, Yuqi Huo, Yizhao Gao, Zhiwu Lu, and Ji-Rong Wen. 2022. Cots: Collaborative two-stream vision-language pre-training model for cross-modal retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15692–15701.
- Sourab Mangrulkar, Sylvain Gugger, Lysandre Debut, Younes Belkada, Sayak Paul, and Benjamin Bossan. 2022. Peft: State-of-the-art parameter-efficient fine-tuning methods. <https://github.com/huggingface/peft>.
- 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.
- Jae Sung Park, Jack Hessel, Khyathi Chandu, Paul Pu Liang, Ximing Lu, Peter West, Youngjae Yu, Qiuyuan Huang, Jianfeng Gao, Ali Farhadi, et al. 2024. Localized symbolic knowledge distillation for visual commonsense models. *Advances in Neural Information Processing Systems*, 36.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.

- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992.
- Zhihong Shao, Minlie Huang, Jiangtao Wen, Wenfei Xu, and Xiaoyan Zhu. 2019. Long and diverse text generation with planning-based hierarchical variational model. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3257–3268.
- Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, et al. 2023a. Aligning large multimodal models with factually augmented rlhf. *arXiv preprint arXiv:2309.14525*.
- Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, et al. 2023b. Aligning large multimodal models with factually augmented rlhf. *arXiv preprint arXiv:2309.14525*.
- Tianyi Tang, Junyi Li, Wayne Xin Zhao, and Ji-Rong Wen. 2023. Mvp: Multi-task supervised pre-training for natural language generation. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8758–8794.
- Bin Wang, Fan Wu, Xiao Han, Jiahui Peng, Huaping Zhong, Pan Zhang, Xiaoyi Dong, Weijia Li, Wei Li, Jiaqi Wang, et al. 2024. Vigc: Visual instruction generation and correction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 5309–5317.
- Junyang Wang, Yiyang Zhou, Guohai Xu, Pengcheng Shi, Chenlin Zhao, Haiyang Xu, Qinghao Ye, Ming Yan, Ji Zhang, Jihua Zhu, et al. 2023a. Evaluation and analysis of hallucination in large vision-language models. *arXiv preprint arXiv:2308.15126*.
- Junyang Wang, Yiyang Zhou, Guohai Xu, Pengcheng Shi, Chenlin Zhao, Haiyang Xu, Qinghao Ye, Ming Yan, Ji Zhang, Jihua Zhu, et al. 2023b. Evaluation and analysis of hallucination in large vision-language models. *arXiv e-prints*, pages arXiv–2308.
- Bang Yang, Fenglin Liu, Zheng Li, Qingyu Yin, Chenyu You, Bing Yin, and Yuexian Zou. 2023. Multimodal prompt learning for product title generation with extremely limited labels. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2652–2665.
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. 2023. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*.
- Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. 2023a. A survey on multimodal large language models. *arXiv preprint arXiv:2306.13549*.
- Shukang Yin, Chaoyou Fu, Sirui Zhao, Tong Xu, Hao Wang, Dianbo Sui, Yunhang Shen, Ke Li, Xing Sun, and Enhong Chen. 2023b. Woodpecker: Hallucination correction for multimodal large language models. *arXiv preprint arXiv:2310.16045*.
- Qifan Yu, Juncheng Li, Longhui Wei, Liang Pang, Wentao Ye, Bosheng Qin, Siliang Tang, Qi Tian, and Yueting Zhuang. 2023. Hallucidoctor: Mitigating hallucinatory toxicity in visual instruction data. *arXiv preprint arXiv:2311.13614*.
- Bohan Zhai, Shijia Yang, Xiangchen Zhao, Chenfeng Xu, Sheng Shen, Dongdi Zhao, Kurt Keutzer, Manling Li, Tan Yan, and Xiangjun Fan. 2023. Hallsplit: Rethinking and controlling object existence hallucinations in large vision language models for detailed caption. *arXiv preprint arXiv:2310.01779*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Yiyang Zhou, Chenhang Cui, Jaehong Yoon, Linjun Zhang, Zhun Deng, Chelsea Finn, Mohit Bansal, and Huaxiu Yao. 2024. Analyzing and mitigating object hallucination in large vision-language models. In *The Twelfth International Conference on Learning Representations*.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.

A Prompt Template for Text Generation from LMM

The two prompt templates for House dataset are shown in Figure 3 and in Figure 4

B Prompt Template for Training Data Generation

(1) **Extracting features from text:** Prompt template for Extracting features from the sentence of the text is shown in Figure: 5.

(2) **Listing visible and non-visible features:** Prompt template for listing visible and not visible features from the list of features is shown in Figure: 6.

(3) **Listing hallucinated features:** Prompt template for listing hallucinated features is shown in Figure: 7.

(4) **Listing salient and non-salient features:** Prompt template for listing salient and not salient features is shown in Figure: 8.

(5) **Generating rationale for features:** The prompt template for generating rationale is shown in Figure 9.

C Prompt Design for Post-hoc Text Editing using LLM

The prompt template is shown in Figure: 10 for updating the text using LLM GPT-3.5 according to the feedback of the critic model.

D Prompt Design for Preparing Salient and Faithful Ground-truth Text

We use the prompt in Figure: 11 to extract the faithful and salient features from the hallucinated ground-truth text. Then using the prompt in Figure: 12, we generate a final salient and faithful ground-truth text.

E Data Split Ratio for Training the Critic Model

Table 6: House Data split ratio for Critic Model-3a

Data	#Num. of Instances
Training Samples	9517
-Hallucination Features	5478
-Salient Features	2703
-Not-Salient Features	1336
Validation Samples	1418
-Hallucination Features	845
-Salient Features	409
-Not-Salient Features	164

Table 7: House Data split ratio for Critic Model-3b

Data	#Num. of Instances
Training Samples	10,359
Validation Samples	1155

Table 8: Product Data split ratio for the Critic Model-3a

Data	#Num. of Instances
Training Samples	13202
-Hallucination Features	2529
-Salient Features	6795
-Not-Salient Features	3878
Validation Samples	2219
-Hallucination Features	393
-Salient Features	1243
-Not-Salient Features	583

F Generated Samples

Figure 13 and Figure 14 show qualitative examples of sample graph-images, sample tabular data-image, the pre-processed ground-truth texts, and the texts generated by different models on the House dataset and Product dataset, respectively.

Table 9: Product Data split ratio for Critic Model-3b

Data	#Num. of Instances
Training Samples	4235
Validation Samples	471

#INPUT Prompt:
Consider the uploaded Image of a House
List the Key features of this image for House Advertising.

#OUTPUT:
Key features of image:1:
The key features of this image are: a large glass wall, a wooden deck with chairs, a swimming pool, and a view of the surrounding city.

Key features of image:2:
This image shows a rooftop patio with a large table and chairs, a small table with two chairs, and a large planter with plants. The view from the rooftop is of a cityscape with skyscrapers and trees in the foreground. The sky is clear and sunny, and there is a blue and white striped awning over the patio. The patio is surrounded by a railing made of metal and glass, and there are several large windows on the side of the building. The building is made of concrete and glass, with a modern design.

Key features of image:3:
The image shows a modern, minimalist bathroom with a white wall tiles, a grey toilet, and a mirror above the sink. The toilet is mounted on the wall and there is a white towel rack next to it. The sink is made of white porcelain and there are two towels hanging from the rack. The shower area is visible in the background, and there is a grey tiled floor. The walls have large white tiles and there is a small window with a white curtain. The overall impression is of a clean, modern, and functional bathroom.

.....
.....

Figure 3: Prompt Template for LMM to generate key features of the image for House dataset

#INPUT Prompt:
Consider the House Linearized Graph:
This house location Manningham. This house property type Apartment. This house address 642/654 Doncaster Rd, Manningham VIC 3108, Australia. This house sold price 526000. This house leased date Sat 23-May-15. This house agency name Vicprop - MELBOURNE. This house agency address Level 10, 179 Queen Street Melbourne Vic 3000 . This house number of bedrooms 2.

Also consider the key features of the images of the House:

Key features of image:1:
The key features of this image are: a large glass wall, a wooden deck with chairs, a swimming pool, and a view of the surrounding city.

Key features of image:2:
This image shows a rooftop patio with a large table and chairs, a small table with two chairs, and a large planter with plants. The view from the rooftop is of a cityscape with skyscrapers and trees in the foreground. The sky is clear and sunny, and there is a blue and white striped awning over the patio. The patio is surrounded by a railing made of metal and glass, and there are several large windows on the side of the building. The building is made of concrete and glass, with a modern design.

.....
.....

Now Generate a Real-estate advertising summary with an appropriate title that summarizes all the features of the House Linearized Graph and the key features of the images of the House.

#OUTPUT:
Title: Luxurious 2-bedroom Apartment in the Heart of Melbourne
This stunning 2-bedroom apartment is located in the heart of Melbourne, just minutes away from the city's top attractions. With a prime location, this apartment offers breathtaking views of the surrounding city and is surrounded by tall trees and other buildings. The apartment has ducted heating, air-conditioning, a dishwasher, a dining room, a large glass wall.....

Figure 4: Prompt Template for LMM to generate text for House dataset

Prompt
You are given a Sentence. You have to list the features one by one from the Sentence.

Examples:
#Few In Context learning Examples

Sample INPUT:
Sentence: This fantastic home boasts 3 spacious bedrooms, modern fully tiled bathroom with spa, and separate laundry with additional toilet, fully landscaped rear garden with water tank leading to the remote control double lock up garage.

Sample OUTPUT:
List of Features:
 1. 3 bedrooms
 2. Spacious bedrooms
 3. Fully tiled bathroom
 4. Bathroom with spa
 5. Separate laundry
 6. Additional toilet
 7. Fully landscaped rear garden
 8. Garden with water tank
 9. Remote control double lock up garage

Figure 5: Prompt Template for LLM to extract list of features from a sentence

Prompt
*You are given a List of Features and a set of uploaded images.
 You have to verify whether a particular feature is "Visible" or "Not Visible" in the images.*

Examples:
#Few In Context learning Examples

Sample INPUT:
 Is the feature-'Lace veranda' is visible in the uploaded images? Answer 'Yes' or 'No'. If 'Yes', then mention the image number.
 Is the feature-'Corner lot' is visible in the uploaded images? Answer 'Yes' or 'No'. If 'Yes', then mention the image number.
 Is the feature-'Hardwood floors' is visible in the uploaded images? Answer 'Yes' or 'No'. If 'Yes', then mention the image number.

Sample OUTPUT:
 Lace veranda: Yes, Visible in Image 1
 Corner lot: No.
 Hardwood floors: Yes, Visible in Image 1, Image 2, Image 4

Figure 6: Prompt Template for GPT-4V to list visible and not visible image features

Prompt
*You are given a Features List and a House Linearized Graph.
 You have to only list those features from the 'Feature List' that are not supported by the 'House Linearized Graph'.*

Examples:
#Few In Context learning Examples

INPUT:
*#Not Visible Image Features List
 #House Linearized Graph*

OUTPUT:
#List of Hallucinated Features

Figure 7: Prompt Template for LLM to list hallucinated features

Prompt
 You are given two lists: List-1, List-2 and a House Linearized Graph. Your task is to **Classify the features of List-2 into two groups:**

Salient Features: The features from List-2 which are similar with any of the features of List-1 or supported by the 'House Linearized Graph'
Not-Salient Features: All the remaining features from List-2 that are not classified as Salient Features

INPUT:
 #List-1 :: Visible Image Features List from Ground-Truth Text (Salient Image Features List)
 #List-2 :: Visible Image Features List from the LMMs (Salient and Not-Salient Image Features List)
 #House Linearized Graph

OUTPUT:
 #List of Salient Image Features and Not-Salient Image Features

Figure 8: Prompt Template for LLM to list salient and not salient image features

Prompt Input to LLM:
 The feature of a House image: "**Double-fronted block facade home**" is an important feature for a real-estate advertising, because ..? "Complete within one sentence".

Output from Zero-Shot LLM (ChatGPT-3.5):
 The feature of a House image: "**Double-fronted block facade home**" is an important feature for real-estate advertising because it signifies architectural distinction, spaciousness, and curb appeal, drawing attention and interest from potential buyers.

Prompt Input to LLM:
 The feature of a House image: "**Black Roof**" is not an important feature for a real-estate advertising, because ..? "Complete within one sentence".

Output from Zero-Shot LLM (ChatGPT-3.5):
 The feature of a House image: "**Black Roof**" is not an important feature for real-estate advertising because it is a relatively common and neutral characteristic that does not significantly enhance the property's appeal or value.

Figure 9: Prompt Template for LLM to generate rationale for salient features and not-salient features

Consider the following Text:
 #Text:.....

Remove the following features from the Text :
 #List of Not-salient and Hallucinated Features

Update the summary by considering the following features:
 #List of Salient Features

Note: Do not make any other modification of the Summary.

Figure 10: Prompt Template for LLM to do post-hoc correction of the generated text according to the feedback of the critic model

Consider the Linearized Graph/Tabular Data and the Images.
 You are given a "List of Features".

Your task is to list only those features from the "List of Features", that are not aligned with the Linearized Graph/Tabular Data or the Images.

INPUT:
 #Linearized Graph/Tabular Data
 #Images
 #List of Features

OUTPUT:
 #List of Faithful Features

Figure 11: Prompt Template for GPT-4V to list faithful and salient features from the ground-truth text

Consider the following "List of Features".

Now generate a real-estate advertising summary with all the features from the "List of Features".

INPUT:

#List of Features

OUTPUT:

#Advertising Text

Figure 12: Prompt Template for LLM to make a paragraph with the faithful and salient features

House Knowledge Graph and Images:

The Knowledge Graph is a central node 'House' connected to various features. Features include: location (Essendon), sold price (\$1380000), num. of bedrooms (3), num. of bathrooms (1), num. of parking spaces (2), courtyard, nearest park, dining room, 9 mins walking distance to closest railway station, leased date (Sat 30-May-15), agency name (Nelson Alexander Real Estate - Essendon), address (36 Edward St, Essendon VIC 3040, Australia), nearest school (Essendon East Keilor District College Senior Campus, Co-Ed), nearest train station (Essendon Railway Station, Secondary), and agency address (1019 Mount Alexander Road Essendon Vic 3040).

The Images section contains three photographs: an exterior view of a yellow house with a white porch, an interior view of a bright living room with large windows, and a close-up of the front porch with a white railing and small garden.

Pre-processed Ground-truth text:

Stunning Edwardian Charm in Prestige Essendon Location!: This grand home features three spacious bedrooms, a central bathroom, elegant formal lounge, dining room, and adjoining kitchen. Enjoy high ceilings, lead-light windows, a cellar, hallway storerooms, and a wrap-around verandah overlooking the expansive north-facing backyard. With a garage, carport, and walking distance to local parks and Essendon Station, this elevated corner block boasts a total land area of 603sqm. Don't miss out on this rare opportunity to own a piece of Edwardian history in one of Essendon's most sought-after neighborhoods!

Our model:

Spacious 3-bedroom House in Essendon with Picket Fence and White Porch: This advertisement is for a charming classical Victorian home in the urban area of Essendon, Melbourne. The house has a picket fence surrounding the front yard, a small porch with a white railing and two small windows, a small tree, and a small garden. The key features of the image are the yellow exterior and white trim, small garden, and white wooden porch with large windows and wooden frames. Located on a generous block, this charming home boasts original period features that add to its character. The spacious living room is perfect for families looking for a cozy, well-maintained, and well-located house in Essendon. Additional information provided includes the address (36 Edward St, Essendon VIC 3040, Australia), sold price (\$1380000), leased date (Sat 30-May-15), agency name (Nelson Alexander Real Estate - Essendon), agency address (1019 Mount Alexander Road Essendon Vic 3040), number of bedrooms (3), number of bathrooms (1), number of parking spaces (2), whether it has a courtyard, a dining room, a nearby park, the number of carport spaces (1), the number of garage spaces (1), proximity to the nearest train station (9 mins walking distance), and the nearest school (Essendon East Keilor District College Senior Campus). The image key features are also highlighted, making this property an ideal choice for those seeking a Victorian home with period charm in Essendon.

Woodpecker model:

Charming 3-Bedroom House in Essendon with Picket Fence and White Porch: This advertisement is for a spacious 3-bedroom house located in the desirable area of Essendon, Melbourne. The house features a charming yellow exterior with white trim, giving it a warm and inviting feel. The front yard is enclosed by a classic picket fence, adding to the overall charm of the property. Upon entering the house, you are greeted by a small porch with a white railing, providing the perfect spot to relax and enjoy the surroundings. The porch is adorned with two large windows, allowing for plenty of natural light to fill the interior space. The house is well-maintained and in good condition, with a small garden and a tree adding to the overall appeal. The key features of the image include the yellow exterior with white trim, the small wooden front door, and the white wooden porch with large windows and wooden frames. In addition to the visual features, the House Linearized Graph provides important information about the property. The house is located at 36 Edward St, Essendon VIC 3040, Australia. It was sold for \$1380000 and was leased on Sat 30-May-15. The agency responsible for the property is Nelson Alexander Real Estate - Essendon, located at 1019 Mount Alexander Road Essendon Vic 3040. The house offers 3 bedrooms, 1 bathroom, and 2 parking spaces, including 1 carport and 1 garage. It also features a courtyard and a dining room. The nearest park is within walking distance, providing a convenient outdoor space for recreation. For families looking for a cozy and well-located house, this advertisement is perfect. The property offers a spacious interior, a picket fence, and a small garden, creating a welcoming and comfortable atmosphere. The proximity to Essendon Railway Station, which is just a 9-minute walk away, adds to the convenience of the location. The nearest school, Essendon East Keilor District College Senior Campus, is also within close proximity, making it an ideal choice for families with school-aged children. Don't miss out on the opportunity to make this charming house in Essendon your new home. Contact Nelson Alexander Real Estate - Essendon today to arrange a viewing.

LURE model:

The house is a 3-bedroom house with 1 bathroom and 2 parking spaces, located in a quiet neighborhood with easy access to parks and public transportation. It has a beautiful outdoor living area and a modern kitchen and living room with plenty of storage space. The bedrooms are spacious and bright, and the house is equipped with ducted heating and cooling. The house is located in Essendon, just minutes away from local shops, cafes, and restaurants. The area has several parks and reserves, perfect for recreation. The location is just a few minutes' walk away from Essendon Railway Station, which is a major transport hub. The house is in close proximity to several schools, including Essendon East Keilor District College Senior Campus, which is a CoEd school. Overall, this house is a great option for those looking for a charming and convenient place to live.

MiniGPT4 model:

Spacious 3-bedroom House in Essendon with Picket Fence and White Porch: This advertisement is for a charming, small yellow house in the urban area of Essendon, Melbourne. The house has a picket fence surrounding the front yard, a small porch with a white railing and two small windows, a small tree, and a small garden. The house is well-maintained and in good condition. The key features of the image are the yellow exterior and white trim, small garden, small wooden front door, and white wooden porch with large windows and wooden frames. The House Linearized Graph provides additional information about the property, including its address (36 Edward St, Essendon VIC 3040, Australia), sold price (\$1380000), leased date (Sat 30-May-15), agency name (Nelson Alexander Real Estate - Essendon), agency address (1019 Mount Alexander Road Essendon Vic 3040), number of bedrooms (3), number of bathrooms (1), number of parking spaces (2), whether it has a courtyard, a dining room, a nearby park, the number of carport spaces (1), the number of garage spaces (1), proximity to the nearest train station (9 mins walking distance), and the nearest school (Essendon East Keilor District College Senior Campus). The image key features are also highlighted. This advertisement is perfect for families looking for a cozy, well-maintained, and well-located house in Essendon with a picket fence and a small garden.

Figure 13: Example of generated text on the House dataset.

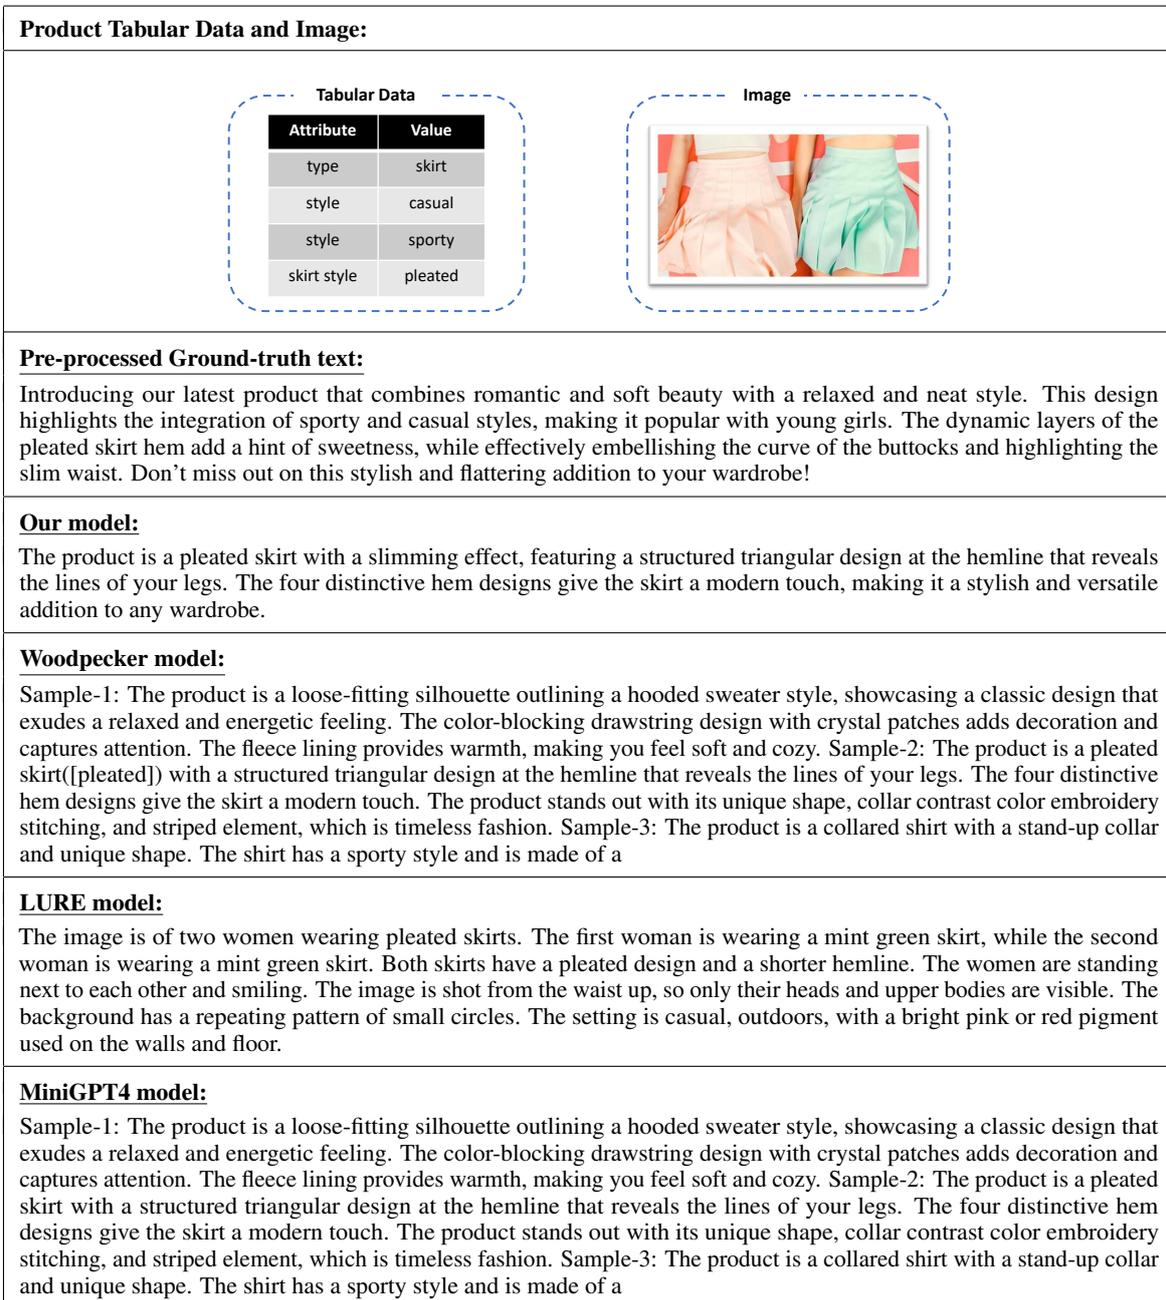


Figure 14: Example of generated text on the Product dataset.