

Augmenting Intra-Modal Understanding in MLLMs for Robust Multimodal Keyphrase Generation

Jiajun Cao^{1,5}, Qinggang Zhang^{3*}, Yunbo Tang², Zhishang Xiang², Chang Yang³, Jinsong Su^{2,4,5*}

¹Department of Digital Media Technology, Xiamen University

²School of Informatics, Xiamen University

³The Hong Kong Polytechnic University

⁴Shanghai Artificial Intelligence Laboratory

⁵Key Laboratory of Digital Protection and Intelligent Processing of Intangible Cultural Heritage of Fujian and Taiwan (Xiamen University), Ministry of Culture and Tourism, China

caojiajun@stu.xmu.edu.cn, zqg.zhang@hotmail.com, jssu@xmu.edu.cn

Abstract

Multimodal keyphrase generation (MKP) aims to extract a concise set of keyphrases that capture the essential meaning of paired image-text inputs, enabling structured understanding, indexing, and retrieval of multimedia data across the web and social platforms. Success in this task demands effectively bridging the semantic gap between heterogeneous modalities. While multimodal large language models (MLLMs) achieve superior cross-modal understanding by leveraging massive pretraining on image-text corpora, we observe that they often struggle with modality bias and fine-grained intra-modal feature extraction. This oversight leads to a lack of robustness in real-world scenarios where multimedia data is noisy, along with incomplete or misaligned modalities. To address this problem, we propose AimKP, a novel framework that explicitly reinforces intra-modal semantic learning in MLLMs while preserving cross-modal alignment. AimKP incorporates two core innovations: (i) Progressive Modality Masking, which forces fine-grained feature extraction from corrupted inputs by progressively masking modality information during training; (ii) Gradient-based Filtering, that identifies and discards noisy samples, preventing them from corrupting the model’s core cross-modal learning. Extensive experiments validate AimKP’s effectiveness in multimodal keyphrase generation and its robustness across different scenarios.

Code — <https://github.com/XMUDeepLIT/AimKP>

Introduction

With the explosive growth of multimedia content across the web and social platforms, there is an increasing demand for advanced techniques to understand and organize multimodal data. Multimodal keyphrase generation (MKP) addresses this critical need by generating concise, semantically rich keyphrases that encapsulate the essential meaning of multimodal inputs, enabling structured understanding, efficient indexing, and cross-modal retrieval. For instance, consider the example in the left panel of Figure 1, where the

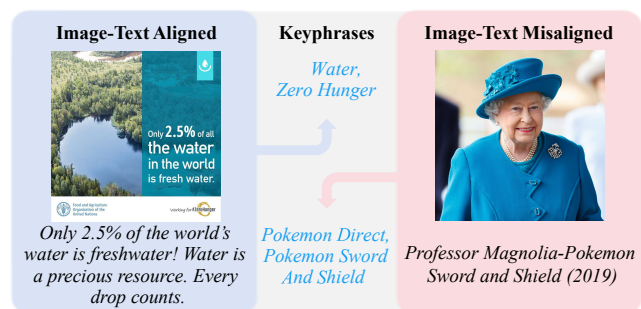


Figure 1: Examples of MKP, demonstrating cases of image-text aligned (left) and image-text misaligned (right) pairs.

text emphasizes global freshwater scarcity and the image depicts a freshwater lake with related slogans and organizational logos. An effective MKP system should generate both the explicit keyphrase *Water* and the implicit thematic keyphrase *Zero Hunger*. This capability enables critical applications such as opinion mining and content recommendation, where complementary multimodal features are needed to yield accurate and human-aligned keyphrases.

Compared to traditional text-based keyphrase generation (Chen et al. 2018; Yuan et al. 2020; Ye et al. 2021), MKP requires the model to achieve both granular comprehension of modality-specific semantics for anchoring critical cues, and cross-modal integration for aligned semantic fusion. Earlier MKP methods predominantly focus on cross-modal alignment via attention mechanisms (Gong and Zhang 2016; Zhang et al. 2019), frequently incorporating external tools such as OCR systems, object detectors (Wang et al. 2020), or APIs (Dong et al. 2023). Although utilizing these auxiliary resources enhances multimodal semantic understanding, such methods are fundamentally limited by the base models’ reasoning capabilities. As shown in the right panel of Figure 1, when presented with textual descriptions of Pokémon’s fictional Professor Magnolia alongside images depicting the real-world Queen Elizabeth, these lightweight models struggle to disambiguate cross-modal entities, leading to erroneous keyphrase generation.

*Corresponding author.

Recently, the advent of multimodal large language models (MLLMs) (Alayrac et al. 2022; Li et al. 2023; Liu et al. 2023) has revolutionized multimodal understanding. By leveraging massive pretraining on image-text corpora, MLLMs exhibit remarkable capabilities in text recognition and visual grounding, and have set new benchmarks in tasks like image captioning and visual question answering (OpenAI 2023; Liu et al. 2024; Qwen et al. 2024). Despite recent advances, directly deploying MLLMs for MKP is challenging due to their divergent objectives: MKP requires a fine-grained understanding of modality-specific semantics for keyphrase generation, whereas MLLMs prioritize cross-modal alignment, inherently sacrificing granular semantics.

This gap is quite evident in practice. The preliminary study on LLaVA-1.5 (Figure 2) shows that this representative MLLM achieves competitive performance on multimodal data, but suffers from severe degradation when processing single-modality context. More critically, it underperforms specialized single-modality models by 4%-8%, with the largest discrepancy (8%) occurring in image-only scenarios. These observations reveal two critical points: (i) MLLMs struggle with intra-modal understanding. Existing MLLMs are trained on tightly aligned multimodal data, their cross-attention mechanism encourages the model to prioritize high-level cross-modal associations over fine-grained, modality-specific details. This inadvertently suppresses modality-specific reasoning capabilities to anchor keyphrases in specific visual or textual cues, which are essential for keyphrase generation. (ii) MLLMs always suffer from modality bias. Most MLLMs exhibit a strong preference for a specific modality (Parcalabescu and Frank 2025; Zhang et al. 2025b,c; Zheng et al. 2025). For example, LLaVA exhibits a textual bias from its language-heavy pretraining, tending to over-weight textual signals and overlook subtle visual cues. Such imbalance violates the core requirement of MKP for adaptive modality fusion.

These limitations become more obvious in practical scenarios, as real-world multimedia data typically exhibits noise along with incomplete or misaligned modalities. To address this problem, we introduce AimKP, a unified training framework that adapts MLLMs for MKP through two innovations: (i) *Progressive Modality Masking* that forces fine-grained feature extraction from corrupted inputs by progressively masking of modality information, and (ii) *Gradient-Based Filtering* dynamically prunes masked samples based on their gradients, preventing conflicting signals from harmful corruptions. To the best of our knowledge, we are the first to propose a framework that systematically adapts MLLMs to MKP task. Our contributions are summarized as follows:

- We identify key limitations of MLLMs in MKP and, based on these findings, propose AimKP, a novel framework to adapt MLLMs for the task.
- AimKP first introduces Progressive Modality Masking, a scheme that systematically masks modality information to force fine-grained feature extraction.
- To stabilize training, AimKP also incorporates Gradient-Based Filtering, which measures the similarity of gradients to prune uninformative or harmful masked samples.

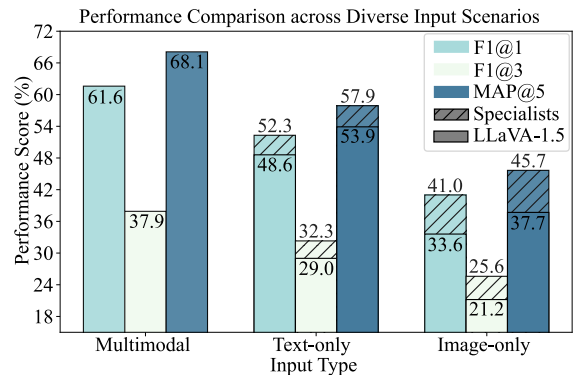


Figure 2: Performance comparison of MLLMs fine-tuned on multimodal vs. unimodal contexts across three input settings: full multimodal input, text-only input, and image-only input, with metrics: F1@1, F1@3, MAP@5.

- Extensive experiments on the benchmark dataset demonstrate that AimKP substantially improves MLLMs’ intra-modal understanding and achieves a new state-of-the-art in overall MKP performance.

Preliminary Study

Before going into the details of AimKP, we first conduct a preliminary study to explore the potential and challenges of applying MLLMs in MKP. This study serves two core purposes: (i) to verify whether MLLMs, with their strong cross-modal alignment capabilities, can serve as a viable foundation for MKP; and (ii) to identify critical limitations in their current performance that demand targeted improvements, laying the groundwork for the design of our framework.

MKP with MLLMs

To delve into MLLMs for MKP, we leverage a representative MLLM (LLaVA-1.5), which consists of a CLIP vision encoder (Ilharco et al. 2021), a lightweight visual adapter, and a Vicuna (Chiang et al. 2023) language model backbone. As illustrated in Figure 4(a), the image X_V is divided into 24×24 non-overlapping patches, which are then flattened into a 1D sequence. These patches are encoded into visual embeddings via the vision encoder and adapter, with each patch’s embedding functioning as a token in the language model. Concurrently, the text X_T is appended with a task prompt, tokenized, and embedded using the model’s text encoder. We then concatenate the visual embeddings and the textual embeddings into a unified multimodal input following the instruction tuning paradigm. The model is trained to autoregressively generate the full keyphrase sequence $Y = \{y_1, \dots, y_{|Y|}\}$ conditioned on the inputs, maximizing the likelihood of ground-truth keyphrases.

For the fine-tuning setup, the vision encoder is frozen. We train the visual adapter and language model jointly, with Low-Rank Adaptation (LoRA) applied to the language

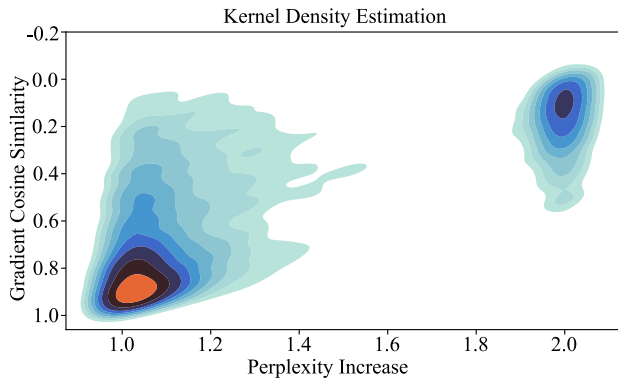


Figure 3: Kernel density plot of cosine similarity of gradients vs. perplexity increase (the ratio of masked-sample keyphrase perplexity to original-sample perplexity).

model. The training loss is standard cross-entropy loss:

$$\mathcal{L} = - \mathbb{E}_{(X_V, X_T, Y) \sim \mathcal{D}} \left[\sum_{t=1}^{|Y|} \log p_{\theta}(y_t | X_V, X_T, y_{<t}) \right]. \quad (1)$$

Identifying the Intra-Modal Deficit

We focus on evaluating MLLMs’ performance across diverse scenarios, particularly focusing on how they handle modality-specific semantics. Empirical results in Figure 2 show that LLaVA achieves strong performance on multi-modal inputs, outperforming existing MKP methods with an F1@1 of 61.6% and MAP@5 of 68.1%. This confirms MLLMs’ capacity for cross-modal understanding. However, when either the image or the text is missing, the model’s performance drops significantly. Further comparing with single-modality specialists (i.e., LLaVA fine-tuned and performing inference solely on text or images for keyphrase generation), LLaVA lags behind the text specialist by 3.7% in F1@1 and 4.0% in MAP@5 on text-only inputs, and 7.4% in F1@1 and 8.0% in MAP@5 behind the image specialist on image-only inputs, leaving its **intra-modal performance far from its theoretical ceiling**.

We attribute these results to modality bias and a critical behavioral shortcut: on well-aligned training data, the model becomes **over-reliant on cross-modal associations** because this strategy is often the easiest path to minimize loss. When its understanding of one modality is insufficient, it compensates with cues from the other rather than developing robust intra-modal understanding. Compounding its **inherent textual preference**, this combination weakens the model’s ability to grasp visual details, widening the gap in image-only scenarios. This strategy, however, becomes fragile in real-world scenarios where one modality may be uninformative, noisy, or even misleading. In these cases, underdeveloped intra-modal reasoning provides no reliable fallback, leading to severe performance degradation when fusing these problematic cross-modal signals.

Motivation for Gradient-Based Filtering Strategy

To address this problem, a natural intuition is to mask one modality during training, forcing the model to reason more from the unmasked modality. However, naive modality masking is risky: when core intra-modal cues are masked, the samples become uninformative noise and can undermine training. Hence, we require a mechanism that flags when a masked sample would steer the model toward divergent directions. Drawing on gradient balancing from multi-task learning (Wei and Hu 2024), we compute the cosine similarity between the gradient of the original loss and that of each masked variant. Figure 3 shows a clear negative relationship between gradient similarity and the increase in perplexity caused by masking. Samples whose masking barely raises perplexity (preserving key information for keyphrase generation) tend to have high gradient similarity. In contrast, those with large increases in perplexity exhibit low gradient alignment. This empirical observation supports our hypothesis that **gradient similarity can serve as a reliable flag to identify uninformative masked samples**, enabling us to filter them out and stabilize training.

The Framework of AimKP

In this section, we propose a unified framework to enhance intra-modal learning without sacrificing cross-modal alignment. As illustrated in Figure 4(b), AimKP comprises: (i) *Progressive Modality Masking*, which forces the model to reason deeply within one modality by progressive masking information of the other modality; (ii) *Gradient-Based Filtering*, which filters out uninformative masked samples, avoiding conflict with the core learning objective.

Progressive Modality Masking

The underdeveloped intra-modal reasoning in MLLMs arises from their modality bias and over-reliance on cross-modal associations as a training shortcut. To address the intra-modal reasoning deficit, we progressively mask information from one modality, compelling the model to extract rich semantics from the corrupted input. These masked samples are further filtered to retain only informative ones, as described in the next section.

For each training pair (X_V, X_T) , we apply structured, gradually increasing masks to both image and text inputs by setting the corresponding regions of the attention mask to zero for masked areas. Both modalities undergo masking at the embedding level: text masking retains tokens at fixed intervals along the sequence, while image masking preserves tokens based on their pre-flattened spatial positions (i.e., height and width in the original grid), ensuring alignment with the 2D structure of images. Specifically, given a stride parameter γ , we define binary masks over 2D visual features and 1D textual tokens:

$$M_{2D}(i, j) = \begin{cases} 1, & (i \bmod \gamma = 0) \wedge (j \bmod \gamma = 0), \\ 0, & \text{otherwise,} \end{cases} \quad (2)$$

$$M_{1D}(t) = \begin{cases} 1, & t \bmod \gamma = 0, \\ 0, & \text{otherwise.} \end{cases}$$

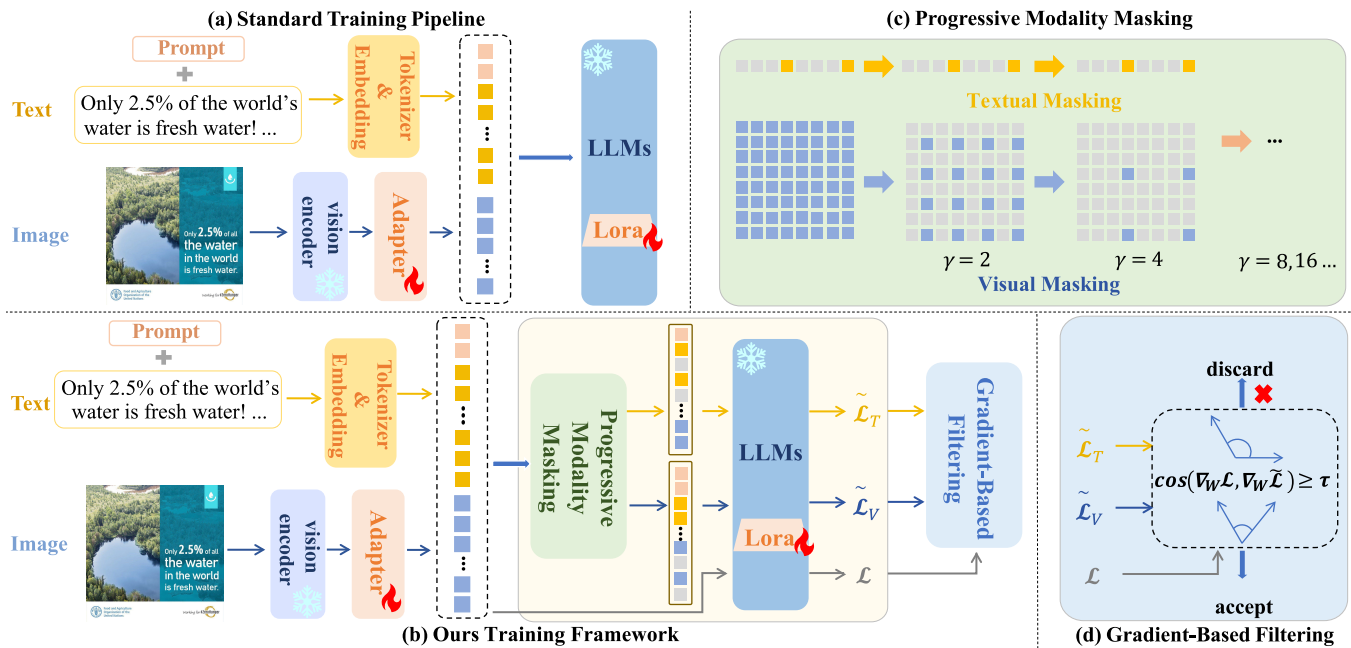


Figure 4: The Framework of AimKP. (a) Standard multimodal fine-tuning. (b) Our intra-modal enhancement framework, which (c) progressively masks modality information at increasing rates to force the model to reason deeply within one modality, and (d) dynamically prunes masked samples based on their gradients, preventing conflicting signals from harmful corruptions.

In practice, we retain the last token within each stride. For a given stride $\gamma = k$, the retention ratio is $1/k^2$ for image tokens and $1/k$ for text tokens. We then apply these masks to the original sequences to obtain the masked inputs \tilde{X}_V and \tilde{X}_T via token-wise masking. γ is initialized to 2 for both modalities and doubled each epoch, systematically increasing masking intensity. As γ grows, the amount of retained information decreases in a controllable manner. Starting with mild masking (e.g., $\gamma = 2$, retaining 50% text tokens or 25% image tokens) to help the model fundamentally adapt to partially missing inputs; then as masking intensifies (e.g., $\gamma = 4$, retaining 25% text tokens or 6.25% image tokens), the model is forced to mine gradually deeper intra-modal understanding. This aligns with curriculum learning principles, where the model adapts to progressively more challenging tasks. We further introduce refinements based on sample informativeness to γ detailed in the following section.

Gradient-Based Filtering

While progressive masking enhances the model’s intra-modal learning and escalates task difficulty, overly aggressive masking can introduce noise or remove essential signals about the task, destabilizing training. As noted in our preliminary analysis, gradient similarity serves as an effective proxy for the informativeness of a given masked sample. To filter out uninformative variants, we compute the cosine similarity between the gradient of the original loss \mathcal{L} and that of its masked counterpart $\tilde{\mathcal{L}}$: $s = \cos(\nabla_W \mathcal{L}, \nabla_W \tilde{\mathcal{L}})$. High similarity score suggests the masking is effective and produces informative variants, whereas low similarity suggests the masked sample introduces conflicting gradients that can

harm the optimization of the primary objective \mathcal{L} . We apply a threshold τ to decide each variant’s fate:

- If $s \geq \tau$, we include the loss of the masked inputs as an auxiliary loss in the training objective. This high similarity can also be interpreted as an opportunity to challenge the model further, so we double the corresponding stride for the next epoch, further intensifying masking.
- If $s < \tau$, we set the auxiliary loss weight to zero, excluding the over-masked variant and halve the stride for the next epoch (minimum 2) to reduce the masking intensity.

This dynamic adjustment filters out samples with conflicting gradients, ensuring that only masked variants aligned with the primary objective contribute to training. The adaptive masking intensity regulation further allows us to strike a fine balance between exploring the model’s intra-modal potential and avoiding unproductive training. Ultimately, this combination enhances the models’ intra-modal abilities without compromising its core learning on the complete inputs.

Training Objective

Let \mathcal{L} be the loss for generating target keyphrases Y from the original, unmasked image-text inputs (X_V, X_T) ; we define two auxiliary losses to reinforce intra-modal learning. The first, $\tilde{\mathcal{L}}_V$, requires the model to generate the target Y from the masked visual features \tilde{X}_V and the complete text features X_T . The second, $\tilde{\mathcal{L}}_T$, is analogous, computed using (X_V, \tilde{X}_T) . The full objective combines the original loss with the two auxiliary losses, each controlled by a dynamic 0-1 switch λ_V and λ_T :

$$\mathcal{L}_{\text{total}} = \mathcal{L} + \lambda_V \tilde{\mathcal{L}}_V + \lambda_T \tilde{\mathcal{L}}_T. \quad (3)$$

We apply separate thresholds τ_V and τ_T for the image-masked and text-masked variants, the switches λ_V and λ_T are indicator functions determined by the gradient similarity score s and the threshold τ :

$$\lambda_V = \mathbf{1}\{s_V \geq \tau_V\}, \quad \lambda_T = \mathbf{1}\{s_T \geq \tau_T\},$$

where $\mathbf{1}\{\cdot\}$ is the indicator function (1 if true, 0 otherwise).

Experiment

Experiment Setup

Datasets Following previous studies (Wang et al. 2020; Dong et al. 2023), we carry out experiments on the CMKP dataset collected by (Wang et al. 2020). This dataset consists of 53,701 English tweets collected from Twitter, each containing a unique text-image pair with user-generated hashtags as keyphrases, and split into 8:1:1 train-val-test sets.

Evaluation Metrics Following (Wang et al. 2020; Dong et al. 2023), we adopt identical evaluation metrics: (i) **F1@K**: Macro-F1 score for the top-K keyphrase predictions, (ii) **MAP@K**: mean average precision on top-K predictions. For scenarios where the model generates $n < K$ keyphrases, we pad the remaining $(K - n)$ positions with empty labels \emptyset when computing F1 scores, and dynamically set $K' = \min(n, K)$ during the calculation of MAP scores. Notably, in sequential generation settings, predictions are ordered by their decoding sequence rather than confidence scores. Specifically, the order in which keyphrases are generated directly serves as their rank.

Baselines To validate the effectiveness of AimKP, we conduct a comprehensive comparison against a range of strong baselines. These baselines are divided into two main groups. First, we benchmark against existing models specifically designed for or adapted to MKP, namely **CO-ATT** (Zhang et al. 2017), **FLAVA** (Singh et al. 2022), **M³H-ATT** (Wang et al. 2020), **MM-MKP** (Dong et al. 2023), and text-only models adapted to MKP by leveraging image-associated text **BART-large** (Wolf et al. 2020) and **CopyBART** (Yu, Gao, and Zhang 2024). Second, and serving as our main comparison models, we include powerful MLLMs under standard fine-tuning, specifically **LLaVA-1.5-7B** (Liu et al. 2024) and **Qwen2-VL-7B** (Wang et al. 2024).

Implementation Details Our primary experimental setup is centered on the **LLaVA-1.5-7B** model, which we fine-tune using LoRA and optimized using Adam (Kingma and Ba 2015). In the final loss formulation, the original loss and the two auxiliary losses are equally weighted after filtering. We have also experimented with alternative weighting strategies, but observed no consistent improvement across settings. For the gradient-based filtering, we set modality-specific thresholds $\tau_V = 0.4$ and $\tau_T = 0.1$. To establish a foundational capability, we first train the model for one epoch on normal data, and then apply progressive modality masking and the gradient-based filtering. The training is conducted on four NVIDIA A6000 GPUs with a learning rate of $2e-4$, and a total batch size of 64. We perform validation at the end of each epoch and select the model

Models	F1@1	F1@3	MAP@5
<i>Image-only</i>			
LLaVA (image specialist)	41.02	25.64	45.66
LLaVA	33.57	21.17	37.68
LLaVA-AimKP	37.41	23.77	41.94
<i>Text-only</i>			
LLaVA (text specialist)	52.33	32.34	57.93
LLaVA	48.56	28.99	53.92
CopyBART	49.67	33.89	53.95
LLaVA-AimKP	50.04	30.63	55.45
<i>Multimodal</i>			
CO-ATT	42.12	31.55	48.39
FLAVA	46.05	31.23	49.30
M ³ H-ATT	47.06	33.11	52.07
MM-MKP	48.19	33.86	53.28
BART-large	50.47	34.69	55.11
CopyBART	51.42	36.54	57.35
LLaVA	61.58	37.90	68.07
LLaVA-AimKP	63.16	39.00	69.96
Qwen2-VL	63.08	38.43	69.89
Qwen2-VL-AimKP	64.18	38.73	71.00

Table 1: Performance on image-only, text-only, and multimodal inputs on the CMKP dataset. ‘‘Specialist’’ denotes models fine-tuned exclusively on a single modality. AimKP refers to models trained under our framework.

checkpoint that yields the best composite score on the validation set for final testing. All experiments are performed with three random seeds, and we report the averaged results. During inference, we decode outputs using beam search with sampling with a beam size of 5 and a temperature of 0.5, repeating the decoding process three times and taking the average of the results. Prior to evaluation, both the generated predictions and ground-truth keyphrases are stemmed with the Porter Stemmer (Porter 2006) and subsequently deduplicated. We also conduct additional experiments implementing AimKP on **Qwen2-VL-7B** to validate the effectiveness of our method across different architectures.

Main Results

General MKP Table 1 presents comparative results on the CMKP test dataset, yielding two key observations: First, standard fine-tuned MLLMs (e.g., LLaVA-1.5, Qwen2-VL) outperform small models by a significant margin. For instance, LLaVA-1.5 achieves 61.58% F1@1 (+10.16%) and 68.07% MAP@5 (+10.72%) over the strongest CopyBART baseline. This underscores MLLMs’ powerful inherent visual understanding and cross-modal integration capabilities, while highlighting the critical role of their pretrained cross-modal knowledge in MKP. Second, even on already high-performing MLLMs, AimKP yields consistent gains: LLaVA-1.5-AimKP improves F1@1 by 1.58% (63.16% vs. 61.58%) and MAP@5 by 1.89% (69.96% vs. 68.07%), while Qwen2-VL-AimKP gains 1.10% in F1@1 and 1.11%

Models	F1@1	F1@3	MAP@5
AimKP	63.16	39.00	69.96
w/o masking on text	62.41	38.37	69.12
w/o masking on image	62.74	38.66	69.43
w/o filtering	62.79	38.75	69.47
fixed masking ($\gamma = 2$)	63.11	38.91	69.93
fixed masking ($\gamma = 4$)	63.15	38.61	69.93

Table 2: Ablation study of AimKP’s key components, w/o denotes removing this component.

in MAP@5. These results validate that progressive masking combined with gradient-based filters strengthens MLLMs’ general keyphrase generation performance.

Intra-Modal Understanding Despite the gains on general MKP, AimKP also mitigates the significant gap remaining in single-modality performance. On text-only inputs, AimKP boosts MAP@5 from 53.92% to 55.45%, shrinking the deficit relative to the text specialist (57.93%) from 4.01 to 2.48 points. Similarly, on image-only inputs, LLaVA-1.5’s MAP@5 score rises from 37.68% to 41.94% with AimKP, reducing the gap to the image specialist (45.66%) from 7.98 to 3.72 points. These improvements, particularly the larger gains in image-only scenarios, demonstrate that AimKP enables the model to continuously deepen its intra-modal reasoning while alleviating its inherent textual bias.

Moreover, AimKP yields smaller gains on F1@3: only 1.1% on LLaVA-1.5 and 0.3% on Qwen2-VL. Even the text specialist (32.34%) fails to outperform CopyBART (33.89%) on text-only inputs. We hypothesize that stronger models tend to generate only high-certainty keyphrases, and incomplete information further worsens this, as MLLMs’ advanced capabilities make them more sensitive to such incompleteness. Since the F1@3 metric penalizes outputs with fewer than three keyphrases by padding with \emptyset tokens, these cautious behaviors result in lower scores.

Ablation Study

To validate the effectiveness and analyze the contributions of each component within AimKP, we conduct a series of ablation studies. The results are summarized in Table 2.

Effectiveness of Bimodal Masking Having established that our progressive masking strategy effectively improves performance, we now seek to verify the importance of applying this augmentation to both modalities. To this end, we conduct experiments where we disable masking on either the text or image inputs. The results in lines 3-4 reveal a significant performance degradation in both scenarios. Specifically, removing intra-modal augmentation of the image modality (w/o masking on text) or the text modality (w/o masking on image) causes a drop of 0.84 and 0.53 points in MAP@5, respectively. This difference aligns closely with our earlier observation of MLLMs’ inherent textual preference, which leaves image understanding relatively underdeveloped. The consistent performance drops when masking is removed from either modality confirm that our intra-modal

augmentation strategy is effective for both modalities, validating its ability to strengthen modality-specific semantics.

Effectiveness of Gradient-Based Filtering The gradient-based filtering is designed to prevent noisy or counterproductive updates from harming the training process. To validate its effectiveness, we disable it by setting the threshold $\tau = -1$, which accepts all masked variants regardless of their gradient similarity. Line 5 shows that this leads to a notable performance drop across all metrics (e.g., F1@1 decreases from 63.16% to 62.79%). This result indicates that naively applying aggressive masking can introduce samples with corrupted semantics that generate conflicting gradients. Our filter is crucial for identifying and discarding these harmful updates, thereby safeguarding the primary learning objective and ensuring stable performance gains.

Progressive Masking vs. Fixed Masking We also compare our progressive masking strategy against fixed masking ($\gamma = 2$ and $\gamma = 4$) throughout training. As shown in lines 6-7, our progressive approach consistently outperforms both fixed strategies across all metrics and exhibits greater stability. An easy ratio may not provide a sufficient training signal, while a consistently hard one can introduce excessive noise. Our progressive scheme dynamically adapts the difficulty, demonstrating the benefit of a curriculum-like approach.

Case Study

To further validate the effectiveness of AimKP, we compare its performance with the baseline models MM-MKP and LLaVA using selected cases in Figure 5.

As shown in Case (a), MM-MKP misidentifies the animated character ensemble as *Fire Emblem*, while LLaVA and LLaVA-AimKP recognize *Spider Man*-related themes, showcasing MLLMs’ edge in world knowledge utilization. LLaVA-AimKP further recognizes the specific character ensemble and contextual clues unique to *Spider Verse* (i.e., Spider Man Universe). Similarly, in Case (b), both LLaVA and LLaVA-AimKP link “vader costume” to *Star Wars* lore, unlike small models lacking such contextual awareness. However, LLaVA fails to deeply grasp intra-modal information and generates *May The 4th Be With You* (referring to a specific fan holiday), which is absent from the inputs.

Case (c) highlights LLaVA-AimKP’s fine-grained feature extraction. The inputs depict a community gathering for brand promotion, where the key cues lie in shirts with the “Black Rifle Coffee Company” (BRCC) logo in the image. LLaVA-AimKP, benefiting from enhanced intra-modal feature extraction, accurately identifies these brand-related visual cues and generates *America’s Coffee, BRCC*. In contrast, LLaVA outputs *Black Bull Whitetails*, a phrase clearly unrelated to any content. Furthermore, case (d) indicates that AimKP empowers the model to incorporate fine-grained features into more semantically complete predictions. The input counts down to the Racer football season opener, indicating 9 weeks until the game. While LLaVA superficially identifies *Racer Football* with the text, LLaVA-AimKP captures the underlying context and emotion using information from both modalities, generating the right keyphrases *Go Racers, Shoes Up* (the team’s spirit and motto).





	Case (a)	Case (b)	Case (c)	Case (d)
Image				
Text	<i>Another commission off to a happy home! This one was a bit daunting with all the characters but it was super fun drawing into this world again.</i>	<i>My girlfriend asked me to make a surprise appearance at her kindergarten class in my vader costume. My only requirement was that we take this photo.</i>	<i>Two great days out at the ranch with a great group of people. Stay tuned for some epic content.</i>	<i>Less than 9 weeks until we have racer football!</i>
Keyphrases	<i>Spider Verse</i>	<i>Star Wars</i>	<i>America's Coffee, BRCC</i>	<i>Go Racers, Shoes Up</i>
MM-MKP	<i>Fire Emblem</i>	<i>Game Of Thrones</i>	<i>Sas Who Dares Wins, Days</i>	<i>Super Bowl, Fly Eagles Fly</i>
LLaVA	<i>Spider Man</i>	<i>May The 4th Be With You</i>	<i>Black Bull Whitetails</i>	<i>Racer Football</i>
LLaVA-AimKP	<i>Spider Verse</i>	<i>Star Wars</i>	<i>America's Coffee, BRCC</i>	<i>Go Racers, Shoes Up</i>

Figure 5: Case study comparing keyphrase outputs of MM-MKP, LLaVA, and LLaVA-AimKP on four examples.

Related Work

Multimodal Large Language Models MLLMs (Alayrac et al. 2022; Li et al. 2023; OpenAI 2023; Liu et al. 2023) have demonstrated strong capabilities in visual question answering, image captioning, and cross-modal reasoning. Building on these foundations, recent works (Team et al. 2023; Zhu et al. 2024; Lan et al. 2025; Li et al. 2025) extend MLLMs to more complex visual tasks, improving fine-grained understanding (Zhang et al. 2024, 2025a) and reasoning over intricate scenes (Xiang et al. 2025).

Multimodal Keyphrase Generation Keyphrase generation (KPG) has been a significant area of focus within natural language processing (Zhuang et al. 2025). Current neural KPG models can be broadly categorized into three paradigms (Xie et al. 2023): (i) ONE2ONE (Chen et al. 2018; Meng et al. 2017; Chen et al. 2019), which converts a training sample containing multiple keyphrases into several training instances. Each instance pairs the source input with a single keyphrase. (ii) ONE2SEQ (Yuan et al. 2020; Chen et al. 2020; Kulkarni et al. 2022) treats KPG as a sequence-to-sequence task and concatenates all ground-truth keyphrases into a single target sequence according to a predefined order. (iii) ONE2SET (Ye et al. 2021; Xie et al. 2022; Shao et al. 2024), which models KPG as a set generation problem, generating keyphrases as an unordered set in parallel. While these works focus primarily on text-based KPG, a growing body of work is exploring the multimodal domain. Common approaches use co-attention networks to integrate text and visual information (Gong and Zhang 2016; Zhang et al. 2019). Wang et al. (2020) incorporates explicit optical characters and implicit image attributes from external tools, developing a MKP encoder-decoder model with multi-head attention mechanism. Dong et al. (2023) further enhances textual inputs with visual entities from external APIs and mitigates image noise through multi-granularity filtering.

Intra-Modal Augmentation Intra-modal augmentation addresses modality imbalance by strengthening models’

comprehension on each modality. A common approach frames this as multi-task learning, applying unimodal losses to each encoder. For instance, Self-MM (Yu et al. 2021) generates dynamic unimodal labels as auxiliary supervision; UMT (Du et al. 2021) employs a teacher–student framework, combining fusion loss and distillation loss to align each encoder with a unimodal teacher. Taking a different approach, EAU (Gao et al. 2024) learns unimodal representations by explicitly modeling data uncertainty during contrastive learning. Others focus on balancing the optimization process directly. OGM-GE (Peng et al. 2022) and MM-Pareto (Wei and Hu 2024) dynamically reweight gradients from the primary and auxiliary losses to correct imbalances.

To the best of our knowledge, we are the first to adapt MLLMs for MKP while mitigating modality bias and over-reliance on cross-modal shortcuts. Unlike existing methods which use static unimodal losses or focus on task-level gradient balancing, we introduce progressive modality masking to dynamically raise intra-modal difficulty, and sample-level gradient-based filtering to retain only informative masked samples. This ensures MLLMs build robust intra-modal understanding while maintaining cross-modal strengths.

Conclusion

In this paper, we introduce AimKP, a novel framework that addresses inherent modality bias and underdeveloped intra-modal understanding of MLLMs in MKP. AimKP leverages progressive modality masking to compel fine-grained feature extraction within each modality, and employs gradient-based filtering to remove uninformative masked samples, thereby stabilizing the training process. Extensive experiments and analyses demonstrate that AimKP not only strengthens MLLMs’ intra-modal capabilities but also achieves a new state-of-the-art in overall performance.

Future work will focus on adapting AimKP to other multimodal tasks where modality imbalance similarly degrades performance and investigating transition learning (Zhou et al. 2023) to better leverage diverse modalities.

Acknowledgements

The project was supported by National Key R&D Program of China (No.2022ZD0160501), Natural Science Foundation of Fujian Province of China (No.2024J011001), and the Public Technology Service Platform Project of Xiamen (No.3502Z20231043). We also thank the reviewers for their insightful comments.

References

- Alayrac, J.-B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millican, K.; et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35: 23716–23736.
- Chen, J.; Zhang, X.; Wu, Y.; Yan, Z.; and Li, Z. 2018. Keyphrase Generation with Correlation Constraints. In Riloff, E.; Chiang, D.; Hockenmaier, J.; and Tsujii, J., eds., *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, 4057–4066.
- Chen, W.; Chan, H. P.; Li, P.; and King, I. 2020. Exclusive Hierarchical Decoding for Deep Keyphrase Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, 1095–1105.
- Chen, W.; Gao, Y.; Zhang, J.; King, I.; and Lyu, M. R. 2019. Title-Guided Encoding for Keyphrase Generation. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI 2019)*, 6268–6275.
- Chiang, W.-L.; Li, Z.; Lin, Z.; Sheng, Y.; Wu, Z.; et al. 2023. Vicuna: An Open-Source Chatbot Impressing GPT-4 with 90%* ChatGPT Quality. <https://lmsys.org/blog/2023-03-30-vicuna/>. Accessed: 2025-01-05.
- Dong, Y.; Wu, S.; Meng, F.; Zhou, J.; Wang, X.; Lin, J.; and Su, J. 2023. Towards Better Multi-modal Keyphrase Generation via Visual Entity Enhancement and Multi-granularity Image Noise Filtering. In *Proceedings of the 31st ACM International Conference on Multimedia, MM '23*, 3897–3907.
- Du, C.; Li, T.; Liu, Y.; Wen, Z.; Hua, T.; Wang, Y.; and Zhao, H. 2021. Improving Multi-Modal Learning with Uni-Modal Teachers. *CoRR*, abs/2106.11059.
- Gao, Z.; Jiang, X.; Xu, X.; Shen, F.; Li, Y.; and Shen, H. T. 2024. Embracing Unimodal Aleatoric Uncertainty for Robust Multimodal Fusion. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, 26866–26875.
- Gong, Y.; and Zhang, Q. 2016. Hashtag Recommendation Using Attention-Based Convolutional Neural Network. In Kambhampati, S., ed., *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, 2782–2788.
- Ilharco, G.; Wortsman, M.; Wightman, R.; Gordon, C.; Carlini, N.; Taori, R.; Dave, A.; Shankar, V.; et al. 2021. OpenCLIP. <https://doi.org/10.5281/zenodo.5143773>. Accessed: 2025-01-05.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Kulkarni, M.; Mahata, D.; Arora, R.; and Bhowmik, R. 2022. Learning Rich Representation of Keyphrases from Text. In Carpuat, M.; de Marneffe, M.; and Ruíz, I. V. M., eds., *Findings of the Association for Computational Linguistics: NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, 891–906.
- Lan, Z.; Niu, L.; Meng, F.; Li, W.; Zhou, J.; and Su, J. 2025. AVG-LLaVA: An Efficient Large Multimodal Model with Adaptive Visual Granularity. In Che, W.; Nabende, J.; Shutova, E.; and Pilehvar, M. T., eds., *Findings of the Association for Computational Linguistics, ACL 2025, Vienna, Austria, July 27 - August 1, 2025*, 16852–16869.
- Li, F.; Zhang, R.; Zhang, H.; Zhang, Y.; Li, B.; Li, W.; Ma, Z.; and Li, C. 2025. LLaVA-Interleave: Tackling Multi-image, Video, and 3D in Large Multimodal Models. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. C. H. 2023. BLIP-2: Bootstrapping Language-Image Pre-training with Frozen Image Encoders and Large Language Models. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, 19730–19742.
- Liu, H.; Li, C.; Li, Y.; and Lee, Y. J. 2024. Improved Baselines with Visual Instruction Tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 26296–26306.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual Instruction Tuning. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 34892–34916.
- Meng, R.; Zhao, S.; Han, S.; He, D.; Brusilovsky, P.; and Chi, Y. 2017. Deep Keyphrase Generation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, 582–592.
- OpenAI. 2023. GPT-4 Technical Report. *CoRR*, abs/2303.08774.
- Parcalabescu, L.; and Frank, A. 2025. Do Vision & Language Decoders use Images and Text equally? How Self-consistent are their Explanations? In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*.
- Peng, X.; Wei, Y.; Deng, A.; Wang, D.; and Hu, D. 2022. Balanced multimodal learning via on-the-fly gradient modulation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8238–8247.
- Porter, M. F. 2006. An algorithm for suffix stripping. *Program*.

- Qwen; ; Yang, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Li, C.; et al. 2024. Qwen2.5 Technical Report. arXiv:2412.15115.
- Shao, L.; Zhang, L.; Peng, M.; Ma, G.; Yue, H.; Sun, M.; and Su, J. 2024. One2Set + Large Language Model: Best Partners for Keyphrase Generation. In *AI-Onaizan, Y.; Bansal, M.; and Chen, Y., eds., Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, 11140–11153.
- Singh, A.; Hu, R.; Goswami, V.; Couairon, G.; Galuba, W.; Rohrbach, M.; and Kiela, D. 2022. FLAVA: A Foundational Language And Vision Alignment Model. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, 15617–15629.
- Team, G.; Anil, R.; Borgeaud, S.; Alayrac, J.-B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A. M.; Hauth, A.; Millican, K.; et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Wang, P.; Bai, S.; Tan, S.; Wang, S.; Fan, Z.; Bai, J.; Chen, K.; Liu, X.; et al. 2024. Qwen2-VL: Enhancing Vision-Language Model’s Perception of the World at Any Resolution. arXiv:2409.12191.
- Wang, Y.; Li, J.; Lyu, M.; and King, I. 2020. Cross-Media Keyphrase Prediction: A Unified Framework with Multi-Modality Multi-Head Attention and Image Wordings. In *Webber, B.; Cohn, T.; He, Y.; and Liu, Y., eds., Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 3311–3324. Online.
- Wei, Y.; and Hu, D. 2024. MMPareto: Boosting Multimodal Learning with Innocent Unimodal Assistance. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; et al. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Liu, Q.; and Schlangen, D., eds., Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, EMNLP 2020 - Demos, Online, November 16-20, 2020*, 38–45.
- Xiang, Z.; Wu, C.; Zhang, Q.; Chen, S.; Hong, Z.; Huang, X.; and Su, J. 2025. When to use graphs in rag: A comprehensive analysis for graph retrieval-augmented generation. *arXiv preprint arXiv:2506.05690*.
- Xie, B.; Song, J.; Shao, L.; Wu, S.; Wei, X.; Yang, B.; Lin, H.; Xie, J.; and Su, J. 2023. From statistical methods to deep learning, automatic keyphrase prediction: A survey. *Inf. Process. Manag.*, 60(4): 103382.
- Xie, B.; Wei, X.; Yang, B.; Lin, H.; Xie, J.; Wang, X.; Zhang, M.; and Su, J. 2022. WR-One2Set: Towards Well-Calibrated Keyphrase Generation. In *EMNLP*.
- Ye, J.; Gui, T.; Luo, Y.; Xu, Y.; and Zhang, Q. 2021. One2Set: Generating Diverse Keyphrases as a Set. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, 4598–4608.
- Yu, B.; Gao, C.; and Zhang, S. 2024. Training with One2MultiSeq: CopyBART for social media keyphrase generation. *J. Supercomput.*, 80(11): 15517–15544.
- Yu, W.; Xu, H.; Yuan, Z.; and Wu, J. 2021. Learning modality-specific representations with self-supervised multi-task learning for multimodal sentiment analysis. In *Proceedings of the AAAI conference on artificial intelligence*, 12, 10790–10797.
- Yuan, X.; Wang, T.; Meng, R.; Thaker, K.; Brusilovsky, P.; He, D.; and Trischler, A. 2020. One Size Does Not Fit All: Generating and Evaluating Variable Number of Keyphrases. In *ACL*.
- Zhang, Q.; Dong, J.; Chen, H.; Zha, D.; Yu, Z.; and Huang, X. 2024. Knowgpt: Knowledge graph based prompting for large language models. *Advances in Neural Information Processing Systems*, 37: 6052–6080.
- Zhang, Q.; Wang, J.; Huang, H.; Huang, X.; and Gong, Y. 2017. Hashtag Recommendation for Multimodal Microblog Using Co-Attention Network. In *Sierra, C., ed., Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, 3420–3426.
- Zhang, Q.; Xiang, Z.; Xiao, Y.; Wang, L.; Li, J.; Wang, X.; and Su, J. 2025a. FaithfulRAG: Fact-Level Conflict Modeling for Context-Faithful Retrieval-Augmented Generation. *arXiv preprint arXiv:2506.08938*.
- Zhang, S.; Yao, Y.; Xu, F.; Tong, H.; Yan, X.; and Lu, J. 2019. Hashtag Recommendation for Photo Sharing Services. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, 5805–5812.
- Zhang, Y.; Ma, J.; Hou, Y.; Bai, X.; Chen, K.; Xiang, Y.; Yu, J.; and Zhang, M. 2025b. Evaluating and Steering Modality Preferences in Multimodal Large Language Model. *CoRR*, abs/2505.20977.
- Zhang, Z.; Tang, H.; Sheng, J.; Zhang, Z.; Ren, Y.; et al. 2025c. Debiasing Multimodal Large Language Models via Noise-Aware Preference Optimization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025*, 9423–9433.
- Zheng, X.; Liao, C.; Fu, Y.; Lei, K.; Lyu, Y.; Jiang, L.; Ren, B.; Chen, J.; Wang, J.; et al. 2025. MLLMs are Deeply Affected by Modality Bias. *CoRR*, abs/2505.18657.
- Zhou, C.; Liang, Y.; Meng, F.; Zhou, J.; Xu, J.; Wang, H.; Zhang, M.; and Su, J. 2023. A Multi-Task Multi-Stage Transitional Training Framework for Neural Chat Translation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(7): 7970–7985.
- Zhu, D.; Chen, J.; Shen, X.; Li, X.; and Elhoseiny, M. 2024. MiniGPT-4: Enhancing Vision-Language Understanding with Advanced Large Language Models. In *ICLR*.
- Zhuang, L.; Chen, S.; Xiao, Y.; Zhou, H.; Zhang, Y.; Chen, H.; Zhang, Q.; and Huang, X. 2025. LinearRAG: Linear Graph Retrieval Augmented Generation on Large-scale Corpora. *arXiv preprint arXiv:2510.10114*.