

MAGE: Multimodal Alignment and Generation Enhancement via Bridging Visual and Semantic Spaces

Shaojun E^{1,3,*}, Yuchen Yang^{2,*}, Jiaheng Wu², Yan Zhang¹, Tiejun Zhao², Ziyang Chen^{1,†}

¹ Global Tone Communication Technology Co., Ltd., Beijing, China

² Faculty of computing, Harbin Institute of Technology, Harbin, China

³ School of Computer Science and Technology, Beijing Jiaotong University, Beijing, China

{eshaojun, zhangyan01, chenzyang}@gtcom.com.cn,
{23S003057, 23S136126}@stu.hit.edu.cn; tjzhao@hit.edu.cn

Abstract

In the latest advancements in multimodal learning, effectively addressing the spatial and semantic losses of visual data after encoding remains a critical challenge. This is because the performance of large multimodal models is positively correlated with the coupling between visual encoders and large language models. Existing approaches often face issues such as vector gaps or semantic disparities, resulting in information loss during the propagation process. To address these issues, we propose *MAGE* (Multimodal Alignment and Generation Enhancement), a novel framework that bridges the semantic spaces of vision and text through an innovative alignment mechanism. By introducing the Intelligent Alignment Network (IAN), *MAGE* achieves dimensional and semantic alignment. To reduce the gap between synonymous heterogeneous data, we employ a training strategy that combines cross-entropy and mean squared error, significantly enhancing the alignment effect. Moreover, to enhance *MAGE*’s “Any-to-Any” capability, we developed a fine-tuning dataset for multimodal tool-calling instructions to expand the model’s output capability boundaries. Finally, our proposed multimodal large model architecture, *MAGE*, achieved significantly better performance compared to similar works across various evaluation benchmarks, including MME, MMBench, and SEED. Complete code and appendix are available at: <https://github.com/GTCOM-NLP/MAGE>

1 Introduction

In recent years, with the remarkable progress of large language models (LLMs) in natural language processing tasks [Brown *et al.*, 2020; Vaswani, 2017], researchers have gradually extended their powerful capabilities to multimodal domains, enabling these models to simultaneously comprehend and generate both visual and textual information [Anthropic, 2024; Team *et al.*, 2023; OpenAI, 2024]. fig. 1 shows the performance of common projection mechanisms. However, achieving effective integration of visual encoders and LLMs

still faces significant challenges in cross-modal alignment, including semantic gaps and dimensional mismatches [Yu *et al.*, 2022; Yuan *et al.*, 2021]. These challenges directly impact the performance and efficiency of the models in multimodal tasks, limiting the practical applications of multimodal large language models (MLLMs).

Existing methods for aligning visual and linguistic modalities face several limitations [Dai *et al.*, 2024]. First, linear projectors preserve the integrity of visual features through simple mappings but struggle to model semantic relationships between modalities, often resulting in semantic information loss in complex tasks. Second, projector-based mechanisms, such as Resamplers [Alayrac *et al.*, 2022] and Q-Formers [Li *et al.*, 2023b], improve alignment through query mechanisms and cross-modal attention but overlook the global semantic information of visual features, leading to suboptimal performance in deep semantic tasks. Finally, recent approaches like TokenPacker [Li *et al.*, 2024b] and Honeybee [Cha *et al.*, 2024] have made progress by enhancing efficiency and preserving local features, yet they insufficiently address the role of loss functions in alignment. Most methods rely on single-generation or alignment losses, failing to fully capture deep semantic relationships between visual and linguistic modalities [Wang *et al.*, 2024; Dai *et al.*, 2023].

The limitations of these approaches indicate that existing research has yet to fully address the dual challenges of dimensional consistency and semantic integrity in the alignment of visual and linguistic modalities. Therefore, there is an urgent need for a solution that simultaneously optimizes both efficiency and semantic alignment.

To address the challenges of semantic and dimensional gaps between visual encoders and large language models (LLMs), while simultaneously meeting the requirements for generation and alignment in multimodal tasks, we propose a novel multimodal large model architecture, named *MAGE*, as illustrated in fig. 2. *MAGE* aims to achieve efficient and flexible cross-modal generation and processing through innovative designs in three aspects: semantic alignment, training strategy, and interaction framework.

The core architecture of *MAGE* comprises three innovative modules. First, we introduce the Intelligent Alignment Network module, which consists of two submodules: the Vector Alignment Module and the Semantic Enhancement Module.




Figure 1: (Left) an example of an attention map from the IAN. (middle) The attention matrix of a visual token obtained after passing through IAN in the left figure. and (right) A comparison of the spatial understanding capability between IAN and common projectors, where Avg^N represents the average performance computed from six common understanding tasks, including MME, MMB, and SEED.

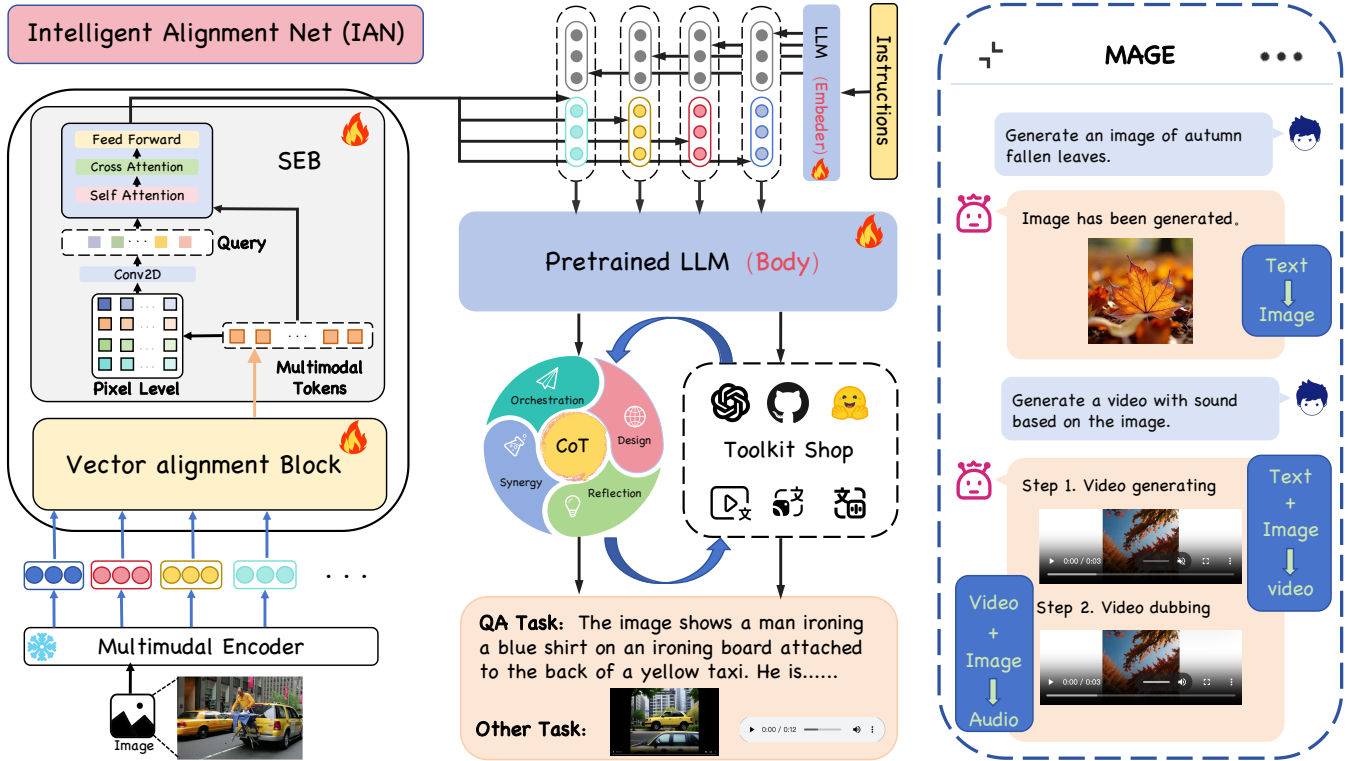


Figure 2: **The detailed aspects of MAGE.** IAN consists of two components: vector alignment and semantic enhancement. The encoded vectors encapsulate rich semantic knowledge, enabling the LLM to decompose and plan instructions into different subtasks, thereby scheduling various tool models to fulfill user instructions.

The Vector Alignment Module maps visual feature vectors to a dimensional space consistent with the LLM input, ensuring structural dimensional alignment and eliminating modality gaps. After alignment, the Semantic Enhancement Module injects high-level semantic information into the visual features, thereby enhancing their expressive power in the semantic space of the language model and providing strong support for cross-modal tasks.

Second, to further improve the effectiveness of cross-modal alignment, we propose a dual-loss training strategy. This strategy includes the Image-Text Generation loss and

the Image-Text Distance Minimization loss. The ITG loss guides the model to generate natural language descriptions from visual inputs, enhancing cross-modal alignment from the perspective of generative tasks. Meanwhile, the ITDM loss minimizes the semantic distance between visual and textual features, further reinforcing deep semantic consistency between the two modalities in alignment tasks. These two loss components complement each other, effectively improving the model’s performance in both generation and alignment tasks.

Finally, to adapt to complex task scenarios, we expand

the output end into a scheduling terminal supporting multiple agents and construct a multi-task scheduling dataset. Through training, MAGE can understand user instructions, perform task decomposition and planning, and call different tool models to achieve the goal. Unlike traditional agent systems driven by large language models (e.g., HuggingGPT), MAGE is a multi-tool invocation model driven by multimodal large language model. Its advantage lies in integrating multiple modalities, enabling more accurate understanding and generation of cross-modal task instructions. At the same time, through flexible tool invocation, it achieves more efficient task execution and supports richer application scenarios.

In summary, the proposed MAGE achieves significant performance improvements in multiple multimodal benchmarks by employing a superior semantic alignment strategy, surpassing existing methods. Experimental results demonstrate that even with a reduction in visual tokens, MAGE outperforms current methods in task performance, resource efficiency, and output flexibility.

The main contributions of this work are summarized as follows:

- We proposed a novel cross-modal alignment architecture, IAN (Intelligent Alignment Network), which integrates a Vector Alignment Module and a Semantic Enhancement Module. Experiments show that this architecture effectively addresses the challenges of dimensional mismatch and semantic gaps between visual and linguistic modalities.
- A dual-loss training strategy was designed, comprising the Image-Text Generation loss and the Image-Text Distance Minimization loss, enabling end-to-end optimization of the semantic alignment process.
- We constructed a multimodal tool-scheduling dataset, HMDSets (a dialogue dataset), and developed an innovative framework. This framework dynamically organizes structured configuration information to accomplish composite generation tasks, significantly extending the model’s capability boundaries.

2 Related Work

2.1 Multimodal Large Language Models

In recent years, multimodal large models such as Flamingo [Alayrac *et al.*, 2022], BLIP-2 [Li *et al.*, 2023b], LLaVA [Liu *et al.*, 2024b], and their subsequent versions (e.g., LLaVA-v1.5 [Liu *et al.*, 2024a] and LLaVA-next [Li *et al.*, 2024a]) have achieved significant progress. Flamingo integrates visual features with language models through the introduction of the Perceiver Resampler and cross-modal attention mechanisms, demonstrating exceptional performance in few-shot learning scenarios. BLIP-2 employs a lightweight Q-Former module to bridge the modality gap between visual encoders and frozen LLMs, thereby improving zero-shot image-to-text generation capabilities. The LLaVA series further enhances multimodal capabilities in vision-language tasks through innovative cross-modal alignment strategies and large-scale image-text pretraining.

2.2 Visual Projector in VLMs

In multimodal large language models (VLMs), the visual projection layer is a critical module for achieving cross-modal alignment between visual information and language models. classical projection layers can generally be categorized into two types: MLP-based projectors and query transformer (Q-Former)-based projectors [Li *et al.*, 2023b; Dai *et al.*, 2023].

In recent years, emerging methods [Cha *et al.*, 2024; Li *et al.*, 2024b; Zhao *et al.*, 2024] have aimed to strike a balance between efficiency and performance. However, existing approaches fail to adequately consider semantic integrity and alignment precision during the alignment process, which provides a strong motivation to explore more efficient and precise visual-language alignment strategies.

2.3 Any-to-Any Model

Recent efforts aim to build models capable of arbitrary modality transformation, mimicking human flexibility. Models like NExT-GPT and CoDi show promise, but struggle with quality and precision in complex tasks. Agent systems such as HuggingGPT [Shen *et al.*, 2024] and MetaGPT [Hong *et al.*, 2023] integrate tools with LLMs to support dynamic multimodal tasks, yet still depend on separate models for understanding and generation—limiting their ability to handle implicit or complex instructions. LLaVA-PLUS [Liu *et al.*, 2025a] aims to enhance multimodal language understanding, focusing on improving LLMs’ ability in language comprehension through the Agent approach. However, it fails to fully address the scalability issue of agents in handling multimodal outputs.

3 Method

3.1 IAN: an Intelligent Alignment Network

The architecture of IAN consists of two core components: the VAB (Vector Alignment Block) and the SEB (Semantic Enhancement Block). These two modules were introduced to address key limitations in prior research, where approaches often relied on simplistic strategies like a single linear layer or a Q-former initialized with random learnable parameters [Liu *et al.*, 2024b; Li *et al.*, 2023b; Cha *et al.*, 2024; Yao *et al.*, 2024]. Such methods struggled to achieve effective semantic and dimensional alignment simultaneously.

Based on previous research, we found that reducing visual tokens doesn’t require complex visual-semantic extractors like Q-former [Yao *et al.*, 2024]. Simple 2D pooling operations at the patch level can effectively reduce tokens while maintaining strong performance. Building on this, we proposed replacing fixed-parameter pooling layers with learnable convolutional layers to enable adaptive learning of mapping patterns. This idea is key to SEB, where trainable convolutional layers refine and enhance visual embeddings to better align with the language model. To address challenges in dimensional alignment and semantic enhancement, we designed IAN with two modules: VAB for alignment and SEB for enhancement. Unlike previous methods that used a single linear layer or Q-former for both tasks—often oversimplifying or complicating the process—our modular approach

decouples these functions, allowing VAB to align dimensions and SEB to refine semantic representations more effectively.

VAB Module We design a flexible linear projection network combined with multilayer nonlinear transformations, enabling the efficient alignment of high-dimensional embeddings from the multimodal encoder with the low-dimensional embedding space of the language model while preserving the original semantic contextual information of the multimodal encoder. Specifically, VAB receives embeddings $\mathbf{V} \in \mathbb{R}^{N \times D_v}$ from the multimodal encoder, where N represents the number of visual features and D_v denotes the embedding dimension of the multimodal encoder. These embeddings undergo nonlinear transformations through MLP and learnable normalization layers, ultimately producing aligned embeddings $\mathbf{V}' \in \mathbb{R}^{N \times D_l}$, where D_l is the embedding dimension of the language model.

SEB Module We propose the SEB to enhance semantic consistency between visual embeddings and the language model via cross-modal self-attention. The image is processed by the CLIP [Radford *et al.*, 2021a] and aligned by VAB to generate a global visual embedding $\mathbf{a} \in \mathbb{R}^{N \times D_l}$, capturing high-level features. To introduce local detail, the image is divided into patches and passed through a CNN [Wu, 2017] to obtain a new embedding $\mathbf{b} \in \mathbb{R}^{N \times \tilde{D}_b}$, which serves as the query. In the attention mechanism, \mathbf{b} acts as the query, while \mathbf{a} is used as both key and value. The attention computation integrates visual and semantic information, resulting in a refined embedding. The output embedding $\mathbf{V}'' \in \mathbb{R}^{N \times D_l}$ incorporates both global and local visual features, effectively bridging the gap between the visual encoder and the large language model.

3.2 Aligning IAN to LLM

In visual language models (VLMs), effective alignment between images and text is essential for enhancing performance on multimodal tasks and improving semantic understanding. While existing methods typically use cross-entropy loss for optimization, this approach has limitations that restrict the model’s multimodal capabilities. We propose combining cross-entropy loss with mean squared error (MSE) loss to better align visual encodings and language embeddings.

Motivation for the New Alignment Strategy

Current methods predominantly rely on cross-entropy loss to optimize language models for accurate text generation [Liu *et al.*, 2024b; Li *et al.*, 2023b; Cha *et al.*, 2024; Li *et al.*, 2024b]. However, this approach does not directly minimize the distance between visual and language embeddings, leading to weak alignment and suboptimal performance in multimodal tasks. Furthermore, cross-entropy loss lacks explicit guidance on the relationship between image and text embeddings, which can result in inconsistent representations. To address these limitations, we propose combining Mean Squared Error (MSE) loss with cross-entropy loss, where once the image features are mapped by the projection layer, they should align with the image caption encoded by the large language model [Zhao *et al.*, 2024]. Additionally, we introduce an image-text distance minimization (ITDM) loss alongside the ITG loss to

directly minimize the distance between visual and text embeddings, thereby enhancing both language generation and multimodal alignment.

Specific Alignment Strategy

Based on the above motivations, our proposed alignment strategy employs two primary loss functions: Image-Guided Text Generation (ITG) loss and Image-Text Distance Minimization (ITDM) loss.

The ITG loss is designed to guide the model in generating corresponding text based on the input image. Specifically, the input image is first processed by the CLIP model to obtain its representation vector. This vector is then passed through the IAN module, where it is mapped into the text embedding space of the large language model. Subsequently, the ITG loss is computed in an autoregressive manner. The loss function can be expressed as follows:

$$\mathcal{L}_{\text{ITG}} = -\frac{1}{L} \sum_{i=1}^L \log p_{\theta}(y_i | y_{<i}, \mathbf{X}_{\text{img}}, \mathbf{X}_{\text{text}}), \quad (1)$$

where $Y = \{y_i\}_{i=1}^L$ denotes the aligned visual embedding produced by IAN, \mathbf{X}_{img} denotes the image input, while \mathbf{X}_{text} represents the textual instructions. The model parameters are represented by θ , which also include the parameters of the projection layer, denoted as IAN.

The ITDM loss function aims to minimize the distance between the image vector processed by IAN and the text vector encoded by the LLM. Specifically, the image vector processed by IAN is represented as \mathbf{d}_{ian} , and the text vector encoded by the LLM is represented as \mathbf{d}_{llm} . The ITDM loss is calculated using the following formula:

$$\mathcal{L}_{\text{ITDM}} = \frac{1}{N} \sum_{i=1}^N \|\mathbf{d}_{\text{ian}} - \mathbf{d}_{\text{llm}}\|_2^2, \quad (2)$$

Here, N represents the batch size. To fully leverage the advantages of both the cross-entropy loss and the mean squared error loss, we define a composite loss function as follows:

$$\mathcal{L} = \mathcal{L}_{\text{ITG}} + \mathcal{L}_{\text{ITDM}}, \quad (3)$$

The composite loss function improves language generation and aligns visual and language embeddings. Through contrastive learning, multimodal vectors are aligned with the semantic space of large language models. After fine-tuning, the model better processes multimodal information and enhances its outputs.

3.3 Extended SKILLS with Multimodal Tool Use

To expand the input boundaries of the model, we propose a novel architectural design that supports multimodal outputs and enhances the model’s task planning capabilities to solve complex tasks.

HMDSet Dataset Construction We designed a human-machine dialogue dataset (HMDSet) that includes a wide range of input and output modalities, covering text, images, audio, video, and complex task handling. This enables MAGE to tackle complex cross-modal understanding and generation tasks.

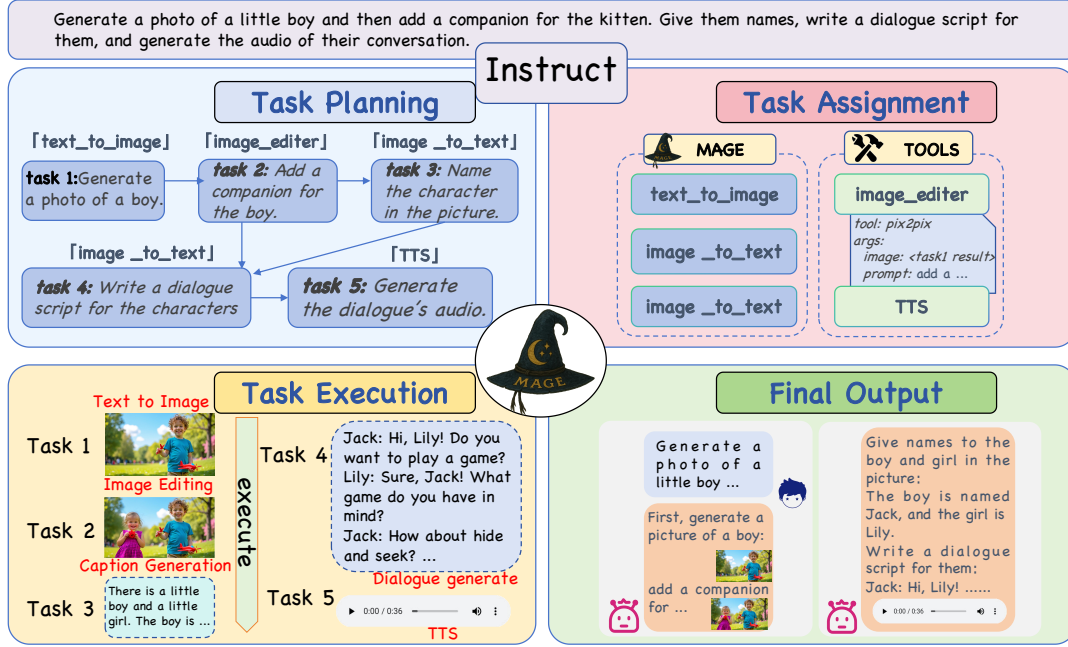


Figure 3: **Hierarchical Workflow of the Proposed Agent Framework.** This figure illustrates a hierarchical process for generating an image of a boy with a companion. The process includes task planning, task assignment, task execution, and final output.).

Hierarchical Architecture and Output Format of the Agent The unique aspect of this multimodal agent lies in its ability to call the required toolset in a structured JSON format based on natural language output generated by a large language model (LLM), as shown in fig. 3. We employ a hierarchical strategy in the task planning and prediction process to efficiently determine the specific tools and parameters.

4 Experiments

4.1 Implementation Details

In this section, we describe the experimental setup used in our study to evaluate the performance of our proposed model. The following outlines the implementation details, datasets, and benchmarks that were used to ensure robustness and comprehensiveness of the results.

Model Configuration. We used two variants of the Vicuna-v1.5 [Zheng *et al.*, 2023] large language model (7B and 13B) and CLIP ViT-L/14 [Radford *et al.*, 2021b] (336px) as the vision encoder, with input images set to a resolution of 336×336 pixels. Visual features are derived from the penultimate layer of the CLIP model, excluding the CLS token. During training, we fine-tuned all model parameters (including CLIP and LLM) without using parameter-efficient techniques like LoRA [Hu *et al.*, 2021], ensuring the model learns complete representations. Training was conducted using 48 A6000 GPUs. All the training data is presented in table 2. Our training process is divided into three phases:

Stage 1: Pretraining Details. For the initial training phase, we use the CC-558K dataset, to perform semantic alignment across modalities. In this phase, we freeze the parameters

of the visual encoder (CLIP) and the LLM, focusing solely on training the IAN. This approach enables us to align the visual and linguistic features without altering the core model components.

Stage 2: Instruction-tuning Details. In the second phase, we scale up the dataset to approximately 980K samples by combining the 665K mixture dataset with the ALLaVA-Instruct-VFLAN-4V [Chen *et al.*, 2024] dataset. This extended dataset includes a rich set of question-answer pairs, simulating real-world user interactions. For training in this phase, only 90% of this dataset was used. This phase enables fine-tuning of the model for more natural and contextually relevant multimodal responses.

Stage 3: Self-Cognition Training Details We refer to the third stage of training as "Self-Awareness Planning Training," aimed at enhancing the model's ability to handle complex cross-modal understanding and generation. The training data combines the remaining 10% of the fine-tuning dataset from the second stage with the newly constructed self-awareness dataset, which includes multimodal human-machine dialogues involving text, images, audio, and video. The output is a JSON task plan specifying tool categories and parameters. The dialogue examples are generated by GPT-4 and manually reviewed to ensure diversity and accuracy. This stage of training enables the model to develop a chain-of-thought capability, allowing it to proactively invoke tools and autonomously plan and execute complex tasks.

4.2 Experimental Setting

We introduce experimental settings, including the benchmarks and evaluation metrics.

Method	LLM	Vision Encoder	Projector	#Token	Res.	MMB	MMB ^{cn}	SEED ^l	MME	POPE
Approaches using 7B LLM										
Qwen-VL-Chat [Bai <i>et al.</i> , 2023]	Qwen-7B	OpenCLIP ViT-bigG	Resampler	256	448	60.6	56.3	58.2	1487/360	-
InstructBLIP [Dai <i>et al.</i> , 2023]	Vicuna-7B	EVA-CLIP ViT-G	Q-former	64	224	36.0	23.7	53.4	-	-
LLaVA-TokenPacker-HD [Li <i>et al.</i> , 2024b]	Vicuna-7B	CLIP ViT-L/14	TokenPacker	64	336	67.4	-	-	1489/338	86.3
LLaVA-1.5 [Liu <i>et al.</i> , 2024a]	Vicuna-7B	CLIP ViT-L/14	Linear	576	336	64.3	58.3	58.6	1510/-	85.9
LLaVA-NeXT [Li <i>et al.</i> , 2024a]	Vicuna-7B	CLIP ViT-L/14	Linear	576	336	67.4	62.3	-	1519/332	86.5
C-Abstractor [Cha <i>et al.</i> , 2024]	Vicuna-7B	CLIP ViT-L/14	C-Abstractor	144	224	70.1	-	64.5	1584 /307	-
D-Abstractor [Cha <i>et al.</i> , 2024]	Vicuna-7B	CLIP ViT-L/14	D-Abstractor	144	224	70.8	-	63.8	1544/291	-
MAGE(Ours)	Vicuna-7B	CLIP ViT-L/14	IAN	144	336	71.2	70.3	65.3	1572 / 323	86.7
Approaches using 13B LLM										
InstructBLIP [Dai <i>et al.</i> , 2023]	Vicuna-13B	EVA-CLIP ViT-G	Q-former	64	224	-	-	-	1212/-	78.9
LLaVA-TokenPacker-HD [Li <i>et al.</i> , 2024b]	Vicuna-13B	CLIP ViT-L/14	TokenPacker	64	336	69.5	-	-	1595/ 356	88.1
LLaVA-1.5 [Liu <i>et al.</i> , 2024a]	Vicuna-13B	CLIP ViT-L/14	Linear	576	336	67.7	63.6	61.6	1531/-	85.9
LLaVA-NeXT [Li <i>et al.</i> , 2024a]	Vicuna-13B	CLIP ViT-L/14	Linear	576	336	70	68.5	-	1575/326	86.2
C-Abstractor [Cha <i>et al.</i> , 2024]	Vicuna-13B	CLIP ViT-L/14	C-Abstractor	256	336	73.2	-	68.2	1629/315	-
D-Abstractor [Cha <i>et al.</i> , 2024]	Vicuna-13B	CLIP ViT-L/14	D-Abstractor	256	336	73.5	-	66.6	1632 /333	-
MAGE(Ours)	Vicuna-13B	CLIP ViT-L/14	IAN	144	336	73.9	73.0	67.8	1632 / 312	89.9

Table 1: **Comparison with other SOTA.** Res means the resolution of the image encoder, and #Token represents the number of visual tokens. The best results are shown in bold, and the second-best are underlined.

Task	Datasets	#samples
Pretraining dataset	BlipCapFilt [Li <i>et al.</i> , 2022]	558K
General Instruction	LLaVA150K [Liu <i>et al.</i> , 2024b], ShareGPT [Chiang <i>et al.</i> , 2023]	665K
Doc/Chart/Screen	Doc-VQA [Mathew <i>et al.</i> , 2021], ChartQA [Masry <i>et al.</i> , 2022], DVQA [Kafle <i>et al.</i> , 2018]	43K
Math/Reasoning	GeoQA+ [Chen <i>et al.</i> , 2021],	17K
General OCR	SynthDog-EN [Kim <i>et al.</i> , 2022]	40K
Tool Use	HMDSet	31K

Table 2: List of all training datasets.

Benchmarks. We evaluate the model’s performance using four multimodal large language model (MLLM) benchmarks, each targeting different aspects of multimodal understanding. MME [Fu *et al.*, 2024] assesses perception and cognition capabilities using a binary yes/no format; MMBench [Liu *et al.*, 2025b] evaluates reasoning and multimodal tasks through multiple-choice questions; SEED-Bench [Li *et al.*, 2023a] tests the model’s ability to handle complex multimodal queries; and POPE [Li *et al.*, 2023c] detects model hallucinations using binary questions to assess the accuracy of providing correct information.

Evaluation Metrics. We use the official evaluation scripts provided for each benchmark to measure the model’s performance. For MMBench and SEED-Bench, the evaluation metric is accuracy. POPE’s performance is assessed using the F1 score, balancing precision and recall. MME scores are calculated using official scripts, with separate results for the Perception and Cognition tasks. Additionally, we report the normalized average score (Avg^N [Chen *et al.*, 2020; Li *et al.*, 2021]), facilitating fair comparisons across different tasks.

Model	MMB	MMB ^{cn}	SEED ^l	MME	POPE
MAGE	71.2	70.3	65.3	1572/323	86.7
w/o IAN proj.	68.1	67.8	62.9	1538/301	86.0
w/o align.	69.7	68.3	61.4	1541/307	86.2
w/o IAN & align.	67.0	65.8	60.9	1521/294	85.9

Table 3: **Ablation study evaluating the impact of the IAN projector and alignment strategies in the proposed MAGE model.** The table demonstrates how the performance decreases when the IAN projector (w/o IAN proj.), the alignment strategy (w/o align.), or both (w/o IAN & align.) are removed.

4.3 Overall Results

The results, as shown in table 1, demonstrate that the MAGE model proposed in this study performs excellently in a 7B parameter scale LLM, particularly when using only 144 visual tokens. The MAGE model outperforms those using 256 and 576 visual tokens. For instance, MAGE scores 71.2 on the MMBench benchmark, significantly surpassing LLaVA-TokenPacker-HD (67.4) and LLaVA-1.5 (64.3). In MMBench-cn, MAGE scores 70.3, leading LLaVA-NeXT (62.3) and Qwen-VL-Chat (56.3). On the SEED benchmark, MAGE scores 65.3, surpassing C-Abstractor (64.5) and D-Abstractor (63.8).

Furthermore, MAGE 14B also excels on the POPE benchmarks, with a score of 89.9 on POPE, outperforming LLaVA-1.5 (85.9) and LLaVA-TokenPacker-HD (88.1), and scoring 73.9 on MMBench, surpassing LLaVA-TokenPacker-HD (69.5) and D-Abstractor (73.5). These results indicate that MAGE efficiently extracts visual features, significantly reduces computational costs while improving performance, demonstrating its competitiveness in multimodal tasks.

4.4 Ablation Study

The impact of the IAN projection module (IAN proj.) and alignment strategy (align.) on model performance was

	Projector	#Token	MME POS	MMB SR	MMB OL	SEED PR	SEED SR	SEED IL	Avg ^N
B1	Linear	576	140	24.4	40.7	70.8	48.9	60.9	52.6
B2	Resampler	144	75.0	22.2	43.2	62.5	47.5	50.6	43.9
B3	C-Abstractor	144	135	24.4	54.3	66.7	49.0	58.8	53.5
B4	IAN	144	135	33.3	51.9	67.0	50.3	55.4	54.2
B5	Resampler	256	73.3	24.4	37.0	79.2	44.4	51.8	45.6
B6	C-Abstractor	256	136.7	26.7	55.6	75.0	52.7	59.3	56.3
B7	IAN	256	140	39.2	53.1	76.5	54.7	56.1	58.3

Table 4: **Comparison of Different Projectors** The table presents results for various tasks like Position (POS), Spatial Relationship (SR), Object Localization (OL), and Physical Relation (PR) for MMBench, and SR and Instance Location (IL) for SEED-Bench. Avg^N represents the normalized average across all six tasks. The number of visual tokens (#Token) is also reported. Bold values indicate the best results.

evaluated through ablation experiments on five benchmark datasets: MMBench, MMBench-cn, SEED, MME, and POPE, with results summarized in table 3.

The full MAGE model achieved the best performance across all benchmarks. Removing the IAN projection module led to a performance drop, particularly on SEED (from 65.3 to 62.9). Further removal of the alignment strategy caused an additional decrease, with the SEED score dropping to 61.4. The lowest performance was observed when both components were removed, with the SEED score falling to 60.9. These results emphasize the critical role of the IAN module and alignment strategy in improving model performance.

4.5 Analysis on IAN

The experimental results, as shown in table 4, demonstrate that MAGE outperforms other designs across multiple tasks. In the MME POS task, MAGE scored 140, surpassing other bridging layers. In the MMBench SR task, MAGE achieved the best score of 39.2. Similarly, in the SEED-Bench SR task, MAGE ranked first with a score of 54.7 and an average score of 58.3, outperforming all other models.

Even with only 144 visual tokens, MAGE scored 33.3 in the MMBench SR task, outperforming models using 256 visual tokens, such as Resampler (22.2) and C-Abstractor (24.4), demonstrating its efficient reasoning under resource constraints.

Furthermore, the results in table 5 show that increasing the number of visual tokens and model size leads to improved performance. In the MMB task, the 7B model improved from a score of 70.9 with 64 visual tokens to 71.4 with 256 visual tokens, while the 13B model improved from 73.1 to 74.4, indicating stable growth. The 13B model outperformed the 7B model in most tasks, suggesting that larger models with more visual tokens effectively enhance performance.

4.6 Tool Use Expand the Output Modality of the VLM

Figure 4 illustrates three specific examples, covering two basic tasks: text-to-image and text-to-Video, as well as a more complex task: generating audio and video based on an uploaded image. In these tasks, the model automatically se-

Model	#Token	MMB	MMB ^{cn}	MME ^P	MME	SEED ^I	POPE
7B	64	70.9	70.0	1544	1851	64.8	86.3
	144	71.2	70.3	1572	1895	65.3	86.7
	256	71.4	70.6	1580	1913	65.6	86.9
13B	64	73.1	71.9	1595	1907	67.4	89.1
	144	73.9	73.0	1632	1944	67.8	89.9
	256	74.4	73.3	1643	1969	68.1	90.3

Table 5: **Ablation study on the effect of visual token settings.** The table compares two model sizes across different tokens (#Token).



Figure 4: The figure shows an example of how our model extends multimodal applications by utilizing tools, including two basic tasks: text-to-image and text-to-video, as well as a composite task that generates images by combining both images and videos.

lects the most appropriate tool for processing based on the input text description or multimodal content, thereby generating high-quality outputs.

In particular, for situations where direct modality generation is not possible, the model coordinates the invocation of other tools for inference. As shown in case 3 in Figure 4, when there is no available tool in the toolset to directly generate audio and video from an image, the model first calls the image-to-text tool to convert the image into a textual description, which is then used as the prompt for generating audio and video. This mechanism significantly enhances the adaptability and flexibility of VLMs when handling diverse tasks.

5 Conclusion

This study proposes MAGE (Multimodal Alignment and Generation Enhancement), a novel architecture that addresses semantic and dimensional gaps between visual encoders and large language models (LLMs). MAGE achieves efficient alignment through the Intelligent Alignment Network (IAN), optimizes generation and alignment using a dual-loss strategy, and supports flexible multimodal task outputs, such as image, audio, and video generation. Experimental results demonstrate that MAGE performs exceptionally well across various benchmarks, achieving state-of-the-art efficiency and accuracy.

Acknowledgements

This research was supported by the Beijing Municipal Science and Technology Project and the GTCOM 2030 AI Research Institute (AIRIS).

Contribution Statement

Shaojun E and Yuchen Yang contributed equally. Ziyang Chen is the corresponding author.

References

- [Alayrac *et al.*, 2022] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736, 2022.
- [Anthropic, 2024] Anthropic. Claude 3.5 sonnet. <https://www.anthropic.com/news/claude-3-5-sonnet>, 2024. Accessed: 2024-7-21.
- [Bai *et al.*, 2023] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 1(2):3, 2023.
- [Brown *et al.*, 2020] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [Cha *et al.*, 2024] Junbum Cha, Wooyoung Kang, Jonghwan Mun, and Byungseok Roh. Honeybee: Locality-enhanced projector for multimodal llm. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13817–13827, 2024.
- [Chen *et al.*, 2020] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer, 2020.
- [Chen *et al.*, 2021] Jiaqi Chen, Jianheng Tang, Jinghui Qin, Xiaodan Liang, Lingbo Liu, Eric P Xing, and Liang Lin. Geoqa: A geometric question answering benchmark towards multimodal numerical reasoning. *arXiv preprint arXiv:2105.14517*, 2021.
- [Chen *et al.*, 2024] Guiming Hardy Chen, Shunian Chen, Ruifei Zhang, Junying Chen, Xiangbo Wu, Zhiyi Zhang, Zhihong Chen, Jianquan Li, Xiang Wan, and Benyou Wang. Allava: Harnessing gpt4v-synthesized data for a lite vision-language model. *arXiv preprint arXiv:2402.11684*, 2024.
- [Chiang *et al.*, 2023] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, March 2023.
- [Dai *et al.*, 2023] 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, 2023.
- [Dai *et al.*, 2024] Wenliang Dai, Nayeon Lee, Boxin Wang, Zhuolin Yang, Zihan Liu, Jon Barker, Tuomas Rintamäki, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. Nvlm: Open frontier-class multimodal llms. *arXiv preprint arXiv:2409.11402*, 2024.
- [Fu *et al.*, 2024] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiaowu Zheng, Ke Li, Xing Sun, Yunsheng Wu, and Rongrong Ji. Mme: A comprehensive evaluation benchmark for multimodal large language models, 2024.
- [Hong *et al.*, 2023] Sirui Hong, Xiaowu Zheng, Jonathan Chen, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, et al. Metagpt: Meta programming for multi-agent collaborative framework. *arXiv preprint arXiv:2308.00352*, 2023.
- [Hu *et al.*, 2021] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- [Kafle *et al.*, 2018] Kushal Kafle, Brian Price, Scott Cohen, and Christopher Kanan. Dvqa: Understanding data visualizations via question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5648–5656, 2018.
- [Kim *et al.*, 2022] Geewook Kim, Teakgyu Hong, Moonbin Yim, Jeongyeon Nam, Jinyoung Park, Jinyeong Yim, Wonseok Hwang, Sangdoo Yun, Dongyoon Han, and Seunghyun Park. Ocr-free document understanding transformer. In *European Conference on Computer Vision*, pages 498–517. Springer, 2022.
- [Li *et al.*, 2021] Linjie Li, Jie Lei, Zhe Gan, Licheng Yu, Yen-Chun Chen, Rohit Pillai, Yu Cheng, Luowei Zhou, Xin Eric Wang, William Yang Wang, et al. Value: A multi-task benchmark for video-and-language understanding evaluation. *arXiv preprint arXiv:2106.04632*, 2021.
- [Li *et al.*, 2022] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, pages 12888–12900. PMLR, 2022.
- [Li *et al.*, 2023a] Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. Seed-bench: Benchmarking multimodal llms with generative comprehension. *arXiv preprint arXiv:2307.16125*, 2023.
- [Li *et al.*, 2023b] 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 *International conference on machine learning*, pages 19730–19742. PMLR, 2023.
- [Li *et al.*, 2023c] Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*, 2023.
- [Li *et al.*, 2024a] 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.
- [Li *et al.*, 2024b] Wentong Li, Yuqian Yuan, Jian Liu, Dongqi Tang, Song Wang, Jie Qin, Jianke Zhu, and Lei Zhang. Tokenpacker: Efficient visual projector for multimodal llm. *arXiv preprint arXiv:2407.02392*, 2024.
- [Liu *et al.*, 2024a] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306, 2024.
- [Liu *et al.*, 2024b] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024.
- [Liu *et al.*, 2025a] 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 *European Conference on Computer Vision*, pages 126–142. Springer, 2025.
- [Liu *et al.*, 2025b] 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 *European Conference on Computer Vision*, pages 216–233. Springer, 2025.
- [Masry *et al.*, 2022] Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv preprint arXiv:2203.10244*, 2022.
- [Mathew *et al.*, 2021] Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2200–2209, 2021.
- [OpenAI, 2024] OpenAI. Hello gpt-4o. <https://openai.com/index/hello-gpt-4o/>, 2024. Accessed: 2024-5-13.
- [Radford *et al.*, 2021a] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [Radford *et al.*, 2021b] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [Shen *et al.*, 2024] Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. Hugginggpt: Solving ai tasks with chatgpt and its friends in hugging face. *Advances in Neural Information Processing Systems*, 36, 2024.
- [Team *et al.*, 2023] Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- [Vaswani, 2017] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
- [Wang *et al.*, 2024] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024.
- [Wu, 2017] Jianxin Wu. Introduction to convolutional neural networks. *National Key Lab for Novel Software Technology. Nanjing University. China*, 5(23):495, 2017.
- [Yao *et al.*, 2024] Linli Yao, Lei Li, Shuhuai Ren, Lean Wang, Yuanxin Liu, Xu Sun, and Lu Hou. Deco: Decoupling token compression from semantic abstraction in multimodal large language models. *arXiv preprint arXiv:2405.20985*, 2024.
- [Yu *et al.*, 2022] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022.
- [Yuan *et al.*, 2021] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021.
- [Zhao *et al.*, 2024] Xiangyu Zhao, Bo Liu, Qijiong Liu, Guangyuan Shi, and Xiao-Ming Wu. Easygen: Easing multimodal generation with bidiffuser and llms. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1351–1370, 2024.
- [Zheng *et al.*, 2023] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623, 2023.