

Recognizing Everything from All Modalities at Once: Grounded Multimodal Universal Information Extraction

Meishan Zhang¹, Hao Fei^{2*}, Bin Wang¹, Shengqiong Wu², Yixin Cao³, Fei Li⁴, Min Zhang¹

¹ Harbin Institute of Technology (Shenzhen), ² National University of Singapore,

³ School of Computer Science, Fudan University, ⁴ Wuhan University

zhangmeishan@hit.edu.cn, haofei37@nus.edu.sg, 23s051047@stu.hit.edu.cn, swu@u.nus.edu

caoyixin2011@gmail.com, lifei_csnlp@whu.edu.cn, zhangmin2021@hit.edu.cn

Abstract

In the field of information extraction (IE), tasks across a wide range of modalities and their combinations have been traditionally studied in isolation, leaving a gap in deeply recognizing and analyzing cross-modal information. To address this, this work for the first time introduces the concept of *grounded Multimodal Universal Information Extraction (MUIE)*, providing a unified task framework to analyze any IE tasks over various modalities, along with their fine-grained groundings. To tackle MUIE, we tailor a multimodal large language model (MLLM), REAMO, capable of extracting and grounding information from all modalities, i.e., ‘*recognizing everything from all modalities at once*’. REAMO is updated via varied tuning strategies, equipping it with powerful capabilities for information recognition and fine-grained multimodal grounding. To address the absence of a suitable benchmark for grounded MUIE, we curate a high-quality, diverse, and challenging test set, which encompasses IE tasks across 9 common modality combinations with the corresponding multimodal groundings. The extensive comparison of REAMO with existing MLLMs integrated into pipeline approaches demonstrates its advantages across all evaluation dimensions, establishing a strong benchmark for the follow-up research. Our resources are publicly released at <https://haofei.vip/MUIE>.

1 Introduction

IE is a pivotal topic (Li et al., 2022a), encompassing subtasks such as Named Entity Recognition (NER; Nadeau and Sekine, 2007), Relation Extraction (RE; Miwa and Bansal, 2016), and Event Extraction (EE; Ahn, 2006), which plays a crucial role in constructing domain-specific knowledge bases (Bosselut et al., 2019) and in facilitating deep semantic understanding of data (Satyapanich et al.,

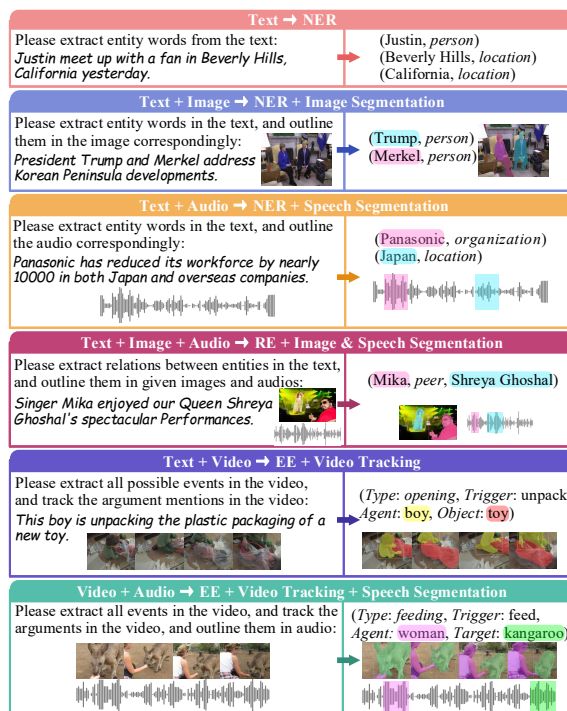


Figure 1: Examples of grounded multimodal universal information extraction (MUIE).

2020). In reality, a vast amount of information is conveyed through modalities beyond text. Consequently, research in IE has evolved from focusing solely on textual data to embracing various other modalities, leading to the development of multimodal IE (MIE; Liu et al., 2019), e.g., images, videos, and audio. Despite the growing research efforts dedicated to MIE, the exploration in this area remains insufficiently developed. We argue that several critical aspects must be fully considered for future MIE research trends.

Firstly, current studies primarily investigate MIE tasks within individual modalities (or certain modality combinations) (Sun et al., 2021; Chen et al., 2022b). With the existence of several modality categories and diverse definitions for different IE tasks, studying each modality separately to construct specialized MIE models would inevitably

*Corresponding author

lead to resource wastage and inefficiency. In real-world applications, there is a constant need for building unified systems with “one-for-all” robust generalizability for faster practical deployment. In light of the recent success of textual universal IE (UIE; Lu et al., 2022; Fei et al., 2022), MIE unification should also be promising. **Second**, the majority of existing studies (Zhang et al., 2017) exhibit a bias toward text-centric IE outputs, necessitating the decoding of detailed textual IE labels and inherently prioritizing text. While they often treat other modalities as auxiliaries and do not produce outputs for them, this practice does not align with reality, because all modalities can equally carry important information. For example, even infants who have not yet learned to speak can learn about entities from vision. Thus, each modality should be treated equally, and detect fine-grained information from all given modalities. **Last**, most current MIE (Zheng et al., 2021) involving multiple modalities (e.g., Image&Text) tends to extract the modality-aligned part of the information under the assumption that different modalities associate with each other. However, in practical scenarios, the information carried by different modalities can be either shared (Li et al., 2022b), or unrelated (Wu et al., 2023a). This suggests that information should be flexibly recognized from any modality sources.

In response to these challenges, this paper is dedicated to pioneering a novel task, grounded Multimodal Universal Information Extraction (MUIE). As illustrated in Fig. 1, MUIE aims to unify the modeling of various IE tasks (e.g., NER, RE, EE) with any (or combination) inputs across the most common modalities (e.g., text, audio, image, and video), and produce fine-grained multimodally grounded IE results. To solve MUIE, we consider taking advantage of the existing generative LLMs (OpenAI, 2022; Chung et al., 2022; Chiang et al., 2023) with in-context instructions (Dong et al., 2022). We develop a novel multimodal LLM, REAMO, achieving “*Recognizing Everything from All Modalities at Once*”. REAMO not only outputs all possible textual IE labels but also identifies corresponding groundings across other modalities: 1) statically, by segmenting visual objects and audio speeches, and 2) dynamically, by tracking textual or vocal events in videos. Technically, REAMO employs a Vicuna (Chiang et al., 2023) LLM as its core semantic reasoner, utilizing ImageBind (Girdhar et al., 2023) as a multimodal encoder to project image, video and audio inputs into

LLM-understandable signals. At the decoding end, we integrate the SEEM (Zou et al., 2023) for visual grounding&tracking, and the SHAS (Tsiamas et al., 2022) for audio segmentation, where the messages are passing from LLM to decoders through structured meta-response effectively. Given input multimodal information, REAMO is able to output UIE label tokens as well as fine-grained groundings recurrently.

We then design learning objectives to tune REAMO to endow it with robust MUIE and cross-modal grounding capabilities. First, we repurpose existing textual UIE annotation into instruction format, and use it to tune the backbone LLM for activating the UIE ability. Then, we perform both coarse-grained instance-level and fine-grained grounding-aware cross-modal alignment learning, enhancing the REAMO’s capability in fine-grained multimodal semantic understanding. Furthermore, we instruction-tune REAMO on specific corpus, to build its working behavior of generating structured meta-response texts.

In response to the absence of standard evaluation data for grounded MUIE, we further introduce an evaluation benchmark, where we annotate a high-quality testing set of 3,000 instances covering NER, RE, and EE tasks under 9 common modality combinations. The data further advances by annotating both modality-shared/-specific content to simulate aligned and misaligned modality scenarios. Extensive zero-shot experiments on these benchmarks demonstrate that REAMO shows strong performance over existing MLLMs with respect to IE tasks and multimodal grounding.

Overall, we make three key contributions:

- To our knowledge, this is the first to propose a grounded MUIE setting, unifying all IE tasks across modalities, further with fine-grained multimodally grounded targets.
- We introduce an MLLM for the task, REAMO, excelling in MUIE prediction and achieving cross-modal grounding of static objects and dynamic events.
- We contribute a high-quality, diverse, and challenging dataset, setting an evaluation benchmark for follow-up grounded MUIE research.

2 Related Works

IE (Lample et al., 2016; Wei et al., 2020; Fei et al., 2021; Li et al., 2021; Wang et al., 2022; Cao et al., 2022) has long been a significant research direction, consistently attracting substantial interest and focus

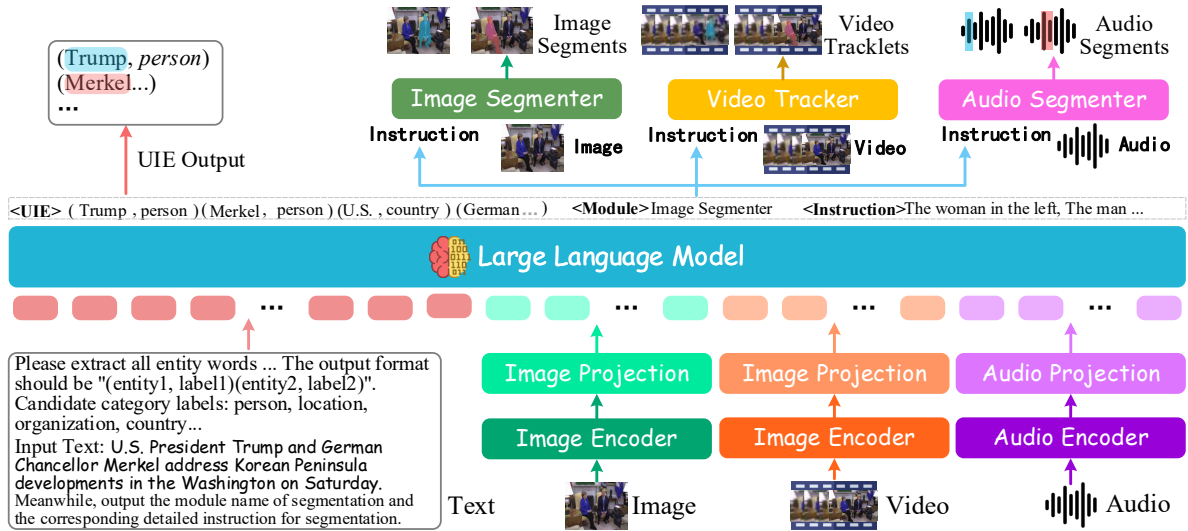


Figure 2: An overview of the proposed REAMO MLLM architecture for grounded MUIE.

for decades. As the world contains various modalities of information, MIE has been consequently introduced, e.g., multimodal NER (Sun et al., 2021; Yu et al., 2023; Wang et al., 2023a), multimodal RE (Chen et al., 2022b), and multimodal EE (Li et al., 2020). However, the majority of existing well-established benchmarks for MIE still predominantly focus on texts, supplemented with images (Zhang et al., 2017; Wu et al., 2023a).

Historically, IE research has treated different tasks as separate studies for a long (Sun et al., 2021; Chen et al., 2022b; Liu et al., 2024; Zhang et al., 2024). Recently, Lu et al. (2022) pioneer UIE, proposing to unify all IE tasks under a single generative model to produce all IE results, significantly reducing the maintenance cost for individual tasks. With the latest rapid development of LLMs, the latest advancements have utilized LLMs with in-context prompting for UIE, achieving promising zero-shot performance. Similarly, the swift progress of MLLMs (Fei et al., 2023) should also ignite hope for MUIE. Yet research in MUIE remains under-explored.

To our knowledge, the work most closely related to this paper is by Sun et al. (2023), who leverage existing MLLMs to unify various MIE tasks through a two-stage process of span extraction and classification in a multimodal QA format. Yet, we identify clear limitations in their approach that fall short of achieving comprehensive MIE unification. Firstly, while their work only considers text and image modalities, lacking comprehensiveness, we broadly cover the four most common modalities. Secondly, beyond addressing these limitations, we also introduce a novel MLLM tailored for grounded

MUIE and contribute a new benchmark dataset for MUIE research. These efforts aim to pioneer the next stage of MUIE research.

3 Task Definition: Grounded Multimodal Universal Information Extraction

We now give a formal definition of grounded MUIE. Suppose the inputs are any of a text T , an image I , an audio A , a video V , or their combination.

NER task seeks to predict all possible textual labels of entities $\{E^{ner}\}$, with pre-defined entity types $C^{ner} \in \mathcal{C}^{ner}$ (e.g., person, location and organization), where each E may correspond to a span within T , or visual region within I , or a speech segment within A . We denote the visual grounding mask as M_{img} and the speech segment as M_{aud} .

RE task aims to first identify all possible entities $\{E^{re}\}$ following the NER step, and then determine a pre-defined relation label $R^{RE} \in \mathcal{R}^{RE}$ for two entities $\langle E_i^{RE}, E_j^{RE} \rangle$ that should be paired. Also E^{RE} should correspond to T , I , or A , as in NER.

EE task detects all possible structured event records that consist of event trigger E^{et} , event type $C^{et} \in \mathcal{C}^{et}$, event argument E^{ea} and event argument role $C^{er} \in \mathcal{C}^{er}$. Here E^{et} and E^{ea} correspond to a continuous span within T or a speech segment within A . Also E^{ea} might refer to the visual region within I , or the temporal dynamic tracklet in video V (i.e., object tracking). We denote the video tracking mask as M_{vid} . C^{et} and C^{er} are pre-defined label sets.

Following Wang et al. (2023b), we employ in-context learning (ICL; Dong et al., 2022) to prompt LLMs for MUIE, with the specific task executed depending on the user’s intention. In the bottom

left of Fig. 2 we simply illustrate the ICL prompt.

4 Our Proposed Model

4.1 MLLM Framework of REAMO

Fig. 2 presents a schematic overview of REAMO MLLM, which consists of three main parts: multi-modal encoder, LLM, and decoder for UIE prediction & multimodal grounding.

Multimodal Encoding. While REAMO takes four types of modality sources, except texts that are directly input to LLM, the image, audio, and video inputs should be encoded. Following Wu et al. (2023b); Fei et al., 2024b), we leverage the high-performance ImageBind (Girdhar et al., 2023) as a unified multimodal encoder. Then, via a projection layer, different input representations are aligned into language-like embeddings that are understandable to the LLM.

LLM Reasoner. An LLM serves as the center unit of REAMO for content semantics understanding and reasoning. Specifically, we follow the most common practice, using the Vicuna-v1.5 (Chiang et al., 2023) as the backbone LLM.



MUIE Decoding with Grounding. The central LLM takes on the crucial role of decision-making. Based on the input prompt, LLM will produce textual responses, containing the UIE task results as well as the meta-response that will be used to call the downstream modules to generate fine-grained multimodal groundings. Here, consider integrating the existing high-performance SEEM model (Zou et al., 2023) for image segmentation and

video tracking (Li et al., 2024; Wu et al., 2024; Fei et al., 2024a), and the SHAS model (Tsiamas et al., 2022) for audio segmentation. Specifically, the meta-response includes three parts, as shown above. The corresponding '<Module>' and '<Instruction>' information will be passed to the corresponding modality segmenter(s), to activate them to generate grounding(s).

4.2 MUIE Fine-tuning for REAMO

With REAMO at hand, we now consider fine-tuning it through multiple objectives to enable REAMO with strong MUIE capability.

UIE Instruction Tuning. Our initial goal is to equip the system with the fundamental capability for UIE in the text modality. To achieve this, we consider tuning the backbone LLM specifically for UIE. Following the practices of Wang et al. (2023b), we repurpose existing annotation data to form a set of instruction-tuning datasets for UIE. To avoid the huge cost of fully updating the LLM, we leverage the LoRA technique (Hu et al., 2022), achieving the goal by tuning only a small subset of parameters, without altering the overall LLM.

Multimodal Alignment Learning. REAMO integrates the ImageBind encoder to enable the LLM to comprehend basic multimodal signals. Following this, we proceed with multimodal alignment learning. We consider using the language-centric LLM as the core, requiring only the alignment of other modalities to text. We mainly utilize the vast array of available 'X-caption' pair data ('X' stands for image, audio, or video). We adopt an 'X-to-text' generation, where the input is 'X', and the LLM generates the corresponding caption. During this process, we fix the ImageBind and LLM while only updating the projection layer.

Fine-grained Cross-modal Grounding-aware Tuning. Above we merely enable REAMO with a coarse-grained multimodal understanding. Yet our goal is to attain a subtle modality comprehension. Thus, we further engage in fine-grained multimodal grounding. Our primary approach revolves around utilizing existing 'X-to-text' phrase grounding data, e.g., MS-COCO (Lin et al., 2014), where REAMO's input consists of textual phrases and grounded regional modality features, and then prompt LLM to determine their match.

Invocation-based Meta-response Tuning. To teach the LLM to produce the correct invocation meta-responses, we need to create data specifically

Modality	Tasks		
	NER	RE	EE
I	PASCAL-C (Mottaghi et al., 2014)	VRD (Lu et al., 2016)	imSitu (Yatskar et al., 2016)
V			VidSitu (Sadhu et al., 2021)
A	ACE05-Aud (Walker et al., 2011)	ReTACRED (Wu et al., 2022)	
T+I	Twt17 (Lu et al., 2018)	MNRE (Zheng et al., 2021)	M ² E ² (Li et al., 2020)
T+V			VidSitu-Txt (Sadhu et al., 2021)
T+A	ACE05-Aud (Walker et al., 2011)	ReTACRED (Wu et al., 2022)	
I+A		MNRE-Aud (Zheng et al., 2021)	
T+I+A	Twt17-Aud (Lu et al., 2018)		
V+A			VidSitu-Aud (Sadhu et al., 2021)

Table 1: Summary of the grounded MUIE test data. Items in the light yellow background mean they are the data after preprocessing, i.e., via modality translation, where the colored postfix means the target modality.

designed for instruction tuning. Technically, we make use of the existing annotated datasets for various vision tasks included in this work. For each task under specific different user input scenarios, with the corresponding data, we design various template dialogue-format examples. Based on these examples we then prompt the GPT-4 to generate more samples under various topics and enriched scenarios.

5 A Benchmark for Grounded MUIE

To evaluate the performance of our grounded MUIE system, we develop a benchmark testing set.

5.1 Data Source

We select 9 existing datasets from different modalities (or combinations thereof) for IE/MIE tasks. Table 1 summarizes these datasets of the raw sources. We then process these datasets, such as Text \leftrightarrow Speech, to create 6 new datasets under new multimodal (combination) scenarios. Before annotation, we carefully select 200 instances from their corresponding testing sets, ensuring each instance contained as much IE information as possible.

- **PASCAL-C**¹: is an object detection dataset for evaluating the robustness.
- **VRD**², Visual Relationship Dataset: is designed to assess the precision of detecting interactions among pairs of objects. Comprising 5,000 images, the dataset encompasses 100 object categories and 70 predicates.
- **imSitu**³: serves as a resource for facilitating situation recognition, a task concerned with generating a succinct depiction of the scenario

¹<https://github.com/bethgelab/robust-detection-benchmark>

²<https://cs.stanford.edu/people/ranjaykrishna/vrd/>

³<http://imsitu.org/>

portrayed in images. This includes (1) the main activity, (2) the participating actors, objects, substances, and locations and most importantly (3) the roles these participants play in the activity.

- **ACE2005**⁴: is extensively used in information extraction. It comprises annotated news articles in English, covering a diverse range of topics and events, with annotations including named entities, relations, and events.
- **ReTACRED**⁵: is a revised version of TACRED for relation detection, containing over 91k sentences spread across 40 relations.
- **VidSitu**⁶: is a large-scale dataset containing diverse 10-second videos from movies depicting complex situations (a collection of related events). Events in the video are richly annotated at 2-second intervals with verbs, semantic-roles, entity co-references, and event relations.
- **Twt17**⁷, Twitter-17: is a publicly available Twitter dataset for NER. It encompasses 723 test tweets, with annotations covering four entity types, namely, *person*, *location*, *organization*, *miscellaneous*.
- **MNRE**⁸, Multimodal Neural Relation Extraction: comprises 15,484 samples and 9,201 accompanying images across 23 distinct relation categories, partitioning into training, development, and testing subsets, consisting of 12,247, 1,624, and 1,614 samples respectively.
- **M²E²**⁹: is comprised of 245 multimedia news articles meticulously annotated with

⁴<http://projects.ldc.upenn.edu/ace/>

⁵<https://github.com/gstoica27/Re-TACRED>

⁶<https://vidsitu.org/>

⁷<https://github.com/jefferyYu/UMT>

⁸<https://github.com/thecharm/MNREp>

⁹<https://github.com/limanling/m2e2>

events and their corresponding arguments.

Pre-processing And Modality Translation Before we begin the annotation work for grounding, we enrich the types of modality combinations. This is essential because most MIE datasets currently focus on the combination of images and text. However, we aim to simulate a variety of common modality combinations that could occur in real-world scenarios. To achieve this, we transform and preprocess existing datasets. Specifically, our approach involves cross-modal parallel translation to generate data in another modality. We preprocess the following datasets:

- **VidSitu-Aud**: we start by captioning videos from the VidSitu dataset, then use the given video event annotations combined with the captions to have ChatGPT generate a coherent sentence, serving as the paired text for each piece of data.
- **VidSitu-Aud**: Based on VidSitu-Txt, we convert each sentence into speech using Text-To-Speech (TTS) tools. We employ state-of-the-art open-source TTS models: Bark¹⁰ and Edge-TTS¹¹.
- **ACE-Aud**: we take original textual sentences from the ACE dataset and record speech using TTS technology.
- **MNRE-Aud**: we record speech for sentences from the MNRE using TTS.
- **Tw17-Aud**: Similarly, we record speech for sentences from the Tw17 dataset using TTS.

However, we emphasize that such parallel data generation can only produce modality-aligned content. To create diverse content, we plan to introduce randomness by adding noise to some instances. For example, we might alter parts of the original text before synthesizing the speech with TTS. Considering the aim to only annotate a test set for the sake of conserving labor and reducing costs, we plan to select subsets from the test sets of various original datasets obtained. We aim to select 200 from each. Our selection criterion focuses on ensuring a high quantity of entities and objects in the content, and the final labels should cover a rich vocabulary. Given there are 15 combinations of modalities and tasks (including those augmented through our post-processing), we will have a total of 3,000 data entries of grounded MUIE test instances. Next, we revisit the annotation information, specifically

re-annotating instances from the combined modality datasets where cross-modal content is not fully aligned. This ensures that the dataset covers both modality-shared and modality-specific instances.

6 Experiments

6.1 Settings

We measure the system performance by following most practices of end-to-end UIE: F1 of entity span with type for NER, F1 of all subject&object entities and their relation label for RE. For EE, we consider event trigger (ET) F1 including both trigger and event type, and event argument F1 including both arguments with role types. W.r.t. the multimodal grounding, for both image and audio segmentation, we consider the mean Intersection over Union (mIoU); for video segmentation, we use the average Jaccard (J).

The encoder projection is a linear layer with a hidden size of 4,096. As no prior method is designed for grounded MUIE, we implement pipeline systems. Specifically, we first employ an MLLM to do UIE on multimodal input. And then we pass the raw multimodal source and the necessary UIE label to the SEEM or SHAS model for fine-grained grounding of image, video and audio. We consider the following existing well-exposed MLLMs. For image-related UIE: InstructBLIP (Dai et al., 2023), LLaVA (Liu et al., 2023), and also MiniGPT-v2 (Chen et al., 2023) that can output image segmentation end-to-end. For or video-related UIE: VideoChat (Li et al., 2023), Video-LLaVA (Lin et al., 2023). SpeechGPT (Zhang et al., 2023a) for audio-related UIE. Video-LLaMA (Zhang et al., 2023b) supporting video+audio; NExT-GPT (Wu et al., 2023b) supporting all four modalities. All these systems take the 7B LLM, unless otherwise specified. For fairness, all baselines are further tuned using the same UIE instruction-tuning data as ours. All system takes zero-shot inference, without tuning on in-house datasets.

6.2 Zero-shot Results on Image-related MUIE

Table 2 presents the results of different models on our MUIE dataset under both text+image and pure image conditions. From the data, we observe that overall, the end-to-end approach (MiniGPT-v2) outperforms pipeline methods. Besides, our system demonstrates a clear and consistent advantage consistently across all subtasks.

¹⁰<https://github.com/suno-ai/bark>

¹¹<https://github.com/rany2/edge-tts>

Method	T+I Input							I Input						
	Twt17		MNRE		M ² E ²			PASCAL-C		VRD		imSitu		
	NER	I-Seg	RE	I-Seg	ET	EA	I-Seg	NER	I-Seg	RE	I-Seg	ET	EA	I-Seg
LLaVA+SEEM	23.0	45.8	15.4	51.8	22.8	13.5	48.3	17.8	26.1	10.4	36.9	19.5	8.2	29.8
InstructBLIP+SEEM	26.9	52.0	17.0	54.8	23.0	20.1	52.5	21.0	39.7	11.6	39.1	18.5	11.3	30.0
MiniGPT-v2	45.4	48.7	22.4	56.2	27.3	16.3	54.8	41.8	62.0	18.6	38.3	37.0	13.3	32.2
REAMO	47.4	53.5	24.6	56.9	30.2	25.6	60.1	43.0	64.6	26.0	43.9	41.5	16.3	39.6

Table 2: Zero-shot performance in the UIE scenario of text+image or standalone image input. I-Seg: grounding by image segmentation. Performance of REAMO is with blue background.

Method	T+A Input				A Input			
	ACE05-Aud		ReTACRED		ACE05-Aud		ReTACRED	
	NER	A-Seg	RE	A-Seg	NER	A-Seg	RE	A-Seg
SpeechGPT	26.7	21.4	45.4	27.5	14.0	13.3	30.4	21.0
NExT-GPT+SHAS	19.6	15.6	37.5	20.4	8.3	10.2	25.1	12.4
REAMO	28.5	24.3	46.8	29.1	17.4	16.7	33.4	25.1

Table 3: Zero-shot performance in the text+audio or standalone audio input scenarios.

Method	T+V (VidSitu-Txt)			V (VidSitu)		
	ET	ER	V-Trck	ET	ER	V-Trck
VideoChat+SEEM	28.8	18.5	28.1	14.3	9.2	20.9
Video-LLaVA+SEEM	31.0	22.4	31.4	18.6	8.8	20.6
REAMO	32.8	23.1	34.4	22.3	14.5	23.2

Table 4: Zero-shot results in the text+video or standalone video input scenarios.




Figure 3: Performance gap between modality-shared (aligned) and modality-specific (unaligned) MUIE.

6.3 Zero-shot Results on Audio-related MUIE

In Table 3, we present the performance of various models under text+audio and pure audio input, respectively. The trends observed are similar to those seen in the previous table: 1) End-to-end approaches (SpeechGPT) demonstrate stronger performance compared to the pipeline method (NExT-GPT+SHAS), effectively mitigating the issue of error propagation. 2) Our REAMO consistently outperforms others across all subtasks and scenarios.

6.4 Zero-shot Results on Video-related MUIE

In Table 4, we present the final set of results for EE task based on text+video and pure video input. The




Figure 4: Impact of different object/entity numbers.

overall trend observed here again aligns with that of the previous tables, with our model achieving the strongest performance. Following, we can confirm that for single modalities (such as image, audio, video), without the assistance of textual modality, the effectiveness of MUIE significantly diminishes.

6.5 Results on Modality-Compound MUIE

Finally, we present the performance of our model across a broader range of complex mixed-modality scenarios for MUIE. As shown in Table 5, our REAMO consistently delivers the best results than the baseline solutions. This becomes even more pronounced in complex mixed-modality contexts.

Method	T+I+A (Twf17-Aud)			I+A (MNRE-Aud)			V+A (VidSitu-Aud)			
	NER	I-Seg	A-Seg	RE	I-Seg	A-Seg	ET	ER	V-Trck	A-Seg
Video-LLaMA+SEEM/+SHAS	-	-	-	-	-	-	12.0	4.8	12.7	8.4
NExT-GPT+SEEM/+SHAS	30.7	32.4	13.9	15.4	46.5	18.8	19.3	13.7	19.9	15.0
REAMO	37.4	33.3	15.1	21.8	53.4	21.8	24.2	18.5	22.0	20.9

Table 5: Zero-shot performance in more complex modality-hybrid scenarios of MUIE.

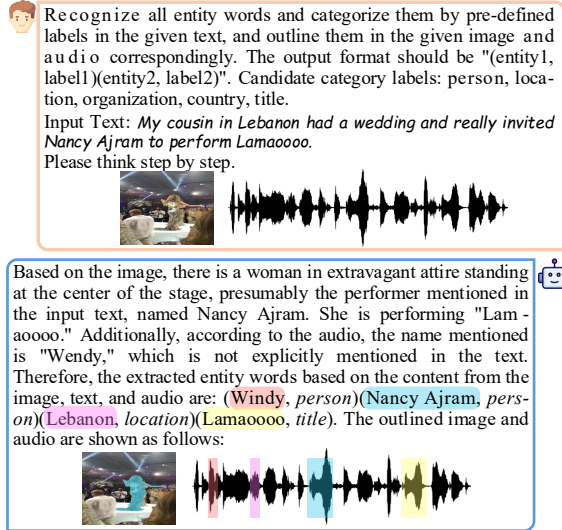


Figure 5: Qualitative result A on MUIE (NER) with modality-specific case via reasoning.

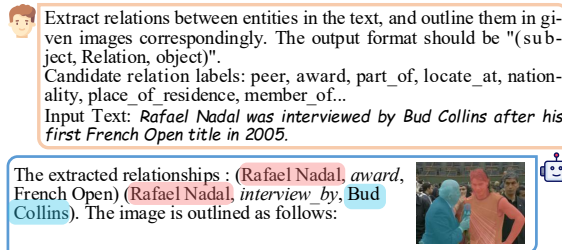


Figure 6: Qualitative result B on MUIE (RE) with grounding rationale via reasoning.

6.6 Gap between Modality-Shared vs. Modality-Specific MUIE

Previously, we mentioned that in the MUIE context, different models may encounter both aligned shared information and misalignments between modalities. Now, we have categorized cases in the data as either Modality-Shared or Modality-Specific and compared the performance of various models under these two scenarios. As illustrated in Fig. 3, REAMO outperforms all baseline models in addressing challenges arising from misalignments between modalities.

6.7 Influence of Entity/Object Numbers

Fig. 4 further illustrates the impact of the number of entities (or objects) in an instance on the performance of the MUIE system. It is evident that our REAMO system maintains commendable

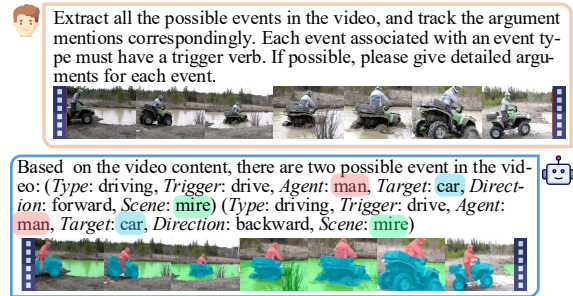


Figure 7: Qualitative result C on MUIE (EE) with commonsense-aware cognitive reasoning.

performance in extracting an increasing number of objects across T+I/T+A/T+V scenarios, both in terms of MUIE results and fine-grained grounding, being clearly superior to baseline pipeline systems.

6.8 Case Study

Then, we provide visualizations of case studies, through which we aim to offer a more intuitive and comprehensive demonstration of our MUIE system. Fig. 5, 6, and 7 each displays an example for NER, RE, and EE tasks, respectively, where REAMO yields correct answers. In case A, our system demonstrates how it can engage in thought processes and flexibly determine MUIE labels from different modalities. In case B, the model can provide accurate relation extractions and precise segmentation results that correspond to the entities in the input text. In case C, most impressively, it manages to link semantic information to factual commonsense understanding based on accurate fine-grained cross-modal grounding, and correctly answer the question.

6.9 Error Analysis

Finally, we delve into our model itself, and analyze the possible shortcomings for shedding light on future explorations. Via our experiments, we here summarize several error types:

- **Repetition of Extracted Content:** When text and information from other modalities are not strictly consistent, our method may output different entity names, arguments, or relationships. However, upon integrating information from different modalities, they should corre-

spond to the same entity names, arguments, or relationships.

- **Incomplete Information Extraction:** The outcomes of information extraction are incomplete, such as incomplete named entity recognition, failure to identify relations involving in-depth reasoning, or incomplete identification of event arguments.
- **Incorrect Grounding Match:** The entity or arguments do not always match with the grounding results. For instance, when the text mentions ‘Obama’ and ‘Trump’ and the image depicts both individuals, the image object segmenter fails to ascertain who is ‘Obama’ and ‘Trump’, resulting in an erroneous grounding match.
- **Miss-grounding:** Our model may output entities or arguments without successfully grounding the corresponding regions in the respective image, video, or audio.
- **Over-grounding:** The model may generate multiple instructions and perform grounding in the image, video, or audio, yet no corresponding regions actually exist in the visual or auditory content.
- **Error Propagation:** Since our system operates as a pipeline process, where a meta-response is first produced and then used to invoke functional modules, this sequence introduces error propagation. If there is an issue with the content of the meta-response, the outcomes from the subsequent modules will be incorrect. To address this, we need to develop more advanced end-to-end MLLMs.

7 Conclusion

This paper introduces a novel task of multimodal information extraction setting: grounded Multimodal Universal Information Extraction (MUIE). First, MUIE definition unifies all IE tasks across various modalities, including text, audio, images, and video, with fine-grained multimodally grounded targets. To solve MUIE, we devise REAMO, a novel MLLM that can extract and ground information from all modalities. REAMO is tuned through various strategies to achieve proficiency in recognizing and grounding multimodal information. Further we introduce a high-quality, diverse, and challenging benchmark dataset for evaluating MUIE systems. Experimental results demonstrate REAMO’s stronger performance in extracting and grounding, setting a strong benchmark for the following

grounded MUIE research.

Acknowledgments

This work is supported by the CCF-Baidu Open Fund, and the Research on automatic construction of cross-domain knowledge graph with multimodal information (No. U23B2055)

Limitations

The limitations of this work primarily arise from the following two aspects. From model perspective, our MLLM, designed for Grounded MUIE, handles four common modalities (text, images, video, audio) and may struggle to integrate new ones without significant retraining. Challenges persist in extracting complex implicit information and in cognitive reasoning, with current capabilities mainly supporting video tracking for event extraction. Further research is needed to enhance and expand tracking for other tasks like NER and RE in videos.

From data perspective, we introduced a dataset solely comprising a test set, limiting in-house training opportunities. Future work will focus on enlarging this dataset to include a training set and expanding annotations to cover more modality combinations for the three IE tasks: NER, RE, and EE.

Ethics Statement

The development and application of MUIE systems potentially raise several potential ethical considerations or risks that should be properly treated to ensure responsible research and deployment.

Privacy and Consent. Multimodal data in MUIE systems can include sensitive personal information. Ensuring data is collected and used with explicit consent and in compliance with data protection laws is critical.

Bias and Fairness. MLLMs in MUIE systems may inherit biases from their training data, potentially causing discriminatory effects. Active efforts are required to mitigate such biases to guarantee fairness across diverse groups.

Transparency and Accountability. MUIE systems often lack clear decision-making transparency due to their complexity. Promoting explainable AI with detailed documentation can enhance user and stakeholder trust.

Misuse Potential. The ability of MUIE systems to process diverse data can be exploited for harmful purposes like misinformation or unauthorized surveillance. Establishing robust legal and ethical safeguards is essential to prevent misuse.

References

- David Ahn. 2006. The stages of event extraction. In *Proceedings of the Workshop on Annotating and Reasoning about Time and Events*, pages 1–8.
- Max Bain, Arsha Nagrani, Gül Varol, and Andrew Senior. 2021. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *Proceedings of the ICCV*, pages 1708–1718.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. Comet: Commonsense transformers for automatic knowledge graph construction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4762–4779.
- Hu Cao, Jingye Li, Fangfang Su, Fei Li, Hao Fei, Shengqiong Wu, Bobo Li, Liang Zhao, and Donghong Ji. 2022. OneEE: A one-stage framework for fast overlapping and nested event extraction. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1953–1964.
- Boli Chen, Guangwei Xu, Xiaobin Wang, Pengjun Xie, Meishan Zhang, and Fei Huang. 2022a. AISHELL-NER: named entity recognition from chinese speech. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2022, Virtual and Singapore, 23-27 May 2022*, pages 8352–8356.
- Jun Chen, Deyao Zhu, Xiaoqian Shen, Xiang Li, Zechun Liu, Pengchuan Zhang, Raghuraman Krishnamoorthi, Vikas Chandra, Yunyang Xiong, and Mohamed Elhoseiny. 2023. Minigt-v2: large language model as a unified interface for vision-language multi-task learning. *arXiv preprint arXiv:2310.09478*.
- Xiang Chen, Ningyu Zhang, Lei Li, Yunzhi Yao, Shumin Deng, Chuanqi Tan, Fei Huang, Luo Si, and Huajun Chen. 2022b. Good visual guidance make a better extractor: Hierarchical visual prefix for multimodal entity and relation extraction. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1607–1618.
- 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. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, and Jason Wei. 2022. Scaling instruction-finetuned language models. *CoRR*, abs/2210.11416.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven C. H. Hoi. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning. *CoRR*, abs/2305.06500.
- Achal Dave, Tarasha Khurana, Pavel Tokmakov, Cordelia Schmid, and Deva Ramanan. 2020. TAO: A large-scale benchmark for tracking any object. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part V*, pages 436–454.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey for in-context learning. *arXiv preprint arXiv:2301.00234*.
- Hao Fei, Qian Liu, Meishan Zhang, Min Zhang, and Tat-Seng Chua. 2023. Scene graph as pivoting: Inference-time image-free unsupervised multimodal machine translation with visual scene hallucination. In *Proceedings of the ACL*, pages 5980–5994.
- Hao Fei, Shengqiong Wu, Wei Ji, Hanwang Zhang, Meishan Zhang, Mong Li Lee, and Wynne Hsu. Video-of-thought: Step-by-step video reasoning from perception to cognition. In *Proceedings of the International Conference on Machine Learning, ICML*, pages 6373–6391.
- Hao Fei, Shengqiong Wu, Jingye Li, Bobo Li, Fei Li, Libo Qin, Meishan Zhang, Min Zhang, and Tat-Seng Chua. 2022. Lasuie: Unifying information extraction with latent adaptive structure-aware generative language model. In *Proceedings of the Advances in Neural Information Processing Systems, NeurIPS 2022*, pages 15460–15475.
- Hao Fei, Shengqiong Wu, Yafeng Ren, Fei Li, and Donghong Ji. 2021. Better combine them together! integrating syntactic constituency and dependency representations for semantic role labeling. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 549–559.
- Hao Fei, Shengqiong Wu, Hanwang Zhang, Tat-Seng Chua, and Shuicheng Yan. 2024a. Vitron: A unified pixel-level vision llm for understanding, generating, segmenting, editing. *CoRR*.
- Hao Fei, Shengqiong Wu, Meishan Zhang, Min Zhang, Tat-Seng Chua, and Shuicheng Yan. 2024b. Enhancing video-language representations with structural spatio-temporal alignment. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Manant Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. 2023. Imagebind: One embedding space to bind them all. *CoRR*, abs/2305.05665.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. Lora: Low-rank adaptation of large language models. In *Proceedings of the ICLR*.
- Chris Dongjoo Kim, Byeongchang Kim, Hyunmin Lee, and Gunhee Kim. 2019. Audiocaps: Generating captions for audios in the wild. In *Proceedings of the NAACL*, pages 119–132.

- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. In *Proceedings of the NAACL*, pages 260–270.
- Jingye Li, Hao Fei, Jiang Liu, Shengqiong Wu, Meishan Zhang, Chong Teng, Donghong Ji, and Fei Li. 2022a. Unified named entity recognition as word-word relation classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10965–10973.
- Jingye Li, Kang Xu, Fei Li, Hao Fei, Yafeng Ren, and Donghong Ji. 2021. MRN: A locally and globally mention-based reasoning network for document-level relation extraction. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1359–1370.
- Kunchang Li, Yanan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. 2023. Videochat: Chat-centric video understanding. *CoRR*, abs/2305.06355.
- Manling Li, Ruochen Xu, Shuohang Wang, Luwei Zhou, Xudong Lin, Chenguang Zhu, Michael Zeng, Heng Ji, and Shih-Fu Chang. 2022b. Clip-event: Connecting text and images with event structures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16420–16429.
- Manling Li, Alireza Zareian, Qi Zeng, Spencer Whitehead, Di Lu, Heng Ji, and Shih-Fu Chang. 2020. Cross-media structured common space for multimedia event extraction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2557–2568.
- Xiangtai Li, Haobo Yuan, Wei Li, Henghui Ding, Size Wu, Wenwei Zhang, Yining Li, Kai Chen, and Chen Change Loy. 2024. Omg-seg: Is one model good enough for all segmentation? In *CVPR*.
- Bin Lin, Yang Ye, Bin Zhu, Jiayi Cui, Munan Ning, Peng Jin, and Li Yuan. 2023. Video-llava: Learning united visual representation by alignment before projection. *CoRR*, abs/2311.10122.
- 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 *Proceedings of the ECCV*, pages 740–755.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. *CoRR*, abs/2304.08485.
- Jiang Liu, Hao Fei, Fei Li, Jingye Li, Bobo Li, Liang Zhao, Chong Teng, and Donghong Ji. 2024. Tkdp: Threefold knowledge-enriched deep prompt tuning for few-shot named entity recognition. *IEEE Transactions on Knowledge and Data Engineering*.
- Xiaoqing Liu, Feiyu Gao, Qiong Zhang, and Huasha Zhao. 2019. Graph convolution for multimodal information extraction from visually rich documents. In *Proceedings of the NAACL*, pages 32–39.
- Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. 2016. Visual relationship detection with language priors. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 852–869. Springer.
- Di Lu, Leonardo Neves, Vitor Carvalho, Ning Zhang, and Heng Ji. 2018. Visual attention model for name tagging in multimodal social media. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1990–1999.
- Yaojie Lu, Qing Liu, Dai Dai, Xinyan Xiao, Hongyu Lin, Xianpei Han, Le Sun, and Hua Wu. 2022. Unified structure generation for universal information extraction. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5755–5772.
- Makoto Miwa and Mohit Bansal. 2016. End-to-end relation extraction using lstms on sequences and tree structures. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics.
- Roosbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Namgyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. 2014. The role of context for object detection and semantic segmentation in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 891–898.
- David Nadeau and Satoshi Sekine. 2007. A survey of named entity recognition and classification. *Linguistic Investigations*, 30(1):3–26.
- OpenAI. 2022. Introducing chatgpt.
- Sebastian Riedel, Limin Yao, and Andrew McCallum. 2010. Modeling relations and their mentions without labeled text. In *Proceedings of the ECML*, pages 148–163.
- Arka Sadhu, Tanmay Gupta, Mark Yatskar, Ram Nevetia, and Aniruddha Kembhavi. 2021. Visual semantic role labeling for video understanding. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Taneeya Satyapanich, Francis Ferraro, and Tim Finin. 2020. CASIE: extracting cybersecurity event information from text. In *Proceedings of the AAAI*, pages 8749–8757.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual captions: A cleaned,

- hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the ACL*, pages 2556–2565.
- Lin Sun, Jiquan Wang, Kai Zhang, Yindu Su, and Fangsheng Weng. 2021. Rpbert: a text-image relation propagation-based bert model for multimodal ner. In *Proceedings of the AAAI conference on artificial intelligence*, 15, pages 13860–13868.
- Yuxuan Sun, Kai Zhang, and Yu Su. 2023. Multimodal question answering for unified information extraction. *arXiv preprint arXiv:2310.03017*.
- Ryuichi Takanobu, Tianyang Zhang, Jiexi Liu, and Minlie Huang. 2019. A hierarchical framework for relation extraction with reinforcement learning. In *Proceedings of the AAAI*, pages 7072–7079.
- Ioannis Tsiamas, Gerard I. Gállego, José A. R. Fonolosa, and Marta R. Costa-jussà. 2022. SHAS: approaching optimal segmentation for end-to-end speech translation. In *Interspeech 2022, 23rd Annual Conference of the International Speech Communication Association, Incheon, Korea, 18-22 September 2022*, pages 106–110.
- Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2011. Ace 2005 multilingual training corpus. *LDC2006T06, Philadelphia, Penn.: Linguistic Data Consortium*.
- Fengqi Wang, Fei Li, Hao Fei, Jingye Li, Shengqiong Wu, Fangfang Su, Wenxuan Shi, Donghong Ji, and Bo Cai. 2022. Entity-centered cross-document relation extraction. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9871–9881.
- Jieming Wang, Ziyang Li, Jianfei Yu, Li Yang, and Rui Xia. 2023a. Fine-grained multimodal named entity recognition and grounding with a generative framework. In *Proceedings of the ACM MM*, pages 3934–3943.
- Xiao Wang, Weikang Zhou, Can Zu, Han Xia, Tianze Chen, Yuansen Zhang, Rui Zheng, Junjie Ye, Qi Zhang, Tao Gui, et al. 2023b. Instructuie: Multi-task instruction tuning for unified information extraction. *arXiv preprint arXiv:2304.08085*.
- Xiaozhi Wang, Ziqi Wang, Xu Han, Wangyi Jiang, Rong Han, Zhiyuan Liu, Juanzi Li, Peng Li, Yankai Lin, and Jie Zhou. 2020. MAVEN: A massive general domain event detection dataset. In *Proceedings of the EMNLP*, pages 1652–1671. Association for Computational Linguistics.
- Zhepei Wei, Jianlin Su, Yue Wang, Yuan Tian, and Yi Chang. 2020. A novel cascade binary tagging framework for relational triple extraction. In *Proceedings of the ACL*, pages 1476–1488.
- Shengqiong Wu, Hao Fei, Yixin Cao, Lidong Bing, and Tat-Seng Chua. 2023a. Information screening whilst exploiting! multimodal relation extraction with feature denoising and multimodal topic modeling. In *Proceedings of the ACL*, pages 14734–14751.
- Shengqiong Wu, Hao Fei, Xiangtai Li, Jiayi Ji, Tat-Seng Chua, Hanwang Zhang, and Shuicheng Yan. 2024. Towards semantic equivalence of tokenization in multimodal llm. *CoRR*.
- Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. 2023b. Next-gpt: Any-to-any multimodal llm. *arXiv preprint arXiv:2309.05519*.
- Tongtong Wu, Guitao Wang, Jinming Zhao, Zhaoran Liu, Guilin Qi, Yuan-Fang Li, and Gholamreza Haffari. 2022. Towards relation extraction from speech. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10751–10762.
- Mark Yatskar, Luke Zettlemoyer, and Ali Farhadi. 2016. Situation recognition: Visual semantic role labeling for image understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5534–5542.
- Jianfei Yu, Ziyang Li, Jieming Wang, and Rui Xia. 2023. Grounded multimodal named entity recognition on social media. In *Proceedings of the ACL*, pages 9141–9154.
- Dong Zhang, Shimin Li, Xin Zhang, Jun Zhan, Pengyu Wang, Yaqian Zhou, and Xipeng Qiu. 2023a. Speechgpt: Empowering large language models with intrinsic cross-modal conversational abilities. *CoRR*, abs/2305.11000.
- Hang Zhang, Xin Li, and Lidong Bing. 2023b. Video-llama: An instruction-tuned audio-visual language model for video understanding. *CoRR*, abs/2306.02858.
- Meishan Zhang, Bin Wang, Hao Fei, and Min Zhang. 2024. In-context learning for few-shot nested named entity recognition. *arXiv preprint arXiv:2402.01182*.
- Tongtao Zhang, Spencer Whitehead, Hanwang Zhang, Hongzhi Li, Joseph Ellis, Lifu Huang, Wei Liu, Heng Ji, and Shih-Fu Chang. 2017. Improving event extraction via multimodal integration. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 270–278.
- Changmeng Zheng, Junhao Feng, Ze Fu, Yi Cai, Qing Li, and Tao Wang. 2021. Multimodal relation extraction with efficient graph alignment. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 5298–5306.
- Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Gao, and Yong Jae Lee. 2023. Segment everything everywhere all at once. *arXiv preprint arXiv:2304.06718*.

A Model Prompt Specification

Here, we show the prompts for each subtask (i.e., NER, RE, EE) in MUIE:

► Prompt for Named Entity Recognition:

Please recognize all entity words and categorize them by pre-defined labels in the given text, and outline them in the given image or video or audio correspondingly. The output format should be “(entity1, label1)(entity2, label2)”. If an entity possibly has a counterpart in the given image or video or audio, please generate a token “<concept>” after the entity word, for subsequent cross-modal grounding.

Candidate category labels: *person, location, organization, country...*

Input Text/Image/Video/Audio:   


► Prompt for Relation Extraction:

Please extract all relations between named entities, and outline them in the given image or video or audio correspondingly. The output format should be “(subject entity, relation, object entity)”. If an entity possibly has a counterpart in the given image or video or audio, please generate a token “<concept>” after the entity word, for subsequent cross-modal grounding.

Candidate relation labels: *peer, award, part_of, locate_at, nationality, place_of_residence, member_of...*

Input Text/Image/Video/Audio:   


► Prompt for Event Extraction:

Extract all the possible events in the video, and track the argument mentions correspondingly. Each event associated with an event type must have a trigger verb. If possible, please give detailed arguments for each event.

Candidate event types: *Marry, Attack, Injure, Be-born, Meet, Transport, Start-position...*

Candidate event argument types: *Agent, Target, Direction, Time, Place, Instrument, Organization, Duration...*

Input Text/Image/Video/Audio:   


B Specification of MUIE Fine-tuning

B.1 UIE Instruction Tuning

In this step, we train only the core LLM. To avoid the significant cost associated with fully updating

the LLM, we employ the LoRA technique, which allows us to achieve our objectives by tuning only a small subset of parameters, thus leaving the overall architecture of the LLM unchanged. The datasets we utilize are sourced from existing IE datasets. We list these datasets in Table 6. These data are converted into an instruction format, following the practices outlined by Wang et al. (2023b).

Task	Data Source	Amount
NER	OntoNotes 5.0 ¹²	76,714
	CoNLL2003 ¹³	20,744
RE	NYT (Riedel et al., 2010)	56,196
	NYT11 HRL (Takanobu et al., 2019)	62,648
EE	MAVEN (Wang et al., 2020)	49,873

Table 6: Datasets used for UIE tuning.

B.2 Multimodal Alignment Learning

To accomplish the alignment, we adopt an ‘X-to-text’ generation task trained on the ‘X-caption’ pair (‘X’ stands for image, audio, or video) data from existing corpus and benchmarks, i.e., given the representation of an ‘X’, to prompt the frozen LLM to generate the corresponding text description. Specifically, we utilize three types of ‘X-caption’ pair data, including: 1) ‘Video-caption’ pair dataset: Webvid-2M (Bain et al., 2021), a large-scale dataset of short videos with textual description sourced from stock footage sites, 2) ‘Image-caption’ pair dataset: CC3M (Sharma et al., 2018), contains over 3 million images accompanied by diverse styles of natural-language descriptions, and 3) ‘Audio-caption’ pair dataset: AudioCaps (Kim et al., 2019), an extensive dataset of approximately 46k audio clips paired with human-written textual descriptions collected via crowdsourcing.

B.3 Fine-grained Cross-modal Grounding-aware Tuning

For this step, our focus is on achieving fine-grained, concept-level cross-modal alignment, primarily aligning the other three modalities to the textual modality. Our primary method involves utilizing existing ‘X-to-text’ phrase grounding datasets. REAMO’s input comprises textual phrases and grounded regional modality features, prompting the LLM to determine their match. For fine-grained image-to-text alignment, we consider the MS-COCO dataset (Lin et al., 2014). For video-to-text alignment, we turn to the TAO dataset (Dave et al., 2020). For entity-level audio-to-text

alignment, we primarily utilize the dataset for the Speech NER task (Chen et al., 2022a). Since there is a lack of dataset in English speech, we thus use TTS tools to generate the CoNLL 2003 textual NER data into speech NER data.

Invocation-based Meta-response Tuning. For this step, we aim to teach the LLM to produce the correct format of invocation meta-responses. Thus, we need to create data specifically designed for instruction tuning. Technically, we make use of the existing annotated datasets for various vision tasks included in this work. For each task under specific different user input scenarios, with the corresponding data, we design various template dialogue-format examples. Based on these examples we then prompt the GPT-4 to generate more samples under various topics and enriched scenarios. To ensure data diversification, the resulting instruction data includes a rich number of instances (10k), under different simulations of modalities and tasks.

C Grounded MUIE Evaluation

Our grounded MUIE evaluation dataset involves predictions for three tasks, including UIE label prediction, multimodal grounding prediction, and cognitive QA task prediction. Here, we provide detailed evaluation metrics for these three subtasks.

C.1 UIE Evaluation Metrics

To evaluate textual UIE results of the model, we use span-based offset Micro-F1 as the primary metric.

- For NER task, we follow a span-level evaluation setting, where the entity boundary and entity type must be correctly predicted.
- For RE task, a relation triple is correct if the model correctly predicts the boundaries of the subject entity, the object entity, and the entity relation.
- For EE task, we report two evaluation metrics:
 - Event Trigger (ET): an event trigger is correct if the event type and the trigger word are correctly predicted.
 - Event Argument (EA): an event argument is correct if its role type and event type match a reference argument mention.

C.2 Modality Grounding Evaluation Metrics

For the evaluation of the fine-grained modality grounding accuracy, the key idea is to measure the mean Intersection over Union (mIoU).

Image Segmentation. Let us denote by $\hat{M}_{img} = \{M_g\}_{g=1}^G$ the ground truth set of G regions, and $M_{img} = \{M_k\}_{k=1}^K$ the set of K predictions. Inspired by prior work, if $K \neq G$, we employ padding with \emptyset to equalize the sizes of both sets, resulting in a final size of $P = \max(G, K)$. Then, we find a bipartite matching between these two sets by searching for a permutation of P elements, $\sigma \in \mathcal{S}_P$, with the lowest cost:

$$\hat{\sigma} = \operatorname{argmin}_{\sigma \in \mathcal{S}_P} \sum_i^P \mathcal{L}_{match}(\hat{M}_i, M_{\sigma(i)}), \quad (1)$$

where $\mathcal{L}_{match}(\hat{M}_i, M_{\sigma(i)})$ is a pairwise matching cost between ground truth M_i and a prediction with index $\sigma(i)$. We compute this optimal assignment efficiently with the Hungarian algorithm. We define $\mathcal{L}_{match}(\hat{M}_i, M_{\sigma(i)})$ as $\mathcal{L}_{bce}(M_i, M_{\sigma(i)}) + \mathcal{L}_{dice}(\hat{M}_i, M_{\sigma(i)})$. The final IoU of each prediction is:

$$\text{IoU} = \frac{\text{Area of Overlap}}{\text{Area of Union}} \quad (2)$$

Based on the IoU scores, we can calculate mIoU metric by referring image segmentation dataset.

Video Tracking. For videos, we compute the Jaccard Index (a.k.a, mIoU score) for each frame via the above calculations, and then average them.

Audio Segmentation. Similarly, the mIoU score for each audio segment is computed to evaluate the quality of speech segmentation results. We measure the 1D span of the extracted segments and the 1D span of gold segments.