

# Towards Neuro-Symbolic Video Understanding

Minkyu Choi<sup>\*1,2</sup>, Harsh Goel<sup>†1,2</sup>, Mohammad Omama<sup>†1,2</sup>,  
Yunhao Yang<sup>1</sup>, Sahil Shah<sup>1,2</sup>, and Sandeep Chinchali<sup>1,2</sup>

<sup>1</sup> The University of Texas at Austin TX 78712, USA

<sup>2</sup> UT SWARM Lab, Department of Electrical and Computer Engineering  
<https://utaustin-swarmlab.github.io>

\*Corresponding author: [minkyu.choi@utexas.edu](mailto:minkyu.choi@utexas.edu)

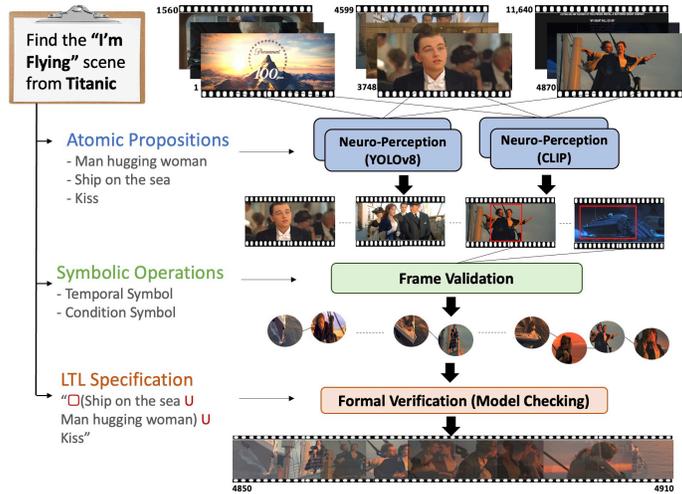
<sup>†</sup>Contributed equally to this work.

**Abstract.** The unprecedented surge in video data production in recent years necessitates efficient tools to extract meaningful frames from videos for downstream tasks. Long-term temporal reasoning is a key desideratum for frame retrieval systems. While state-of-the-art foundation models, like VideoLLaMA and ViCLIP, are proficient in short-term semantic understanding, they surprisingly fail at long-term reasoning across frames. A key reason for this failure is that they intertwine per-frame perception and temporal reasoning into a single deep network. Hence, decoupling but co-designing the semantic understanding and temporal reasoning is essential for efficient scene identification. We propose a system that leverages vision-language models for semantic understanding of individual frames and effectively reasons about the long-term evolution of events using state machines and temporal logic (TL) formulae that inherently capture memory. Our TL-based reasoning improves the F1 score of complex event identification by 9 – 15%, compared to benchmarks that use GPT-4 for reasoning, on state-of-the-art self-driving datasets such as Waymo and NuScenes. The source code is available at <https://github.com/UTAustin-SwarmLab/Neuro-Symbolic-Video-Search-Temporal-Logic>.

**Keywords:** Video Understanding · Video Reasoning · Neuro Symbolic AI · Temporal Logic · Formal Methods

## 1 Introduction

There is a significant increase in video data production, with platforms such as YouTube receiving 500 hours of uploads every minute. Additionally, autonomous vehicle companies such as Waymo generate 10-100 TB [27] of data, and worldwide security cameras record around 500 PB [26] of data daily. Consequently, we require tools with sophisticated query capabilities to navigate this immense volume of video content. For instance, a query such as “Find me all scenes where event A happened, event B did not occur, and event C occurs hours later” requires advanced methods capable of long-term temporal reasoning. Surprisingly, we find that today’s state-of-the-art (SOTA) video and language foundation



**Fig. 1: NSVS-TL Pipeline.** The input query — “Find the ‘I’m Flying scene from Titanic” — is first decomposed into semantically meaningful atomic propositions such as “man hugging woman”, “ship on the sea”, and “kiss” from a high-level user query. SOTA vision and vision-language models are then employed to annotate the existence of these atomic propositions in each video frame. Subsequently, we construct a probabilistic automaton that models the video’s temporal evolution based on the list of per-frame atomic propositions detected in the video. Finally, we evaluate when and where this automaton satisfies the user’s query. We do this by expressing it in a formal specification language that incorporates temporal logic. The TL equivalent of the above query is ALWAYS ( $\square$ ) “man hugging woman” UNTIL ( $\cup$ ) “ship on the sea” UNTIL ( $\cup$ ) “kiss”. Formal verification techniques are utilized on the automaton to retrieve scenes that satisfy the TL specification.

models fail to identify complex events, especially when key-frames are rare in a video or when the input query is complex (see Fig. 2).

Our key insight is that these foundation models aggregate per-frame semantics into a latent vector from which precise scene identification is difficult, especially over long videos. Instead, this paper innovates a novel neuro-symbolic approach that composes neural perception with temporal logic reasoning. We use vision-language models for per-frame semantic understanding but employ state machines and temporal logic to reason about the long-term evolution of events since they inherently capture memory.

Imagine we aim to pinpoint the iconic “I’m flying” scene in the movie Titanic. How do we identify a 5-minute-long scene precisely within the 3-hour and 14-minute movie? Our proposed method, Neuro-Symbolic Video Search with Temporal Logic (*NSVS-TL*), presents a solution using a neuro-symbolic approach for video understanding (see Fig. 1). Our key contributions can be summarized as:

1. **Real-time Neuro-Symbolic Scene Identification:** We propose a framework that segregates temporal reasoning from perception, overcoming the limitations of video-language models in localizing video segments in long

videos. We employ perception models to identify frames relevant to the temporal logic query. Subsequently, we construct an automaton to verify the satisfiability of the temporal logic query. This segregation enables us to search for temporally extended scenes over long videos.

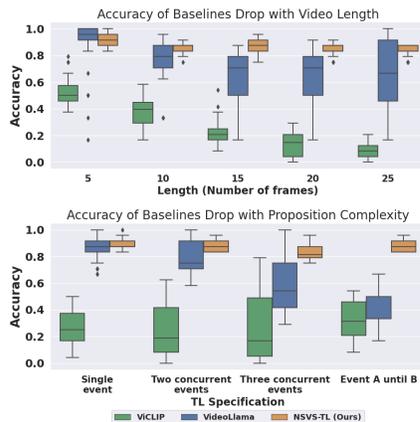
2. **Formal Verification of Video Frame Automaton against a TL Specification:** Our algorithm seamlessly provides a confidence measure for predictions on localized scenes by leveraging formal verification on a TL specification. We integrate the probabilistic model checker STORM [13] to formally verify the automaton constructed during scene identification and to measure the probability of satisfying the TL specification.
3. **Datasets:** We propose Temporal Logic Video (TLV) datasets. These include a synthetic annotated video dataset created by stitching together images of objects from the ImageNet [3] and COCO [17] datasets to verify frame search algorithms. Additionally, we annotate autonomous vehicle datasets – Waymo [27] and NuScenes [9] – with TL specifications derived from ground truth object annotations such as “bus”, “pedestrian”, etc., in the scenes.

## 2 Related Work

### 2.1 Video Search and Event Detection

Existing research in video event detection predominantly focuses on tracking the spatio-temporal information of objects to identify events by employing deep neural networks for learning latent representations like objects’ motion and position [4, 7, 8, 12, 16, 21, 33, 40]. These methods learn latent representations of videos for downstream event detection or classification, which require substantial computational resources for training neural networks [7, 8, 16]. In contrast, our approach leverages SOTA vision or vision-language models for per-frame semantic recognition, reducing the need for such infrastructure.

Additionally, some studies focus on searching and detecting events described in natural language [4, 34]. Video language models — VideoLlama [39] and Video-ChatGPT [18] — integrate language foundation models like GPT-4 [22] and Llama [29] for video question answering and zero-shot event recognition. However, their aggregation of temporal video information



**Fig. 2: Comparative Performance on the Event Identification Task: Video Language Models versus *NSVS-TL*.** The accuracy of event identification with Video Language Models (Blue/Green) drops as video length or query complexity increases. In contrast, *NSVS-TL* (Orange) shows consistent performance irrespective of video length or query complexity.

hinders accurate frame identification in long videos. To address this, some methods have introduced temporal logic for event detection [2, 35]. In contrast, our proposed method focuses on scene identification of any temporal length.

## 2.2 Symbolic Video Understanding

Many works explore approaches to building symbolic representations for video understanding. Such symbolic representations of videos are useful for video classification [6, 30], event detection [16, 21, 33], video question-answering [28, 36], etc. Existing methods either construct graph structures [20, 32, 37] or use latent-space representations as symbolic representations of videos [1, 15, 25]. Among the existing works, several of them propose methods of learning symbolic representation to understand long videos [11, 28, 31, 32]. However, these approaches also rely on latent-space representations, which lack interpretability. In contrast, we use a fully interpretable automaton-based representation that provides formal guarantees to downstream tasks like video search.

## 3 Preliminaries

We introduce a practical example to assist the reader in understanding our proposed methodology. Let’s consider a scenario where our objective is to ensure that an autonomous vehicle follows a safety protocol when children are present in school zones. The key query for this test is: “Find all scenes where the school zone sign is visible and children are present, continuing until the vehicle exits the school zone”.

### 3.1 Neural Perception Model

*Understanding the semantics of a frame:* The initial step involves identifying frames corresponding to propositions in the query — such as “school zone sign” and “children”. Therefore, we require the use of vision or vision-language neural perception models. Our method incorporates neural perception models to extract information from video frames, focusing on object detection capabilities. The confidence scores produced by these models are crucial as they indicate the level of certainty in detection results, which significantly impacts the accuracy of the subsequent temporal logic analysis and scene identification.

### 3.2 Temporal Logic

*Reasoning about events across time:* Temporal logic (TL) provides a structured framework for describing and reasoning about the temporal properties of sequences or processes [5, 19]. It extends classical logic with temporal operators to express propositions about the flow of time. In video analysis, temporal logic can be used to define complex conditions or sequences of events within a video stream. The syntax of temporal logic is composed of a set of first-order logic

operators, a set of temporal operators, and a set of propositions. The set of first-order logic operators includes AND ( $\wedge$ ), OR ( $\vee$ ), NOT ( $\neg$ ), Implies ( $\implies$ ), etc. The set of temporal operators includes Always ( $\square$ ), Eventually/Exist ( $\diamond$ ), Next ( $\mathbf{X}$ ), Until ( $\mathbf{U}$ ), etc. Intuitively, for events  $a$  and  $b$ ,  $\square a$  means  $a = \text{True}$  for every step,  $\diamond b$  means there exists one step where  $b = \text{True}$ , and  $a\mathbf{U}b$  means  $a = \text{True}$  for all the steps before  $b = \text{True}$ , then  $b = \text{True}$  for all the subsequent steps. A minimal syntax for temporal logic is specified as follows:

$$\Phi ::= p_k \mid \neg\Phi \mid \Phi \vee \Phi \mid \Phi \wedge \Phi \mid \mathbf{X}\Phi \mid \Phi \mathbf{U} \Phi, \quad (1)$$

where  $\Phi$  is a temporal logic specification,  $p_k \in P$  is an atomic proposition — a statement or assertion that must be true or false. We denote  $P$  as the set of all atomic propositions. Temporal logic includes linear temporal logic (LTL), computation tree logic (CTL), probabilistic computation tree logic (PCTL), etc. We present some of these terminologies in the Appendix.

For instance, in our running example, the set of propositions  $P$  includes atomic propositions  $p_k$  such as “SchoolZoneSign” and “Children”, and the temporal logic specification for the query is formulated as:

$$\Phi = (\text{SchoolZoneSign} \wedge \text{children}) \mathbf{U} \neg\text{SchoolZoneSign}. \quad (2)$$

### 3.3 Probabilistic Automaton

*Constructing a model of video: A Probabilistic Automaton (PA)* is a mathematical model used to represent dynamic systems transitioning across various states with associated probabilities in response to certain inputs. The PA is defined as a 5-tuple  $\mathcal{A} = (Q, \Omega, \delta, Q_0, F)$ , where:

- $Q$  is a finite set of states in which the automaton can reside.
- $\Omega$  is a finite set of symbols, or inputs, that trigger transitions between states.
- $\delta : Q \times \Omega \times Q \rightarrow [0, 1]$  is the transition function that maps a transition from one state to another given an input symbol with the probability of that transition. The sum of the probabilities for all outgoing transitions from a state  $q$  equals 1, ensuring a complete probability distribution.
- $Q_0 \subset Q \wedge Q_0 \neq \emptyset$  is the set of initial states from which the automaton begins its operation.
- $F \subseteq Q$  is the set of acceptance or terminal states representing successful terminations.

In the running example,  $Q_0$  is defined by the initial detection of “SchoolZoneSign” and “Children”. The terminal states  $F$  might contain conditions where the “SchoolZoneSign” is no longer detected, indicating the vehicle’s exit from the school zone. The PA is a sophisticated model that transitions through states based on the temporal logic specification  $\Phi$  (see Eq. (2)). The transition probabilities between these states are quantified by  $\delta$ , with  $Q_0$  and  $F$  denoting the start and end of the scenes of interest.  $Q$  and  $\delta$  are illustrated in Fig. 3 in details. With this constructed PA, we now conduct formal verification to ascertain whether the PA conforms to our predefined specification  $\Phi$ .

### 3.4 Formal Verification

*Verifying frames of interest against pre-defined specifications:* Formal verification involves mathematically proving that a system adheres to specified properties. One such technique under formal verification is *Model Checking*, a formal verification technique leveraged to check if a model of a system adheres to a set of specified properties. Given a model  $\mathcal{M}$  and a specification  $\Phi$  expressed in a suitable logic, model checking systematically explores the state space of  $\mathcal{M}$  to verify if  $\Phi$  holds. If  $\mathcal{M}$  satisfies  $\Phi$ , denoted as  $\mathcal{M} \models \Phi$ , the model checking algorithm returns true; otherwise, it returns false. In the context of our running example, this step ensures that the PA precisely adheres to the vehicle’s compliance with safety protocols while driving in the school zone.

## 4 Methodology

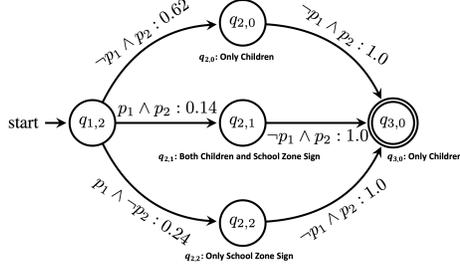
We introduce a novel way to identify scenes of interest using a neuro-symbolic approach. Given video streams or clips and a TL specification  $\Phi$ , Neuro-Symbolic Visual Search with Temporal Logic (*NSVS-TL*) identifies scenes of interest through four steps:

- **Step 1:** We calibrate the confidence of neural perception models to ensure precise object detection. This calibration enables the detection of relevant propositions in a given frame to construct a PA.
- **Step 2:** Subsequently, each frame undergoes a validation process with two distinct validation functions. This step ensures that only frames containing relevant visual information proceed to the next phase of the method.
- **Step 3:** Upon validation, we construct a probabilistic automaton to encapsulate the temporal and logical relations between successive frames.
- **Step 4:** Finally, we utilize model checking to determine whether a constructed automaton satisfies the TL specification. If an automaton passes this check, then a sequence of frames within the automaton is identified as a scene of interest by the given TL specification.

*Calibrating Neural Perception Model.* We first calibrate neural perception models,  $f_v$ , which detect atomic propositions  $p_k$  in a frame  $\mathcal{F}_t$ , where  $t$  represents the frame’s time index with a confidence  $\hat{y}$  for probabilistic verification. We use a generalized logistic function as a mapping estimation function  $z(\hat{y})$ , which calibrates the confidence of models into an accuracy metric, as defined in Eq. (3):

$$z(\hat{y}) = \frac{1}{1 + e^{-k(\hat{y} - \hat{y}_0)}}, \quad (3)$$

where  $k$  is a scaling factor that adjusts the sensitivity and  $\hat{y}_0$  is an inflection point. We optimize  $k$  and  $\hat{y}_0$  to calibrate the neural perception models as presented in Fig. 4. We also empirically select optimal thresholds for true positives



**Fig. 3: Sample Automaton of the Running Example.** Illustrates transitions from  $\mathcal{F}_1$  to  $\mathcal{F}_3$ . Key transitions include:  $q_{1,2} \rightarrow q_{2,0}$  with a 0.62 probability for only children,  $q_{1,2} \rightarrow q_{2,1}$  with a 0.14 probability for both sign and children, and  $q_{1,2} \rightarrow q_{2,2}$  with a 0.24 probability for only the sign. There are only children in  $\mathcal{F}_3$ . Therefore, all states in  $\mathcal{F}_2$  are connected to  $q_3$  with 1.0 probability.

( $\gamma_{tp}$ ) and false positives ( $\gamma_{fp}$ ). The calibration function, which incorporates these thresholds along with a mapping estimation function, is presented in Eq. (4):

$$g(\hat{y}; \gamma_{fp}, \gamma_{tp}) = \begin{cases} 0 & \text{if } \hat{y} < \gamma_{fp}, \\ 1 & \text{if } \hat{y} > \gamma_{tp}, \\ z(\hat{y}) & \text{otherwise.} \end{cases} \quad (4)$$

For confidence below  $\gamma_{fp}$ , it outputs 0, indicating low reliability. Above  $\gamma_{tp}$ , it outputs 1, reflecting high accuracy. For intermediate confidence levels, it utilizes a mapping estimation  $z(\hat{y})$  as illustrated in Fig. 4.

*Frame Validation.* Following the calibration step, the process transitions to frame validation. To ensure the presence of atomic propositions, we first validate a frame  $\mathcal{F}_t$  defined by a detection verification function:

$$V_c(\mathcal{F}_t, P) = \begin{cases} 1 & \text{if } g(f_v(\mathcal{F}_t, p_k); \gamma_{fp}, \gamma_{tp}) > 0, \forall p_k \in P \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

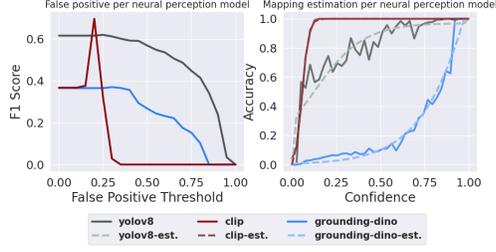
A frame  $\mathcal{F}_t$  is considered valid for the automaton if  $V_c(\mathcal{F}_t, P) = 1$ . However, this is not enough to consider the inclusion of  $\mathcal{F}_t$  in the automaton. For instance, consider a slight modification to our running example:

$$\Phi = (\text{SchoolZoneSign} \wedge \text{children}) \cup \neg \text{adults}. \quad (6)$$

In this case, we are interested in children walking around the school zone unaccompanied by adults. If we solely rely on the validation function from Eq. (5), frames that include a school sign, children, and adults are added to the automaton. However, including frames with adults contradicts our requirement to exclude them. To address this, we implement *symbolic verification* to ensure a frame  $\mathcal{F}_t$  aligns with our TL specification  $\Phi$ . This process involves evaluating frames against both first-logic operators (denoted as  $\Psi$ ) and temporal operators (denoted as  $\Theta$ ) defined in  $\Phi$  as follows:

$$\Psi = \{\psi \mid \psi \in \{\wedge, \vee, \neg, \dots\} \wedge \psi \text{ appears in } \Phi\}, \quad (7a)$$

$$\Theta = \{\theta \mid \theta \in \{\square, \diamond, \mathbf{U}, \dots\} \wedge \theta \text{ appears in } \Phi\}. \quad (7b)$$



**Fig. 4: Calibrating Neural Perception Models.** We empirically select the optimal false positive threshold as shown in the left figure, while the right figure illustrates mapping estimation functions optimized for each neural perception model.

During a symbolic verification process, each frame initially evaluates the set of first-order logic operators  $\Psi$ . Formally,  $\mathcal{F}_t$  is satisfied on  $\Psi$ , denoted as  $\mathcal{F}_t \models \Psi$ , if and only if the following conditions hold:

- **Conjunction** ( $\wedge$ ): A frame  $\mathcal{F}_t$  is satisfied on  $\Psi$  for the conjunction only if both propositions  $p_1$  and  $p_2$  associated by  $\wedge$  are present in  $\mathcal{F}_t$ .
- **Disjunction** ( $\vee$ ): A frame  $\mathcal{F}_t$  is satisfied on  $\Psi$  for the disjunction if at least one of the propositions  $p_1$  and  $p_2$  associated by  $\vee$  is present in  $\mathcal{F}_t$ .
- **Negation** ( $\neg$ ): A frame  $\mathcal{F}_t$  is satisfied on  $\Psi$  for the negation if the proposition  $p_k$  is not present in  $\mathcal{F}_t$ .

In our running example (see Eq. (6)), a frame  $\mathcal{F}_t$  must have a school zone and children, but not adults. This is represented by a combination of connection ( $\wedge$ ) and negation ( $\neg$ ) operators. A frame that includes a school sign and children satisfies the conjunction requirement ( $\wedge$ ). However, if adults are also present in the same frame, it violates the negation criterion ( $\neg$ ), as our query explicitly requires the exclusion of adults.

While the first part of symbolic verification, using logical operators ( $\Psi$ ) helps identify frames where children are present without adults ( $\neg$ adult), it can ignore the important temporal aspects of our query. Consider a scenario in which our objective is to observe children without adults in the school zone during the period before any adult arrives. If we only validate frames on  $\Psi$ , a frame meets the immediate logical requirements (children without adults) yet fails to capture the essential temporal aspects — the arrival of adults. This deficiency highlights the importance of integrating the temporal operators ( $\Theta$ ) in our evaluation. Here, we examine the temporal relevance of each proposition in  $\Phi$  as follows:

$$\tau(p_k) = \begin{cases} 1 & \text{if } \exists p_k \text{ such that } \theta p_k, \\ 1 & \text{if } \exists p_k, p_{k+1} \text{ such that } p_k \theta p_{k+1}, \\ 0 & \text{otherwise,} \end{cases} \quad (8)$$

where  $\theta \in \Theta$  (see Eq. (7b)). In conclusion, by the symbolic verification rules defined above, a frame  $\mathcal{F}_t$  is validated by the symbolic verification function expressed as:

$$V_{sv}(\mathcal{F}_t, P) = \begin{cases} 1 & \text{if } (\exists p_k \in P : \tau(p_k) = 1) \vee (\mathcal{F}_t \models \Psi), \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

Therefore, combining the symbolic verification function with the detection verification function (see Eq. (5)), the frame validation function is defined as follows:

$$V(\mathcal{F}_t, P) = \begin{cases} 1 & \text{if } (V_c(\mathcal{F}_t, P) = 1) \wedge (V_{sv}(\mathcal{F}_t, P) = 1), \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

A frame  $\mathcal{F}_t$  is valid for the automaton if  $V(\mathcal{F}_t, P) = 1$ .

*Dynamic Automaton Construction.* Once a frame  $\mathcal{F}_t$  is validated by Eq. (10), we construct a PA as a model of the video with validated frames. We call this process *Dynamic Automaton Construction* because we dynamically construct the automaton based on frame validation instead of building it for an entire video. We first generate  $2^{|P|}$  new automaton states, where  $|P|$  is the number of atomic propositions. Each of these states corresponds to a specific combination of propositions from the set  $P$ , thereby allowing a comprehensive representation of all possible combinations of propositions within the automaton. In our running example (see Eq. (2)), we construct four states representing all possible combinations of atomic propositions or their negation with the proposition set  $P = \{p_1, p_2\}$ , where  $p_1 = \text{SchoolZoneSign}$  and  $p_2 = \text{Children}$ . The labels for these states would be  $p_1 \wedge p_2$ ,  $\neg p_1 \wedge p_2$ ,  $p_1 \wedge \neg p_2$ , and  $\neg p_1 \wedge \neg p_2$ , respectively.

After creating new states, we compute the transition probabilities from previous states to new states  $q_{t,\omega}$  where  $\omega \in 2^{|P|}$  as follows:

$$\mathbb{P}_{t,\omega} = \prod_{p_k} g(f_v(\mathcal{F}_t, p_k); \gamma_{fp}, \gamma_{tp}) \quad \forall p_k \in P. \quad (11)$$

The probability of the negation of a proposition ( $\neg p_k$ ) is computed as  $1 - g(f_v(\mathcal{F}_t, p_k); \gamma_{fp}, \gamma_{tp})$ . For every new state  $q_{t,\omega}$  with a probability  $\mathbb{P}_{t,\omega} > 0$ , we construct transitions from the previous states  $q_{t-1,\omega}$  to  $q_{t,\omega}$ . In this case,  $\mathbb{P}_{t,\omega}$  represents the probability of the transition from any state in frame  $\mathcal{F}_{t-1}$  to the new state  $q_{t,\omega}$  of  $\mathcal{F}_t$ . We illustrate how the automaton is created for our running example in Fig. 3.

*Model Checking.* Once we have constructed the automaton (model of a video), we apply model checking using STORM [14], a probabilistic model checking method, to determine whether it satisfies our specification formulated by probabilistic computation tree logic (PCTL). Our method iteratively processes frames and dynamically builds the automaton. This ensures that its evolving structure aligns with the desired TL specifications. This verification step checks the automaton’s transitions and states to ensure they comply with the probabilistic temporal properties defined in the TL specification.

When the automaton  $\mathcal{A}$  (see Sec. 3) satisfies  $\Phi$ , it indicates that all frames in  $\mathcal{A}$  correspond to our defined scenes of interest. Subsequently, all frames representing the automaton’s states are added to the scenes of interest set  $\mathcal{S}$ . After this addition, the automaton  $\mathcal{A}$  is reset ( $\mathcal{A} = \emptyset$ ), preparing the system for processing subsequent frames. The process then proceeds with frame  $\mathcal{F}_{t+1}$ , repeating the previous steps and including frames that satisfy  $\Phi$  into  $\mathcal{S}$ . In the context of our

running example, all frames are added to  $\mathcal{S}$  when an autonomous vehicle exits a school zone where children are present.

## 5 TLV: Temporal Logic Video Datasets

Given the lack of SOTA video datasets for long-horizon, temporally extended activity and object detection, we introduce the Temporal Logic Video (TLV) datasets. The synthetic TLV datasets are compiled by stitching together static images from image datasets like COCO [17] and ImageNet [3]. This enables the artificial introduction of a wide range of TL specifications. Additionally, we create two video datasets annotated with TL specifications based on the open-source autonomous vehicle (AV) driving datasets NuScenes [9] and Waymo [27]. The dataset is available on GitHub<sup>3</sup>.

### 5.1 Dataset Compilation

The synthetic TLV datasets are compiled by combining static images randomly selected from the COCO and ImageNet datasets. The propositions selected from the ground truth labels of the original datasets are used to generate TL specifications. For example, if ‘person’ and ‘car’ are selected from the original COCO dataset, they become the propositions for the dataset. Possible TL specifications could be either “person and car” or “person until car”. Applying this approach to autonomous driving, we similarly annotate the Waymo and NuScenes datasets, using detected ground truth object labels to create TL annotations. The total number of frames of TLV datasets is detailed in Tab. 1.

Ground Truth TL Specifications	Synthetic TLV Dataset		Real TLV Dataset	
	COCO	ImageNet	Waymo	Nuscenes
Eventually Event A	15,750	15,750	-	-
Always Event A	15,750	15,750	-	-
Event A And Event B	31,500	-	-	-
Event A Until Event B	15,750	15,750	8,736	19,808
(Event A And Event B) Until Event C	31,500	-	5,789	7,459

**Table 1:** The number of frames in Synthetic TLV and Real TLV Datasets.

### 5.2 Ground Truth Temporal Logic Specifications

1. **Always/Eventually event A.** These temporal logic queries can be expressed as either  $\Phi = \Box p_k$  or  $\Diamond p_k$ , identifying occurrences of a single event within a video. For instance, a natural language equivalent of this query is “Identify all frames depicting a car accident” or “Locate all frames showcasing an animal’s presence”.

<sup>3</sup> <https://github.com/UTAustin-SwarmLab/Temporal-Logic-Video-Dataset>

2. **Event A and B.** This scenario involves the concurrent occurrence of two distinct events streams, expressed as  $\Phi = p_1 \wedge p_2$ . Natural language specifications such as “Locate all frames that have a pedestrian and a vehicle” or “Locate scenes that show a car accident and a pedestrian crossing” can be represented by this TL query.
3. **Event A until event B.** It evaluates scenarios where a specific event  $p_1$  is to be identified and monitored until another event  $p_2$  occurs, expressed as  $\Phi = p_1 \cup p_2$ . For instance, “Identify all scenes where a pedestrian is walking until a vehicle passes by”.
4. **Event A and B until C.** This specification encapsulates a video where two simultaneously occurring events are monitored up to the point a third event occurs; for example, “Recognize scenes that show a child playing and dog barking until an adult appears”. It is expressed by  $\Phi = (p_1 \wedge p_2) \cup p_3$ .

The synthetic TLV dataset is constructed based on all the TL specifications mentioned above. In contrast, the AV datasets are more specifically focused on the TL specifications of “Event A until B” and “Event A and B until C”. This is due to the fact that generating some TL specifications for real datasets would require altering the original video sequences, a scenario we aim to avoid.

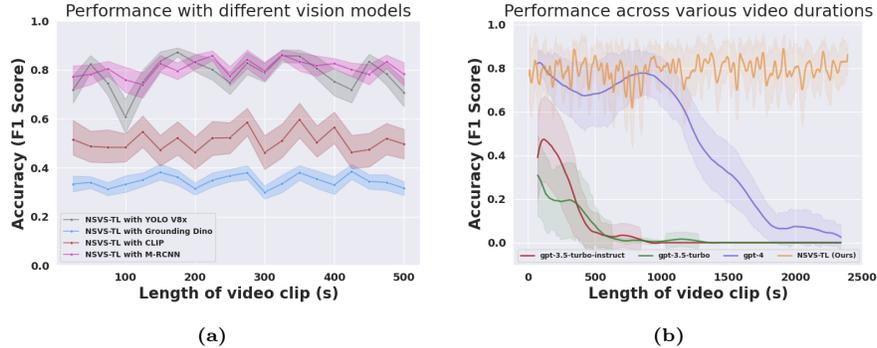
## 6 Experiment

In our experiments, we assess *NSVS-TL* on diverse scene identification tasks using multiple datasets, aiming to answer key questions:

- How significantly does the choice of a *neural perception model* influence the performance of scene identification in the *NSVS-TL* system?
- Given the optimal perception model, how does the efficacy of *NSVS-TL*, which employs TL to reason over detected propositions per frame, compare to that of a Large Language Model (LLM) across various TL specifications?
- How does the duration of a video affect the performance of our method when the event is rare (in a long-duration video) or is temporally extended?

### 6.1 Neural Perception Models

We utilize various neural perception models to assess the impact of different perception models on scene identification. We leverage SOTA computer vision models such as YOLO V8 [24], Grounding Dino [38], Masked R-CNN [10], and Contrastive Language-Image Pre-training (CLIP) [23]. We use an image detection model instead of a video detection model for several reasons. First, we define a video as a sequence of frames, so detecting objects in each frame effectively means detecting objects throughout the entire video. Second, image detection can process per-frame input faster than video object detectors, making it essential for real-time video understanding, which is one of our key contributions.



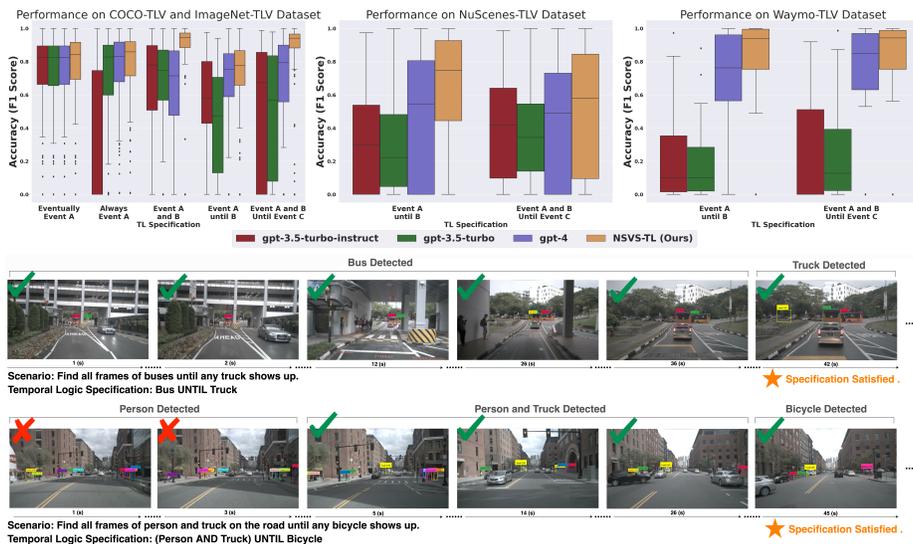
**Fig. 5: Performance of *NSVS-TL* Across Different Video Lengths.** Fig. 5a demonstrates the impact of various neural perception models on scene identification performance. Additionally, Fig. 5b illustrates the F1 scores for scene retrieval against the video length, fulfilling the “A until B” temporal specification.

## 6.2 Benchmarks for Scene Identification

Analysis of our method, benchmarked against video-language foundation models such as Video-Llama and ViCLIP, reveals their inherent limitations in handling long-term and complex temporal queries (see Fig. 2). While these models perform well with shorter sequences, they lack robust temporal reasoning capabilities and struggle with extended video lengths. This limitation makes them less effective for long-term event detection tasks. Moreover, no current models exist that are capable of retrieving precise scenes or frames based on textual description or TL specification, despite their ability to understand the semantics of video content. Hence, we design benchmarks for a balanced comparison that uses LLMs for reasoning over per-frame annotations instead of TL. Initially, neural perception models provide per-frame annotations to LLMs such as GPT. These LLMs are then prompted to identify scenes of interest according to specified propositions and TL specifications. We use these benchmarks to evaluate the impact of video length on scene identification performance. The prompts and methodologies applied in these comparisons are elaborated in the Appendix.

Model	COCO-TLV + ImageNet-TLV					Waymo-TLV		NuScenes-TLV		
	◇ A	□ A	A ∧ B	A ∪ B	(A ∧ B) ∪ C	A ∪ B	(A ∧ B) ∪ C	A ∪ B	(A ∧ B) ∪ C	
Grounding Dino	0.30	0.50	0.28	0.30	0.33	-	-	-	-	
CLIP-ViT-B-32	0.69	0.68	0.12	0.60	0.47	-	-	-	-	
YOLOv8n	0.55	0.51	0.66	0.43	0.58	0.20	0.18	0.17	0.05	
YOLOv8m	0.75	0.72	0.84	0.67	0.77	0.38	0.28	0.31	0.16	
YOLOv8x	0.79	0.77	<b>0.85</b>	0.72	<b>0.80</b>	0.48	0.38	0.32	0.17	
Masked R-CNN	<b>0.81</b>	<b>0.84</b>	<b>0.85</b>	<b>0.75</b>	0.74	<b>0.83</b>	<b>0.80</b>	<b>0.64</b>	<b>0.52</b>	

**Table 2: Comprehensive Performance Comparison Across Models and TLV Datasets.** This illustrates the impact of different neural perception models across various datasets. The meaning of each symbol is described in Sec. 5.2.



**Fig. 6: Comparative Performance of *NSVS-TL* Across Complex Temporal Logic Specifications.** The upper box plots demonstrate *NSVS-TL*'s performance to benchmarks across different TL specifications. The lower section displays selected scenes from the NuScenes-TLV dataset, highlighting practical applications and results.

### 6.3 Evaluation Metrics

We evaluate our method using multi-class classification metrics such as precision, recall, and F1 score for scene identification. The F1 score was used to evaluate our method with a balanced consideration of both precision (correctly identified frames / all frame identifications) and recall (correctly identified frames / correct ground truth frames), aiming to minimize false positives and false negatives. This measure is crucial for accurately detecting and retrieving relevant frames, indicating the effectiveness of our method in pinpointing specific scenes. The TLV Dataset, with its ground truth annotations, was used for this evaluation.

## 7 Results

*Impact of different Neural Perception Models.* Our work introduces a novel methodological approach that combines a neural perceptual model with symbolic methods. This integration provides a new perspective in understanding long-horizon videos. In fact, we propose a framework that allows for the integration of any neural perceptual models into our work, enhancing the capability to understand videos and enabling us to localize frames of interest with respect to queries. We show a comparison between different neural perceptual models in Tab. 2 and Fig. 5a. We observe that *NSVS-TL* with various neural perception

models performs differently depending on the complexity of the TL specification and datasets. The variation underscores the importance of these models’ inherent visual detection capabilities for scene identification.

*Comparative Analysis of Temporal Logic and LLM-based Reasoning for Frame Retrieval.* We evaluate TL-based reasoning against LLM-based reasoning over per-frame semantic annotations derived from neural perception models. The evaluation is conducted on four datasets: COCO-TLV, ImageNet-TLV, NuScenes-TLV, and Waymo-TLV. The results are collectively presented in Fig. 6. For single event scenarios such as “*Always/Eventually Event A Occurs*”, both our method and LLM-based reasoning perform reasonably well since these events do not require complex reasoning. In multi-event scenarios (“*Event A and Event B*”, “*Event A until Event B*,” and “*Event A and Event B until Event C*”), our TL-based reasoning outperforms all LLM-based baselines.

*Impact of Video Length and Temporally Extended Scenarios on scene identification.* We evaluate multi-event sequences with temporally extended gaps, such as those in the specification of events A until B. For these scenarios, we design a scene identification task that extends the video length: event A starts at the beginning and continues until event B occurs at the end. This tests the reasoning capacity of the LLM benchmarks as the temporal distance between events increases. The scene identification performance of GPT-3.5 and GPT-3.5 Turbo Instruct degrade within videos that are 500 seconds long as illustrated in Fig. 5b. GPT-4, on the other hand, suffers from a sharp decline in performance for videos whose lengths range from 1000 seconds and more. In contrast, *NSVS-TL* maintains consistent performance throughout. This consistency is observed even in videos spanning up to 40 minutes or 2400 seconds, indicating *NSVS-TL*’s potential reliability in handling even longer videos.

## 8 Conclusion

*NSVS-TL*, a neuro-symbolic method, enhances video understanding and scene identification in contexts demanding extended temporal reasoning. *NSVS-TL*, advances scene identification across various video durations by integrating semantic understanding with temporal reasoning. Unlike other methods focusing on short-duration video reasoning, we propose a novel approach for comprehending long-horizon, temporally extended videos.

**Limitations and future work.** *NSVS-TL*’s per-frame neural perception is limited in capturing multi-frame semantics, especially in scenarios that require understanding sequences, like “man falling from a horse”, though it performs well with simpler scenarios such as “man on a horse”. Future work will focus on interpreting multi-frame events, thereby extending its use in more complex video understanding tasks.

## Acknowledgments

This material is based upon work supported in part by the Office of Naval Research (ONR) under Grant No. N00014-22-1-2254 and N00014-24-1-2097. Additionally, this work was supported by the Defense Advanced Research Projects Agency (DARPA) contract FA8750-23-C-1018 and DARPA ANSR: RTX CW2231110. Approved for Public Release, Distribution Unlimited.

## References

1. Bertasius, G., Wang, H., Torresani, L.: Is space-time attention all you need for video understanding? In: Meila, M., Zhang, T. (eds.) International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 139, pp. 813–824. PMLR (2021)
2. Cheng, Y., Fan, Q., Pankanti, S., Choudhary, A.N.: Temporal sequence modeling for video event detection. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 2235–2242. IEEE Computer Society, Columbus, OH, USA (2014)
3. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. pp. 248–255 (2009). <https://doi.org/10.1109/CVPR.2009.5206848>
4. Doshi, K., Yilmaz, Y.: Towards interpretable video anomaly detection. In: IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2654–2663. IEEE, Waikoloa, HI, USA (2023)
5. Emerson, E.A.: Temporal and modal logic. In: Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics (1991), <https://api.semanticscholar.org/CorpusID:6062082>
6. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: IEEE/CVF International Conference on Computer Vision. pp. 6201–6210. IEEE (2019)
7. Feichtenhofer, C., Pinz, A., Wildes, R.P.: Spatiotemporal multiplier networks for video action recognition. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 7445–7454. IEEE Computer Society, Honolulu, HI, USA (2017)
8. Feichtenhofer, C., Pinz, A., Zisserman, A.: Convolutional two-stream network fusion for video action recognition. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 1933–1941. IEEE Computer Society, Las Vegas, NV, USA (2016)
9. Guo, Y., Caesar, H., Oscar, B., Phillion, J., Fidler, S.: The efficacy of neural planning metrics: A meta-analysis of pkl on nuscenec. In: IROS 2020 Workshop on Benchmarking Progress in Autonomous Driving (2020)
10. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn (2017)
11. Huang, Q., Xiong, Y., Rao, A., Wang, J., Lin, D.: Movienet: A holistic dataset for movie understanding. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J. (eds.) European Conference on Computer Vision. Lecture Notes in Computer Science, vol. 12349, pp. 709–727. Springer (2020)
12. Jiang, F., Yuan, J., Tsafaris, S.A., Katsaggelos, A.K.: Anomalous video event detection using spatiotemporal context. *Comput. Vis. Image Underst.* **115**(3), 323–333 (2011). <https://doi.org/10.1016/J.CVIU.2010.10.008>, <https://doi.org/10.1016/j.cviu.2010.10.008>
13. Junges, S., Volk, M.: Stormpy - python bindings for storm (2021)

14. Junges, S., Volk, M.: Stormpy - python bindings for storm (2021), [github.com/moves-rwth/stormpy](https://github.com/moves-rwth/stormpy)
15. Kroshchanka, A., Golovko, V., Mikhno, E., Kovalev, M., Zahariev, V., Zagorskij, A.: A neural-symbolic approach to computer vision. In: International Conference on Open Semantic Technologies for Intelligent Systems. pp. 282–309. Springer (2021)
16. Li, N., Chang, F., Liu, C.: Human-related anomalous event detection via spatial-temporal graph convolutional autoencoder with embedded long short-term memory network. *Neurocomputing* **490**, 482–494 (2022)
17. Lin, T.Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C.L., Dollár, P.: Microsoft coco: Common objects in context (2014)
18. Maaz, M., Rasheed, H., Khan, S., Khan, F.S.: Video-chatgpt: Towards detailed video understanding via large vision and language models. arXiv preprint arXiv:2306.05424 (2023)
19. Manna, Z., Pnueli, A.: The Temporal Logic of Reactive and Concurrent Systems: Specification. Springer-Verlag (1992)
20. Mavroudi, E., Haro, B.B., Vidal, R.: Representation learning on visual-symbolic graphs for video understanding. In: European Conference on Computer Vision. Lecture Notes in Computer Science, vol. 12374, pp. 71–90. Springer (2020)
21. Medioni, G.G., Cohen, I., Brémond, F., Hongeng, S., Nevatia, R.: Event detection and analysis from video streams. *IEEE Trans. Pattern Anal. Mach. Intell.* **23**(8), 873–889 (2001)
22. OpenAI: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023)
23. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision (2021)
24. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection (2016)
25. Sarkar, S., Lore, K.G., Sarkar, S.: Early detection of combustion instability by neural-symbolic analysis on hi-speed video. In: Proceedings of the NIPS Workshop on Cognitive Computation: Integrating Neural and Symbolic Approaches co-located with the 29th Annual Conference on Neural Information Processing Systems. CEUR Workshop Proceedings, vol. 1583. CEUR-WS.org, Montreal, Canada (2015)
26. SecurityInfoWatch: Data generated by new surveillance cameras to increase exponentially in the coming years (2016), <https://www.securityinfowatch.com/video-surveillance/news/12160483/data-generated-by-new-surveillance-cameras-to-increase-exponentially-in-the-coming-years>, accessed: 2023-09-30
27. Sun, P., Kretschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., Vasudevan, V., Han, W., Ngiam, J., Zhao, H., Timofeev, A., Ettinger, S., Krivokon, M., Gao, A., Joshi, A., Zhang, Y., Shlens, J., Chen, Z., Anguelov, D.: Scalability in perception for autonomous driving: Waymo open dataset (June 2020)
28. Tapaswi, M., Zhu, Y., Stiefelhagen, R., Torralba, A., Urtasun, R., Fidler, S.: Movieqa: Understanding stories in movies through question-answering. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 4631–4640. IEEE Computer Society (2016)
29. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al.: Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288 (2023)

30. Tran, D., Wang, H., Feiszli, M., Torresani, L.: Video classification with channel-separated convolutional networks. In: *IEEE/CVF International Conference on Computer Vision*. pp. 5551–5560. IEEE (2019)
31. Wu, C., Krähenbühl, P.: Towards long-form video understanding. In: *IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1884–1894. Computer Vision Foundation / IEEE (2021)
32. Xiong, Y., Huang, Q., Guo, L., Zhou, H., Zhou, B., Lin, D.: A graph-based framework to bridge movies and synopses. In: *IEEE/CVF International Conference on Computer Vision*. pp. 4591–4600. IEEE (2019)
33. Xu, Z., Yang, Y., Hauptmann, A.G.: A discriminative CNN video representation for event detection. In: *IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1798–1807. IEEE Computer Society, Boston, MA, USA (2015)
34. Yang, G., Li, M., Zhang, J., Lin, X., Ji, H., Chang, S.: Video event extraction via tracking visual states of arguments. In: Williams, B., Chen, Y., Neville, J. (eds.) *AAAI Conference on Artificial Intelligence*. pp. 3136–3144. AAAI Press, Washington, DC, USA (2023)
35. Yang, Y., Gaglione, J.R., Chinchali, S., Topcu, U.: Specification-driven video search via foundation models and formal verification. *arXiv preprint arXiv:2309.10171* (2023)
36. Yi, K., Wu, J., Gan, C., Torralba, A., Kohli, P., Tenenbaum, J.: Neural-symbolic VQA: disentangling reasoning from vision and language understanding. In: Bengio, S., Wallach, H.M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. pp. 1039–1050 (2018)
37. Yu, D., Yang, B., Wei, Q., Li, A., Pan, S.: A probabilistic graphical model based on neural-symbolic reasoning for visual relationship detection. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10599–10608. IEEE, New Orleans, LA, USA (2022)
38. and Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Li, C., Yang, J., Su, H., Zhu, J., et al.: Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499* (2023)
39. Zhang, H., Li, X., Bing, L.: Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858* (2023), <https://arxiv.org/abs/2306.02858>
40. Zheng, X., Zhang, Y., Zheng, Y., Luo, F., Lu, X.: Abnormal event detection by a weakly supervised temporal attention network. *CAAI Transactions on Intelligence Technology* **7**(3), 419–431 (2022)