

Toward Foundation Models for Online Complex Event Detection in CPS-IoT: A Case Study

Liying Han
liying98@ucla.edu
UCLA

Gaofeng Dong
gfdong@g.ucla.edu
UCLA

Xiaomin Ouyang*
xmouyang@cse.ust.hk
HKUST

Lance Kaplan
lance.m.kaplan.civ@army.mil
DEVCOM Army Research Laboratory

Federico Cerutti
federico.cerutti@unibs.it
University of Brescia

Mani Srivastava†
mbs@ucla.edu
UCLA

Abstract

Complex events (CEs) play a crucial role in CPS-IoT applications, enabling high-level decision-making in domains such as smart monitoring and autonomous systems. However, most existing models focus on short-span perception tasks, lacking the long-term reasoning required for *CE* detection. *CEs* consist of sequences of short-time *atomic events (AEs)* governed by spatiotemporal dependencies. Detecting them is difficult due to long, noisy sensor data and the challenge of filtering out irrelevant *AEs* while capturing meaningful patterns. This work explores *CE* detection as a case study for CPS-IoT foundation models capable of long-term reasoning. We evaluate three approaches: (1) leveraging large language models (LLMs), (2) employing various neural architectures that learn *CE* rules from data, and (3) adopting a neurosymbolic approach that integrates neural models with symbolic engines embedding human knowledge. Our results show that the state-space model, Mamba, which belongs to the second category, outperforms all methods in accuracy and generalization to longer, unseen sensor traces. These findings suggest that state-space models could be a strong backbone for CPS-IoT foundation models for long-span reasoning tasks.

CCS Concepts

• **Computing methodologies** → **Machine learning**; • **Computer systems organization** → **Embedded and cyber-physical systems**; • **Artificial intelligence** → **Knowledge representation and reasoning**; **Hybrid symbolic-numeric methods**.

Keywords

CPS-IoT, Foundation Models, Complex Event Reasoning, LLMs, State-space Models, Mamba, Neurosymbolic AI

ACM Reference Format:

Liying Han, Gaofeng Dong, Xiaomin Ouyang, Lance Kaplan, Federico Cerutti, and Mani Srivastava. 2025. Toward Foundation Models for Online Complex Event Detection in CPS-IoT: A Case Study. In *The 2nd International Workshop on Foundation Models for Cyber-Physical Systems & Internet of*

*This work was done while the author was at UCLA.

†The author holds concurrent appointments as an Amazon Scholar and a Professor at UCLA, but the work in this paper is unrelated to Amazon.



This work is licensed under a Creative Commons Attribution 4.0 International License. *FMSys '25, Irvine, CA, USA*

© 2025 Copyright held by the owner/author(s).

ACM ISBN /2025/05

<https://doi.org/10.1145/3722565.3727198>

Things (FMSys '25), May 6–9, 2025, Irvine, CA, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3722565.3727198>

1 Introduction

Current CPS-IoT applications excel in short-span perception tasks like human activity recognition and object detection, which require only a few seconds of sensor data. However, real-world scenarios often demand high-level contextual understanding over long periods for complex decision-making, a capability largely overlooked. For instance, an elderly person wearing a smartwatch for continuous health monitoring generates IMU data over hours. An 8-hour recording at 100 Hz produces 2.88 million data points, making it crucial to efficiently compress and retain key information to identify long-term, meaningful patterns related to health conditions.

To effectively capture these long-term patterns, it is useful to represent them as **Complex Events (CEs)**, which consist of sequences of short-time **Atomic Events (AEs)** governed by spatiotemporal dependencies. They are crucial for decision-making in CPS-IoT applications such as smart monitoring and autonomous systems that require long-term contextual understanding, as illustrated in Fig. 1(a) and (b). **Detecting CEs is challenging** for three reasons. First, the model must identify relevant *AEs* while ignoring irrelevant ones, increasing sequence-matching complexity. For instance, a sanitary protocol can be represented as “Use restroom → X → Wash hands → X → Eat,” where “X” includes irrelevant *AEs* like “walking” or “sitting,” complicating detection. Second, *CEs* involve long-term dependencies with varying temporal gaps, making it difficult to model temporal and logical rules accurately. Third, many CPS-IoT applications require online detection for real-time responses to safety-critical events, necessitating efficient and accurate inference.

Two main approaches exist for *CE* detection. Data-driven methods use neural networks to learn patterns from data but struggle with long-term dependencies due to fixed context windows or inefficient memory mechanisms. For example, Recurrent Neural Networks (RNNs) and Transformers capture context through sequential processing and attention mechanisms but face memory fading and computational inefficiency over long durations. Recent advances in state-space models (SSMs), such as Mamba [4, 8, 9], show promise by modeling event progression and capturing temporal dependencies. However, their effectiveness in noisy, high-dimensional sensor data for *CE* detection remains unexplored.

Neurosymbolic frameworks, as the second approach, combine neural networks with symbolic reasoning to leverage predefined rules for *CE* patterns, reducing the need for extensive labeled data.




Figure 1: (a) An intelligent smartwatch assistant monitors workplace hygiene and alerts users of potential violations. (b) In a smart facility, a surveillance system detects suspicious activities, such as unusual parcel hand-offs, using distributed cameras. (c) An illustration of the online CED task showing sensor processing and CE labeling.

However, symbolic engines are sensitive to noise, and probabilistic reasoning engines [12, 14, 19, 24], although more robust through probabilistic programming, require complex temporal rule implementation and demand specialized expertise.

Large Language Models (LLMs) are emerging as competitive alternatives due to their large-scale pre-training and enhanced reasoning abilities, enabling them to model long-term dependencies. Recent works like [16, 22] apply LLMs in sensor-based reasoning for CPS-IoT. However, their tasks do not address complex spatiotemporal dependencies or long temporal rule-based reasoning required for CE detection.

This work explores online CE detection as a case study for building CPS-IoT foundation models capable of long-term reasoning, enabling improved situational awareness for downstream tasks. We investigate three approaches: (1) leveraging LLMs as CE detectors, (2) employing neural architectures that learn CE rules directly from data, and (3) integrating neural models with symbolic engines embedding human knowledge. Our evaluation shows that the state-space model, Mamba, from the second category, outperforms all methods in both accuracy and generalization to longer, unseen sensor traces. These findings highlight the potential of state-space models as robust backbones for CPS-IoT foundation models designed for long-term complex event detection.

2 Related Works

Complex Event Detection (CED) has been extensively studied in traditional stream processing for structured databases. Early works [3, 6, 18] employed event engines such as Finite State Machines (FSMs) to identify complex events based on predefined patterns. More recently, research has shifted towards CED over unstructured, high-dimensional data. Approaches such as [17, 21] integrate neural detectors with symbolic rule-based systems, enabling back-propagation through logical rules. However, these methods suffer from computational scalability issues. Moreover, most existing approaches focus on complex activities within short time spans [7, 10, 11], limiting their applicability to longer sensor traces with extended temporal dependencies.

State-Space Models (SSMs) capture long-span temporal dependencies through state transitions, maintaining state representations that implicitly encode temporal progression. Recent advances, such

as Mamba [4, 8, 9], show superior performance in long-term contextual understanding by using selection mechanisms to process sequences more efficiently and adaptively model different temporal dependencies. However, SSMs remain underexplored in noisy, high-dimensional sensor environments, particularly for long-span temporal events in CPS-IoT applications.

3 Online Complex Event Detection

3.1 Complex Event Definitions

Definition 3.1. Atomic events (AEs) are short-duration, low-level events that serve as building blocks for complex events. They occur instantaneously or within a small time window and are directly detectable by models like image classification, object detection, or activity recognition.

Definition 3.2. Complex events (CEs) are high-level events defined as sequences or patterns of atomic events (AEs) occurring in specific temporal or logical relationships.

Let $A = a_1, a_2, \dots, a_n$ be the set of all AEs, where each a_i has a start and end time. Similarly, let $E = e_1, e_2, \dots, e_k$ be the set of all CEs of interest.

Each complex event (CE) $e_i \in E$ is defined as:

$$e_i = R_i(A_i) = R_i(a_i^1, a_i^2, \dots, a_i^{n_i}),$$

where $A_i \subseteq A$ is the subset of AEs relevant to e_i , R_i is a *pattern function* defining the temporal or logical relationship among the AEs in A_i , and $n_i = |A_i|$ is the number of AEs involved in defining e_i . Each e_i has a time t_{e_i} , which is the moment (or interval) when the pattern R_i is satisfied, indicating the occurrence of e_i .

Pattern Function (R_i) maps A_i to e_i by defining patterns among the AEs. We consider four main categories, some of which include subcategories:

- (1) **Sequential Patterns: Relaxed Order** – Key AEs must occur in order but may include unrelated AEs in between.
- (2) **Temporal Patterns:** (a) *Duration Based* – Measuring the duration of specific AEs; (b) *Timing Relationship* – Capturing relative timing constraints, such as min and max delays between AEs.
- (3) **Repetition Patterns:** (a) *Frequency Based* – Counting occurrences of AEs over a time window; (b) *Contextual Count* – Counting AEs relative to other AEs' timing.

(4) **Combination Patterns:** Merging *Sequential* and *Temporal Patterns* to express complex relationships.

Importantly, all patterns considered in this work are *bounded to finite states*, enabling representation by finite state machines (FSMs).

3.2 Online Detection Task Formalization

Assume a system receives a raw data stream \mathbf{X} from a single sensor with modalities M at a sampling rate r . The system processes the stream using a non-overlapping sliding window of length Δt . At the t th window, the data segment is:

$$\mathbf{D}_t = \mathbf{X}(t), \quad \mathbf{X}(t) \in \mathbb{R}^{(r \times \Delta t) \times m} \quad (1)$$

where m is the feature dimension of the sensor data of modality M .

Each window t has a corresponding ground-truth *CE* label y_t , which depends on the *AEs* from previous windows $t - 1$ and the current window t . As illustrated in Fig. 1(c), if a *CE* spans from t_1 to t_2 , then only y_{t_2} is labeled as the *CE*, while all y_{t_1} to y_{t_2-1} are labeled as “0” to indicate no *CE* is detected before t_2 .

For up to T sliding windows, the objective of the real-time CED model f is to minimize the difference between the predicted *CE* label \hat{y}_t and the ground-truth label y_t at every window:

$$\min |\hat{y}_t - y_t|, \quad \text{where } \hat{y}_t = f(\mathbf{D}_t), \quad 1 \leq t \leq T, \quad (2)$$

This forms a *multi-label multi-class classification* problem, where the objective can be expressed in vector form as:

$$\min \|f(\mathbf{D}) - \mathbf{y}\|, \quad \text{where } \mathbf{D} = \{\mathbf{X}(1), \dots, \mathbf{X}(T)\}. \quad (3)$$

For data-driven methods, the task uses **high-level, coarse CE labels** without fine-grained *AE* labels during training, requiring the model to learn *AE* semantics and *CE* rules simultaneously. This represents a **unique and challenging** combination of *distant supervision* (event-level only labeling) and *weak supervision* (sparse and high-level labeling).

3.3 Multimodal CED Dataset

As no large-scale dataset exists for online CED, we develop a multimodal dataset in a smart health monitoring setting, leveraging multimodal sensors for rich contextual information in realistic *CE* tasks. The dataset includes 10 *CE* classes spanning various categories, with detailed rule definitions in Table 1.

3.3.1 Sensor Data. To generate *CE* sensor traces, we built a stochastic simulator that mimics daily human behaviors by generating a random *AE* label every 5-second window following realistic distributions. We use 9 *AE* classes: “walk”, “sit”, “brush”, “click mouse”, “drink”, “eat”, “type”, “flush toilet”, and “wash”. To ensure the generated *AE* sequences are consistent with the *CE* patterns, we define *AE* transitions between windows. For example, to mimic the scenario where a person forgets to clean hands after using the restroom, “wash” has a 0.7 probability after “flush_toilet”. We also include traces without *CE* patterns to balance the dataset and prevent biased *CE* rule learning, ensuring the model is exposed to scenarios with or without each *CE*.

3.3.2 Labeling. These *AE* traces are converted to sensor traces by sampling the corresponding segments reflecting the activity from **WISDM** [20] for IMU data and **ESC-70** [15] for audio data. Ground-truth *CE* labels are generated using human-defined finite

Table 1: Complex Event Classes and Definitions

Complex Events	Definitions	Category
Default (e_0)	When no complex events of interest take place.	
Workspace sanitary protocol violation (e_1)	A violation occurs if a person starts working (click or type) without 20 seconds of consecutive handwashing after using the restroom. Upon violation, trigger an alert and reset the system.	Sequential + Temporal
Sanitary habit violation (e_2)	Hands are not cleaned if no 20-second consecutive handwashing occurs within 2 minutes before a meal session. A meal session starts when eating or drinking begins and ends when any activity other than eat, drink, or sit is detected.	Sequential + Temporal
Inadequate brushing time (e_3)	Brushing teeth for less than 2 minutes. If brushing stops, wait for 10 seconds; otherwise, report violation and reset the system.	Temporal (Relative + Duration)
Routine Sequence (e_4)	brush $\rightarrow u^* \rightarrow \text{eat} \rightarrow u^* \rightarrow \text{drink} \mid \text{brush} \rightarrow u^* \rightarrow \text{drink} \rightarrow u^* \rightarrow \text{eat}$, where $u = A \setminus \{\text{brush, eat, drink}\} \dagger$.	Sequential - Related
Start working and then take a break (e_5)	sit $\rightarrow u^* \rightarrow \text{type/click} \rightarrow v^* \rightarrow \text{walk}$, where $u = A \setminus \{\text{sit, type, click, walk}\}$, and $v = A \setminus \{\text{type, click, walk}\} \dagger$.	Sequential - Related
Sufficient Washing Reminder (e_6)	When washing lasts for 30 seconds consecutively.	Temporal - Duration
Adequate brushing time (e_7)	When brushing lasts a total of 2 minutes. The timer pauses if brushing stops but resumes if brushing restarts. Once the 2-minute threshold is reached, the event is reported, and the timer resets.	Temporal (Relative + Duration)
Post-Meal Rest (e_8)	After eating, wait for at least 3 minutes to work.	Temporal - Relative
Active Typing Session (e_9)	The event occurs if at least 3 typing sessions (start typing, stop typing) happen within 60 seconds of the first session’s start.	Repetition - Frequency
Focused Work Start (e_{10})	The event is triggered by sitting after being seated, as long as no walking occurs during this time. The event is reported after exactly 5 clicks after sitting and before walking.	Repetition - Contextual

Notes: †Here A represents the set of all *atomic events*.

state machines (FSMs), one for each *CE* class. The simulator first generates *AE* labels for each sequence, which are then processed by FSMs to determine the corresponding online *CE* labels.

3.3.3 Training & Test Data. The dataset consists of 5-minute *CE* sensor traces synthesized using inertial and audio data from multiple subjects, with each *CE* sequence containing 60 windows ($5\text{min} \times 60\text{sec}/\text{min} \div 5\text{sec} = 60$). We generated 10,000 training and 2,000 validation examples, including sequences with no *CE* occurrences to reflect real-world sparsity. The test dataset uses sensor traces from an unseen subject and includes 2,000 examples of 5-minute sequences. Additionally, we created out-of-distribution (OOD) datasets with 2,000 examples each for 15-minute and 30-minute sequences, maintaining the same *CE* patterns but with longer *AE* durations and wider temporal gaps between key *AEs* to introduce additional challenges for generalization.

4 Online CED Pipeline Overview

We propose a two-module online processing system as shown in Fig. 2. The system uses a non-overlapping sliding window of size W (5 seconds) to segment raw sensor streams into data segments $\{s_t\}_{t=1}^T$, which are then processed by the following modules.

4.1 Pretrained Feature Encoder

This module encodes each segment s_t into a high-dimensional embedding vector v_t , producing a sequence of embedding vectors $\{v_t\}_{t=1}^T$. The encoder is pretrained to extract latent features

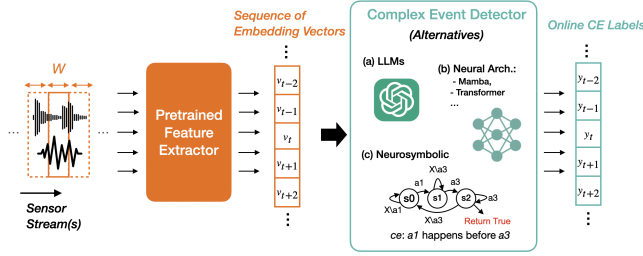


Figure 2: Overview of the online CED pipeline.

from raw data and then frozen to ensure consistent embeddings across all downstream complex event detector models, enabling fair comparisons on *CE* pattern reasoning. For our multimodal CED task, we develop a multimodal encoder for s_t using early fusion. We use BEATs [2] and LIMU-Bert [23] to extract audio and IMU features, which are then processed through separate GRU layers, concatenated and fused in a fusion layer to form a 128-dimensional embedding v_t . An MLP layer is added as the *AE* classifier head for supervised training with cross-entropy loss. The encoder is then frozen, and the joint embedding v_t is used for *CE* detectors.

4.2 Complex Event Detector

This module takes the embedding vectors $\{v_t\}_{t=1}^T$ and detects *CE* patterns, producing an online *CE* label sequence $\{y_t\}_{t=1}^T$, where y_t is the *CE* label at window t . This component supports **LLM-based, neural, and neurosymbolic alternatives** for comparison. Let the pretrained feature encoder be m , and the complex event detector be g . The online CED task objective defined in Eq. 2 becomes:

$$\min |\hat{y}_t - y_t|, \quad \text{where } \hat{y}_t = g(m(D_t)), \quad 1 \leq t \leq T. \quad (4)$$

Here, the feature encoder m remains fixed, while the complex event detector g is varied to compare different methods and architectures.

5 Exploring LLMs as *CE* Detectors

5.1 Simplified Task

We evaluate LLMs as the complex event detector g using ground truth *AE* sequences as inputs, replacing the encoder m with a perfect *AE* classifier. This **simplifies** the CED task by hiding sensor noise to focus solely on *CE* reasoning. We further simplify by using only three *CE* (e_1 , e_2 , and e_3). Each *AE* label represents the activity in a 5-second window (e.g., [“walk”, “walk”, “sit”] means walking for 10 seconds and then sitting). These events occur over 5-minute intervals, yielding 60 *AE* labels per example. LLMs also receive the CED task context and *CE* definitions from Table 1, and predict a *CE* label every 5-second window. We evaluate 2,000 examples.¹

5.2 Evaluation

5.2.1 *Metrics.* We evaluate LLM performance using three metrics:

- **Length Accuracy:** Measures the match rate between the length of the predicted *CE* labels and input *AE* sequences in the “ n -to- n ” sequence prediction task required for the Online CED task.
- **Conditional F1 Score:** Calculates element-wise *F1* score for complex event labels, assessing both event type and timing accuracy,

¹Prompts available: /r/LLM-CED-Prompts-CD6C/

Table 2: Evaluation results of LLMs.

	Length Acc.	Coarse F1				Conditional F1			
		e_1	e_2	e_3	Avg.	e_1	e_2	e_3	Avg.
<i>Zero-shot</i>									
Qwen2.5-14B	0.12	0.60	0.66	0.57	0.61	0.14	0.15	0.04	0.11
GPT-4o-mini	0.04	0.0	0.0	0.64	0.21	0.0	0.0	0.0	0.0
GPT-4o	0.12	0.87	0.78	0.80	0.82	0.14	0.59	0.03	0.25
o1-mini	0.13	0.78	0.84	0.80	0.80	0.0	0.72	0.02	0.25
o3-mini	0.33	0.90	0.85	0.88	0.87	0.05	0.78	0.24	0.35
<i>Few-shot (k = 3)</i>									
Qwen2.5-14B	0.14	0.62	0.66	0.60	0.63	0.13	0.13	0.03	0.10
GPT-4o-mini	0.03	0.0	0.0	0.58	0.19	0.0	0.0	0.0	0.0
GPT-4o	0.16	0.87	0.81	0.81	0.83	0.13	0.63	0.14	0.30
o1-mini	0.25	0.81	0.84	0.82	0.82	0.0	0.71	0.17	0.29
o3-mini	0.44	0.94	0.86	0.88	0.89	0.33	0.84	0.35	0.50

conditioned on sequences with the correct length T and averaged over timestamps 1 to T .

- **Coarse F1 Score:** Evaluates *CE* labeling at a high level by checking if the correct event type is recognized within a 5-minute sample, without requiring a precise timestamp match.

We test LLMs on both zero-shot and few-shot tasks, with three input-output examples included in the prompt for few-shot experiments.

5.2.2 *Results.* We evaluated five SOTA LLM models on complex events 1, 2, and 3, as shown in Table 2. All models performed poorly, even on the simplified CED task. The low length accuracy suggests hallucination in long-chain reasoning, as LLMs often generated sequences between 55 and 65 labels instead of the required 60. Since the *F1* score cannot be computed when sequence lengths mismatch, we calculated the conditional *F1* score only for outputs with the correct length. Among zero-shot models, o3-mini performed slightly better but remained unsatisfactory. We also computed the coarse *F1* score, which only evaluates whether a *CE* occurred within a 5-minute window without requiring exact timing. In this relaxed setting, LLMs showed significant improvement, suggesting they can partially recognize *CE* patterns but struggle with precise timing. Few-shot examples improved performance for stronger reasoning models like o3-mini but had little effect on others, indicating that o3-mini better utilizes *CE* examples for self-checking its understanding. Although we did not evaluate o1 and o3 models, which may have stronger reasoning abilities, their high cost and inference time make them impractical for online CED, which requires frequent, low-latency API requests. In summary, while LLMs show potential for online CED, they suffer from hallucinations, poor long-chain reasoning, and high latency due to their transformer-based architecture, leading to delays in real-time inference for longer sensor traces.

6 Neural and Neurosymbolic Methods

Given the limitations of LLMs, we design and train various neural and neurosymbolic architectures as alternatives for the complex event detector g following some requirements.

6.1 Architectures

6.1.1 Requirements. Causal Structure: Online CED requires models to predict complex events at each time step t using only past observations (0 to t) to ensure no future information is accessed. **Minimum Receptive Field:** Neural networks must have a receptive field larger than the longest temporal patterns of complex events in the training data to capture full event patterns.

6.1.2 Variants. We explore three types of architectures:

End-to-end Neural Architectures. These models take high-dimensional sensor embeddings as input and are trained end-to-end for CE detection. We compare: (1) a 5-layer *Unidirectional LSTM* (hidden size 256, $\approx 2.5M$ parameters), (2) a *Causal TCN* [1] with dilation rate 32, kernel size 3, and 128 filters per layer (8-min receptive field, $\approx 4.6M$ parameters), (3) a 6-layer *Causal Transformer* with triangular attention mask and positional encoding (hidden size 128, 8-head attention, $\approx 4.2M$ parameters), and (4) a 12-block *Mamba* state-space model (hidden size 128, state size 64, $\approx 1.8M$ parameters). This configuration is chosen because the original Mamba paper [8] shows that two SSM blocks are equivalent to one Transformer layer.

Two-stage Concept-based Architectures. Denoted as *Neural AE + X*, these models first use a *neural AE* classifier to map each window of sensor embedding to an *AE* class, which are then processed by a neural backbone model X for CE detection. Variants include *Neural AE + LSTM*, *Neural AE + TCN*, *Neural AE + Transformer*, and *Neural AE + Mamba*. The backbone models are identical to those in the end-to-end architectures but take one-hot *AE* class traces as input.

Neurosymbolic Architecture: We design a neurosymbolic model, *Neural AE + FSM*, that integrates human knowledge of CE rules. It uses the same neural *AE* classifier and a user-defined symbolic reasoner, employing an FSM for each complex event rule, identical to those used in CE labeling. We also explored a probabilistic FSM in ProbLog [5], leveraging softmax embedding for probabilistic reasoning over CE sequences. However, it showed only marginal improvements and was excluded due to design complexity and reliance on expert knowledge.

6.2 Training Loss

The online CED task faces severe class imbalance due to the temporal sparsity of CE labels, where the majority class is “0” (no event), dominating the label distribution. To address this, we use Focal Loss (FL) [13], which emphasizes errors on rare but critical classes. The loss is defined as:

$$\min_{\theta} L_{FL}(\theta) = - \sum_{i=1}^N \sum_{t=1}^T \alpha_{y_i(t)} (1 - p_{y_i(t)})^{\gamma} \log(p_{y_i(t)}), \quad (5)$$

where $p_{y_i(t)}$ is the predicted probability of class y_i at time t , γ reduces the impact of frequent classes, and α_y balances class weights. After a hyperparameter grid search, we set $\gamma = 2$, $\alpha_0 = 0.005$ for the most frequent class “0”, and $\alpha_y = 0.25$ for rare but critical CE classes.

6.3 Evaluation

6.3.1 Experimental Setup. All neural models are trained using the AdamW optimizer with Focal Loss, a learning rate of 1×10^{-3} ,




Figure 3: Boxplot of positive F_1 scores on complex events with different training data sizes. Solid line in the box shows median; dashed line in the box shows mean.

weight decay of 0.1, and batch size of 256. Early stopping is applied based on validation loss, with a maximum of 5000 training epochs. Results are averaged over 10 random seeds. The *Neural AE* classifier used for *Neural AE + X* models is trained with the pretrained feature encoder, achieving 95% *AE* classification accuracy on the test set.

6.3.2 Metrics. We evaluate performance using the F_1 score for each CE class e_i and report *Positive F1* ($F1_{pos}$), which is the average F_1 over positive event classes (e_1 to e_{10}), excluding the less important “negative” label e_0 . A higher F_1 score indicates a better precision-recall balance, reflecting correctness and completeness.

6.3.3 Results. We evaluate model performance on CE s defined in Table 1 across different training set sizes, as shown in Fig. 3. The results indicate that *Mamba achieves the best performance*, followed by LSTM. The *Neural + X* models underperform compared to end-to-end models, likely due to errors and noise introduced by the *Neural AE* classifier. This also explains why *Neural AE + FSM* performs worse despite incorporating correct human-defined CE rules. Additionally, we test model generalization on out-of-distribution (OOD) complex events lasting 15 and 30 minutes, which follow the same CE rules but extend their temporal spans. As shown in Fig. 4, Mamba exhibits the best generalization among both end-to-end neural models and *Neural AE + FSM*. Table 3 further shows that *Mamba, when trained with more labeled sensor data on 5-minute sequences, improves generalization to longer unseen traces*, suggesting that it efficiently learns CE rules as training data increases. These findings indicate that Mamba may serve as a strong backbone for foundation models targeting long and complex temporal reasoning.

7 Conclusion

We formalized the online CED task and developed a dataset in a CPS-IoT scenario to evaluate various methods. Although LLMs showed potential, they failed on the simplified CED task due to poor precise reasoning over long temporal patterns. We then explored neural and neurosymbolic alternatives as CE detectors. The end-to-end Mamba model outperformed all others, demonstrating smaller




Figure 4: Positive F1 scores of different models tested on 5-min, 15-min, and 30-min CE sensor data.

Table 3: Positive F1 scores of end-to-end Mamba model tested on 5-min, 15-min, and 30-min CE sensor data (with a 2-sigma confidence interval).

Training Data Size	Positive F1		
	5min	15min (OOD)	30min (OOD)
2000	.79 ± .04	.68 ± .06	.55 ± .07
4000	.88 ± .08	.78 ± .09	.69 ± .14
6000	.88 ± .06	.78 ± .09	.70 ± .09
8000	.90 ± .08	.81 ± .12	.74 ± .13
10000	.92 ± .06	.82 ± .08	.75 ± .09

model size and better generalization to unseen longer CE traces. This highlights its effectiveness in online CED tasks and its potential as a suitable backbone for complex temporal reasoning foundation models in CPS-IoT applications.

Acknowledgments

This research was sponsored in part by the DEVCOM ARL (award # W911NF1720196), the AFOSR (awards # FA95502210193 and # FA95502310559), the NIH (award # 1P41EB028242), the NSF (award # 2325956), and the EOARD (award # FA8655-22-1-7017).

References

- [1] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. 2018. An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. *CoRR* abs/1803.01271 (2018). arXiv:1803.01271 <http://arxiv.org/abs/1803.01271>
- [2] Sanyuan Chen, Yu Wu, Chengyi Wang, Shujie Liu, Daniel Tompkins, Zhuo Chen, and Furu Wei. 2022. BEATs: Audio Pre-Training with Acoustic Tokenizers. arXiv:2212.09058 [eess.AS]
- [3] Gianpaolo Cugola and Alessandro Margara. 2012. Processing Flows of Information: From Data Stream to Complex Event Processing. *ACM Comput. Surv.* 44, 3, Article 15 (jun 2012), 62 pages. doi:10.1145/2187671.2187677
- [4] Tri Dao and Albert Gu. 2024. Transformers are SSMs: Generalized Models and Efficient Algorithms Through Structured State Space Duality. arXiv:2405.21060 [cs.LG] <https://arxiv.org/abs/2405.21060>
- [5] Luc De Raedt, Angelika Kimmig, and Hannu Toivonen. 2007. ProbLog: a probabilistic prolog and its application in link discovery. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (Hyderabad, India) (IJ-CAI'07)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2468–2473.
- [6] Hervé Debar and Andreas Wespi. 2001. Aggregation and Correlation of Intrusion-Detection Alerts. In *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 85–103. doi:10.1007/3-540-45474-8_6
- [7] Bernard Ghanem Fabian Caba Heilbron, Victor Escorcía and Juan Carlos Nibbles. 2015. ActivityNet: A Large-Scale Video Benchmark for Human Activity Understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 961–970.
- [8] Albert Gu and Tri Dao. 2024. Mamba: Linear-Time Sequence Modeling with Selective State Spaces. arXiv:2312.00752 [cs.LG] <https://arxiv.org/abs/2312.00752>
- [9] Albert Gu, Karan Goel, and Christopher Ré. 2022. Efficiently Modeling Long Sequences with Structured State Spaces. arXiv:2111.00396 [cs.LG] <https://arxiv.org/abs/2111.00396>
- [10] Y.-G. Jiang, J. Liu, A. Roshan Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar. 2014. THUMOS Challenge: Action Recognition with a Large Number of Classes. <http://csrcv.ucf.edu/THUMOS14/>.
- [11] Salman Khan, Izzeddin Teeti, Andrew Bradley, Mohamed Elhoseiny, and Fabio Cuzzolin. 2023. A Hybrid Graph Network for Complex Activity Detection in Video. arXiv:2310.17493 [cs.CV] <https://arxiv.org/abs/2310.17493>
- [12] Ziyang Li, Jiani Huang, and Mayur Naik. 2023. Scallop: A Language for Neurosymbolic Programming. arXiv:2304.04812 [cs.PL] <https://arxiv.org/abs/2304.04812>
- [13] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. 2017. Focal Loss for Dense Object Detection. *CoRR* abs/1708.02002 (2017). arXiv:1708.02002 <http://arxiv.org/abs/1708.02002>
- [14] Robin Manhaeve, Sebastijan Dumancic, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. 2018. DeepProbLog: Neural Probabilistic Logic Programming. In *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2018/file/dc5d637ed5e62c36ecb73b654b05ba2a-Paper.pdf>
- [15] Marc Moreaux, Michael Garcia Ortiz, Isabelle Ferrané, and Frederic Lerasle. 2019. Benchmark for Kitchen20, a daily life dataset for audio-based human action recognition. In *2019 International Conference on Content-Based Multimedia Indexing (CBMI)*. 1–6. doi:10.1109/CBMI.2019.8877429
- [16] Xiaomin Ouyang and Mani Srivastava. 2024. LLMsense: Harnessing LLMs for High-level Reasoning Over Spatiotemporal Sensor Traces. arXiv:2403.19857 [cs.AI] <https://arxiv.org/abs/2403.19857>
- [17] Marc Roig Vilamala, Tianwei Xing, Harrison Taylor, Luis Garcia, Mani Srivastava, Lance Kaplan, Alun Preece, Angelika Kimmig, and Federico Cerutti. 2023. DeepProbCEP: A neuro-symbolic approach for complex event processing in adversarial settings. *Expert Systems with Applications* 215 (2023), 119376. doi:10.1016/j.eswa.2022.119376
- [18] Nicholas Schultz-Møller, Matteo Migliavacca, and Peter Pietzuch. 2009. Distributed Complex Event Processing with Query Rewriting. In *Proceedings of the Third ACM International Conference on Distributed Event-Based Systems*. doi:10.1145/1619258.1619264
- [19] Arseny Skryagin, Wolfgang Stammer, Daniel Ochs, Devendra Singh Dhami, and Kristian Kersting. 2021. SLASH: Embracing Probabilistic Circuits into Neural Answer Set Programming. arXiv:2110.03395 [cs.AI] <https://arxiv.org/abs/2110.03395>
- [20] Gary Weiss. 2019. WISDM Smartphone and Smartwatch Activity and Biometrics Dataset. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5HK59>.
- [21] Tianwei Xing, Luis Garcia, Marc Roig Vilamala, Federico Cerutti, Lance Kaplan, Alun Preece, and Mani Srivastava. 2020. Neuroplex: Learning to Detect Complex Events in Sensor Networks through Knowledge Injection. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems (Virtual Event, Japan) (SenSys '20)*. Association for Computing Machinery, New York, NY, USA, 489–502. doi:10.1145/3384419.3431158
- [22] Huatao Xu, Liying Han, Qirui Yang, Mo Li, and Mani Srivastava. 2024. Penetrative AI: Making LLMs Comprehend the Physical World. arXiv:2310.09605 [cs.AI] <https://arxiv.org/abs/2310.09605>
- [23] Huatao Xu, Pengfei Zhou, Rui Tan, Mo Li, and Guobin Shen. 2021. LIMU-BERT: Unleashing the Potential of Unlabeled Data for IMU Sensing Applications. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems (Coimbra, Portugal) (SenSys '21)*. Association for Computing Machinery, New York, NY, USA, 220–233. doi:10.1145/3485730.3485937
- [24] Zhun Yang, Adam Ishay, and Joohyung Lee. 2023. NeurASP: Embracing Neural Networks into Answer Set Programming. arXiv:2307.07700 [cs.AI] <https://arxiv.org/abs/2307.07700>