

# NeuS-QA: Grounding Long-Form Video Understanding in Temporal Logic and Neuro-Symbolic Reasoning

Sahil Shah<sup>1</sup>, S P Sharan<sup>\*1</sup>, Harsh Goel<sup>\*1</sup>, Minkyu Choi<sup>1</sup>, Mustafa Munir<sup>1</sup>,  
Manvik Pasula<sup>2</sup>, Radu Marculescu<sup>1</sup>, Sandeep Chinchali<sup>1</sup>

<sup>1</sup>The University of Texas at Austin, Texas, USA

<sup>2</sup>Independent Researcher, USA

{ss96869, spsharan, harshg99, minkyu.choi, mmunir}@utexas.edu

## Abstract

While vision-language models (VLMs) excel at tasks involving single images or short videos, they still struggle with Long Video Question Answering (LVQA) due to its demand for complex multi-step temporal reasoning. Vanilla approaches, which simply sample frames uniformly and feed them to a VLM along with the question, incur significant token overhead. This forces aggressive downsampling of long videos, causing models to miss fine-grained visual structure, subtle event transitions, and key temporal cues. Recent works attempt to overcome these limitations through heuristic approaches; however, they lack explicit mechanisms for encoding temporal relationships and fail to provide any formal guarantees that the sampled context actually encodes the compositional or causal logic required by the question. To address these foundational gaps, we introduce *NeuS-QA*, a training-free, plug-and-play neuro-symbolic pipeline for LVQA. *NeuS-QA* first translates a natural language question into a logic specification that models the temporal relationship between frame-level events. Next, we construct a video automaton to model the video’s frame-by-frame event progression, and finally employ model checking to compare the automaton against the specification to identify all video segments that satisfy the question’s logical requirements. Only these logic-verified segments are submitted to the VLM, thus improving interpretability, reducing hallucinations, and enabling compositional reasoning without modifying or fine-tuning the model. Experiments on the LongVideoBench and CinePile benchmarks show that *NeuS-QA* significantly improves performance by **over 10%**, particularly on questions involving event ordering, causality, and multi-step reasoning.

**Code** — <https://utaustin-swarmlab.github.io/NeuS-QA/>

## 1 Introduction

Recent advances in video-language models (VLMs) have led to significant progress in visual question answering (VQA) for static images and short video clips (Park et al. 2024; Wang et al. 2024b; Ye et al. 2025; Wang et al. 2025; Choudhury et al. 2023). However, as real-world applications shift toward longer-form content such as surveillance streams, egocentric vlogs, and movie scenes, users are increasingly asking more complex, temporally grounded questions that

span multiple events (Wu et al. 2024; Rawal et al. 2024; Choi et al. 2024; Wang et al. 2024a; Tan et al. 2025). For example, as illustrated in Figure 1, a user watching a one-hour travel vlog might ask: “*After the man boarded the plane and sat down, what did he grab?*”

Answering such questions demands more than surface-level perception; it requires **① semantic grounding** to *identify* entities and their interactions, **② temporal reasoning** to *track* and *interpret* the sequence of events, and **③ compositionality** to *integrate* these temporally distinct sub-events into a coherent reasoning chain that aligns with the structure and intent of the query.

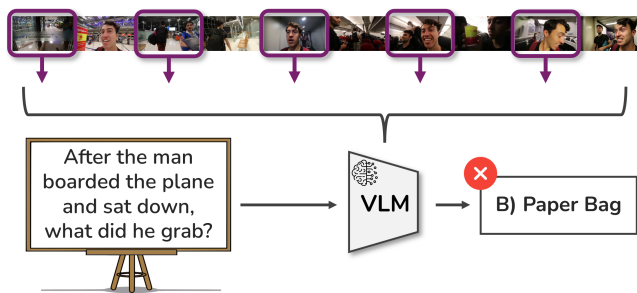
A common strategy in video question answering with VLMs is to uniformly sample frames across the video. While this method is sufficient for short clips, it becomes problematic as video length increases: the volume of visual input grows rapidly, forcing models to either exceed token limits or aggressively downsample videos. In doing so, they risk skipping large portions of the video and missing *fine-grained visual details, subtle event transitions, and key temporal cues* crucial for answering complex temporal queries.

Long-form videos often follow a *narrative arc*, with multiple scenes unfolding sequentially; yet many questions about them pertain to only a specific portion of that story. Figure 1 illustrates a vacation vlog that depicts a man’s full journey from entering the airport to de-boarding to finally checking into a hotel room. But the user’s query requires information solely from the in-flight segment. This is a common pattern in long-form video question answering (LVQA): although the content is broad, most queries target specific temporal windows.

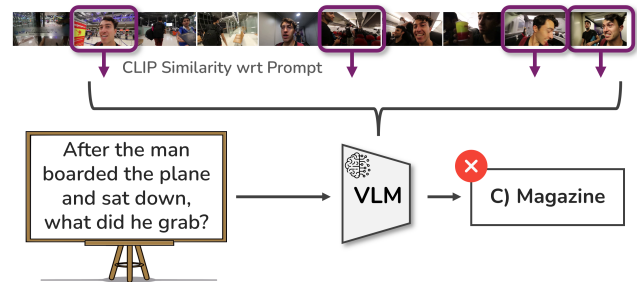
Therefore, a *natural alternative* to uniform sampling is to direct the model’s attention to a small, targeted segment of the video that is most relevant to the query. However, it raises a core challenge: *How can we retrieve the correct segment that is semantically and temporally aligned with the question, without already knowing the answer?*

To address this, we argue that temporal logic (TL) offers a powerful solution. It allows us to formalize *how* events unfold over time by symbolically representing their structure using operators such as “until,” “eventually,” and “always.” This makes it possible to precisely specify the kinds of temporal patterns a question refers to (like an event occurring

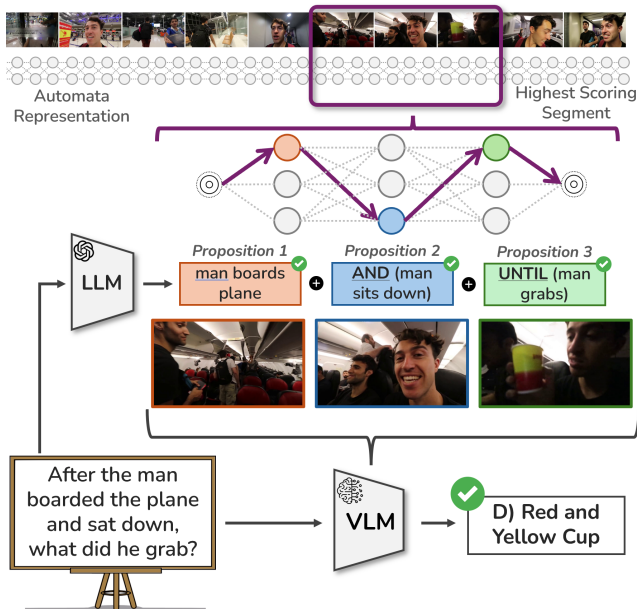
<sup>\*</sup>These authors contributed equally.



(a) **Vanilla VLM Prompting** uniformly samples frames across the video regardless of relevance, often missing key temporal context.



(b) **Heuristic Retrieval** ranks frames based on semantic similarity to the query but lacks temporal structure and formal guarantees.



(c) **Ours (*NeuS-QA*)** builds a video automaton over grounded propositions, constructs a temporal logic specification from the query, and uses model checking to retrieve logic-satisfying video segments leading to more accurate and interpretable answers.

Figure 1: Comparison of frame selection strategies to answer *temporally-grounded* questions over *long-form* videos.

after another), and to systematically filter for video segments that satisfy these constraints.

To this end, we propose *NeuS-QA*, a neuro-symbolic pipeline that combines the structure of TL with the perceptual capabilities of VLMs to accurately answer *temporally*

*complex questions over long-form videos*. Given a question in natural language, *NeuS-QA* first translates it into a TL specification that encodes the desired sequence of events. From the video, it then constructs a video automaton by assigning scores for semantic propositions in individual frames, and applies model checking to identify video segments that satisfy this specification. Finally, a VLM is queried on these verified segments, enabling precise, temporally grounded answering without exposing the model to irrelevant or misleading context.

## 2 Related Works

**Video Question Answering.** Early VQA systems (Antol et al. 2015) typically relied on fully supervised architectures that encoded entire videos using convolutional and recurrent networks, encoder-decoder pipelines, or attention mechanisms over video features (Zhu et al. 2016; Ye et al. 2017; Jang et al. 2017). However, these models lacked scalability to longer videos and were constrained by limited visual-linguistic reasoning. The emergence of VLMs (Bai et al. 2025; OpenGVLab 2024; Achiam et al. 2023; Li et al. 2024a,b) has enabled zero-shot VQA across diverse domains with stronger generalization. However, these models typically operate by consuming a fixed set of frames sampled from the entire video, limiting their ability to handle long-form videos with complex temporal structures. Their reliance on global frame pooling or simple attention mechanisms often leads to irrelevant or diluted context when presented with multi-event sequences (Wu et al. 2024; Wang et al. 2024a), resulting in degraded performance.

To address the limitations of global frame sampling, recent state-of-the-art (SOTA) approaches adopt retrieval-augmented strategies that identify and condition on the most relevant video segments (Park et al. 2024; Wang et al. 2024b; Ye et al. 2025; Song et al. 2024; Wang et al. 2025; Islam et al. 2025; Choudhury et al. 2023). While effective at reducing irrelevant context, these methods often rely on coarse heuristics or textual summarization, which sacrifice visual fidelity and lack formal structure. In contrast, *NeuS-QA* uses temporal logic to guide frame selection, yielding a more interpretable and precise grounding of video content.

**Symbolic Representations for Video.** Symbolic representations have been explored across fields such as robotics (Puranic, Deshmukh, and Nikolaidis 2021; Hasanbeig et al. 2019; Shoukry et al. 2017), autonomous vehicles (Zheng et al. 2025; Mehdipour et al. 2023), generative models (Shah et al. 2025; Sharan et al. 2025; Choi et al. 2025), and structured neural networks (Manginas, Paliouras, and De Raedt 2024; Hashemi et al. 2023). In the context of vision, symbolic methods have been employed to support long-form understanding (Tapaswi et al. 2016; Huang et al. 2020; Xiong et al. 2019; Wu and Krähenbühl 2021); however, these typically rely on latent-space embeddings, sacrificing interpretability. Yang et al. (2023) and Choi et al. (2024) move toward representing videos as formal models to enable structured reasoning. *NeuS-QA* follows a similar philosophy, modeling videos as verifiable automata to enable interpretable and temporally grounded analysis.

### 3 Preliminaries

To illustrate our approach, we begin by introducing a running example. Suppose, given a video, we ask “*After the woman pours hot water over granola and spoons yogurt into the bowl, what does she place as a topping?*”

**Temporal Logic.** Temporal Logic (TL) is a formal language used to describe the progression of events through a combination of logical and temporal operators. In TL, individual events are represented as atomic propositions, statements that evaluate to either `True` or `False`. These propositions are combined using logical operators such as AND ( $\wedge$ ), OR ( $\vee$ ), NOT ( $\neg$ ), and IMPLY ( $\implies$ ), as well as temporal operators like ALWAYS ( $\square$ ), EVENTUALLY ( $\diamond$ ), NEXT ( $X$ ), and UNTIL ( $U$ ).

Returning to our running example, we define a set of atomic propositions  $\mathcal{P}$  and a corresponding TL specification  $\varphi$  as follows:

$$\begin{aligned} \mathcal{P} = \{ & \text{woman pours hot water over granola,} \\ & \text{woman spoons yogurt into bowl,} \\ & \text{woman places topping} \}, \\ \varphi = & (\text{woman pours hot water over granola} \wedge \\ & \text{woman spoons yogurt into bowl}) \wedge \diamond \\ & \text{woman places topping.} \end{aligned}$$

This specification encodes the query where the woman first pours hot water over granola, *and* adds yogurt to the bowl, and *eventually* adds some topping at the end.

**Video Automaton.** We represent a video as a stochastic finite-state automaton where each state corresponds to a sampled frame, and transitions capture temporal progression. As videos are a finite sequence of discrete frames with inherent temporal structure, we model them as a discrete-time Markov chain (Norris 1998; Kemeny and Snell 1960).

Formally, the video automaton  $\mathcal{A}_V$  is defined as the tuple:

$$\mathcal{A}_V = (Q, q_0, \delta, \lambda), \quad (1)$$

where  $Q$  is a finite set of states,  $q_0 \in Q$  is the initial state,  $\delta : Q \times Q \rightarrow [0, 1]$  is a transition function such that  $\delta(q_i, q_j)$  gives the probability of transitioning from state  $q_i$  to state  $q_j$ , and  $\lambda : Q \rightarrow 2^{|\mathcal{P}|}$  is a labeling function that assigns a set of atomic propositions to each state. This structure allows TL formulas to be verified over the video by treating sequences of frames as labeled paths.

**Formal Verification.** With formal verification, we assess whether a video automaton satisfies a given TL specification (Clarke, Grumberg, and Peled 1999; Huth and Ryan 2004). We define a path  $\pi$  through the automaton  $\mathcal{A}_V$  as an infinite sequence of states  $\pi = q_0 q_1 q_2 \dots$  such that  $\delta(q_i, q_{i+1}) > 0$  for all  $i \geq 0$ . Each path induces a trace—the sequence of label sets assigned to the states along the path, denoted as  $\text{trace}(\pi) = \lambda(q_0)\lambda(q_1)\lambda(q_2)\dots$ . This trace captures the progression of observable propositions over time, as illustrated in Figure 2. We apply probabilistic model checking (Baier and Katoen 2008) to compute the satisfaction probability  $\mathbb{P}[V \models \varphi]$ , quantifying how well the video trace satisfies the TL specification. This provides a formal and interpretable framework for identifying event sequences in video.

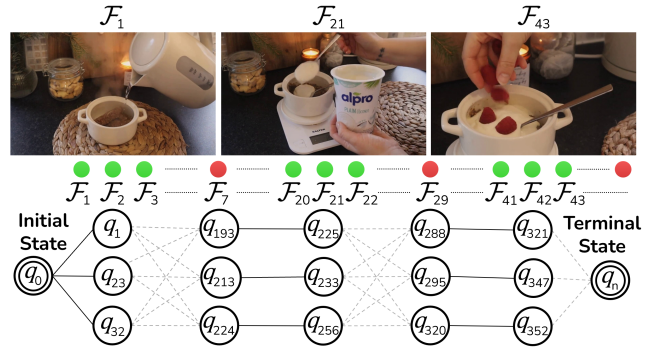


Figure 2: **Video automaton for the running example.** Top: representative frames  $\mathcal{F}_1$  (*woman pours hot water over granola*),  $\mathcal{F}_{21}$  (*woman spoons yogurt into bowl*), and  $\mathcal{F}_{43}$  (*woman places topping*). Bottom: finite-state automaton synthesized from the TL specification in  $\varphi$  of the running example. **Green circles** indicate *valid frames*—those in which at least one domain-specific atomic proposition (e.g., “woman pours hot water over granola”) is detected and whose inclusion preserves all ordering constraints accumulated *so far*. **Red circles** mark frames that (i) contain no relevant propositions or (ii) would force a transition contradicting the specification (for instance, “woman grabs spoon but doesn’t spoon yogurt into bowl”). We verify the automaton incrementally: whenever a new frame is labelled green, we extend the current run  $\rho = q_0 \rightarrow \dots \rightarrow q_t$  and immediately re-evaluate the property  $\rho \models \varphi$ . If the property still holds, the frame remains green; otherwise, it is re-labelled red and excluded. The procedure terminates once the accepting state  $q_n$  is reached, guaranteeing that the sequence of green frames is a witness trace satisfying the specification.

### 4 Methodology

We present an overview of *Neus-QA*, which extracts temporally grounded scenes of interest from the video based on a natural language query to guide vision-language models. *Neus-QA* first identifies the key events and semantics necessary to answer the question. These are then formalized as a temporal logic specification that encodes the expected structure and ordering of events. To evaluate which parts of the video satisfy this specification, we construct an automaton whose states represent frame-level propositions predicted by a VLM. As the VLM iteratively processes consecutive video segments, we label frames with relevant propositions and incrementally build the automaton. Formal verification is then applied over this automaton to identify sequences that satisfy the query, yielding the final set of frames of interest. We outline this process in Algorithm 1.

**$\mathcal{L}_{Q2TL}$ : Question to Temporal Logic.**  $\mathcal{L}_{Q2TL}$  provides a structured two-shot framework for translating natural language questions into TL specifications using prompting strategies that capture both spatial and temporal relationships. Given a question  $\mathbf{q}$ ,  $\mathcal{L}_{Q2TL}$  extracts a set of atomic propositions  $\mathcal{P} = \{p_1, p_2, \dots, p_n\}$  corresponding to key objects, actions, and relationships in the query. This extraction is guided by optimized prompts with carefully selected few-

---

**Algorithm 1: NeuS-QA Core Logic**

---

**Require:** Proposition extractor  $\mathcal{L}_{\text{Q2TL}}(\cdot)$ , frame-level detector  $\mathcal{M}_{\text{prop}}(\cdot)$ , frame step size  $\kappa$ , automaton builder  $\Theta(\cdot)$ , probabilistic model checker  $\mu(\cdot)$ , smoothing function  $\hat{F}(\cdot)$ , temporal expander  $\mathcal{L}_{\text{ext}}(\cdot)$ , trimming function  $f_{\text{trim}}$ , vision-language answerer  $\mathcal{M}_{\text{VQA}}(\cdot)$ , satisfaction threshold  $\tau$

**Input :** Video  $\mathcal{V}$ , question  $\mathbf{q}$ , answer choices  $\mathcal{C}$ , ground-truth answer  $c^*$

**Output :** Boolean result  $y \in \{\text{True}, \text{False}\}$

```
1  $\mathcal{P}, \varphi \leftarrow \mathcal{L}_{\text{Q2TL}}(\mathbf{q})$  // Question to TL
2  $Z \leftarrow \emptyset$  // Create detection matrix
3  $\mathcal{A}_{\mathcal{V},0} \leftarrow \Theta(\emptyset, \mathcal{P}, Z)$  // Init automaton
4 for  $t \leftarrow 0$  to  $\text{length}(\mathcal{V}) - \kappa$  do
5    $\mathcal{F}_t \leftarrow \{\mathcal{V}_t, \dots, \mathcal{V}_{t+\kappa}\}$  // Extract frames
6   for  $p_i \in \mathcal{P}$  do
7      $Z_{t,p_i} \leftarrow \mathcal{M}_{\text{prop}}(p_i, \mathcal{F}_t)$  // Find
      frames by proposition
8    $\mathcal{A}_{\mathcal{V},t} \leftarrow \Theta(\mathcal{A}_{\mathcal{V},t-1}, \mathcal{P}, Z_{0:t})$  // Extend
      automaton
9    $\mathbb{P}[\mathcal{A}_{\mathcal{V},t} \models \varphi] \leftarrow \mu(\mathcal{V}, \varphi)$  // Compute
      satisfaction probability
10  if  $\hat{F}(\mathbb{P}[\mathcal{A}_{\mathcal{V},t} \models \varphi]) > \tau$  then
11    break // Calibrated probability
      is above threshold
12  $(t_{\text{start}}, t_{\text{end}}) \leftarrow \mathcal{V}_{[\text{satisfy}]}$  // Obtain
      automaton frames of interest
13  $(\alpha, \beta) \leftarrow \mathcal{L}_{\text{ext}}(\mathbf{q})$  // Find extensions
14  $\mathcal{V}' \leftarrow f_{\text{trim}}(\mathcal{V}, t_{\text{start}} + \alpha, t_{\text{end}} + \beta)$  // Trim to
      extended frames of interest
15  $\hat{c} \leftarrow \mathcal{M}_{\text{VQA}}(\mathbf{q}, \mathcal{C}, \mathcal{V}')$  // Compute VQA
16  $y \leftarrow (\hat{c} = c^*)$  // Determine match
```

---

shot examples to ensure semantic and event-level coverage.

After the propositions are identified,  $\mathcal{L}_{\text{Q2TL}}$  constructs a TL specification  $\varphi$  by mapping the temporal structure of the query onto formal operators, linking the extracted propositions accordingly. The resulting propositions and specification are detailed in Section 3.

**Automaton Representation of Video.** Given a TL specification  $\varphi$  from  $\mathcal{L}_{\text{Q2TL}}$ , we construct a formal representation of the video by incrementally building a discrete-time Markov chain automaton  $\mathcal{A}_{\mathcal{V}}$  grounded in event-based atomic propositions. Starting from an empty automaton  $\mathcal{A}_{\mathcal{V},0}$  (L3), we process the video in frame windows  $\mathcal{F}_t$  of size  $\kappa$  (L5), building a probabilistic state-based model that captures the evolution of identified events over time.

For each frame sequence  $\mathcal{F}_t$ , the system computes probabilistic detections  $Z_{t,p_i} \in [0, 1]$ , indicating the likelihood that a proposition  $p_i \in \mathcal{P}$  is present, using a vision-language model  $\mathcal{M}_{\text{prop}}(p_i, \mathcal{F}_t)$  (L7). These values are derived from the model’s Yes/No logits, calibrated using a held-out validation set (see Appendix) to improve the reliability of semantic assessments. Based on the cumulative detections  $Z_{0:t}$  and the current automaton state  $\mathcal{A}_{\mathcal{V},t-1}$ , we update the automaton via the transition builder  $\Theta(\mathcal{A}_{\mathcal{V},t-1}, \mathcal{P}, Z_{0:t})$  (L8). This automaton incrementally models the temporal progres-

sion of event-based propositions throughout the video.

**Formal Verification for Scenes of Interest.** We use formal verification to identify all time intervals where the TL specification  $\varphi$  is satisfied. This process employs established model-checking techniques to enable rigorous temporal analysis over the video automaton (L9).

At each timestep  $t$ , given the current automaton  $\mathcal{A}_{\mathcal{V},t}$ , we compute the satisfaction probability as:

$$\mathbb{P}_t = \mathbb{P}[\mathcal{A}_{\mathcal{V},t} \models \varphi] = \mu(\mathcal{V}, \varphi). \quad (2)$$

This satisfaction probability is computed using probabilistic model checking, implemented via the STORM model checker (Hensel et al. 2020; Junges and Volk 2021). This process operates under the Probabilistic Computation Tree Logic (PCTL) framework, where  $\mu(\cdot)$  analyzes the state labels and transition probabilities in  $\mathcal{A}_{\mathcal{V},t}$  to assess how likely the specification  $\varphi$  is satisfied.

To obtain a stable and calibrated signal, we apply a smoothing function  $\hat{F}(\cdot)$  that maps raw satisfaction probabilities into per-frame detection scores, defined as:

$$\hat{F}(c) = (1 + e^{-\gamma(c-\tau)})^{-1}, \quad (3)$$

where  $c$  is the raw satisfaction probability,  $\tau$  is the confidence threshold, and  $\gamma$  controls the steepness of the sigmoid curve. This transformation suppresses noise below a defined confidence threshold while preserving high-confidence detections. By modulating the steepness of the transition around the threshold  $\tau$ , the function produces a smooth, interpretable confidence curve. The output yields a minimal satisfying interval (e.g., frames 615  $\rightarrow$  820) that captures query-relevant events with high temporal precision.

**Temporal Extension.** We extend the identified interval to ensure a complete temporal context for robust question answering (L13). The context window extension process extends the interval forward and backward by a fixed or learned temporal window  $\alpha, \beta$ , ensuring that sufficient temporal context is preserved for comprehensive understanding. We leverage the VLM to account for event lead-up and aftermath periods following temporal keywords such as “before” and “after” in the query  $\mathbf{q}$ . The final extended window becomes the filtered input to the VLM for question answering, providing a temporally focused, yet contextually complete video segment for optimal performance. The temporal extension is formalized in L14.

**Vision Language Model Answering.** The final stage uses this temporally filtered video segment as context to generate accurate responses through VLM prompting.

## 5 Experimental Setup

During automaton construction, we use InternVL2-8B (OpenGVLab 2024) as the VLM to generate frame-level detections at 3 frames per second. To translate natural language questions into TL specifications, we use GPT-o1-mini (Achiam et al. 2023) with a few-shot prompting strategy and a carefully designed system prompt to capture common temporal structures and patterns. We apply Stormpy (Hensel

	Method	Frames	Category				Overall (↑)
			T3E	E3E	T3O	O3O	
Vision Foundation Models	LLaVA-OneVision-Qwen2-7B (Li et al. 2024a)	32	43.40	64.86	54.39	43.18	53.07
	Aria (Li et al. 2024b)	64	52.83	58.11	59.65	36.36	53.07
	Qwen2.5-VL-7B-Instruct (Bai et al. 2025)	32	39.62	58.11	54.39	45.46	50.44
	GPT-4o-2024-08-06 (Achiam et al. 2023)	32	43.40	58.11	47.37	45.45	49.56
	Qwen2-VL-7B-Instruct (Yang et al. 2024)	32	45.28	54.05	45.61	47.73	48.68
	InternVL2-8B (OpenGVLab 2024)	32	39.62	48.65	49.12	43.18	45.61
Structured Reasoning Frameworks	VideoTree (Wang et al. 2025)	8-32	<b>56.38</b>	50.00	50.00	43.84	50.47
	LVNet (Park et al. 2024)	12	41.94	54.17	35.48	46.15	45.59
	VideoAgent (Wang et al. 2024b)	32	45.45	52.63	27.27	27.27	40.38
	T* (Ye et al. 2025)	8	28.95	35.56	48.33	22.22	35.75
<i>NeuS-QA</i> (Ours)	<i>NeuS-QA</i> + Qwen2.5-VL-7B-Instruct	32	54.72	62.16	<b>66.67</b>	<b>54.55</b>	<b>60.09</b>
	<i>NeuS-QA</i> + Aria	32	54.72	64.86	63.16	47.73	58.77
	<i>NeuS-QA</i> + LLaVA-OneVision-Qwen2-7B-ov	32	50.94	<b>70.27</b>	56.14	52.27	58.77
	<i>NeuS-QA</i> + GPT-4o-2024-08-06	32	49.06	68.92	59.65	50.00	58.33

Table 1: **Benchmarking Accuracy (%) of SOTA VQA Models.** *NeuS-QA*, when used alongside different VLM backbones, outperforms raw foundation models and structured reasoning frameworks on LongVideoBench.

et al. 2020; Junges and Volk 2021), a probabilistic model checker, to identify video segments that satisfy these TL specifications. We present the results of *NeuS-QA* as a plug-and-play module with various VLMs in Section 6.

**Datasets.** For empirical evaluation of *NeuS-QA* we use two LVQA datasets: LongVideoBench (Wu et al. 2024) and CinePile (Rawal et al. 2024). LongVideoBench serves as our primary benchmark with videos up to 60 minutes long across diverse domains such as instructional content, egocentric vlogs, and educational streams. CinePile is used as a secondary dataset to demonstrate the performance of *NeuS-QA* on long-form movie scenes that require reasoning about narrative structure and character interaction. For both datasets, we focus on *subsets* of question-video pairs that require *temporal ordering*, *causal inference*, and *compositional logic*, filtering out simple examples based solely on static visual recognition. In LongVideoBench, these reasoning-oriented queries correspond to four categories: **T3E** (event before/after text), **E3E** (event before/after event), **T3O** (object before/after text), and **O3O** (object before/after object). To streamline the evaluation, we burn the provided subtitles into the video, allowing *NeuS-QA* and consequently any VLM to directly infer and answer questions involving dialogue without requiring separate time-aligned subtitles.

**Evaluation Metrics and Benchmarking.** We evaluate *NeuS-QA* on multiple-choice LVQA tasks using *accuracy* as the primary metric across both LongVideoBench and CinePile. Since *NeuS-QA* serves as a plug-and-play key-segment retrieval layer, we apply it modularly on top of several state-of-the-art VLMs, including Qwen2.5 (Bai et al. 2025), GPT-4o (Achiam et al. 2023), and LLaVA-OneVision (Li et al. 2024a). For each model, we report performance with and without our retrieval pipeline to measure the impact of symbolic retrieval. In addition to these baselines, we compare against recent structured LVQA methods such as T\*, VideoTree, and LVNet (Ye et al. 2025; Wang et al. 2025; Park et al. 2024), as discussed in Section 2. All benchmarks

are conducted using the `lmms-eval` library (Zhang et al. 2024), which provides a unified framework for evaluating LVQA models across datasets.

## 6 Results

Building on our experimental setup, we now turn to empirical evaluation. Our goal is to answer the following key questions regarding neuro-symbolic LVQA:

1. To what extent does *NeuS-QA* improve VQA performance on significantly longer videos?
2. Do the frames selected by *NeuS-QA* add more relevant context that improves VQA performance?
3. How crucial is the neurosymbolic reasoning component in *NeuS-QA* for effective long video understanding?

### 6.1 Neuro-Symbolic LVQA

**Evaluations on LongVideoBench.** We evaluate *NeuS-QA* on LongVideoBench by comparing it against both VLMs and structured reasoning baselines. As shown in Table 1, *NeuS-QA* improves VQA accuracy across all question types. Paired with Qwen2.5-VL, it achieves 60.09% overall accuracy—nearly 10% over the base model and outperforms all structured reasoning baselines, including VideoTree and LVNet. Qualitative examples in Figure 4 further illustrate *NeuS-QA*’s ability to identify the correct frames of interest. These results highlight the benefits of symbolic temporal reasoning for long video understanding and demonstrate that dynamically selecting video segments based on the initial query is more effective than fixed or heuristic approaches.

We also assess scalability with increasing video duration as shown in Figure 3, performance remains stable or improves, whereas prompting-based methods show substantial degradation. Notably, our method even outperforms human-annotated ground truth segments, achieving 59% vs. 52% when both are used as input to the same VLM. We attribute this to the automaton’s ability to isolate not only the direct answer span but also semantically and temporally relevant context (e.g., events before/after).

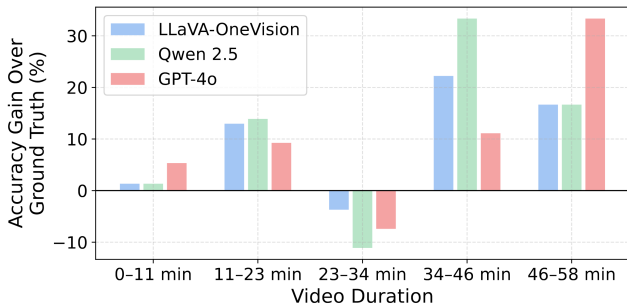


Figure 3: **Accuracy (%) gains of *NeuS-QA* over ground truth frame annotations.** We visualize the per-model improvement of *NeuS-QA* in terms of absolute accuracy difference relative to ground truth annotations across different video durations. Positive values indicate that *NeuS-QA* outperforms ground truth frame annotations.

Strategy		Overall (↑)
Vision Foundation Models	Qwen2.5-VL-7B-Instruct	50.73
	GPT-4o-2024-08-06	50.24
	Llama-3.2-11B-Vision-Instruct	20.98
	LLaVA-OneVision-Qwen2-7B	31.71
	Idefics2-8B	25.85
<i>NeuS-QA</i> (Ours)	Qwen2.5-VL-7B-Instruct	<b>53.66</b>

Table 2: **Generalization to Narrative Video.** *NeuS-QA* achieves the highest accuracy (%) on CinePile, demonstrating its effectiveness in abstract, story-driven video domains.

**Performance on CinePile.** We evaluate *NeuS-QA* on CinePile, a long-video benchmark designed for narrative-driven understanding. Unlike instructional or surveillance-style datasets, CinePile emphasizes abstract, story-based reasoning. As shown in Table 2, *NeuS-QA* achieves 53.66% accuracy when paired with Qwen2.5-VL, yielding notable gains over zero-shot prompting alone. These results indicate that our method generalizes effectively to more complex video domains, demonstrating its ability to identify contextually relevant content even without task-specific tuning.

## 6.2 Ablation on *NeuS-QA* Architectural Decisions

**Frame Presentation Strategy.** Frame presentation affects how much context the model can leverage when answering questions. We compare the default *NeuS-QA* approach of trimming the video to only the most relevant segments against alternative approaches that retain full temporal context while varying frame sampling density, sampling densely within relevant intervals and sparsely elsewhere. As shown in Table 4, these adaptive strategies recover some incorrect answers but result in lower overall accuracy. This suggests that while context retention offers some benefit, tightly focusing on relevant segments remains more effective.

**Impact of Automata VLM Choice.** The choice of vision-language model used to identify atomic propositions during automaton construction directly affects the quality of the extracted frames of interest. We compare InternVL2-8B and Qwen2.5-VL for this step, keeping the downstream

Choice of VLM	Median Clip Size (s)	Accuracy (↑)
Qwen2.5-VL-7B-Instruct	40	52.63
InternVL2-8B	40	<b>58.33</b>

Table 3: **Impact of VLM backbone choice on accuracy (%).** Using InternVL2-8B for automaton construction results in the highest accuracy when evaluated with GPT-4o.

Strategy	Accuracy (↑)	Improvements (↑)	Regressions (↓)
Original (Full)	49.56	–	–
<i>NeuS-QA</i> (Trim)	<b>58.33</b>	<b>18.42</b>	9.65
Dynamic Scaling	53.51	12.28	8.33
Speed Bias	52.19	10.53	<b>7.90</b>

Table 4: **Impact of frame sampling strategy.** Trimming to relevant segments with *NeuS-QA* yields the highest accuracy (+9.77%) over the full-video baseline, driven by a substantial increase in correct predictions (+18.42%) and a moderate rise in regressions (+9.65%). Dynamic scaling and speed bias strategies offer intermediate trade-offs, illustrating the balance between context retention and focus.

VQA model fixed. As shown in Table 3, InternVL2-8B leads to higher VQA accuracy, even though both models produce similar median clip lengths. This suggests that stronger VLMs that more accurately identify propositions yield more faithful automatons, improving the precision and relevance of the frames selected for reasoning.

## 6.3 How Important is Neuro-Symbolic Feedback on Frame Selection?

To isolate the impact of formal reasoning, we evaluate a simplified variant of *NeuS-QA* that skips temporal logic and automaton construction. Instead, it relies *solely* on prompting a VLM to judge whether a segment is relevant to the question. As shown in Table 5, this approach underperforms significantly: without structured semantics or temporal constraints, **frame selection becomes noisy** and accuracy drops. These findings underscore the importance of neuro-symbolic feedback in guiding segment retrieval with interpretable, temporally grounded structure.


## 7 Discussion

**Applicability Beyond Explicit Temporal Prompts.** While *NeuS-QA* is designed to handle complex temporal queries, its applicability extends even to cases where temporal structure is implicit. Since we use an LLM to translate questions into temporal logic, it can infer event order from linguistic cues, even if the prompt *does not* contain explicit temporal operators such as “before” or “after.” In simpler or ambiguous queries, the system effectively treats the entire prompt as a single proposition, reducing the reasoning task to identifying the most relevant clip. In such degenerate cases, performance aligns with standard VLM-based frame retrieval approaches, as shown in Section 6.3. In Table 1, we focus on four temporally compositional categories from LongVideoBench where our approach offers the greatest benefits. Results for other categories (see Appendix) also show improvements due to our logic-based retrieval, though

Q) In a white room, there is a glass door with a black frame and a man wearing a green coat and earphones. After he says 'story', who is the first person to appear?

NeuS-QA & VideoTree

- A) A woman in a blue shirt  
 B) A man in glasses and a black top ✓  
 C) A woman with black hair tied up ✗  
 D) A woman in a blue top  
 E) A man with dreadlocks



Q) There are instruments on the shelves of a laboratory. A man with black hair, wearing a gray vest, picks up a blue cup from the shelf. What does he do after?

NeuS-QA & VideoTree

- A) Kneeled down  
 B) Talked while turning around ✓  
 C) Drank water  
 D) Ate something  
 E) Talked while running ✗




Figure 4: **Qualitative Examples from the NeuS-QA pipeline.** Compared to other structured reasoning frameworks, *NeuS-QA* more accurately identifies the correct frames of interest in a long video.

Metric	NeuS-QA (Ours)	VLM-only
Accuracy (0s, 600s]	<b>61.33</b>	42.67
Accuracy (600s, 3600s]	<b>56.86</b>	38.82
Overall Accuracy (↑)	<b>58.33</b>	40.09

Table 5: **Effect of neuro-symbolic feedback on frame selection.** Replacing temporal logic and automaton construction with direct VLM queries significantly degrades accuracy (%) across both short and long video segments.

gains are less pronounced due to the reduced temporal complexity of those queries.

**Choice of Temporal Logic Formalism.** Another natural question is: “Why temporal logic?” Is it the only or best choice for structuring complex queries? We adopt TL because it strikes a balance between expressivity and compatibility with existing formal verification tools such as Storm and CARL (Hensel et al. 2020; Kremer et al. 2025). TL captures temporal constructs like ordering, simultaneity, and eventuality while remaining amenable to model checking. However, *NeuS-QA* is **logic-agnostic** too: alternatives like Linear Temporal Logic (LTL), Signal Temporal Logic (STL), or *any* domain-specific structured representation could be substituted, as long as they provide a mechanism to create and verify the logical conditions of a query.

This flexibility makes *NeuS-QA* a **general framework** for interpretable long-video reasoning. Our contribution is **not** about choosing these *design* and *architectural* choices, but about enabling **symbolic alignment** between video content and query structure as a strategy for aiding LVQA.

**Limitations and Failure Modes.** A key limitation of *NeuS-QA* lies in the coupling between its “neuro” and “symbolic” components. Although the symbolic verifier is *exact* by construction (guaranteeing correctness once propositions are grounded), the quality of the output depends heavily on the VLM’s ability to *detect* relevant events. In cases with multiple complex propositions, a single missed detection can cause the TL specification to be evaluated as unsatisfied, leading to false negatives during clip retrieval. These

failure cases are most common when visual cues are subtle, occluded, or appear briefly in the video.

**Future Work.** Looking forward, we see opportunities for extending *NeuS-QA* beyond one-shot retrieval and VQA. One direction is to explore *agentic* neuro-symbolic systems that reason over long videos *iteratively* using memory, planning, and intermediate logic representations to decompose and refine queries. Additionally, while our video automaton representation scales well with video length, constructing it remains computationally expensive. In theory, this cost could be amortized by pre-building a general-purpose video automaton and applying it to multiple downstream queries in a coarse-to-fine retrieval loop.

**Broader Implications.** While *NeuS-QA* achieves strong improvements and sets new state-of-the-art results on LVQA benchmarks (Table 1), overall scores including those from VLMs prompted with ground truth clips (Figure 3) remain modest. This highlights the broader limitations of current perception models. As VLMs *continue to improve*, the quality of both frame-level grounding and automaton construction *will scale accordingly*.

## 8 Conclusion

We introduce *NeuS-QA*, a plug-and-play, training-free neuro-symbolic framework for long-form video question answering. By translating natural language queries into temporal logic and verifying them over a video automaton, *NeuS-QA* isolates segments semantically and temporally aligned with the question before querying a VLM. Our results show strong gains over standard prompting and recent retrieval-based methods (Wang et al. 2025; Park et al. 2024), especially on temporally complex benchmarks like LongVideoBench and CinePile (Wu et al. 2024; Rawal et al. 2024). As VLMs improve, the fidelity of our symbolic alignment will scale in tandem. *NeuS-QA* is a general framework for structured, interpretable long-video reasoning. The key insight lies not in any particular system design or architecture, but symbolically aligning video content with queries as an effective and principled approach to LVQA.

## Acknowledgments

This material is based upon work supported in part by the Office of Naval Research (ONR) under Grant No. N00014-22-1-2254. Additionally, this work was supported by the Defense Advanced Research Projects Agency (DARPA) contract DARPA ANSR: RTX CW2231110. Approved for Public Release, Distribution Unlimited.

## References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Zitnick, C. L.; and Parikh, D. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, 2425–2433.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Baier, C.; and Katoen, J.-P. 2008. *Principles of Model Checking*. The MIT Press.
- Choi, M.; Goel, H.; Omama, M.; Yang, Y.; Shah, S.; and Chinchali, S. 2024. Towards neuro-symbolic video understanding. In *European Conference on Computer Vision*, 220–236. Springer.
- Choi, M.; Sharan, S.; Goel, H.; Shah, S.; and Chinchali, S. 2025. We'll Fix it in Post: Improving Text-to-Video Generation with Neuro-Symbolic Feedback. *arXiv preprint arXiv:2504.17180*.
- Choudhury, R.; Niinuma, K.; Kitani, K. M.; and Jeni, L. A. 2023. Zero-shot video question answering with procedural programs. *arXiv preprint arXiv:2312.00937*.
- Clarke, E. M.; Grumberg, O.; and Peled, D. A. 1999. *Model Checking*. MIT Press.
- Hasanbeig, M.; Kantaros, Y.; Abate, A.; Kroening, D.; Pappas, G. J.; and Lee, I. 2019. Reinforcement learning for temporal logic control synthesis with probabilistic satisfaction guarantees. In *2019 IEEE 58th conference on decision and control (CDC)*, 5338–5343. IEEE.
- Hashemi, N.; Hoxha, B.; Yamaguchi, T.; Prokhorov, D.; Fainekos, G.; and Deshmukh, J. 2023. A neurosymbolic approach to the verification of temporal logic properties of learning-enabled control systems. In *Proceedings of the ACM/IEEE 14th International Conference on Cyber-Physical Systems (with CPS-IoT Week 2023)*, 98–109.
- Hensel, C.; Junges, S.; Katoen, J.; Quatmann, T.; and Volk, M. 2020. The Probabilistic Model Checker Storm. *CoRR*, abs/2002.07080.
- Huang, Q.; Xiong, Y.; Rao, A.; Wang, J.; and Lin, D. 2020. MovieNet: A Holistic Dataset for Movie Understanding. In Vedaldi, A.; Bischof, H.; Brox, T.; and Frahm, J., eds., *European Conference on Computer Vision*, volume 12349 of *Lecture Notes in Computer Science*, 709–727. Springer.
- Huth, M.; and Ryan, M. 2004. *Logic in Computer Science: Modelling and Reasoning about Systems*. Cambridge University Press.
- Islam, M. M.; Nagarajan, T.; Wang, H.; Bertasius, G.; and Torresani, L. 2025. Bimba: Selective-scan compression for long-range video question answering. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 29096–29107.
- Jang, Y.; Song, Y.; Yu, Y.; Kim, Y.; and Kim, G. 2017. Tgif-qa: Toward spatio-temporal reasoning in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2758–2766.
- Junges, S.; and Volk, M. 2021. Stormpy - Python bindings for Storm.
- Kemeny, J.; and Snell, J. 1960. *Finite Markov Chains*. Springer.
- Kremer, G.; Volk, M.; Junges, S.; Schupp, S.; Corzilius, F.; Winkler, T.; Mazzinatus; Bene; Netz; Krüger, A.; Hensel, C.; MartaGr; Neuhaus, J.; Quatmann, T.; Berger, J.; Janson, T.; Klein, J.; glatteis; and Raoul. 2025. CARL-Storm: Computer Arithmetic and Logic Library for the Storm Model Checker, version 14.30. Software archived on Zenodo; published June 13, 2025.
- Li, B.; Zhang, Y.; Guo, D.; Zhang, R.; Li, F.; Zhang, H.; Zhang, K.; Zhang, P.; Li, Y.; Liu, Z.; et al. 2024a. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*.
- Li, D.; Liu, Y.; Wu, H.; Wang, Y.; Shen, Z.; Qu, B.; Niu, X.; Zhou, F.; Huang, C.; Li, Y.; et al. 2024b. Aria: An open multimodal native mixture-of-experts model. *arXiv preprint arXiv:2410.05993*.
- Manginas, N.; Paliouras, G.; and De Raedt, L. 2024. NeSyA: Neurosymbolic Automata. *arXiv preprint arXiv:2412.07331*.
- Mehdipour, N.; Althoff, M.; Tebbens, R. D.; and Belta, C. 2023. Formal methods to comply with rules of the road in autonomous driving: State of the art and grand challenges. *Automatica*, 152: 110692.
- Norris, J. 1998. *Markov Chains*. Cambridge University Press.
- OpenGVLab. 2024. InternVL 2.0: A Suite of Multimodal Large Language Models for Vision and Language Tasks. *OpenGVLab Blog*. Available at <https://internvl.github.io/blog/2024-07-02-InternVL-2.0/>.
- Park, J.; Ranasinghe, K.; Kahatapitiya, K.; Ryu, W.; Kim, D.; and Ryoo, M. S. 2024. Too many frames, not all useful: Efficient strategies for long-form video qa. *arXiv preprint arXiv:2406.09396*.
- Puranic, A.; Deshmukh, J.; and Nikolaidis, S. 2021. Learning from demonstrations using signal temporal logic. In *Conference on Robot Learning*, 2228–2242. PMLR.
- Rawal, R.; Saifullah, K.; Farré, M.; Basri, R.; Jacobs, D.; Somepalli, G.; and Goldstein, T. 2024. Cinepile: A long video question answering dataset and benchmark. *arXiv preprint arXiv:2405.08813*.

- Shah, S.; Goel, H.; Narasimhan, S. S.; Choi, M.; Sharan, S.; Akcin, O.; and Chinchali, S. 2025. A Challenge to Build Neuro-Symbolic Video Agents. *arXiv preprint arXiv:2505.13851*.
- Sharan, S.; Choi, M.; Shah, S.; Goel, H.; Omama, M.; and Chinchali, S. 2025. Neuro-symbolic evaluation of text-to-video models using formal verification. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 8395–8405.
- Shoukry, Y.; Nuzzo, P.; Balkan, A.; Saha, I.; Sangiovanni-Vincentelli, A. L.; Seshia, S. A.; Pappas, G. J.; and Tabuada, P. 2017. Linear temporal logic motion planning for teams of underactuated robots using satisfiability modulo convex programming. In *2017 IEEE 56th annual conference on decision and control (CDC)*, 1132–1137. IEEE.
- Song, E.; Chai, W.; Ye, T.; Hwang, J.-N.; Li, X.; and Wang, G. 2024. Moviechat+: Question-aware sparse memory for long video question answering. *arXiv preprint arXiv:2404.17176*.
- Tan, X.; Luo, Y.; Ye, Y.; Liu, F.; and Cai, Z. 2025. AL-LVB: All-in-One Long Video Understanding Benchmark. *arXiv:2503.07298*.
- Tapaswi, M.; Zhu, Y.; Stiefelhagen, R.; Torralba, A.; Urta-sun, R.; and Fidler, S. 2016. MovieQA: Understanding Stories in Movies through Question-Answering. In *IEEE Conference on Computer Vision and Pattern Recognition*, 4631–4640. IEEE Computer Society.
- Wang, W.; He, Z.; Hong, W.; Cheng, Y.; Zhang, X.; Qi, J.; Gu, X.; Huang, S.; Xu, B.; Dong, Y.; et al. 2024a. Lvbench: An extreme long video understanding benchmark. *arXiv preprint arXiv:2406.08035*.
- Wang, X.; Zhang, Y.; Zohar, O.; and Yeung-Levy, S. 2024b. Videoagent: Long-form video understanding with large language model as agent. In *European Conference on Computer Vision*, 58–76. Springer.
- Wang, Z.; Yu, S.; Stengel-Eskin, E.; Yoon, J.; Cheng, F.; Bertasius, G.; and Bansal, M. 2025. Videotree: Adaptive tree-based video representation for llm reasoning on long videos. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 3272–3283.
- Wu, C.; and Krähenbühl, P. 2021. Towards Long-Form Video Understanding. In *IEEE Conference on Computer Vision and Pattern Recognition*, 1884–1894. Computer Vision Foundation / IEEE.
- Wu, H.; Li, D.; Chen, B.; and Li, J. 2024. Longvideobench: A benchmark for long-context interleaved video-language understanding. *Advances in Neural Information Processing Systems*, 37: 28828–28857.
- Xiong, Y.; Huang, Q.; Guo, L.; Zhou, H.; Zhou, B.; and Lin, D. 2019. A Graph-Based Framework to Bridge Movies and Synopses. In *IEEE/CVF International Conference on Computer Vision*, 4591–4600. IEEE.
- Yang, A.; Yang, B.; Hui, B.; Zheng, B.; Yu, B.; Zhou, C.; Li, C.; Li, C.; Liu, D.; Huang, F.; Dong, G.; Wei, H.; Lin, H.; Tang, J.; Wang, J.; Yang, J.; Tu, J.; Zhang, J.; Ma, J.; Yang, J.; Xu, J.; Zhou, J.; Bai, J.; He, J.; Lin, J.; Dang, K.; Lu, K.; Chen, K.; Yang, K.; Li, M.; Xue, M.; Ni, N.; Zhang, P.; Wang, P.; Peng, R.; Men, R.; Gao, R.; Lin, R.; Wang, S.; Bai, S.; Tan, S.; Zhu, T.; Li, T.; Liu, T.; Ge, W.; Deng, X.; Zhou, X.; Ren, X.; Zhang, X.; Wei, X.; Ren, X.; Liu, X.; Fan, Y.; Yao, Y.; Zhang, Y.; Wan, Y.; Chu, Y.; Liu, Y.; Cui, Z.; Zhang, Z.; Guo, Z.; and Fan, Z. 2024. Qwen2 Technical Report. *arXiv:2407.10671*.
- Yang, Y.; Gaglione, J.-R.; Chinchali, S.; and Topcu, U. 2023. Specification-driven video search via foundation models and formal verification. *arXiv preprint arXiv:2309.10171*.
- Ye, J.; Wang, Z.; Sun, H.; Chandrasegaran, K.; Durante, Z.; Eyzaguirre, C.; Bisk, Y.; Niebles, J. C.; Adeli, E.; Fei-Fei, L.; et al. 2025. Re-thinking temporal search for long-form video understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 8579–8591.
- Ye, Y.; Zhao, Z.; Li, Y.; Chen, L.; Xiao, J.; and Zhuang, Y. 2017. Video question answering via attribute-augmented attention network learning. In *Proceedings of the 40th International ACM SIGIR conference on Research and Development in Information Retrieval*, 829–832.
- Zhang, K.; Li, B.; Zhang, P.; Pu, F.; Cahyono, J. A.; Hu, K.; Liu, S.; Zhang, Y.; Yang, J.; Li, C.; et al. 2024. Lmms-eval: Reality check on the evaluation of large multimodal models. *arXiv preprint arXiv:2407.12772*.
- Zheng, X.; Li, Z.; Ruchkin, I.; Piskac, R.; and Pajic, M. 2025. Neurostrata: Harnessing neurosymbolic paradigms for improved design, testability, and verifiability of autonomous cps. *arXiv preprint arXiv:2502.12267*.
- Zhu, Y.; Groth, O.; Bernstein, M.; and Fei-Fei, L. 2016. Visual7w: Grounded question answering in images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4995–5004.