

Sparrow: Text-Anchored Window Attention with Visual-Semantic Glimpsing for Speculative Decoding in Video LLMs

Libo Zhang Zhaoning Zhang* Wangyang Hong
Peng Qiao Dongsheng Li

National Key Laboratory of Parallel and Distributed Computing
College of Computer Science and Technology
National University of Defense Technology, Changsha, China.
{zhanglibo, zhangzhaoning, hongwanyang9527,
pengqiao, dsli}@nudt.edu.cn

Abstract

Although speculative decoding is widely used to accelerate Vision-Language Models (VLMs) inference, it faces severe performance collapse when applied to Video Large Language Models (Vid-LLMs). The draft model typically falls into the trap of attention dilution and negative visual gain due to key-value cache explosion and context window mismatches. We observe a visual semantic internalization phenomenon in Vid-LLMs, indicating that critical visual semantics are implicitly encoded into text hidden states during deep-layer interactions, which renders raw visual inputs structurally redundant during deep inference. To address this, we propose the Sparrow¹ framework, which first utilizes visually-aware text-anchored window attention via hidden state reuse to fully offload visual computation to the target model, and leverages intermediate-layer visual state bridging to train the draft model with semantic-rich intermediate states, thereby filtering out low-level visual noise. Additionally, a multi-token prediction strategy is introduced to bridge the training-inference distribution shift. Experiments show that Sparrow achieves an average speedup of $2.82\times$ even with 25k visual tokens, effectively resolving the performance degradation in long sequences and offering a practical solution for real-time long video tasks.

1 Introduction

Video Large Language Models (Vid-LLMs) have achieved remarkable breakthroughs in spatiotemporal understanding and generation tasks (Li et al., 2024a; Bai et al., 2025; Maaz et al., 2024; Zhang et al., 2024). However, this leap in capability incurs substantial computational costs: the rich spatiotemporal semantics inherent in long videos are mapped into extremely long input sequences, resulting in prohibitive GPU memory consumption and computational overhead. The inference latency bottleneck

* indicates corresponding authors.

¹<https://github.com/ddInference/Sparrow>

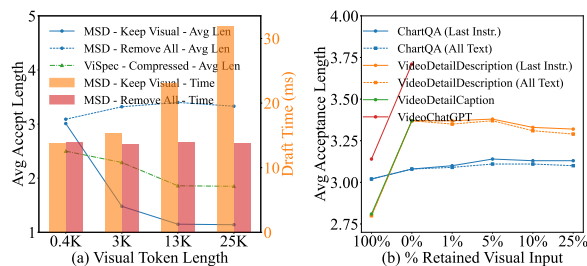


Figure 1: Analysis of the impact of visual token length on performance. (a) Comparison of average accepted length and drafting latency between MSD and ViSpec on the VideoDetailCaption dataset across varying visual sequence lengths. (b) Performance on image and video tasks with MSD at different visual retention rates. The Last Instr. and All Text strategies retain the top- $x\%$ visual tokens based on the attention scores of the final instruction token and the average attention scores of all text tokens towards visual regions, respectively. ViSpec and MSD employ Qwen2.5-VL-7B and Qwen2-VL-7B as target models, respectively, with both utilizing the official released weights for their draft models.

of current Vid-LLMs stems primarily from the dual constraints imposed by the generation paradigm and input representation. At the generation level, mainstream autoregressive decoding is constrained by the memory wall, where its token-by-token sequential mechanism imposes an inherent speed bottleneck (Fu et al., 2024). At the representation level, video data, once encoded (Wang et al., 2025b; Zhang et al., 2023), produces massive visual tokens. This not only significantly exacerbates forward computational loads but also precipitates a rapid expansion of the key-value (KV) cache during long-sequence generation, posing a critical obstacle to large-scale deployment. To mitigate this bottleneck, previous studies have explored multi-granular visual pruning strategies (Chen et al., 2024; Liu et al., 2024f; Xing et al., 2024; Tao et al., 2025;

Zhang et al., 2025a), premised on the prior that visual tokens exhibit significant redundancy. Although such methods reduce memory usage during the prefill stage, they inherently constitute lossy compression: the direct removal of input information inevitably compromises the model’s fine-grained perception capabilities. In contrast, speculative decoding (Hu et al., 2025a; Chen et al., 2023; Leviathan et al., 2023) provides a superior, lossless alternative by leveraging a lightweight draft model for rapid prediction followed by parallel verification by the target model.

However, adapting speculative decoding to long-video scenarios presents formidable challenges. Existing multimodal speculative decoding, primarily tailored for short-sequence image tasks, suffer from a precipitous decline in acceleration gains when applied to long videos (as illustrated in Figure 1, the average accepted lengths of MSD (Lin et al., 2025) and ViSpec (Kang et al., 2025) plummet by 63% and 30%, respectively). This performance collapse stems from two critical obstacles. First is the memory access bottleneck within the draft model. The KV cache explosion induced by long videos significantly increases the inference latency of the lightweight model, effectively negating the temporal dividends gained from speculation. Second is the inherent conflict between model architecture and sequence length. On one hand, long video sequences (often exceeding 10k tokens) far surpass the pre-training context windows of mainstream lightweight models (typically capped at 2,048 tokens), leading to length mismatch and forced information truncation. On the other hand, excessively long sequences of visual tokens trigger severe attention dilution. The naïve application of existing multimodal speculative decoding methods leaves capacity-constrained draft models struggling to extract critical features from massive visual inputs, ultimately resulting in a drastic degradation of speculative accuracy and overall efficiency.

Investigating the root causes of this phenomenon, we identify a critical factor in long-video scenarios: the negative gain of visual information. For capacity-constrained draft models, retaining a massive number of visual tokens proves more detrimental than beneficial; the tens of thousands of visual features effectively manifest as computational noise, causing speculative performance to degrade significantly as the visual sequence length increases. Concurrently, an analysis of the information flow within Vid-LLMs reveals a pivotal mech-

anism: visual semantic internalization. Through multi-stage inter-layer interactions, Vid-LLMs implicitly internalize essential visual semantics, encoding them directly into the hidden states corresponding to text positions. Consequently, during the deep inference stage, the explicit stream of raw visual inputs becomes structurally redundant.

Based on these insights, we propose the visually-aware text-anchored window attention via hidden state reuse (HSR-VATA) strategy. The core of this method is computation offloading, aiming to transfer the intensive visual processing tasks entirely to the target model. We observe that the target model inherently integrates comprehensive visual information and encodes it into text hidden states; therefore, the draft model need only reuse these ready-made states via the HSR mechanism. This reuse acts as an efficient glimpse into the massive visual stream, enabling the draft model to rapidly capture key semantics and avoid the high overhead of processing raw visual sequences. Coupled with VATA to confine attention to text anchor positions, this framework allows the draft model to maintain its visual perception capabilities without explicitly receiving or processing raw visual tokens.

Furthermore, addressing the difficulty of the draft model’s minimalist architecture in handling low-level raw visual noise, we propose intermediate-layer visual state bridging (IVSB) during the training phase. By directly extracting intermediate-layer visual states from the target model, specifically those where semantic interaction is most active and alignment is implicitly established, as the visual information source for training the draft model, we leverage the multi-stage processing characteristics of Vid-LLMs. This approach preserves high-level semantic features while effectively filtering out low-level redundancy and noise from raw visual signals, achieving efficient cross-modal alignment during training. Simultaneously, by incorporating a multi-token prediction strategy, we effectively bridge the distribution discrepancy between training and inference.

Contributions. The main contributions of this paper are as follows:

- To the best of our knowledge, this is the first work to apply a lightweight draft model to Vid-LLMs. We reveal attention dilution and negative gain in long-video speculative decoding, validating visual offloading feasibility via the target model’s deep hidden states.

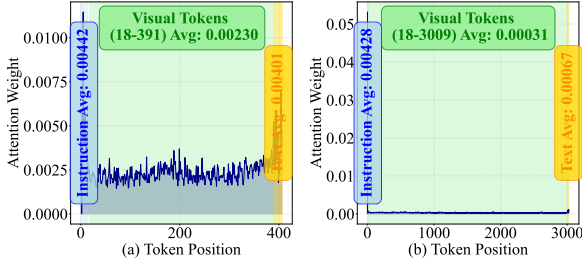


Figure 2: (a) & (b) Visualization of the average attention weight distribution from the last instruction token to preceding Instruction, Visual, and Text tokens in short (0.4k) and long (3k) sequence tasks.

- We propose the Sparrow framework. By integrating HSR-VATA, we eliminate visual redundancy and attention dilution. Furthermore, by combining IVSB with multi-token prediction, we filter visual noise while reinforcing cross-modal alignment, effectively resolving the training-inference distribution shift.
- Experiments demonstrate that even with ultra-long visual inputs of 25k tokens, our method achieves an average speedup of $2.82\times$, while effectively overcoming the bottleneck of performance degradation in long-video scenarios.

2 Insight

Current research on multimodal speculative decoding (Hu et al., 2025b; Huo et al., 2025; Lin et al., 2025; Lee et al., 2024; Ganesan et al., 2025; Kang et al., 2025) primarily focuses on image tasks, where draft models typically employ either full (e.g., MSD) or compressed (e.g., ViSpec) visual information as input. In LLaVA (Liu et al., 2024c), for instance, visual tokens typically number around 576, accounting for merely 64% of the total input sequence (Chen et al., 2024). However, as the domain transitions towards long video tasks, the input distribution shifts drastically. Taking Qwen2.5-VL-7B as an example, video encoding can generate up to 25k visual tokens, with their proportion surging to over 99%. This severe distributional shift compels us to re-examine a critical question: Can parameter-constrained draft models effectively process such a massive volume of visual information?

2.1 Performance Bottlenecks in Long-Video Scenarios

We first investigate the scalability of existing strategies in long video scenarios. As shown in Fig-

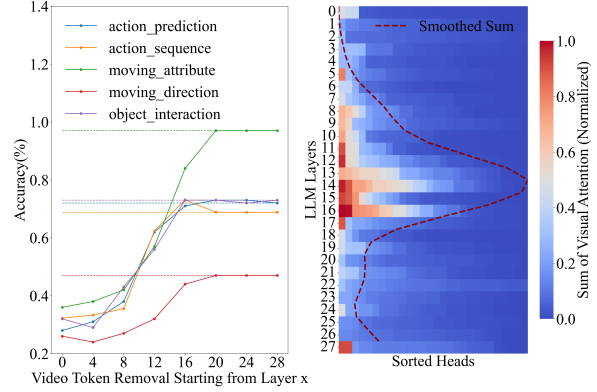


Figure 3: Layer-wise Importance Analysis of Visual Tokens in Qwen2.5-VL-7B. (a) Comparison of task accuracy (solid lines) after removing visual tokens starting from layer x versus the native baseline (dashed lines). (b) Sum of attention received by all visual tokens from the last instruction token across different attention heads and model layers.

ure 1 (a), current methods encounter severe bottlenecks as sequence length grows. The full-input strategy (e.g., MSD) suffers a critical efficiency regression: as video token length expands from 0.4k to 25k, the average accepted length decays by 63% while latency doubles, resulting in a negative speedup. Meanwhile, compression-based strategies (e.g., ViSpec), constrained by their training on short sequences, fail to capture complex spatiotemporal dynamics; this leads to a significant 30% performance drop despite the reduced computational overhead.

2.2 Redundancy of Visual Information and Negative Gain Phenomenon

To unveil the true contribution of visual information during inference, we implemented multi-granularity visual token pruning based on attention distributions, contrasting the distinct response patterns between image and video tasks (see Figure 1 (a) & (b) and Table 7). Experimental results indicate that the average accepted length in image tasks is insensitive to pruning ratios, suggesting a neutral characteristic of visual redundancy. Conversely, video tasks exhibit a counter-intuitive negative correlation: the average accepted length increases with fewer visual tokens. Furthermore, the model demonstrates exceptional robustness under zero-visual-input settings, remaining virtually unaffected by the original video token length (0.4k to 25k). This significant discrepancy provides compelling evidence that excessive video sequences in

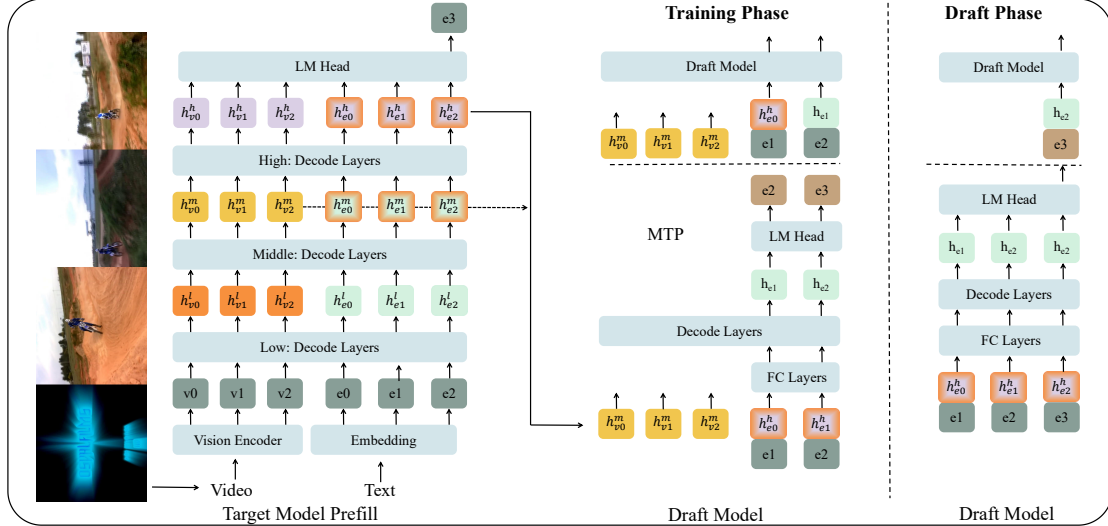


Figure 4: Illustration of the Sparrow framework. The left and right panels illustrate the training and inference phases, respectively. Initially, the target model performs multi-stage fusion on visual embeddings v_i and text embeddings e_i , yielding a noise-filtered visual hidden state $h_{v_i}^m$ and a visually-infused text hidden state $h_{e_i}^h$. In the training phase, $h_{v_i}^m$ and $h_{e_i}^h$ are concatenated and fed into the draft model to produce $h_{e_i}^h$. Subsequently, $h_{v_i}^m$ is concatenated with $h_{e_i}^h$ and re-input into the draft model. Losses from both stages are jointly computed to mitigate the discrepancy between training and inference. In the inference phase, the draft model directly takes $h_{e_i}^h$ as input.

the drafting phase are not only highly redundant but effectively constitute noise that suppresses language generation.

2.3 Attribution Analysis I: Attention Dilution Under Limited Capacity

Why does visual information transition into a detrimental factor in long videos? Integrating the analysis of Figure 2 (a) & (b), we attribute the root cause to the limited attention capacity of small-scale draft models. In short-sequence scenarios (images/short videos), while the performance stability upon removing visual inputs confirms high redundancy, this redundancy proves non-fatal: the model, leveraging the selectivity of self-attention, successfully anchors onto key regions (Figure 2 (a)) to automatically suppress interference. Conversely, the situation deteriorates drastically for long videos. Tens of thousands of visual tokens not only breach the effective boundaries of positional encoding but also precipitate severe attention dilution. Under the barrage of massive visual noise, the capacity-constrained model falls into cognitive overload, where attention resources are involuntarily scattered across irrelevant details, failing to distinguish valid information (Figure 2 (b)). This mechanistic failure directly undermines sequence modeling capability, resulting in a distinct monotonic decline in average accepted length as video

token length increases (Figure 1 (a)).

2.4 Attribution Analysis II: Layer-wise Interaction and Deep Internalization of Visual Semantics

Draft models typically utilize the second-to-last layer hidden states of the target model as input. To investigate why stripping visual information during inference does not compromise performance, we analyzed the internal information flow via a layer-wise truncation experiment (as shown in Figure 3 (a)). The results reveal a distinct stage-wise dependency: while explicit visual signals are indispensable in the early network stages, the prediction performance remains robust beyond the 20th layer, even after completely removing the visual flow and relying solely on text hidden states. We hypothesize that as network depth increases, visual semantics are gradually fused and encoded into the text hidden states, rendering the original visual tokens a form of structural redundancy. This mechanism plausibly explains the phenomenon observed in Figure 1 (a), where physically removing visual input effectively eliminates redundancy-induced interference in high-noise scenarios such as long videos, thereby enhancing model performance.

Given the apparent low dependence of deep layers on explicit visual information, we further employed attention heatmaps to locate cross-modal

Table 1: Impact of visual information injection on draft model performance during training.

Training & Inference Input	ChartQA	AI2D
Visual + Text	4.25	4.28
0% Visual + Text	4.23	4.25
50% Visual + Text	4.21	4.24

interaction zones (as shown in Figure 3 (b)), forming strong mutual corroboration with the aforementioned truncation experiment. Visual analysis indicates that the intermediate layers of the LLM serve as the core arena for visual-text interaction, whereas the shallow and deep layers focus predominantly on text processing. This finding is highly consistent with the performance trajectory in Figure 3 (a): the model’s predictive capability rises sharply during the intense interactions of the middle layers and stabilizes synchronously as interaction intensity wanes in the deep layers. This correspondence further supports the hypothesis that visual information is integrated into text representations primarily within the intermediate stages.

3 Method

3.1 Visually-Aware Text-Anchored Window Attention via Hidden State Reuse

Section 2.4 reveals the phenomenon of visual semantic internalization : deep visual information is fused into text hidden states, turning raw visual signals into structural redundancy. Based on this, we propose the cross-modal task offloading strategy. This strategy completely offloads the expensive visual feature extraction and cross-modal alignment tasks to the target model, allowing the draft model to focus solely on prediction using text hidden states that carry visual context, thereby bypassing its structural limitations.

To this end, we design a hidden state reuse (HSR) mechanism to reconstruct the input. Instead of processing the original long sequence, the draft model receives an augmented input z_t . As illustrated in Figure 4, we utilize a projection layer FC to perform dimensionality reduction on the concatenated features of the current text embedding e_t and the target model’s text hidden state from the previous timestep $h_{e_{t-1}}^h$, which encapsulates the visual context:

$$z_t = FC \left(e_t \oplus h_{e_{t-1}}^h \right) \quad (1)$$

where e_t provides explicit token identity to elim-

inate uncertainty (Li et al., 2024d), while h_{t-1}^h inherits cross-modal representations deeply integrated by the target model. This direct state reuse is akin to casting an efficient Glimpse at the vast stream of visual information. Through this glimpse, the draft model acquires relevant visual information, thereby achieving visual offloading at the physical level. This enables the model to circumvent redundant computations on the raw sequence while retaining the capacity to capture the necessary visual context. Note that the superscript h denotes the second-to-last layer of the target model.

Given that z_t has internalized visual semantics, we propose visually-aware text-anchored window attention (VATA). VATA discards the visual KV cache, forcing attention closure strictly within the text domain \mathcal{T} :

$$Attention_{VATA} = Softmax \left(\frac{Q_t K_{\mathcal{T}}^T}{\sqrt{d}} \right) V_{\mathcal{T}} \quad (2)$$

Specifically, within the equation, Q_t denotes the query vector at the current timestep t , and d represents the hidden dimension used for scaling dot-product attention, while $K_{\mathcal{T}}$ and $V_{\mathcal{T}}$ correspond to the Key and Value matrices restricted strictly to the text domain \mathcal{T} . VATA is based on the premise of visual semantic internalization, limiting computation to text-anchored positions. Crucially, after HSR interaction, representations at these positions have evolved into visual-text fusion representations, rather than raw textual information. This design ensures the draft model infers directly based on deeply integrated cross-modal representations, drastically reducing complexity from $O((L_{vis} + L_{txt})^2)$ to the pure text level $O(L_{txt}^2)$, where L_{vis} and L_{txt} denote the sequence lengths of visual and textual tokens, respectively.

3.2 Intermediate-Layer Visual State Bridging

Although visual input is masked during inference to mitigate noise, Table 1 shows that the average accepted length decreases during training when visual information is either fully excluded or partially pruned. This indicates that even if the visual modality manifests as noise during inference, the explicit context it provides remains indispensable for parameter optimization. Consequently, effectively incorporating visual supervision during the training phase becomes pivotal.

Regarding the integration of visual information, MSD advocates for directly utilizing raw visual em-

Table 2: We evaluate the performance of LLaVA-OneVision-7B on long-video benchmarks using NVIDIA L20 GPUs, and report the average accepted length (τ), the end-to-end wall-clock speedup ratio (ESR), and the decoding speedup ratio (DSR) for the generation phase, with the visual token length fixed at 25K across all datasets.

Method	VideoDetailCaption			MVBench			LongVideoBench			VideoMME			Avg		
	τ	ESR	DSR	τ	ESR	DSR	τ	ESR	DSR	τ	ESR	DSR	τ	ESR	DSR
Temp = 0															
SpecVLM	4.15	1.25×	1.47×	4.20	1.26×	1.43×	4.04	1.23×	1.39×	3.91	1.21×	1.34×	4.07	1.24×	1.41×
MSD	2.69	1.47×	1.76×	2.65	1.43×	1.71×	2.63	1.43×	1.71×	2.68	1.49×	1.77×	2.66	1.46×	1.74×
Ours	3.88	1.96×	2.84×	3.67	1.83×	2.69×	3.78	1.90×	2.77×	3.98	2.02×	2.96×	3.83	1.93×	2.82×
Temp = 1															
MSD	2.20	1.27×	1.43×	2.31	1.31×	1.50×	2.17	1.28×	1.42×	2.05	1.22×	1.35×	2.18	1.27×	1.43×
Ours	2.80	1.60×	2.06×	2.90	1.62×	2.13×	2.77	1.60×	2.05×	2.57	1.55×	1.91×	2.76	1.59×	2.04×

beddings. Their core motivation is to preserve pure information unassimilated by the language modeling process, thereby preventing visual features from being polluted by textual semantics during deep interactions. However, this strategy overlooks the capacity bottleneck of the draft model. Constrained by its extremely shallow network structure, the model struggles to extract effective semantics from highly redundant and unaligned raw representations, which restricts its prediction accuracy.

To address this, we propose intermediate-layer visual state bridging (IVSB), which extracts the visual hidden state $h_{e_{vis}}^{m^*}$ from the interaction-active layer m^* of the target model. As illustrated in Figure 3 (b), this layer offers a dual advantage: on the one hand, it has achieved implicit semantic alignment and filtered out low-level noise; on the other hand, it avoids the loss of fine-grained details caused by over-compression in the final layers. This purified, high-quality feature is well-adapted to the limited processing capabilities of the draft model. Consequently, we set m^* to half of the target model’s total layers.

As illustrated in Figure 4, We fuse this visual state with the text embedding e_{txt} and the target model’s second-to-last layer text state $h_{e_{txt}}^{(h)}$ (where superscript h denotes the second-to-last layer of the target model, which has already internalized the preceding visual information). The initial input for the draft model, z_{init} , is formalized as:

$$z_{init} = h_{e_{vis}}^{m^*} \oplus FC(e_{txt} \oplus h_{e_{txt}}^h) \quad (3)$$

By synergistically utilizing the explicitly purified visual features and the implicit visual context, the draft model avoids constructing complex low-level mappings, allowing it to focus on precise prediction based on well-adapted semantic representations.

To bridge the discrepancy between training and inference, we further introduce a multi-token pre-

diction mechanism (Liu et al., 2024a) (MTP). Given that during the inference phase, the draft model must rely on its own preceding outputs rather than the perfect states provided by the target model, we construct a recursive training pipeline. Specifically, starting from the initial input z_{init} , the draft model generates a textual hidden state $\hat{h}_{e_{txt}}$. Subsequently, this state is concatenated with the fixed intermediate visual anchor $h_{e_{vis}}^{m^*}$ to construct the recursive input z_{rec} :

$$z_{rec} = h_{e_{vis}}^{m^*} \oplus FC(e_{txt} \oplus \hat{h}_{e_{txt}}) \quad (4)$$

This design compels the model to adapt to its self-generated distribution during training, effectively alleviating exposure bias. Crucially, the continuous reuse of the high-quality visual anchor within z_{rec} provides a stable semantic reference for the autoregressive generation process.

4 Experiment

Hardware. All evaluation experiments are conducted on 2x NVIDIA L20 and 2x NVIDIA A800 GPUs.

Tasks and Benchmarks. We evaluate performance on video captioning and video description tasks, which typically require generating long-form text. The benchmarks employed include VideoDetailCaption (LMMs-Lab, 2024a), MVBench (Li et al., 2024c), LongVideoBench (Wu et al., 2024), and VideoMME (Fu et al., 2025b).

Target Models We select two widely used VideoLLMs series: LLaVA-OneVision-7B (Li et al., 2024a) and Qwen2.5-VL-7B (Ahmed et al., 2025).

Baselines We include MSD (Lin et al., 2025), ViSpec (orig.) (Kang et al., 2025), ViSpec, and SpecVLM (Ji et al., 2025) as our baselines. To ensure a fair comparison, except for ViSpec (orig.)

Table 3: Performance results of Qwen2.5-VL-7B on video benchmarks on NVIDIA L20 GPUs. We report average accepted length (τ) and decoding speedup ratio (DSR). Detailed visual token lengths: VideoDetailCaption ($\sim 17k$), MVBench ($\sim 13k$), LongVideoBench ($\sim 15k$).

Method	VideoDetailCaption		MVBench		LongVideoBench		Avg	
	τ	DSR	τ	DSR	τ	DSR	τ	DSR
Temp = 0								
MSD	1.11	0.52 \times	1.86	0.95 \times	1.75	0.88 \times	1.57	0.78 \times
SpecVLM	3.91	1.20 \times	3.80	1.09 \times	3.82	1.15 \times	3.84	1.15 \times
ViSpec (orig.)	2.79	1.53 \times	2.83	1.52 \times	2.75	1.47 \times	2.79	1.51 \times
ViSpec	3.52	1.87 \times	3.60	1.92 \times	3.55	1.91 \times	3.56	1.90 \times
Ours	3.89	2.04\times	3.87	2.06\times	3.78	2.02\times	3.85	2.04\times
Temp = 1								
MSD	1.10	0.48 \times	1.63	0.70 \times	1.52	0.74 \times	1.42	0.64 \times
ViSpec	2.24	1.21 \times	2.44	1.31 \times	2.46	1.34 \times	2.38	1.29 \times
Ours	2.37	1.27\times	2.55	1.36\times	2.54	1.37\times	2.49	1.33\times

and SpecVLM which utilize their officially released weights, all other baselines were reproduced and retrained with a batch size of 16. Note that ViSpec (orig.) adopts a batch size of 8. Further details are provided in Appendix C.

Metrics We assess acceleration performance using: (i) the speedup ratios of wall time and decoding latency relative to vanilla autoregressive decoding; and (ii) the average accepted length τ , which directly reflects speculation accuracy.

Given that our method is lossless, we do not evaluate output quality. With the exception of SpecVLM, which employs a static tree structure, our method and all other baselines utilize a dynamic tree structure. For the general setup, the dynamic configuration (denoted as total tokens-depth-width) is set to 30-4-8 on L20 GPUs (following ViSpec) and 48-5-10 on A800 GPUs. Specifically, to ensure a fair comparison in Figure 2, SpecVLM uses its default static setting (total budget 25, depth 5), while our method and other baselines are unified to a configuration of 25-5-8.

4.1 Main Results

The experimental results in Table 2 and Table 3 demonstrate that our method establishes a significant performance advantage across all benchmarks when using LLaVA-OneVision-7B and Qwen2.5-VL-7B as target models. Specifically, MSD encounters severe performance bottlenecks when facing long sequences with an average visual length of 17k tokens. The attention dilution caused by full visual inputs makes it difficult to focus on core information, resulting in an average accepted length of only 1.11 on the VideoDetailedCaption task of Qwen2.5-VL-7B, causing a negative speedup of 0.52 \times . Although SpecVLM achieves an opti-

Table 4: Performance of Qwen2.5-VL-7B on VideoDetailCaption across varying visual token lengths, evaluated on A800 GPUs, in terms of average accepted length (τ) and decoding speedup ratio (DSR).

Method	0.5k		1.5k		17k		25k	
	τ	DSR	τ	DSR	τ	DSR	τ	DSR
MSD	4.12	2.96 \times	3.97	2.87 \times	1.11	0.48 \times	1.04	0.42 \times
ViSpec (orig.)	3.45	2.67 \times	3.35	2.52 \times	2.96	1.31 \times	2.81	1.20 \times
ViSpec	4.29	3.22 \times	4.18	3.10 \times	3.74	1.64 \times	3.42	1.48 \times
MSD+VATA	4.11	2.98 \times	4.01	2.92 \times	3.94	1.72 \times	4.01	1.68 \times
Ours	4.51	3.30\times	4.36	3.18\times	4.32	1.86\times	4.37	1.82\times

mal average accepted length of 4.07 on LLaVA-OneVision by leveraging the target model reuse strategy, the heavy computational cost of the draft model severely offsets the gains from the high acceptance rate. Its decoding speedup of 1.41 \times is far lower than the 2.82 \times achieved by our method, making it difficult to meet real-time requirements. Furthermore, while ViSpec improves efficiency through visual token compression, its acceleration gains are limited by the difficulty in capturing the complex spatiotemporal details of long videos; its average speedup of 1.90 \times on Qwen2.5-VL-7B still lags behind the 2.04 \times achieved by our method.

To validate effectiveness in real-world scenarios, Table 5 presents a performance analysis of long-video inference. As visual inputs scale from 4.5K to 25K, prefill latency increases from 1.20s to 11.46s, causing its share of the total inference time to rise from 13.8% to 38.7%. Since speculative decoding cannot optimize the prefill phase, this increased overhead constrains the overall end-to-end speedup. Simultaneously, as shown in Table 2, visual noise induced by long sequences degrades the performance of conventional draft models, causing the average accepted length of MSD to drop from 3.07 to 2.69; meanwhile, constrained by the computational burden of reusing the target model, SpecVLM sees its speedup decline from 1.41 \times to 1.24 \times . In contrast, our method demonstrates strong robustness, maintaining an average accepted length of 3.83 even with 25K visual tokens. Consequently, it achieves a 1.96 \times end-to-end speedup based on a 2.84 \times decoding speedup, verifying its effectiveness in long-sequence tasks.

Table 4 presents the acceleration performance evaluation across varying input lengths, compellingly demonstrating that our method combines efficiency in short-sequence scenarios with robustness in long-sequence contexts. Specifically, in the

Table 5: Performance breakdown of LLaVA-OneVision-7B on VideoDetailedCaption across varying visual input lengths using L20 GPUs.

Visual Input Length	4.5K	15K	25K
Average Accepted Length	3.07	2.78	2.69
Generated Tokens	418.86	450.53	455.10
Target Model Calls	136.23	161.58	168.69
Wall (End-to-end) Time (s)	8.71	19.28	29.59
Total Decode Time (s)	7.51	13.52	18.13
Prefill Time (s)	1.20	5.76	11.46
Prefill Ratio	13.80%	29.90%	38.70%
Latency Per Decode Step (s)	0.06	0.08	0.11

0.5k short-sequence setting, we achieved a peak speedup of $3.30\times$ and an average accepted length of 4.50, surpassing ViSpec’s $3.22\times$ and MSD’s $2.96\times$, respectively. More critically, as the sequence length extends to 25k, traditional methods encounter severe challenges: MSD suffers a performance collapse with inference speed plummeting to $0.42\times$, while ViSpec’s average accepted length drops to 3.42. In contrast, our method robustly maintains an accepted length of 4.35 and a speedup of $1.82\times$, fully demonstrating its consistent superiority across both short and long sequences.

4.2 Ablation

As shown in Table 6, MTP training effectively mitigates distribution shifts and significantly enhances fundamental alignment capability, increasing the average effective length on the MME benchmark from 3.40 (Row 1) to 3.75 (Row 3). Building on this, introducing IVSB (Intermediate Visual Features) further exploits the model’s potential, achieving a peak performance of 4.53 on short sequences (0.5k) in Row 5. However, constrained by noise accumulation from attention dilution, the performance of this configuration suffers a precipitous decline to 1.21 on 25k long sequences. The VATA mechanism effectively resolves this issue by offloading redundant visual computations, allowing our full method (Ours, Row 6) to stage a strong recovery to 4.37 at 25k. Notably, a comparison of Row 6 across different lengths demonstrates that VATA renders the model’s performance nearly invariant to the negative impacts of increasing sequence length.

Experimental observations reveal that the cumulative noise associated with long sequences often masks the model’s true capability limit; therefore, performance on short sequences serves as a superior indicator of Fundamental Capability.

Table 6: We conduct an ablation study on MTP, intermediate-layer visual state bridging (IVSB), and visually-aware text-anchored window attention (VATA) using Qwen2.5-VL-7B deployed on A800 GPUs as the base model on the MME and VideoDetailedCaption benchmarks, reporting the average accepted length across varying visual token lengths.

Method Components			MME	VideoDetailedCaption				
MTP	IVSB	VATA	0.4k	0.5k	1.5k	6k	17k	25k
X	X	X	3.40	3.95	3.87	3.29	2.97	2.95
X	X	✓	3.41	3.95	3.85	3.83	3.77	3.86
✓	X	X	3.75	4.44	4.31	3.87	3.58	3.51
✓	X	✓	3.75	4.45	4.34	4.24	4.24	4.29
✓	✓	X	3.82	4.53	4.38	3.08	1.91	1.21
✓	✓	✓	3.79	4.51	4.36	4.29	4.32	4.37

A comparison of configurations utilizing VATA (Rows 2, 4, and 6) demonstrates a clear correlation: the stronger the model’s fundamental capability on short sequences, the more superior its long-sequence performance remains after applying VATA. This is particularly evident in Row 6, which achieves global optimality by leveraging the high baseline capability established by IVSB in conjunction with VATA. This finding validates the effectiveness of our synergistic design philosophy: prioritizing the enhancement of fundamental capability first, and then employing VATA to ensure long-sequence robustness.

5 Conclusion

In this paper, we delve into the root causes of speculative decoding inefficacy in long-video scenarios, identifying attention dilution and negative visual gain as the critical bottlenecks hindering lightweight draft models from handling ultra-long sequences. Based on the empirical validation of the visual semantic internalization mechanism within the target model, we propose the Sparrow framework. Through visual computational offloading via HSR-VATA, Sparrow not only effectively overcomes the capacity constraints of the draft model but also decouples visual processing costs from sequence length. Meanwhile, the synergistic integration of IVSB and the Multi-Token Prediction strategy further strengthens the predictive performance of the draft model. Benefiting from these designs, our method achieves a $2.82\times$ decoding speedup on sequences of 25k visual tokens.

6 Limitations

While our method exhibits superior acceleration performance during the long-sequence decoding phase, the computational bottleneck in the prefill phase remains a primary limitation. As the visual input length scales, prefill latency increases substantially. Since speculative decoding inherently optimizes only the autoregressive generation process without addressing the prefill phase, this escalating overhead constrains the ceiling of the overall end-to-end speedup. Future work will aim to explore prefill acceleration techniques for long videos, such as incorporating visual token pruning. Although pruning strategies may introduce potential information loss, we believe that by designing refined importance-evaluation mechanisms, we can strike an optimal balance between inference speed and model accuracy while significantly reducing prefill costs, thereby achieving more efficient end-to-end inference.

7 Acknowledgments

This work is sponsored in part by the National Natural Science Foundation of China (Grant No. 62421002) and the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China (Grant No. JYB2025XDXM202).

References

- Imtiaz Ahmed, Sadman Islam, Partha Protim Datta, Imran Kabir, Naseef Ur Rahman Chowdhury, and Ahshanul Haque. 2025. Qwen 2.5: A comprehensive review of the leading resource-efficient llm with potential to surpass all competitors. *Authorea Preprints*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Charlie Chen, Sebastian Borgeaud, Geoffrey Irving, Jean-Baptiste Lespiau, Laurent Sifre, and John Jumper. 2023. Accelerating large language model decoding with speculative sampling. *arXiv preprint arXiv:2302.01318*.
- Liang Chen, Haozhe Zhao, Tianyu Liu, Shuai Bai, Junyang Lin, Chang Zhou, and Baobao Chang. 2024. An image is worth 1/2 tokens after layer 2: Plug-and-play inference acceleration for large vision-language models. In *European Conference on Computer Vision*, pages 19–35. Springer.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- Yingqi Fan, Anhao Zhao, Jinlan Fu, Junlong Tong, Hui Su, Yijie Pan, Wei Zhang, and Xiaoyu Shen. 2025. Visipruner: Decoding discontinuous cross-modal dynamics for efficient multimodal llms. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 18896–18913.
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, and 1 others. 2025a. Mme: A comprehensive evaluation benchmark for multimodal large language models. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Chaoyou Fu, Yuhan Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, and 1 others. 2025b. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. In *CVPR*.
- Yichao Fu, Peter Bailis, Ion Stoica, and Hao Zhang. 2024. Break the sequential dependency of llm inference using lookahead decoding. *arXiv preprint arXiv:2402.02057*.
- Mugilan Ganesan, Shane Segal, Ankur Aggarwal, Nish Sinnadurai, Sean Lie, and Vithursan Thangarasa. 2025. Massv: Multimodal adaptation and self-data distillation for speculative decoding of vision-language models. *arXiv preprint arXiv:2505.10526*.
- Coleman Hooper, Sehoon Kim, Hiva Mohammadzadeh, Hasan Genc, Kurt Keutzer, Amir Gholami, and Yakun Sophia Shao. 2025. Speed: Speculative pipelined execution for efficient decoding. In *Enhancing LLM Performance: Efficacy, Fine-Tuning, and Inference Techniques*, pages 19–32. Springer.
- Yunhai Hu, Zining Liu, Zhenyuan Dong, Tianfan Peng, Bradley McDanel, and Sai Qian Zhang. 2025a. Speculative decoding and beyond: An in-depth survey of techniques. *arXiv preprint arXiv:2502.19732*.
- Yunhai Hu, Tianhua Xia, Zining Liu, Rahul Raman, Xingyu Liu, Bo Bao, Eric Sather, Vithursan Thangarasa, and Sai Qian Zhang. 2025b. Dream: Drafting with refined target features and entropy-adaptive cross-attention fusion for multimodal speculative decoding. *arXiv preprint arXiv:2505.19201*.
- Mingxiao Huo, Jiayi Zhang, Hewei Wang, Jinfeng Xu, Zheyu Chen, Huilin Tai, and Yijun Chen. 2025. Spec-llava: Accelerating vision-language models with dynamic tree-based speculative decoding. *arXiv preprint arXiv:2509.11961*.
- Yicheng Ji, Jun Zhang, Jinpeng Chen, Cong Wang, Lidan Shou, Gang Chen, and Huan Li. 2026. See the

- forest for the trees: Loosely speculative decoding via visual-semantic guidance for efficient inference of video llms. *arXiv preprint arXiv:2604.05650*.
- Yicheng Ji, Jun Zhang, Heming Xia, Jinpeng Chen, Lidan Shou, Gang Chen, and Huan Li. 2025. SpecVLM: Enhancing speculative decoding of video llms via verifier-guided token pruning. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 7216–7230.
- Jialiang Kang, Han Shu, Wenshuo Li, Yingjie Zhai, and Xinghao Chen. 2025. Vispec: Accelerating vision-language models with vision-aware speculative decoding. *arXiv preprint arXiv:2509.15235*.
- Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. 2016. A diagram is worth a dozen images. In *European conference on computer vision*, pages 235–251. Springer.
- Minjae Lee, Wonjun Kang, Minghao Yan, Christian Classen, Hyung Il Koo, and Kangwook Lee. 2024. In-batch ensemble drafting: Toward fast and robust speculative decoding for multimodal language models. *OpenReview.net*.
- Yaniv Leviathan, Matan Kalman, and Yossi Matias. 2023. Fast inference from transformers via speculative decoding. In *International Conference on Machine Learning*, pages 19274–19286. PMLR.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, and 1 others. 2024a. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*.
- Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, and Chunyuan Li. 2024b. Llava-next-interleave: Tackling multi-image, video, and 3d in large multimodal models. *arXiv preprint arXiv:2407.07895*.
- Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, and 1 others. 2024c. Mvbench: A comprehensive multi-modal video understanding benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22195–22206.
- Yuhui Li, Fangyun Wei, Chao Zhang, and Hongyang Zhang. 2024d. Eagle: Speculative sampling requires rethinking feature uncertainty. *arXiv preprint arXiv:2401.15077*.
- Luxi Lin, Zhihang Lin, Zhanpeng Zeng, and Rongrong Ji. 2025. Speculative decoding reimaged for multimodal large language models. *arXiv preprint arXiv:2505.14260*.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, and 1 others. 2024a. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Fangcheng Liu, Yehui Tang, Zhenhua Liu, Yunsheng Ni, Kai Han, and Yunhe Wang. 2024b. Kangaroo: Lossless self-speculative decoding via double early exiting. *arXiv preprint arXiv:2404.18911*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024c. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 26296–26306.
- Jiahao Liu, Qifan Wang, Jingang Wang, and Xunliang Cai. 2024d. Speculative decoding via early-exiting for faster llm inference with thompson sampling control mechanism. *arXiv preprint arXiv:2406.03853*.
- Tianyu Liu, Yun Li, Qitan Lv, Kai Liu, Jianchen Zhu, and Winston Hu. 2024e. Parallel speculative decoding with adaptive draft length. *arXiv e-prints*, pages arXiv–2408.
- Ting Liu, Liangtao Shi, Richang Hong, Yue Hu, Qunjun Yin, and Linfeng Zhang. 2024f. Multi-stage vision token dropping: Towards efficient multimodal large language model. *arXiv preprint arXiv:2411.10803*.
- LMMs-Lab. 2024a. Video detail caption. *Accessed: 2024-11*.
- LMMs-Lab. 2024b. Videochatgpt. *Accessed: 2024-04*.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521.
- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Khan. 2024. Video-chatgpt: Towards detailed video understanding via large vision and language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12585–12602.
- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. 2023. Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv:2306.05424*.
- Ahmed Masry, Xuan Long Do, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. 2022. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. In *Findings of the association for computational linguistics: ACL 2022*, pages 2263–2279.
- Giovanni Monea, Armand Joulin, and Edouard Grave. 2023. Pass: Parallel speculative sampling. *arXiv preprint arXiv:2311.13581*.

- Hanshi Sun, Zhuoming Chen, Xinyu Yang, Yuandong Tian, and Beidi Chen. 2024. Triforce: Lossless acceleration of long sequence generation with hierarchical speculative decoding. *arXiv preprint arXiv:2404.11912*.
- Keda Tao, Can Qin, Haoxuan You, Yang Sui, and Huan Wang. 2025. Dycoke: Dynamic compression of tokens for fast video large language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 18992–19001.
- Jikai Wang, Yi Su, Juntao Li, Qingrong Xia, Zi Ye, Xinyu Duan, Zhefeng Wang, and Min Zhang. 2025a. Opt-tree: Speculative decoding with adaptive draft tree structure. *Transactions of the Association for Computational Linguistics*, 13:188–199.
- Yi Wang, Xinhao Li, Ziang Yan, Yinan He, Jiashuo Yu, Xiangyu Zeng, Chenting Wang, Changlian Ma, Haian Huang, Jianfei Gao, and 1 others. 2025b. Internvideo2. 5: Empowering video mllms with long and rich context modeling. *arXiv preprint arXiv:2501.12386*.
- Zilong Wang, Zifeng Wang, Long Le, Huaixiu Steven Zheng, Swaroop Mishra, Vincent Perot, Yuwei Zhang, Anush Mattapalli, Ankur Taly, Jingbo Shang, and 1 others. 2024. Speculative rag: Enhancing retrieval augmented generation through drafting. *arXiv preprint arXiv:2407.08223*.
- Haoning Wu, Dongxu Li, Bei Chen, and Junnan Li. 2024. Longvideobench: A benchmark for long-context interleaved video-language understanding. *Advances in Neural Information Processing Systems*, 37:28828–28857.
- Zhinan Xie, Peisong Wang, Shuang Qiu, and Jian Cheng. 2025. Hivis: Hiding visual tokens from the drafter for speculative decoding in vision-language models. *arXiv preprint arXiv:2509.23928*.
- Long Xing, Qidong Huang, Xiaoyi Dong, Jiajie Lu, Pan Zhang, Yuhang Zang, Yuhang Cao, Conghui He, Jiaqi Wang, Feng Wu, and 1 others. 2024. Pyramiddrop: Accelerating your large vision-language models via pyramid visual redundancy reduction. *arXiv preprint arXiv:2410.17247*.
- Hao Yin, Guangzong Si, and Zilei Wang. 2025. Lifting the veil on visual information flow in mllms: Unlocking pathways to faster inference. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 9382–9391.
- Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. 2023. Mm-vet: Evaluating large multimodal models for integrated capabilities. *arXiv preprint arXiv:2308.02490*.
- Ce Zhang, Kaixin Ma, Tianqing Fang, Wenhao Yu, Hongming Zhang, Zhisong Zhang, Yaqi Xie, Kattia Sycara, Haitao Mi, and Dong Yu. 2025a. Vs-can: Rethinking visual token reduction for efficient large vision-language models. *arXiv preprint arXiv:2505.22654*.
- Hang Zhang, Xin Li, and Lidong Bing. 2023. Videollama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*.
- Jun Zhang, Yicheng Ji, Feiyang Ren, Yihang Li, Bowen Zeng, Zonghao Chen, Ke Chen, Lidan Shou, Gang Chen, and Huan Li. 2026. Efficient inference for large vision-language models: Bottlenecks, techniques, and prospects. *arXiv preprint arXiv:2604.05546*.
- Libo Zhang, Zhaoning Zhang, Rui Li, Zhiliang Tian, Songzhu Mei, Dongsheng Li, and 1 others. 2025b. Dovetail: A cpu/gpu heterogeneous speculative decoding for llm inference. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 17393–17406.
- Xuan Zhang, Cunxiao Du, Sicheng Yu, Jiawei Wu, Fengzhuo Zhang, Wei Gao, and Qian Liu. 2025c. Sparse-to-dense: A free lunch for lossless acceleration of video understanding in llms. *arXiv preprint arXiv:2505.19155*.
- Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. 2024. Video instruction tuning with synthetic data. *arXiv preprint arXiv:2410.02713*.
- Zhi Zhang, Srishti Yadav, Fengze Han, and Ekaterina Shutova. 2025d. Cross-modal information flow in multimodal large language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 19781–19791.

A Related Work

Speculative Decoding Speculative decoding has established itself as a pivotal acceleration paradigm in the realm of LLMs, significantly enhancing autoregressive generation efficiency while strictly adhering to the target model’s original output distribution. Operating on a draft-and-verify mechanism, it employs a lightweight draft model to rapidly generate candidate token sequences, which are subsequently verified and corrected by the target model via a single parallel forward pass (Leviathan et al., 2023; Chen et al., 2023). Theoretically grounded in distribution consistency conditions (Leviathan et al., 2023), this technique fundamentally ensures zero deviation in the inference process. In recent years, the field has witnessed continuous evolution with the emergence of variants such as tree-structured speculation (Li et al., 2024d; Wang et al., 2025a), self-speculative decoding (Hooper et al., 2025; Liu et al., 2024b,d), and retrieval-augmented strategies (Wang et al., 2024). Complemented by system-level optimization frameworks (Sun et al., 2024; Monea et al., 2023; Liu et al., 2024e; Zhang et al., 2025b), these advancements continually push the boundaries of inference efficiency.

With the ubiquity of VLMs, speculative decoding (Huo et al., 2025; Lin et al., 2025; Ganesan et al., 2025; Kang et al., 2025; Zhang et al., 2025c; Ji et al., 2025; Hu et al., 2025b; Xie et al., 2025; Ji et al., 2026; Zhang et al., 2026) has been extended to multimodal scenarios. Existing methodologies are primarily categorized into full visual input and compressed visual input strategies, depending on how the draft model handles the visual stream. Full visual input approaches (Huo et al., 2025; Lin et al., 2025; Ganesan et al., 2025) typically reuse the target model’s complete visual encoding pipeline. Conversely, compressed visual input approaches aim to reduce KV cache usage and latency by minimizing the scale of visual tokens. These methods generally follow two technical trajectories: (1) Token Aggregation and Mapping, exemplified by the adapter-based compression in ViSpec (Kang et al., 2025); and (2) Attention-Guided Dynamic Selection, which exploits attention sparsity to filter redundant tokens, with representative works including STD (Zhang et al., 2025c), SPECVLM (Ji et al., 2025), and DREAM (Hu et al., 2025b).

Despite differing perspectives, both our work and HiViS (Xie et al., 2025) arrive at the same key insight: the draft model does not require explicit

visual tokens and can effectively reason using only the text-side hidden states of the target model, in which visual semantics are implicitly encoded. The key difference lies in their usage: HiViS completely discards visual tokens during both training and inference, whereas we retain them during training to provide essential visual supervision and only mask them during inference. Moreover, HiViS targets short-sequence tasks to reduce redundancy, while our method addresses long-video scenarios, mitigating the visual negative gain caused by attention dilution in ultra-long visual sequences.

Stage-wise Visual Processing in VLMs Research on the visual processing mechanisms of VLMs has evolved from phenomenological observation to mechanistic elucidation. Early works, such as FastV (Chen et al., 2024), uncovered the non-uniformity of vision-text fusion based on attention distributions, highlighting that visual weights in deep layers are negligible (merely 0.21%) and proposing single-stage filtering to eliminate redundancy. With the deepening of mechanistic research, (Zhang et al., 2025d) and (Yin et al., 2025) utilized attention knockout experiments to reveal the law of Layer-wise Functional Differentiation in cross-modal fusion: shallow layers primarily inject global features, middle layers serve as the core window for fine-grained semantic alignment, and deep layers govern final prediction. Building on this, HiMAP (Yin et al., 2025) introduced a dynamic pruning strategy specifically targeting weak interactions in shallow layers and high aggregation in deep layers. Recent studies (Fan et al., 2025; Zhang et al., 2025a) have further synthesized this into a three-stage framework: shallow layers function mainly as "Attention Sinks" or focus on text processing; middle layers act as the pivotal window for cross-modal fusion; and deep layers are dedicated to linguistic structure optimization. Despite minor discrepancies regarding the specific role of shallow layers, a strong consensus has emerged: the substantive contribution of visual information is highly concentrated in the intermediate layers. Once the multimodal representation is constructed in the middle layers, raw visual tokens in the deep layers become redundant due to semantic saturation.

B Task and Benchmark Details

Tasks and Benchmark Details. Given that Sparrow is primarily dedicated to achieving lossless accel-

Table 7: Impact of visual information injection on draft model performance during inference.

Draft Model Input	ChartQA	AI2D
Visual + Text	4.25	4.28
50% Visual + Text	4.26	4.30
12.5% Visual + Text	4.25	4.30
0% Visual + Text	4.27	4.31

eration during the decoding stage, we select video captioning and video description as our core evaluation tasks. These tasks require the model to deeply understand specific subjects and actions within the video and generate long textual descriptions rich in visual information. Unlike conventional Video Question Answering (VQA), which typically involves direct answer selection, these long-form generation tasks better highlight the acceleration advantages of the model in long-sequence decoding.

For the datasets in our main experiments, to ensure comprehensive coverage of varying video durations and application scenarios, we randomly sampled 100 instances for performance evaluation from mainstream benchmarks, including VideoDetailCaption (LMMs-Lab, 2024a), MVBench (Li et al., 2024c), LongVideoBench (Wu et al., 2024), and VideoMME (Fu et al., 2025b).

Furthermore, prior to the main experiments, to verify the universality of the phenomena observed in this work across different tasks and modalities, we also conducted sampling analysis on VideoChatGPT (LMMs-Lab, 2024b) and VideoDetailDescription (Maaz et al., 2023), as well as the visual question answering task AI2D (Kembhavi et al., 2016), the document and chart understanding benchmark ChartQA (Masry et al., 2022), and the comprehensive capability benchmark MME (Fu et al., 2025a).

C Implementation Details

SpecVLM adopts a self-speculative strategy, wherein the draft model directly reuses the weights of the target model. By aggressively pruning redundant visual tokens at the input level, it significantly reduces KV cache occupancy and computational overhead. This input-level sparsification enables inference acceleration without altering the underlying model structure. In contrast, MSD, Vispec, and our proposed method opt to construct an indepen-

dent single-layer draft model. Architecturally, this model mirrors a single decoder layer of the target model, aiming to maximize speculative efficiency through extreme structural lightweighting while ensuring feature space alignment and compatibility.

C.1 Visual Information Processing Strategies in Draft Models

SpecVLM leverages the draft model’s insensitivity to missing visual information by proposing a verifier-guided two-stage video token pruning strategy. Guided by the attention signals of the target model (the verifier), this method prioritizes retaining highly informative tokens while performing spatially uniform pruning on low-attention redundant tokens. Ultimately, this approach reduces the visual input for the draft model by approximately 90%.

MSD posits that visual and textual characteristics are distinct, thus adopting an architectural decoupling approach by directly using raw visual embeddings as input. Unlike features processed by the LLM, these raw embeddings remain uncontaminated by the language modeling process. Consequently, they preserve the most pristine visual information, providing more accurate guidance for the draft model.

ViSpec efficiently processes visual information through a dual fusion mechanism. First, it utilizes a lightweight visual adapter to compress redundant image tokens into compact representations and integrates them into the draft model’s attention layers, thereby retaining key spatial information while reducing computational overhead. Second, it extracts a global image feature vector and injects it into every subsequently generated text token. This provides the model with continuous and robust global visual context, significantly enhancing the accuracy of multimodal predictions.

C.2 Selection of Training Data

Building on the analysis in Sec. 2, directly transferring a small model trained solely on image-text pairs to long-video scenarios inevitably leads to a double collapse. The excessively long visual sequences (>10k tokens) not only exceed the pre-trained positional encoding window but also introduce significant background noise. This distracts the draft model, causing critical fused representations to be submerged amidst the noise. However, a turning point emerged when applying the VATA strategy—which physically masks raw vi-




Figure 5: Layer-wise importance analysis of visual tokens in LLaVA-OneVision-7B. The figure compares task accuracy after removing visual tokens starting from layer x (solid lines) with the native baseline performance (dashed lines).

sual inputs to force reliance on hidden states. We surprisingly observed that the average acceptance length no longer decays as video token length increases. This phenomenon provides decisive empirical support for our hypothesis: once deep fused hidden states are available, raw visual inputs become purely redundant noise during inference.

Consequently, we adopt a minimalist data strategy, eliminating the need for expensive and scarce long-video datasets for draft model training. The rationale lies in the HSR strategy, which effectively offloads the complex computations of visual perception, spatiotemporal modeling, and cross-modal alignment to the Target Model. Whether the input is a single-frame image or a long video, the draft model’s task is unified and simplified to merely decoding the deep fused representations from the Target Model. Therefore, we train exclusively using low-cost Image-Text Pairs. This approach not only circumvents the massive computational overhead of processing contexts exceeding 20k tokens but also enables the small model to focus solely on mapping deep fused features to the output vocabulary. This facilitates a seamless transfer of capabilities learned from image-text data to long-video scenarios, achieving robust cross-modal generalization with minimal resource costs.

C.3 Datasets and Training Details

Following the data construction paradigm of ViSpec, we established a two-stage training dataset. In the first phase (Text Alignment), we randomly




Figure 6: Layer-wise information retention analysis of visual and text modalities in Qwen2.5-VL-7B. The curves illustrate the retention rates calculated based on similarity, with red denoting visual features and blue denoting text features. The vertical dashed line marks the critical Layer 20.

sampled 68,000 multi-turn conversation samples from the ShareGPT dataset to ensure a robust linguistic foundation. The second phase (Multimodal Fine-tuning) utilized the LLaVA Visual Instruct Pretrain LCS dataset (Li et al., 2024b), with an equal sample size of 68,000. Crucially, we implemented an instruction-driven data augmentation strategy during this phase. By appending the command "Please answer with at least 1000 words" to the original prompts, we induced the target VLM to generate long-form synthetic responses enriched with detailed explanations and complex reasoning chains, thereby constructing a high-quality dataset for multimodal instruction fine-tuning.

MSD adopts a two-stage training strategy spanning a total of 40 epochs. The first stage utilizes pure text data, while the second stage employs a linear dynamic mixing strategy that progressively decreases the proportion of text data and increases the visual data as training proceeds. In contrast, both ViSpec and our method follow a different two-stage paradigm: the first stage consists of text-only training, followed by a direct switch to exclusive multimodal data training in the second stage. For these methods, each stage is trained for 20 epochs. All methods were trained with a batch size of 16.

D Supplementary Insights and Analysis

To further elucidate the phenomenon observed in the main text—where removing visual information after Layer 20 does not compromise performance—we introduce the information retention rate as a supplementary perspective, specifically defined as the average cosine similarity between the hidden states at layer l and the original input embeddings. As shown in Figure 6, the visual

Table 8: Performance results of Qwen2.5-VL-7B on benchmarks on NVIDIA GPUs. We report average accepted length (τ), decoding speed (Tokens/Sec), and speedup ratio (ESR).

Method	COCO caps			MME			SQA			MM-Vet			Avg.	
	τ	Tokens/Sec	ESR	τ	Tokens/Sec	ESR	τ	Tokens/Sec	ESR	τ	Tokens/Sec	ESR	τ	ESR
T = 0														
MSD	3.77	79.63	2.00×	3.49	68.37	1.83×	3.75	76.66	1.95×	3.60	70.14	1.87×	3.65	1.91×
ViSpec(orig.)	3.37	72.94	1.83×	3.19	64.85	1.71×	3.26	68.73	1.75×	3.20	64.81	1.72×	3.26	1.75×
ViSpec	4.00	84.93	2.13×	3.62	71.44	1.88×	3.78	78.01	1.98×	3.72	72.83	1.94×	3.78	1.98×
Ours - VATA	4.13	85.61	2.15×	3.78	72.59	1.92×	3.86	77.50	1.97×	3.82	73.42	1.95×	3.90	2.00×
Ours	4.11	85.80	2.15×	3.73	72.17	1.90×	3.84	78.19	1.99×	3.80	74.16	1.97×	3.87	2.00×
T = 1														
MSD	2.85	61.83	1.55×	2.73	56.78	1.49×	3.24	67.30	1.71×	2.94	62.7	1.66×	2.94	1.60×
ViSpec(orig.)	2.68	59.63	1.50×	2.52	53.90	1.42×	2.85	61.61	1.56×	2.70	58.25	1.55×	2.69	1.51×
ViSpec	2.94	65.18	1.64×	2.78	58.84	1.55×	3.21	68.00	1.72×	2.91	63.52	1.69×	2.96	1.65×
Ours - VATA	3.00	64.82	1.63×	2.95	59.84	1.57×	3.23	67.72	1.72×	2.80	58.86	1.56×	3.00	1.62×
Ours	2.99	65.44	1.64×	2.84	59.11	1.55×	3.23	68.58	1.74×	3.00	64.02	1.70×	3.02	1.66×

retention rate exhibits a precipitous decline in the shallow and middle layers, standing in sharp contrast to the relatively flat text curve. This indicates that visual features undergo intense transformation and reconstruction during the early stages to align with the text modality. More critically, the similarity score drops below 0.25 near Layer 20, corresponding precisely with the results of the truncation experiment. Consequently, 0.25 can be regarded as a critical Empirical Threshold: it signifies the completion of visual information fusion and provides quantitative guidance for future Visual Pruning strategies—suggesting that once the similarity falls below this threshold, visual tokens can be safely removed without causing performance degradation.

We present the layer-wise truncation experiment based on the LLaVA-OneVision-7B model. As shown in Figure 5, the results reproduce the phenomenon observed in the main text: while explicit visual signals are indispensable in the early network stages, prediction performance remains robust beyond the 20th layer, even after completely removing the visual flow and relying solely on text hidden states.

E Extension to Image Tasks and Short-Sequence Generalization

Table 2 and Table 3 have demonstrated the superiority of our method in long-video understanding tasks. To evaluate its generalization capability, we conduct extended experiments on image-based tasks and compare it against efficient inference methods specifically designed for images or short-sequence visual inputs, including MSD, ViSpec, and ViSpec (orig.). The results show that our method maintains high inference efficiency while

exhibiting stronger generality.

As shown in Table 8, our method achieves the best average speedup ratio across four image-related benchmark tasks. The experiments were conducted on a single NVIDIA virtual GPU (vGPU-32GB) rented from AutoDL, which provides up to 32 GB of video memory, 52.22 TFLOPS of FP32 compute performance, and approximately 103 Tensor TFLOPS of FP16 compute performance, paired with an Intel(R) Xeon(R) Gold 6459C CPU. When VATA is enabled (Ours), visual features are entirely skipped during draft generation. Although this design reduces the average accepted length, it further shortens the draft model’s inference latency. On COCO Captions (Chen et al., 2015)(COCO Caps), ScienceQA (Lu et al., 2022) (SQA), and MM-Vet (Yu et al., 2023), the time saved during drafting sufficiently compensates for the additional decoding steps incurred by the lower acceptance length, resulting in an overall speedup. In contrast, on the MME (Fu et al., 2025a) benchmark, the reduction in drafting time does not offset the cost of extra decoding steps, leading to a slight slowdown. Notably, in multimodal short-sequence settings, explicitly removing visual inputs does not significantly degrade performance, owing to the visual semantic internalization phenomenon described in Section 3.1. Our method consistently outperforms ViSpec and MSD across all tasks, demonstrating robustness and broad applicability across different modalities and task configurations.