

TEMPCORE: Are Video QA Benchmarks Temporally Grounded?

Hyunjong Ok Jaeho Lee
POSTECH

hyunjong.ok@gmail.com, jaeho.lee@postech.ac.kr

Abstract

Vision-language models (VLMs) can ingest only a limited number of video frames, making frame selection a practical necessity. But do current Video QA benchmarks genuinely require temporal frame selection, or can most questions be answered regardless of which frames are shown? We introduce **Frame Selection Sensitivity** (FSS), a per-sample diagnostic that measures how much VLM accuracy changes when the most relevant frames are replaced with the least relevant ones. Across six benchmarks and eight VLMs, we find that a large majority of samples are frame-agnostic: only a minority are genuinely sensitive to frame choice. Combining FSS with a Language Independence Score (LIS) reveals that merely 5.5–31% of samples are *Temporally Sensitive*. We construct TEMPCORE, compact evaluation subsets that isolate these temporal samples from existing benchmarks, and will release code and per-sample annotations upon publication.

1 Introduction

Large language models (LLMs) have transformed natural language processing, and their extension to vision-language models (VLMs) has brought comparable advances to image understanding (Li et al., 2025; Zhu et al., 2025; Bai et al., 2025b). These developments have spurred growing interest in applying VLMs to *video* understanding, where temporal dynamics—event ordering, causal transitions, and changes over time—are fundamental.

Extending VLMs to video processing, however, is fundamentally constrained by input capacity: videos may contain thousands of frames, but current models can process only around 32 at a time (Lin et al., 2024; Tang et al., 2025). The standard approach is to rank frames by relevance using encoders such as SigLIP (Zhai et al., 2023; Tschannen et al., 2025) or CLIP (Radford et al., 2021), and retain the top- K frames (Tang et al.,

2025; Himakunthala et al., 2023; Shen et al., 2025). Accordingly, a growing line of work explores more sophisticated selection strategies—including combinatorial ranking, spatiotemporal saliency, hierarchical clustering, and generative index prediction (Yu et al., 2025; Hu et al., 2025; Wang et al., 2025; Yao et al., 2025)—all based on the premise that selecting the right frames is critical.

While these methods report gains on standard benchmarks, a closer look complicates the picture. On some benchmarks, simple uniform sampling performs comparably to learned selection (Brkic et al., 2025), and a growing literature suggests that VLMs are largely “temporally blind,” exhibiting limited sensitivity to event ordering and duration (Upadhyay et al., 2026; Jung et al., 2025; Cai et al., 2024). If temporally insensitive models can still achieve high scores, the benchmarks themselves may not genuinely require temporal reasoning—and if most questions can be answered from a single frame or from text alone, then frame selection is irrelevant regardless of the method. This tension raises two intertwined questions:

1. *Does key-frame selection actually improve Video QA performance?*
2. *Do current benchmarks genuinely require temporal reasoning?*

To address these questions, we examine whether encoder-based frame relevance translates into meaningful differences in downstream VLM accuracy. Concretely, we score every video frame with a SigLIP encoder and compare VLM accuracy when the model receives the most relevant (MAXPROB) versus least relevant (MINPROB) frames. This yields **Frame Selection Sensitivity** (FSS), a per-sample diagnostic that partitions each benchmark sample into *Key* (temporally sensitive), *Non-key* (frame-agnostic), or *Anti-key* (hurt by “relevant” frames). Applying FSS to six benchmarks with 8 VLMs ranging from 256M to 8B parameters—

the largest systematic study of frame selection to date—we answer:

Answer to Q1: Frame selection works, but only on long videos. On long-video benchmarks, MAXPROB consistently outperforms MINPROB, while on short-video benchmarks, the effect is near zero. Our approach uses a general-purpose image-text encoder (SigLIP) in question-answer (QA-mode) scoring, which leverages full question-answer context for more accurate temporal classification; encoder improvements—such as question-aware models—could narrow this gap further. The large oracle headroom confirms that substantial temporal signal exists in the video but remains unexploited by current selection methods.

Answer to Q2: Most existing benchmarks are not temporally grounded. Only 6–36% of samples are Key; the vast majority are Non-key or Anti-key. A complementary **Language Independence Score** (LIS) reveals that only 5.5–31.0% of samples are *Temporally Sensitive*—both vision-dependent and frame-sensitive (§7.1). No-frame baselines reach up to 50% accuracy on language priors alone, and aggregate scores are dominated by Non-key samples, allowing models with strong language priors to appear competitive even when they fail on frame-sensitive temporal questions.

Based on this analysis, we construct TEMP-CORE, temporally grounded evaluation subsets that retain only these samples (§7), providing a focused measure of frame-selection sensitivity.

Contributions.

1. We propose FSS, a per-sample diagnostic that quantifies temporal sensitivity, and conduct a systematic frame selection study (§3, §5).
2. We introduce a two-dimensional FSS×LIS taxonomy that reveals only 5.5–31.0% of samples are *Temporally Sensitive* (§6, §7).
3. We construct TEMP-CORE, temporally grounded evaluation subsets extracted from existing benchmarks without manual re-annotation (§7).

2 Related Work

Large language models and multimodal LLMs. The expansion of large language models (LLMs) has enabled strong performance across various general-purpose language tasks (Achiam et al., 2023; Grattafiori et al., 2024; Yang et al., 2024;

Gemma Team et al., 2024). Building upon this foundation, research has extended such capabilities to other modalities by leveraging the knowledge of LLMs (Alayrac et al., 2022; Liu et al., 2023; Chu et al., 2024; Wang et al., 2024; Zhu et al., 2025). Video MLLMs, in particular, have achieved strong performance in video question answering, captioning, and summarization (Zhang et al., 2023; Lin et al., 2024; Zhu et al., 2024; Maaz et al., 2024; Shen et al., 2025; Zhang et al., 2025a).

Key frame selection. Key frame selection reduces the computational cost of long-video understanding by identifying the most relevant frames for a given query. Encoder-based methods score individual frames via vision-language similarity: Tang et al. (2025) select frames that are both query-relevant and visually diverse to avoid redundancy, Yu et al. (2025) rank frame subsets by how well they complement each other rather than scoring frames independently, and Hu et al. (2025) leverage multimodal LLM prompting to derive spatial importance and temporal selection signals for frame ranking. A second line of work explores structured selection strategies: Wang et al. (2025) build a hierarchical tree of visual segments that enables query-adaptive retrieval from long videos, Buch et al. (2025) learn to dynamically adjust how many frames to select per query instead of using a fixed budget, and Sun et al. (2025) optimize frame orderings using a list-wise ranking objective that considers the full sequence jointly. These methods *improve* selection; we *diagnose* when selection matters and construct evaluation subsets that isolate temporal ability.

Temporal diagnostics & benchmarks. Several studies document that VLMs struggle with temporal reasoning: near-zero accuracy on ordering tasks (Upadhyay et al., 2026), inconsistency across equivalent temporal queries (Jung et al., 2025), and failures in fine-grained temporal understanding (Cai et al., 2024). Complementary work shows that many Video QA questions can be solved from a single frame or without temporal reasoning: Cores et al. (2025) show that most MVBench tasks can be solved without temporal reasoning, with single-frame and text-only baselines remaining competitive; Lei et al. (2023) demonstrate similar single-frame biases across multiple video-language benchmarks. On the diagnostic side, counterfactual interventions (Niu et al., 2021), text-bias reduction (Kulgod et al., 2026), confidence intervals (Du,

2026), and sampling strategy comparisons (Brkic et al., 2025) have been proposed to isolate specific biases. Dedicated temporal benchmarks take complementary approaches. TempCompass (Liu et al., 2024) constructs synthetic test cases by applying video-level manipulations—reversal, spatial concatenation, and temporal concatenation—to probe whether models truly perceive temporal order rather than relying on static appearance. TemporalBench (Cai et al., 2024) instead enlists expert annotators to label fine-grained temporal events within each clip, producing densely annotated samples that test ordering, duration, and causal sequencing. TEMPCORE differs on two fronts: (i) it provides a fully automated, model-based diagnostic pipeline that requires no manual annotation or synthetic video generation, and (ii) it draws from six widely adopted benchmarks covering egocentric, perceptual, and causal-reasoning domains, yielding a compact yet diverse evaluation set.

To our knowledge, FSS–LIS is the first model-side diagnostic that assigns per-sample temporal-sensitivity labels to existing Video QA benchmarks without additional annotation. Rather than asking how to select frames more effectively, as prior frame-selection methods do, we use the MAXPROB/MINPROB contrast together with a no-frame language-prior axis to audit whether benchmark samples actually depend on temporal visual evidence. Dedicated temporal benchmarks instead construct new test cases or annotations; our approach identifies temporally grounded samples directly within benchmarks already in widespread use.

3 Diagnostic Framework

3.1 Frame Scoring

Given a video $V = (f_1, f_2, \dots, f_L)$ of L frames uniformly extracted at 10 FPS, and a text query t (question, or concatenation with answer), we compute a per-frame relevance score using a pretrained SigLIP vision-language encoder (Zhai et al., 2023; Tschannen et al., 2025):

$$s(f_i, t) = \sigma(\cos(E_V(f_i), E_T(t))), \quad (1)$$

where E_V and E_T are the vision and text encoders, $\cos(\cdot, \cdot)$ denotes cosine similarity, and σ is the sigmoid function.

Text modes. We employ two modes for t : in Q-mode t is the question text, and in QA-mode t is

the concatenation of the question and the correct answer. QA-mode requires oracle answer labels and serves as an upper bound on encoder-based selection.

3.2 Frame Selection Strategies

Using the frame scores $\{s(f_i, t)\}_{i=1}^L$, we define five different strategies:

MAXPROB. Select the top- K scoring frames:

$$\mathcal{F}^+ = \text{top-K}(f_1, \dots, f_L; s(\cdot, t)). \quad (2)$$

MINPROB. Select the bottom- K scoring frames:

$$\mathcal{F}^- = \text{bot-K}(f_1, \dots, f_L; s(\cdot, t)). \quad (3)$$

This adversarial baseline serves as the diagnostic counterpart to MAXPROB.

Temporal diversity constraint. Adjacent frames within the same shot often receive nearly identical scores, causing unconstrained top- K selection to concentrate on only a few temporal regions. This effect is especially severe for MINPROB, and can therefore be mistaken for genuine frame sensitivity. To ensure comparable temporal coverage, we group frames by integer second and select them in rounds. In each round, we take the next highest-scoring frame from every second, so that the first $\min(K, \lceil L/\text{fps} \rceil)$ selected frames come from distinct seconds. MINPROB uses the same procedure with the ranking reversed.

Uniform. Select K frames uniformly distributed over the video timeline. This is the standard baseline in most deployed systems.

Window oracle. Slide a window of K frames with stride S across the video, run the VLM on each window, and report the best accuracy. When the total number of windows exceeds a budget $W_{\max}$ (default 50), we uniformly subsample $W_{\max}$ windows to bound computation. This measures the theoretical upper bound achievable by any temporal selection strategy given window size K .

No-frame. Replace all input frames with a uniform gray dummy image. The VLM processes the question but receives no visual content.

3.3 Frame Selection Sensitivity (FSS)

To measure how much each sample depends on receiving the right frames, we define *Frame Selection Sensitivity* (FSS). Let $\mathcal{M}$ denote the set of models,

and let $\text{acc}_{m,i}^+$ and $\text{acc}_{m,i}^-$ be the binary correctness (1 or 0) of model m on sample i when given MAX-PROB and MINPROB frames, respectively.

$$\text{FSS}_i = \frac{1}{|\mathcal{M}|} \sum_{m \in \mathcal{M}} (\text{acc}_{m,i}^+ - \text{acc}_{m,i}^-). \quad (4)$$

$\text{FSS}_i > 0$ indicates that the sample benefits from relevant frames; $\text{FSS}_i < 0$ indicates that relevant frames *hurt*—a visual appearance bias.

Sample categories. Given a threshold τ (default $\tau = 0.15$), we classify each sample as:

- *Key* ($\text{FSS}_i > \tau$): temporally sensitive; benefits from correct frame selection.
- *Non-key* ($|\text{FSS}_i| \leq \tau$): frame-agnostic; answered regardless of frame choice.
- *Anti-key* ($\text{FSS}_i < -\tau$): appearance-biased; “relevant” frames are detrimental.

This classification is an *operational* definition of temporal sensitivity: it captures whether frame selection affects model behavior, which is a necessary—but not necessarily sufficient—condition for a question to require temporal reasoning. FSS-based pruning retains samples with $|\text{FSS}| > \tau$, yielding an evaluation subset in which frame selection is consequential.

Choice of τ . We set $\tau = 0.15$, so that a Key or Anti-key label requires more than one model’s net flip. Agreement with human temporal judgments and ablations over τ are in Section B.2.

Default configuration. We adopt QA-mode as the primary FSS scoring mode, as it more accurately identifies temporally sensitive questions by leveraging full question-answer context. Since FSS is an analytical diagnostic, maximizing accuracy takes priority over label-free operation.

3.4 Language Independence Score (LIS)

Let $\text{correct}_{m,i}^{\text{nf}}$ be the binary correctness of model m on sample i under the no-frame strategy.

$$\text{LIS}_i = 1 - \frac{1}{|\mathcal{M}|} \sum_{m \in \mathcal{M}} \text{correct}_{m,i}^{\text{nf}}. \quad (5)$$

$\text{LIS}_i = 1$ when no model answers correctly without meaningful visual input, indicating maximal vision dependence; $\text{LIS}_i = 0$ when all models answer correctly, indicating that the item is solvable from language priors alone.

4 Experimental Setup

4.1 Models

We evaluate eight VLMs from five families: LLaVA-OneVision (Li et al., 2025), VideoL-LaMA3 (Zhang et al., 2025a), SmolVLM2 (Marafioti et al., 2025), Qwen3-VL (Bai et al., 2025a), and InternVL3 (Zhu et al., 2025), spanning 256M to 8B parameters. All models receive $K = 32$ frames in a zero-shot multiple-choice setting; ablations over $K \in \{16, 32, 48\}$ appear in Section C.1. We compute frame relevance scores (Equation (1)) with SigLIP2-400M (Tschannen et al., 2025) as the default; ablations with SigLIP-400M and CLIP ViT-L/14 appear in Section C.6. See Table 5 for full details.

4.2 Benchmarks

We use six Video QA benchmarks grouped by video length. *Short* (≤ 3 min): MVBench (Li et al., 2024), NExTQA (Xiao et al., 2021), and EgoSchema (Mangalam et al., 2023). *Long* (> 5 min): MLVU (Zhou et al., 2025), LongVideoBench (Wu et al., 2024), and VideoMME (Fu et al., 2025). Detailed information and full statistics appear in Table 6. Design intent and multi-axis linguistic profiles of the six benchmarks appear in Section A.2.

Roadmap. We first examine whether frame selection affects VLM accuracy (§5), then characterize how temporally grounded current benchmarks are (§6), and finally construct TEMPCORE, temporally grounded evaluation subsets (§7).

5 Does Frame Selection Work?

5.1 RQ1: Are Vision Encoders Robust to Interrogative Cues?

Vision-language similarity is an intuitive basis for frame scoring, yet VQA queries are inherently interrogative: they contain words like “who,” “what,” and “how many” that carry little visual grounding. Since encoders such as SigLIP are trained on declarative image–caption pairs, their similarity estimates may be poorly calibrated for question-form inputs. We test this by comparing Q-mode (question only) with QA-mode (question + ground-truth answer), where the latter supplies a declarative phrase that better matches the encoder’s training distribution. Table 1 summarizes benchmark-averaged accuracy.

QA-mode outperforms Q-mode on all six benchmarks, with an average gain of +2.3% (Table 1),

Table 1: **QA-mode scoring beats Q-mode on all six benchmarks.** MAXPROB accuracy (%) averaged across the eight VLMs, with $\Delta = \text{QA-mode} - \text{Q-mode}$. Supplying the answer phrase gives the encoder a declarative target closer to its training distribution.

Benchmark	Q-mode	QA-mode	Δ
<i>Short-video</i>			
MVBench	48.4	49.0	+0.6
NExTQA	58.5	59.2	+0.7
EgoSchema	34.2	39.2	+5.0
<i>Long-video</i>			
MLVU	54.9	57.3	+2.4
LongVideoBench	46.7	47.8	+1.1
Video-MME	46.0	49.7	+3.7
Average	48.1	50.4	+2.3



Figure 1: **Grad-CAM attention maps for question-text vs. answer-text scoring.** Question-text attention is diffuse, while answer-text attention concentrates on the relevant region, illustrating the noise introduced by interrogative cues.

indicating that encoders are weakened by interrogative cues. The smallest gain among the long-video benchmarks falls to LongVideoBench, whose questions are up to nearly four times longer *in word count* than those of other benchmarks and already contain rich visual descriptions, making Q-mode nearly sufficient on its own (see Section C.7 for statistics and examples).

Figure 1 further illustrates this vulnerability via Grad-CAM (Selvaraju et al., 2017) attention maps: when scoring with question text, attention is diffuse across the frame, whereas answer text concentrates attention on the task-relevant region. This suggests that ambiguous, interrogative queries yield noisier similarity estimates.

Takeaway. These results suggest that text-based frame selection methods must be robust to interrogative and ambiguous input cues. Developing encoders that handle question-form queries as effectively as declarative captions remains an important direction for future work.

5.2 RQ2: MaxProb vs. MinProb

Despite the interrogative-cue limitation identified above, does vision-language similarity still pro-




Figure 2: **Frame selection helps on every benchmark, but three times more on long videos.** MAXPROB vs. MINPROB accuracy (8-model average). The mean gap is +3.5 pp on short-video benchmarks against +10.2 pp on long-video ones. LVB denotes LongVideoBench.

vide an effective frame selection signal? Figure 2 compares MAXPROB and MINPROB accuracy (8-model average) across six benchmarks (full per-model results in Table 21).

On long-video benchmarks, MAXPROB consistently outperforms MINPROB, with an average gain of +10.2% across models. On short-video benchmarks (MVBench, NExTQA), the gain is markedly smaller: in short clips, most $K = 32$ windows already contain the relevant content, so the scorer’s ranking provides little additional benefit.

Takeaway. Vision-language similarity is an effective frame selection signal for long videos, where relevant content is temporally sparse. On short videos, the benefit shrinks to a few points—temporal information demand is inherently low, and any $K = 32$ window already covers most of the relevant content.

5.3 RQ3: Oracle Gap

How much room remains for better frame selection? Figure 3 compares window-oracle accuracy—the best accuracy achievable across all K -frame sliding windows—against a uniform-sampling baseline (full per-model results in Table 22).

The gaps are substantial: long-video benchmarks average +25.2% while short-video benchmarks average +9.9%. However, the oracle selects the single best window per sample, and different windows may elicit correct answers through distinct reasoning paths, potentially inflating the gap beyond pure temporal content.

To control for this concern, Table 2 decomposes oracle gaps by FSS category. Key samples exhibit consistently larger gaps than Non-key samples on all six benchmarks, confirming that the more temporal a sample is, the more headroom remains for improved frame selection.




Figure 3: **Oracle gaps reveal large untapped headroom for frame selection.** Uniform-sampling vs. window-oracle accuracy (model average). Long-video gaps far exceed short-video gaps, confirming that temporal frame selection is most consequential for extended videos. LVB denotes LongVideoBench.




Figure 4: **FSS scores concentrate near zero, confirming non-temporal dominance.** Aggregated FSS distributions for short-video and long-video benchmarks. Dashed lines mark classification thresholds ($\tau = \pm 0.15$). The heavy mass near zero shows that frame selection has negligible effect on the majority of samples. Per-benchmark breakdowns are in Section D.1.

Takeaway. Large oracle gaps indicate substantial room for better frame selection, and this gap is concentrated in temporally sensitive samples. Longer benchmarks with higher temporal purity stand to benefit the most from advances in frame retrieval.

6 How Temporal Are Video QA Benchmarks?

6.1 FSS Distributions

FSS is computed from MAXPROB and MINPROB accuracy (defined in §3). Figure 4 shows the distribution of FSS scores across samples in each benchmark. The overwhelming majority of samples cluster near zero, indicating that most Video QA items are insensitive to frame selection.

Quantitative summary. Table 3 reports Key and Anti-key fractions for each benchmark. Key fractions order by video length: MVBench is lowest at 5.6% and MLVU highest at 36.0%, with EgoSchema between the two groups at its intermediate clip length. Anti-key fractions stay small

Benchmark	Split	N	Uniform	Oracle	Gap
<i>Short-video</i>					
MVBench	All	3786	48.7	54.5	+5.8
	Key	212	52.2	65.4	+13.2
	Non-key	3502	48.4	53.7	+5.3
	Anti-key	72	50.9	61.3	+10.4
NExTQA	All	8438	58.9	68.0	+9.1
	Key	656	51.6	69.2	+17.5
	Non-key	7588	59.7	67.9	+8.2
	Anti-key	194	51.7	68.4	+16.7
EgoSchema	All	500	38.1	52.8	+14.7
	Key	94	41.4	61.7	+20.3
	Non-key	390	37.1	50.2	+13.1
	Anti-key	16	43.0	64.1	+21.1
<i>Long-video</i>					
MLVU	All	2173	51.5	77.3	+25.8
	Key	782	50.6	81.1	+30.4
	Non-key	1212	52.3	74.5	+22.2
	Anti-key	179	50.1	79.8	+29.7
LongVideoBench	All	1122	46.9	67.5	+20.5
	Key	229	46.7	76.1	+29.3
	Non-key	823	47.2	64.9	+17.7
	Anti-key	69	44.9	71.3	+26.4
Video-MME	All	2625	46.4	73.9	+27.6
	Key	783	50.9	82.6	+31.6
	Non-key	1673	44.2	69.3	+25.1
	Anti-key	168	46.4	79.6	+33.2

Table 2: **Key and Anti-key samples exhibit larger oracle gaps than Non-key samples on all six benchmarks.** Oracle gap (%) decomposed by FSS category, on the samples with both window-oracle and uniform-strategy results. Largest per-benchmark gap in bold, demonstrating that FSS-classified samples are precisely those where temporal frame selection is consequential.

throughout, so relevant frames help far more often than they hurt.

6.2 Implications for Benchmark Design

Our analysis calls for a reinterpretation of current Video QA benchmarks. Only a small fraction of samples in any benchmark are temporally sensitive to frame selection; the remainder are resolved by language priors, visual appearance, or other non-temporal cues. This does not invalidate existing benchmarks—appearance and semantic comprehension are important capabilities—but it underscores that *temporal* performance should be reported separately. FSS-based pruning (retaining $|FSS| > \tau$) provides a lightweight, model-agnostic mechanism to do so. We address this by combining FSS with the Language Independence Score (LIS) to construct TEMPCORE, temporally grounded evaluation subsets.




Figure 5: **Overview of the TEMPCORE construction pipeline.** Stage 1 filters out Trivial samples (solvable from language priors, $LIS \leq \lambda$). Stage 2 classifies the remaining vision-dependent samples along the FSS axis into three categories: Temporally Grounded ($FSS > \tau$), Frame-Agnostic ($|FSS| \leq \tau$), and Visual Bias ($FSS < -\tau$). The Temporal Purity Index (TPI) summarizes the fraction of vision-dependent samples that are temporally sensitive.

Benchmark	Key (%)	Anti-key (%)	Key/AK
<i>Short-video</i>			
MVBench	5.6	1.9	2.94
NExTQA	7.7	2.3	3.35
EgoSchema	18.8	3.2	5.88
<i>Long-video</i>			
MLVU	36.0	8.2	4.37
LongVideoBench	21.3	6.4	3.31
Video-MME	29.7	6.4	4.63

Table 3: **Key fractions order by video length** ($\tau = 0.15$). “Key/AK” is the ratio of Key to Anti-key samples, above 2.9 on every benchmark: relevant frames help far more often than they hurt.

7 TEMPCORE

FSS captures sensitivity to frame selection but cannot distinguish genuinely temporal samples from those solvable by language priors alone—a high-FSS question may simply contain answer-leaking text cues. To filter out such cases, we combine FSS with the Language Independence Score (LIS, introduced in §3.4) and construct TEMPCORE, temporally grounded evaluation subsets extracted from existing benchmarks without manual re-annotation (Figure 5).

7.1 Two-Dimensional Taxonomy

Combining LIS and FSS yields a three-tier taxonomy, with Temporally Sensitive samples further divided into two subcategories. Given thresholds τ (FSS, default 0.15) and λ (LIS, default 0.50): **Trivial** samples ($LIS \leq \lambda$) are solvable from text alone and are removed first. Among the remaining vision-

dependent samples ($LIS > \lambda$), **Frame-Agnostic** samples ($|FSS| \leq \tau$) require visual input but are insensitive to which frames are shown, whereas **Temporally Sensitive** samples ($|FSS| > \tau$) are both vision-dependent and frame-sensitive. We further distinguish the latter by the sign of FSS: *Temporally Grounded* samples ($FSS > \tau$) benefit from relevant frames, while *Visual Bias* samples ($FSS < -\tau$) are harmed by them, indicating that visually matching frames can nevertheless be temporally misleading.

LIS-FSS orthogonality. A key design requirement is that LIS and FSS capture distinct properties. Spearman rank correlations between the two scores are near zero across all six benchmarks ($|\rho| < 0.10$), establishing that language independence and frame sensitivity are empirically orthogonal dimensions (see Figure 13).

7.2 Benchmark Quality Profiles

We summarize each benchmark with the **Temporal Purity Index (TPI)**:

$$TPI = \frac{|\text{Temporally Sensitive}|}{N - |\text{Trivial}|}, \quad (6)$$

where Temporally Sensitive = Key $\cup$ Anti-key (i.e., $|FSS| > \tau$ and $LIS > \lambda$), and Trivial = $\{i : LIS_i \leq \lambda\}$ encompasses all samples solvable from language priors. TPI thus measures the fraction of *vision-dependent* samples for which temporal frame selection is consequential.

Table 4 reports the category breakdown. Several patterns emerge: (i) **Temporally Sensitive**

Table 4: **At most a third of any benchmark is both vision-dependent and frame-sensitive.** Benchmark quality profiles under two-stage (LIS, FSS) classification. LIS is computed from no-frame (dummy visual input) predictions. *Trivial*: solvable from text alone ($\text{LIS} \leq \lambda$). *Fr.Agn*: vision-dependent but frame-insensitive. *Vis.Bias*: harmed by encoder-selected frames ($\text{FSS} < -\tau$). *Temp.Grnd*: benefits from relevant frames ($\text{FSS} > \tau$). $\text{TempSens} = \text{Vis.Bias} + \text{Temp.Grnd}$. $\text{TPI} = \text{TempSens} / (N - \text{Trivial})$.

Benchmark	N	Trivial (%)	Fr.Agn (%)	Vis.Bias (%)	Temp.Grnd (%)	TempSens (%)	TPI
<i>Short-video</i>							
MVBench	3786	33.4	61.0	1.2	4.3	5.5	0.083
NExTQA	8564	37.0	55.7	1.5	5.8	7.3	0.117
EgoSchema	500	20.2	60.6	2.6	16.6	19.2	0.241
<i>Long-video</i>							
MLVU	2173	35.6	33.4	5.8	25.2	31.0	0.481
LongVideoBench	1337	36.6	43.9	4.2	15.3	19.5	0.308
Video-MME	2700	31.3	43.0	3.9	21.8	25.6	0.373

fractions range from 5.5% (MVBench) to 31.0% (MLVU), indicating that even long-video benchmarks contain $< 32\%$ temporally sensitive, vision-dependent samples. (ii) **Frame-Agnostic** samples dominate MVBench (61.0%) and EgoSchema (60.6%), suggesting these benchmarks heavily test semantic scene understanding rather than temporal reasoning. (iii) **Trivial** fractions (20–37%) are substantial. (iv) The Temporal Purity Index ranges from 0.083 to 0.481, providing a single-number summary of each benchmark’s temporal grounding quality.

7.3 Validation

Fully validating temporal necessity is inherently difficult, as it is a latent property that cannot be observed directly from annotations alone. We therefore treat validation as converging evidence across complementary signals: human judgments, oracle-gap patterns, and internal and external correlations.

Anti-key temporal dependence. A crucial design choice in TEMPCORE is including Anti-key samples ($\text{FSS} < -\tau$) alongside Key samples as temporally sensitive. Figure 6 aggregates oracle gap across all six benchmarks by FSS bin, revealing a clear U-shape: both positive (Key) and negative (Anti-key) tails exhibit elevated oracle gaps compared to the near-zero center, confirming that Anti-key samples—like Key samples—benefit substantially from oracle frame selection.

Human evaluation corroborates this finding. Figure 7 shows that $|\text{FSS}|$ achieves a stronger correlation with human temporal scores than FSS ($\rho = +0.391$ vs. $+0.376$), and Anti-key samples receive higher temporal ratings than Non-key samples (3.55 vs. 2.98). Anti-key samples are visually




Figure 6: **Oracle gap exhibits a U-shape across FSS bins (cross-benchmark average).** Both Key and Anti-key tails show higher oracle gaps than the Non-key center, confirming that Anti-key samples are temporally sensitive. Shaded region: ± 1 std across benchmarks.




Figure 7: **FSS vs. human score.** (a) FSS shows positive correlation. (b) $|\text{FSS}|$ yields stronger correlation, capturing temporal dependence at both extremes.

misled but temporally dependent: SigLIP’s highest-confidence frames capture appearance rather than temporal content, and removing them paradoxically helps because the model falls back on more temporally informative frames. See additional analysis and details in Section B.

Internal consistency. Oracle gaps on Key samples exceed those on Non-key samples on all six benchmarks (mean $+23.7$ vs. $+15.3$; Table 2), and human annotators rate Key samples highest, followed by Anti-key and Non-key (see Figure 8), confirming that TEMPCORE retains those samples where temporal frame selection is consequential.

External correlation. Non-key model rankings correlate more strongly with six image VLM benchmarks than Key rankings (mean $\rho = 0.79$ vs. 0.59 across benchmarks; Section B.3), confirming that Key performance captures a distinct temporal skill rather than general VLM capability.

Practical usage. TEMPCORE subsets contain samples with $|\text{FSS}| > \tau$ and $\text{LIS} > \lambda$, comprising 5.5–31.0% of the original evaluation set. They can be used as a complementary metric: reporting TEMPCORE-subset accuracy alongside overall benchmark scores separates frame-sensitive from




Figure 8: **Human temporal scores by FSS category** ($\tau = 0.15$). Key samples receive the highest ratings; Anti-key exceeds Non-key, supporting their inclusion in Temporally Sensitive. Boxes: mean $\pm$ 1 SE; whiskers: $\pm$ 1 SD; centre line marks the mean.

frame-agnostic performance. The pruning removes the 69–94% of samples that are not frame-sensitive, yielding a focused evaluation of frame-selection sensitivity.

8 Conclusion

We presented three contributions toward understanding temporal grounding in Video QA evaluation. First, we proposed FSS, a per-sample diagnostic that quantifies dependence on temporal frame selection, and conducted a large-scale systematic study across multiple models, benchmarks, and strategies. Encoder-based frame selection yields substantial gains on long-video benchmarks, while oracle ceilings confirm large untapped headroom. Second, combining FSS with a LIS reveals that only a minority of benchmark samples are temporally sensitive—both vision-dependent and frame-sensitive. Third, we constructed TEMPCORE, temporally grounded evaluation subsets, providing a focused measure of frame-selection sensitivity.

Limitations

While FSS effectively identifies frame-sensitive samples, it captures only one facet of temporal reasoning—dependence on which frames are selected. Other temporal requirements, such as causal ordering, event duration estimation, or state-change detection, may not manifest as frame-selection sensitivity and are therefore invisible to FSS. Incorporating complementary temporal diagnostics (e.g., temporal ordering probes, action phase segmentation metrics) could yield a more comprehensive characterization of benchmark temporality. We leave the development of such multi-signal temporal taxonomies to future work.

We therefore do not claim that TEMPCORE measures temporal reasoning. A TEMPCORE subset isolates the samples whose answers depend on which frames are shown—a necessary condition for temporal reasoning (§3), but not a sufficient one.

Acknowledgments

This work was supported in part by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grants funded by the Korea government (MSIT) in 2026 (Grant No. RS-2026-25531289, Project Title: Development of Incrementally Adaptive On-device Learning Technology for New Environments, Contribution: 40%; Grant No. RS-2024-00457882, Project Title: AI Research Hub Project, Contribution: 10%; Grant No. RS-2019-II191906, Project Title: Artificial Intelligence Graduate School Program (POSTECH), Contribution: 10%) and in part by the National Research Foundation of Korea (NRF) grants funded by the Korea government (MSIT) in 2026 (Grant No. RS-2026-25494004, Project Title: Extremely Efficient Model Optimization for Large Multimodal Reasoning Foundation Models, Contribution: 20%; Grant No. RS-2024-00453301, Project Title: MASC: Convergence of Multimodal AI and Semantic Communications for Hyper-Connected Hyper-Intelligent Public Safety Systems, Contribution: 20%).

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, and 1 others. 2022. Flamingo: A visual language model for few-shot learning. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- 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, and 1 others. 2025a. Qwen3-VL technical report. *arXiv preprint arXiv:2511.21631*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shi-

- jie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, and 1 others. 2025b. Qwen2.5-VL technical report. *arXiv preprint arXiv:2502.13923*.
- Marija Brkic, Anas Filali Razzouki, Yannis Tevissen, Khalil Guetari, and Mounim A. El Yacoubi. 2025. Frame sampling strategies matter: A benchmark for small vision language models. *arXiv preprint arXiv:2509.14769*.
- Shyamal Buch, Arsha Nagrani, Anurag Arnab, and Cordelia Schmid. 2025. Flexible frame selection for efficient video reasoning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Mu Cai, Reuben Tan, Jianrui Zhang, Bocheng Zou, Kai Zhang, Yao Feng, Fangrui Zhu, Jing Gu, Yiwu Zhong, Yuzhang Shang, Yao Dou, Jaden Park, Jianfeng Gao, Yong Jae Lee, and Jianwei Yang. 2024. [Temporalbench: Benchmarking fine-grained temporal understanding for multimodal video models](#). In *Workshop on Video-Language Models @ NeurIPS 2024*.
- Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei, Zhifang Guo, Yichong Leng, Yuanjun Lv, Jinzheng He, Junyang Lin, Chang Zhou, and Jingren Zhou. 2024. Qwen2-Audio technical report. *arXiv preprint arXiv:2407.10759*.
- Daniel Cores, Michael Dorkenwald, Manuel Mucientes, Cees G. M. Snoek, and Yuki M. Asano. 2025. Lost in time: A new temporal benchmark for VideoLLMs. In *British Machine Vision Conference (BMVC)*.
- Wenzhang Du. 2026. When +1% is not enough: A paired bootstrap protocol for evaluating small improvements. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 519–531. Springer.
- Chaoyou Fu, Yuhan Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, and 1 others. 2025. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, and 1 others. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, and 1 others. 2024. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Vaishnavi Himakunthala, Andy Ouyang, Daniel Rose, Ryan He, Alex Mei, Yujie Lu, Chinmay Sonar, Michael Saxon, and William Yang Wang. 2023. Let’s think frame by frame with VIP: A video infilling and prediction dataset for evaluating video chain-of-thought. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Kai Hu, Feng Gao, Xiaohan Nie, Peng Zhou, Son Tran, Tal Neiman, Lingyun Wang, Mubarak Shah, Raffay Hamid, Bing Yin, and Trishul Chilimbi. 2025. M-LLM based video frame selection for efficient video understanding. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Minjoon Jung, Junbin Xiao, Byoung-Tak Zhang, and Angela Yao. 2025. On the consistency of video large language models in temporal comprehension. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Sutej Kulgod, Sean Ye, Sanchit Tanwar, and Christoffer Heckman. 2026. Reducing text bias in synthetically generated MCQAs for VLMs in autonomous driving. *arXiv preprint arXiv:2602.17677*.
- Jie Lei, Tamara Berg, and Mohit Bansal. 2023. Revealing single frame bias for video-and-language learning. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2025. [LLaVA-onevision: Easy visual task transfer](#). *Transactions on Machine Learning Research*.
- Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, Limin Wang, and Yu Qiao. 2024. MVBench: A comprehensive multi-modal video understanding benchmark. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Bin Lin, Yang Ye, Bin Zhu, Jiayi Cui, Munan Ning, Peng Jin, and Li Yuan. 2024. Video-LLaVA: Learning united visual representation by alignment before projection. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Yuanxin Liu, Shicheng Li, Yi Liu, Yuxiang Wang, Shuhuai Ren, Lei Li, Shihuo Chen, Xu Sun, and Lu Hou. 2024. TempCompass: Do video LLMs really understand videos? In *Findings of the Association for Computational Linguistics*.
- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. 2024. Video-ChatGPT: Towards detailed video understanding via large vision and language models. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.

- Karttikeya Mangalam, Raiymbek Akshulakov, and Jitendra Malik. 2023. EgoSchema: A diagnostic benchmark for very long-form video language understanding. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Andrés Marafioti, Orr Zohar, Miquel Farré, Merve noyan, Elie Bakouch, Pedro Manuel Cuenca Jiménez, Cyril Zakka, Loubna Ben allal, Anton Lozhkov, Nouamane Tazi, Vaibhav Srivastav, Joshua Lochner, Hugo Larcher, Mathieu Morlon, Lewis Tunstall, Leandro Von Werra, and Thomas Wolf. 2025. [SmolVLM: Redefining small and efficient multi-modal models](#). In *Second Conference on Language Modeling*.
- Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. 2021. Counterfactual VQA: A cause-effect look at language bias. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*.
- Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Xiaoqian Shen, Yongyang Xiong, Changsheng Zhao, Lemeng Wu, Jun Chen, Chenchen Zhu, Zechun Liu, Fanyi Xiao, Balakrishnan Varadarajan, Florian Bordes, Zhuang Liu, Hu Xu, Hyunwoo J. Kim, Bilge Soran, Raghuraman Krishnamoorthi, and 1 others. 2025. LongVU: Spatiotemporal adaptive compression for long video-language understanding. In *International Conference on Machine Learning (ICML)*.
- Hui Sun, Shiyin Lu, Huanyu Wang, Qing-Guo Chen, Zhao Xu, Weihua Luo, Kaifu Zhang, and Ming Li. 2025. MDP3: A training-free approach for list-wise frame selection in video-LLMs. In *IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Xi Tang, Jihao Qiu, Lingxi Xie, Yunjie Tian, Jianbin Jiao, and Qixiang Ye. 2025. Adaptive keyframe sampling for long video understanding. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, Olivier Hénaff, Jeremiah Harmsen, Andreas Steiner, and Xiaohua Zhai. 2025. SigLIP 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv preprint arXiv:2502.14786*.
- Ujjwal Upadhyay, Mukul Ranjan, Zhiqiang Shen, and Mohamed Elhoseiny. 2026. Time blindness: Why video-language models can’t see what humans can? In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, and 1 others. 2024. Qwen2-VL: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Ziyang Wang, Shoubin Yu, Elias Stengel-Eskin, Jaehong Yoon, Feng Cheng, Gedas Bertasius, and Mohit Bansal. 2025. VideoTree: Adaptive tree-based video representation for LLM reasoning on long videos. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Haoning Wu, Dongxu Li, Bei Chen, and Junnan Li. 2024. LongVideoBench: A benchmark for long-context interleaved video-language understanding. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. 2021. NExT-QA: Next phase of question-answering to explaining temporal actions. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, and 1 others. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.
- Linli Yao, Haoning Wu, Kun Ouyang, Yuanxing Zhang, Caiming Xiong, Bei Chen, Xu Sun, and Junnan Li. 2025. Generative frame sampler for long video understanding. In *Findings of the Association for Computational Linguistics*.
- Sicheng Yu, Chengkai Jin, Huanyu Wang, Zhenghao Chen, Sheng Jin, Zhongrong Zuo, Xiaolei Xu, Zhenbang Sun, Bingni Zhang, Jiawei Wu, Hao Zhang, and Qianru Sun. 2025. Frame-Voyager: Learning to query frames for video large language models. In *International Conference on Learning Representations (ICLR)*.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023. Sigmoid loss for language image pre-training. In *IEEE/CVF International Conference on Computer Vision (ICCV)*.
- 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. VideoLLaMA3: Frontier multimodal foundation models for image and video understanding. *arXiv preprint arXiv:2501.13106*.
- Hang Zhang, Xin Li, and Lidong Bing. 2023. Video-LLaMA: 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*.
- Kaichen Zhang, Bo Li, Peiyuan Zhang, Fanyi Pu, Joshua Adrian Cahyono, Kairui Hu, Shuai Liu, Yuanhan Zhang, Jingkang Yang, Chunyuan Li, and 1 others. 2025b. Lmms-eval: Reality check on the evaluation of large multimodal models. In *Findings of the Association for Computational Linguistics: NAACL 2025*.
- Junjie Zhou, Yan Shu, Bo Zhao, Boya Wu, Zhengyang Liang, Shitao Xiao, Minghao Qin, Xi Yang, Yongping Xiong, Bo Zhang, and 1 others. 2025. Mlvu: Benchmarking multi-task long video understanding. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
- Bin Zhu, Bin Lin, Munan Ning, Yang Yan, Jiaxi Cui, HongFa Wang, Yatian Pang, Wenhao Jiang, Junwu Zhang, Zongwei Li, Cai Zhang, Zhifeng Li, Wei Liu, and Li Yuan. 2024. LanguageBind: Extending video-language pretraining to n-modality by language-based semantic alignment. In *International Conference on Learning Representations (ICLR)*.
- Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, Zhangwei Gao, Erfei Cui, Xuehui Wang, Yue Cao, Yangzhou Liu, and 1 others. 2025. InternVL3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*.
- Orr Zohar, Xiaohan Wang, Yann Dubois, Nikhil Mehta, Tong Xiao, Philippe Hansen-Estruch, Licheng Yu, Xiaofang Wang, Felix Juefei-Xu, Ning Zhang, Serena Yeung-Levy, and Xide Xia. 2025. Apollo: An exploration of video understanding in large multimodal models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Model	Architecture	Params
LLaVA-OV-7B	LLaVA-OneVision (Li et al., 2025)	7B
LLaVA-OV-0.5B	LLaVA-OneVision (Li et al., 2025)	0.5B
VideoLLaMA3-2B	VideoLLaMA3 (Zhang et al., 2025a)	2B
SmolVLM2-2.2B	SmolVLM2 (Marafioti et al., 2025)	2.2B
SmolVLM2-500M	SmolVLM2 (Marafioti et al., 2025)	500M
SmolVLM2-256M	SmolVLM2 (Marafioti et al., 2025)	256M
Qwen3-VL-8B	Qwen3-VL (Bai et al., 2025a)	8B
InternVL3.5-8B	InternVL3 (Zhu et al., 2025)	8B

Table 5: **Eight open-source VLMs spanning three size tiers (256M–8B).** Five model families evaluated in zero-shot, multiple-choice setting with 32 input frames.

Appendix overview.

- Section A – Experimental setup
- Section B – FSS and TempCore validation
- Section C – Method ablations
- Section D – Detailed results

A Experimental Setup

A.1 Model Details

We evaluate eight open-source VLMs spanning five model families and three size tiers (256M–8B parameters). The selection covers both established architectures (LLaVA-OneVision, InternVL) and recent lightweight models (SmolVLM2, VideoLLaMA3), ensuring diversity in design choices such as vision encoder integration, temporal token handling, and training data composition. All models are evaluated in a zero-shot, multiple-choice setting with 32 input frames.

A.2 Benchmark Details

We select six Video QA benchmarks that span a range of video lengths, question types, and evaluation protocols. Short-video benchmarks (EgoSchema, MVBench, NExTQA) feature clips under three minutes and test fine-grained actions, scene transitions, and causal reasoning. Long-video benchmarks (MLVU, LongVideoBench, Video-MME) contain videos ranging from several minutes up to an hour and require multi-hop temporal reasoning over extended narratives. All benchmarks use multiple-choice format, enabling deterministic accuracy evaluation without LLM-based judging.

MVBench (Li et al., 2024) assembles twenty temporal task types by converting existing video annotations into multiple-choice items, deliberately targeting skills a single frame should not satisfy. **NExTQA** (Xiao et al., 2021) is organised around causal and temporal questions over daily-life clips,

Benchmark	# QA	Median Len.	Type
EgoSchema (Mangalam et al., 2023)	500	~3 min	Short
MVBench (Li et al., 2024)	4,000	<1 min	Short
NExTQA (Xiao et al., 2021)	8,564	~1 min	Short
MLVU (Zhou et al., 2025)	2,175	~8 min	Long
LongVideoBench (Wu et al., 2024)	1,337	~8 min	Long
Video-MME (Fu et al., 2025)	2,700	various	Mixed

Table 6: **Six Video QA benchmarks spanning seconds to tens of minutes.** “Length” is the median video duration; “# QA” is the full evaluation set size.

asking why an action occurred or what followed it. **EgoSchema** (Mangalam et al., 2023) extends clips to three minutes of egocentric footage and pairs them with narrative-rich options, so answering requires integrating a long span rather than recognising one moment. **MLVU** (Zhou et al., 2025) spans nine task types across movies, surveillance, egocentric video, cartoons and games, and is widely used to expose degradation as length grows. **LongVideoBench** (Wu et al., 2024) builds each question around a referring query pointing at a specific moment in videos up to an hour long, which is why its questions are markedly longer than the others’. **Video-MME** (Fu et al., 2025) is designed for full-spectrum coverage, sampling short, medium and long videos across six domains. The linguistic consequences of these choices are quantified in Table 13.

B FSS and TempCore Validation

B.1 Anti-key Samples as Temporally Sensitive

Our main analysis includes Anti-key samples in TEMP CORE alongside Key samples (§7). The following evidence motivates this design decision: *Anti-key samples are also temporally sensitive*. The key insight is that SigLIP’s max-prob frames can be *misleading*—selecting visually matching but temporally wrong content (e.g., for “Did the yellow object leave the scene?”, SigLIP selects frames *with* the object rather than frames showing its disappearance).

We hypothesize two mechanisms underlying Anti-key behavior: (i) for *static-appearance tasks* (e.g., Object Existence, Fine-grained Action), SigLIP preferentially selects close-up shots of salient objects, but the question asks about *absence* or subtle state changes that are better captured by contextual background frames; (ii) for *temporal-sequence tasks* (e.g., Action Sequence), the highest-similarity frames depict the endpoint




Figure 9: **Oracle gap exhibits a U-shape across FSS bins.** Both positive (Key) and negative (Anti-key) FSS tails show higher oracle gaps than the Non-key center, confirming that Anti-key samples are also temporally sensitive—they require specific temporal segments, but SigLIP selects misleading frames. Bins with fewer than 5 samples are omitted.

of an action, depriving the model of the preceding context needed to infer order. In MLVU, the Anti-key fraction (8.2%) is substantial despite MLVU having the highest Key fraction (36.0%), indicating that visual appearance bias affects even temporally demanding benchmarks.

Oracle gap vs. FSS bin. Figure 9 bins samples by quantized FSS values and plots mean oracle gap per bin. Both positive (Key) and negative (Anti-key) tails show elevated oracle gaps compared to the near-zero center (Non-key), producing a U-shaped curve. This confirms that *both* Key and Anti-key samples benefit substantially from oracle frame selection, while Non-key samples—which respond similarly regardless of frame selection strategy—show lower oracle gaps.

Oracle gap vs. $|FSS|$. Figure 10 plots per-sample oracle gap against $|FSS|$. The binned mean (red line) rises overall: higher $|FSS|$ correlates with larger oracle gaps regardless of sign, reinforcing that both Key and Anti-key samples are temporally localized.

Correct window centroid position. Table 7 analyzes where in the video the correct-answer windows are located. Key samples show a higher centroid standard deviation than Non-key on five of

six benchmarks, so their correct windows sit at varied points in the video rather than defaulting to the middle. Anti-key exceeds Non-key on three of six, a weaker and less consistent effect.

Together with the oracle gap subset analysis (Table 2), these results establish that Anti-key samples are temporally sensitive but misled by visually plausible-but-temporally-wrong frames. The true dichotomy is Temporal (Key + Anti-key) vs. Non-temporal (Non-key).

B.2 Human Temporal Judgment Validation

To validate FSS against human perception, we conducted a preliminary human evaluation study. An annotator assessed stratified samples spanning all six benchmarks and the three FSS categories (Key, Non-key, Anti-key at $\tau = 0.15$) on a 1–5 temporal necessity scale, passing where the necessity judgment felt unclear; the rated set contains $N = 100$ samples (32 Key, 48 Non-key, 20 Anti-key), with 9 further Non-key ratings excluded uniformly at random for balance. For each sample, annotators were shown the video, the question, and the answer options, and asked to assign a score indicating the extent to which temporal information was necessary to answer the question.




Figure 10: **Absolute FSS predicts oracle gap magnitude.** Per-sample scatter with binned means (red). The positive trend confirms that both high-FSS (Key) and low-FSS (Anti-key) samples benefit from temporal frame selection.

Metric correlations. Table 8 reports individual metric correlations with human temporal scores. Crucially, $|FSS|$ (absolute value) achieves higher Spearman ρ than signed FSS (+0.391 vs. +0.376), confirming that *both* Key and Anti-key samples exhibit temporal dependence.

FSS category analysis. Table 9 and Figure 8 group human temporal scores by FSS category. Key samples score highest (4.66 mean, 96.9% rated temporal), Anti-key follow (3.55, 60.0%) and Non-key are lowest (2.98, 47.9%). The Temporally Sensitive class ($\text{Key} \cup \text{Anti-key}$) exceeds Non-key ($p < 0.001$, one-sided Mann–Whitney); Anti-key alone is directionally higher but not significant at this sample size ($p = 0.18$).

τ -sweep against human ground truth. Table 10 compares two temporal classification methods at varying FSS thresholds, using human scores ≥ 4 as ground truth. Method A (Key-only: $FSS > \tau$) achieves higher precision, while Method B (Key+Anti-key: $|FSS| > \tau$) gains in recall. Method B matches or exceeds Method A at every τ , confirming the inclusion of Anti-key in TEMPCORE. We adopt Method B ($|FSS| > \tau$) as the default classification throughout this work.

B.3 External Correlation with Image VLM Benchmarks

To validate that FSS correctly separates temporal from non-temporal performance, we correlate Key-subset and Non-key-subset model rankings with six image VLM benchmarks (MMM U, MMM U-Pro, MME, AI2D, TextVQA, OCRBench) evaluated using lmms-eval (Zhang et al., 2025b). If FSS correctly identifies frame-sensitive samples, Non-key accuracy should correlate more strongly with image benchmarks (measuring general VLM capability), while Key accuracy should diverge.

Table 11 shows the expected direction on every benchmark: Key correlation never exceeds Non-key. The gap is widest on MLVU (mean $\rho = 0.33$ against 0.78) and Video-MME (0.44 against 0.78), the two benchmarks with the highest temporal purity (Table 4), and narrows to +0.26 on MVBench and +0.10 on EgoSchema and LongVideoBench.

NExTQA is the exception: its Key and Non-key rankings coincide exactly ($\rho = 0.78$ for both), so the split carries no information about model ranking there. With only eight models a Spearman correlation is coarse, and NExTQA’s Key set is small (7.7%); we read the pattern as holding on the benchmarks with substantial temporal content rather than universally. This provides external evidence that FSS captures distinct, frame-sensitive capabilities rather than general VLM strength.

Benchmark	Category	Begin (%)	Middle (%)	End (%)	Std
<i>Short-video</i> MVBench	Key	10.2	70.5	19.3	0.160
	Non-key	5.3	85.8	8.9	0.119
	Anti-key	0.0	92.7	7.3	0.090
NExTQA	Key	7.0	83.0	10.0	0.127
	Non-key	2.6	93.9	3.5	0.083
	Anti-key	4.2	90.0	5.8	0.097
EgoSchema	Key	10.6	76.6	12.8	0.135
	Non-key	4.5	91.8	3.7	0.094
	Anti-key	6.2	68.8	25.0	0.107
<i>Long-video</i> MLVU	Key	7.3	87.1	5.6	0.110
	Non-key	2.4	95.9	1.7	0.078
	Anti-key	1.7	92.7	5.6	0.096
LongVideoBench	Key	5.0	90.3	4.7	0.105
	Non-key	4.9	90.4	4.7	0.108
	Anti-key	3.8	91.2	5.0	0.099
Video-MME	Key	7.7	86.0	6.2	0.112
	Non-key	4.0	92.3	3.8	0.093
	Anti-key	2.9	93.6	3.5	0.085

Table 7: **Key samples’ correct windows are spread across the video; Anti-key’s are not consistently so.** Distribution of correct-window centroids across video thirds (Begin: $[0, 1/3]$, Middle: $[1/3, 2/3]$, End: $[2/3, 1]$). Std is the standard deviation of centroid position, so a higher value means correct windows sit at varied points in the video rather than defaulting to the middle. Key exceeds Non-key on five of six benchmarks (LongVideoBench is the exception); Anti-key exceeds it on three.

Metric	Spearman ρ	PB r	AUC
FSS	+0.376	+0.327	0.685
FSS	+0.391	+0.345	0.726

Table 8: **Individual metric correlations with human temporal scores.** |FSS| captures both Key and Anti-key temporal dependence, yielding higher correlation than signed FSS. Threshold for binarization: score ≥ 4 . (PB r : Point-Biserial correlation.)

C Method Ablations

C.1 Frame Count (K) Ablation

We ablate the number of selected frames $K \in \{16, 32, 48\}$ on EgoSchema with LLaVA-OV-7B and LLaVA-OV-0.5B (Table 12). Accuracy differences across K values are modest ($\leq 5\%$), suggesting that MAXPROB is robust to the frame budget. We use $K = 32$ throughout as a practical default that balances temporal coverage with computational cost.

C.2 Benchmark Question Profiles

Five-axis profile of the six benchmarks, referenced in Section 4. Descriptive density, WH-type entropy

FSS Bin	N	Mean	Temporal%	High Conf.%
Key	32	4.66	96.9%	100%
Non-key	48	2.98	47.9%	100%
Anti-key	20	3.55	60.0%	100%

Table 9: **Human temporal scores by FSS category** ($\tau = 0.15$). Key and Anti-key samples both exhibit elevated temporal scores compared to Non-key, supporting the U-shape hypothesis. Temporal%: score ≥ 4 .

τ	Key-Only (FSS $> \tau$)			Key+Anti-Key (FSS $> \tau$)		
	Prec.	Rec.	F1	Prec.	Rec.	F1
0.050	0.822	0.561	0.667	0.753	0.833	0.791
0.100	0.822	0.561	0.667	0.753	0.833	0.791
0.125	0.969	0.470	0.633	0.827	0.652	0.729
0.150	0.969	0.470	0.633	0.827	0.652	0.729
0.200	0.969	0.470	0.633	0.827	0.652	0.729
0.250	1.000	0.303	0.465	0.852	0.348	0.495
0.375	1.000	0.212	0.350	0.889	0.242	0.381
0.500	1.000	0.136	0.240	0.909	0.152	0.260

Table 10: τ -sweep: **Key-only vs. Key+Anti-key temporal classification against human ground truth.** Including Anti-key samples (right) improves recall and F1 across most τ values, confirming that anti-key frames also indicate temporal dependence.

and type-token ratio are computed over the full question set; ambiguity is a mean rating on a five-point scale.

C.3 Frame Coverage and Sensitivity

Table 14 reports mean |FSS| binned by frame coverage K/L , the ratio of the frame budget to the number of scored frames. Only MVBench reaches the highest coverage bin; EgoSchema and MLVU sit entirely below $K/L = 0.10$. “–” marks an empty bin.

C.4 Contamination Probe

Two controls test whether no-frame (LIS) accuracy reflects benchmark contamination acquired during video-multimodal training. First, we extend the LIS model panel with three text-only LLMs, for which video-specific contamination is impossible by construction (Table 15); benchmark-level LIS moves by at most 2.9 pp and by under 1 pp on three of six. These probes cover a 300-sample subset per benchmark, so this is a robustness check on the aggregate, and per-sample LIS throughout the paper remains the 8-VLM quantity.

Second, a same-backbone paired control (Table 16) compares Qwen3-8B against Qwen3-VL-

Split	Spearman ρ vs. Image VLM Benchmark						Mean
	MMMU	MMMU-Pro	MME	AI2D	TextVQA	OCRBench	
<i>Short-video</i>							
<i>MVBench</i>							
Key	.64	.64	.62	.55	.57	.31	.56
Non-key	.86	.86	.90	.83	.81	.67	.82
<i>NExTQA</i>							
Key	.83	.83	.86	.79	.76	.60	.78
Non-key	.83	.83	.86	.79	.76	.60	.78
<i>EgoSchema</i>							
Key	.71	.71	.74	.67	.69	.45	.66
Non-key	.81	.81	.83	.76	.71	.64	.76
<i>Long-video</i>							
<i>MLVU</i>							
Key	.43	.43	.40	.31	.33	.10	.33
Non-key	.83	.83	.86	.79	.76	.60	.78
<i>LongVideoBench</i>							
Key	.79	.79	.83	.76	.74	.55	.74
Non-key	.90	.90	.90	.86	.83	.64	.84
<i>Video-MME</i>							
Key	.52	.52	.57	.45	.43	.17	.44
Non-key	.83	.83	.86	.79	.76	.60	.78

Table 11: **Non-key accuracy correlates more strongly with image VLM benchmarks than Key accuracy.** Spearman ρ between Key and Non-key (including Anti-key) model rankings and six image benchmarks evaluated via lmms-eval. Key correlation is never higher than Non-key, and the gap is widest on MLVU and Video-MME, indicating that Key performance captures a distinct temporal skill rather than general VLM capability. $N = 8$ models.

Model	$K=16$	$K=32$	$K=48$
LLaVA-OV-7B	50.6	53.8	55.0
LLaVA-OV-0.5B	23.8	24.2	25.0

Table 12: **MaxProb accuracy is robust to frame budget K .** Q-mode accuracy (%) on EgoSchema across $K \in \{16, 32, 48\}$ under the temporal-diversity selection of §3 for two representative models. Differences are modest, supporting $K=32$ as a practical default. **Bold** marks the default $K=32$.

8B-Instruct, which adds video-multimodal post-training on top of it. Both see the same questions with no usable visual input, so the difference bounds what that post-training contributes, contamination included. On MVBench, NExTQA, EgoSchema and MLVU it is +1.0 to +2.3 pp. On LongVideoBench (+13.3) and Video-MME (+8.7) it is not small, and our probes cannot separate benchmark exposure from ordinary gains in multiple-choice instruction following; the Trivial fractions of those two benchmarks should be read with that reservation.

C.5 Random and Uniform Baselines

C.6 Per-Encoder Comparison

We compare MAXPROB accuracy with three vision-language encoders: SigLIP2-400M (default), SigLIP-400M, and CLIP ViT-L/14. Table 18 reports Q-mode accuracy on the Perception benchmark (19,140 samples) for two LLaVA models. All three encoders produce nearly identical accuracy, with differences within 0.5 pp. This confirms that frame selection quality is robust to the choice of scoring encoder, consistent with the finding of Zohar et al. (2025) that modern vision-language encoders converge in frame-level ranking ability.

C.7 Question and Answer Text Statistics

Table 19 reports the mean question and answer option lengths across six benchmarks. LongVideoBench questions average 44.1 words, up to nearly four times longer than those of other benchmarks (12–25 words), because they embed detailed scene descriptions—clothing, spatial layout, character actions—before posing a query. This explains why QA-mode adds least on LongVideoBench among the long-video benchmarks (§5.1): the question text alone provides sufficient visual grounding for the encoder, and appending answer text adds noise rather than signal.

LongVideoBench case study. The following examples illustrate the descriptive nature of LongVideoBench questions:

LongVideoBench (77 words): “In a room decorated with various paintings, there are two men standing in the middle. On the left is a man with short hair wearing a red short-sleeve shirt, holding a beer. On the right is a man with a black beard, wearing a gray short-sleeve shirt. Behind the man in red is a white door. When the subtitle mentions ‘over 2600 meters yeah there you go we’, what did the man in the red shirt do?”

LongVideoBench (79 words): “In the video, there are two men in a room. The man on the left is wearing a yellow shirt, a hat backwards, and a watch on his left hand. The man on the right is wearing a shirt with red flowers and green leaves and has a bracelet on his left hand. When the subtitle mentions ‘interested enough to join us but first,’ what change happens to the man wearing the shirt with red flowers and green leaves?”

MLVU (4 words): “Did I drink water?”

The contrast is striking: LongVideoBench questions specify clothing colors, accessories, and spatial positions in enough detail for a vision-language encoder to localize the relevant frames from the

Benchmark	N	Q-words	Desc.%	WH-ent	TTR	Ambig.
MVBench	4,000	13.1	1.81	2.04	0.038	1.63
NExTQA	8,564	11.5	3.93	1.84	0.028	1.61
EgoSchema	500	24.5	0.14	1.24	0.114	2.71
MLVU	2,174	15.4	1.97	2.32	0.071	1.29
LongVideoBench	1,337	44.1	8.47	0.71	0.104	1.33
Video-MME	2,700	13.6	1.25	2.27	0.151	1.47

Table 13: **The six benchmarks differ on every axis we can measure, so median length alone does not characterise them.** “N”: evaluation-set size; “Q-words”: mean question length. “Desc.%”: share of question tokens that are visual descriptors (colour, shape, object). “WH-ent”: Shannon entropy in bits over 11 WH-type bins, so higher means WH-types are evenly mixed and lower means one type dominates. “TTR”: type-token ratio over question text, a lexical-diversity index. “Ambig.”: mean of two independent LLM 1–5 under-specification scores (1 = clear). Column maxima in bold.

Benchmark	≥ 0.5	0.25–0.5	0.10–0.25	< 0.10
MVBench	4.5	7.2	7.0	9.5
NExTQA	–	5.6	6.0	8.3
EgoSchema	–	–	–	11.9
MLVU	–	–	–	20.9
LongVideoBench	–	6.1	6.6	15.8
Video-MME	–	2.1	7.8	18.4

Table 14: **The scarcest frame budget yields the largest |FSS| in every benchmark.** Cell: 8-model average |FSS| (%) for samples binned by K/L, the ratio of the K=32 frame budget to the number of scored frames L. The < 0.10 bin holds the row maximum everywhere, exceeding the best-covered bin by 2.7–16.3 pp in the four benchmarks that populate more than one bin. Frame selection therefore matters in proportion to how much of the video the budget must discard. “–” indicates no samples in that bin.

question alone. In contrast, short questions like the MLVU example provide almost no visual cues, making the answer text a valuable source of grounding signal.

D Detailed Results

D.1 Per-Benchmark Figures

Figure 11 presents the per-benchmark FSS distributions corresponding to the aggregated view in Figure 4. Figure 12 shows the cross-benchmark model ranking correlation matrices; Figure 13 shows the per-benchmark LIS–FSS scatter plots.

D.2 Performance Tables

Table 20–Table 28 provide complete per-model results for each evaluation strategy and benchmark.

Benchmark	8 VLMs	+ 3 LLMs	Δ (pp)
MVBench	29.1	30.8	+1.6
NExTQA	34.2	37.1	+2.9
EgoSchema	25.7	27.5	+1.8
MLVU	36.1	36.3	+0.2
LongVideoBench	35.3	35.1	–0.2
Video-MME	33.1	33.2	+0.0

Table 15: **Adding models that have never seen a video barely moves LIS.** Mean no-frame accuracy (%) over the 8-VLM panel versus the same panel plus 3 text-only LLMs (Qwen2.5-7B, Llama-3.1-8B, Qwen3-8B), for which video-specific contamination is impossible by construction. The largest shift is +2.9 pp and three of six benchmarks move by under 1 pp, so the language-prior signal LIS measures is not an artefact of video-multimodal training. The LLM probes cover a 300-sample subset per benchmark, so this is a benchmark-level robustness check; per-sample LIS throughout the paper remains the 8-VLM quantity.

Benchmark	N	Qwen3-8B (text-only)	Qwen3-VL-8B (no-frame)	Δ
MVBench	300	37.3	39.3	+2.0
NExTQA	300	46.7	49.0	+2.3
EgoSchema	300	33.7	35.7	+2.0
MLVU	300	38.7	39.7	+1.0
LongVideoBench	300	32.0	45.3	+13.3
Video-MME	300	35.3	44.0	+8.7
Mean				+4.9

Table 16: **Video post-training leaves no-frame accuracy nearly unchanged on four benchmarks but adds substantially on two.** Qwen3-8B is a text-only LLM; Qwen3-VL-8B-Instruct shares that backbone and adds video-multimodal post-training. Both answer the same questions with no usable visual input, so Δ upper-bounds what that post-training contributes—including any benchmark exposure it carried, but also ordinary gains in multiple-choice instruction following. Columns are computed on identical samples; every probe item was checked to carry the dataset’s own gold letter.




Figure 11: **Per-benchmark FSS distributions** ($\tau = 0.15$). All six benchmarks show a dominant near-zero peak, with MLVU exhibiting the broadest spread.




Figure 12: **Cross-benchmark model ranking correlations.** 1×4 Spearman ρ matrices for Key×Key, Non-key×Non-key, Key×Non-key, and Non-key×Key subsets.

Benchmark	MinProb	Random	Uniform	MaxProb	Max – Random
MVBench	47.4	48.7	48.7	49.0	+0.3
NExTQA	56.8	58.4	59.4	59.2	+0.8
EgoSchema	32.7	37.8	38.1	39.2	+1.4
MLVU	44.1	50.9	51.5	57.3	+6.4
LongVideoBench	41.4	44.3	45.2	47.8	+3.5
Video-MME	38.6	46.1	46.3	49.7	+3.7

Table 17: **Random sampling baseline confirms the encoder-based selection signal.** Mean accuracy across 8 VLMs (%), QA-mode scoring with $K = 32$ frames, on the samples with complete results under all four strategies. MaxProb exceeds Random on all 6 benchmarks (largest margin +6.4 pp on MLVU) and Random exceeds MinProb on 6 of 6, so the gap reflects ranking quality rather than a sampling-protocol artifact. “Uniform” samples K frames at equal intervals; it tracks Random closely and edges MaxProb only on NExTQA (59.4 vs. 59.2).

Model	SigLIP2-400M	SigLIP-400M	CLIP ViT-L
LLaVA-OV-7B	57.5	57.5	57.9
LLaVA-OV-0.5B	48.8	48.8	48.9
Δ (max – min)		≤ 0.5 pp	

Table 18: **Encoder choice has negligible impact on MaxProb accuracy.** Q-mode MaxProb accuracy (%) on Perception (19,140 samples) with three vision-language encoders. All differences are within 0.5 pp, confirming that frame selection quality is robust to encoder choice.




Figure 13: **Per-benchmark LIS–FSS scatter plots.** Each panel shows the two-dimensional (LIS, FSS) distribution with four-category coloring. Dashed lines mark thresholds ($\tau = 0.15$, $\lambda = 0.50$).

Benchmark	Q (words)	A (words)	#Opts
MVBench	13.1	4.0	3.7
NExTQA	11.5	2.8	5.0
EgoSchema	24.5	22.5	5.0
MLVU	15.4	2.5	4.0
LongVideoBench	44.1	9.0	4.7
Video-MME	13.6	5.1	4.0

Table 19: **LongVideoBench questions are up to nearly $4\times$ longer than other benchmarks’.** Mean question length (Q), answer option length (A), and number of options (#Opts) across six benchmarks. The descriptive nature of LongVideoBench questions explains why QA-mode adds least there among the long-video benchmarks.

Table 20: **QA-mode scoring wins in 42 of 48 model×benchmark cells.** Per-model Q-mode vs. QA-mode MAXPROB accuracy (%) across all six benchmarks; Table 1 reports the model-averaged summary. Five of the six exceptions are VideoLLaMA3-2B, and none exceeds 1.4 pp, so the advantage is broad rather than driven by a few models.

Model	Short-video						Long-video					
	MVBench		NExTQA		EgoSchema		MLVU		LongVideoBench		Video-MME	
	Q	QA	Q	QA	Q	QA	Q	QA	Q	QA	Q	QA
InternVL3.5-8B	68.7	69.5	79.5	80.7	55.2	66.2	68.7	72.6	59.2	61.1	59.3	62.8
LLaVA-OV-7B	55.8	57.0	78.0	78.7	53.8	62.8	67.6	70.8	58.3	59.3	56.8	62.2
Qwen3-VL-8B	56.4	57.6	70.9	71.2	50.6	58.2	57.4	60.1	50.2	52.4	50.5	54.1
SmolVLM2-2.2B	46.5	47.4	60.8	62.3	25.8	32.0	58.0	59.9	50.2	50.6	47.5	53.2
LLaVA-OV-0.5B	46.3	47.4	57.2	58.5	24.2	27.4	57.7	59.9	45.3	46.7	44.5	49.3
SmolVLM2-500M	42.0	42.3	46.9	48.3	23.4	26.6	51.4	55.2	42.5	43.8	40.5	45.7
VidLLaMA3-2B	40.2	38.8	46.9	46.4	23.2	23.0	42.9	42.8	41.1	40.9	36.7	36.8
SmolVLM2-256M	31.6	31.9	27.6	27.4	17.0	17.2	35.6	37.4	26.8	27.7	32.5	33.7

Model	Short-video									Long-video									Avg. Δ	
	MVBench			NExTQA			EgoSchema			MLVU			LongVideoBench			Video-MME			Short	Long
	Max	Min	Δ	Max	Min	Δ	Max	Min	Δ	Max	Min	Δ	Max	Min	Δ	Max	Min	Δ		
InternVL3.5-8B	69.5	67.2	+2.3	80.7	77.0	+3.6	66.2	53.2	+13.0	72.6	53.4	+19.2	61.1	50.3	+10.8	62.8	47.9	+15.0	+6.3	+15.0
LLaVA-OV-7B	57.0	54.8	+2.2	78.7	75.0	+3.7	62.8	49.0	+13.8	70.8	49.9	+20.9	59.3	47.9	+11.4	62.2	43.4	+18.8	+6.6	+17.0
Qwen3-VL-8B	57.6	56.2	+1.4	71.2	68.9	+2.3	58.2	50.2	+8.0	60.1	48.6	+11.5	52.4	46.1	+6.4	54.1	44.5	+9.6	+3.9	+9.1
SmolVLM2-2.2B	47.4	45.5	+2.0	62.3	58.8	+3.5	32.0	24.8	+7.2	59.9	45.5	+14.4	50.6	43.0	+7.6	53.2	39.1	+14.1	+4.2	+12.0
LLaVA-OV-0.5B	47.4	44.4	+3.0	58.5	54.7	+3.8	27.4	21.4	+6.0	59.9	43.4	+16.6	46.7	38.6	+8.1	49.3	34.7	+14.6	+4.2	+13.1
SmolVLM2-500M	42.3	40.7	+1.6	48.3	45.6	+2.6	26.6	23.2	+3.4	55.2	41.0	+14.2	43.8	38.1	+5.7	45.7	33.9	+11.8	+2.5	+10.6
VidLLaMA3-2B	38.8	39.1	-0.2	46.4	46.8	-0.4	23.0	23.4	-0.4	42.8	41.8	+1.0	40.9	40.5	+0.4	36.8	37.4	-0.6	-0.4	+0.3
SmolVLM2-256M	31.9	31.4	+0.5	27.4	27.1	+0.3	17.2	16.2	+1.0	37.4	29.6	+7.8	27.7	27.1	+0.7	33.7	27.7	+6.0	+0.6	+4.8
Average	49.0	47.4	+1.6	59.2	56.8	+2.4	39.2	32.7	+6.5	57.3	44.1	+13.2	47.8	41.4	+6.4	49.7	38.6	+11.2	+3.5	+10.2

Table 21: **MaxProb outperforms MinProb on long-video benchmarks in 23 of 24 model–benchmark cells.** Per-model accuracy (%) with Short and Long average Δ columns. Best value per column in bold.

Model	Short-video									Long-video									Avg. Gap	
	MVBench			NExTQA			EgoSchema			MLVU			LongVideoBench			Video-MME			Short	Long
	Uniform	Oracle	Gap	Uniform	Oracle	Gap	Uniform	Oracle	Gap	Uniform	Oracle	Gap	Uniform	Oracle	Gap	Uniform	Oracle	Gap		
InternVL3.5-8B	69.1	76.7	+7.5	80.2	87.7	+7.4	64.6	83.8	+19.2	65.6	93.7	+28.1	57.3	83.9	+26.6	61.3	85.2	+24.0	+11.4	+26.2
LLaVA-OV-7B	57.0	65.1	+8.1	78.1	88.2	+10.0	63.0	83.4	+20.4	63.9	92.7	+28.8	56.5	81.8	+25.3	57.7	87.9	+30.1	+12.8	+28.1
Qwen3-VL-8B	57.2	66.4	+9.2	70.5	84.4	+13.9	54.0	81.0	+27.0	52.0	87.3	+35.3	48.8	80.9	+32.1	49.7	85.6	+36.0	+16.7	+34.5
LLaVA-OV-0.5B	46.4	52.1	+5.7	57.0	69.9	+13.0	26.4	40.2	+13.8	51.3	82.8	+31.5	43.0	63.8	+20.8	44.3	75.0	+30.7	+10.8	+27.7
SmolVLM2-2.2B	46.7	52.1	+5.4	61.1	70.4	+9.3	29.4	41.8	+12.4	54.6	77.0	+22.5	47.3	67.8	+20.6	49.4	73.8	+24.4	+9.1	+22.5
SmolVLM2-500M	42.1	47.7	+5.6	47.2	57.4	+10.3	26.2	44.4	+18.2	48.6	78.5	+30.0	40.9	69.2	+28.3	40.3	75.1	+34.8	+11.4	+31.0
SmolVLM2-256M	31.5	33.6	+2.1	27.9	37.2	+9.3	17.0	22.2	+5.2	33.3	60.4	+27.1	26.8	48.9	+22.1	30.8	67.5	+36.7	+5.5	+28.6
VidLLaMA3-2B	39.5	41.0	+1.5	47.0	48.9	+1.9	24.4	25.8	+1.4	42.9	46.1	+3.2	41.1	43.5	+2.4	37.1	41.1	+3.9	+1.6	+3.2
Average	48.7	54.3	+5.6	58.6	68.0	+9.4	38.1	52.8	+14.7	51.5	77.3	+25.8	45.2	67.5	+22.3	46.3	73.9	+27.6	+9.9	+25.2

Table 22: **Large oracle gaps reveal substantial headroom for frame selection.** Uniform-sampling vs. window-oracle accuracy (%). Long-video gaps far exceed short-video gaps, confirming that temporal frame selection is most consequential for extended videos. Best value per column in bold.

Model	No Frame	Uniform	Oracle	MaxProb	MinProb
InternVL3.5-8B	38.9	69.1	76.7	69.5	67.2
Qwen3-VL-8B	37.5	57.2	66.4	57.6	56.2
LLaVA-OV-7B	36.5	57.0	65.1	57.0	54.8
LLaVA-OV-0.5B	33.0	46.4	52.1	47.4	44.4
SmolVLM2-2.2B	30.2	46.7	52.1	47.4	45.5
SmolVLM2-500M	31.2	42.1	47.7	42.3	40.7
VidLLaMA3-2B	31.9	39.5	41.0	38.8	39.1
SmolVLM2-256M	29.3	31.5	33.6	31.9	31.4

Table 23: **Window oracle dominates all strategies on MVBench.** Cross-strategy accuracy (%) across 8 VLMs. Oracle substantially outperforms all other strategies, while MaxProb and uniform sampling perform comparably on this short-video benchmark. Best value per column in bold.

Model	No Frame	Uniform	Oracle	MaxProb	MinProb
InternVL3.5-8B	48.8	80.2	87.7	80.7	77.0
LLaVA-OV-7B	49.2	78.1	88.2	78.7	75.0
Qwen3-VL-8B	49.8	70.5	84.4	71.2	68.9
LLaVA-OV-0.5B	35.7	57.0	69.9	58.5	54.7
SmolVLM2-2.2B	22.6	61.1	70.4	62.3	58.8
SmolVLM2-500M	25.2	47.2	57.4	48.3	45.6
VidLLaMA3-2B	30.0	47.0	48.9	46.4	46.8
SmolVLM2-256M	21.1	27.9	37.2	27.4	27.1

Table 24: **Window-oracle reveals large untapped potential on NExtQA.** Cross-strategy accuracy (%) across 8 VLMs. MaxProb and uniform sampling remain closely matched. Best value per column in bold.

Model	No Frame	Uniform	Oracle	MaxProb	MinProb
InternVL3.5-8B	33.0	64.6	83.8	66.2	53.2
LLaVA-OV-7B	34.6	63.0	83.4	62.8	49.0
Qwen3-VL-8B	36.6	54.0	81.0	58.2	50.2
SmolVLM2-2.2B	20.2	29.4	41.8	32.0	24.8
SmolVLM2-500M	19.4	26.2	44.4	26.6	23.2
LLaVA-OV-0.5B	15.6	26.4	40.2	27.4	21.4
VidLLaMA3-2B	25.6	24.4	25.8	23.0	23.4
SmolVLM2-256M	20.2	17.0	22.2	17.2	16.2

Table 25: **MAXPROB leads MINPROB on EgoSchema for 7 of 8 models.** Cross-strategy accuracy (%) across the eight VLMs. The mean gap is +6.5 pp, between the short-video and long-video groups, matching EgoSchema’s three-minute clips. Best value per column in bold.

Model	No Frame	Uniform	Oracle	MaxProb	MinProb
InternVL3.5-8B	44.9	65.6	93.7	72.6	53.4
LLaVA-OV-7B	45.1	63.9	92.7	70.8	49.9
Qwen3-VL-8B	41.4	52.0	87.3	60.1	48.6
LLaVA-OV-0.5B	35.8	51.3	82.8	59.9	43.4
SmolVLM2-2.2B	30.2	54.6	77.0	59.9	45.5
SmolVLM2-500M	31.1	48.6	78.5	55.2	41.0
VidLLaMA3-2B	34.2	42.9	46.1	42.8	41.8
SmolVLM2-256M	25.7	33.3	60.4	37.4	29.6

Table 26: **MLVU shows the strongest strategy differentiation among all benchmarks.** Cross-strategy accuracy (%) across 8 VLMs, reflecting MLVU’s high temporal demand. Best value per column in bold.

Model	No Frame	Uniform	Oracle	MaxProb	MinProb
InternVL3.5-8B	44.6	57.3	83.9	61.1	50.3
LLaVA-OV-7B	39.7	56.5	81.8	59.3	47.9
Qwen3-VL-8B	42.8	48.8	80.9	52.4	46.1
SmolVLM2-2.2B	29.2	47.3	67.8	50.6	43.0
LLaVA-OV-0.5B	33.7	43.0	63.8	46.7	38.6
SmolVLM2-500M	31.0	40.9	69.2	43.8	38.1
VidLLaMA3-2B	36.5	41.1	43.5	40.9	40.5
SmolVLM2-256M	25.1	26.8	48.9	27.7	27.1

Table 27: **MaxProb provides meaningful gains over MinProb on LongVideoBench.** Cross-strategy accuracy (%) across 8 VLMs, demonstrating effective frame retrieval on long videos. Best value per column in bold.

Model	No Frame	Uniform	Oracle	MaxProb	MinProb
InternVL3.5-8B	44.5	61.3	85.2	62.8	47.9
LLaVA-OV-7B	38.2	57.7	87.9	62.2	43.4
Qwen3-VL-8B	40.6	49.7	85.6	54.1	44.5
SmolVLM2-2.2B	28.1	49.4	73.8	53.2	39.1
LLaVA-OV-0.5B	30.7	44.3	75.0	49.3	34.7
SmolVLM2-500M	29.5	40.3	75.1	45.7	33.9
SmolVLM2-256M	25.4	30.8	67.5	33.7	27.7
VidLLaMA3-2B	28.2	37.1	41.1	36.8	37.4

Table 28: **Video-MME exhibits the largest oracle gaps among all benchmarks.** Cross-strategy accuracy (%) across 8 VLMs. Best value per column in bold.