

MULTIMODAL QUESTION ANSWERING FOR UNIFIED INFORMATION EXTRACTION

Anonymous authors

Paper under double-blind review

ABSTRACT

Multimodal information extraction (MIE) aims to extract structured information from unstructured multimedia content. Due to the diversity of tasks and settings, most current MIE models are task-specific and data-intensive, which limits their generalization to real-world scenarios with diverse task requirements and limited labeled data. To address these issues, we propose a novel multimodal question answering (MQA) framework to unify three MIE tasks by reformulating them into a unified span extraction and multi-choice QA pipeline. Extensive experiments on six datasets show that: 1) Our MQA framework consistently and significantly improves the performances of various off-the-shelf large multimodal models (LMM) on MIE tasks, compared to vanilla prompting. 2) In the zero-shot setting, MQA outperforms previous state-of-the-art baselines by a large margin. In addition, the effectiveness of our framework can successfully transfer to the few-shot setting, enhancing LMMs on a scale of 10B parameters to be competitive or outperform much larger language models such as ChatGPT and GPT-4. Our MQA framework can serve as a general principle of utilizing LMMs to better solve MIE and potentially other downstream multimodal tasks.¹

1 INTRODUCTION

Multimodal information extraction (MIE) aims to extract structured information from unstructured multimedia sources, which has drawn increasing attention as social media platforms are flooded with multimedia contents. Typically, with text and images as input, MIE can be categorized into three specific sub-tasks including multimodal named entity recognition (MNER; Sun et al. (2021)), multimodal relation extraction (MRE; Chen et al. (2022b)), and multimodal event detection (MED; Li et al. (2020a)), each with its own output format, as depicted in Figure 1a.

In addition to varying task formats, another major challenge of MIE lies in the diversity of task settings. Taking MED as an example, event triggers can be text-only, with images serving as supplementary materials (Zhang et al., 2017). Conversely, images can be used as the main component, while texts play an auxiliary role (Li et al., 2022). Finally, event triggers can also show up in both text and images simultaneously (Li et al., 2020a).

Given the diversity of tasks and settings, current works design task-specific models to address each challenge separately. Moreover, these methods require a large amount of labeled data for training, which may not always be available in many real-world scenarios. As a result, such a “*one task, one model*” paradigm is not only data-intensive and time-consuming, but also limits model’s ability of generalizing to new tasks or datasets. This becomes a major bottleneck in situations with limited resources or where fast adaptation is necessary.

To address these issues, in this work, we propose a novel multimodal question answering (MQA) framework to unify all aforementioned MIE tasks under diverse settings, as briefly illustrated in Figure 1b. Building upon off-the-shelf large multimodal models (LMMs), MQA framework can uniformly address these tasks in a zero-shot fashion, thus keeping its generalization capabilities. Specifically, we decompose all MIE tasks into cascaded atomic task forms, namely span extraction (optional) and classification tasks. Furthermore, we design multi-choice question answering

¹Code will be released upon acceptance.

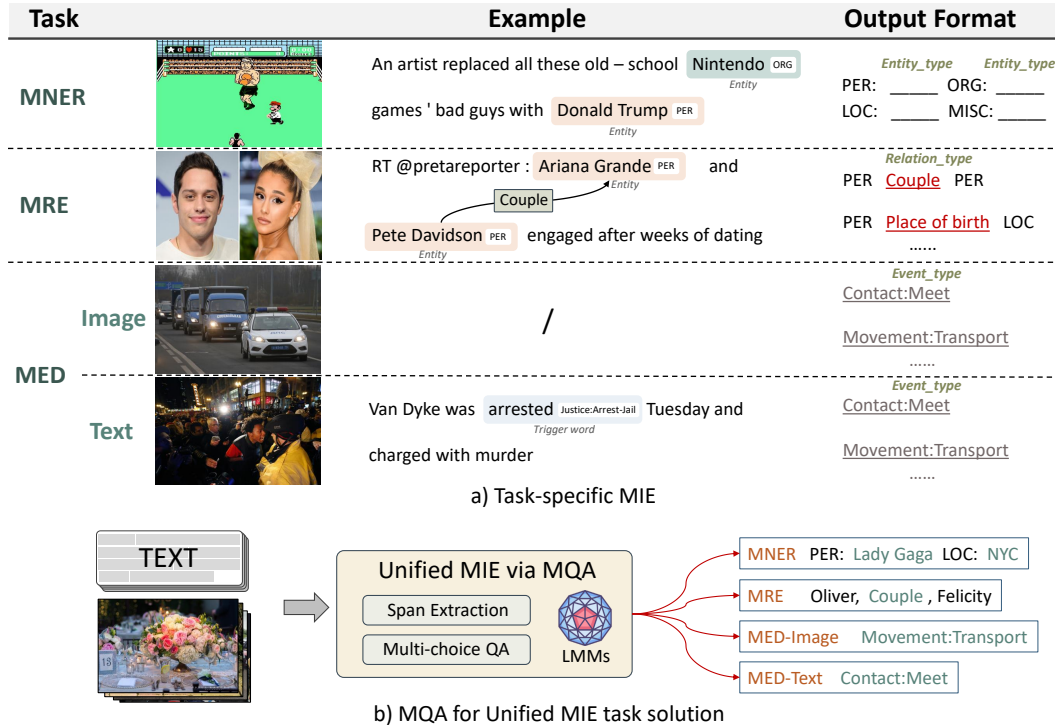


Figure 1: a) Illustration of MIE tasks, including input examples and output format. b) Our proposed MQA framework for LMMs to unify various MIE tasks.

(QA) to elicit the classification abilities of various LMMs (Li et al., 2023b; Dai et al., 2023) by aligning (Zhang et al., 2023a) classification tasks with pre-trained task forms (QA) of LMMs.

Comprehensive experiments on six LMMs across six datasets from three MIE tasks show that the MQA framework outperforms vanilla MIE prompting strategy by a considerable margin. Moreover, without any training data (zero-shot), MQA achieves significantly better results than previous state-of-the-art (SOTA) zero-shot and few-shot methods tailored for each task, and much larger language models including ChatGPT (OpenAI, 2022) and GPT-4 (OpenAI, 2023) on most of the datasets. Such empirical results demonstrate the effectiveness of MQA framework on MIE tasks and its general applicability on LMMs, and underscore the potential of task reformulation for better adapting LMMs to other downstream multimodal tasks. To summarize, our contributions are as follows:

- To address the task diversity and generalization issue, we propose a novel MQA framework to uniformly decompose three MIE tasks into cascaded atomic tasks: span extraction and classification. Also, we design multi-choice QA reformulation to elicit the classification abilities of various LMMs. To the best of our knowledge, we make the first attempt to unify MNER, MRE, and MED tasks.
- Extensive experiments of six LMMs on six datasets from three MIE tasks show that our MQA framework significantly outperforms vanilla prompting strategies. With MQA, LMMs surpasses previous low-resource SOTA and much larger language models including ChatGPT and GPT-4 on most datasets. Further in-depth analysis shows strong robustness of our MQA framework and consistent effectiveness in the few-shot setting.
- Finally, our work provides discussions and insights on recent LMMs from a MIE perspective. Also, our study indicates the potential of task reformulation, serving as a promising general principle for better adopting LMMs to other downstream multimodal tasks.

2 RELATED WORK

Multimodal Information Extraction. MIE tasks mainly include multimodal named entity recognition (MNER), multimodal relation extraction (MRE), and multimodal event detection (MED) with

images and texts as input. **MNER** (Moon et al., 2018; Lu et al., 2018; Arshad et al., 2019; Yu et al., 2020; Sun et al., 2021; Wang et al., 2022; Chen et al., 2023) aims to identify mentions and classify them into pre-defined categories from texts, using images as additional information. **MRE** (Zhang et al., 2021b;a; Chen et al., 2022b) aims to infer the relationship between head and tail entities based on the given image. **MED** (Zhang et al., 2017; Li et al., 2020a; 2022; Liu et al., 2022) is a task of identifying event triggers and their corresponding event types. In context of text, event triggers typically refer to specific words or phrases that signify the occurrence of an event. According to Li et al. (2020a), event triggers can be spans in the given texts or be the given images per se, showing the challenging nature of MED task.

Low-Resource Information Extraction. To address the challenge of data scarcity in the domain of multimodal information extraction, several methods are proposed to leverage transfer learning and data augmentation techniques. A number of works (Yu et al., 2020; Zhang et al., 2021; Lu et al., 2022a; Chen et al., 2023) train the models with one dataset from the same domain and evaluate them on another (e.g., from Twitter15 to Twitter17). However, such transfer learning still requires a sufficient alignment between the source and target domains, which is crucial for attaining optimal performance. PGIM (Li et al., 2023a) utilizes ChatGPT (OpenAI, 2022) for additional knowledge generation, thus augmenting the downstream models for MNER tasks in order to achieve better few-shot performances. Chen et al. (2022a) propose MKGFormer to bridge the modality gap between text and image inside Transformers (Vaswani et al., 2017), achieving decent performances in low-resource scenarios of MRE and MNER. Li et al. (2020a) adopt annotated uni-modal corpora to separately train textual and visual event detection models, and bridge their modality gap with an image-caption. In a contrastive learning fashion, Li et al. (2022) pre-train a vision-language model by connecting events and arguments across different modalities. These weak supervision methods achieve impressive MED results in a zero-shot setting.

Large Multimodal Models. Pioneering multimodal models (Lu et al., 2019; Tan & Bansal, 2019; Zhou et al., 2020; Chen et al., 2020; Li et al., 2020b; 2022) are typically pre-trained on vision and language data to establish connections between modalities. As evidenced by the remarkable capabilities demonstrated by large language models (Brown et al., 2020; Chowdhery et al., 2022; Touvron et al., 2023a;b; Chung et al., 2022) especially when increasing the model size and training tokens, recent multimodal works focus on equipping off-the-shelf language models with visual abilities via lightweight fine-tuning (Li et al., 2023b; Dai et al., 2023; Liu et al., 2023; Zhang et al., 2023b; Zhu et al., 2023; Ye et al., 2023). Meanwhile, instruction tuning (Lou et al., 2023; Ouyang et al., 2022) empowers these models to follow instructions, thus enabling them to generalize to unseen tasks. However, Zhang et al. (2023a) shows that instruction-tuned models fail to deliver decent IE results and highlight the potential benefits of reformulating IE as a QA task. Our work mainly builds on these prior efforts while we extend it into multimodal scenarios and unify three MIE tasks.

3 MULTIMODAL QUESTION ANSWERING FRAMEWORK

3.1 PRELIMINARY

We provide formal definitions for MNER, MRE, and MED tasks as follows.

Multimodal Named Entity Recognition. Given a sentence T and an associated image I , the task of MNER requires models to identify disjoint continuous spans within the sentence T and subsequently classify each span into one of the pre-defined entity types, such as person and location. The image I mainly serves as an additional clue to enhance the extraction of textual entities.

Multimodal Relation Extraction. MRE aims to infer the relationship between two entities based on both a sentence and its paired image. Concretely, each instance of a relation contains a sentence T and an associated image I , along with a head entity E_h and a tail entity E_t within T . Given a relation example (S, I, E_h, E_t) , models are required to identify the relation between E_h and E_t expressed in T from a set of pre-defined relation types with the auxiliary information from image I .

Multimodal Event Detection. Based on the modality (image or text) in which the event trigger is detected, MED can be further categorized into multimodal image-centric event extraction (MIED) and multimodal text-centric event extraction (MTED). MIED classifies the given image I into event

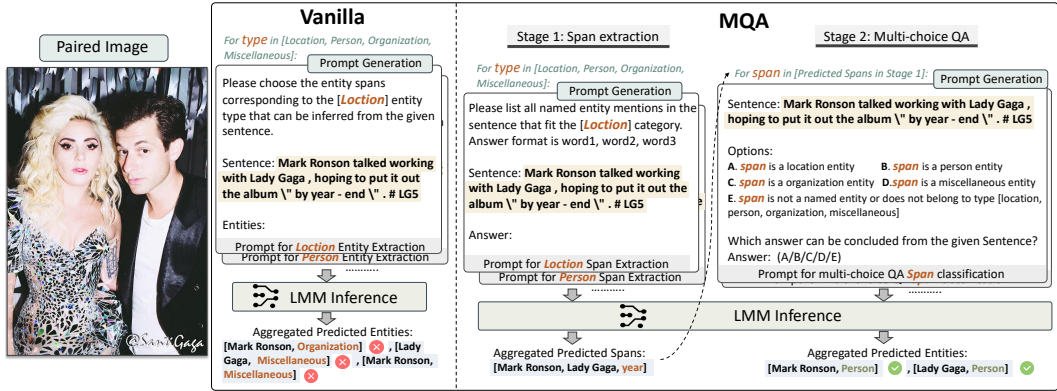


Figure 2: This figure compares the vanilla method and our proposed MQA framework for MIE tasks, using the MNER task as an illustrative example. With the vanilla prompting strategy (Gutiérrez et al., 2022), LMMs directly identify entities. This complicated and error-prone process may lead to inferior results (e.g., one same span can be classified as two different entity types). In contrast, our MQA framework decomposes a MIE task into two cascaded phases: span extraction and multi-choice QA. Spans are extracted as candidates for later multi-choice QA. Each candidate span can be classified into pre-defined categories and one additional *none-of-the-above* option (E) to discard false positives from the span extraction stage.

types. Conversely, similarly to MNER, with image I serving as an auxiliary cue, MTED extracts a set of event mentions e from a given sentence T and categorizes them into event types separately.

3.2 TASK DECOMPOSITION

Both MNER and MED can be decomposed into span extraction and span/image classification tasks, and MRE can be regarded as sentence classification conditioned on entities and images. Therefore, we unify these three MIE tasks as optional span extraction and classification tasks.

Span Extraction. Given a sentence with n tokens $T = \{t_1, \dots, t_n\}$, models are tasked with identifying a set of non-overlapping spans. (e.g., $s = \{t_i, \dots, t_j\}$ where $1 \leq i$ and $j \leq n$). For MNER and MTED, these identified spans subsequently serve as candidates for later classification to determine the exact entity/event types individually.

Multi-choice QA. Given a span/image/sentence, LMMs need to classify it into pre-defined categories corresponding to the specific task. By providing appropriate instructions and label space in the prompt, LMMs can perform these classification tasks uniformly by generating label names (Gutiérrez et al., 2022). In particular, for MNER and MED, we combine *each* extracted span s with the original sentence T to perform the classification task, with the image I as a visual clue.

3.3 MULTI-CHOICE QUESTION ANSWERING

To standardize all classification tasks, we reformulate them as multi-choice QA tasks. As illustrated in the second stage of the MQA framework in Figure 2, we provide LMMs with question answering instructions and inference examples with each possible class as an answer option, enabling LMMs to perform classification by generating an answer index. This empowers the LMMs to harness their extensive VQA abilities gained from either multimodal pretraining or visual instruction tuning to enhance classification performance, as evidenced by Zhang et al. (2023a). Furthermore, by adopting the MQA framework, we streamline the output format diversity across various tasks and alleviate the decoding complexity, resulting in a more unified and effective approach.

3.4 ANSWER OPTION CONSTRUCTION

Inspired by prior works (Zhang et al., 2023a; Ma et al., 2023), we represent each class as an answer option. Specifically, since the extracted span candidates may not belong to any pre-defined meaningful classes in MNER and MED, we design an irrelevant class called *none-of-the-above* (NOTA), as

exemplified by option E in the multi-choice QA prompt template depicted in Figure 2. Subsequently, we remove spans that are categorized as irrelevant by LMMs to mitigate the potential occurrence of false positives from the span extraction step before multi-choice QA stage.

4 EXPERIMENT SETUP

4.1 DATASETS

We evaluate our methods on MNER, MRE, and MED tasks, covering six datasets. **For MNER**, we consider (1) Twitter15 (Zhang et al., 2018) and (2) Twitter17 (Lu et al., 2018), both of which are constructed from social media content and have four entity types including person, location, organization, and miscellaneous. We report micro-averaged F1 following previous works (Sun et al., 2021; Chen et al., 2023). **For MRE**, we adopt (3) MRE-V1 (Zheng et al., 2021b) and (4) MRE-V2 (Zheng et al., 2021a). Upon Twitter15 (Zhang et al., 2018), Twitter17 (Lu et al., 2018), and crawled Twitter data, these two MNER datasets are manually annotated with the relationship between two textual entities based on the visual evidence shown in the paired image. Following Chen et al. (2022a), we use micro-averaged F1 with *None-of-the-Above* (NOTA) relation excluded for fair comparison in low-resource setting. **For MED**, we consider (5) image-centric and (6) text-centric split of M²E² (Li et al., 2020a). Following prior work (Li et al., 2020a; 2022), we report micro-averaged F1 without NOTA event type in image-centric setting and micro F1 in text-centric setting.

4.2 LARGE MULTIMODAL MODELS

To comprehensively evaluate the MQA framework on various LMMs, we consider two series of LMMs, encompassing a total of six models. Concretely, BLIP-2 (Li et al., 2023b) with various-sized FLAN-T5 (Chung et al., 2022) as the language component, and InstructBLIP (Dai et al., 2023) with different-sized FLAN-T5 and Vicuna (Chiang et al., 2023) as language backbones. Both series of models are trained based on a frozen image encoder and a frozen LLM. BLIP-2 (Li et al., 2023b) incorporates a Q-Former designed to extract visual features, serving as soft prompts for frozen large language models to facilitate text generation. Follow-up work, InstructBLIP (Dai et al., 2023), further enhances BLIP-2 by performing instruction tuning on various vision-language tasks and modifying the Q-Former module to extract instruction-aware visual features.

4.3 IMPLEMENTATION DETAILS

In our prompt engineering efforts, we explore various prompt formats and task instructions for both vanilla and MQA methods, utilizing BLIP-2 (Li et al., 2023b) on a subset of 400 examples from each task’s validation set. These optimized task instructions and prompt formats are then employed for all the models in subsequent experiments. Detailed information on the prompt formats for vanilla and MQA are available in Appendix A.1. In addition, we design experiments for robustness assessment and few-shot fine-tuning to further validate the superiority of our method over the vanilla approach.

4.3.1 VANILLA

With the vanilla approach, models generate verbalized answers directly (Gutiérrez et al., 2022). For MNER task, we instruct models to extract relevant spans corresponding to each entity type. In the case of MRE, following Zhang et al. (2023a), we adopt entity type constraints to narrow down the label space. Given the sentence and two entities, models are required to generate the relation name from the valid relation set, based on the relationship between two entities expressed in the sentence. For MIED task, models are tasked with categorizing the provided image into a pre-defined event type by generating the corresponding event name. Similar to MNER, for the MTED task, we specify the event type and instruct the model to extract words related to the event.

4.3.2 MQA FRAMEWORK

We break down and unify the MIE tasks into span extraction and span/sentence/image classification. Span extraction can be regarded as optional depending on the specifics of the given task. By

conducting multiple-choice QA tasks, LLMs are equipped to generate an answer index for classification, rather than a complete label name. In practice, we employ two-stage prompts to deconstruct the initial MIE tasks, detailed in Table 1.

MNER/MTED. These two tasks can be decomposed into a two-stage process, encompassing span extraction and multi-choice QA. As for the MNER task, the initial stage of MQA focuses on identifying entity spans. This stage is designed to allow the LLMs to generate a considerable number of entity candidates, thereby permitting a certain level of false positive entities in a strategic trade-off. In the second stage, we apply a multiple-choice QA to categorize each extracted span into a predetermined set, while any false positive entities can be classified as NOTA and consequently eliminated. Similarly, for MTED, LLMs are first prompted to extract text spans, and then a multi-choice QA is utilized to categorize the event associated with these spans more accurately.

Table 1: Details of task decomposition.

MIE Task	Stage 1	Stage 2
MNER/MTED	Span Extraction	Multi-choice QA
MRE/MIED	Multi-choice QA	-

MIED/MRE. These two MIE tasks are relatively straightforward, each encompassing a single-stage classification task, and we design appropriate prompts for both tasks respectively. In MIED, the model is prompted to identify the correct event type based on the provided image. For MRE, we specify the head and tail entities with their respective entity types and provide both the image and the entire sentence, requiring the model to select from a pre-defined set of relations.

4.4 BASELINES

We adopt several baselines to benchmark our proposed MQA framework. These baselines fall into two main groups: (1) prior SOTA baselines and (2) the most powerful larger language models to date. For prior low-resource SOTA methods, we consider PGIM (Li et al., 2023a) for MNER, which attains its performance through 50-shot few-shot fine-tuning. We adopt MKGformer (Chen et al., 2022a) for MRE, which utilizes 40-shot fine-tuning. For two MED tasks, we adapt the SOTA zero-shot model WASE (Li et al., 2020a) as our baseline.

We also consider the most advanced large-scale language models available to date, exemplified by ChatGPT (OpenAI, 2022) and GPT-4 (OpenAI, 2023). Similar to the prompting engineering on open-source LLMs, we meticulously craft and evaluate prompts for these two language models on the validation set of all datasets except the MIED task, which can not be solved by text-only models. Please refer to Table 18 in Appendix for detailed prompts.

4.5 FEW-SHOT FINE-TUNING

To assess the effectiveness of MQA framework in the few-shot setting, following Li et al. (2023a), we conduct a series of 50-shot few-shot fine-tuning experiments across all six datasets.

Few-shot Data Sampling Strategy. To ensure the distribution consistency between the subset and full set, we initially calculate the proportion of each category within the training set. Subsequently, we determine each category’s sample numbers among the 50 samples. While ensuring at least one sample for each category, we conduct random sampling from the dataset following the ascending order of sample numbers per category. As for MNER and MRE tasks, we follow the original data split. In the case of the MIED and MTED tasks, where specific training splits are not provided, we employ the aforementioned method, sampling 50 training samples from the entire dataset as the training set, 200 samples as the validation set, with the remaining serving as the test set.

Training Hyperparameters. We perform few-shot fine-tuning experiments based on BLIP-2 with Flan-T5 XXL due to its exceptional performance in the zero-shot setting. The fine-tuning process is carried out on a single NVIDIA A100 GPU with 80G of memory. Specifically, for the MNER, MRE, and MTED datasets, we freeze only the Flan-T5 language component while fine-tuning other parts. For the MIED dataset, we fine-tune Qformer with both the LLM and vision encoder frozen. We employ the Adam optimizer with a learning rate of $2e-5$ and execute the training over 10 epochs with a batch size of 8. We select the model checkpoint with the highest micro-F1 score on the validation set for subsequent evaluation on the test set.

Table 2: Experimental results on six MIE datasets (%). The prior SOTA results for each task are achieved by different models, rather than a single unified model. In this context, we use †, ‡, and * to denote these models, where † represents the PGIM (Li et al., 2023a), ‡ symbolizes the MKGformer (Chen et al., 2022a), and * denotes the WASE (Li et al., 2020a). Moreover, MQA** represents the results of multi-choice QA with ground truth spans from the span extraction stage, serving as a performance upper bound. We highlight the best results in **bold** and mark the F1 score improvements of our MQA framework over vanilla method in **green**.

Methods		MNER		MRE		MED		Average
		Twitter15	Twitter17	MNRE-V1	MNRE-V2	Image	Text	
<i>Baselines</i>								
Prior SOTA (Zero-shot)		-	-	-	-	49.9*	50.6*	-
Prior SOTA (Few-shot)		52.2 †	50.7 †	-	40.9 ‡	-	-	-
ChatGPT		39.7	45.2	41.2	48.3	-	13.9	-
GPT-4		42.3	54.7	61.7	61.3	-	32.5	-
<i>BLIP-2 series (Zero-shot)</i>								
Flan-T5 XL	Vanilla	22.2	21.9	40.8	41.6	15.3	13.9	26.0
	MQA	43.8 (+21.6)	56.8 (+34.9)	51.7 (+10.9)	56.5 (+14.9)	52.5 (+37.2)	27.1 (+13.2)	48.1 (+22.1)
	MQA**	83.4	84.7	-	-	-	73.2	-
Flan-T5 XXL	Vanilla	26.6	29.1	42.9	46.4	42.9	17.1	34.2
	MQA	50.6 (+24.0)	62.6 (+33.5)	53.5 (+10.6)	61.6 (+15.2)	55.9 (+13.0)	53.3 (+36.2)	56.3 (+22.1)
	MQA**	83.9	86.0	-	-	-	86.6	-
<i>InstructBLIP series (Zero-shot)</i>								
Flan-T5 XL	Vanilla	23.0	25.7	44.9	51.7	38.5	10.0	32.3
	MQA	39.7 (+16.7)	49.9 (+24.2)	53.3 (+8.4)	59.0 (+7.3)	52.9 (+14.4)	22.7 (+12.7)	46.3 (+14.0)
	MQA**	81.2	80.9	-	-	-	70.3	-
Flan-T5 XXL	Vanilla	23.2	28.6	33.7	40.2	33.1	15.6	29.1
	MQA	48.1 (+24.9)	58.8 (+30.2)	55.2 (+21.5)	61.5 (+21.3)	52.6 (+19.5)	39.8 (+24.2)	52.7 (+23.6)
	MQA**	83.6	86.3	-	-	-	83.0	-
Vicuna 7B	Vanilla	9.8	17.8	18.5	24.9	13.2	2.7	14.5
	MQA	9.8 (+0.0)	14.0 (-3.8)	27.6 (+9.1)	26.3 (+1.4)	20.9 (+7.7)	0.5 (-2.2)	16.5 (+2.0)
	MQA**	33.6	35.3	-	-	-	26.7	-
Vicuna 13B	Vanilla	14.0	15.0	17.5	26.1	8.1	1.4	13.7
	MQA	15.2 (+1.2)	19.9 (+4.9)	31.8 (+14.3)	38.5 (+12.4)	21.7 (+13.6)	0.1 (-1.3)	21.2 (+7.5)
	MQA**	64.6	61.3	-	-	-	34.3	-

5 EXPERIMENTS

5.1 MAIN RESULTS

In zero-shot setting, our results derived from six MIE datasets are shown in Table 2, from which we make the following observations:

Firstly, across six datasets with six different models, MQA framework demonstrates significant performance improvements compared to vanilla method in 32 out of 36 evaluations, where most of them show an increase of over 10% in the F1 score. Particularly, the BLIP-2 series showcases substantial improvements, boasting an average increase of over 20% in F1 scores. Notably, the model built on Flan-5 XL demonstrates a remarkable increase of 37.2% in the F1 score on the MIED dataset, signifying an extraordinary 243% relative performance improvement compared to the vanilla method. These results corroborate the wide-ranging effectiveness of our proposed MQA framework.

Secondly, by employing MQA approach, LMMs notably outperform the most advanced large-scale language models (ChatGPT and GPT-4) with 175B parameters and beyond, across most datasets. This superiority becomes particularly evident when using the 10B parameter-level BLIP-2 with FlanT5 XXL on the MNER and MED tasks, where we demonstrate a significant performance advantage. Furthermore, our model consistently surpasses previous zero-shot SOTA models across various datasets and achieves comparable or better results than the 50-shot fine-tuning SOTA results on MNER. For comparison, our MQA’s 50-shot fine-tuning results can be found in Section 5.2.

Thirdly, with either vanilla or MQA methods, the performance of InstructBLIP generally falls short compared to similarly scaled BLIP-2 models, despite that it is instruction-tuned with more data. We hypothesize that this disparity is due to the limited amount of task types for instruction tuning. Consequently, there are limited benefits when compared to its BLIP-2 upon receiving new instructions.

Table 3: Results of Flan-T5 XXL fine-tuned on 50 samples, with vanilla and MQA methods across six MIE datasets (%). We highlight F1 score improvements of MQA over vanilla in green.

Methods		MNER		MRE		MED	
		Twitter15	Twitter17	MNRE-V1	MNRE-V2	Image	Text
0-shot	Vanilla	26.6	29.1	42.9	46.4	42.9	17.1
	MQA	50.6 (+24.0)	62.6 (+33.5)	53.5 (+10.6)	61.6 (+15.2)	54.2 (+11.3)	53.6 (+36.5)
50-shot	Vanilla	28.9	31.5	45.1	48.0	45.2	16.8
	MQA	55.5 (+26.6)	70.6 (+39.1)	54.3 (+9.2)	61.7 (+13.7)	55.4 (+10.2)	54.6 (+37.8)

Fourthly, models based on Vicuna significantly underperform those based on Flan-T5, both with vanilla and MQA methods. This can be attributed to the fact that Vicuna is fine-tuned with open-ended conversations derived from user interactions with ChatGPT, instead of with traditional NLP tasks like Flan-T5, thereby restricting Vicuna’s ability on MIE tasks.² That being said, our MQA still brings decent performance gains to Vicuna-based LMMs on average.

Lastly, we conduct an additional experiment by providing golden spans in MNER and MTED tasks (MQA** in Table 2). When golden spans are presented, the MQA results for all models demonstrate further enhancements. For instance, Flan-T5 XXL achieves F1 scores of 83.9%, 86.0%, and 86.6% on the Twitter15, Twitter17, and MTED datasets, respectively. The results, especially in Twitter15, are comparable with those of SOTA model (Chen & Feng, 2023) fine-tuned with full dataset, showing performance bottleneck of span extraction and the effectiveness of multi-choice QA per se.

5.2 FEW-SHOT FINE-TUNING RESULT

The 50-shot fine-tuning results are shown in Table 3. From these results, we can draw the following conclusions: (1) The performance of our MQA remains significantly superior compared to the vanilla method across the board. (2) The MQA method showcases improvements across all datasets and tasks. This is particularly noticeable in the MNER task, where it demonstrates a 4.9% and 8% increase in F1 score on the Twitter15 and Twitter17 datasets, respectively. Conversely, the model with vanilla method undergoes a slight performance degradation on the MTED task. (3) Following the 50-shot few-shot fine-tuning, our MQA method outperforms PGIM (Li et al., 2023a) on the Twitter15 dataset, which also undergoes 50-shot fine-tuning on Twitter, allowing our MQA method to achieve SOTA results across all six datasets.

5.3 EVALUATION OF MODEL ROBUSTNESS

Prompt robustness refers to the model’s ability to accurately understand and respond to different prompts expressing the same task intent. Prior works (Lou et al., 2023; Sun et al., 2023) highlight that different instructions may lead to considerable performance fluctuations in language models. Therefore, to evaluate the instruction-following robustness of vanilla and MQA methods, we manually rewrite four different instructions for each of them. In addition, Lu et al. (2022b) find the order of the input examples also has a significant impact on model performance. To verify the robustness of the input order of multi-choice QA, we randomly shuffle the answer options. All experiments are conducted using BLIP-2 model with Flan-T5 XXL language component.

Evaluation of Robustness to Instruction Variants. To evaluate the robustness of instruction-following, we select one representative dataset from each of the four tasks: MNER, MRE, MIED, and MTED. For each dataset, we manually write four diverse yet semantically equivalent instructions by rephrasing the initial instruction. Please refer to Appendix A.2 for details of these instructions for vanilla and MQA methods.

From the results in Table 4, we have the following observations: (1) Across various instructions on different tasks, MQA consistently outperforms vanilla prompting significantly. (2) MQA exhibits better overall robustness. To be more specific, in the MNER and MRE datasets, the model obtains a considerably low sample standard deviation of 0.1% and 0.4%, respectively, whereas vanilla, even under lower performance, shows a higher sample standard deviations of 6.1% and 3.8%. However,

²Vicuna is fine-tuned with data from <https://sharegpt.com>.

Table 4: Comparison of the robustness to instruction variants between vanilla and MQA methods (%). The final row presents the mean and sample standard deviation of the model performance under four instructions.

Instruction Variant		MNER-17			MRE-v2			MIED			MTED		
		P	R	F1	P	R	F1	P	R	F1	P	R	F1
Instruction1	Vanilla	20.0	53.7	29.1	33.8	74.1	46.4	37.6	49.9	42.9	9.7	69.4	17.1
	MQA	61.2	64.0	62.6	50.2	79.7	61.6	60.4	52.0	55.9	49.8	57.2	53.3
Instruction2	Vanilla	15.9	52.2	24.4	29.7	62.7	40.3	30.3	58.3	39.8	9.6	73.6	17.0
	MQA	61.2	64.0	62.6	48.2	81.6	60.6	29.9	71.1	42.1	51.4	56.1	53.7
Instruction3	Vanilla	11.5	50.7	18.8	31.0	66.9	42.3	28.4	66.9	39.9	9.3	71.1	16.5
	MQA	61.2	64.0	62.6	49.2	79.2	60.7	32.8	78.0	46.2	53.9	51.2	52.5
Instruction4	Vanilla	9.1	47.9	15.2	29.4	61.4	37.5	28.3	67.5	39.9	9.4	72.0	16.5
	MQA	60.9	63.8	62.3	48.8	80.3	61.0	66.1	31.2	42.4	54.0	50.4	52.2
$u \pm \sigma$	Vanilla	21.9 ± 6.1			41.6 ± 3.8			40.6 ± 1.5			16.8 ± 0.3		
	MQA	62.5 ± 0.1			61.0 ± 0.4			46.6 ± 6.4			52.9 ± 0.7		

Table 5: Robustness of our MQA method to input option order variants (%). The final row displays the mean and sample standard deviation of the performances across the four variants.

Option Order	P	MNER-17			MRE-v2			MIED			MTED		
		P	R	F1	P	R	F1	P	R	F1	P	R	F1
Order1	61.2	64.0	62.6	50.2	79.7	61.6	60.4	52.0	55.9	49.8	57.2	53.3	
Order2	63.4	62.8	63.1	49.6	80.2	61.3	55.9	53.5	54.7	47.7	52.1	49.8	
Order3	60.4	63.7	62.0	49.8	79.8	61.3	57.9	53.0	55.3	49.2	55.1	52.0	
Order4	60.4	64.0	62.1	50.2	80.5	61.9	57.3	55.4	56.3	48.8	52.5	50.6	
$u \pm \sigma$	62.5 ± 0.5			61.5 ± 0.3			55.5 ± 0.7			51.4 ± 1.5			

an exception is observed in the MIED dataset wherein our model exhibited a relatively high variation. We hypothesize this might be due to the fact that MIED task solely relies on images and BLIP-2 series models have not been exposed to image classification prompts during pre-training.

Evaluation of Robustness to Input Order. To evaluate our MQA framework’s robustness to varying input orders, we design four distinct input formats by permuting the arrangement of multi-choice options in a random manner. For comprehensiveness, we also evaluate the robustness on four datasets from four MIE tasks with results shown in Table 5. We find that our MQA exhibits a high degree of robustness to varying input orders. The performance shows little fluctuation, demonstrating the model’s effectiveness in handling changes in the order of multi-choice options. Notably, our approach achieves less than 1% sample standard deviation on the MNER, MRE, and MIED datasets. These results confirm the MQA framework’s high robustness to variations in input order.

The robustness to instruction-following and input order of MQA framework indicates less prompt engineering effort and high practicalness in the real-world scenarios.

6 CONCLUSION

In this study, we present a novel MQA framework, aiming to unify three diverse MIE tasks across various settings, effectively addressing concerns related to task generalization and data scarcity. Upon six off-the-shelf LMMs, MQA consistently achieves significant performance improvements compared to vanilla prompting strategies. Remarkably, without necessitating any fine-tuning, the MQA framework surpasses prior task-specific low-resource SOTA methods across six datasets and achieves comparable or even better results than much larger language models such as ChatGPT and GPT-4 on most of the datasets. In addition, MQA demonstrates robustness to various instructions and input orders, and exhibits effectiveness in the few-shot setting. All these desirable properties indicate that the proposed MQA framework is a powerful and unified tool for MIE tasks. More importantly, our experiments suggest that, aligning diverse downstream multimodal task formats (e.g., MIE) with the task forms that LMMs have been pre-trained on (e.g., VQA), holds a promising direction for enhancing the utility of LMMs in downstream tasks.

REFERENCES

- O Arshad, I Gallo, S Nawaz, and A Calefati. Aiding intra-text representations with visual context for multimodal named entity recognition. In *Proceedings of ICDAR*, 2019.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *In Proceedings of NeurIPS*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfc4967418bfb8ac142f64a-Abstract.html>.
- Feng Chen and Yujian Feng. Chain-of-thought prompt distillation for multimodal named entity and multimodal relation extraction, 2023.
- Feng Chen, Jiajia Liu, Kaixiang Ji, Wang Ren, Jian Wang, and Jingdong Wang. Learning implicit entity-object relations by bidirectional generative alignment for multimodal ner. In *Proceedings of AMC Multimedia*, 2023.
- Xiang Chen, Ningyu Zhang, Lei Li, Shumin Deng, Chuanqi Tan, Changliang Xu, Fei Huang, Luo Si, and Huajun Chen. Hybrid transformer with multi-level fusion for multimodal knowledge graph completion. In *Proceedings of SIGIR*, 2022a. URL <https://doi.org/10.1145/3477495.3531992>.
- Xiang Chen, Ningyu Zhang, Lei Li, Yunzhi Yao, Shumin Deng, Chuanqi Tan, Fei Huang, Luo Si, and Huajun Chen. Good visual guidance make a better extractor: Hierarchical visual prefix for multimodal entity and relation extraction. In *Findings of NAACL*, 2022b.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *Proceedings of ECCV*, 2020.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, March 2023. URL <https://lmsys.org/blog/2023-03-30-vicuna/>.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. Palm: Scaling language modeling with pathways. *arxiv preprint arXiv:2204.02311*, 2022.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Y. Zhao, Yanping Huang, Andrew M. Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. Scaling instruction-finetuned language models. *CoRR*, abs/2210.11416, 2022. doi: 10.48550/arXiv.2210.11416. URL <https://doi.org/10.48550/arXiv.2210.11416>.

- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. *arXiv preprint arXiv:2305.06500*, 2023.
- Bernal Jiménez Gutiérrez, Nikolas McNeal, Clay Washington, You Chen, Lang Li, Huan Sun, and Yu Su. Thinking about gpt-3 in-context learning for biomedical ie? think again. In *Findings of EMNLP 2022*, 2022.
- Jinyuan Li, Han Li, Zhuo Pan, and Gang Pan. Prompt chatgpt in mner: Improved multimodal named entity recognition method based on auxiliary refining knowledge from chatgpt. *arXiv preprint 2305.12212*, 2023a.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023b.
- Manling Li, Alireza Zareian, Qi Zeng, Spencer Whitehead, Di Lu, Heng Ji, and Shih-Fu Chang. Cross-media structured common space for multimedia event extraction. In *Proceedings of ACL*, 2020a.
- Manling Li, Ruochen Xu, Shuohang Wang, Luwei Zhou, Xudong Lin, Chenguang Zhu, Michael Zeng, Heng Ji, and Shih-Fu Chang. Clip-event: Connecting text and images with event structures. In *Proceedings of CVPR*, 2022.
- Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. Oscar: Object-semantics aligned pre-training for vision-language tasks. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (eds.), *Proceedings of ECCV*, 2020b. URL https://doi.org/10.1007/978-3-030-58577-8_8.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arxiv preprint arXiv:2304.08485*, 2023.
- Jian Liu, Yufeng Chen, and Jinan Xu. Multimedia event extraction from news with a unified contrastive learning framework. In *Proceedings of ACM Multimedia*, 2022.
- Renze Lou, Kai Zhang, and Wenpeng Yin. Is prompt all you need? no. a comprehensive and broader view of instruction learning. *arXiv preprint arXiv:2303.10475*, 2023.
- Di Lu, Leonardo Neves, Vitor Carvalho, Ning Zhang, and Heng Ji. Visual attention model for name tagging in multimodal social media. In *Proceedings of ACL*, 2018.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Proceedings of NeurIPS*, 2019.
- Junyu Lu, Dixiang Zhang, Jiaying Zhang, and Pingjian Zhang. Flat multi-modal interaction transformer for named entity recognition. In *Proceedings of COLING*, 2022a.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. In *Proceedings ACL*, May 2022b. URL <https://aclanthology.org/2022.acl-long.556>.
- Yubo Ma, Yixin Cao, YongChing Hong, and Aixin Sun. Large language model is not a good few-shot information extractor, but a good reranker for hard samples!, 2023.
- Seungwhan Moon, Leonardo Neves, and Vitor Carvalho. Multimodal named entity recognition for short social media posts. In *Proceedings of NAACL*, 2018.
- OpenAI. Introducing chatgpt. <https://openai.com/blog/chatgpt>, 2022.
- OpenAI. Gpt-4 technical report, 2023.

- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. *CoRR*, abs/2203.02155, 2022. doi: 10.48550/arXiv.2203.02155. URL <https://doi.org/10.48550/arXiv.2203.02155>.
- Jiuding Sun, Chantal Shaib, and Byron C. Wallace. Evaluating the zero-shot robustness of instruction-tuned language models, 2023.
- Lin Sun, Jiquan Wang, Kai Zhang, Yindu Su, and Fangsheng Weng. Rpbert: a text-image relation propagation-based bert model for multimodal ner. In *Proceedings of AAAI*, 2021.
- Hao Tan and Mohit Bansal. LXMERT: Learning cross-modality encoder representations from transformers. In *Proceedings of EMNLP-IJCNLP*, 2019. URL <https://aclanthology.org/D19-1514>.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023a.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023b.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of NeurIPS*, 2017.
- Xinyu Wang, Min Gui, Yong Jiang, Zixia Jia, Nguyen Bach, Tao Wang, Zhongqiang Huang, and Kewei Tu. Ita: Image-text alignments for multi-modal named entity recognition. In *Proceedings of NAACL*, 2022.
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023.
- Jianfei Yu, Jing Jiang, Li Yang, and Rui Xia. Improving multimodal named entity recognition via entity span detection with unified multimodal transformer. In *Proceedings of ACL*, 2020. URL <https://aclanthology.org/2020.acl-main.306>.
- Dong Zhang, Suzhong Wei, Shoushan Li, Hanqian Wu, Qiaoming Zhu, and Guodong Zhou. Multimodal graph fusion for named entity recognition with targeted visual guidance. In *Proceedings of AAAI*, 2021.
- Kai Zhang, Bernal Jiménez Gutiérrez, and Yu Su. Aligning instruction tasks unlocks large language models as zero-shot relation extractors. In *Findings of ACL*, 2023a.
- Qi Zhang, Jinlan Fu, Xiaoyu Liu, and Xuanjing Huang. Adaptive co-attention network for named entity recognition in tweets. In *Proceedings of AAAI*, 2018. URL <https://api.semanticscholar.org/CorpusID:19214156>.

Renrui Zhang, Jiaming Han, Chris Liu, Peng Gao, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, and Yu Qiao. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199*, 2023b.

Tongtao Zhang, Spencer Whitehead, Hanwang Zhang, Hongzhi Li, Joseph G. Ellis, Lifu Huang, Wei Liu, Heng Ji, and Shih-Fu Chang. Improving event extraction via multimodal integration. In *Proceedings of ACM Multimedia*, 2017. URL <https://doi.org/10.1145/3123266.3123294>.

Changmeng Zheng, Junhao Feng, Ze Fu, Yi Cai, Qing Li, and Tao Wang. Multimodal relation extraction with efficient graph alignment. In *Proceedings of ACM Multimedia*, 2021a.

Changmeng Zheng, Zhiwei Wu, Junhao Feng, Ze Fu, and Yi Cai. Mnre: A challenge multimodal dataset for neural relation extraction with visual evidence in social media posts. In *Proceedings of ICME*, 2021b.

Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason J. Corso, and Jianfeng Gao. Unified vision-language pre-training for image captioning and VQA. In *Proceedings of AAAI*, 2020. URL <https://ojs.aaai.org/index.php/AAAI/article/view/7005>.

Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023.

A APPENDIX

A.1 PROMPTS UTILIZED FOR VANILLA AND MQA APPROACHES

The specific prompts employed in the Vanilla and MQA approaches for each task are presented in Table. 6 (MNER), Table. 7 (MIED), Table. 8 (MRE), Table. 9 (MTED).

A.2 INSTRUCTIONS VARIANTS WITHIN THE PROMPT USED FOR THE ROBUSTNESS EVALUATION OF VANILLA AND MQA

In the robustness evaluation experiment, various instruction variants within the prompts designed for vanilla can be observed in Table 10 (MNER), Table 11 (MIED), Table 12 (MRE), and Table 13 (MTED). Meanwhile, instruction variants for the MQA during the multi-choice QA stage are presented in Table 14 (MNER), Table 15 (MIED), Table 16 (MRE), and Table 17 (MTED).

A.3 PROMPTS FOR CHATGPT/GPT4

We leverage the OpenAI API to execute the calls, utilizing the model codenamed *gpt3.5-turbo* for ChatGPT, and *gpt4* for GPT-4. The prompts used to assess the performance of ChatGPT/GPT-4 on text-based MIE tasks are displayed in Table 18.

Table 6: Prompt formats used in Vanilla and MQA approaches for MNER task.

Formulations	Prompts
Vanilla	<p>Please choose the entity spans corresponding to the [Entity category E_c] entity type that can be inferred from the given sentence.</p> <p>Entities:</p>
MQA - Span Extraction	<p>Please list all named entity mentions in the sentence that fits the [Entity category E_c] category. Answer format is entity1, entity2, entity3.</p> <p>Sentence: [Sentence S]</p> <p>Answer:</p>
MQA - Multi-choice QA	<p>Choose the right answer about the entity that can be inferred from the given sentence.</p> <p>Sentence: [Sentence S]</p> <p>Options:</p> <p>A. [Predicted Entity Span E_{ps}] is a location entity B. [Predicted Entity Span E_{ps}] is a person entity C. [Predicted Entity Span E_{ps}] is an organization entity D. [Predicted Entity Span E_{ps}] is not a named entity or does not belong to type [location, person, organization, miscellaneous]</p> <p>Which answer can be inferred from the given Sentence? Answer:</p>

Table 7: Prompt formats used in Vanilla and MQA approaches for MIED task.

Formulations	Prompts
Vanilla	<p>Given an image and a sentence, classify which type of activity can be inferred.</p> <p>Sentence: This is an image attached to a news article.</p> <p>All possible types are listed below:</p> <ul style="list-style-type: none"> - Movement:Transport - Contact:Phone-Write - Conflict:Attack - Contact:Meet - Justice:Arrest-Jail - Conflict:Demonstrate - Life:Die - Transaction:Transfer-Money - None <p>Type:</p>
MQA - Multi-choice QA	<p>Given an image, classify which type of activity can be inferred.</p> <p>Sentence: This is an image attached to a news article.</p> <p>All possible types are listed below:</p> <p>Options:</p> <p>A. The image with the sentence describes the Transportation of Movement event B. The image with the sentence describes the Phone or Write of Contact event C. The image with the sentence describes the Attack with no death of Conflict event D. The image with the sentence describes the Meet of Contact event E. The image with the sentence describes the Arrest of Criminal event F. The image with the sentence describes the Demonstration of Conflict event G. The image with the sentence describes the Death of Life event H. The image with the sentence describes the Money Transfer of Transaction event I. The image describes no event</p> <p>Which answer can be inferred from the image? Answer:</p>

Table 8: Prompt formats used in Vanilla and MQA approaches for MRE task. We adopt the relation template in MQA - Multi-choice QA following Zhang et al. (2023a)

Formulations	Prompts
Vanilla	<p data-bbox="548 884 1313 951">Given the image and the text, select the relation between Entity 1 and Entity 2 that can be inferred from the given sentence. The image may provide fine-grained information about the entities. All possible relations are listed below:</p> <p data-bbox="548 974 732 1041">-[Possible Relation 1] -[Possible Relation 2] -None</p> <p data-bbox="548 1064 743 1108">Sentence: [Sentence S] Relation:</p>
MQA - Multi-choice QA	<p data-bbox="548 1119 1313 1163">In light of the provided sentence, determine which option is the most feasible inference. The image may present detailed information about the entities.</p> <p data-bbox="548 1186 743 1209">Sentence: [Sentence S]</p> <p data-bbox="548 1232 881 1323">Options: A. [Entities in Relation 1 Template] B. [Entities in Relation 2 Template] C. [Entities in NoTA Relation Template]</p> <p data-bbox="548 1346 919 1390">Which option is the most possible inference? Option:</p>

Table 9: Prompt formats used in Vanilla and MQA approaches for MTED task.

Formulations	Prompts
Vanilla	<p>Given an image and a sentence, determine which word can infer the [Event category E_c] activity.</p> <p>Sentence: [Sentence S]</p> <p>Words:</p>
MQA - Span Extraction (Pre-process)	<p>Determine which option can be inferred from the given sentence.</p> <p>Sentence: [Sentence S]</p> <p>Which option can be inferred from the given Sentence?</p> <p>Options:</p> <ul style="list-style-type: none"> A. Activities involving the movement or transportation of people or goods from one place to another B. Interactions between individuals through phone calls or written communication C. Aggressive actions or assaults by one party against another D. Instances where individuals physically meet or come into contact with each other E. Incidents involving the arrest and subsequent detention in jail or custody of individuals F. Public displays of disagreement or protest to express opinions or demands G. The life of a person ends H. The exchange of money or financial resources between parties
MQA - Span Extraction	<p>Please choose the most possible trigger word from the verbs and nouns in the sentence that reflect the [Predicted Event category E_{pe}] event. Note that trigger words can only be noun or verb. Answer format is word!</p> <p>Sentence: [Sentence S]</p> <p>Answer:</p>
MQA - Multi-choice QA	<p>Determine which option can be inferred from the given sentence and image.</p> <p>Sentence: [Sentence S]</p> <p>Answer candidates:</p> <ul style="list-style-type: none"> A. The word [Predicted Trigger Word T_p] is a common word and does not reflect any of the other event B. The word [Predicted Trigger Word T_p] is the key of the Transport action, which is a subtype of Movement event C. The word [Predicted Trigger Word T_p] is the key of the PhoneWrite action, which is a subtype of Contact event D. The word [Predicted Trigger Word T_p] is the key of the conflict but no death action, which is a subtype of Conflict event E. The word [Predicted Trigger Word T_p] is the key of the Meeting action, which is a subtype of Contact event F. The word [Predicted Trigger Word T_p] is the key of the Crime Arrest or sent into Jail action, which is a subtype of Justice event G. The word [Predicted Trigger Word T_p] is the key of the Demonstrate action, which is a subtype of Conflict event H. The word [Predicted Trigger Word T_p] is the key of the Die action, which is a subtype of Life event I. The word [Predicted Trigger Word T_p] is the key of the Transfer-Money action, which is a subtype of Transaction event <p>Which answer can be inferred?</p> <p>Answer:</p>

Table 10: Instruction variants in prompts used in the Vanilla approach for robustness evaluation in MNER task. Variations in the instructions are highlighted in pink.

Formulations	Prompts
Instruction 1	Please choose the entity spans corresponding to the [Entity category E_c] entity type that can be inferred from the given sentence. Entities:
Instruction 2	Choose the spans of [Entity category E_c] entity that can be inferred from the given sentence. Entities:
Instruction 3	Please choose the spans related to the [Entity category E_c] category that can be deduced from the presented sentence. Entities:
Instruction 4	Decide on the spans associated with the [Entity category E_c] entity category that can be inferred from the given sentence. Entities:

Table 11: Instruction variants in prompts used in the Vanilla approach for robustness evaluation in MIED task. Variations in the instructions are highlighted in pink. The content of possible types of events are identical across several instruction variants, thus they are abbreviated with “.....”. Please refer to Table 7 for a detailed overview.

Formulations	Prompts
Instruction 1	Given an image and a sentence, classify which type of activity can be inferred. Sentence: This is an image attached to an news article. All possible types are listed below: Type:
Instruction 2	Given an image, classify which type of activity can be inferred. Sentence: This is an image attached to a news article. All possible types are listed below: Type:
Instruction 3	Given an image, identify the event-related type that can be deduced from the provided image. Sentence: This is an image attached to an news article. All possible types are listed below: Type:
Instruction 3	Given the image, please choose which type of activity can be inferred. Sentence: This is an image attached to a news article. All possible types are listed below: Type:

Table 12: Instruction variants in prompts used in the Vanilla approach for robustness evaluation in MRE task. Variations in the instructions are highlighted in pink.

Formulations	Prompts
Instruction 1	<p>Given the image and the text, select the relation between Entity 1 and Entity 2 that can be inferred from the given sentence, The image may provide fine-grained information about the entities. All possible relations are listed below:</p> <p>-[Possible Relation 1] -[Possible Relation 2] -None</p> <p>Sentence: [Sentence S] Relation:</p>
Instruction 2	<p>Given the image and the text, determine which relation between Entity 1 and Entity 2 can be deduced. The image might offer detailed insights about the entities. All possible relations are listed below:</p> <p>-[Possible Relation 1] -[Possible Relation 2] -None</p> <p>Sentence: [Sentence S] Relation:</p>
Instruction 3	<p>Based on the image and the text, determine which relation is the most possible between Entity 1 and Entity 2. The image could offer intricate details about the entities. All possible relations are listed below:</p> <p>-[Possible Relation 1] -[Possible Relation 2] -None</p> <p>Sentence: [Sentence S] Relation:</p>
Instruction 4	<p>In light of the image and the text, determine what relation between Entity 1 and Entity 2 is the most feasible. The image may present detailed information about the entities. All possible relations are listed below:</p> <p>-[Possible Relation 1] -[Possible Relation 2] -None</p> <p>Sentence: [Sentence S] Relation:</p>

Table 13: Instruction variants in prompts used in the Vanilla approach for robustness evaluation in MTED task. Variations in the instructions are highlighted in pink.

Formulations	Prompts
Instruction 1	<p>Given an image and a sentence, determine which word can infer the [Event category E_c] activity.</p> <p>Sentence: [Sentence S]</p> <p>Words:</p>
Instruction 2	<p>Given an image and a sentence, classify the right word in the sentence that can infer the [Event category E_c] activity.</p> <p>Sentence: [Sentence S]</p> <p>Words:</p>
Instruction 3	<p>Given an image and a sentence, determine the right word in the sentence that can infer the [Event category E_c] event.</p> <p>Sentence: [Sentence S]</p> <p>Words:</p>
Instruction 4	<p>Given an image and a sentence, select the correct word within the sentence that can infer the [Event category E_c] activity.</p> <p>Sentence: [Sentence S]</p> <p>Words:</p>

Table 14: Instruction variants in prompts used in the Multi-choice QA stage of MQA approach for robustness evaluation in MNER task. Variations in the instructions are highlighted in pink.

Formulations	Prompts
Instruction 1	<p>Choose the right answer about the entity that can be inferred from the given sentence.</p> <p>Sentence: [Sentence S]</p> <p>Options:</p> <p>A. [Entity Span] is a location entity B. [Entity Span] is a person entity C. [Entity Span] is an organization entity D. [Entity Span] is not a named entity or does not belong to type [location, person, organization, miscellaneous]</p> <p>Which answer can be inferred from the given sentence? Answer:</p>
Instruction 2	<p>Select the correct option regarding the entity that can be deduced from the provided sentence.</p> <p>Sentence: [Sentence S]</p> <p>Options:</p> <p>A. [Entity Span] is a location entity B. [Entity Span] is a person entity C. [Entity Span] is an organization entity D. [Entity Span] is not a named entity or does not belong to type [location, person, organization, miscellaneous]</p> <p>Which answer can be inferred from the given sentence? Answer:</p>
Instruction 3	<p>Pick the accurate answer concerning the entity that can be implied from the presented sentence.</p> <p>Sentence: [Sentence S]</p> <p>Options:</p> <p>A. [Entity Span] is a location entity B. [Entity Span] is a person entity C. [Entity Span] is an organization entity D. [Entity Span] is not a named entity or does not belong to type [location, person, organization, miscellaneous]</p> <p>Which answer can be inferred from the given sentence? Answer:</p>
Instruction 4	<p>Choose the right answer pertaining to the entity that can be concluded from the given sentence.</p> <p>Sentence: [Sentence S]</p> <p>Options:</p> <p>A. [Entity Span] is a location entity B. [Entity Span] is a person entity C. [Entity Span] is an organization entity D. [Entity Span] is not a named entity or does not belong to type [location, person, organization, miscellaneous]</p> <p>Which answer can be concluded from the given sentence? Answer:</p>

Table 15: Instruction variants in prompts used in the Multi-choice QA stage of MQA approach for robustness evaluation in MIED task. Variations in the instructions are highlighted in pink. The options are identical across several instruction variants, thus they are abbreviated with “.....”. Please refer to Table 7 for a detailed overview.

Formulations	Prompts
Instruction 1	<p>Given an image, classify which type of activity can be inferred.</p> <p>Sentence: This is an image attached to an news article.</p> <p>Options:</p> <p>Which answer can be inferred from the image?</p> <p>Answer:</p>
Instruction 2	<p>Determine which choice of activity is relevant to the image.</p> <p>Options:</p> <p>Which answer is relevant to the image?</p> <p>Answer:</p>
Instruction 3	<p>Please determine which choice of activity can be inferred from the image.</p> <p>Options:</p> <p>Which option can be inferred from the image?</p> <p>Answer:</p>
Instruction 4	<p>Identify the option that can be deduced from the picture.</p> <p>Sentence: This is an image attached to an news article.</p> <p>Options:</p> <p>What can be deduced from the picture?</p> <p>Answer:</p>

Table 16: Instruction variants in prompts used in the Multi-choice QA stage of MQA approach for robustness evaluation in MRE task. Variations in the instructions are highlighted in pink.

Formulations	Prompts
Instruction 1	<p>In light of the provided sentence, determine which option is the most feasible inference. The image may present detailed information about the entities.</p> <p>Sentence: [Sentence S]</p> <p>Options: A. [Entities in Relation 1 Template] B. [Entities in Relation 2 Template] C. [Entities in NoTA Relation Template]</p> <p>Which option is the most possible inference? Option:</p>
Instruction 2	<p>Select the option that can be inferred from the given sentence. The image may provide fine-grained information about the entities.</p> <p>Sentence: [Sentence S]</p> <p>Options: A. [Entities in Relation 1 Template] B. [Entities in Relation 2 Template] C. [Entities in NoTA Relation Template]</p> <p>Which option can be inferred from the given sentence and image? Option:</p>
Instruction 3	<p>From the provided sentence, determine which option can be deduced. The image might offer detailed insights about the entities.</p> <p>Sentence: [Sentence S]</p> <p>Options: A. [Entities in Relation 1 Template] B. [Entities in Relation 2 Template] C. [Entities in NoTA Relation Template]</p> <p>Which option can be deduced from the given sentence and image? Option:</p>
Instruction 4	<p>Based on the given sentence, determine which option is the most possible deduction. The image could offer intricate details about the entities.</p> <p>Sentence: [Sentence S]</p> <p>Options: A. [Entities in Relation 1 Template] B. [Entities in Relation 2 Template] C. [Entities in NoTA Relation Template]</p> <p>Which option is most possible? Option:</p>

Table 17: Instruction variants in prompts used in the Multi-choice QA stage of MQA approach for robustness evaluation in MTED task. The answer candidates are identical across several instruction variants, thus they are abbreviated with “.....”. Please refer to Table 9 for a detailed overview. Variations in the instructions are highlighted in pink.

Formulations	Prompts
Instruction 1	<p>Determine which option can be inferred from the given sentence and image.</p> <p>Sentence: [Sentence S]</p> <p>Answer candidates:</p> <p>Which answer can be inferred?</p> <p>Answer:</p>
Instruction 2	<p>Identify the choice that can be deduced from the provided sentence and image.</p> <p>Sentence: [Sentence S]</p> <p>Answer candidates:</p> <p>Which answer can be inferred?</p> <p>Answer:</p>
Instruction 3	<p>Choose the right answer about a word in the sentence that can be inferred.</p> <p>Sentence: [Sentence S]</p> <p>Answer candidates:</p> <p>Which answer can be inferred?</p> <p>Answer:</p>
Instruction 4	<p>Select the correct answer concerning a word within the sentence that can be inferred.</p> <p>Sentence: [Sentence S]</p> <p>Answer candidates:</p> <p>Which answer can be inferred?</p> <p>Answer:</p>

Table 18: Prompts employed for ChatGPT/GPT4 in the context of text-based MIE tasks.

Task	Prompts
MNER	<p>Extract the entity span that belongs to [Entity category E_c] entity category from the provided sentence. Output format is entity1. If no [Entity category E_c] entity can be inferred, respond with "".</p> <p>Output:</p>
MRE	<p>Determine the what relation between Entity 1 and Entity 2 is according to the text. All possible relations are listed below:</p> <ul style="list-style-type: none"> -[Possible Relation 1] -[Possible Relation 2] -None <p>Sentence: [Sentence S] Relation:</p>
MTED	<p>Extract the most one trigger word from the provided sentence that infer the [Event category E_c] activity. Note that the trigger word can only be a noun or verb. The answer format should be word1. If no trigger word reflecting the [Event category E_c] activity can be identified, respond with "".</p> <p>Sentence: [Sentence S] Words:</p>