

StreamingEval: A Unified Evaluation Framework towards Realistic Streaming Video Understanding

Guowei Tang¹, Tianwen Qian^{1†}, Huanran Zheng¹, Yifei Wang¹, Xiaoling Wang¹

¹East China Normal University, Shanghai, China

51274404096@stu.ecnu.edu.cn, twqian@cs.ecnu.edu.cn

Abstract

Real-time, continuous understanding of visual signals is essential for real-world interactive AI applications, and poses a fundamental system-level challenge. Existing research on streaming video understanding, however, typically focuses on isolated aspects such as question-answering accuracy under limited visual context or improvements in encoding efficiency, while largely overlooking practical deployability under realistic resource constraints. To bridge this gap, we introduce StreamingEval, a unified evaluation framework for assessing the streaming video understanding capabilities of Video-LLMs under realistic constraints. StreamingEval benchmarks both mainstream offline models and recent online video models under a standardized protocol, explicitly characterizing the trade-off between efficiency, storage and accuracy. Specifically, we adopt a fixed-capacity memory bank to normalize accessible historical visual context, and jointly evaluate visual encoding efficiency, text decoding latency, and task performance to quantify overall system deployability. Extensive experiments across multiple datasets reveal substantial gaps between current Video-LLMs and the requirements of realistic streaming applications, providing a systematic basis for future research in this direction. Codes will be released at <https://github.com/wwgTang-111/StreamingEval1>.

1 Introduction

As video large language models (Video-LLMs) (Maaz et al., 2024; Zhang et al., 2023; Song et al., 2024; Jin et al., 2024) continue to advance, applications such as embodied robots (Driess et al., 2023; Brohan et al., 2023; Wang et al., 2026), live-streaming assistants (Chen et al., 2024; Xu et al., 2025), and autonomous

[†]Corresponding author.

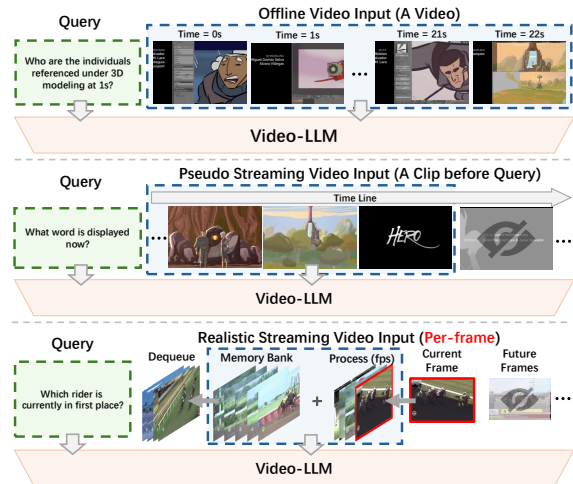


Figure 1: Illustration of the conventional offline video understanding paradigm versus the streaming paradigm. Top: offline inference with full access to the video. Middle: pseudo-streaming inference, which truncates videos at query timestamps but still processes each clip in an offline manner. Bottom: realistic streaming inference with real-time incremental input and limited memory bank.

driving systems (Levinson et al., 2011; Qian et al., 2024; Brödermann et al., 2025) are becoming increasingly feasible. In these scenarios, visual inputs arrive continuously in an incremental manner, requiring models to process them in real time under strict latency and resource constraints. However, existing mainstream Video-LLMs are designed and evaluated in offline settings, where the input videos are pre-recorded and fully accessible by the model. In contrast, streaming video understanding requires the model to process continuously arriving inputs without access to future frames, while sustaining instant responsiveness over potentially unbounded time horizons.

This fundamental mismatch between offline and streaming settings gives rise to several technical challenges that are largely absent in conventional video understanding tasks. First, streaming mod-

els must reason with incomplete and evolving visual context, as only past and current frames are available at any time. Second, streaming video is characterized by unbounded temporal duration, and long-running applications require persistent visual memory, but limited GPU resources inevitably lead to out-of-memory failures. Third, streaming systems impose stringent efficiency requirements. If visual encoding or memory updates fall behind the incoming frame rate, frames will accumulate and break real-time responsiveness. Similarly, excessive decoding latency directly degrades the interactive experience. Figure 1 illustrates the differences between offline and streaming video understanding paradigms.

Recent studies have begun to investigate individual aspects of these challenges in isolation. For example, OVO-Bench (Li et al., 2025) introduces an evaluation benchmark aligned with streaming context constraints and analyzes the model performance across different timestamps of questions (Backward Tracing, Real-Time Visual Perception, and Forward Active Responding). Several recent studies investigate the impact of visual encoding efficiency on online video processing (Yao et al., 2025; Wang et al., 2025c; Zeng et al., 2025; Chen et al., 2025; Ning et al., 2025), or quantify the effects of text decoding latency (Reddi et al., 2019). However, the diversity of evaluation methodologies has so far prevented the emergence of a unified and systematic evaluation framework. In particular, some benchmarks rely on pseudo-streaming settings, where videos are truncated at the query timestamp but still processed offline, limiting fair, reproducible, and deployment-oriented comparisons across models.

To bridge this gap, we introduce StreamingEval, a unified evaluation framework for video understanding under realistic streaming constraints. StreamingEval formalizes streaming video understanding as a system-level problem and provides a standardized evaluation protocol with metrics that jointly characterize accuracy, latency, and memory usage. Concretely, this framework implements a standardized streaming pipeline with three modules: (i) a frame player that continuously emits raw video frames at a fixed frame rate; (ii) an encoding-and-memory processor that performs per-frame visual encoding and memory updates according to each model’s design; (iii) and a response generator that, upon receiving a query, encodes the query, loads the current visual memory, and invokes the

language model to generate answers. Based on this framework, we evaluate 12 representative online and offline Video-LLMs. Native online models retain their original streaming mechanisms and configurations, while offline models are modified with a fixed-capacity visual memory bank managed via a first-in-first-out (FIFO) policy, standardizing the accessible historical context during streaming inference.

Extensive experiments across multiple datasets demonstrate that current methods claimed to be “online” fail to operate reliably under strict streaming constraints. Moreover, under identical streaming settings, mainstream offline models often outperform specialized online models at the cost of higher resource consumption. Performance can be improved by enlarging memory capacity or raising visual input resolution, but this entails a trade-off with efficiency and deployability. Overall, our results reveal substantial gaps between current Video-LLMs and the requirements of realistic streaming applications, highlighting open challenges for future research. In summary, our main contributions are as follows:

- We propose StreamingEval, a framework for assessing both the holistic capabilities of models and their deployability in streaming settings.
- We establish a scalable suite of online evaluation metrics and a unified protocol, enabling fair comparisons across models under consistent constraints.
- We conduct systematic empirical evaluations and analyses of representative state-of-the-art online and offline models, providing a reusable benchmark and clear directions for future research.

2 Related Work

2.1 Video LLMs

Offline and online VideoLLMs have emerged as two major paradigms in contemporary large-scale video model research (Huang et al., 2025; Chen et al., 2024). Offline approaches typically assume access to the entire video clip at inference time, and obtain video representations via strategies such as sparse sampling (Feichtenhofer et al., 2018; Bertasius et al., 2021), global aggregation (Arnab et al., 2021), or long-video compression (Ryoo et al., 2021), which are then aligned

with a language model for reasoning. (Maaz et al., 2024; Zhang et al., 2023; Lin et al., 2024a; Jin et al., 2024; Song et al., 2024; Weng et al., 2024) Although these methods achieve strong performance on standard benchmarks, their underlying assumption that the full video is available before inference often deviates from real streaming applications. In contrast, online models are explicitly designed for streaming inputs and interactive use: they must incrementally encode newly arriving frames, continuously update memory/state, and respond immediately when a user query is issued. As deployment demands grow, research attention has increasingly shifted from optimizing offline accuracy alone to improving online usability and interactive experience (Chen et al., 2024; Huang et al., 2025; Ning et al., 2025; Wang et al., 2025c).

2.2 Streaming Benchmark

Previous evaluations of streaming video understanding have largely emphasized answer accuracy. Early work such as VStream-QA (Zhang et al., 2024) simulates realistic streaming queries by timestamping each question and restricting it to depend only on video content observed up to that moment, and it mainly measures correctness via accuracy and automatic scoring. StreamingBench (Lin et al., 2024b) further broadens task coverage and modality dimensions, using overall performance as the primary metric to assess streaming understanding. OVOBench (Li et al., 2025) constructs multiple online VQA subtasks under three temporal contexts (past, present, and future). Recently, StreamEQA (Wang et al., 2025b) establishes a streaming benchmark specifically tailored for embodied scenarios. However, in real-world online applications, latency is equally critical: it not only shapes user experience but is also tightly coupled with throughput, compute budgets, and serving costs. Reporting accuracy alone can therefore obscure a model’s practical usability under strict online constraints and the inherent accuracy-latency trade-off. To bridge this gap, our evaluation framework incorporates latency metrics alongside accuracy under a rigorous online setting, enabling reproducible and fair comparative analyses across streaming models.

3 Streaming Evaluation Protocol

In this subsection, we present StreamingEval in detail. We first provide a rigorous definition of

streaming online inference, and then describe the concrete implementation pipeline of the evaluation framework. Next, we explain how to construct fair and reproducible comparison settings for both offline and online models. Finally, we introduce a set of evaluation metrics that are critical for online inference in streaming scenarios.

3.1 Online Task Definition

We define *streaming online question answering* as a real-time, interactive multimodal dialogue setting in which the model continuously receives video frames along the temporal axis. At any time t , the input may consist of three components: (1) the interaction history between the user and the model up to time t , denoted as C_t ; (2) the video frame acquired at time t , denoted as V_t ; and (3) the user query issued at time t , denoted as Q_t . When the encoded query arrives at the model at time t_1 , the backbone language model f autoregressively generates a response R_{t_1} conditioned on $(C_{t_1}, V_{\text{enc}[0, t_1]}, Q_{t_1})$, i.e., by maximizing the conditional probability $p(R_{t_1} | C_{t_1}, V_{[0, t_1]}, Q_{t_1})$. Here, $V_{\text{enc}[0, t_1]}$ denotes all video frames received up to time t_1 that have already been encoded by the model. This setting emphasizes the model’s ability to perform inference and respond at arbitrary time points.

3.2 StreamingEval Framework

To enable streaming online inference, we implement the execution backend of StreamingEval as an asynchronous, time-causal pipeline composed of three decoupled processes that run in parallel, the framework is illustrated in Figure 2. Specifically, it consists of a **Frame Player**, an **Encoder-and-Memory Updater**, and a **Responder**. These processes communicate via inter-process queues and/or shared buffers, emulating the behavior of an online system in which video frames arrive continuously, the model updates continually, and user queries may occur at any time, without introducing additional synchronization-induced blocking (Ma et al., 2021).

Input Stream and Frame Player. For an arbitrary video stream, we represent it as $\{(v_i, \tau_i)\}_{i=1}^{\infty}$, where v_i denotes the i -th frame and τ_i denotes its arrival time. The frame player samples frames at a fixed interval ρ and streams them to downstream processes.

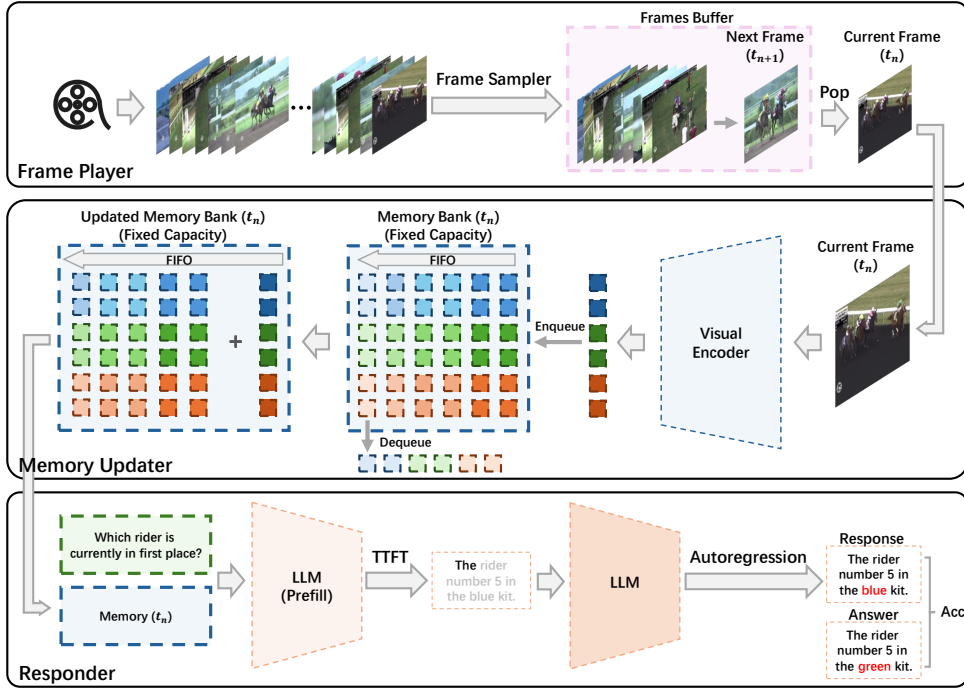


Figure 2: Overview of the StreamingEval framework. The framework standardizes streaming video understanding by modeling continuous input ingestion, incremental visual memory updates, and query-driven inference within a unified protocol.

Encoder & Memory Updater. The encoder first maps each incoming frame to a visual representation:

$$z_i = g_\theta(v_i), \quad (1)$$

where g_θ denotes the visual encoding backbone. The memory updater then maintains an online memory state M . For online models, the memory is updated according to the model-specific update rule \mathcal{U} by writing the new representation into memory:

$$M_{\tau_i^+} = \mathcal{U}(M_{\tau_i^-}, z_i; B, \pi), \quad (2)$$

where τ_i^- and τ_i^+ denote the memory states immediately before and after the arrival of the i -th frame at time τ_i , respectively. B specifies the memory budget (capacity constraint), and π denotes the corresponding write/eviction policy. For offline models, we apply a projection layer to map visual features into the embedding space aligned with the language model, and then store them in a fixed-length visual context window maintained with a first-in-first-out (FIFO) policy.

Responder. A user may launch a query q_{t_0} at any time t_0 . Once the query is triggered, the responder first encodes it, and denotes the time when encoding finishes as t_1 . The responder then reads the memory snapshot available at time t_1 , denoted

by M_{t_1} , and conditions on the dialogue context C_{t_1} together with the query q_{t_0} to autoregressively generate an answer:

$$R_{t_1} \sim p_\phi(\cdot | q_{t_0}, C_{t_1}, M_{t_1}), \quad (3)$$

where C_{t_1} represents the interaction history up to time t_1 , M_{t_1} denotes the information updated up to time t_1 , and p_ϕ is parameterized by the backbone language model.

3.3 Comparable Online/Offline Setups

In streaming settings, a fair comparison between native online models and general multimodal models is not straightforward: multimodal models typically assume access to the full video, whereas online models must operate under strict causal constraints and rely only on historical information that has already arrived and been processed. Moreover, visual token embeddings differ in dimensionality across models; therefore, constraining the history length solely by the *number* of tokens can yield inconsistent actual memory footprints. To address this, StreamingEval adopts a “two settings, one unified budget” strategy: we preserve the native mechanisms of online models as much as possible, while introducing a resource-constrained offline adapter for multimodal models that constructs con-

text under the same resource budget, enabling fair comparisons.

Native Online-model Setting. For native online models, we follow the original online mechanisms and configurations in their papers as closely as possible, including incremental encoding, memory/state updates, retrieval policies, and default input resolution and preprocessing. Models run within our multi-process emulator: frames arrive continuously at a fixed frame rate, the model updates its memory on the fly, and when a query is issued at time t , the responder can only access the memory snapshot that is *available up to t* (i.e., processed before t). This ensures strict causality while staying faithful to each model’s intended design.

Multimodal-model Adapter Setting. For multimodal models, we introduce a unified bounded-memory adapter. As video frames arrive, each model produces visual representations in its native manner; after a projection layer aligns them to the language model embedding space, the resulting representations are written into a fixed-capacity memory bank. When the memory bank exceeds the budget, we evict the oldest content using a deterministic FIFO policy. When a query is issued, we concatenate the current memory bank with the dialogue context as the model input, thereby simulating a deployable version of multimodal models under strict online constraints.

For offline models, we adopt a fixed-capacity memory bank with a FIFO eviction policy to enable as neutral and reproducible an evaluation as possible under realistic streaming constraints, while avoiding the additional algorithmic gains introduced by compression or summarization modules that could compromise fairness. At the same time, StreamingEval is open to different memory-management strategies: methods such as clustering-based compression, learned summarization, and KV compression can all be readily plugged in, and their effects will be naturally reflected in accuracy, latency, memory usage, and the final StreamingScore.

Unified Resource Budget. To avoid unfair comparisons where different models incur different GPU memory footprints under the same number of visual tokens due to embedding-dimensionality mismatch, we enforce a byte-level resource budget and cap the storage of historical context at

M . We account only for the two components that scale with context length: the projected visual-token representations cached in the memory bank, and the language Transformer KV cache associated with these visual tokens for incremental inference. (Kwon et al., 2023; Dao et al., 2022) Implementation details of the computations are provided in the appendix.

3.4 Evaluation Metrics

Online streaming multimodal question answering is simultaneously constrained by *throughput, interactive latency, the effective historical-context budget, and task correctness*; no single metric suffices to characterize real-world deployability. Accordingly, we introduce the StreamingEval suite: we define four core metrics from four complementary system perspectives encoding, decoding, memory, and task and additionally report an overall score for *summary reporting*, to more clearly reflect the trade-offs among accuracy, latency, throughput, and resource usage across methods.

Visual Encoding: A streaming system must keep up with the incoming input rate; otherwise frame backlog accumulates and latencies quickly become unstable. Therefore, we use MaxFPS to measure the maximum input frame rate the model can *sustain* without persistent backlog, capturing the real-time throughput ceiling of the visual encoding and memory-update pipeline.

Text Decoding: User experience is largely determined by the latency of the response; therefore, we report TTFT (time-to-first-token), defined as the elapsed time from query arrival to the generation of the first token. (MLCommons, 2025; MLCommons Inference Working Group)

Memory Overhead: A key constraint in streaming settings is how much historical context can be retained under limited resources, which directly impacts deployability and coverage; therefore, we use Memory_bank to denote the budget of online-available visual history cache, aligning the effective amount of historical context that different models can leverage in streaming scenarios.

Task Performance: The most straightforward measure of model performance is answer accuracy; therefore, we use question-answering accuracy to evaluate a model’s capability on streaming online QA.

Comprehensive Metric: To more accurately quantify a models overall capability and its performance under streaming constraints, we integrate

the metrics from the major contributing factors in streaming scenarios and define a composite metric, *StreamingScore*:

$$\text{StreamingScore}(\mathbf{w}) \triangleq \frac{\text{MaxFPS}^{w_f} \cdot \text{Acc}^{w_a}}{\text{TTFT}^{w_t} \cdot M^{w_r}}, \quad (4)$$

$$M \triangleq \text{Mem} \cdot \ln(\text{Params}), \quad (5)$$

$$w_f, w_a, w_t, w_r \geq 0, \quad w_f + w_a + w_t + w_r = 1. \quad (6)$$

This metric is designed to capture the trade-offs among accuracy, latency, throughput, and resource usage in streaming settings. A higher *StreamingScore* indicates a model that maintains higher throughput while achieving better accuracy with lower time-to-first-token latency and lower resource consumption.

4 Experiments

We evaluate 12 mainstream multimodal and online models under the StreamingEval framework. We first describe the experimental setup and datasets, then report overall results, followed by analyses under different memory_bank budgets and input image resolution ranges, and finally summarize the key findings.

4.1 Experiment Settings

We evaluate representative multimodal models and online video models within the StreamingEval framework. All experiments are conducted on a single RTX 4090 (48GB) GPU with BF16 inference, using a unified prompt and decoding configuration. TTFT is measured in wall-clock time. Streaming inputs are fed at 1fps. Due to space limitations, more experimental setups are shown in the Appendix A.2.

4.2 Datasets

We evaluate StreamingEval on two widely used benchmarks for online/streaming video understanding: OVO-Bench (Huang et al., 2025) and StreamingBench (Lin et al., 2024b). Both benchmarks cover diverse video understanding and question answering formats, designed to test temporal reasoning, event recognition, and multimodal inference.

OVO-Bench. OVO-Bench is designed to evaluate online video understanding capabilities and assess models in three typical online scenarios: Backward Tracing, Real-Time Understanding, and

Forward Active Responding. It comprises 12 tasks and approximately 2,800 fine-grained annotations with precise timestamps, and supports systematic query-based evaluation along the video timeline.

StreamingBench. StreamingBench organizes its task taxonomy around three dimensions: real-time visual understanding, multi-source (audiovisual) understanding, and contextual understanding. The benchmark comprises 18 tasks, 900 video clips, and 4,500 human-annotated QA pairs, and issues multiple queries at different timestamps within the same video to simulate continuous streaming inputs.

4.3 Main Results

Table 1 summarizes the results of 12 mainstream multimodal and online models under StreamingEval. We report task-level performance, system-level deployability metrics, and provide a normalized composite score, *StreamingScore*, for convenient comparison.

Encoding Efficiency. Most models achieve a peak throughput (MaxFPS) that satisfies the minimum requirement for streaming online inference (1 FPS), indicating basic real-time feasibility on the input side. VideoChatOnline is a notable exception: its MaxFPS is substantially below 1 FPS, revealing a severe bottleneck in the video encoding/state-update pipeline. We attribute this to its relatively complex memory-update and cross-frame maintenance mechanism, which significantly reduces the number of frames that can be processed per unit time. At such throughput, the system cannot keep pace with the incoming video stream. In practice, this would induce persistent backlog and accumulating latency, rendering the model difficult to deploy in real-world streaming settings.

Time-to-First-Token. For most models, the TTFT is primarily governed by the size of the visual-memory context window and the overhead of cross-modal context construction, yet it remains below 1.5 seconds in nearly all cases. Consequently, assuming the encoding throughput is sufficient, this level of decoding startup latency is unlikely to constitute a primary obstacle to practical deployment.

Overall Capability Comparison: Offline vs. Online Models As shown in Table 1, under our strict streaming protocol, offline VideoLLMs and

Model	MaxFPS	Mem(GB)	TTFT	OVO-Bench					StreamingBench	
				Real-Time	Backward	Forward	Overall	StreamingScore	Real-time	StreamingScore
Open-source Offline VideoLLMs										
Qwen3-VL-8B (Bai et al., 2025)	8	0.5	0.20	78.73	51.82	43.46	58.00	2.21	77.31	2.37
InternVL3.5-8B (Wang et al., 2025a)	7	0.5	0.21	74.31	40.89	46.14	53.78	2.04	77.96	2.24
Llava-OV1.5-8B (An et al., 2025)	7	0.5	0.21	75.51	45.80	45.44	55.58	2.05	76.19	2.22
MiniCPM-V4.5-8B (Yu et al., 2025)	6	0.5	0.18	74.55	54.52	42.05	57.04	2.07	76.55	2.23
VideoLLaMA3-7B (Zhang et al., 2025a)	7	0.5	0.56	55.73	45.32	45.88	48.44	1.58	68.90	1.72
VideoChat-7B (Li et al., 2023)	9	0.5	0.52	69.77	40.25	43.27	49.95	1.73	72.22	1.89
Open-source Online Video-LLMs										
Flash-VStream-7B (Zhang et al., 2024)	8	0.35	0.12	29.86	25.35	44.23	33.15	2.34	23.23	2.18
Flash-VStream-7B* (Zhang et al., 2025b)	1	0.66	1.31	59.88	46.43	47.41	50.31	0.74	74.48	0.81
ReKV-7B (Di et al., 2025)	5	0.50	1.29	57.34	44.69	44.35	48.00	1.18	64.53	1.27
StreamForest-7B (Zeng et al., 2025)	4	0.46	0.98	61.20	52.02	53.49	55.57	1.26	77.26	1.37
TimeChat-Online-7B (Yao et al., 2025)	7	0.34	0.62	58.60	42.00	36.40	45.67	1.67	73.64	1.88
VideoChatOnline-4B (Huang et al., 2025)	0.14	0.32	0.18	43.97	37.40	39.69	40.40	0.92	58.81	1.19

Table 1: Overall results on OVO-Bench and StreamingBench. For OVO-Bench, we report the average scores of the Real-Time, Backward, and Forward task clusters; for StreamingBench, we report the overall Real-time average score, where the streaming score is computed following the definition described in the main text. Detailed results are provided in Table 3.

native online models exhibit a relatively consistent gap in question-answering accuracy. Overall, offline models tend to achieve higher accuracy on the aggregate scores of OVO-Bench and StreamingBench, reflecting their stronger representation and reasoning capacity, as well as their advantage in more fully leveraging visual evidence during answer generation. In contrast, to satisfy causal constraints and limited state updates, native online models typically rely on incremental memory mechanisms that compress, truncate, or selectively retrieve historical information, which is more prone to information loss on tasks requiring long-range temporal dependencies or fine-grained cues, thereby resulting in lower accuracy. Moreover, from the perspective of streaming interactivity, although offline models are often superior in accuracy, their streaming scores on OVO/SB do not necessarily lead accordingly; instead, some native online models attain higher streaming scores by offering faster responses and a more stable online pace under causal constraints, further revealing a systematic trade-off between strict streaming deployability and task accuracy. Overall, the results indicate a practical tension between strict online deployability and task accuracy: retaining critical information more effectively under limited resources remains a key direction for improving accuracy in streaming scenarios.

Performance Differences Across Task Categories. The per-category results in Table 1 reveal systematic performance divergences between

offline and online VideoLLMs across task types.

For Backward Tracing, which relies on long-horizon temporal integration and reasoning over extended visual history, offline models are typically stronger, suggesting that full-context modeling better supports retrospective inference and evidence aggregation.

By contrast, for Forward Active Responding—which emphasizes rapid decision making and incremental interaction under strict streaming constraints—online models are largely on par with offline counterparts, and can even outperform them in some cases, highlighting the structural benefits of native online mechanisms for continuous updates and timely responses.

For Real-Time Visual Perception, offline models remain generally superior overall. This indicates that even when the task is more “current-segment” oriented, online approaches may still suffer from information compression and noise accumulation induced by limited buffering and incremental updates, which can undermine the stability of fine-grained visual representations compared to offline models with unified context encoding.

4.4 Further Discussion

Sensitivity to the Memory_bank Budget. To evaluate model performance under different historical-cache budgets, we keep all other protocol settings fixed and vary the Memory_bank budget over {1.5G, 1.0G, 0.5G, 0.3G, 0.1G}, reporting the resulting changes in Overall Accuracy (Fig-

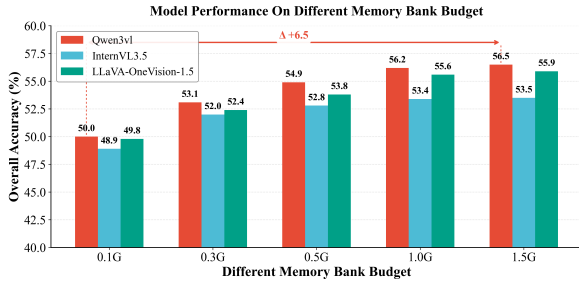


Figure 3: Overall accuracy versus memory_bank budget for three representative models. Detailed results are provided in Table 6.

ure 3). As shown, accuracy nearly saturates in the relatively ample regime of 1.0G–1.5G: increasing the budget from 1.0G to 1.5G yields only marginal gains, indicating diminishing returns from further cache expansion. In contrast, as the budget is tightened from 1.0G to 0.5G, 0.3G, and finally 0.1G, all three models exhibit a clear accuracy drop, suggesting that performance is largely constrained by a historical-information bottleneck. While the overall ranking remains consistent, the scores converge substantially under the most restrictive 0.1G setting, implying that different methods are more strongly capped by the shared cache constraint. Meanwhile, Qwen3-VL still retains relatively higher accuracy at this extreme budget, reflecting more efficient extraction and utilization of limited historical information and better suitability for resource-constrained deployment.

Different Input Resolution. Different input resolution determines how much visual detail is available per frame. With Memory_bank fixed (set to the default in Table 2) and all other protocol settings unchanged, we evaluate three common input sizes (448×448 , 336×336 , and 224×224), using a unified video-frame resizing pipeline to ensure consistent processing (Table 2). As shown in Table 2, all three representative models achieve their best accuracy at the highest resolution of 448×448 ; as the resolution decreases, performance drops to varying degrees, indicating a stable degradation when downscaling inputs. The models also differ in robustness to resolution reduction: Qwen3-VL degrades more mildly, whereas InternVL 3.5 shows a more pronounced drop, suggesting stronger reliance on the additional fine-grained details provided by higher resolution; Flash-VStream also deteriorates as resolution decreases and remains substantially behind

Input Res.	Qwen3-VL	InternVL 3.5	FlashVstream
224×224	51.99	48.93	33.15
336×336	52.95	51.70	34.67
448×448	53.34	52.75	35.73

Table 2: Overall results under different input resolutions on OVO-Bench. We report the overall score for each model under 224×224 , 336×336 , and 448×448 inputs. Detailed results are provided in Table 7.

the other two in absolute accuracy. From the perspective of “downscaling loss,” the drop from 448×448 to 336×336 is relatively small, while the drop from 336×336 to 224×224 is more evident, implying that further downscaling more quickly hits a detail bottleneck. Overall, higher resolution improves accuracy but comes with higher computation and latency costs; under deployment budget constraints, 336×336 is often a more robust trade-off between effectiveness and cost, whereas 224×224 is cheaper but incurs a notably larger performance decrease.

4.5 Outlook and Future Directions

Based on a systematic experimental analysis using StreamingEval, we observe that in online video understanding settings, a model’s offline accuracy does not necessarily translate into a usable online experience. Our evaluation indicates that different models exhibit significant and stable trade-offs along four dimensions—accuracy, throughput, latency, and limited history. In practice, the key bottlenecks that shape user experience often stem from response waiting time, performance degradation under throughput constraints, and information loss and error accumulation induced by restricted historical context.

Accordingly, we argue that future research on online video understanding models should more explicitly incorporate system constraints into model design objectives. First, at the architectural level, priority should be given to representation and memory mechanisms with streaming, incremental update capabilities, enabling models to reliably maintain critical states under limited history. Second, at the inference level, it is important to explore latency-accuracy controllable computation allocation strategies, so as to provide effective responses for real-time interactive use.

5 Conclusions

We propose StreamingEval, a unified evaluation protocol for video question answering models under strictly streaming, online constraints. Beyond accuracy, we emphasize deployability-oriented metrics, including time to first token (TTFT), MaxFPS, and a constrained memory bank, and introduce a resource-budget adapter to enable fair comparisons between online models and offline VideoLLMs. Using this protocol, we evaluate 12 representative models and show that being online does not necessarily translate into practical deployability; moreover, changes in the memory bank and input resolution lead to corresponding shifts in accuracy. We hope StreamingEval will facilitate reproducible, system-level evaluation and inspire future work to optimize for the end-to-end online user experience.

6 Limitations

Due to the lack of publicly available code or reproducible interfaces for some closed-source/proprietary models, we cannot reliably standardize their inference configurations or faithfully measure their true performance under a unified evaluation protocol and runtime environment. Consequently, our evaluation does not provide a comprehensive comparison across all mainstream closed-source models. Limited computational resources and our focus on mobile/edge deployment constrain our experiments mainly to 7B/8B scale models. As a result, we are unable to draw firm conclusions about the performance/efficiency trade-offs of larger models under the same streaming setting. In future work, we will seek additional GPU resources and expand our reproducible coverage to evaluate larger models. We will also continuously update the project website with new results and analyses.

Potential Risks This work proposes an evaluation protocol only and does not include or release any dataset; therefore, the privacy and licensing risks are limited. The main risks are that the protocol may be over-generalized as a universal standard, or may incentivize over-optimization for the reported metrics that does not reflect real-world performance. We recommend validating the protocol across diverse domains/datasets/scenarios with robustness and subgroup analyses, and caution against using a single evaluation result to justify

deployment or decision-making in high-stakes settings.

Acknowledgements

Project supported by Shanghai Municipal Science and Technology Major Project 2025SHZDZX025G16.

References

- Xiang An, Yin Xie, Kaicheng Yang, Wenkang Zhang, Xiuwei Zhao, Zheng Cheng, Yirui Wang, Songcen Xu, Changrui Chen, Chunsheng Wu, Huajie Tan, Chunyuan Li, Jing Yang, Jie Yu, Xiyao Wang, Bin Qin, Yumeng Wang, Zizhen Yan, Ziyong Feng, and 3 others. 2025. [Llava-onevision-1.5: Fully open framework for democratized multimodal training](#). *Preprint*, arXiv:2509.23661.
- Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lucic, and Cordelia Schmid. 2021. [Vivit: A video vision transformer](#). *arXiv preprint arXiv:2103.15691*.
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhi-fang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Ming-sheng Li, and 45 others. 2025. [Qwen3-vl technical report](#). *Preprint*, arXiv:2511.21631.
- Gedas Bertasius, Heng Wang, and Lorenzo Torresani. 2021. [Is space-time attention all you need for video understanding?](#) *arXiv preprint arXiv:2102.05095*.
- Tim Brödermann, Christos Sakaridis, Yuqian Fu, and Luc Van Gool. 2025. [Cafuser: Condition-aware multimodal fusion for robust semantic perception of driving scenes](#). *IEEE Robotics and Automation Letters*.
- Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, Pete Florence, Chuyuan Fu, Montse Gonzalez Arenas, Keerthana Gopalakrishnan, Kehang Han, Karol Hausman, Alexander Herzog, Jasmine Hsu, Brian Ichter, and 35 others. 2023. [Rt-2: Vision-language-action models transfer web knowledge to robotic control](#). In *Proceedings of The 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, pages 2165–2183. PMLR. ArXiv:2307.15818.
- Joya Chen, Zhaoyang Lv, Shiwei Wu, Kevin Qinghong Lin, Chenan Song, Difei Gao, Jia-Wei Liu, Ziteng Gao, Dongxing Mao, and Mike Zheng Shou. 2024. [Videollm-online: Online video large language model for streaming video](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18407–18418.

- Yilong Chen, Xiang Bai, Zhibin Wang, Chengyu Bai, Yuhan Dai, Ming Lu, and Shanghang Zhang. 2025. [Streamkv: Streaming video question-answering with segment-based kv cache retrieval and compression](#). *Preprint*, arXiv:2511.07278.
- Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. [Flashattention: Fast and memory-efficient exact attention with IO-awareness](#). *CoRR*, abs/2205.14135.
- Shangzhe Di, Zhelun Yu, Guanghao Zhang, Haoyuan Li, Tao Zhong, Hao Cheng, Bolin Li, Wanggui He, Fangxun Shu, and Hao Jiang. 2025. [Streaming video question-answering with in-context video kv-cache retrieval](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, and 3 others. 2023. [Palm-e: An embodied multimodal language model](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 8469–8488. PMLR.
- Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. 2018. Slowfast networks for video recognition. *arXiv preprint arXiv:1812.03982*.
- Zhenpeng Huang, Xinhao Li, Jiaqi Li, Jing Wang, Xiangyu Zeng, Cheng Liang, Tao Wu, Xi Chen, Liang Li, and Limin Wang. 2025. [Online video understanding: Ovbench and videochat-online](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3328–3338.
- Peng Jin, Ryuichi Takanobu, Wancai Zhang, Xiaochun Cao, and Li Yuan. 2024. [Chat-univi: Unified visual representation empowers large language models with image and video understanding](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13700–13710.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with pagedattention](#). *CoRR*, abs/2309.06180.
- Jesse Levinson, Jake Askeland, Jan Becker, Jennifer Dolson, David Held, Soeren Kammel, J Zico Kolter, Dirk Langer, Oliver Pink, Vaughan Pratt, and 1 others. 2011. [Towards fully autonomous driving: Systems and algorithms](#). In *2011 IEEE intelligent vehicles symposium (IV)*, pages 163–168. IEEE.
- KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. 2023. [Videochat: Chat-centric video understanding](#). *Preprint*, arXiv:2305.06355.
- Yifei Li, Junbo Niu, Ziyang Miao, Chunjiang Ge, Yuanhang Zhou, Qihao He, Xiaoyi Dong, Haodong Duan, Shuangrui Ding, Rui Qian, Pan Zhang, Yuhang Zang, Yuhang Cao, Conghui He, and Jiaqi Wang. 2025. [Ovo-bench: How far is your video-llms from real-world online video understanding?](#) *Preprint*, arXiv:2501.05510.
- Bin Lin, Yang Ye, Bin Zhu, Jiayi Cui, Munan Ning, Peng Jin, and Li Yuan. 2024a. [Video-LLaVA: Learning united visual representation by alignment before projection](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5971–5984, Miami, Florida, USA. Association for Computational Linguistics.
- Junming Lin, Zheng Fang, Chi Chen, Zihao Wan, Fuwen Luo, Peng Li, Yang Liu, and Maosong Sun. 2024b. [Streamingbench: Assessing the gap for mllms to achieve streaming video understanding](#). *Preprint*, arXiv:2411.03628.
- Xutai Ma, Jiateng Xie, Juan Miguel Pino, Philipp Koehn, and Graham Neubig. 2021. [Simuleval: An evaluation toolkit for simultaneous translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 144–150. Association for Computational Linguistics.
- 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, Bangkok, Thailand. Association for Computational Linguistics.
- MLCommons. 2025. [MLPerf Client Benchmark](#). Software and documentation. Accessed 2026-01-05. Used for TTFT (time-to-first-token) measurement/benchmarking.
- MLCommons Inference Working Group. Definition of ttft and tpot latency (github issue #1593). <https://github.com/mlcommons/inference/issues/1593>. Accessed: 2026-01-04.
- Zhenyu Ning, Guangda Liu, Qihao Jin, Wenchao Ding, Minyi Guo, and Jieru Zhao. 2025. [Livevlm: Efficient online video understanding via streaming-oriented KV cache and retrieval](#). *CoRR*, abs/2505.15269.
- Tianwen Qian, Jingjing Chen, Linhai Zhuo, Yang Jiao, and Yu-Gang Jiang. 2024. [Nuscenes-qa: A multi-modal visual question answering benchmark for autonomous driving scenario](#). In *Proceedings of*

- the AAAI Conference on Artificial Intelligence*, volume 38, pages 4542–4550.
- Vijay Janapa Reddi, Christine Cheng, David Kanter, Peter Mattson, Guenther Schmuelling, Carole-Jean Wu, Brian Anderson, Maximilien Breughe, Mark Charlebois, William Chou, Ramesh Chukka, Cody Coleman, Sam Davis, Pan Deng, Greg Diamos, Jared Duke, Dave Fick, J. Scott Gardner, Itay Hubara, and 28 others. 2019. [MLPerf Inference Benchmark](#). *CoRR*, abs/1911.02549.
- Michael S. Ryoo, A. J. Piergiovanni, Anurag Arnab, Mostafa Dehghani, and Anelia Angelova. 2021. Tokenlearner: What can 8 learned tokens do for images and videos? *arXiv preprint arXiv:2106.11297*.
- Enxin Song, Wenhao Chai, Guan hong Wang, Yucheng Zhang, Haoyang Zhou, Feiyang Wu, Haozhe Chi, Xun Guo, Tian Ye, Yanting Zhang, Yan Lu, Jenq-Neng Hwang, and Gaoang Wang. 2024. Moviechat: From dense token to sparse memory for long video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18221–18232.
- Kuanning Wang, Ke Fan, Yuqian Fu, Siyu Lin, Hu Luo, Daniel Seita, Yanwei Fu, Yu-Gang Jiang, and Xiangyang Xue. 2026. Ocr: Object-centric learning with 3d and tactile priors for human-to-robot action transfer. *arXiv preprint arXiv:2603.14401*.
- Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, Zhaokai Wang, Zhe Chen, Hongjie Zhang, Ganlin Yang, Haomin Wang, Qi Wei, Jinhui Yin, Wenhao Li, Erfei Cui, and 56 others. 2025a. [InternV3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency](#). *Preprint*, arXiv:2508.18265.
- Yifei Wang, Zhenkai Li, Tianwen Qian, Huanran Zheng, Zheng Wang, Yuqian Fu, and Xiaoling Wang. 2025b. [Streameqa: Towards streaming video understanding for embodied scenarios](#). *arXiv preprint arXiv:2512.04451*.
- Yiyu Wang, Xuyang Liu, Xiyan Gui, Xinying Lin, Boxue Yang, Chenfei Liao, Tailai Chen, and Linfeng Zhang. 2025c. [Accelerating streaming video large language models via hierarchical token compression](#). *Preprint*, arXiv:2512.00891.
- Yuetian Weng, Mingfei Han, Haoyu He, Xiaojun Chang, and Bohan Zhuang. 2024. [Longvlm: Efficient long video understanding via large language models](#). *CoRR*, abs/2404.03384.
- Ruyi Xu, Guangxuan Xiao, Yukang Chen, Liuning He, Kelly Peng, Yao Lu, and Song Han. 2025. [Streamingvlm: Real-time understanding for infinite video streams](#). *CoRR*, abs/2510.09608.
- Linli Yao, Yicheng Li, Yuancheng Wei, Lei Li, Shuhuai Ren, Yuanxin Liu, Kun Ouyang, Lean Wang, Shicheng Li, Sida Li, Lingpeng Kong, Qi Liu, Yuanxing Zhang, and Xu Sun. 2025. [Timechat-online: 80% visual tokens are naturally redundant in streaming videos](#). *Preprint*, arXiv:2504.17343.
- Tianyu Yu, Zefan Wang, Chongyi Wang, Fuwei Huang, Wenshuo Ma, Zhihui He, Tianchi Cai, Weize Chen, Yuxiang Huang, Yuanqian Zhao, Bokai Xu, Junbo Cui, Yingjing Xu, Liqing Ruan, Luoyuan Zhang, Hanyu Liu, Jingkun Tang, Hongyuan Liu, Qining Guo, and 15 others. 2025. [Minicpm-v 4.5: Cooking efficient mllms via architecture, data, and training recipe](#). *Preprint*, arXiv:2509.18154.
- Xiangyu Zeng, Kefan Qiu, Qingyu Zhang, Xinhao Li, Jing Wang, Jiabin Li, Ziang Yan, Kun Tian, Meng Tian, Xinhai Zhao, Yi Wang, and Limin Wang. 2025. [Streamforest: Efficient online video understanding with persistent event memory](#). *Preprint*, arXiv:2509.24871.
- Boqiang Zhang, Kehan Li, Zesen Cheng, Zhiqiang Hu, Yuqian Yuan, Guanzheng Chen, Sicong Leng, Yuming Jiang, Hang Zhang, Xin Li, Peng Jin, Wenqi Zhang, Fan Wang, Lidong Bing, and Deli Zhao. 2025a. [Videollama 3: Frontier multimodal foundation models for image and video understanding](#). *Preprint*, arXiv:2501.13106.
- Hang Zhang, Xin Li, and Lidong Bing. 2023. [VideoLLaMA: An instruction-tuned audio-visual language model for video understanding](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 543–553, Singapore. Association for Computational Linguistics.
- Haoji Zhang, Yiqin Wang, Yansong Tang, Yong Liu, Jiashi Feng, Jifeng Dai, and Xiaojie Jin. 2024. [Flash-vstream: Memory-based real-time understanding for long video streams](#). *arXiv preprint arXiv:2406.08085*.
- Haoji Zhang, Yiqin Wang, Yansong Tang, Yong Liu, Jiashi Feng, and Xiaojie Jin. 2025b. [Flash-vstream-qwen \(official implementation\)](#). GitHub repository. Flash-VStream variant built on Qwen2-VL. Accessed: 2026-01-06.

A Appendix

We include supplementary material in the appendix to facilitate a deeper understanding of our experimental setup and results. Due to space constraints, we focus on the most essential details in each section. The appendix follows the structure of the main paper: Appendix I elaborates on the evaluation protocol, Appendix II provides additional experimental details.

A.1 Detailed Evaluation Protocol

For the i -th model, let B denote the number of visual tokens retained. We define the linearly growing storage cost as

$$\text{Mem}_i(B) = B \cdot d_i \cdot s_{\text{emb}} + B \cdot 2L_i \cdot h_i^{\text{kv}} \cdot s_{\text{kv}},$$

where d_i is the embedding dimension of the projected visual tokens and s_{emb} is the bytes per element for these embeddings; L_i is the number of layers in the language backbone; h_i^{kv} is the per-layer KV channel width (e.g., $h_i^{\text{kv}} = n_i^{\text{kv}} \cdot d_i^{\text{head}}$, with n_i^{kv} the number of KV heads and d_i^{head} the per-head dimension). The factor 2 accounts for both Key and Value tensors, and s_{kv} is the bytes per element for KV cache (e.g., $s_{\text{kv}}=2$ for BF16). Accordingly, we convert the byte budget M (denoting $M_{\text{bytes}} = M \cdot 10^9$) into the model-specific visual-token cap:

$$B_i = \left\lfloor \frac{M_{\text{bytes}}}{d_i s_{\text{emb}} + 2L_i h_i^{\text{kv}} s_{\text{kv}}} \right\rfloor.$$

A.2 Experimental Details

In this section, we first describe the experimental setup and environment. We then present the measurement results for each subtask in Section 4.3, followed by the experimental results discussed in Section 4.4.

Experiment Details We conducted all experiments on GPUs provided by a cloud server. Notably, the GPU model and its performance characteristics do not correspond to mainstream, widely used GPU models. Therefore, to ensure transparency and reproducibility, we report in 5 the performance specifications provided by the cloud service vendor. We use compatible versions of Transformers and PyTorch, and enable FlashAttention-2 and Accelerate for runtime acceleration. Empirically, the overhead introduced by inter-process communication is negligible, and we observe no statistically significant degradation in end-to-end inference latency.

Experiment Setups. To ensure resource comparability across multimodal models in an online inference setting, we adopt a unified, memory-footprint-based criterion derived from the online model to configure each model’s context window. Specifically, we abstract the storage terms that scale approximately linearly with context length into two components: (i) multimodal representations written to the memory bank (e.g., embeddings of visual/audio tokens), and (ii) the Transformer attention cache (KV cache) associated with these multimodal tokens. Given a shared memory budget M (in bytes), we use the online memory estimation function to compute each model’s incremental storage cost per unit of context and then back-solve for the maximum admissible context length B_i under budget M . Accordingly, we assign each multimodal model a context window that more closely matches its online memory consumption, enabling fairer and more reproducible cross-model comparisons without relying on raw token counts.

Experiment Results. To ensure consistency and comparability in the main results, we adopt a uniform weighting scheme in the main table, i.e., setting all weight coefficients equally as $w_f = w_a = w_t = w_r = 0.25$, and report the corresponding StreamingScore under this setting. We note that this configuration is intended as a default comprehensive evaluation protocol, and does not imply that all real-world deployment scenarios place identical emphasis on throughput, accuracy, time-to-first-token, and resource consumption. In practical applications, different users or tasks often prioritize different system capabilities, such as answer quality, interaction responsiveness, resource efficiency, or throughput. Therefore, assigning larger weights to the metrics that better align with specific deployment preferences can more accurately reflect the practical utility of a model in a given scenario. Based on this consideration, we further report results under different weight configurations in the appendix to demonstrate the applicability and robustness of StreamingScore across diverse deployment preferences. To validate practical utility, we evaluate four representative deployment scenarios by setting the target terms weight to 0.4 and the remaining three to 0.2. The table 4 below shows how representative models rankings change across scenarios. This scenario-aware met-

Model	OVO-Bench													StreamingBench													
	Real-Time Visual Perception						Backward Tracing				Forward Active Responding			Real-time Visual Understanding													
	OCR	ACR	ATR	STU	FPD	OJR	Avg.	EPM	ASI	HLD	Avg.	REC	SSR	CRR	Avg.	CS	OP	ATP	PR	ACP	SU	EU	CT	TR	CR	Avg.	
Open-source Offline VideoLLMs																											
Qwen3-VL-7B	91.95	78.90	79.31	67.98	73.27	80.98	78.73	48.82	61.49	48.92	51.82	21.78	64.86	50.42	43.46	89.27	81.74	74.79	75.00	88.78	72.36	68.55	40.41	85.67	71.88	77.31	
InternVL3.5-8B	88.59	75.23	77.59	61.80	71.29	73.91	74.31	51.18	56.75	11.83	40.89	25.36	67.73	50.00	46.14	89.59	81.74	76.49	78.70	88.78	69.11	82.39	43.01	79.44	76.56	77.96	
Llava-OV1.5-8B	91.28	75.23	81.03	68.54	64.36	72.28	75.51	48.48	46.62	40.86	45.80	20.77	68.36	57.08	45.44	82.33	79.84	74.50	75.00	87.79	73.17	71.07	43.52	83.80	71.09	76.19	
MiniCPM-V4.5-8B	89.93	72.48	75.86	64.04	70.30	75.00	74.55	48.82	54.05	63.98	54.52	19.63	64.39	48.75	42.05	87.07	80.93	73.37	76.85	87.46	70.33	77.36	40.93	81.62	72.66	76.55	
VideoLLaMA3-7B	90.60	68.81	45.22	31.03	20.21	68.16	54.00	48.82	60.81	27.42	45.32	22.35	69.95	51.25	45.88	77.29	79.29	71.39	53.70	77.23	51.22	64.15	31.09	82.87	66.41	68.90	
VideoChat-7B	72.48	77.06	76.72	60.11	73.27	66.30	71.01	48.82	55.41	14.52	40.25	20.49	72.81	32.08	43.27	82.65	78.75	75.64	77.78	87.46	68.29	69.18	38.34	60.12	70.31	72.22	
Open-source Online Video-LLMs																											
Flash-VStream-7B	25.50	32.11	29.31	33.71	29.70	28.80	29.86	36.36	33.78	5.91	25.35	5.44	67.25	60.00	44.23	24.91	25.89	23.87	18.52	23.87	25.20	27.33	48.70	13.08	43.57	23.23	
Flash-VStream-7B*	69.09	50.87	68.07	46.30	66.85	58.09	59.88	55.58	63.71	20.00	46.43	17.73	66.50	58.00	47.41	81.24	79.73	79.40	77.14	71.36	68.28	70.01	50.67	79.18	79.11	74.48	
ReKV-7B	68.44	47.34	63.44	43.64	67.04	54.14	57.34	55.22	63.35	15.50	44.69	13.00	65.00	55.05	44.35	69.11	69.51	70.05	67.07	58.54	59.50	62.65	52.81	65.66	74.44	64.53	
StreamForest-7B	68.46	53.21	71.55	47.75	65.35	60.87	61.20	58.92	64.86	32.26	52.02	32.81	70.59	57.08	53.49	82.65	83.11	84.26	76.85	75.64	69.11	77.50	54.40	78.19	82.81	77.26	
TimeChatOnline-7B	69.80	48.60	64.70	44.90	68.30	55.40	58.60	53.90	62.80	9.10	42.00	32.50	36.50	40.00	36.40	78.86	79.13	80.77	77.78	66.19	67.07	70.44	53.72	77.26	81.25	73.64	
VideoChatOnline-4B	47.22	40.21	46.68	39.20	48.65	41.86	43.97	49.05	54.78	8.37	37.40	19.07	54.00	46.00	39.69	62.99	63.47	63.56	60.35	53.74	54.75	57.76	52.24	58.38	70.17	58.81	

Table 3: Detailed experimental results for the experiments reported in Section 4.3

Table 4: Ranking changes of representative models under different deployment preferences.

Model / Type	Best Answer ($w_a = 0.4$)	Interaction First ($w_i = 0.4$)	Edge Resource-Saving ($w_r = 0.4$)	Throughput First ($w_f = 0.4$)
Qwen3-VL-7B (Offline)	Rank 1	Rank 2	Rank 2	Rank 2
Flash-VStream (Online)	Rank 2	Rank 1	Rank 1	Rank 1
MiniCPM-V4.5 (Offline)	Rank 3	Rank 3	Rank 3	Rank 5
StreamForest (Online)	Rank 9	Rank 10	Rank 9	Rank 9

ric captures strengths under distinct deployment constraints (e.g., the offline model **Qwen3-VL** becomes **Rank 1** in the “best-answer” setting, while the online model **Flash-VStream** is consistently strongest in “interaction-first,” “resource-saving,” and “throughput-first” settings). Importantly, despite scenario-specific re-ordering among top models, the overall ranking trend remains statistically robust across weight configurations (Spearman $\rho \in [0.972, 0.993]$, Kendall $\tau \in [0.909, 0.970]$), which reduces the risk that models with clear holistic weaknesses can game the leaderboard by merely tilting weights.

In-depth Discussion. As shown in Table 6, different methods exhibit clear trade-offs across task clusters on OVO-Bench. Overall, open-source offline Video-LLMs are particularly strong on the “Real-Time Visual Perception” subset: Qwen3-VL achieves the highest average score in this subset (78.73), while InternVL3.5, LLaVA-OneVision, and MiniCTM remain at a comparable level (around 74–76), suggesting that the offline inference paradigm is more robust for fine-grained perception tasks (e.g., OCR, attribute recognition, and state understanding). In contrast, open-source online (streaming/online) Video-LLMs are generally weaker on the same perception subset, indicating that under online settings constrained by temporal input processing and immediate generation requirements low-level visual perception is more

likely to become a bottleneck. Overall, offline methods tend to be stronger in fine-grained visual perception and robustness, while online methods need further improvements to close the gap in perceptual capability.

Notable Points about E2E Latency. Beyond the efficiency metrics reported in the main text, we conduct an additional analysis of end-to-end (E2E) latency. In streaming generation, E2E latency is directly tied to user experience: perceived waiting time is determined not only by the time-to-first-token (TTFT), but also by the efficiency of subsequent decoding and whether the models responses are concise and clear. We define E2E latency as the wall-clock time from query arrival to the completion of the final output, i.e., when the last token is generated and returned.

To measure this metric, we convert the OVO-Bench multiple-choice QA format into an open-ended generation setting and record E2E latency under a unified inference configuration. Importantly, the main-table results are obtained under a multiple-choice setup, whereas E2E latency can only be meaningfully measured in an open-ended generation scenario. Because these two evaluation settings differ fundamentally in output form, reporting E2E alongside the multiple-choice metrics in the main table could mislead readers into assuming they are directly comparable under an identical setup. For clarity and to avoid ambiguity, we therefore present the E2E experiments and discussions in the appendix.

Under the same prompting and decoding configurations, E2E latency is largely influenced by the base language models generation style and stopping behavior; meanwhile, most models compared in the main text share similar base language models. Accordingly, in the appendix we focus

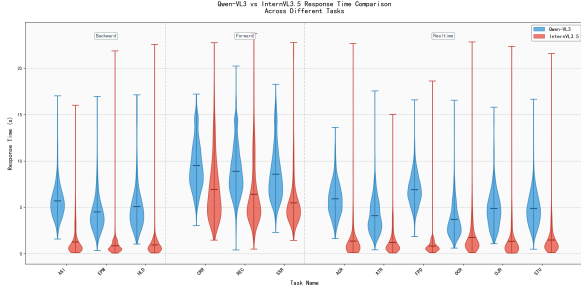


Figure 4: Comparison of two classic paradigms for video understanding.

on two representative families Qwen3-VL and InternVL3.5 and provide a comparative analysis of their E2E latency on OVO-Bench. As shown in Fig. 4, InternVL3.5 (red) consistently achieves lower end-to-end (E2E) latency than Qwen-VL3 (blue) across OVO-Bench tasks, with a tighter distribution and fewer long-tail cases. This trend holds across the Backward (ASI/EPM/HLD) and Realtime (ACR/ATR/FPD/OCR/OJR/STU) categories, where InternVL3.5 exhibits markedly smaller median latency and improved stability. For the Forward tasks (CRR/REC/SSR), both models show increased latency and heavier tails, reflecting the higher uncertainty and longer generations in open-ended settings; nevertheless, InternVL3.5 remains faster in central tendency. Overall, Fig. 4 suggests that InternVL3.5 provides not only lower average latency but also more stable user-perceived responsiveness, whereas Qwen-VL3 is more susceptible to long-tail slowdowns.

Factor	Value
GPU model	NVIDIA GeForce RTX 4090-48G
Peak throughput	40.32 TFLOPS
GPU memory (VRAM)	48 GB
GPU memory bandwidth	1008.10 GB/s
PCIe lanes	16
PCIe bandwidth	31.50 GB/s
CPU model	Intel Xeon Platinum 8570
Available CPU cores	20
Available RAM	48 GB
Disk model	SAMSUNG MZ7LH480 / SAMSUNG MZ7LH480
Disk bandwidth	330.08 MB/s
Free disk space	200 GB

Table 5: Hardware factors affecting performance

Detailed Analysis across Different Memory Budget. As shown in Table 6, the three mod-

els exhibit consistent strengths across the three OVO-Bench task clusters. For Real-Time Visual Perception (OCR/ACR/ATR/STU/FPD/OJR), Qwen3-VL delivers the strongest and most stable overall performance, reaching a peak average of 78.73 at 0.5G and maintaining a high level (above 77) at larger scales; Llava-OV1.5 is typically second best, while InternVL3.5 is comparatively weaker and shows a mild degradation at the largest scale (e.g., 71.21 at 1.5G). For Backward Tracing (EPM/ASI/HLD), Qwen3-VL consistently leads (approximately 50.9–53.1 in average), whereas InternVL3.5 lags behind, largely due to persistently low HLD scores across scales. In contrast, the ranking on Forward Active Responding (REC/SSR/CRR) is more dynamic: InternVL3.5 surpasses Qwen3-VL at several scales and remains consistently stronger on REC, while Llava-OV1.5 benefits more from scaling and achieves the best cluster average at 1.0G/1.5G (49.72/50.00). Overall, Qwen3-VL is more competitive on perception and backward-tracing tasks, InternVL3.5 is comparatively stronger on active responding, and Llava-OV1.5 exhibits the largest scaling gains on forward tasks.

Detailed Analysis across Different Resolutions.

As shown in Table 7, increasing the input resolution ($224 \times 224 \rightarrow 336 \times 336 \rightarrow 448 \times 448$) leads to a structurally different impact across task clusters, and the scaling benefits vary by model. For Real-Time Visual Perception, higher resolution consistently yields the largest and most stable gains: Qwen3-VL improves from 67.98 to 73.36 and further to 75.75 in average score, while InternVL3.5 shows an even larger increase from 61.65 to 68.70 and then to 74.31, indicating that fine-grained visual cues (e.g., text details and local attributes) substantially benefit from higher-resolution inputs. In contrast, the effect on Backward Tracing is modest and less consistent: Qwen3-VL only slightly improves ($49.76 \rightarrow 50.55 \rightarrow 51.35$), and InternVL3.5 remains nearly flat ($41.68 \rightarrow 40.73 \rightarrow 40.89$), suggesting that the bottleneck of backward tracing is less dominated by visual detail. For Forward Active Responding, resolution scaling comdoes not provide a uniform advantage and can even be neutral or negative (e.g., Qwen3-VL slightly drops from 44.35 to 43.01 and 42.18), implying that these tasks are more sensitive to language-side generation and decision behaviors than to input resolution. Overall, resolution scal-

ing primarily benefits perception-oriented tasks, while its impact on backward and forward tasks is comparatively limited and model-dependent.

Model	OVO-Bench															
	Real-Time Visual Perception							Backward Tracing				Forward Active Responding				
	OCR	ACR	ATR	STU	FPD	OJR	Avg.	EPM	ASI	HLD	Avg.	REC	SSR	CRR	Avg.	
0.1G																
Qwen3-VL	84.56	73.39	76.72	62.36	70.30	75.00	73.48	42.76	54.73	60.75	50.87	14.90	59.30	43.33	37.08	
InternVL3.5	87.25	76.15	75.86	66.29	69.31	70.65	73.95	49.16	56.08	15.59	40.89	19.77	57.07	45.42	38.67	
Llava-OV1.5	86.91	72.94	78.87	67.98	61.89	67.66	72.71	44.44	42.90	48.66	45.33	11.87	57.58	48.58	39.34	
0.3G																
Qwen3-VL	92.62	77.98	77.59	69.66	70.30	76.63	77.54	45.45	57.43	47.31	48.81	19.63	64.07	47.50	41.74	
InternVL3.5	87.92	77.06	73.28	60.11	69.31	73.91	73.24	50.84	56.08	13.44	41.05	23.07	67.57	49.58	44.99	
Llava-OV1.5	91.28	75.69	78.02	68.54	61.89	70.10	74.25	46.62	44.25	40.86	43.91	15.40	64.74	52.26	44.14	
0.5G																
Qwen3-VL	91.95	78.90	79.31	67.98	73.27	80.98	78.73	48.82	61.49	48.92	51.82	21.78	64.86	50.42	43.46	
InternVL3.5	88.59	75.23	77.59	61.80	71.29	73.91	74.31	51.18	56.76	11.83	40.89	25.36	67.73	50.00	46.14	
Llava-OV1.5	91.28	75.23	81.03	68.54	64.36	72.28	75.51	48.48	46.62	40.86	45.80	20.77	68.36	57.08	45.44	
1.0G																
Qwen3-VL	91.95	76.15	80.17	66.29	71.29	78.80	77.42	50.84	64.86	47.31	53.09	24.64	67.41	52.92	46.14	
InternVL3.5	87.25	72.48	75.86	60.67	71.29	72.28	72.88	51.18	57.43	10.22	40.57	29.37	66.61	54.58	48.18	
Llava-OV1.5	90.61	72.48	80.60	67.13	63.37	70.38	74.09	49.49	48.64	39.25	45.79	22.62	67.49	59.04	49.72	
1.5G																
Qwen3-VL	91.28	77.06	79.31	67.98	74.26	77.17	77.66	51.85	65.54	43.01	52.46	25.93	67.73	53.33	46.90	
InternVL3.5	85.91	69.72	75.86	60.11	73.27	66.85	71.21	56.57	61.49	11.83	44.53	31.09	63.75	53.75	47.67	
Llava-OV1.5	89.60	71.56	80.16	67.70	65.85	66.85	73.62	52.69	51.01	37.91	47.20	24.40	66.50	59.10	50.00	

Table 6: Detailed experimental results for the experiments reported in Section 4.4

Model	OVO-Bench															
	Real-Time Visual Perception							Backward Tracing				Forward Active Responding				
	OCR	ACR	ATR	STU	FPD	OJR	Avg.	EPM	ASI	HLD	Avg.	REC	SSR	CRR	Avg.	
224*224																
Qwen3-VL	69.80	68.81	76.72	59.55	67.33	69.02	67.98	50.51	60.81	39.78	49.76	22.78	67.25	47.08	44.35	
InternVL3.5	65.77	59.63	66.38	58.99	62.38	58.70	61.65	51.85	59.46	11.29	41.68	24.64	62.16	59.58	45.05	
FlashVStream-LLaVA	25.50	32.11	29.31	33.71	29.70	28.80	29.86	36.36	33.78	5.91	25.35	5.44	67.25	60.00	44.23	
336*336																
Qwen3-VL	81.88	78.90	76.72	64.61	70.30	71.20	73.36	49.83	59.46	44.62	50.55	20.92	66.45	45.83	43.01	
InternVL3.5	79.19	67.89	71.55	58.43	72.28	66.85	68.70	52.53	56.76	9.14	40.73	27.51	65.02	56.67	47.03	
FlashVStream-LLaVA	30.31	36.69	30.45	34.98	32.71	31.25	32.73	36.35	32.64	5.71	24.90	5.53	68.40	57.74	43.89	
448*448																
Qwen3-VL	86.58	77.06	75.00	66.29	73.27	77.17	75.75	51.52	58.11	45.70	51.35	19.34	66.45	45.00	42.18	
InternVL3.5	88.59	75.23	77.59	61.80	71.29	73.91	74.31	51.18	56.76	11.83	40.89	25.36	67.73	50.00	46.14	
FlashVStream-LLaVA	32.99	38.24	31.46	36.42	33.13	34.23	34.41	36.49	32.26	6.49	25.08	5.11	69.86	53.85	42.94	

Table 7: Detailed experimental results for the experiments reported in Section 4.4