SAMA: Towards Multi-Turn Referential Grounded Video Chat with Large Language Models

Ye Sun¹, Hao Zhang², Henghui Ding¹, Tiehua Zhang³, Xingjun Ma¹, Yu-Gang Jiang¹

¹Fudan University, ²HKUST, ³Tongji University https://github.com/sunye23/SAMA

Abstract

Achieving fine-grained spatio-temporal understanding in videos remains a major challenge for current Video Large Multimodal Models (Video LMMs). Addressing this challenge requires mastering two core capabilities: video referring understanding, which captures the semantics of video regions, and video grounding, which segments object regions based on natural language descriptions. However, most existing approaches tackle these tasks in isolation, limiting progress toward unified, referentially grounded video interaction. We identify a key bottleneck in the lack of high-quality, unified video instruction data and a comprehensive benchmark for evaluating referentially grounded video chat. To address these challenges, we contribute in three core aspects: dataset, model, and benchmark. First, we introduce SAMA-239K, a large-scale dataset comprising 15K videos specifically curated to enable joint learning of video referring understanding, grounding, and multi-turn video chat. Second, we propose the SAMA model, which incorporates a versatile spatio-temporal context aggregator and a Segment Anything Model to jointly enhance fine-grained video comprehension and precise grounding capabilities. Finally, we establish SAMA-Bench, a meticulously designed benchmark consisting of 5,067 questions from 522 videos, to comprehensively evaluate the integrated capabilities of Video LMMs in multi-turn, spatio-temporal referring understanding and grounded dialogue. Extensive experiments and benchmarking results show that SAMA not only achieves strong performance on SAMA-Bench but also sets a new state-of-the-art on general grounding benchmarks, while maintaining highly competitive performance on standard visual understanding benchmarks.

1 Introduction

Recent years have witnessed remarkable progress in Large Multimodal Models (LMMs) [45, 56], driving significant advances in general vision-language understanding [36, 74, 30, 13]. However, extending these capabilities effectively to the video domain remains a critical and open research challenge [34, 68, 41, 31]. Unlike static images, videos introduce additional complexity through temporal dynamics, requiring models not only to perform holistic scene comprehension but also to achieve fine-grained alignment between language and temporally evolving objects and actions within a continuous spatio-temporal context [2]. While current Video LMMs [34, 68, 41, 31, 29, 32] demonstrate strong performance in global scene understanding, they often fall short in fine-grained spatio-temporal reasoning. This limitation is particularly evident when models are required to interpret references to specific entities and deliver precisely grounded information amid complex and dynamic activities—capabilities that are essential for applications such as detailed event forensics, interactive robotic assistance, and nuanced content analysis.

[†]Corresponding authors: henghui.ding@gmail.com, xingjunma@fudan.edu.cn

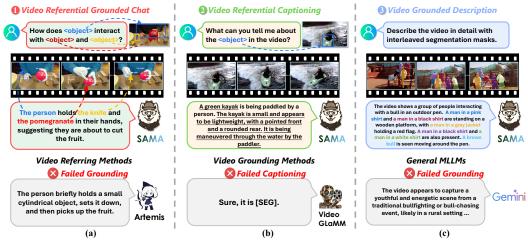


Figure 1: Comparisons with prior specialized and general MLLMs. Our SAMA model excels in multiple fine-grained video understanding tasks that require both referring and grounding capabilities: complex video referential grounded chat, video referential captioning, and video grounded description.

Enabling fine-grained video understanding hinges on mastering two fundamental and deeply interconnected capabilities: video referring understanding—the ability to comprehend semantics associated with designated video regions [47, 67, 19]—and **video grounding**—the capacity to accurately segment specific video regions based on natural language descriptions [12, 9, 66, 44, 1, 63]. In the image domain, substantial progress has been made with multimodal large language models (MLLMs) such as Ferret [64], Shikra [3], GPT4RoI [71], LLaVA-Grounding [69], and CogVLM [58], which integrate object-level representations into MLLMs to enable coordinate output or segmentation masks for fine-grained image comprehension. However, analogous advances in the video domain remain limited and face unique challenges. For video referring understanding, recent works like Artemis [47], VideoRefer [67], and OmniRGPT [19] demonstrate reasonable regional understanding by leveraging object prompts. Meanwhile, in video grounding, specialized LLM-based models such as VISA [60], VideoLISA [1], and VideoGLaMM [44] have been developed to segment target objects based on textual descriptions. Yet, a critical limitation persists: most of these efforts treat referring understanding and grounding as distinct tasks, neglecting their intrinsic interdependence—namely, that both fundamentally require precise spatio-temporal alignment with semantic cues. This separation prevents current Video LMMs from achieving a unified, human-like capability to both comprehend nuanced object details and fluidly interpret dynamic scenes, as shown in Figure 1.

We identify a key bottleneck hindering further progress: the lack of high-quality, unified video instruction data specifically tailored to enable the joint learning of referring understanding and grounding, along with the absence of a comprehensive benchmark for evaluating integrated performance [67]. While existing video understanding datasets typically focus on isolated tasks such as action recognition or captioning [21, 26], and foundational grounding datasets offer valuable segmentation or localization annotations for visual entities [9, 11, 52, 43, 73, 10, 14], a critical gap remains: the lack of resources that combine these components to support advanced referential grounded video dialogue. Specifically, current datasets fall short in providing: 1) rich, multi-turn referring-based video dialogues; 2) precise spatio-temporal grounding of objects and actions mentioned in those dialogues; and 3) explicitly structured joint supervision to facilitate the simultaneous learning of referring understanding and grounding within an interactive framework.

To address these challenges, we contribute in 3 key aspects: dataset, model, and benchmark.

• *Dataset*. We present SAMA-239K, a large-scale video instruction-tuning dataset specifically curated to facilitate the joint learning of video referring understanding and grounding. To construct this dataset, we developed a LLM-based automated annotation pipeline that first enriches video data by generating accurate segmentation masks using HQ-SAM [23]. Building upon these grounding annotations and leveraging advanced prompt engineering with Gemini-1.5 Pro [54], we generate 67K detailed object-level descriptions and 172K multi-turn conversational QA pairs to support training. Unlike previous single-turn video QA datasets, SAMA-239K is designed to emulate natural human

interaction, evolving from basic observations to complex reasoning and contextual inference, while explicitly linking conversational entities to their corresponding spatio-temporal masks.

- *Model*. We propose SAMA, a novel Video LMM architecture that seamlessly integrates fine-grained spatial and temporal understanding. SAMA introduces a versatile spatio-temporal aggregator to capture video-level temporal dynamics, alongside a powerful SAM2 [49] for extracting frame-level pixel-level semantics. This unified design enables the model to fully exploit the rich supervision in SAMA-239K, supporting end-to-end learning across video referring understanding, video grounding, and multi-turn referential video dialogue within a single cohesive framework.
- *Benchmark*. We introduce SAMA-Bench, a comprehensive benchmark comprising 522 diverse videos curated from four established datasets, totaling 5,067 challenging questions. It consists of two main components: SAMA-Bench^G, which evaluates referential grounded dialogue through a broad range of question types involving object behaviors, interactions, attributes, and spatial relationships; and SAMA-Bench^C, which assesses video referring region captioning by requiring models to generate detailed captions for specified object categories, including persons, animals, tools, and vehicles.
- *Evaluation*. Extensive benchmarking and experiments show that our SAMA model outperforms existing grounding LMMs on SAMA-Bench and also achieves strong results on established grounding benchmarks. To the best of our knowledge, SAMA is the first Video LMM to successfully unify fine-grained referential understanding and grounded dialogue, achieving state-of-the-art performance on both image and video referring segmentation tasks.

2 Related Work

Multimodal Large Language Models. The remarkable capabilities of LLMs [45, 56, 40] across a wide range of language-centric tasks have reshaped the landscape of artificial intelligence (AI), establishing a powerful foundation for extending reasoning abilities to visual understanding. In response, research has rapidly converged on the development of image-based MLLMs. Pioneering works such as LLaVA [36], MiniGPT-4 [74], and InstructBLIP [8] have demonstrated impressive visual understanding by effectively integrating visual inputs with LLMs. Building upon these successes in processing static images, attention has increasingly shifted toward Video LMMs [32, 34, 68, 41, 29]. To capture the rich spatio-temporal information in videos, Video LMMs typically extract sequential visual features from frames using pre-trained vision backbones. These features are often processed via direct token concatenation [5, 34], and then interleaved with textual embeddings as input to the LLM to generate responses. Despite promising results on tasks requiring holistic understanding—such as general video question answering and captioning—existing Video LMMs still face notable challenges in fine-grained, spatio-temporal object-level grounding and understanding [64, 3, 58, 69].

MLLMs for Referring and Grounding. Enabling fine-grained interaction in multimodal systems requires MLLMs to accurately refer to and ground specific entities within visual inputs. In the image domain, notable progress has been made. One major direction, represented by models such as Kosmos-2 [46], Shikra [3], Ferret [64], GPT4RoI [71], and CogVLM [58], equips MLLMs with the ability to process and generate spatial location information, typically in the form of bounding boxes or point coordinates expressed numerically or as discrete tokens. Another active direction integrates segmentation capabilities, enabling models to produce pixel-level masks aligned with textual references. This is often achieved by connecting LLMs with segmentation models like SAM [25] through specialized integration mechanisms [69, 27, 48, 55].

Extending referring and grounding capabilities to the video domain introduces new challenges, particularly due to temporal complexity and the need for consistent object tracking. Recent works like Artemis [47], VideoRefer [67], Omni-RGPT [19], and DAM [33] have advanced fine-grained video understanding through region-level descriptions conditioned on prompts. However, they do not support dynamic segmentation mask generation within conversational outputs. In contrast, VideoGLaMM [44] achieves strong referring segmentation in videos but focuses on grounded descriptions, lacking support for interactive, fine-grained video dialogue. This highlights a key gap: the lack of models that jointly support precise object-level grounding and nuanced referring in interactive video conversations. We bridge this gap through improved datasets, model designs, and benchmarks, achieving rich, grounded, and interactive video dialogue.

3 SAMA

3.1 SAMA-239K Data Creation

As illustrated in Figure 2, our data construction leverages a sophisticated LLM-based automated annotation pipeline, centered around Gemini-1.5 Pro [54], which systematically transforms existing video referring segmentation or grounding datasets into rich, grounded conversational data.

Object-level Description Generation. Our initial aim is to obtain precise, object-centric textual descriptions. For grounding datasets that only offer

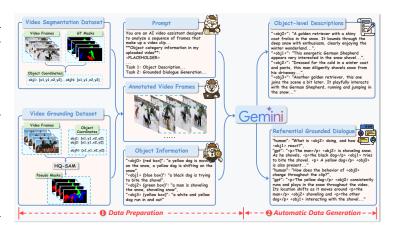


Figure 2: Data creation pipeline of SAMA-239K. Best viewed in zoom.

bounding boxes but lack pixel-level masks, we first employ HQ-SAM [23] to generate high-quality pseudo masks from the existing box information. Inspired by SoM [61], we then leverage segmentation mask information to highlight objects of interest within sampled video frames using distinctly colored bounding boxes. These visually cued frames, along with relevant categorical or action information for the highlighted objects, are provided to Gemini to generate fine-grained, descriptive captions for each identified object instance.

Referential Grounded Dialogue Generation. Building upon the rich, localized descriptions, we further task Gemini with creating multi-turn question-answer pairs. These dialogues are designed to emulate natural human conversational patterns, initiating with basic observational questions about one or more objects and progressively advancing to more complex reasoning that might involve inter-object relationships or inferential steps. A critical instruction is that the generated answers must explicitly incorporate grounding information, linking back to the previously established spatiotemporal masks of the referred entities.

Finally, through an automated filtering process to ensure the correctness of the generated data format, our proposed SAMA-239K has 172,296 referential grounded video-QA triplets along with 67,005 object-level descriptions in total.

3.2 SAMA-239K Data Characteristics

Spatio-Temporal Dynamics. For understanding complex object motion and achieving precise, temporally-consistent grounding, we select training data from referring video object segmentation datasets, primarily MeViS [9] and Ref-YouTube-VOS [52], which provide video sequences with dense annotations of objects undergoing significant motion changes.

Large Vocabulary Knowledge. To ensure SAMA comprehends a broad range of objects and their diverse textual descriptions, we incorporate LV-VIS [57], a video instance segmentation dataset that contributes a wide vocabulary of object categories, vital for open-world generalization.

Complex Scene Robustness. Real-world video understanding demands robustness against complex environmental conditions. To cultivate this, we incorporate data from SAV [49], which contributes challenging scenarios featuring occlusions and motion blur. Furthermore, we utilize video grounding dataset VidSTG [73] to ensure the model maintain consistent tracking and contextual comprehension over long temporal horizons.

3.3 SAMA Model

Here, we present the architecture of SAMA, which builds upon Sa2VA [66] and introduces novel contributions to synergize fine-grained frame-level information with long-range temporal context.

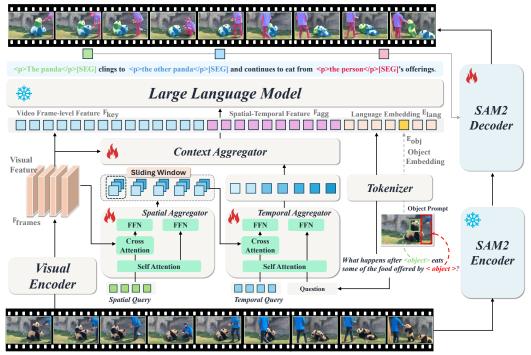


Figure 3: Model architecture of our SAMA for referential grounded video understanding.

3.3.1 Overall Architecture

The overall architecture of the SAMA model is illustrated in Figure 3. Conventional Video LMMs typically process video input via two main strategies: (1) concatenating tokenized features from sparsely sampled keyframes [5, 66, 34], or (2) aggressively pooling or merging features from individual frames into a very small number of tokens to accommodate longer sequences [32, 18, 50, 70]. We argue that both approaches have inherent limitations: the former often overlooks critical long-range temporal dynamics, while the latter tends to lose the fine-grained visual detail necessary for accurate referring segmentation and detailed object understanding.

SAMA is designed to overcome these limitations by effectively integrating both rich, frame-level visual information and robust, long-range temporal context. As depicted, a video input first passes through a Visual Encoder (e.g., a pre-trained ViT [5]) to extract per-frame visual features $\mathbf{F}_{\text{frames}} \in \mathbb{R}^{N \times P \times D_v}$, where N is the number of sampled frames, P is the number of patches per frame, and D_v is the feature dimension. A subset of these frame features, typically from keyframes, $\mathbf{F}_{\text{key}} \subset \mathbf{F}_{\text{frames}}$, are directly flattened and projected to serve as fine-grained visual tokens for the LLM. Concurrently, features from a more extensive set of frames $\mathbf{F}_{\text{long}} \subseteq \mathbf{F}_{\text{frames}}$ (potentially all N frames) are processed by our proposed Versatile Spatial-Temporal-Context Aggregator (detailed in Sec. 3.3.2) to derive a compact yet informative representation of long-range temporal dynamics, denoted as \mathbf{F}_{agg} .

The language input is processed by a tokenizer to produce language embeddings \mathbf{E}_{lang} . For referring understanding tasks involving specific objects, we incorporate object embeddings \mathbf{E}_{obj} into \mathbf{E}_{lang} . Specifically, \mathbf{E}_{obj} is obtained by performing mask pooling on the visual features \mathbf{F}_{key} of the relevant frame using a bounding box prompt for the target object. \mathbf{E}_{obj} is then inserted into the language embeddings \mathbf{E}_{lang} at the position corresponding to the object reference.

Finally, \mathbf{F}_{key} , \mathbf{F}_{agg} , and $[\mathbf{E}_{lang}, \mathbf{E}_{obj}]$ are concatenated and fed into a large language model Φ_{LLM} to generate the textual response. For grounding tasks, if the LLM outputs a special [SEG] token, the hidden state corresponding to this token \mathbf{H}_{SEG} is passed to the decoder of a pre-trained SAM2 model [49] to produce the final segmentation mask \mathbf{M}_{obj} for the referred object.

3.3.2 Versatile Spatial-Temporal-Context Aggregator

To efficiently capture both spatial details and temporal evolution from a potentially large number of video frames \mathbf{F}_{long} , we introduce a three-stage aggregator module.

Spatial Aggregator. Given N_L frames selected for long-range temporal modeling, where each frame i has visual features $\mathbf{F}_{\mathrm{long}}^{(i)} \in \mathbb{R}^{P \times D_v}$, the initial goal of the spatial aggregator is to condense the spatial information from each frame while retaining essential semantics, thereby reducing computational load for subsequent temporal modeling. Inspired by Q-Former [30], for each frame i, we employ a set of K_S learnable spatial queries $\mathbf{Q}_S \in \mathbb{R}^{K_S \times D_v}$. These queries interact with the frame's patch features $\mathbf{F}_{\mathrm{long}}^{(i)}$ through cross-attention mechanisms, followed by feed-forward networks (FFN), to produce a compact representation $\mathbf{Z}_S^{(i)} \in \mathbb{R}^{K_S \times D_v}$:

$$\mathbf{Z}_{S}^{(i)} = \operatorname{SpatialAgg}(\mathbf{Q}_{S}, \mathbf{F}_{\operatorname{long}}^{(i)}). \tag{1}$$

In our implementation, we use $K_S=32$ tokens per frame. This results in a sequence of spatially aggregated tokens $\mathbf{Z}_S=[\mathbf{Z}_S^{(1)},\mathbf{Z}_S^{(2)},...,\mathbf{Z}_S^{(N_L)}]\in\mathbb{R}^{(N_L\cdot K_S)\times D_v}$.

Temporal Aggregator. To capture comprehensive spatio-temporal information relevant to the input query, the sequence of spatially aggregated tokens \mathbf{Z}_S is processed by the temporal aggregator. We employ a sliding window approach: tokens from \mathbf{Z}_S corresponding to a window of W_T frames are iteratively fed into a temporal Q-Former structure. The temporal aggregator takes K_T learnable temporal queries $\mathbf{Q}_T \in \mathbb{R}^{K_T \times D_v}$ and interacts with the windowed spatial tokens $\mathbf{Z}_{S,\text{window}}$. Crucially, the textual question embedding $\mathbf{E}_{\text{question}}$ (and any associated object embeddings \mathbf{E}_{obj} if present within the question context) is also injected into this temporal Q-Former by concatenation with the temporal queries, which guides the aggregation process to focus on temporal dynamics relevant to the query:

$$\mathbf{Z}_{T}^{\text{window}} = \text{TemporalAgg}(\mathbf{Q}_{T}, \mathbf{Z}_{S, \text{window}}, \mathbf{E}_{\text{question}}, [\mathbf{E}_{\text{obj}}]).$$
 (2)

The outputs from each sliding window are then concatenated to yield a query-aware spatio-temporal representation $\mathbf{Z}_T \in \mathbb{R}^{K_{final} \times D_v}$, where K_{final} is the resulting number of temporal tokens.

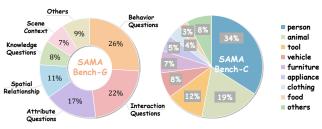
Context Aggregator. Finally, to ensure that the spatio-temporal information \mathbf{Z}_T is effectively aligned and integrated with the original fine-grained visual features \mathbf{F}_{long} , we employ a context aggregator. This module uses the frame features \mathbf{F}_{long} as queries, and the learned query-aware temporal features \mathbf{Z}_T as keys and values in an attention mechanism, allowing the model to selectively enhance the frame representations with relevant long-range temporal context. Subsequently, we apply adaptive average pooling to the spatio-temporally enhanced frame features, compressing each frame into a single token. These tokens are then projected through a linear layer to align with the language embedding space of the LLM. The final aggregated context feature \mathbf{F}_{agg} is obtained through the following process:

$$\mathbf{F}_{\text{agg}} = \text{Mean}(\text{Softmax}((\mathbf{F}_{\text{long}}\mathbf{W}^{Q}) \times (\mathbf{Z}_{T}\mathbf{W}^{K})^{T}/\sqrt{C}) \times \mathbf{Z}_{T}\mathbf{W}^{V})\mathbf{W}^{P}, \mathbf{F}_{\text{agg}} \in \mathbb{R}^{N \times D_{v}}, \quad (3)$$

where \mathbf{W}^Q , \mathbf{W}^K , \mathbf{W}^V , \mathbf{W}^P denote the linear projection layer. \mathbf{F}_{agg} , combined with the fine-grained keyframe features and language features are fed into the LLM to generate the target response.

3.4 SAMA-Bench

To evaluate the integrated referential understanding and grounded dialogue capabilities of Video LMMs, we propose SAMA-Bench, constructed using the same annotation pipeline as SAMA-239K. SAMA-Bench comprises 5,067 questions synthesized from 522 videos across four public validation datasets: MeViS [9], Ref-YouTube-VOS [52], LVVIS [57], and VidSTG [73].



(a) Question types in Bench-G (b) Category list in Bench-C

Figure 4: Data characteristics of SAMA-Bench.

It includes two sub-benchmarks: SAMA-Bench^c (3,036 questions) and SAMA-Bench^c (2,031 questions). The distributions of question types and object categories are shown in Figure 4.

Table 1: Performance comparisons on SAMA-Bench. The best results are **boldfaced**, and second-best results are <u>underlined</u>. "-" denotes that the model does not support the specified output format. Entries in gray represent that the original model is incapable of performing the task. Values in red show SAMA's performance change relative to corresponding Sa2VA variants.

Method			SAMA-Bench ^G				SAMA-Bench ^C	
	mIoU	Recall	METEOR	CIDEr	CLAIR	METEOR	CIDEr	CLAIR
			Generalis	st Models				
InternVL2.5-26B [5]	-	_	0.14	0.33	0.47	0.09	0.07	0.31
Gemini-2.0 Flash [53]	_	_	0.11	0.24	0.53	0.11	0.16	0.52
Gemini-1.5 Pro [53]	-	_	0.15	0.48	0.62	0.13	0.27	0.56
			Specialis	t Models				
Image-level models								
GLaMM [48] + SAM2 [49]	0.28	0.04	0.04	0.03	0.16	0.04	0.02	0.33
Shikra [3] + SAM2 [49]	0.27	0.26	0.08	0.15	0.32	0.04	0.01	0.29
Ferret-7B [64] + SAM2 [49]	0.64	0.44	0.14	0.21	0.37	0.10	0.12	0.31
Ferret-13B [64] + SAM2 [49]	0.64	0.43	0.14	0.20	0.39	0.11	0.10	0.31
Video-level models								
Sa2VA-1B [66]	0.09	0.07	0.10	0.16	0.31	0.06	0.03	0.26
Sa2VA-4B [66]	0.55	0.25	0.05	0.02	0.19	0.00	0.00	0.07
Sa2VA-8B [66]	0.64	0.17	0.02	0.02	0.20	0.00	0.00	0.13
SAMA-1B	0.67 (0.58 [†])	0.53 (0.46†)	0.16 (0.06†)	0.56 (0.40†)	0.53 (0.22 [†])	0.14 (0.087)	<u>0.31</u> (0.28 ↑)	0.45 (0.19
SAMA-4B	0.69 (0.14 [†])	0.55 (0.30†)	0.17 (0.12 [†])	0.65 (0.63 [†])	0.57 (0.38 [†])	<u>0.13</u> (0.13 ↑)	0.30 (0.30†)	0.48 (0.411
SAMA-8B	0.70 (0.06†)	0.55 (0.38†)	0.17 (0.15†)	0.69 (0.671)	$0.58 (0.38\uparrow)$	0.13 (0.13↑)	0.32 (0.327)	0.50 (0.37

SAMA-Bench^G This sub-benchmark evaluates a model's ability to conduct fine-grained, referential, and grounded visual dialogue. It focuses on region-level understanding and responses that require precise spatio-temporal grounding. As illustrated in Figure 4 (a), question types include **object behaviors**, **interactions**, **attributes**, **spatial relationships**, and more, requiring identification and reasoning over grounded entities. Following GLaMM [48, 44], we assess grounding performance using mIoU and Recall, and dialogue quality using METEOR, CIDEr, and CLAIR.

SAMA-Bench^C This sub-benchmark focuses on generating accurate descriptions for specified spatio-temporal video regions. As shown in Figure 4 (b), the benchmark spans diverse object categories, including **person**, **animal**, **tool**, **vehicle**, **furniture**, **appliance**, **clothing**, and **food**. Given a first-frame prompt, models are required to generate detailed captions for a target region. Description quality is evaluated using METEOR, CIDEr, and CLAIR.

4 Experiments

4.1 Experimental Setup

Training Datasets. SAMA is trained on a diverse set of image/video QA and referring segmentation/grounding datasets, including LLaVA-1.5 665K [37], the ChatUniVi [20] dataset, image-based referring/grounding data (refCOCO/+/g [22, 42], GRand-F [48]), and video referring segmentation datasets (Ref-YouTube-VOS [52], MeViS [9], and ReVOS [60]). Critically, training is enhanced with our proposed SAMA-239K dataset, comprising 239K instances of referential grounded dialogue and object-level descriptions.

Implementation Details. We implement SAMA leveraging the XTuner [7] codebase. Model training and inference is conducted on 8 NVIDIA A100 GPUs (80GB). During the instruction tuning phase, we make the parameters of our spatial-temporal-context aggregator and the decoder of the SAM2 model [49] trainable to learn spatial-temporal information and inject referential grounded video chat capability into the base model. The model weights are initialized using the pre-trained Sa2VA [66] to benefit from its existing grounding capabilities. The initial learning rate is set to 4e-5. The maximum sequence length for the LLM is configured to 8,192.

Evaluation Metrics. For SAMA-Bench, grounding performance on SAMA-Bench^G is measured using mIoU and Recall. Text generation quality on both SAMA-Bench^G and SAMA-Bench^C is evaluated using METEOR, CIDEr, and CLAIR. For standard image/video referring segmentation tasks, we adopt widely used metrics such as cIoU and J&F. Chat performance is assessed with VLMEvalKit [15].

4.2 Main Results

Performance Comparisons on SAMA-Bench. Table 1 presents a comprehensive performance comparison on SAMA-Bench, encompassing both the SAMA-Bench^G and SAMA-Bench^C subbenchmarks. To ensure fair evaluation across diverse model types, tailored protocols were applied. For image-level specialist models (e.g., GLaMM [48], Shikra [3], Ferret [64]) that accept coordinate

Table 2: Performance comparisons on referring segmentation in images and videos. **Bold** and <u>underlined</u> values indicate the best and second-best results, respectively. Red highlights SAMA's performance difference from corresponding Sa2VA variants.

Method		Image Seg	mentation		Video Segmentation			
	RefCOCO [22]	RefCOCO+ [22]	RefCOCOg [65]	GCG [48]	MeViS [9]	Ref-DAVIS17 [24]	Ref-YTVOS [52]	ReVOS [60]
Image-level models								
LISA-7B [27]	74.1	62.4	66.4	_	_	_	_	_
PixelLM-7B [51]	73.0	66.3	69.3	_	_	_	_	_
GLaMM-7B [48]	79.5	72.6	74.2	28.9	_	_	_	_
LLaVA-G-7B [69]	77.1	68.8	71.5	_	_	_	_	_
GSVA-13B [59]	79.2	70.3	75.7	_	_	_	_	_
OMG-LLaVA-7B [72]	78.0	69.1	72.9	29.9	_	_	_	_
Video-level models								
VideoGLaMM [44]	_	_	_	_	45.15	69.5	_	_
VISA-13B [60]	72.4	59.8	65.5	_	44.5	70.4	63.0	50.9
VideoLISA-3.8B [1]	73.8	63.4	68.3	_	44.4	68.8	63.7	47.5
Sa2VA-4B [66]	82.4	77.6	79.7	31.0	46.4	73.7	71.3	54.1
Sa2VA-8B [66]	82.6	78.0	80.3	32.2	51.5	75.9	72.3	57.6
SAMA-4B	82.5 (0.1 †)	77.9 (0.3 ↑)	80.3 (0.6 [†])	32.6 (1.6†)	48.3 (1.9†)	74.1 (0.4 ↑)	71.5 (0.2 †)	<u>58.8</u> (4.7 ↑)
SAMA-8B	82.7 (0.1 ↑)	78.1 (0.1 ↑)	80.6 (0.3 ↑)	31.7 (0.5↓)	53.7 (2.2↑)	77.3 (1.4†)	72.8 (0.5\(\dagger)\)	59.1 (1.5 ↑)

Table 3: Performance comparisons on chat benchmarks in images and videos. **Bold** and <u>underlined</u> values indicate the best and second-best results, respectively.

Method			Video Chat				
	MME [16] MMBench [38] SEED-Bench [28] AI2D [28] MMStar [4]					SQA ^{test} [39]	Video-MME [17]
		Ge	eneralist MLLM	5			
LLAVA-1.5-13B [35]	1531	68.8	70.1	-	-	-	-
Video-LLaVA-7B [34]	-	60.9	-	-	-	-	39.9
LLaMA-VID-7B [32]	1521	65.1	59.9	-	-	-	-
mPLUG-Owl3-8B [62]	-	77.6	-	-	-	-	53.5
InternVL2-8B [6]	-	81.7	<u>76.2</u>	-	-	-	54.0
		MLLMs wit	h segmentation	capability			
Image-level models							
PixelLM-7B [51]	309/135	17.4	-	-	-	-	-
GLaMM-7B [48]	14/9	36.8	-	-	-	-	-
OMG-LLaVA-7B [72]	1177/235	47.9	56.5	-	-	-	-
Video-level models							
Sa2VA-8B [66]	1690 / <u>610</u>	84.4	76.5	82.7	62.4	97.4	54.3
SAMA-8B	<u>1639</u> / 621	80.8	<u>76.2</u>	<u>79.4</u>	60.1	<u>95.3</u>	51.8

inputs but not direct video, we adopted a multi-step evaluation. First, these models were queried with questions containing target object coordinates from the initial video frame. Their coordinate-inclusive textual responses were then processed by Gemini to validate object correspondence. Finally, these validated coordinates prompted SAM2 [49] to generate segmentation masks for the video. For generalist and video models lacking explicit coordinate input (e.g., Gemini, InternVL, Sa2VA), visual cues like colored bounding boxes were overlaid on frames for querying their referential capabilities. The results in Table 1 highlight two key findings: (1) SAMA consistently outperforms all baselines on both sub-benchmarks. Notably, SAMA-8B achieves an mIoU of 0.70 and a Recall of 0.55, significantly surpassing Sa2VA, Shikra, and GLaMM, and exceeding the strong Ferret-13B + SAM2 baseline by 6% mIoU and 12% Recall. (2) SAMA also demonstrates strong text generation capabilities. It achieves the highest METEOR and CIDEr scores among all models, even outperforming Gemini-1.5 Pro, and significantly surpassing Ferret. These results underscore SAMA's superior spatio-temporal reasoning and fine-grained grounding abilities.

Performance Comparisons on Referring Segmentation in Images and Videos. The results in Table 2 demonstrate SAMA's robust and consistent performance across both image and video referring segmentation tasks. (1) SAMA achieves superior image segmentation performance. Specifically, SAMA-8B obtains top scores on RefCOCO (82.7), RefCOCO+ (78.1), and RefCOCOg (80.6), outperforming all previous models, including GLaMM-7B and VISA-13B. Furthermore, on the GCG benchmark, SAMA-4B reaches 32.6, surpassing its Sa2VA-4B variant by a notable margin of +1.6. (2) SAMA sets new state-of-the-art results in video segmentation. On video benchmarks, SAMA-8B achieves 53.7 on MeViS, 77.3 on Ref-DAVIS17, and 59.1 on ReVOS—outperforming Sa2VA-8B by 2.2, 1.4, and 1.5 points, respectively. These consistent improvements highlight SAMA's enhanced grounding ability.

Table 4: Ablation studies on spatial-temporal-context aggregator and SAMA-239K dataset.

ethod _			SAMA-Bench ^G			Video Segmentation			
- Luiou	mIoU	Recall	METEOR	CIDEr	CLAIR	MeViS (val_u) [9]	Ref-DAVIS [24]	ReVOS [60	
MA-4B	0.69	0.55	0.17	0.65	0.57	55.4	74.0	58.8	
/o STC	0.66 (0.031)	0.53 (0.021)	0.15 (0.021)	0.63 (0.021)	0.50 (0.071)	57.9 (2.5 [†])	74.5 (0.5 [†])	58.3 (0.3↓)	
/o SAMA-239K	0.46 (0.23 \$\\$)	0.10 (0.45 \$\dagger\$)	0.04 (0.13↓)	0.02 (0.63 \$\\$)	0.22 (0.35 ↓)	56.0 (0.6 ↑)	74.2(0.2 ↑)	55.6 (3.2 ↓)	
ject Prompts	Prediction	on results fr	rom SAMA						
How does	the moveme	ent of <object< td=""><td>t> change over</td><td>r time? Is ther</td><td>e anything els</td><td>se in the scene besi</td><td>ides <object> and</object></td><td><object>?</object></td></object<>	t> change over	r time? Is ther	e anything els	se in the scene besi	ides <object> and</object>	<object>?</object>	
SAMA		ent of <object< td=""><td>t> change over</td><td></td><td></td><td></td><td>ides <object> and</object></td><td></td></object<>	t> change over				ides <object> and</object>		
70 - 60 - 50 -			The movement mid-air, perfor	ct Segmentati of [484, 427, 64 ming tricks off a	on and Associant, 472] is interest concrete wall.		mporal Reasoning ws a skateboarder in the scene are [482,		
70 - 60 -			The movement mid-air, perfor	ct Segmentati of [484, 427, 64 ming tricks off a 2, 456] and [484	on and Associant, 472] is interest concrete wall. 1, 427, 641, 472].	ation S Failed Tentingly dynamic. It show	mporal Reasoning ws a skateboarder in the scene are [482, ns present.		

Figure 5: Visual comparisons between our SAMA and Ferret. Best viewed with zoom.

Performance Comparisons on Chat Benchmarks in Images and Videos To evaluate SAMA's general conversational ability beyond SAMA-Bench, we test it on standard image and video chat benchmarks. As shown in Table 3, SAMA performs robustly overall, with only a slight drop in some general chat metrics compared to Sa2VA. These results warrant contextual interpretation. First, SAMA's primary strength lies in unifying fine-grained referring, grounding, and dialogue within videos—capabilities that most grounding-focused models lack. For instance, compared to image-level grounding models like GLaMM-7B, SAMA-8B achieves significantly higher scores on shared benchmarks such as MME (1665/593 vs. 14/9), reflecting its superior dialogue ability. Second, although SAMA is mainly optimized for grounding, it still performs competitively on general conversation tasks, achieving 80.8 on MMBench, 76.2 on SEED-Bench, and 51.8 on Video-MME, which is comparable to generalist video LMMs not tailored for grounding.

4.3 Ablation Studies

Here, we conduct ablation studies from both model and data perspectives to thoroughly evaluate the effectiveness of the SAMA architecture and the SAMA-239K dataset. As shown in Table 4, removing the spatial-temporal-context aggregator (w/o STC) from SAMA-4B leads to performance drops across all metrics on SAMA-Bench^G, with mIoU decreasing from 0.69 to 0.66 and CLAIR from 0.57 to 0.50. More critically, training SAMA-4B without the SAMA-239K dataset (w/o SAMA-239K) results in a near-total collapse of grounded dialogue capabilities, with Recall plummeting to 0.10 and CIDEr to 0.02. This highlights the dataset's essential role in supporting robust referring understanding and grounded chat. Figure 5 further illustrates that SAMA, trained on SAMA-239K, exhibits substantially stronger referential understanding and temporally grounded dialogue than Ferret. However, neither the STC module nor the SAMA-239K dataset consistently improves performance on MeViS (val_u) and Ref-DAVIS. We recognize that building universally effective components and datasets for all downstream tasks is inherently challenging and plan to address these limitations in future work.

5 Conclusion

In this work, we introduced a comprehensive solution to equip Video Large Multimodal Models with fine-grained referential understanding and grounded dialogue capabilities. Our contributions are threefold: 1) SAMA-239K, a large-scale video instruction dataset designed to unify referential understanding and grounding; 2) SAMA Model, a novel Video LMM architecture that supports detailed video comprehension and segmentation; and 3) SAMA-Bench, a curated benchmark for evaluating integrated referential understanding and grounded dialogue in videos. Extensive experiments and analysis validate the effectiveness of our dataset, model, and benchmark. We hope our work provides

a useful step toward advancing fine-grained grounding and interactive understanding in video-based multimodal systems.

Limitations and Future Work. SAMA has several limitations that offer opportunities for future work. Currently, it only handles box-based prompts, and expanding to other input formats like points or masks could improve its flexibility. While strong at visual grounding, its performance on general dialogue tasks still lags behind specialized conversational models, suggesting room for better reasoning integration. Like prior work, we observe that joint training on grounding and dialogue can cause interference, pointing to a need for more robust multi-task learning approaches. Finally, SAMA-239K is currently limited to short clips; scaling to longer videos would better support research on complex, long-range temporal understanding.

Acknowledgements

This work is in part supported by National Key R&D Program of China (Grant No. 2022ZD0160103), National Natural Science Foundation of China (Grant No. 62276067), and National Natural Science Foundation of China (Grant No. 62472104).

The computations in this research were performed using the CFFF platform of Fudan University.

References

- [1] Zechen Bai, Tong He, Haiyang Mei, Pichao Wang, Ziteng Gao, Joya Chen, Lei Liu, Zheng Zhang, and Mike Zheng Shou. One token to seg them all: Language instructed reasoning segmentation in videos. In *NeurIPS*, 2024.
- [2] Shyamal Buch, Cristóbal Eyzaguirre, Adrien Gaidon, Jiajun Wu, Li Fei-Fei, and Juan Carlos Niebles. Revisiting the" video" in video-language understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2917–2927, 2022.
- [3] Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing multimodal llm's referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023.
- [4] Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, et al. Are we on the right way for evaluating large vision-language models? *arXiv* preprint arXiv:2403.20330, 2024.
- [5] Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv* preprint arXiv:2412.05271, 2024.
- [6] Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, et al. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *arXiv preprint arXiv:2404.16821*, 2024.
- [7] XTuner Contributors. Xtuner: A toolkit for efficiently fine-tuning llm. https://github.com/InternLM/xtuner, 2023.
- [8] 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.
- [9] Henghui Ding, Chang Liu, Shuting He, Xudong Jiang, and Chen Change Loy. MeViS: A large-scale benchmark for video segmentation with motion expressions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2694–2703, 2023.
- [10] Henghui Ding, Chang Liu, Shuting He, Xudong Jiang, Philip HS Torr, and Song Bai. MOSE: A new dataset for video object segmentation in complex scenes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 20224–20234, 2023.
- [11] Henghui Ding, Chang Liu, Shuting He, Kaining Ying, Xudong Jiang, Chen Change Loy, and Yu-Gang Jiang. MeViS: A multi-modal dataset for referring motion expression video segmentation. *TPAMI*, 2025.
- [12] Henghui Ding, Chang Liu, Suchen Wang, and Xudong Jiang. Vision-language transformer and query generation for referring segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16321–16330, 2021.
- [13] Henghui Ding, Song Tang, Shuting He, Chang Liu, Zuxuan Wu, and Yu-Gang Jiang. Multimodal referring segmentation: A survey. *arXiv preprint arXiv:2508.00265*, 2025.

- [14] Henghui Ding, Kaining Ying, Chang Liu, Shuting He, Xudong Jiang, Yu-Gang Jiang, Philip HS Torr, and Song Bai. MOSEv2: A more challenging dataset for video object segmentation in complex scenes. *arXiv* preprint arXiv:2508.05630, 2025.
- [15] Haodong Duan, Junming Yang, Yuxuan Qiao, Xinyu Fang, Lin Chen, Yuan Liu, Xiaoyi Dong, Yuhang Zang, Pan Zhang, Jiaqi Wang, et al. Vlmevalkit: An open-source toolkit for evaluating large multi-modality models. In ACMMM, 2024.
- [16] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, Yunsheng Wu, and Rongrong Ji. Mme: A comprehensive evaluation benchmark for multimodal large language models. arXiv preprint arXiv:2306.13394, 2023.
- [17] Chaoyou Fu, Yuhan Dai, Yondong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, et al. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. arXiv preprint arXiv:2405.21075, 2024.
- [18] Bo He, Hengduo Li, Young Kyun Jang, Menglin Jia, Xuefei Cao, Ashish Shah, Abhinav Shrivastava, and Ser-Nam Lim. Ma-Imm: Memory-augmented large multimodal model for long-term video understanding. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 13504– 13514, 2024.
- [19] Miran Heo, Min-Hung Chen, De-An Huang, Sifei Liu, Subhashree Radhakrishnan, Seon Joo Kim, Yu-Chiang Frank Wang, and Ryo Hachiuma. Omni-rgpt: Unifying image and video region-level understanding via token marks. *arXiv preprint arXiv:2501.08326*, 2025.
- [20] Peng Jin, Ryuichi Takanobu, Wancai Zhang, Xiaochun Cao, and Li Yuan. Chat-univi: Unified visual representation empowers large language models with image and video understanding. In CVPR, 2024.
- [21] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. arXiv preprint arXiv:1705.06950, 2017.
- [22] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *EMNLP*, 2014.
- [23] Lei Ke, Mingqiao Ye, Martin Danelljan, Yu-Wing Tai, Chi-Keung Tang, Fisher Yu, et al. Segment anything in high quality. Advances in Neural Information Processing Systems, 36:29914–29934, 2023.
- [24] Anna Khoreva, Anna Rohrbach, and Bernt Schiele. Video object segmentation with language referring expressions. In *ACCV*, 2018.
- [25] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *ICCV*, 2023.
- [26] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In Proceedings of the IEEE international conference on computer vision, pages 706–715, 2017.
- [27] Xin Lai, Zhuotao Tian, Yukang Chen, Yanwei Li, Yuhui Yuan, Shu Liu, and Jiaya Jia. Lisa: Reasoning segmentation via large language model. In *CVPR*, 2024.
- [28] Bohao Li, Yuying Ge, Yixiao Ge, Guangzhi Wang, Rui Wang, Ruimao Zhang, and Ying Shan. Seed-bench: Benchmarking multimodal large language models. In *CVPR*, 2024.
- [29] Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, and Chunyuan Li. Llava-next-interleave: Tackling multi-image, video, and 3d in large multimodal models. *arXiv* preprint *arXiv*:2407.07895, 2024.
- [30] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, 2023.
- [31] Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, et al. Mvbench: A comprehensive multi-modal video understanding benchmark. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 22195–22206, 2024.
- [32] Yanwei Li, Chengyao Wang, and Jiaya Jia. Llama-vid: An image is worth 2 tokens in large language models. In European Conference on Computer Vision, pages 323–340. Springer, 2024.
- [33] Long Lian, Yifan Ding, Yunhao Ge, Sifei Liu, Hanzi Mao, Boyi Li, Marco Pavone, Ming-Yu Liu, Trevor Darrell, Adam Yala, et al. Describe anything: Detailed localized image and video captioning. arXiv preprint arXiv:2504.16072, 2025.
- [34] Bin Lin, Yang Ye, Bin Zhu, Jiaxi Cui, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning united visual representation by alignment before projection. *arXiv* preprint arXiv:2311.10122, 2023.
- [35] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In CVPR, 2024.
- [36] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023.

- [37] Shilong Liu, Hao Cheng, Haotian Liu, Hao Zhang, Feng Li, Tianhe Ren, Xueyan Zou, Jianwei Yang, Hang Su, Jun Zhu, et al. Llava-plus: Learning to use tools for creating multimodal agents. In ECCV, 2024.
- [38] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? In ECCV, 2024.
- [39] Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *NeurIPS*, 2022.
- [40] Xingjun Ma, Yifeng Gao, Yixu Wang, Ruofan Wang, Xin Wang, Ye Sun, Yifan Ding, Hengyuan Xu, Yunhao Chen, Yunhao Zhao, et al. Safety at scale: A comprehensive survey of large model and agent safety. Foundations and Trends® in Privacy and Security, 8(3-4):254–469, 2025.
- [41] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt: Towards detailed video understanding via large vision and language models. arXiv preprint arXiv:2306.05424, 2023
- [42] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11–20, 2016.
- [43] Jiaxu Miao, Xiaohan Wang, Yu Wu, Wei Li, Xu Zhang, Yunchao Wei, and Yi Yang. Large-scale video panoptic segmentation in the wild: A benchmark. In *CVPR*, 2022.
- [44] Shehan Munasinghe, Hanan Gani, Wenqi Zhu, Jiale Cao, Eric Xing, Fahad Shahbaz Khan, and Salman Khan. Videoglamm: A large multimodal model for pixel-level visual grounding in videos. *arXiv preprint arXiv:2411.04923*, 2024.
- [45] OpenAI. GPT-4 technical report. https://arxiv.org/abs/2303.08774, 2023.
- [46] Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint arXiv:2306.14824*, 2023.
- [47] Jihao Qiu, Yuan Zhang, Xi Tang, Lingxi Xie, Tianren Ma, Pengyu Yan, David Doermann, Qixiang Ye, and Yunjie Tian. Artemis: Towards referential understanding in complex videos. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [48] Hanoona Rasheed, Muhammad Maaz, Sahal Shaji, Abdelrahman Shaker, Salman Khan, Hisham Cholakkal, Rao M Anwer, Eric Xing, Ming-Hsuan Yang, and Fahad S Khan. Glamm: Pixel grounding large multimodal model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13009–13018, 2024.
- [49] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024.
- [50] Shuhuai Ren, Linli Yao, Shicheng Li, Xu Sun, and Lu Hou. Timechat: A time-sensitive multimodal large language model for long video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14313–14323, 2024.
- [51] Zhongwei Ren, Zhicheng Huang, Yunchao Wei, Yao Zhao, Dongmei Fu, Jiashi Feng, and Xiaojie Jin. Pixellm: Pixel reasoning with large multimodal model. In *CVPR*, 2024.
- [52] Seonguk Seo, Joon-Young Lee, and Bohyung Han. URVOS: Unified referring video object segmentation network with a large-scale benchmark. In ECCV, 2020.
- [53] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. arXiv preprint arXiv:2312.11805, 2023.
- [54] Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.
- [55] Yunjie Tian, Tianren Ma, Lingxi Xie, and Qixiang Ye. Chatterbox: Multimodal referring and grounding with chain-of-questions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 7401–7409, 2025.
- [56] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971, 2023.
- [57] Haochen Wang, Cilin Yan, Shuai Wang, Xiaolong Jiang, Xu Tang, Yao Hu, Weidi Xie, and Efstratios Gavves. Towards open-vocabulary video instance segmentation. In proceedings of the IEEE/CVF international conference on computer vision, pages 4057–4066, 2023.

- [58] Weihan Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Song XiXuan, et al. Cogvlm: Visual expert for pretrained language models. Advances in Neural Information Processing Systems, 37:121475–121499, 2024.
- [59] Zhuofan Xia, Dongchen Han, Yizeng Han, Xuran Pan, Shiji Song, and Gao Huang. Gsva: Generalized segmentation via multimodal large language models. In CVPR, 2024.
- [60] Cilin Yan, Haochen Wang, Shilin Yan, Xiaolong Jiang, Yao Hu, Guoliang Kang, Weidi Xie, and Efstratios Gavves. Visa: Reasoning video object segmentation via large language models. In *European Conference* on Computer Vision, pages 98–115. Springer, 2024.
- [61] Jianwei Yang, Hao Zhang, Feng Li, Xueyan Zou, Chunyuan Li, and Jianfeng Gao. Set-of-mark prompting unleashes extraordinary visual grounding in gpt-4v. arXiv preprint arXiv:2310.11441, 2023.
- [62] Jiabo Ye, Haiyang Xu, Haowei Liu, Anwen Hu, Ming Yan, Qi Qian, Ji Zhang, Fei Huang, and Jingren Zhou. mPLUG-Owl3: Towards long image-sequence understanding in multi-modal large language models. *arXiv preprint*, 2024.
- [63] Kaining Ying, Henghui Ding, Guangquan Jie, and Yu-Gang Jiang. Towards omnimodal expressions and reasoning in referring audio-visual segmentation. *arXiv* preprint arXiv:2507.22886, 2025.
- [64] Haoxuan You, Haotian Zhang, Zhe Gan, Xianzhi Du, Bowen Zhang, Zirui Wang, Liangliang Cao, Shih-Fu Chang, and Yinfei Yang. Ferret: Refer and ground anything anywhere at any granularity. ICLR, 2024.
- [65] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In ECCV, 2016.
- [66] Haobo Yuan, Xiangtai Li, Tao Zhang, Zilong Huang, Shilin Xu, Shunping Ji, Yunhai Tong, Lu Qi, Jiashi Feng, and Ming-Hsuan Yang. Sa2va: Marrying sam2 with llava for dense grounded understanding of images and videos. arXiv preprint arXiv:2501.04001, 2025.
- [67] Yuqian Yuan, Hang Zhang, Wentong Li, Zesen Cheng, Boqiang Zhang, Long Li, Xin Li, Deli Zhao, Wenqiao Zhang, Yueting Zhuang, et al. Videorefer suite: Advancing spatial-temporal object understanding with video llm. arXiv preprint arXiv:2501.00599, 2024.
- [68] Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. arXiv preprint arXiv:2306.02858, 2023.
- [69] Hao Zhang, Hongyang Li, Feng Li, Tianhe Ren, Xueyan Zou, Shilong Liu, Shijia Huang, Jianfeng Gao, Chunyuan Li, Jainwei Yang, et al. Llava-grounding: Grounded visual chat with large multimodal models. In ECCV, 2024.
- [70] Shaolei Zhang, Qingkai Fang, Zhe Yang, and Yang Feng. Llava-mini: Efficient image and video large multimodal models with one vision token. arXiv preprint arXiv:2501.03895, 2025.
- [71] Shilong Zhang, Peize Sun, Shoufa Chen, Min Xiao, Wenqi Shao, Wenwei Zhang, Kai Chen, and Ping Luo. Gpt4roi: Instruction tuning large language model on region-of-interest. arXiv preprint arXiv:2307.03601, 2023.
- [72] Tao Zhang, Xiangtai Li, Hao Fei, Haobo Yuan, Shengqiong Wu, Shunping Ji, Change Loy Chen, and Shuicheng Yan. Omg-llava: Bridging image-level, object-level, pixel-level reasoning and understanding. In NeurIPS, 2024.
- [73] Zhu Zhang, Zhou Zhao, Yang Zhao, Qi Wang, Huasheng Liu, and Lianli Gao. Where does it exist: Spatio-temporal video grounding for multi-form sentences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10668–10677, 2020.
- [74] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. arXiv preprint arXiv:2304.10592, 2023.

A Appendix

A.1 Broader Impacts

Our SAMA imposes several positive broader impacts. 1) SAMA introduces a unified framework that enhances the capabilities of Video Large Multimodal Models (LMMs) by bridging fine-grained referential understanding and grounded, multi-turn dialogue-two areas that have traditionally been addressed separately. This integration may encourage a shift in video understanding research toward more context-aware, multimodal systems that unify perception and interaction. 2) SAMA-239K provides a high-quality dataset designed to support joint learning of video grounding and dialogue. It offers a valuable resource for advancing interactive video reasoning and object-centric video comprehension, and may also facilitate research in grounding alignment and conversational multi-object tracking. 3) The SAMA model serves as an early baseline for combining referring understanding with pixel-level grounding in video. It demonstrates that detailed visual grounding and referential dialogue can be effectively optimized within a single cohesive architecture, potentially guiding future designs for complex video-language tasks. 4) SAMA-Bench addresses a key evaluation gap by introducing the first benchmark tailored to jointly assess referential video understanding and grounded dialogue. We believe that the SAMA suite—dataset, model, and benchmark—can help steer the development of future video LMMs and foster progress in multimodal human-AI interaction, where spatial-temporal grounding and communication are tightly coupled.

Ethical Considerations. SAMA-239K is constructed entirely from publicly available research datasets, ensuring no private or personally identifiable data are included. The released data, model, and code will be provided under a research-only license that explicitly prohibits malicious or unethical uses, such as surveillance, disinformation, or privacy violations. We encourage responsible use of SAMA to advance transparent and beneficial research in multimodal video understanding.

A.2 Data Curation

Table 5: SAMA-239K for referential grounded video chat training.

Dataset	Video Clips	Q&A Pairs	Object-level Descriptions
MeViS [9]	1,087	10,257	3,550
Ref-Youtube-VOS [52]	1,976	15,622	4,929
LV-VIS [57]	2,595	30,943	14,268
SAV [49]	2,867	34,729	12,262
VidSTG [73]	6,618	80,745	31,996
SAMA-239K	15,143	172,296	67,005

SAMA-239K. The details of SAMA-239K are summarized in Table 5. To construct the SAMA-239K dataset, we curated and transformed several existing video understanding datasets into a unified referentially grounded video question-answering format. As illustrated in Figure 2, our automatic data generation pipeline integrates both video segmentation and video grounding datasets. For segmentation datasets with ground-truth masks, we extract object coordinates directly from the masks. For grounding datasets lacking masks, we employ HQ-SAM to generate pseudo masks. In both cases, we visualize each object by drawing a uniquely colored bounding box and assigning it a distinct tag. These annotated video frames, along with structured object information and tailored prompts, are fed into Gemini-1.5 Pro [54] to generate fine-grained object-level descriptions and multi-turn referential grounded dialogues. Below, we outline the detailed data curation process for the five primary sources contributing to our dataset: MeViS [9], Ref-YouTube-VOS [52], LV-VIS [57], SAV [49], and VidSTG [73]. MeViS and Ref-YouTube-VOS are large-scale referring video object segmentation datasets with rich pixel-level mask annotations aligned with textual expressions. For the SAV dataset, we adopt object-level referring expressions released by Sa2VA [66]. To ensure visual diversity and avoid trivial cases, we discard videos with only a single object. For the remaining videos, we uniformly sample 16 frames and use segmentation masks to identify referred objects. Each object is visualized with a uniquely colored bounding box, establishing a one-to-one mapping between object and color. These annotated frames, along with referring expressions and color-object mappings, are integrated into prompt templates and fed to Gemini, which then generates fine-grained descriptions and multi-turn dialogues that emulate natural human interaction—from basic observations to complex reasoning about object dynamics and relationships. LV-VIS [57] is a large-vocabulary video instance

segmentation dataset with pixel-level masks and category labels. For each object, its category label is included in the prompt to guide Gemini in generating category-aware descriptions. All other steps—including frame sampling, bounding box visualization, and prompt construction—follow the same procedure. VidSTG [73] is a video grounding dataset that provides bounding box annotations but lacks mask information. We uniformly sample each video at 4-frame intervals to extract JPEG frames. The frame-level bounding boxes are then used as input to HQ-SAM [23] to generate pseudo masks. The resulting data is subsequently processed following the same pipeline as the other datasets to produce training samples.

Table 6: SAMA-Bench^G for video referential grounded chat evaluation.

Dataset	Video Clips	Q&A Pairs
MeViS [9]	41	244
Ref-Youtube-VOS [52]	131	756
LV-VIS [57]	150	1,019
VidSTG [73]	200	1,019
SAMA-Bench ^G	522	3,038

Table 7: SAMA-Bench^C for video referential captioning evaluation.

Dataset	Video Clips	Q&A Pairs
MeViS [9]	41	117
Ref-Youtube-VOS [52]	131	350
LV-VIS [57]	150	589
VidSTG [73]	200	975
SAMA-Bench ^C	522	2,031

SAMA-Bench. SAMA-Bench is constructed using the same annotation pipeline as SAMA-239K, followed by rigorous automatic filtering and manual verification to remove low-quality or ambiguous samples. To ensure comprehensive evaluation across diverse video understanding scenarios, we randomly select a subset of validation videos from four datasets: MeViS, Ref-YouTube-VOS, LV-VIS, and VidSTG. As shown in Table 6, SAMA-Bench^G consists of 3,038 video referential grounded chat questions, with 244 from MeViS, 756 from Ref-YouTube-VOS, 1,019 from LV-VIS, and 1,019 from VidSTG. Similarly, Table 7 summarizes SAMA-Bench^C, which comprises 2,031 video referential captioning questions, sourced from the same set of videos: 117 from MeViS, 350 from Ref-YouTube-VOS, 589 from LV-VIS, and 975 from VidSTG.

Table 8: Model performance under different training data proportions.

Training Data	MME	SEED	AI2D	MMStar	SQAtest	MMBench	Video-MME
50% mix665k	1409	56.2	46.5	39.3	75.1	43.6	39.5
100% mix665k	1451	65.7	57.8	44.5	78.4	55.8	41.3

A.3 More Analysis

Improving Dialogue Performance. We investigate two approaches to improve SAMA's performance on general dialogue. We first validated the impact of expanding conversational training data. Keeping all other factors constant during training, we adjusted the proportion of the LLaVA-v1.5-mix665k dataset used in our instruction-tuning mixture. The results in Table 8 show a clear, positive correlation between the amount of dialogue data and SAMA's performance across various image chat benchmarks, confirming this is a direct and effective improvement strategy. Second, we hypothesize that a discrepancy in output granularity between general dialogue tasks and visual grounding can cause mutual interference during joint training and thus hurt dialogue performance. To probe this, we adjust the next-token prediction loss weight and observe its impact on chat results. As shown in Table 9, we found that increasing the loss weight to 1.5 led to improved performance on most chat benchmarks. This implies that an effective strategy is to re-balance the optimization objective to focus more on the holistic, dialogue-relevant output embeddings, rather than over-emphasizing grounding-specific tokens.

Table 9: Model performance under different loss weight configurations.

Loss Weight	Loss Weight SAMA-Bench ^G							Chat Ber	nchmark		
Loss Weight	mIoU	Recall	METEOR	CIDEr	MME	SEED	AI2D	MMStar	SQAtest	MMBench	Video-MME
1.00	67.1	51.7	14.9	49.7	1451	65.7	57.8	44.5	78.4	55.8	41.3
1.25	66.4	52.2	15.2	51.7	1424	61.4	52.5	42.8	76.4	49.7	39.2
1.50	66.9	51.5	15.0	52.0	1479	66.0	60.0	45.7	80.0	54.2	40.0

Table 10: Model performance under different prompt configurations.

Method		SAM	SAMA-Bench ^C				
Method	mIoU	Recall METEOR		CIDEr	METEOR	CIDEr	
Point (1)	0.60	0.44	0.14	0.46	0.11	0.20	
Point (2)	0.61	0.45	0.14	0.48	0.12	0.24	
Point (4)	0.62	0.46	0.14	0.49	0.13	0.29	
Point (8)	0.64	0.48	0.14	0.48	0.13	0.30	
Box	0.66	0.49	0.15	0.50	0.13	0.30	
Mask	0.65	0.50	0.14	0.46	0.13	0.31	

Prompt Sensitivity. To validate SAMA's flexibility, we conducted new experiments where we extended the instruction tuning to include a mixture of prompt formats. We then evaluated the model across a diverse set of inputs, including bounding boxes, masks, and a varying number of points (1, 2, 4, and 8). The results in Table 10 demonstrate that SAMA maintains strong and consistent performance across all formats.

Table 11: Model performance on SAMA Bench-G VidSTG subset.

Method	mIoU	Recall	METEOR	CIDEr
Ferret13B+SAM2	55.8	36.3	14.3	17.6
SAMA-1B	61.2	45.3	15.5	42.6

Long-Video Evaluation. To demonstrate SAMA's long-video capabilities, we evaluated it on the challenging VidSTG subset of SAMA Bench-G. As shown in Table 11, the results are twofold: first, SAMA-1B achieves satisfactory performance on this long-video-centric benchmark, confirming its ability to handle minute-level videos. Second, our model outperforms the powerful image-domain model, Ferret-13B+SAM2, underscoring the unique challenges of referring and grounding in the video domain.

A.4 Visualization

Figure 6 and Figure 7 illustrate the results of SAMA on the video referential grounded chat task.

Figure 8 illustrates the results of SAMA on the video grounded description task.

Figure 9 illustrates the results of SAMA on the video referential captioning task.

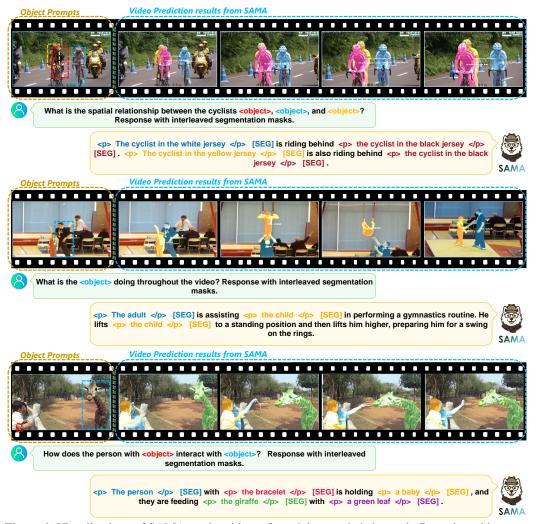


Figure 6: Visualizations of SAMA on the video referential grounded chat task. Best viewed in zoom.

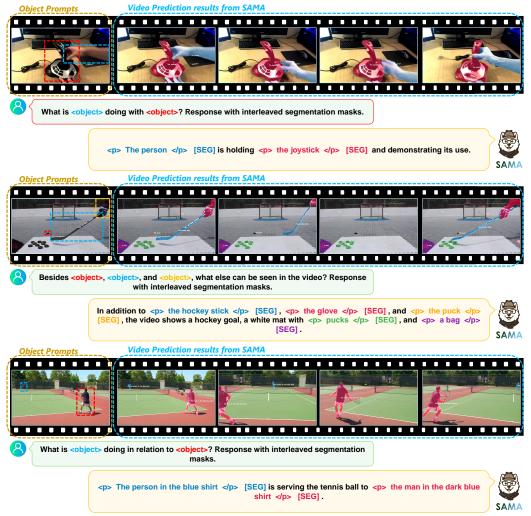


Figure 7: Visualizations of SAMA on the video referential grounded chat task. Best viewed in zoom.



Figure 8: Visualizations of SAMA on the video grounded description task. Best viewed in zoom.



Figure 9: Visualizations of SAMA on the video referential captioning task. Best viewed in zoom.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The contributions of this paper are clearly presented in the abstract and introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the
 contributions made in the paper and important assumptions and limitations. A No or
 NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals
 are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We note that our method still has limitations and room for improvement in our experiments.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This paper does not contain theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Information related to the experiment is detailed in the experiment setup section.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Source code will be released to foster the community.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how
 to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: All relevant settings are described in the experimental setup section.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: The experimental results in this paper require no further error analysis to be reported.

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We describe the training settings for each experiment in the experimental section.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: This research complied fully with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We explore some possible societal impacts of our method in the conclusion and appendix.

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This paper proposes a new task, data and model, so there is no safety risk.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with
 necessary safeguards to allow for controlled use of the model, for example by requiring
 that users adhere to usage guidelines or restrictions to access the model or implementing
 safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Relevant sources used in this paper are appropriately cited.

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: Our code, data and model will be released to foster the community.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The experiments in this paper does not include these contents.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The experiments in this article did not involve human subjects.

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [Yes]

Justification: The paper clearly describes how to use a large language model to generate our training data.

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.