

VideoCuRL: Video Curriculum Reinforcement Learning with Orthogonal Difficulty Decomposition

Hongbo Jin Kuanwei Lin Wenhao Zhang Yichen Jin Ge Li*

School of Electronic and Computer Engineering,

Peking University,

Correspondence: hbjin25@stu.pku.edu.cn

Abstract

Reinforcement Learning (RL) is crucial for empowering Video-LLMs with complex spatiotemporal reasoning. However, current RL paradigms predominantly rely on random data shuffling or naive curriculum strategies based on scalar difficulty metrics. We argue that scalar metrics fail to disentangle two orthogonal challenges in video understanding: Visual-Temporal Perception Load and Cognitive Reasoning Depth. To address this, we propose VideoCuRL, a novel framework that decomposes difficulty into these two axes. We employ efficient, training-free proxies—optical flow/keyframe entropy for visual complexity and Calibrated Surprisal for cognitive complexity—to map data onto a 2D curriculum grid. A competence-aware Diagonal Wavefront strategy then schedules training from base alignment to complex reasoning. Furthermore, we introduce Dynamic Sparse KL and Structured Revisiting to stabilize training against reward collapse and catastrophic forgetting. Extensive experiments show that VideoCuRL surpasses strong RL baselines on reasoning (+2.5% on VSI-Bench) and perception (+2.9% on VideoMME) tasks. Notably, VideoCuRL eliminates the prohibitive inference overhead of generation-based curricula, offering a scalable solution for robust video post-training.

1 Introduction

Large Multimodal Models (LMMs) have catalyzed significant advancements in video understanding, enabling systems capable of following complex instructions, aligning vision and language, and reasoning over extended temporal contexts (Li et al., 2025a, 2024a; Bai et al., 2025). By coupling powerful visual encoders with Large Language Models (LLMs), modern Video-LLMs achieve remarkable performance on benchmarks (Xu et al., 2017; Yu

et al., 2019; Lei et al., 2018; Jang et al., 2017) spanning action recognition, temporal grounding, and video question answering. Nevertheless, robust spatiotemporal video reasoning remains an open challenge. Current models often falter on tasks that demand long-range temporal integration, causal inference, or excessive perception load.

Reinforcement Learning (RL) (Schulman et al., 2017; Ouyang et al., 2022; Shao et al., 2024) has emerged as a promising post-training paradigm to address this gap. In contrast to Supervised Fine-Tuning (SFT), RL incentivizes exploration and has demonstrated clear gains on complex video reasoning tasks (Feng et al., 2025). However, current video reinforcement learning methods lack sophisticated curriculum design and mostly rely on random data shuffling. This design implicitly assumes that all samples are equally suitable throughout training. In practice, this assumption is brittle for video: training on samples misaligned with the model’s evolving competence—either too trivial or too difficult—results in inefficient optimization, unstable gradients, and catastrophic forgetting.

Curriculum Learning (CL) (Narvekar et al., 2020; Wang et al., 2021; Soviany et al., 2022) offers an alternative by organizing training data from easy to hard. Yet, curriculum strategies tailored for video understanding remain underdeveloped. Existing methods mainly rely on simple scalar difficulty metrics, such as sequence length, model perplexity, or response accuracy (Li et al., 2025b; Dong et al., 2025). We argue that such scalar metrics are insufficient for video domain, where difficulty is inherently multi-dimensional. A clip may impose heavy visual-temporal perception load while requiring only shallow reasoning, or conversely, be visually simple yet demand deep causal reasoning. Collapsing these orthogonal challenges into a single scalar obscures the distinction between perceptual and cognitive complexity, inevitably conflating heterogeneous learning demands and resulting in

*corresponding author

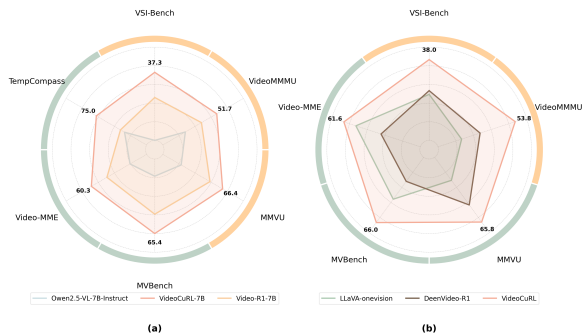


Figure 1: Multidimensional performance comparison. (Left) Comparison with baselines. (Right) Comparison with other open-source models.

unstable training trajectories.

To address this, we introduce **VideoCuRL**, a novel curriculum reinforcement learning framework that explicitly disentangles video difficulty into two orthogonal dimensions: Visual-Temporal Perception Load and Cognitive Reasoning Depth. Instead of relying on expensive model-based generation, VideoCuRL utilizes efficient, training-free proxies to map data onto a 2D curriculum grid. This allows us to employ a competence-aware Diagonal Wavefront strategy, effectively scheduling training from base alignment to complex reasoning while avoiding the prohibitive overhead of existing methods. To guarantee optimization stability, we further integrate Dynamic Sparse KL and Structured Revisiting. These mechanisms counteract the twin challenges of reward sparsity in complex reasoning and catastrophic forgetting of perceptual skills, effectively preventing policy degradation while ensuring the retention of fundamental spatiotemporal capabilities.

We evaluate VideoCuRL on a comprehensive suite of benchmarks, spanning complex video reasoning (Yang et al., 2025; Zhao et al., 2025; Hu et al., 2025) and general video understanding (Fu et al., 2025; Li et al., 2024b; Liu et al., 2024d). Empirical results (see Figure 1) demonstrate that VideoCuRL consistently outperforms strong RL baselines and prior curriculum strategies, achieving significant gains on reasoning-intensive tasks while simultaneously improving perceptual robustness. Notably, VideoCuRL matches or exceeds the performance of generation-based curricula with negligible computational overhead, confirming that our orthogonal difficulty modeling is both effective and scalable.

In summary, our contributions are threefold:

- We propose Orthogonal Difficulty Decomposition, a novel paradigm that challenges traditional scalar metrics by explicitly disentangling video difficulty into Visual-Temporal and Cognitive dimensions.
- We introduce VideoCuRL, a scalable two-dimensional curriculum reinforcement learning framework with a competence-aware Diagonal Wavefront scheduler and robust optimization mechanisms.
- Extensive experiments show that VideoCuRL substantially outperforms existing video RL baselines, securing clear gains on reasoning-intensive tasks while preserving fundamental visual capabilities. It validates the superiority of our 2D curriculum over random shuffling and scalar-based strategies.

2 Related Work

2.1 Large Multimodal Models for Video Understanding

Video understanding has been significantly advanced by Large Multimodal Models (LMMs) (Lin et al., 2024a; Zhang et al., 2024; Weng et al., 2024; Zhang et al., 2025a; Jin et al., 2025). Early approaches adapted image-based architectures via pooling mechanisms (Xu et al., 2024; Liu et al., 2024c) or temporal adapters (Liu et al., 2024b) for basic modal alignment. Current state-of-the-art Video-LLMs (Bai et al., 2025; Wang et al., 2025) integrate visual-temporal encoders with Large Language Models (LLMs) to handle extended temporal contexts. Despite these improvements, a critical disconnect persists between visual perception and cognitive reasoning. Standard Supervised Fine-Tuning (SFT) often fails to elicit robust spatiotemporal reasoning, as models tend to exploit language priors or shallow visual cues rather than grounding logic in temporal dynamics. Consequently, SFT is increasingly seen as insufficient for mastering complex video tasks.

2.2 Reinforcement Learning in Multimodal Post-training

Reinforcement Learning (RL) has become a standard paradigm to enhance reasoning and alignment in LLMs (Schulman et al., 2017; Ouyang et al., 2022; Rafailov et al., 2023; Shao et al., 2024) and has recently extended to multimodal domains (Yu et al., 2024; Wang et al., 2024; Huang et al., 2025;

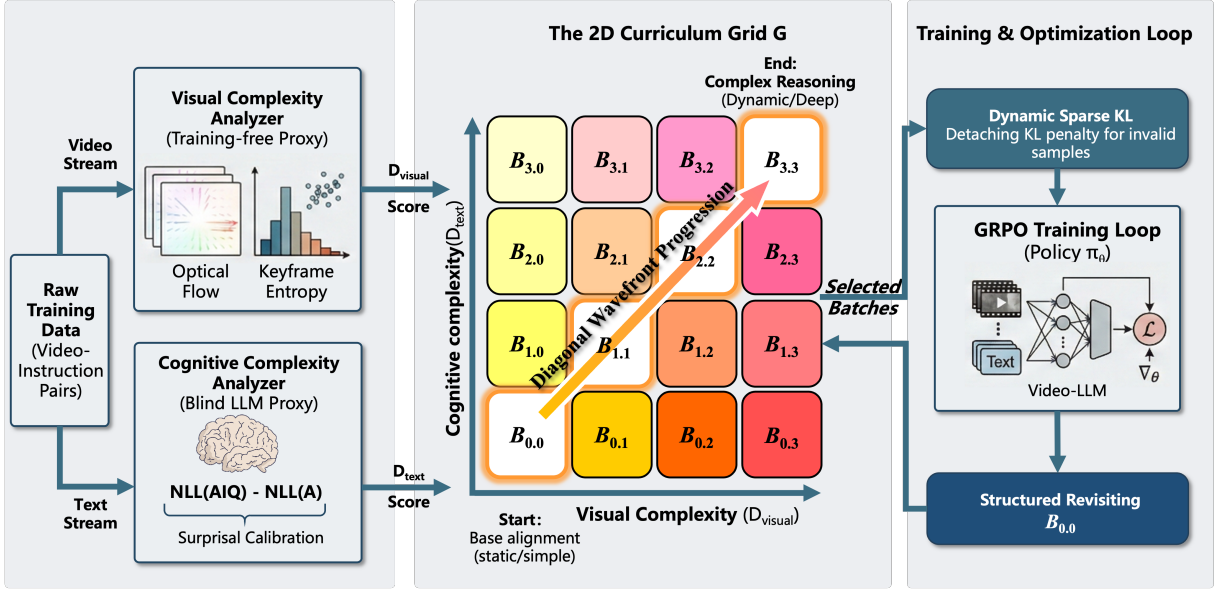


Figure 2: **Overview of VideoCuRL.** VideoCuRL disentangles video difficulty into orthogonal visual-temporal (D_{visual}) and cognitive (D_{text}) dimensions. A competence-aware Diagonal Wavefront strategy schedules training across a 2D grid, supported by Dynamic Sparse KL and Structured Revisiting to stabilize RL and prevent forgetting.

Shen et al., 2025; Jin et al., 2026). By optimizing task-level rewards, RL enables exploration beyond supervised signals. However, its application to video understanding remains challenging. Most existing methods (Feng et al., 2025; Li et al., 2025c; Zhang et al., 2025b) rely on random data sampling, implicitly treating all samples as equally suitable throughout training. This assumption is inefficient for video where samples vary dramatically in both visual perception load and cognitive reasoning demand. To address this, we propose VideoCuRL, the first curriculum reinforcement learning framework that explicitly decouples video difficulty into two orthogonal dimensions: visual and cognitive.

3 Methodology

3.1 Overview

Applying Reinforcement Learning (RL) combined with Curriculum Learning (CL) to Video-LLMs remains unexplored. Current video post-training paradigms predominately rely on Supervised Fine-Tuning (SFT) or RL with random data shuffling, neglecting the progressive nature of human learning. Therefore we propose VideoCuRL, a framework that disentangles difficulty into two orthogonal axes: Visual-Temporal Complexity (D_{visual}) and Cognitive Reasoning Complexity (D_{text}). We construct a 2D curriculum grid and employ a diagonal wavefront scheduling strategy equipped with structured revisiting and adaptive optimization to

robustly enhance reasoning capabilities.

3.2 Visual-Temporal Complexity (D_{visual})

Standard video instruction tuning treats all video clips equally, ignoring that a static shot and a fast-paced montage impose vastly different perceptual loads. We quantify perception load by synthesizing motion dynamics and information redundancy:

Motion Intensity (ϕ_{flow}): To quantitatively measure temporal dynamics, we utilize dense optical flow to capture pixel-wise displacement between consecutive frames. Let a video sequence be denoted as $V = \{F_1, F_2, \dots, F_T\}$, where $F_t \in \mathbb{R}^{H \times W \times 3}$. For each pair of adjacent frames (F_t, F_{t+1}), we compute the dense optical flow field $\mathbf{V}_t \in \mathbb{R}^{H \times W \times 2}$, where each coordinate (x, y) contains a displacement vector $\mathbf{v}_{t,x,y} = (u_{t,x,y}, v_{t,x,y})$ representing the horizontal and vertical shift. We define the frame-level motion score m_t as the spatial average of Euclidean norms of these vectors:

$$\begin{aligned}
 m_t &= \frac{1}{H \times W} \sum_{x=1}^H \sum_{y=1}^W \|\mathbf{v}_{t,x,y}\|_2 \\
 &= \frac{1}{H \times W} \sum_{x,y} \sqrt{u_{t,x,y}^2 + v_{t,x,y}^2}
 \end{aligned} \tag{1}$$

The video-level Motion Intensity ϕ_{flow} is obtained by temporally pooling the frame-level scores, rep-

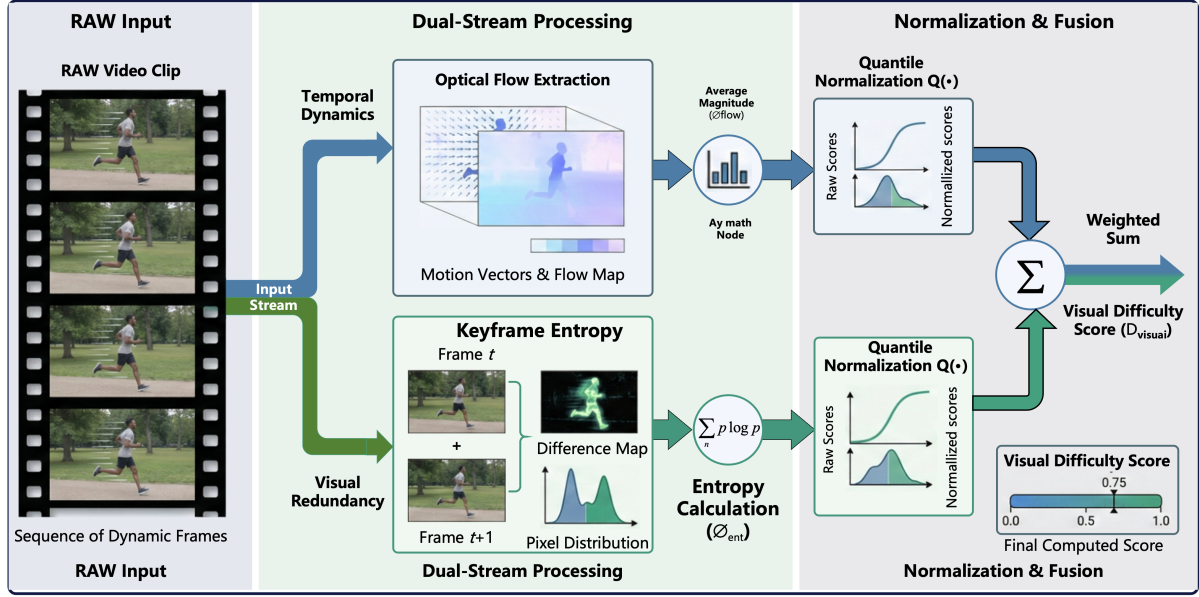


Figure 3: **Visual Complexity Estimation.** Perceptual load is quantified via dual-stream processing: Motion Intensity (ϕ_{flow}) from optical flow and Information Density (ϕ_{ent}) from frame-difference entropy. Scores are fused and quantile-normalized to ensure a balanced difficulty distribution.

representing the global kinetic energy of the clip:

$$\phi_{flow}(V) = \frac{1}{T-1} \sum_{t=1}^{T-1} m_t \quad (2)$$

High ϕ_{flow} values indicate rapid object movement or significant camera shifts, necessitating robust temporal tracking capabilities from the model.

Visual Information Density (ϕ_{ent}): While optical flow captures kinetic magnitude, it does not account for visual redundancy. To quantify the diversity of visual information, we calculate the entropy of pixel-wise differences between consecutive frames. We specifically employ temporal difference entropy rather than spatial frame entropy. This design choice effectively disentangles visual texture complexity (handled well by the pre-trained spatial encoder) from temporal information density (the primary challenge for video reasoning), preventing static scenes with high-frequency textures from being misclassified as difficult samples.

Let $\Delta F_t(x, y)$ be the grayscale intensity of the difference map at coordinates (x, y) , taking values in the integer range $[0, 255]$:

$$\Delta F_t = |\text{Gray}(F_{t+1}) - \text{Gray}(F_t)| \quad (3)$$

We first compute the normalized histogram (probability distribution) p_t of these intensity values. For each intensity level $k \in \{0, \dots, 255\}$, the proba-

bility $p_t(k)$ is given by:

$$p_t(k) = \frac{\text{Count}(\Delta F_t(x, y) = k)}{H \times W} \quad (4)$$

where $\sum_{k=0}^{K-1} p_t(k) = 1$. The visual information entropy E_t for the frame difference is then calculated as:

$$E_t = - \sum_{k=0}^{K-1} p_t(k) \log_2(p_t(k) + \epsilon) \quad (5)$$

(Note: We add a small constant ϵ simply to ensure numerical stability where $p_t(k) = 0$.)

A low E_t implies that most pixels have similar changes (e.g., a static background or uniform motion), indicating high redundancy. Conversely, a high E_t suggests complex, non-uniform visual changes. The final video-level density score ϕ_{ent} is the temporal average of these entropy values:

$$\phi_{ent}(V) = \frac{1}{T-1} \sum_{t=1}^{T-1} E_t \quad (6)$$

This metric effectively filters out visually monotonous clips from information-dense footage, complementing the motion intensity score.

Final Difficulty (D_{visual}): The final visual difficulty score is a weighted combination, normalized via quantile normalization \mathcal{Q} to ensure a uniform distribution across the $[0, 1]$ interval:

$$D_{visual}(i) = \alpha \cdot \mathcal{Q}(\phi_{flow}(v_i)) + (1 - \alpha) \cdot \mathcal{Q}(\phi_{ent}(v_i)) \quad (7)$$

3.3 Cognitive Reasoning Complexity (D_{text})

The Bias of Raw Perplexity. A naive approach for estimating cognitive difficulty is to use the raw perplexity (or Negative Log-Likelihood, NLL) of the ground-truth answer A given the question Q . However, this metric is biased by the *language prior*: answers containing low-frequency vocabulary yield high NLL scores, even if the reasoning is trivial. This leads to a misalignment where "linguistically rare" is mistaken for "logically hard."

Calibrated Surprisal. To mitigate this, we measure the Information Gain provided by the question Q toward the answer A . We define the cognitive complexity score S_{cog} as the reduction in uncertainty, which effectively isolates the reasoning requirement from inherent linguistic unpredictability. Formally, for an answer sequence $A = \{w_1, w_2, \dots, w_n\}$, we calculate the difference between the conditional and unconditional NLL using the chain rule of probability:

$$S_{cog}(A, Q) = \sum_{t=1}^{|A|} \underbrace{-\log \pi_{\phi}(w_t | w_{<t}, Q)}_{\text{Conditional Surprisal}} - \sum_{t=1}^{|A|} \underbrace{(-\log \pi_{\phi}(w_t | w_{<t}))}_{\text{Language Prior}} \quad (8)$$

A high S_{cog} indicates that the answer A is not easily predictable from the question Q alone without strong reasoning, whereas a low score implies high linguistic predictability.

Crucially, this metric relies solely on a "Blind" LLM processing text pairs, avoiding the massive overhead of video encoding or autoregressive video decoding. Finally, we map the raw scores to a uniform distribution via Quantile Normalization to obtain the final cognitive difficulty dimension:

$$D_{text}(i) = \mathcal{Q}(S_{cog}(A_i, Q_i)) \quad (9)$$

By combining D_{text} and D_{visual} , we construct the orthogonal basis for our curriculum grid \mathcal{G} .

3.4 Diagonal Wavefront Scheduling

To effectively navigate the disjoint difficulty landscape, we map the training dataset \mathcal{D} into a $K \times K$ grid \mathcal{G} . Each bucket $B_{i,j}$ encapsulates samples with visual complexity level i and cognitive complexity level j . Instead of a linear progression, which risks overfitting to one modality, we propose a Diagonal Wavefront strategy driven by model competence.

Competence-Aware Gating: We define the Local Competence Score ($LCS_{i,j}$) as the exponential moving average of rewards obtained from samples in bucket $B_{i,j}$ at step t :

$$LCS_{i,j}^{(t)} = \gamma \cdot LCS_{i,j}^{(t-1)} + (1 - \gamma) \cdot \bar{R}_{batch} \quad (10)$$

where \bar{R}_{batch} is the mean reward of the current batch sampled from $B_{i,j}$.

Wavefront Expansion: Let \mathcal{A}_t denote the set of active buckets (the "Wavefront"). Training initiates at $B_{0,0}$ (Static Visuals, Atomic Queries). The curriculum unfolds dynamically: a neighboring harder bucket $B_{i',j'}$ (where $i' \geq i, j' \geq j$) is added to \mathcal{A}_{t+1} if and only if the competence in the current frontier exceeds a threshold \mathcal{T} :

$$\text{Expand if } LCS_{i,j}^{(t)} > \mathcal{T} \quad (11)$$

This mechanism enforces a "Base Alignment" constraint: the model must demonstrate fundamental proficiency in simpler visual-linguistic grounding before attempting samples with high temporal loads or complex causal chains.

3.5 Robust Optimization for Video Reasoning

Training Video-LLMs with RL is notoriously unstable due to the sparsity of valid reasoning paths in complex video tasks. We adapt two mechanisms to stabilize our 2D curriculum.

3.5.1 Dynamic Sparse KL

In high-difficulty regions of our grid (e.g., $B_{K,K}$), the model may initially fail to generate any correct responses, leading to reward collapse (all rewards are 0). Standard RL objectives would force the model to minimize KL divergence with the reference model, causing Policy Degradation. To counter this, we implement a Dynamic Sparse KL mechanism. We dynamically monitor the reward variance σ_r within each group. If $\sigma_r = 0$ (indicating all samples are incorrect or correct), we detach the KL penalty term. This allows the model to explore freely without being penalized for deviating from a reference model that likely also fails on these hard video samples.

3.5.2 Structured Revisiting

To address the risk of catastrophic forgetting common in course learning, we propose a **Structured Revisiting** mechanism. While random replay is a common method for mitigating forgetting, we argue that this indiscriminate sampling approach

Model	Frames	Video Reasoning Benchmarks			General Video Benchmarks		
		VSI-Bench	VideoMMMU	MMVU	MVBench	TempCompass	VideoMME
LLaVA-onevision-7B	32	32.4	33.8	49.2	56.7	-	58.2
LLaVA-video-7B	32	-	34.4	48.8	58.6	-	63.3
VILA-1.5-40B	-	31.2	34.0	-	-	-	60.1
Kangaroo-8B	-	-	-	-	61.1	62.5	56.0
DeepVideo-R1-3B	-	33.0	40.7	59.0	49.6	63.1	51.1
w/o Curriculum							
Qwen2.5-VL-7B-Instruct	16	27.7	47.8	59.2	57.4	72.2	53.1
Video-R1-7B	16	34.6	49.8	64.2	62.7	72.6	57.4
Qwen2.5-VL-7B-Instruct	32	30.1	48.1	60.0	59.0	72.6	56.6
Video-R1-7B	32	35.8	52.3	63.8	63.9	73.2	59.3
Qwen2.5-VL-7B-Instruct	64	31.4	50.4	60.0	59.2	72.9	59.6
Video-R1-7B	64	37.1	52.4	63.8	64.8	73.2	61.4
w/ Curriculum							
Length-based	16	35.1	50.5	64.5	63.0	72.5	60.2
Generation-based	16	36.5	51.0	65.6	64.7	73.8	60.4
VideoCuRL-7B	16	37.3(↑2.7)	51.7(↑1.9)	66.4(↑2.2)	65.4(↑2.7)	75.0(↑2.4)	60.3(↑2.9)
VideoCuRL-7B	32	38.0(↑2.2)	53.8(↑1.5)	65.8(↑2.0)	66.0(↑2.1)	75.8(↑2.6)	61.6(↑2.3)
VideoCuRL-7B	64	39.6(↑2.5)	54.5(↑2.1)	65.7(↑1.9)	66.9(↑2.1)	75.5(↑2.3)	63.9(↑2.5)

Table 1: Main experimental results on video reasoning and general video understanding benchmarks. All models are built on Qwen2.5-VL-7B-Instruct (Bai et al., 2025) and trained on the same Video-R1-260k dataset (Feng et al., 2025) for fair comparison. $\uparrow x$ denotes the performance improvement over Video-R1-7B (same frame setting).

is not optimal for post-video training. Therefore, VideoCuRL does not uniformly sample historical data but instead selectively revisits orthogonal basis buckets from mastered lessons. Specifically, this strategy prioritizes perceptual basis buckets ($B_{K,0}$, high visual complexity and low text difficulty) to maintain the model’s spatiotemporal tracking capabilities, while revisiting inference basis buckets ($B_{0,K}$, low visual complexity and high text difficulty) to solidify logical deduction skills. By prioritizing these basis buckets during replay, the model can efficiently and evenly retain and enhance its capabilities in both perception and reasoning dimensions with minimal computational overhead.

4 Experiments

4.1 Main Results

We present the comprehensive evaluation results in Table 1, comparing VideoCuRL-7B against the base model Qwen2.5-VL-7B-Instruct (Bai et al., 2025) and the state-of-the-art RL baseline Video-R1-7B (Feng et al., 2025). Both baselines and our model are trained on the Video-R1-260k dataset to ensure a fair comparison. The results demonstrate that VideoCuRL consistently outperforms both the SFT baseline and the standard RL approach across reasoning-intensive and general video understanding benchmarks.

Superiority in Complex Reasoning. The core advantage of our orthogonal difficulty decomposition is most evident in reasoning-heavy tasks. On VSI-Bench (Yang et al., 2025), which demands intricate causal deduction, VideoCuRL-7B achieves 37.3% (16 frames) and 39.6% (64 frames), surpassing Video-R1-7B by margins of 2.7% and 2.5% respectively. Similarly, on VideoMMMU (Hu et al., 2025) and MMVU (Zhao et al., 2025), our method consistently leads the leaderboard.

Robustness in General Perception. A common pitfall in RL post-training is "Catastrophic Forgetting" of basic perception skills while chasing higher reasoning rewards. However, VideoCuRL mitigates this effectively. On MVBench (Li et al., 2024b), a comprehensive perception benchmark, our model maintains a clear lead, achieving 65.4% accuracy at 16 frames compared to 62.7% for Video-R1 and 57.4% for the base model. Notably, on TempCompass (Liu et al., 2024d), VideoCuRL achieves 75.0% (16 frames) and 75.8% (32 frames), outperforming Video-R1 by significant margins ($\uparrow 2.4\%$ and $\uparrow 2.6\%$). This empirical evidence confirms that our Structured Revisiting strategy successfully preserves and even refines fine-grained visual tracking skills.

Comparison with Scalar and Generation-based Curricula. As shown in Table 1, Length-

Model Variant	Video Reasoning Benchmarks			General Video Benchmarks		
	VSI-Bench (16f)	VideoMMMU (32f)	MMVU (32f)	MVBench (16f)	TempCompass (32f)	VideoMME (64f)
VideoCuRL-7B (Full Model)	37.3	53.8	65.8	65.4	75.8	63.9
<i>Ablation 1: Single Difficulty Axis</i>						
- Only D_{visual}	35.8 (-1.5)	52.6 (-1.2)	64.5 (-1.3)	65.8 (+0.4)	74.3 (-1.5)	63.2 (-0.7)
- Only D_{text}	36.1 (-1.2)	52.1 (-1.7)	64.2 (-1.6)	62.3 (-3.1)	74.2 (-1.6)	62.8 (-1.1)
<i>Ablation 2: Core Component Removal</i>						
- W/O Structured Revisiting	35.7 (-1.6)	51.9 (-1.9)	63.7 (-2.1)	62.8 (-2.6)	73.5 (-2.3)	62.1 (-1.8)
- W/O Dynamic Sparse KL	34.9 (-2.4)	51.7 (-2.1)	63.9 (-1.9)	63.5 (-1.9)	73.1 (-2.7)	62.2 (-1.7)
<i>Ablation 3: Cognitive Metric Replacement</i>						
- Raw Perplexity ($PPL(A Q)$)	35.5 (-1.8)	52.3 (-1.5)	64.3 (-1.5)	64.2 (-1.2)	74.7 (-1.1)	62.9 (-1.0)

Table 2: Ablation study results. All variants are based on VideoCuRL-7B. Numbers in parentheses denote performance change (\downarrow) or (\uparrow) compared to the full model. "16f/32f/64f" indicate frame settings.

based Curriculum outperforms the random baseline Video-R1, confirming that organizing data from easy to hard is generally beneficial. However, it significantly lags behind VideoCuRL. Furthermore, Generation-based Curriculum achieves competitive performance, slightly behind VideoCuRL. However, this comes at a prohibitive cost: it requires several full inference passes over full dataset before training (Full implementation details are shown in Appendix B), consuming hundreds of GPU hours. In contrast, VideoCuRL achieves superior performance using only training-free proxies (D_{visual} and D_{text}) that are computed offline on CPUs. This demonstrates that our orthogonal decomposition effectively captures the true "hardness" of samples with dramatically higher efficiency.

4.2 Experimental Setup

Implementation Details. We use Qwen2.5-VL-7B (Bai et al., 2025) as our base model. The 2D Curriculum Grid is set to 4×4 ($K = 4$). We provide a detailed sensitivity analysis regarding the choice of K in Appendix D, demonstrating that $K = 4$ strikes an optimal balance between curriculum granularity and training stability. All models are trained using GRPO with a KL coefficient $\beta = 0.04$. Crucially, our proxy metric calculation is performed offline on CPUs, incurring zero GPU inference cost for difficulty estimation. For each pair of adjacent frames (F_t, F_{t+1}), we compute the dense optical flow field V_t using the Gunnar Farneback algorithm (Farneback, 2003) implementation in OpenCV. We use a pyramid scale of 0.5, 3 levels, and a window size of 15 to balance CPU efficiency and motion sensitivity. We compute the entropy on grayscale difference maps. Empirical analysis shows that luminance changes dominate temporal information density in video, enabling efficient CPU-based computation without sacrificing

the metric’s discriminative power.

Training Data & Benchmarks. To strictly evaluate the efficacy of our curriculum strategy, we utilize the exact same data split as the baseline Video-R1, ensuring no data leakage. For evaluation, we employ a comprehensive suite of benchmarks categorized into two groups:

Video Reasoning Benchmarks: We use VSI-Bench (Yang et al., 2025), VideoMMMU (Hu et al., 2025), and MMVU (Zhao et al., 2025) to assess the model’s capability in causal deduction and complex problem-solving.

General Video Benchmarks: We use MVBench (Li et al., 2024b) and TempCompass (Liu et al., 2024d) to evaluate fundamental perception skills (e.g., object tracking, action recognition), and VideoMME (Fu et al., 2025) for long-context understanding.

We use LLaVA-onevision (Li et al., 2024a), LLaVA-video (Zhang et al., 2024), VILA (Lin et al., 2024b), Kangaroo (Liu et al., 2024a), DeepVideo-R1 (Park et al., 2025) for comparison.

4.3 Ablation Studies

To evaluate the contribution of each individual component in VideoCuRL, we conduct comprehensive ablation studies on both reasoning-heavy and perception-oriented benchmarks. The results are summarized in Table 2.

4.3.1 Impact of Orthogonal Difficulty Decomposition

We first investigate the necessity of disentangling visual and cognitive complexities by training two variants: one using only the Visual Complexity axis (D_{visual}) and another using only the Cognitive Reasoning axis (D_{text}).

Visual-only (D_{visual}): While this variant performs relatively well on perception tasks like

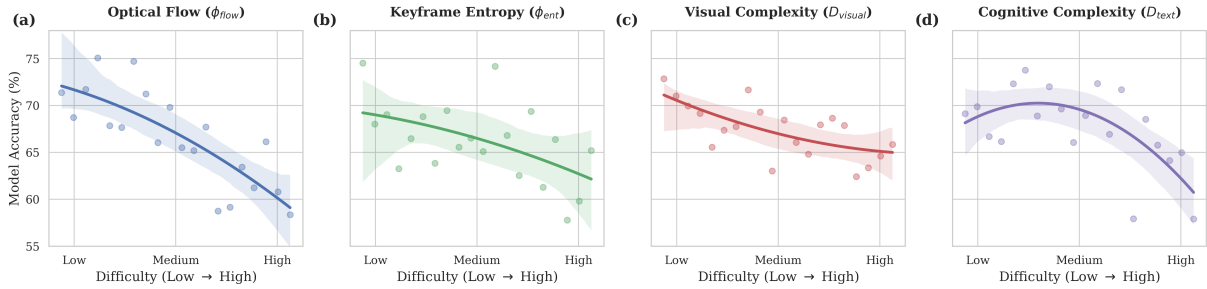


Figure 4: **Quantitative Cross-Validation of Difficulty Proxies.** We analyze the correlation between our training-free difficulty metrics and model accuracy on the training set. (a-b) Individual visual proxies (ϕ_{flow} , ϕ_{ent}) exhibit high variance, indicating that physical features alone are noisy predictors of difficulty. (c) The fused Visual Complexity (D_{visual}) smooths these fluctuations, providing a robust estimator of perceptual load. (d) Cognitive Complexity (D_{text}) reveals the nonlinear relationship between logical challenge and model performance.

MVBench (+0.4% over full model), its performance on reasoning benchmarks (e.g., VSI-Bench) drops significantly by 1.5%. This confirms that perception-based scheduling alone is insufficient for mastering complex logical deduction.

Cognitive-only (D_{text}): Removing the visual dimension leads to a 3.1% decline in MVBench and a 1.6% drop in TempCompass. This suggests that without a visual curriculum, models suffer from "Alignment Tax," where they struggle to ground reasoning in complex temporal cues.

4.3.2 Effectiveness of Robust Optimization Mechanisms

We further examine our two stabilization strategies: Structured Revisiting and Dynamic Sparse KL.

W/O Structured Revisiting: Removing this mechanism leads to a drop across all benchmarks, with the most notable impact on MVBench (-2.6%) and TempCompass (-2.3%). This validates our hypothesis that random replay dilutes critical "basis" samples, leading to catastrophic forgetting of fundamental perception skills.

W/O Dynamic Sparse KL: This variant shows the sharpest decline in high-reasoning tasks, such as VSI-Bench (-2.4%) and VideoMMU (-2.1%). Without detaching the KL penalty in high-difficulty zones, the model suffers from reward collapse and policy degradation, as it is penalized for exploring beyond a failing reference model.

4.3.3 Cognitive Metric: Calibrated Surprisal vs. Raw Perplexity

Finally, we replace our Calibrated Surprisal (PMI-based) with standard Raw Perplexity (PPL). The PPL-based variant yields inferior results across the board, particularly on VSI-Bench (-1.8%) and

VideoMMU (-1.5%). This demonstrates that raw NLL is biased toward linguistic frequency, whereas our calibrated metric more accurately captures the true logical information gain required for spatiotemporal reasoning.

4.4 Empirical Validation of Difficulty Proxies

To evaluate the reliability of our difficulty proxies, we analyze the correlation between difficulty scores and model accuracy. As shown in Figure 4(a-b), raw physical metrics like Optical Flow (ϕ_{flow}) and Keyframe Entropy (ϕ_{ent}) exhibit high variance, suggesting that individual motion features are insufficient for predicting model failure. By fusing them into Visual Complexity D_{visual} , we observe a stabilized inverse trend. For Cognitive Complexity D_{text} , the model accuracy exhibits a general downward trajectory, although local fluctuations persist due to the stochastic nature of autoregressive generation and potential visual-linguistic misalignment. The results of cross-validation confirms that our proxies effectively capture the statistical "hardness" of samples, providing a robust signal for curriculum scheduling. Detailed experimental setups are provided in Appendix C.

5 Conclusion

In this paper, we propose VideoCuRL, a novel Video Curriculum Reinforcement Learning framework. By employing efficient difficulty proxies, VideoCuRL constructs a 2D curriculum grid navigated by a competence-aware Diagonal Wavefront strategy. Extensive evaluations demonstrate that VideoCuRL consistently outperforms strong RL baselines and generation-based curricula without additional training or inference overhead. Through Dynamic Sparse KL and Structured Revisiting, we

also ensure training stability and mitigate catastrophic forgetting of fundamental perceptual skills. Our results underscore that for complex multi-modal reasoning, strategic organization of data is as pivotal as the scale of the data itself.

Limitations

Our framework prioritizes computational scalability by utilizing efficient, training-free difficulty proxies, a design choice that entails a deliberate trade-off between semantic granularity and processing efficiency. While our physical proxies effectively capture dynamic complexity without the prohibitive costs of generation-based oracles, they may under-represent the difficulty of visually subtle yet semantically dense scenes, such as micro-expressions. Similarly, the structural abstraction of difficulty into a discrete grid, while essential for orthogonal separation, introduces coarse-grained boundaries compared to a theoretical continuous difficulty manifold, leaving room for future refinement in continuous curriculum modeling.

Additionally, as a post-training refinement framework, VideoCuRL operates under the standard cold-start assumption inherent to Reinforcement Learning. The methodology presupposes that the base SFT model possesses a foundational level of instruction-following capability to explore valid reasoning paths and initiate the wavefront expansion. Consequently, our approach is positioned as a competence amplifier for aligned models rather than a solution for learning multimodal representations completely

References

- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Lu Dong, Haiyu Zhang, Han Lin, Ziang Yan, Xiangyu Zeng, Hongjie Zhang, Yifei Huang, Yi Wang, Zhenhua Ling, Limin Wang, and 1 others. 2025. Videotgr1: Boosting video temporal grounding via curriculum reinforcement learning on reflected boundary annotations. *arXiv preprint arXiv:2510.23397*.
- Gunnar Farnebäck. 2003. Two-frame motion estimation based on polynomial expansion. In *Scandinavian conference on Image analysis*, pages 363–370. Springer.
- Kaituo Feng, Kaixiong Gong, Bohao Li, Zonghao Guo, Yibing Wang, Tianshuo Peng, Junfei Wu, Xiaoying Zhang, Benyou Wang, and Xiangyu Yue. 2025. Video-r1: Reinforcing video reasoning in mllms. *arXiv preprint arXiv:2503.21776*.
- 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 *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 24108–24118.
- Kairui Hu, Penghao Wu, Fanyi Pu, Wang Xiao, Yuanhan Zhang, Xiang Yue, Bo Li, and Ziwei Liu. 2025. Video-mmmu: Evaluating knowledge acquisition from multi-discipline professional videos. *arXiv preprint arXiv:2501.13826*.
- Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and Shaohui Lin. 2025. Vision-r1: Incentivizing reasoning capability in multimodal large language models. *arXiv preprint arXiv:2503.06749*.
- Yunseok Jang, Yale Song, Youngjae Yu, Youngjin Kim, and Gunhee Kim. 2017. Tgif-qa: Toward spatio-temporal reasoning in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2758–2766.
- Hongbo Jin, Qingyuan Wang, Wenhao Zhang, Yang Liu, and Sijie Cheng. 2025. Videomem: Enhancing ultra-long video understanding via adaptive memory management. *arXiv preprint arXiv:2512.04540*.
- Hongbo Jin, Rongpeng Zhu, Jiayu Ding, Wenhao Zhang, and Ge Li. 2026. Himac: Hierarchical macro-micro learning for long-horizon llm agents. *arXiv preprint arXiv:2603.00977*.
- Jie Lei, Licheng Yu, Mohit Bansal, and Tamara Berg. 2018. Tvqa: Localized, compositional video question answering. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pages 1369–1379.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, and 1 others. 2024a. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*.
- KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhao Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. 2025a. Videochat: Chat-centric video understanding. *Science China Information Sciences*, 68(10):200102.
- Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, and 1 others. 2024b. Mvbench: A comprehensive multi-modal video understanding benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22195–22206.

- Renda Li, Hailang Huang, Fei Wei, Feng Xiong, Yong Wang, and Xiangxiang Chu. 2025b. Adacurl: Adaptive curriculum reinforcement learning with invalid sample mitigation and historical revisiting. *arXiv preprint arXiv:2511.09478*.
- Xinhao Li, Ziang Yan, Desen Meng, Lu Dong, Xiangyu Zeng, Yinan He, Yali Wang, Yu Qiao, Yi Wang, and Limin Wang. 2025c. Videochat-r1: Enhancing spatio-temporal perception via reinforcement fine-tuning. *arXiv preprint arXiv:2504.06958*.
- 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.
- Ji Lin, Hongxu Yin, Wei Ping, Pavlo Molchanov, Mohammad Shoeybi, and Song Han. 2024b. Vila: On pre-training for visual language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 26689–26699.
- Jiajun Liu, Yibing Wang, Hanghang Ma, Xiaoping Wu, Xiaoqi Ma, Xiaoming Wei, Jianbin Jiao, Enhua Wu, and Jie Hu. 2024a. Kangaroo: A powerful video-language model supporting long-context video input. *arXiv preprint arXiv:2408.15542*.
- Ruyang Liu, Chen Li, Yixiao Ge, Thomas H Li, Ying Shan, and Ge Li. 2024b. Bt-adapter: Video conversation is feasible without video instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13658–13667.
- Ruyang Liu, Haoran Tang, Haibo Liu, Yixiao Ge, Ying Shan, Chen Li, and Jiankun Yang. 2024c. Ppllava: Varied video sequence understanding with prompt guidance. *arXiv preprint arXiv:2411.02327*.
- Yuanxin Liu, Shicheng Li, Yi Liu, Yuxiang Wang, Shuhuai Ren, Lei Li, Sishuo Chen, Xu Sun, and Lu Hou. 2024d. Tempcompass: Do video llms really understand videos? *arXiv preprint arXiv:2403.00476*.
- Sanmit Narvekar, Bei Peng, Matteo Leonetti, Jivko Sinapov, Matthew E Taylor, and Peter Stone. 2020. Curriculum learning for reinforcement learning domains: A framework and survey. *Journal of Machine Learning Research*, 21(181):1–50.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Jinyoung Park, Jeehye Na, Jinyoung Kim, and Hyunwoo J Kim. 2025. Deepvideo-r1: Video reinforcement fine-tuning via difficulty-aware regressive grpo. *arXiv preprint arXiv:2506.07464*.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36:53728–53741.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Haozhan Shen, Peng Liu, Jingcheng Li, Chunxin Fang, Yibo Ma, Jiajia Liao, Qiaoli Shen, Zilun Zhang, Kangjia Zhao, Qianqian Zhang, and 1 others. 2025. Vlm-r1: A stable and generalizable r1-style large vision-language model. *arXiv preprint arXiv:2504.07615*.
- Petru Soviany, Radu Tudor Ionescu, Paolo Rota, and Nicu Sebe. 2022. Curriculum learning: A survey. *International Journal of Computer Vision*, 130(6):1526–1565.
- Fei Wang, Wenxuan Zhou, James Y Huang, Nan Xu, Sheng Zhang, Hoifung Poon, and Muhao Chen. 2024. mdpo: Conditional preference optimization for multimodal large language models. *arXiv preprint arXiv:2406.11839*.
- Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, and 1 others. 2025. Internvl3. 5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*.
- Xin Wang, Yudong Chen, and Wenwu Zhu. 2021. A survey on curriculum learning. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):4555–4576.
- Yuetian Weng, Mingfei Han, Haoyu He, Xiaojun Chang, and Bohan Zhuang. 2024. Longvlm: Efficient long video understanding via large language models. In *European Conference on Computer Vision*, pages 453–470. Springer.
- Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. 2017. Video question answering via gradually refined attention over appearance and motion. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 1645–1653.
- Lin Xu, Yilin Zhao, Daquan Zhou, Zhijie Lin, See Kiong Ng, and Jiashi Feng. 2024. Pllava: Parameter-free llava extension from images to videos for video dense captioning. *arXiv preprint arXiv:2404.16994*.

- Jihan Yang, Shusheng Yang, Anjali W Gupta, Rilyn Han, Li Fei-Fei, and Saining Xie. 2025. Thinking in space: How multimodal large language models see, remember, and recall spaces. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 10632–10643.
- Tianyu Yu, Yuan Yao, Haoye Zhang, Taiwen He, Yifeng Han, Ganqu Cui, Jinyi Hu, Zhiyuan Liu, Hai-Tao Zheng, Maosong Sun, and 1 others. 2024. Rlhf-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional human feedback. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13807–13816.
- Zhou Yu, Dejing Xu, Jun Yu, Ting Yu, Zhou Zhao, Yuet-ing Zhuang, and Dacheng Tao. 2019. Activitynet-qa: A dataset for understanding complex web videos via question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9127–9134.
- Boqiang Zhang, Kehan Li, Zesen Cheng, Zhiqiang Hu, Yuqian Yuan, Guanzheng Chen, Sicong Leng, Yuming Jiang, Hang Zhang, Xin Li, and 1 others. 2025a. Videollama 3: Frontier multimodal foundation models for image and video understanding. *arXiv preprint arXiv:2501.13106*.
- Xingjian Zhang, Siwei Wen, Wenjun Wu, and Lei Huang. 2025b. Tynyllava-video-r1: Towards smaller lmms for video reasoning. *arXiv preprint arXiv:2504.09641*.
- Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. 2024. Video instruction tuning with synthetic data. *arXiv preprint arXiv:2410.02713*.
- Yilun Zhao, Haowei Zhang, Lujing Xie, Tongyan Hu, Guo Gan, Yitao Long, Zhiyuan Hu, Weiyuan Chen, Chuhan Li, Zhijian Xu, and 1 others. 2025. Mmvu: Measuring expert-level multi-discipline video understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 8475–8489.

A Preliminaries: Group Relative Policy Optimization (GRPO)

In this work, we employ Group Relative Policy Optimization (GRPO) as our core reinforcement learning algorithm. While Proximal Policy Optimization (PPO) (Schulman et al., 2017) is a standard choice for aligning Large Language Models, it typically requires maintaining a value function (Critic) model roughly the same size as the policy model. In the context of Video-LLMs, where visual encoders and long-context histories already impose significant GPU memory overhead, maintaining a separate Critic is often computationally prohibitive.

GRPO circumvents the need for a value network by estimating the baseline directly from a group of sampled outputs. Formally, the training process is defined as follows:

Group Sampling. Given a batch of inputs consisting of video clips V and textual queries Q , we sample a group of G outputs $\{A_1, A_2, \dots, A_G\}$ from the current policy $\pi_{\theta_{old}}$. This allows the model to explore diverse reasoning paths for the same visual-linguistic context.

Advantage Estimation. Instead of relying on a learned value function $V(S)$ to compute the advantage, GRPO utilizes the group statistics. For each output A_i in the group, we compute a specific reward r_i (based on correctness or rule adherence). The advantage Adv_i is then calculated by normalizing the rewards within the group:

$$Adv_i = \frac{r_i - \text{mean}(\mathbf{r})}{\text{std}(\mathbf{r}) + \epsilon} \quad (12)$$

where $\mathbf{r} = \{r_1, \dots, r_G\}$ is the set of rewards for the group, and ϵ is a small constant for numerical stability. This effectively uses the average performance of the group as the baseline, encouraging the model to reinforce outputs that perform better than the expected average for that specific video-query pair.

Objective Function. The final optimization objective maximizes the surrogate loss while constraining the policy update via KL divergence, similar to PPO. The GRPO objective is formulated

as:

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}_{Q \sim \mathcal{D}, \{A_i\}_{i=1}^G \sim \pi_{\theta_{old}}} \left[\frac{1}{G} \sum_{i=1}^G \left(\min \left(\rho_i Adv_i, \text{clip}(\rho_i, 1 \pm \epsilon) Adv_i \right) - \beta D_{KL}(\pi_{\theta} \parallel \pi_{ref}) \right) \right] \quad (13)$$

where $\rho_i = \frac{\pi_{\theta}(A_i|V,Q)}{\pi_{\theta_{old}}(A_i|V,Q)}$ is the probability ratio, $\text{clip}(\cdot)$ limits the update step size, and D_{KL} represents the KL divergence between the current policy π_{θ} and the reference SFT model π_{ref} to prevent reward hacking.

In our VideoCuRL framework, we further modify this objective by introducing *Dynamic Sparse KL* (as detailed in Section 3.5.1) to handle the reward collapse issues inherent in high-difficulty curriculum stages.

B Implementation of Generation-based Curriculum Baseline

To further validate the efficiency and effectiveness of VideoCuRL, we introduced Generation-based Curriculum as a strong benchmark in our experiments. This approach represents the "model-as-teacher" strategy commonly used in the industry when explicit physical metrics are lacking. Its core logic is to pre-evaluate the training dataset D (i.e., Video-R1-260k) using a pre-trained LMMs, quantifying the difficulty based on the model's actual response to each sample.

Specifically, this method requires a pre-trained Video-LLM to perform complete inference on 260k video-instruction pairs. The difficulty score is determined by the accuracy generated by the model. By calculating the average accuracy of the model generating multiple benchmark answers A , we linearly divide the dataset into multiple stages from easy to difficult and train a single-dimensional linear course using the same learning rate and number of training steps as VideoCuRL.

However, this approach is limited by its extremely high pre-computational overhead. Due to the need for multiple autoregressive decodings of massive amounts of video data, this process consumes hundreds of GPU hours in the preprocessing stage, which is often unsustainable when dealing with large-scale video datasets. In contrast, VideoCuRL achieves equivalent accuracy in difficulty assessment on a CPU using lightweight

Grid Size (K)	VSI-Bench (16f)	MVBench (16f)	VideoMME (64f)	VideoMMU (64f)
$K = 3$ (3×3)	35.8%	63.2%	62.4%	53.1%
$K = 4$ (4×4)	37.3%	65.4%	63.9%	54.5%
$K = 5$ (5×5)	37.5%	65.7%	64.1%	54.3%

Table 3: Sensitivity analysis of the grid size K on video reasoning and perception benchmarks. $K = 4$ provides the optimal balance between granularity and stability.

visual proxy metrics such as optical flow and entropy, along with text-based surprise calibration, without any GPU inference cost.

Experimental results (see Table 1 in the main text) show that although the generative approach achieved a VSI-Bench accuracy of 36.5%, it was still slightly inferior to VideoCuRL-7B in cognitive reasoning and complex perceptual tasks. This indicates that relying solely on scalarized model feedback is insufficient to capture the orthogonality between "perceptual load" and "cognitive depth" in video understanding, while VideoCuRL’s two-dimensional decoupled architecture significantly improves the cost-effectiveness and robustness of training while maintaining high performance.

C Implementation Details of Quantitative Cross-Validation

To ensure transparency, we detail the methodology and the statistical characteristics of the validation results presented in Section 4.4.

Data Sampling and Inference Protocol. We randomly selected a validation subset of 1k video-instruction pairs from the Video-R1-260k dataset. We used the base model (Qwen2.5-VL-7B-Instruct) to perform zero-shot inference. For each sample, the response was scored as 1 (correct) or 0 (incorrect) based on ground-truth verification.

Binning Analysis and Statistical Observations. To visualize the correlation in Figure 4, we adopted a binning approach:

Horizontal Axis: We sorted samples by proxy scores (ϕ_{flow} , ϕ_{ent} , D_{visual} , and D_{text}) and applied Quantile Binning to create 20 equal-sized intervals.

Vertical Axis: We calculated the mean accuracy within each bin.

Observations on Variance: Unlike generation-based estimators that require multiple inference passes, our training-free proxies do not account for the model’s internal state. Consequently, Figure 4 reflects heteroscedasticity—the variance in

model performance remains significant even in low-difficulty bins. This suggests that while D_{text} identifies samples requiring deep logical deduction, failure can still occur on "easy" samples due to unrelated factors such as visual occlusions or linguistic ambiguity.

Interpretation of the Cognitive Proxy. The Calibrated Surprisal (D_{text}) measures the Information Gain required to reach the answer. The observed correlation confirms that as the logical entropy between the question and answer grows, the probability of the model maintaining a correct reasoning chain diminishes. By acknowledging the inherent noise in these proxies, VideoCuRL utilizes them as probabilistic guides rather than deterministic labels, ensuring a robust curriculum that is resilient to individual sample outliers.

D Sensitivity Analysis of Grid Size K

we investigate the impact of the grid resolution K on the overall performance of VideoCuRL. The parameter K determines the granularity of our orthogonal difficulty decomposition, where a $K \times K$ grid partitions the dataset into K^2 distinct buckets.

Experimental Setup. We conducted experiments with $K \in \{3, 4, 5\}$ while keeping all other hyperparameters, such as the KL coefficient $\beta = 0.04$ and total training steps, constant. All variants were trained on the Video-R1-260k dataset using the Diagonal Wavefront scheduling strategy. Performance was evaluated across three representative benchmarks: VSI-Bench (Reasoning), MVBench (Perception), and VideoMME (Long-context).

As shown in Table 3, increasing K from 3 to 4 yields a significant performance gain, particularly in reasoning tasks (+1.5% on VSI-Bench). This suggests that a 3×3 grid is too coarse to effectively disentangle the complex interplay between visual perception load and cognitive depth. While $K = 5$ provides marginal improvements over $K = 4$, the gains are statistically insignificant. However, a 5×5 grid increases the number of buckets from 16

to 25, leading to higher computational overhead in tracking Local Competence Scores (LCS) and managing the wavefront expansion. We observed that $K = 4$ provides the most stable training trajectory. With $K = 5$, some high-difficulty buckets (e.g., $B_{4,4}$) contain fewer samples, occasionally leading to increased reward variance and triggering the Dynamic Sparse KL mechanism more frequently than necessary.

Conclusion. Based on these findings, we selected $K = 4$ as the default configuration for VideoCuRL. It strikes an optimal balance between curriculum precision and optimization stability, ensuring robust spatiotemporal reasoning without excessive partitioning of the training data.