

MHA2MLA-VLM: Enabling DeepSeek’s Economical Multi-Head Latent Attention across Vision-Language Models

Xiaoran Fan^{1*}, Zhichao Sun^{1*}, Tao Ji^{1*}, Lixing Shen², Tao Gui^{1,3,4}

¹Fudan University

²Hikvision Inc

³Shanghai Innovation Institute

⁴Pengcheng Laboratory

{xrfan23, taoji, tgui}@fudan.edu.cn

Abstract

As vision-language models (VLMs) tackle increasingly complex and multimodal tasks, the rapid growth of Key-Value (KV) cache imposes significant memory and computational bottlenecks during inference. While Multi-Head Latent Attention (MLA) offers an effective means to compress the KV cache and accelerate inference, adapting existing VLMs to the MLA architecture without costly pretraining remains largely unexplored. In this work, we present **MHA2MLA-VLM**, a parameter-efficient and multimodal-aware framework for converting off-the-shelf VLMs to MLA. Our approach features two core techniques: (1) a modality-adaptive partial-RoPE strategy that supports both traditional and multimodal settings by selectively masking nonessential dimensions, and (2) a modality-decoupled low-rank approximation method that independently compresses the visual and textual KV spaces. Furthermore, we introduce parameter-efficient fine-tuning to minimize adaptation cost and demonstrate that minimizing output activation error, rather than parameter distance, substantially reduces performance loss. Extensive experiments on three representative VLMs show that MHA2MLA-VLM restores original model performance with minimal supervised data, significantly reduces KV cache footprint, and integrates seamlessly with KV quantization.

Code & Appendix —

<https://github.com/JT-Ushio/MHA2MLA-VLM>

1 Introduction

The Key-Value (KV) cache stores the complete contextual information required by large language models (LLMs), enabling efficient and accurate decoding of the current token. As the tasks handled by LLMs become increasingly complex (e.g. multimodal tasks (Bordes et al. 2024) and deep thinking (Pan et al. 2025)), the context length correspondingly increases. This results in a rapid expansion of the KV cache, which not only occupies large GPU memory but also leads to severe memory access bottlenecks due to the quadratic complexity of the standard attention mechanism (Keles, Wijewardena, and Hegde 2023). Consequently, efficient inference in LLMs, especially in vision-language models (VLMs)

with multimodal contexts, urgently requires cost-effective KV cache management and attention architectures.

A series of studies have identified redundancies in the KV cache (Li et al. 2025). In terms of sequence length (Zhang et al. 2023b; Oren et al. 2024), KV cache pruning removes irrelevant tokens from the cache. Regarding representation precision (Badri and Shaji 2023), KV cache quantization reduces the precision of vector representations. In the vector dimension, modifications such as Grouped/Multi-Query Attention (GQA and MQA) restructure the attention mechanism by enabling a single KV pair to be shared among a group of queries (Ainslie et al. 2023a; Shazeer 2019).

DeepSeek introduced Multi-Head Latent Attention (MLA), an advanced attention mechanism employing low-rank key-value joint compression (DeepSeek-AI et al. 2024). Empirical results show that MLA outperforms standard Multi-Head Attention (MHA, 2017) and its variants, while significantly reducing the KV cache size during inference, thereby enhancing inference efficiency. Ji et al. (2025) proposed MHA2MLA, demonstrating that LLMs originally trained with MHA/GQA can be adapted to leverage MLA during inference. However, whether VLMs can undergo a similar transition to the MLA architecture remains an open question.

MHA2MLA involves two key steps: partial-rope conversion and KV joint low-rank approximation. For partial-rope, text-only LLMs have demonstrated that significantly reducing (e.g., -87.5%) less important rotary frequencies requires only minimal fine-tuning to recover performance. In the case of VLMs, it is necessary to verify whether the retained rotary frequencies are equally effective for both image and text tokens. We address this question and further extend the method to multimodal rope (e.g., used in Qwen2-VL series).

For KV joint low-rank approximation, inspired by SVDLLM V2 (Wang et al. 2025a), we improve the approximation from being applied to parameters ($\min \|W - W'\|_F$, where W' is the low-rank approximation) to output activations ($\min \|XW - XW'\|_F$). This enhancement significantly reduces performance degradation and the amount of fine-tuning data required. Moreover, we observe that the low-rank spaces of image and text tokens are orthogonal, necessitating separate low-rank approximations for each modality.

To reduce the cost of MHA2MLA-VLM adaptation, we introduce parameter-efficient fine-tuning (PEFT, 2023). During the partial-rope phase, only the two projection matrices

* Equal contribution.

✉ Corresponding authors are Tao Ji and Tao Gui.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

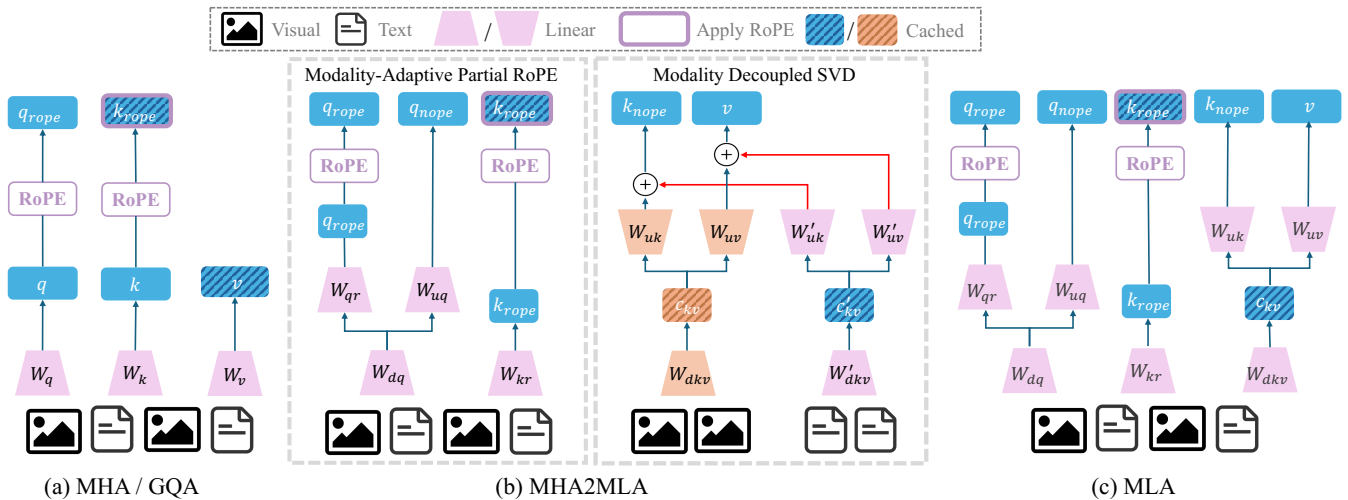


Figure 1: The overview process of converting VLMs from MHA/GQA to MLA using MHA2MLA-VLM. Our method makes the attention inputs match MLA exactly, and low rank compression of the KV cache is consistent with MLA. The modality-decoupled design reduces truncation loss and maximizes the parameter reuse of pretrained weights.

for query and key are fine-tuned, while all other parameters are frozen. For the low-rank approximation phase, only the parameters within MLA are fine-tuned. It reduces the time required by 59% (e.g., the MHA2MLA-VLM of Qwen2.5-VL is shortened from 22 hours to 9 hours). We validate the effectiveness of MHA2MLA-VLM on three representative models: LLaVA-1.5 (2024a), LLaVA-NeXT (2024b), and Qwen2.5-VL (2025). Furthermore, using LLaVA-NeXT, we demonstrate that MLA outperforms the KV cache pruning baseline and integrates seamlessly with KV quantization.

Our main contributions are:

- We successfully extend the MHA2MLA adaptation from text-only LLMs to VLMs, designing multimodal partial-rope and low-rank approximation algorithms.
- By incorporating SVDLLM V2’s minimization of output activation error and introducing PEFT, we significantly reduce the performance degradation and fine-tuning cost.
- We demonstrate the effectiveness of MHA2MLA-VLM in three main VLMs with distinct architectures and demonstrate that it integrates seamlessly with KV quantization.

2 MHA2MLA-VLM

Figure 1 provides an overview of MHA2MLA-VLM, we will describe the details of the two main components of MHA2MLA-VLM: (1) a modality-adaptive partial-RoPE strategy that supports both traditional and multimodal settings by selectively masking nonessential dimensions, and (2) a modality-decoupled low-rank approximation method that independently compresses the visual and textual KV spaces.

2.1 Multimodal Partial-RoPE

To enable efficient migration of MHA/GQA-based VLMs to MLA, we introduce Multimodal Adaptive Partial-RoPE, which adaptively retains the most informative rotary dimensions according to the nature of the input from different modalities, achieving efficient architecture migration.

Current research on partial-RoPE has been limited to unimodal settings. For example, studies such as (Black et al. 2021; Barbero et al. 2025) trained partial-RoPE models from scratch, achieving slightly better perplexity compared to full-RoPE. More recent work (Ji et al. 2025) has explored adapting pre-trained full-RoPE models to partial-RoPE without costly retraining. However, their analyzes are strictly limited to LLM. In multimodal scenarios, for the input, visual and textual information is interleaved; For the VLMs’ forward calculation, visual and text information are jointly entangled in the RoPE dimension. Simply applying the text-based partial RoPE strategy leads to suboptimal allocation of the retention frequency subspace, as visual and textual information exhibit distinct dimensional characteristics are ignored. To overcome these limitations, retain the most informative rotation dimension based on modality-awareness, thereby enabling low-cost and efficient architectural transfer from MHA/GQA to MLA.

Full Vanilla RoPE is a mechanism (Su et al. 2024) for encoding positional information into queries and keys through frequency-specific rotations. Formally, given a query $\mathbf{q}_i \in \mathbb{R}^{d_h}$ and a key $\mathbf{k}_i \in \mathbb{R}^{d_h}$, we split them into 2D chunks:

$$\mathbf{q}_i, \mathbf{k}_i = \left[\mathbf{q}_i^{[2k, 2k+1]} \right]_{0 \leq k < \frac{d_h}{2}}, \left[\mathbf{k}_i^{[2k, 2k+1]} \right]_{0 \leq k < \frac{d_h}{2}}.$$

Formally, for each 2D chunk $\mathbf{q}_i^{[2k, 2k+1]}$ and $\mathbf{k}_i^{[2k, 2k+1]}$, the rotation matrix at position i is defined as:

$$\mathbf{R}_i^{[2k, 2k+1]}(\theta_k) = \begin{bmatrix} \cos(i\theta_k) & -\sin(i\theta_k) \\ \sin(i\theta_k) & \cos(i\theta_k) \end{bmatrix},$$

where $\theta_k = \beta^{-2k/d_h}$ is the frequency of rotation applied to a specific k -th pair of $\mathcal{G}_d \in [0, \frac{d_h}{2})$, and β is the frequency base wavelength. The vanilla RoPE defines a matrix \mathbf{A}_{t_i, t_j} that represents the relative positional encoding between two positions t_i and t_j in a 1D sequence:

$$\mathbf{A}_{t_i, t_j} = (\mathbf{q}_{t_i} \mathbf{R}_{t_i}) (\mathbf{k}_{t_j} \mathbf{R}_{t_j})^\top = \mathbf{q}_{t_i} \mathbf{R}_{\Delta t} \mathbf{k}_{t_j}^\top,$$

where $\Delta t = t_i - t_j$, \mathbf{q}_{t_i} and \mathbf{k}_{t_j} are query and key vectors at positions t_i and t_j , the relative rotation matrix is $\mathbf{R}_{\Delta t}$.

Full Multimodal RoPE There are two common approaches to extending RoPE to multimodal LLMs. One approach directly applies standard RoPE by flattening the visual tokens and treating both text and visual tokens as a single 1D sequence (Liu et al. 2024a; Zhu et al. 2025; Bai et al. 2023). The other considers the unique characteristics of each modality and extends RoPE to multimodal scenarios, such as M-RoPE (Wang et al. 2024; Bai et al. 2025; Wei et al. 2025).

Unlike the vanilla 1D-RoPE in LLMs, which is limited to encoding one-dimensional positional information, M-RoPE effectively models the positional information of multimodal inputs. This is achieved by deconstructing the original rotary embedding into three orthogonal components: temporal (t), height (h), and width (w). For text, these components utilize identical position IDs, making M-RoPE functionally equivalent to standard 1D-RoPE. For images, the temporal of each visual index is held constant, while height and width IDs are assigned distinctly based on the token’s position in the image. For videos, treated as frame sequences, the temporal ID increases with each frame, while the height and width IDs follow the same assignment pattern as for images.

Formally, given a query vector $\mathbf{q}_i \in \mathbb{R}^{d_h}$ and key vector $\mathbf{k}_i \in \mathbb{R}^{d_h}$, we partition them into per-modality components:

$$\mathbf{q}_i = [\mathbf{q}_i^{[t]}; \mathbf{q}_i^{[h]}; \mathbf{q}_i^{[w]}], \mathbf{k}_i = [\mathbf{k}_i^{[t]}; \mathbf{k}_i^{[h]}; \mathbf{k}_i^{[w]}]$$

For each components, the embeddings are rotated separately by using its corresponding 2D rotary position encoding: $\mathcal{G}_t \in [0, \frac{d_h}{8})$, $\mathcal{G}_h \in [\frac{d_h}{8}, \frac{5d_h}{16})$, $\mathcal{G}_w \in [\frac{5d_h}{16}, \frac{d_h}{2})$. The query embeddings after applying M-RoPE are computed as follows:

$$\begin{aligned} \mathbf{q}_{i,rope}^{[t]} &= [\mathbf{R}_{p_t}^{[2k, 2k+1]}(\theta_k) \mathbf{q}_i^{[2k, 2k+1]}]_{k \in \mathcal{K}_t}, \\ \mathbf{q}_{i,rope}^{[h]} &= [\mathbf{R}_{p_h}^{[2k, 2k+1]}(\theta_k) \mathbf{q}_i^{[2k, 2k+1]}]_{k \in \mathcal{K}_h}, \\ \mathbf{q}_{i,rope}^{[w]} &= [\mathbf{R}_{p_w}^{[2k, 2k+1]}(\theta_k) \mathbf{q}_i^{[2k, 2k+1]}]_{k \in \mathcal{K}_w}. \end{aligned}$$

Thus, applying M-RoPE to both queries and keys becomes:

$$\begin{aligned} \mathbf{q}_{i,rope} &= [\mathbf{q}_{i,rope}^{[t]}; \mathbf{q}_{i,rope}^{[h]}; \mathbf{q}_{i,rope}^{[w]}] \in \mathbb{R}^{d_h}, \\ \mathbf{k}_{i,rope} &= [\mathbf{k}_{i,rope}^{[t]}; \mathbf{k}_{i,rope}^{[h]}; \mathbf{k}_{i,rope}^{[w]}] \in \mathbb{R}^{d_h}. \end{aligned}$$

Let $\mathbf{q}_i, \mathbf{k}_i \in \mathbb{R}^{d_h}$ be the query and key vectors at position i . The corresponding relative matrix \mathbf{A} is computed as:

$$\mathbf{A}_{(t_i, h_i, w_i), (t_j, h_j, w_j)} = \mathbf{q}_{(t_i, h_i, w_i)} \mathbf{R}_{\Delta t, \Delta h, \Delta w} \mathbf{k}_{(t_j, h_j, w_j)}^\top,$$

where $\Delta t = t_i - t_j$, $\Delta h = h_i - h_j$, and $\Delta w = w_i - w_j$.

Multimodal Adaptive Partial-RoPE Strategies Given r retained rotational subspaces ($r = \frac{d_r}{2} \ll \text{total subspaces } \frac{d_h}{2}$), the aim is to select which r subspaces preserve RoPE/M-RoPE encoding.

Recent research (Ji et al. 2025) has systematically compared four heuristic partial RoPE methods: High-Frequency Preservation, Low-Frequency Preservation, Uniform Sampling, and Head-wise 2-norm contribution. Experiments showed that the head-wise 2-norm contribution performs better. Specifically, for each head h , Head-wise 2-norm contribution computes the mean 2-norm score for each subspace in an LLM over long sequences. Then rank all subspaces by their 2-norm score and select the top- r :

$$\mathcal{S}_{2\text{-norm}} = \text{top-}r \left(\left\| \mathbf{q}_*^{[2k, 2k+1]} \right\| \left\| \mathbf{k}_*^{[2k, 2k+1]} \right\| \right)_{0 \leq k < \frac{d_h}{2}}. \quad (1)$$

The above method uses heuristic calculations to remove unimportant subspaces and retain r important ones. It cannot effectively reflect the impact of removing specific dimensions on the original whole. Moreover, to enable the migration of VLMs from MHA/GQA to MLA, We propose Contribution-Aware Multimodal Partial-RoPE, based on KL-divergence (MKL). A data-driven and training-free strategy that extends frequency-subspace selection to multimodal inputs.

For each layer l and each attention head h we compute the *frequency-wise KL sensitivity*:

$$\mathcal{I}^{l, h, k} = \mathbb{E}_{\mathcal{D}} \left[\text{KL}(\mathbf{P}_{l, h}^{\text{full}} \parallel \mathbf{P}_{l, h, k}^{\text{masked}}) \right], \quad (2)$$

where $\mathbf{P}_{\text{full}}^{l, h}$ denotes the attention distribution produced by the original full RoPE/M-RoPE model and $\mathbf{P}_{l, h, k}^{\text{masked}}$ is obtained after zero-ablating the k -th subspace in the query and key projections of head h . A large $\mathcal{I}^{l, h, k}$ indicates that subspace k is critical for positional understanding under the current multimodal inputs. The $d_h/2$ subspaces are then ranked in descending order and the top- r indices are retained:

$$\mathcal{S}_{\text{MKL}}^{l, h} = \text{top-}r \mathcal{I}^{l, h, k} \quad_{0 \leq k < \frac{d_h}{2}}. \quad (3)$$

Thus, our modality-adaptive strategy preserves rotation-critical subspaces. Table 4 shows that our Multimodal Adaptive Partial-RoPE strategy (MKL) outperforms the strongest baseline ($\mathcal{S}_{2\text{-norm}}$). We will analyze the effectiveness of the strategy in Section 4.2. Based on our proposed method, non-selected subspaces ($k \notin \mathcal{S}$) become NoPE dimensions, enabling seamless integration with MLA’s latent compression.

2.2 Modality-Decoupled SVD (MD-SVD)

After transforming the VLMs’ Vanilla RoPE/M-RoPE to modality-adaptive partial RoPE, we get the first component \mathbf{k}_{rope} . The next step is to construct $\mathbf{c}_{kv}^{(m)} \in \mathbb{R}^{d_{kv}}$, a low-rank joint embedding of the modality-specific $\mathbf{k}_{nope}^{(m)}$ and $\mathbf{v}^{(m)}$.

Unimodal SVD Baselines Studies on SVD-driven LLM compression can be grouped into two paradigms. The first paradigm operates directly on the models’ weights (Hsu et al. 2022) reduces truncation loss by estimating weight importance and preserving more important weights. (Ji et al. 2025) jointly optimizes the latent space for both keys and values, which validates that joint factorization better preserves pre-trained knowledge. The second paradigm (Wang et al.

Algorithm 1: Pseudocode of Modality-Decoupled SVD

Input: W — Joint KV weight matrix
 $\mathbf{X}_{visual}, \mathbf{X}_{text}$ — Visual / Text activations
 r_{visual}, r_{text} — Target ranks
Output: $\{\mathbf{W}_m^{up}, \mathbf{W}_m^{down}\}_{m \in \{visual, text\}}$

- 1: **procedure** MD_SVD($W, \mathbf{X}_{visual}, \mathbf{X}_{text}, r_{visual}, r_{text}$)
- 2: **for** $m \in \{Visual, Text\}$ **do** ▷ process each modality
- 3: $\mathbf{S}_m \leftarrow \mathbf{X}_m \mathbf{X}_m^\top$
- 4: $[\mathbf{U}_s, \Sigma_s, _] \leftarrow \text{SVD}(\mathbf{S}_m)$
- 5: $\mathbf{D} \leftarrow W \mathbf{U}_s \Sigma_s^{1/2}$
- 6: $[\mathbf{U}_d, \Sigma_d, \mathbf{V}_d] \leftarrow \text{SVD}(\mathbf{D})$
- 7: Keep first r_m components of $(\mathbf{U}_d, \Sigma_d, \mathbf{V}_d)$
- 8: $\mathbf{W}_m^{up} \leftarrow \mathbf{U}_d \Sigma_d^{1/2}$
- 9: $\mathbf{W}_m^{down} \leftarrow \Sigma_d^{1/2} \mathbf{V}_d \Sigma_s^{-1/2} \mathbf{U}_s^{-1}$
- 10: **end for**
- 11: **return** $\{\mathbf{W}_m^{up}, \mathbf{W}_m^{down}\}_m$
- 12: **end procedure**

2025b,a) augments SVD with activation-aware transformations to lower the truncation loss L in the form of Frobenius norm as follows during LLM compression:

$$\mathcal{L}^2 = \|\mathbf{W}\mathbf{X} - \mathbf{W}'\mathbf{X}\|_F^2 \quad (4)$$

Given a single-modality activation matrix \mathbf{X} , we first construct the covariance $\mathbf{S} = \mathbf{X}\mathbf{X}^\top$, second, we perform a second SVD on $\mathbf{D} = \mathbf{W}\mathbf{U}_s \Sigma_s^{1/2}$. (Wang et al. 2025a) proves the compressed weight $\mathbf{W}' = \mathbf{U}_d \text{Trunc.}(\Sigma_d) \mathbf{V}_d \Sigma_s^{-1/2} \mathbf{U}_s^{-1}$ achieves the theoretical minimum truncation loss.

Modality-Decoupled SVD

Motivation Visual and text activations exhibit distinct scales, estimating a single covariance on \mathbf{S}_{visual} or \mathbf{S}_{text} causes the dominant modality to distort the singular value distribution, reducing the quality of other modality after SVD. Our aim is to keep the *shared* KV weight W intact while deriving two modality-decoupled low-rank projections that jointly minimise truncation loss. The pseudocode of our proposed MD-SVD is provided in Algorithm 1.

Theorem 2.1. *For a multimodal token sequence, we denote the joint activation matrix as $\mathbf{X}_{joint} \in \mathbb{R}^{d \times (n_v + n_t)}$ which can be partitioned into two modality-specific components:*

$$\mathbf{X}_{joint} = [\mathbf{X}_{visual}; \mathbf{X}_{text}]$$

where $\mathbf{X}_{visual} \in \mathbb{R}^{d \times n_v}$ and $\mathbf{X}_{text} \in \mathbb{R}^{d \times n_t}$ denote the activations corresponding to visual and text tokens, respectively.

Then the minimum loss of joint-modals is larger than or equal to the sum of the minimum losses of split-modals:

$$\min \mathcal{L}_{joint}^2 \geq \min \mathcal{L}_{visual}^2 + \min \mathcal{L}_{text}^2 \quad (5)$$

Proof. First, Given \mathbf{X}_{joint} and weight matrix \mathbf{W} , inspired by (Wang et al. 2025a), we can obtain an optimal matrix \mathbf{W}' that minimizes the joint loss \mathcal{L}_{joint}^2 :

$$\min \mathcal{L}_{joint}^2 = \|\mathbf{W}\mathbf{X}_{joint} - \mathbf{W}'\mathbf{X}_{joint}\|_F^2 \quad (6)$$

Second, by decomposing \mathbf{X}_{joint} into its constituent modalities, we obtain the following:

$$\begin{aligned} \min \mathcal{L}_{joint}^2 &= \|\mathbf{W}[\mathbf{X}_{visual}; \mathbf{X}_{text}] \\ &\quad - \mathbf{W}'[\mathbf{X}_{visual}; \mathbf{X}_{text}]\|_F^2 \\ &= \|\mathbf{W}\mathbf{X}_{visual} - \mathbf{W}'\mathbf{X}_{visual}\|_F^2 \\ &\quad + \|\mathbf{W}\mathbf{X}_{text} - \mathbf{W}'\mathbf{X}_{text}\|_F^2 \\ &\geq \min \mathcal{L}_{visual}^2 + \min \mathcal{L}_{text}^2. \end{aligned} \quad (7)$$

This inequality holds because joint optimization imposes shared weights between different modalities, which limits the upper bound of model optimization. In contrast, our proposed Modality-Decoupled SVD (MD-SVD) optimizations allow for separate weights, enabling each branch to independently minimize its truncation loss. We provide a quantitative analysis in Section 4.1 to further validate this theoretical insight. \square

3 Experiment

Setups Details of the models, datasets, parameter-efficient and data-efficient strategies, evaluation, and hyperparameter setups are placed in Appendix A.

Based on the above settings, We conduct systematic experiments of MHA2MLA and GQA2MLA under various KV dimensions, our experiments address three critical questions:

1. Can MHA2MLA-VLM maintain multimodal accuracy when training and inference are limited to a tight computation or data budget?
2. What are the characteristics of SVD in multimodal scenarios using MHA2MLA-VLM?
3. Comparison between MHA2MLA-VLM and cache pruning, and can it be combined with cache compression?

3.1 Main Results

As shown in Table 1, compared to the original VLMs after fine-tuning ($d_{kv} = 256$), as d_{kv} decreases, our method significantly reduces the KV cache memory with only slight performance drop. For example, for Qwen2.5-VL, our GQA2MLA method still achieves an overall performance of 79.47 even after reducing the KV memory by 94.64%, which is comparable to the original model or the fine-tuned baseline model.

Second, our architecture migration method demonstrates both parameter and data efficiency. Compared to existing models training from scratch that rely on trillions of training tokens (Bai et al. 2025), our approach enables the migration from MHA to MLA architectures using only within 1.8B tokens. Furthermore, within our two-stage training framework, we fine-tune only $\sim 10\%$ of the VLM’s parameters while still achieving competitive performance. These results highlight the effectiveness of our method in enabling architectural adaptation with limited data or GPU resources, making it highly applicable to resource-constrained scenarios.

In addition, Figure 2 shows the parameter and data efficient fine-tuning losses of MHA2MLA and GQA2MLA in different compression ratios. Even with small data (0.002%) and only $\sim 10\%$ parameter updates, training converges quickly.

Model	Tokens	KV Mem.	Avg.	AI2D	GQA	POPE	SEED ¹	RWQ	MMB ^{EN}	Chart	Doc
LLaVA-1.5 _{7B}			63.89	53.63	61.27	84.03	65.28	56.21	62.89	-	-
- MHA	$d_{kv}=256$		64.13	54.92	61.97	83.96	64.80	56.34	62.80	-	-
	$d_{kv}=64$	0.5B	-62.50%	63.84 _{-0.29}	54.05	62.32	82.74	65.17	54.51	64.26	-
- MHA2MLA	$d_{kv}=32$	(0.025%)	-75.00%	63.58 _{-0.55}	52.49	62.06	83.18	65.04	54.90	63.83	-
	$d_{kv}=16$		-81.30%	62.27 _{-1.86}	48.41	62.08	83.61	65.85	51.37	62.29	-
LLaVA-NeXT _{8B}			70.72	70.56	65.13	87.18	72.46	58.69	71.82	68.44	71.47
- GQA	$d_{kv}=256$		70.78	70.76	65.30	87.22	71.95	59.22	71.74	69.04	70.99
	$d_{kv}=128$	1.8B	-84.38%	70.23 _{-0.55}	69.53	65.25	85.96	71.90	59.08	72.08	68.68
- GQA2MLA	$d_{kv}=64$	(0.012%)	-90.63%	68.75 _{-2.03}	68.20	64.45	86.32	71.57	58.56	69.50	65.56
	$d_{kv}=32$		-93.75%	66.72 _{-4.06}	65.90	63.98	86.56	71.15	56.86	63.32	63.60
Qwen2.5-VL _{7B}			79.47	82.58	60.42	86.22	77.34	68.37	82.82	83.24	94.78
- GQA	$d_{kv}=256$		80.75	83.35	63.26	87.50	76.83	69.67	84.28	86.88	94.29
	$d_{kv}=128$	0.5B	-91.07%	80.63 _{-0.12}	82.71	63.26	87.54	76.83	69.41	84.02	86.80
- GQA2MLA	$d_{kv}=64$	(0.002%)	-94.64%	79.47 _{-1.28}	80.63	63.08	87.85	76.51	68.76	81.35	84.76
	$d_{kv}=32$		-96.43%	77.02 _{-3.73}	78.11	62.87	87.42	74.75	66.54	77.15	82.16

Table 1: Performance of three VLMs with different architectures (e.g., MHA2MLA, GQA2MLA) and d_{kv} . The eight benchmarks include AI2D (2016), GQA (2019), POPE (2023b), SEED-Bench (SEED, 2023a), RealWorldQA (RWQ, 2025), MMBench (MMB, 2024c), ChartQA (2022), DocVQA (2021).

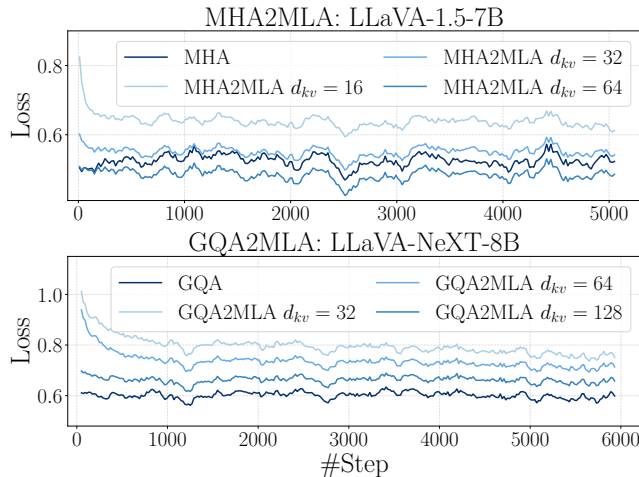


Figure 2: Training loss curves of **MHA2MLA** (LLaVA-1.5) and **GQA2MLA** (LLaVA-NeXT) with different d_{kv} settings.

Greater compression widens the loss gap compared to the uncompressed baseline. When d_{kv} is large (64 or 128), with KV cache compression of 84.38% (LLaVA-NeXT) and 62.50% (LLaVA-1.5), it achieves loss comparable to fine-tuning the original model. Note that the changing trends of these loss curves are the same, indicating that our architecture transfer preserves the VLMs’ internal knowledge to a large extent.

Overall, our MHA2MLA-VLM method generalizes across both MHA and GQA architectures, enabling parameter-efficient and data-efficient training for diverse VLMs.

3.2 MHA2MLA, Cache Pruning and Compression

To demonstrate the efficiency of MHA2MLA-VLM in compressing KV cache, we compare it with representative pruning and quantization methods, as summarized in Table 2.

First, at comparable compression ratios, MHA2MLA achieves significantly better performance than the cache pruning methods. For instance, at a 62.50% reduction in KV cache, the average score reaches 68.75, while cache pruning methods such as H₂O and TOVA yield notably lower scores (e.g., 63.38 and 60.48, respectively). Similar trends can be observed under 75.00% and 81.25% compression ratios.

Second, MHA2MLA can be effectively combined with cache quantization methods to achieve further compression without performance degradation. For all levels of d_{kv} compression in GQA2MLA, applying Int4_{Quanto} and Int4_{HQQ}, even at compression levels of 84.38% ($d_{kv}=128$ with 70.22) and 90.63% ($d_{kv}=64$ with 68.66), the performance is still significantly better than the Int2_{HQQ} of 67.21 based on cache quantization at 87.5%. These results demonstrate that MHA2MLA and quantization techniques are highly compatible and can be effectively combined to achieve compounded compression benefits with minimal loss in performance.

These results validate that MHA2MLA achieves better performance compared to existing cache compression strategies, and can integrate seamlessly with cache quantization methods to achieve higher cache reduction.

3.3 Ablation Study

Effect of Modality-Decoupled SVD To better understand the contribution of our proposed MD-SVD, we disentangle its two core designs (e.g., modality decoupled and SVD-based initialization) and evaluate them separately. As depicted in Table 3, modality decoupled leads to consistent improvements across all d_{kv} (+0.42 at $d_{kv}=32$, +0.45 at $d_{kv}=64$). While MD-SVD Init model exhibits slightly lower performance at $d_{kv}=32$, it achieves a substantial improvement at $d_{kv}=64$ (+1.23). Thus, SD-SVD Init performs better overall.

Moreover, Figure 3 demonstrates that under various KV cache compression conditions, our proposed MD-SVD initialization not only reduces the initial training loss, but also leads

Model	Type	KV Mem.	Avg.
Origin			
LLaVA-NeXT	BF16	100.0%	70.78
Cache Pruning			
	H ₂ O	-81.25%	61.43
	TOVA		54.53
LLaVA-NeXT	H ₂ O	-75.00%	62.34
	TOVA		57.24
	H ₂ O		63.38
	TOVA	-62.50%	60.48
Cache Quantization			
	Int4 _{Quanto}	-75.00%	70.53
LLaVA-NeXT	Int4 _{HQQ}		70.62
	Int2 _{Quanto}		67.21
	Int2 _{HQQ}	-87.50%	60.71
	BF16	-37.50%	70.23
$d_{kv} = 128$	Int4 _{Quanto}	-84.38%	70.22
	Int4 _{HQQ}		70.21
	BF16	-62.50%	68.75
$d_{kv} = 64$	Int4 _{Quanto}	-90.63%	68.66
	Int4 _{HQQ}		68.64
	BF16	-75.00%	66.72
$d_{kv} = 32$	Int4 _{Quanto}	-93.75%	66.71
	Int4 _{HQQ}		66.72

Table 2: Comparison of GQA2MLA with other cache compression strategies, including Cache Pruning and Cache Quantization for LLaVA-NeXT. Bolded scores indicate that GQA2MLA outperforms Cache Pruning or Int2 quantization methods at the same or higher compression levels.

to better convergence throughout training. Overall, these results demonstrate that SVD Init enhances both optimization efficiency and model robustness across a range of low-ranks.

Effect of Two Stage Training Table 3 indicates that two-stage efficient parameter fine-tuning outperforms single-stage tuning at various cache compression rates, demonstrating the effectiveness of two-stage training.

Effect of Partial-RoPE strategies Table 4 presents a comparison of LLaVA-NeXT performance under two different Partial-RoPE strategies. We observe that \mathcal{S}_{MKL} achieves better performance than $\mathcal{S}_{2\text{-norm}}$. To better understand this behavior, we conduct a detailed analysis in Section 4.2.

4 Analysis

4.1 Empirical Validation for MD-SVD

Theorem 2.1 proves that the loss from our proposed Modality-Decoupled SVD strategy is always smaller than the loss from jointly optimizing over visual and text (e.g. $\min \mathcal{L}_{\text{joint}}^2 \geq \min \mathcal{L}_{\text{visual}}^2 + \min \mathcal{L}_{\text{text}}^2$). To empirically validate this theoretical insight, we compute the ratio between split and joint losses, defined as $\mathcal{L}_{\text{split}}^2 / \mathcal{L}_{\text{joint}}^2$ to and report this across all layers for different VLM architectures, as shown in Figure 4.

Model	d_{kv}	KV Mem.	Avg
LLaVA-NeXT _{GQA2MLA}	32	93.75%	67.56
	64	90.63%	69.70
<i>w/o Modality Decoupled</i>	32	93.75%	67.14
	64	90.63%	69.25
<i>w/o MD-SVD Init</i>	32	93.75%	68.21
	64	90.63%	68.47
<i>w/o Two Stage</i>	32	93.75%	67.03
	64	90.63%	68.82

Table 3: Ablation study on the core designs of GQA2MLA on LLaVA-NeXT, including Two Stage training, Modality Decoupled, and MD-SVD Initialization.




Figure 3: Training loss of GQA2MLA on LLaVA-NeXT w and w/o MD-SVD initialization under different d_{kv} .

First, across all models, the ratio curves (three solid lines) consistently fall below the joint baseline (red dotted line) throughout all layers. This observation indicates that processing visual and text modalities decoupled leads to a strictly lower loss compared to joint processing, irrespective of architecture, which empirically validates Theorem 1.

Second, the relative advantage of the decoupled strategy becomes increasingly significant as the layer increases. It indicates that deeper layers are more sensitive to cross-modal interference and benefit more from modality-decoupled optimization. For example, in Qwen2.5-VL, the ratio decreases progressively and reaches a maximum relative loss reduction of 35.85% at the final layer, highlighting the decoupled approach’s compounding benefit in deep architectures.

4.2 Comparison of $\mathcal{S}_{2\text{-norm}}$ and \mathcal{S}_{MKL}

Figure 5 analyze the frequency profiles of the selected key dimensions under two strategies to understand their frequency preferences. \mathcal{S}_{MKL} consistently selects dimensions that con-

Model	Partial-RoPE Strategy	Avg.
LLaVA-NeXT	$\mathcal{S}_{2\text{-norm}}$	70.31
LLaVA-NeXT	\mathcal{S}_{MKL}	70.54

Table 4: Comparison of average performance under different Partial-RoPE strategies for LLaVA-NeXT.

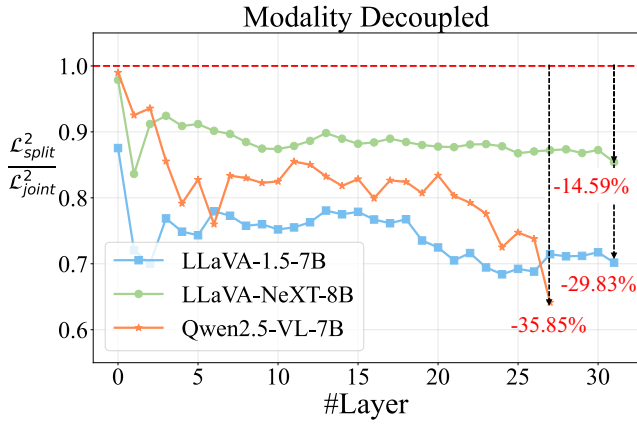


Figure 4: Quantitative analysis of MD-SVD via layer-wise loss ratio. Modality decoupled shows consistent improvements over joint optimization across all models.

centrate more heavily in the high-frequency range, as indicated by medians and tighter interquartile ranges. In contrast, $S_{2\text{-norm}}$ exhibits a broader spread and a strong tendency toward low-frequency selection, which are always noisy and less informative. S_{MKL} explicitly evaluates the impact of each frequency component on the final attention distribution. Frequencies that induce larger changes are considered more informative and thus preferentially selected, allowing the model to retain discriminative high-frequency dimensions and improving the effectiveness of key-value compression.

Empirical results show that this sensitivity-based criterion consistently selects more important key dimensions, and leads to lower attention loss and better downstream performance, especially in multimodal settings where fine-grained alignment is crucial.

5 Related Work

Vision-Language Models Represented by GPT4 (OpenAI et al. 2024), VLMs have shown their strong strength and are increasingly becoming one of the mainstream research directions. They combine visual and language models to achieve cross-modal understanding and reasoning capabilities. Pioneering models such as LLaVA (Liu et al. 2024a) uses a simple projection layer to promote image-text alignment and uses a two-stage training method to improve model capabilities. Furthermore, MouSi (Fan et al. 2024) and Cambrian-1 (Tong et al. 2024) leverage the unique attributes of diverse visual encoders and unify their strengths to enrich the multimodal understanding of VLMs. Recently, the InternLM-XComposer (Zhang et al. 2023a; Dong et al. 2024; Zhang et al. 2024) and InternVL (Chen et al. 2024b,a; Zhu et al. 2025) family of models have shown leading performance.

Moreover, videos or multiple images require more tokens for visual signals. For example, VideoPoet (Konratyuk et al. 2024) and VideoLLaVA (Lin et al. 2024) encode each frame with thousands of tokens, quickly saturating computation budgets. This token explosion makes attention the primary bottleneck and calls for stronger sparsification strategies to

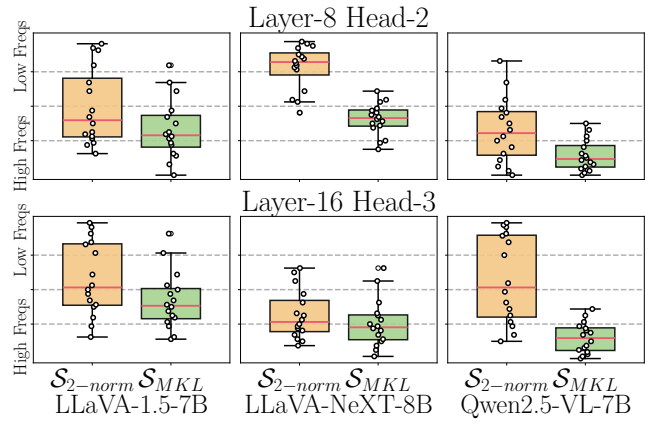


Figure 5: Comparison of multimodal partial-rope selection between $S_{2\text{-norm}}$ and S_{MKL} .

unlock the next level of VLM performance.

Efficient Architectures KV cache consumption is a significant challenge for LLMs and VLMs, especially when dealing with long contexts. The primary methods for KV-cache compression can be broadly categorized into three approaches:

Model-level optimization like MQA (Shazeer 2019) and GQA (Ainslie et al. 2023b) use intra-layer grouping, sharing key-value pairs across heads to save memory. MLA (DeepSeek-AI et al. 2024) further compresses KV caches via a low-rank joint representation across heads and layers, achieving substantial savings and demonstrating effectiveness in large-scale deployments.

Cache quantization techniques, such as KVQuant (Hooper et al. 2024), MassiveActivation (Sun et al. 2024), and HQQ (Badri and Shaji 2023), compress KV cache into low-bit formats (e.g., 4-bit or 2-bit), significantly reducing memory footprint and improving inference efficiency. However, they may suffer from performance degradation under aggressive quantization or cross-modal inputs.

Cache pruning methods, including H₂O (Zhang et al. 2023b), TOVA (Oren et al. 2024), PyramidKV (Cai et al. 2025), and AdaKV (Feng et al. 2025), aim to reduce KV cache size by eliminating less informative tokens or attention heads, but may mistakenly discard contextually crucial information, especially under long-range dependencies.

6 Conclusion

In this work, we present **MHA2MLA-VLM**, a parameter-efficient and multimodal-aware framework for adapting VLMs to DeepSeek’s MLA architecture. Our approach introduces modality-adaptive partial-RoPE and modality-decoupled low-rank approximation, enabling substantial KV cache reduction with minimal fine-tuning across various models. Extensive experiments demonstrate that MHA2MLA-VLM achieves efficient inference while maintaining original model performance, offering a practical solution for scalable multimodal applications.

Acknowledgements

The authors wish to thank the anonymous reviewers for their helpful comments. This work was partially funded by the Major Key Project of PCL under Grant PCL2024A06, National Natural Science Foundation of China (No.62576106, 62376061, 62476061, 62506079), Shanghai Rising-Star Program (23QA1400200), Natural Science Foundation of Shanghai (23ZR1403500) and Fudan Kunpeng & Ascend Center of Cultivation. The computations in this research were performed using the CFFF platform of Fudan University.

References

- Ainslie, J.; Lee-Thorp, J.; de Jong, M.; Zemlyanskiy, Y.; Lebrón, F.; and Sanghai, S. 2023a. GQA: Training Generalized Multi-Query Transformer Models from Multi-Head Checkpoints. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proc. EMNLP 2023*, 4895–4901. Association for Computational Linguistics.
- Ainslie, J.; Lee-Thorp, J.; de Jong, M.; Zemlyanskiy, Y.; Lebrón, F.; and Sanghai, S. 2023b. GQA: Training Generalized Multi-Query Transformer Models from Multi-Head Checkpoints. In Bouamor, H.; Pino, J.; and Bali, K., eds., *EMNLP 2023*, 4895–4901. Association for Computational Linguistics.
- Badri, H.; and Shaji, A. 2023. Half-Quadratic Quantization of Large Machine Learning Models.
- Bai, J.; Bai, S.; Yang, S.; Wang, S.; Tan, S.; Wang, P.; Lin, J.; Zhou, C.; and Zhou, J. 2023. Qwen-VL: A Versatile Vision-Language Model for Understanding, Localization, Text Reading, and Beyond. arXiv:2308.12966.
- Bai, S.; Chen, K.; Liu, X.; and et al. 2025. Qwen2.5-VL Technical Report. arXiv:2502.13923.
- Barbero, F.; Vitvitskiy, A.; Perivolaropoulos, C.; Pascanu, R.; and Veličković, P. 2025. Round and Round We Go! What makes Rotary Positional Encodings useful? arXiv:2410.06205.
- Black, S.; Gao, L.; Wang, P.; Leahy, C.; and Biderman, S. 2021. GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow. If you use this software, please cite it using these metadata.
- Bordes, F.; Pang, R. Y.; Ajay, A.; and et al. 2024. An Introduction to Vision-Language Modeling. arXiv:2405.17247.
- Cai, Z.; Zhang, Y.; Gao, B.; Liu, Y.; Li, Y.; Liu, T.; Lu, K.; Xiong, W.; Dong, Y.; Hu, J.; and Xiao, W. 2025. PyramidKV: Dynamic KV Cache Compression based on Pyramidal Information Funneling. arXiv:2406.02069.
- Chen, Z.; Wang, W.; Tian, H.; Ye, S.; Gao, Z.; Cui, E.; Tong, W.; Hu, K.; Luo, J.; Ma, Z.; Ma, J.; Wang, J.; Dong, X.; Yan, H.; Guo, H.; He, C.; Shi, B.; Jin, Z.; Xu, C.; Wang, B.; Wei, X.; Li, W.; Zhang, W.; Zhang, B.; Cai, P.; Wen, L.; Yan, X.; Dou, M.; Lu, L.; Zhu, X.; Lu, T.; Lin, D.; Qiao, Y.; Dai, J.; and Wang, W. 2024a. How far are we to GPT-4V? Closing the gap to commercial multimodal models with open-source suites. *Sci. China Inf. Sci.*, 67(12).
- Chen, Z.; Wu, J.; Wang, W.; and et al. 2024b. InternVL: Scaling up Vision Foundation Models and Aligning for Generic Visual-Linguistic Tasks. arXiv:2312.14238.
- DeepSeek-AI; Liu, A.; Feng, B.; Wang, B.; and et al. 2024. DeepSeek-V2: A Strong, Economical, and Efficient Mixture-of-Experts Language Model. arXiv:2405.04434.
- Dong, X.; Zhang, P.; Zang, Y.; and et al. 2024. InternLM-XComposer2: Mastering Free-form Text-Image Composition and Comprehension in Vision-Language Large Model. arXiv:2401.16420.
- Fan, X.; Ji, T.; Jiang, C.; Li, S.; Jin, S.; Song, S.; Wang, J.; Hong, B.; Chen, L.; Zheng, G.; Zhang, M.; Huang, C.; Zheng, R.; Xi, Z.; Zhou, Y.; Dou, S.; Ye, J.; Yan, H.; Gui, T.; Zhang, Q.; Qiu, X.; Huang, X.; Wu, Z.; and Jiang, Y.-G. 2024. MouSi: Poly-Visual-Expert Vision-Language Models. arXiv:2401.17221.
- Feng, Y.; Lv, J.; Cao, Y.; Xie, X.; and Zhou, S. K. 2025. Ada-KV: Optimizing KV Cache Eviction by Adaptive Budget Allocation for Efficient LLM Inference. arXiv:2407.11550.
- Hooper, C.; Kim, S.; Mohammadzadeh, H.; Mahoney, M. W.; Shao, Y. S.; Keutzer, K.; and Gholami, A. 2024. KVQuant: Towards 10 Million Context Length LLM Inference with KV Cache Quantization. In Globersons, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J. M.; and Zhang, C., eds., *NeurIPS 2024*.
- Hsu, Y.; Hua, T.; Chang, S.; Lou, Q.; Shen, Y.; and Jin, H. 2022. Language model compression with weighted low-rank factorization. In *ICLR 2022*. OpenReview.net.
- Hudson, D. A.; and Manning, C. D. 2019. GQA: A New Dataset for Real-World Visual Reasoning and Compositional Question Answering. In *CVPR 2019*, 6700–6709. Computer Vision Foundation / IEEE.
- Ji, T.; Guo, B.; Wu, Y.; Guo, Q.; Shenlixing, S.; Chenzhan, C.; Qiu, X.; Zhang, Q.; and Gui, T. 2025. Towards Economical Inference: Enabling DeepSeek’s Multi-Head Latent Attention in Any Transformer-based LLMs. In Che, W.; Nabende, J.; Shutova, E.; and Pilehvar, M. T., eds., *ACL 2025*, 33313–33328. Association for Computational Linguistics.
- Keles, F. D.; Wijewardena, P. M.; and Hegde, C. 2023. On The Computational Complexity of Self-Attention. In Agrawal, S.; and Orabona, F., eds., *ALT 2023*, volume 201 of *Proceedings of Machine Learning Research*, 597–619. PMLR.
- Kembhavi, A.; Salvato, M.; Kolve, E.; Seo, M. J.; Hajishirzi, H.; and Farhadi, A. 2016. A Diagram is Worth a Dozen Images. In Leibe, B.; Matas, J.; Sebe, N.; and Welling, M., eds., *Proc. ECCV 2016*, volume 9908 of *Lecture Notes in Computer Science*, 235–251. Springer.
- Kondratyuk, D.; Yu, L.; Gu, X.; Lezama, J.; Huang, J.; Schindler, G.; Hornung, R.; Birodkar, V.; Yan, J.; Chiu, M.; Somandepalli, K.; Akbari, H.; Alon, Y.; Cheng, Y.; Dillon, J. V.; Gupta, A.; Hahn, M.; Hauth, A.; Hendon, D.; Martinez, A.; Minnen, D.; Sirotenko, M.; Sohn, K.; Yang, X.; Adam, H.; Yang, M.; Essa, I.; Wang, H.; Ross, D. A.; Seybold, B.; and Jiang, L. 2024. VideoPoet: A Large Language Model for Zero-Shot Video Generation. In *ICML 2024*. OpenReview.net.
- Li, B.; Wang, R.; Wang, G.; Ge, Y.; Ge, Y.; and Shan, Y. 2023a. SEED-Bench: Benchmarking Multimodal LLMs with Generative Comprehension. arXiv:2307.16125.

- Li, H.; Li, Y.; Tian, A.; Tang, T.; Xu, Z.; Chen, X.; Hu, N.; Dong, W.; Li, Q.; and Chen, L. 2025. A Survey on Large Language Model Acceleration based on KV Cache Management. *Trans. Mach. Learn. Res.*, 2025.
- Li, Y.; Du, Y.; Zhou, K.; Wang, J.; Zhao, W. X.; and Wen, J. 2023b. Evaluating Object Hallucination in Large Vision-Language Models. In Bouamor, H.; Pino, J.; and Bali, K., eds., *EMNLP 2023*, 292–305. Association for Computational Linguistics.
- Lin, B.; Ye, Y.; Zhu, B.; Cui, J.; Ning, M.; Jin, P.; and Yuan, L. 2024. Video-LLaVA: Learning United Visual Representation by Alignment Before Projection. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y., eds., *EMNLP 2024*, 5971–5984. Association for Computational Linguistics.
- Liu, H.; Li, C.; Li, Y.; and Lee, Y. J. 2024a. Improved Baselines with Visual Instruction Tuning. In *CVPR 2024*, 26286–26296. IEEE.
- Liu, H.; Li, C.; Li, Y.; Li, B.; Zhang, Y.; Shen, S.; and Lee, Y. J. 2024b. LLaVA-NeXT: Improved reasoning, OCR, and world knowledge.
- Liu, Y.; Duan, H.; Zhang, Y.; Li, B.; Zhang, S.; Zhao, W.; Yuan, Y.; Wang, J.; He, C.; Liu, Z.; Chen, K.; and Lin, D. 2024c. MMBench: Is Your Multi-modal Model an All-Around Player? In Leonaridis, A.; Ricci, E.; Roth, S.; Ruskakovsky, O.; Sattler, T.; and Varol, G., eds., *ECCV 2024*, volume 15064 of *Lecture Notes in Computer Science*, 216–233. Springer.
- Masry, A.; Long, D. X.; Tan, J. Q.; Joty, S. R.; and Hoque, E. 2022. ChartQA: A Benchmark for Question Answering about Charts with Visual and Logical Reasoning. In Muresan, S.; Nakov, P.; and Villavicencio, A., eds., *Findings of ACL 2022*, 2263–2279. Association for Computational Linguistics.
- Mathew, M.; Karatzas, D.; and Jawahar, C. V. 2021. DocVQA: A Dataset for VQA on Document Images. In *WACV 2021*, 2199–2208. IEEE.
- OpenAI; Achiam, J.; Adler, S.; Agarwal, S.; and et al. 2024. GPT-4 Technical Report. arXiv:2303.08774.
- Oren, M.; Hassid, M.; Nir, Y.; Adi, Y.; and Schwartz, R. 2024. Transformers are Multi-State RNNs. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y., eds., *EMNLP 2024*, 18724–18741. Association for Computational Linguistics.
- Pan, Q.; Ji, W.; Ding, Y.; Li, J.; Chen, S.; Wang, J.; Zhou, J.; Chen, Q.; Zhang, M.; Wu, Y.; and He, L. 2025. A Survey of Slow Thinking-based Reasoning LLMs using Reinforced Learning and Inference-time Scaling Law. arXiv:2505.02665.
- Shazeer, N. 2019. Fast Transformer Decoding: One Write-Head is All You Need. arXiv:1911.02150.
- Su, J.; Ahmed, M. H. M.; Lu, Y.; Pan, S.; Bo, W.; and Liu, Y. 2024. RoFormer: Enhanced transformer with Rotary Position Embedding. *Neurocomputing*, 568: 127063.
- Sun, M.; Chen, X.; Kolter, J. Z.; and Liu, Z. 2024. Massive Activations in Large Language Models. arXiv:2402.17762.
- Tong, P.; Brown, E.; Wu, P.; Woo, S.; Iyer, A.; Akula, S. C.; Yang, S.; Yang, J.; Middepogu, M.; Wang, Z.; Pan, X.; Fergus, R.; LeCun, Y.; and Xie, S. 2024. Cambrian-1: A Fully Open, Vision-Centric Exploration of Multimodal LLMs. In Globersons, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J. M.; and Zhang, C., eds., *NeurIPS 2024*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention Is All You Need. arXiv:1706.03762.
- Wang, P.; Bai, S.; Tan, S.; and et al. 2024. Qwen2-VL: Enhancing Vision-Language Model’s Perception of the World at Any Resolution. arXiv:2409.12191.
- Wang, X.; Alam, S.; Wan, Z.; Shen, H.; and Zhang, M. 2025a. SVD-LLM V2: Optimizing Singular Value Truncation for Large Language Model Compression. In Chiruzzo, L.; Ritter, A.; and Wang, L., eds., *NAACL 2025*, 4287–4296. Association for Computational Linguistics.
- Wang, X.; Zheng, Y.; Wan, Z.; and Zhang, M. 2025b. SVD-LLM: Truncation-aware Singular Value Decomposition for Large Language Model Compression. In *ICLR 2025*. OpenReview.net.
- Wei, X.; Liu, X.; Zang, Y.; Dong, X.; Zhang, P.; Cao, Y.; Tong, J.; Duan, H.; Guo, Q.; Wang, J.; Qiu, X.; and Lin, D. 2025. VideoRoPE: What Makes for Good Video Rotary Position Embedding? arXiv:2502.05173.
- Xu, L.; Xie, H.; Qin, S.-Z. J.; Tao, X.; and Wang, F. L. 2023. Parameter-Efficient Fine-Tuning Methods for Pre-trained Language Models: A Critical Review and Assessment. arXiv:2312.12148.
- Zhang, P.; Dong, X.; Wang, B.; and et al. 2023a. InternLM-XComposer: A Vision-Language Large Model for Advanced Text-image Comprehension and Composition. arXiv:2309.15112.
- Zhang, P.; Dong, X.; Zang, Y.; and et al. 2024. InternLM-XComposer-2.5: A Versatile Large Vision Language Model Supporting Long-Contextual Input and Output. arXiv:2407.03320.
- Zhang, Y.; Zhang, H.; Tian, H.; Fu, C.; Zhang, S.; Wu, J.; Li, F.; Wang, K.; Wen, Q.; Zhang, Z.; Wang, L.; and Jin, R. 2025. MME-RealWorld: Could Your Multimodal LLM Challenge High-Resolution Real-World Scenarios that are Difficult for Humans? In *ICLR 2025*. OpenReview.net.
- Zhang, Z.; Sheng, Y.; Zhou, T.; Chen, T.; Zheng, L.; Cai, R.; Song, Z.; Tian, Y.; Ré, C.; Barrett, C. W.; Wang, Z.; and Chen, B. 2023b. H2O: Heavy-Hitter Oracle for Efficient Generative Inference of Large Language Models. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *NeurIPS 2023*.
- Zhu, J.; Wang, W.; Chen, Z.; and et al. 2025. InternVL3: Exploring Advanced Training and Test-Time Recipes for Open-Source Multimodal Models. arXiv:2504.10479.