

Using DeepProbLog to perform Complex Event Processing on an Audio Stream

Marc Roig Vilamala^{1,*}, Tianwei Xing², Harrison Taylor¹, Luis Garcia², Mani Srivastava²,
Lance Kaplan³, Alun Preece¹, Angelika Kimmig⁴, Federico Cerutti^{1,5}

¹Cardiff University

²University of California, Los Angeles

³Army Research Laboratory

⁴KU Leuven, Department of Computer Science; Leuven.AI

⁵University of Brescia

*Corresponding author: RoigVilamalaM@cardiff.ac.uk

Abstract

In this paper, we present an approach to Complex Event Processing (CEP) that is based on DeepProbLog. This approach has the following objectives: (i) allowing the use of subsymbolic data as an input, (ii) retaining the flexibility and modularity on the definitions of complex event rules, (iii) allowing the system to be trained in an end-to-end manner and (iv) being robust against adversarial conditions. Our approach makes use of DeepProbLog to use a hybrid neuro-symbolic architecture that combines a neural network to process the symbolic data with a probabilistic logic layer to allow the user to define the rules for the complex events. We demonstrate that our approach is capable of detecting complex events from an audio stream. We also demonstrate that our approach is fairly robust against adversarial conditions by training it with datasets under different levels of poisoning attacks.

1 Introduction

Complex Event Processing (CEP) systems process data streams and detect situations of interest, or *complex events*, which aggregate atomic *events*, or *simple events*. CEP systems detect spatio-temporal relationships between sets of simple events, which form complex events. CEP systems have been applied in many different areas, such as business activity monitoring (Teymourian, Rohde, and Paschke 2012), sensor networks (Anicic et al. 2012b) and weather reports (Anicic et al. 2012a). Most CEP approaches allow the user to define rules which express the conditions under which a complex event occurs. Then, the CEP system uses those rules to detect when those circumstances happen in the given stream of input data. However, defining rules over raw streams of data can be challenging. For example, it is not feasible to define rules directly over raw images, audios or videos. In this paper, we will refer to these types of data for which we cannot (easily) manually define rules to extract the information we want as *subsymbolic data*.

Some new CEP approaches (Roldán et al. 2020; Roig Vilamala et al. 2019) have been created to incorporate the use of subsymbolic data. However, as we will discuss in Section 3, they require pre-trained neural networks to work, which are not always available. While it is possible to train these neural networks separately, it can be costly to obtain training data

for that case. As such, we want an approach that can train in an end-to-end manner. This means that we want a system that can be trained using only labels for the complex events. While some approaches already allow for such end-to-end training (Xing et al. 2020), they significantly limit the flexibility and modularity offered when defining the rules for the complex events. This makes it difficult, or even impossible, for the user to precisely express the conditions under which a complex event occurs, particularly for the more complex situations in which it may happen. In this paper, we aim to propose an approach to CEP that can be trained to use new types of subsymbolic data without limiting the flexibility and modularity of the rule definitions.

Another aim for our approach was to make it robust against adversarial conditions, according to current AI ethical principles (Board 2019). These indicate that AI systems should be reliable, meaning that the systems should behave as expected even in sub-optimal conditions. For this paper, we will be focusing on situations in which the training data is noisy, meaning that part of the labels in the training dataset are incorrect.

As such, we wanted to create an AI system that is capable of performing CEP while fulfilling the following objectives:

1. Being able to operate on subsymbolic data streams.
2. Retaining flexibility and modularity in rule definitions.
3. Being able to perform end-to-end training.
4. Being robust against adversarial conditions.

Currently, none of the approaches to this type of problem cover all such objectives. In Section 3 we will explain the limitations of existing approaches.

In this paper, we propose an approach based on DeepProbLog (Manhaeve et al. 2018; 2021) to detect complex events from an audio stream. DeepProbLog allows us to combine a neural network with probabilistic logic rule definitions. As such, the neural network can be used to process the subsymbolic data, which can then be used within the probabilistic logic to detect the patterns that form complex events. Furthermore, the probabilistic logic allows users to easily define the rules for the complex events. DeepProbLog also allows us to train the system in an end-to-end manner, thus fulfilling the first three objectives. For a background explanation of DeepProbLog, see Section 2. Meanwhile, Sec-

tion 4 explains how we have used DeepProbLog to perform complex event detection.

In order to evaluate the performance of our approach we have generated synthetic datasets. For more details on the dataset generation, see Section 5. Then, in Section 6 we evaluate the performance of our approach after training with the generated datasets. First, we demonstrate that our approach is capable of detecting complex events from an audio stream. Then, we also demonstrate that our approach is fairly robust against a poisoning attack on the training data. More specifically, we demonstrate that it is able to reach almost the same performance even if 20% of the training data is incorrectly labelled.

Finally, in Section 7 we provide conclusions on the results obtained in this paper and discuss potential areas for future research.

2 Background

In this section, we provide background information on complex event processing (CEP). We also give a general overview of ProbLog and DeepProbLog, which are used in our approach.

2.1 Complex Event Processing

Following (Luckham 2002), an *event* is an object that can be subjected to computer processing and it signifies, or is a record of, an activity that has happened. Each *event* has three main aspects:

- *Form*: the form of an event is an object with particular attributes or data components, for instance the time period of the activity;
- *Significance*: an event signifies an activity, hence an event’s form usually contains data describing the activity it signifies;
- *Relativity*: an activity is related to other activities. Events have the same relationships to one another as the activities they signify. The *relativity* of an event refers to the set of relationships between that event and other events. An event’s form usually encodes its relativities, i.e., methods to reconstruct the relationships with other events.

It is therefore important to notice that an event is not just a message or a record of an activity: the forms of events may be messages, but the events also have significance and relativity. In particular, the three main partial, transitive, and antisymmetric relationships between events are:

- *Time*: a relationship that orders events.
- *Cause*: a dependence relationship between activities. An activity (event) depends upon other activities (events) if it happened only because the other activities (events) happened. If event B depends upon event A , then A *caused* B . If neither caused the other, they are *independent*.¹
- *Aggregation*: if event A signifies an activity that consists of the activities of a set of events B_1, B_2, \dots, B_n , then A is an *aggregation* of all the events B_i . Conversely, B_i are *members* of A . Aggregation is an abstraction relationship:

¹This computational notion of causality is ostensibly more limited than the notion of causality in philosophy and science in general: an interested reader is referred to (Pearl 2009)

usually event A is created when a set of events $\{B_i\}$ happens. A is a higher-level event and we call it a *complex event*. A ’s members are the events that *caused* it. Aggregation can be referred to also as *vertical causality*.

In the context of this paper, CEP aims at identifying such aggregation rules, so to make the activities in a complex system understandable to humans. More specifically, we will be using an audio feed as the source of the simple events, which will take the form of short segments of audio (1 second long). The incoming events will have a timestamp attached to them, which will indicate at which point in time the corresponding segment of audio started. As an example, in this paper we will be aggregating events into a complex event if the same type of sound (based on a set of pre-defined classes) repeats within a certain window of time. However, other aggregation rules could be used with our approach to fit the intended functionality.

2.2 ProbLog

ProbLog (De Raedt, Kimmig, and Toivonen 2007) is a probabilistic logic programming language. ProbLog allows users to encode complex interactions between different components. A ProbLog program consists of a set of probabilistic facts F and a set of rules R . Facts have the form $p :: f$ where p is a value between 0 and 1, which represents the likelihood of the fact being true, and f is an atom. Atoms are expressions of the form $q(t_1, \dots, t_n)$ where q is a predicate and t_i are terms. Rules have the form $h :- b_1, \dots, b_n$ where h is an atom and b_i are literals. A literal is an atom or the negation of an atom.

One convenient syntactic extension is an annotated disjunction (AD), which is an expression of the form $p_1 :: h_1; \dots; p_n :: h_n :- b_1, \dots, b_m$, where the p_i are probabilities so that $\sum p_i = 1$, and h_i and b_j are atoms. The meaning of an AD is that whenever all b_i hold, h_j will be true with probability p_j , with all other h_i false (unless other parts of the program make them true). This is convenient to model choices between different categorical variables. ProbLog programs with annotated disjunctions can be transformed into equivalent ProbLog programs without annotated disjunctions (De Raedt and Kimmig 2015).

2.3 DeepProbLog

DeepProbLog (Manhaeve et al. 2018; 2021) is a neural probabilistic logic programming language that allows the user to create hybrid neuro-symbolic architectures. DeepProbLog allows the user to train the neural networks in these architectures as part of the system in an end-to-end manner.

A DeepProbLog program is a ProbLog program that is extended with a set of ground neural ADs (nADs) of the form $nn(m_q, [X_1, \dots, X_k], O, [y_1, \dots, y_n]) :: q(X_1, \dots, X_k, O)$. Here, nn indicates that the following is an nAD and m_q is a neural network identifier. The neural network m_q will be provided the input vector $[X_1, \dots, X_k]$ and output a probability distribution over the domain $O \in [y_1, \dots, y_n]$. nADs work similarly to ADs in the sense that they provide a mutually-exclusive distribution of probabilities over a set of atoms. In nADs, however, these probabilities are generated from the output

of a neural network, instead of being manually defined. The sum of the probabilities over the domain O must equal 1. In neural networks for multiclass classification, this is typically done by applying a softmax layer to real-valued output scores, a choice we also adopt in our experiments.

After defining the structure of the neural network and the logic level, it is possible to use DeepProbLog to infer the answers to our queries. To perform this inference, DeepProbLog transforms the logic layer into an arithmetic circuit and obtains the required probabilities from the neural network. This arithmetic circuit can then be used to calculate the probability that the query is true, based on the output of the neural network.

In order to train the neural network, the system first performs inference as described above. Then, DeepProbLog is able to perform gradient-based learning. First, the arithmetic circuit used during the inference is also used to perform the gradient computations. Since this arithmetic circuit is composed of addition and multiplication operations, this means that it is differentiable. This allows DeepProbLog to compute the gradient with respect to the probabilistic logic program. This gradient can then be used to train the neural network using backpropagation. For a more detailed explanation on the technical aspects of DeepProbLog’s inference and learning, see (Manhaeve et al. 2021).

3 Critical analysis of related work

In this section, we will explore the existing CEP approaches that are able to use subsymbolic data. There are three main types of approaches: (i) using pre-trained neural networks to extract the symbolic information from the subsymbolic data, (ii) using a purely statistical approach and (iii) hybrid neuro-symbolic approaches. In the following sections we will describe the limitations of each of those approaches, which our approach aims to solve.

3.1 Pre-trained neural networks approaches

Some CEP approaches use a pre-trained neural network to transform high-bandwidth data into symbolic information, allowing the user to define rules on it. For example, in (Roldán et al. 2020) the authors show that this allows them to reduce the number of false positives in a system when detecting IoT security attacks. They use a neural network to predict the length of the suspected packets. If the predicted length does not match the actual length of the packet, a complex event is generated indicating that an attack might be happening. In (Roig Vilamala et al. 2019), we present a system that is able to detect different violent activities from a CCTV feed. A pre-trained neural network is used to process short segments of video (16 frames, about half a second) detecting potential violent acts. Another pre-trained neural network is used to detect people from the same video feed. The outputs of these neural networks are then combined using a probabilistic logic program to detect the complex, violent events.

Both (Roldán et al. 2020) and (Roig Vilamala et al. 2019) use pre-trained neural networks to parse the simple events. In this paper, instead, we assume that no such pre-trained

neural networks exist. As such, we assume that only end-to-end training is possible. This means that we only have training labels for when the complex events are happening, and not for the simple events. While this does make the training problem harder, it is undeniably easier to obtain labels for the complex events, thus reducing the costs associating to create the training set.

3.2 Purely statistical approaches

One possible approach is to view the whole CEP problem as a classification problem, and—for instance—use neural networks to detect when complex events occur. These approaches remove the manual definitions of complex events, and instead attempt to train the neural network to identify those definitions at the same time as it learns to classify the subsymbolic data. Due to the relevance of time in the definition of complex events, a Long Short Term Memory (LSTM) (Mishra et al. 2018) or a Convolutional 3D layer (C3D) (Liu et al. 2018) can be used. However, due to the necessity of learning the complex event rules, these approaches need very large amounts of data to train. Furthermore, the complexity of the rules that define the complex events is limited, due to the fact that the neural networks need to learn those rules.

3.3 Hybrid neuro-symbolic approaches

The current state of the art in CEP with subsymbolic data is Neuroplex (Xing et al. 2020). Neuroplex is a hybrid neuro-symbolic approach that makes use of human knowledge in order to reduce the amount of training data required when compared to purely statistical approaches. This is done by dividing the problem into two levels; low-level perception and high-level reasoning. The high-level reasoning is responsible for detecting the complex events based on manually defined rules, while the low-level perception is responsible for parsing the subsymbolic data into a set of classes that can be used when defining the rules.

In Neuroplex, the user defines the rules for the complex events. Then, a neural network is trained to emulate a logic layer that recognizes those rules. This allows users to inject human knowledge into the system. The neural network that emulates those rules is then used as the high-level reasoning. This is combined with another neural network, which performs the task of the low-level perception. Then, the high-level reasoning layer is frozen, meaning that the weights for this layer will not be modified by further training. Finally, the system is trained in an end-to-end manner. This trains the low-level perception neural network to recognize the simple events into the classes used to define the complex events.

Using a neural network to emulate the user defined rules is what allows Neuroplex to train in an end-to-end manner. However, it also introduces some limitations. Firstly, the reasoning neural network needs to be trained each time that the rules for the complex events are updated. As such, the whole system needs to be trained even if there only is a small change to the rules. Secondly, the ways in which complex events can be defined are, currently, substantially limited when compared to other CEP approaches. While improvements could be made to the system that trains the high-level

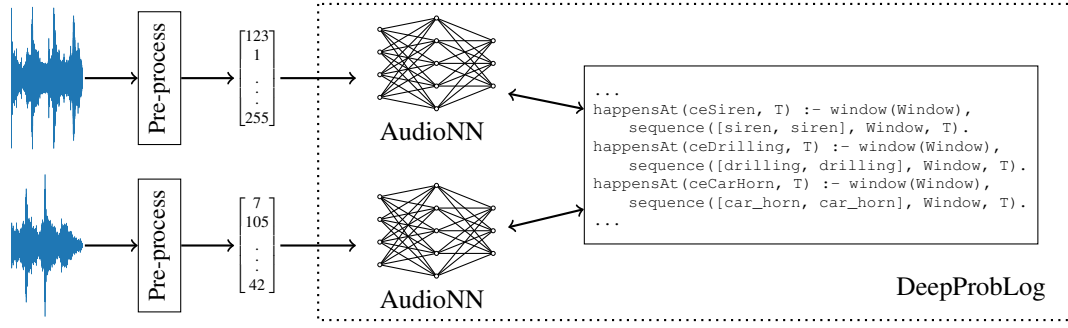


Figure 1: Overall architecture of our approach for the experiments performed in this paper.

reasoning to be more flexible, this would require a significant amount of work. At the moment, the high-level neural network can only be trained to recognize patterns of simple events within a given window. Finally, while Neuroplex can generate synthetic data to train the neural network to emulate the rules, it is not possible for the user to know that the neural network will behave exactly as the rules define in all situations. This is due to the nature of the neural network, which may give an unexpected answer if the given situation has not been seen in the training data. The only way to guarantee that the neural network will always behave as expected is to evaluate every possible situation, which becomes unfeasible as the complexity of the problem increases. As a result, there is a risk Neuroplex will not be robust against some adversarial conditions. In this paper, we propose an architecture that aims to solve these issues.

4 Our approach

In this section, we describe how we have used DeepProbLog to implement a hybrid neuro-symbolic approach to CEP. Our approach allows users to inject human knowledge into the system by manually defining rules for the complex events. At the same time, it allows us to perform end-to-end training in order to make use of subsymbolic data such as audio. This is achieved by dividing the tasks into two distinct levels; (i) a perception level, where a neural network is used to classify subsymbolic data in order to extract the symbolic information and (ii) a reasoning level, where probabilistic logic programming is used to define the complex event rules.

As explained above in Section 3.3, Neuroplex (Xing et al. 2020) also divides the problem into perception and reasoning levels. However, as explained above, Neuroplex uses a neural network to emulate the rules, instead of using an explicit logic layer. By using an explicit logic layer in our approach we remove the need of training a neural network to emulate the functionality of the logic layer, which makes it easier to update the complex event rules. Furthermore, we also remove the risk of the neural network behaving in an unexpected manner, thus providing a higher robustness.

In this paper, we use audio files as an input to the system. For processing purposes, the input audio is divided into one second segments, each of which is considered a simple event. Then, we try to detect patterns where a sound of the same class occurs twice within the given window. More

specifically, we look for cases in which the same class occurs at the last position in the window size and at some other position within the window. When this happens, a complex event is generated. The type of complex event will depend on the class of the repeated simple event.

Figure 1 shows the diagram used for our experimentation. Firstly, the input audio is divided into one second segments and pre-processed. For this, we use VGGish (Hershey et al. 2017), a state-of-the-art feature extractor for audio classification models². VGGish performs a feature extraction process which results in a matrix of size $128 \times N$, where N is the length of the input audio file in seconds. Each position in the matrix contains a value between 1 and 255. After performing this pre-processing, the resulting matrix is fed into our system. The vector resulting from each 1 second segment is fed into a multilayer perceptron (MLP) neural network, AudioNN in the diagram. This neural network classifies the segment into one of the 10 classes that appear in our dataset. The MLP used in our experimentation has 5 layers with 100, 80, 50, 25 and 10 neurons, in this order. A ReLU activation function is used between each of the layers, and a Softmax activation function is applied at the end.

Finally, the logic layer makes use of the output values from the neural network to predict whether or not a certain complex event is happening at a certain point in time. In order to determine this, the rules provided by the user are used. The diagram also shows a snippet of the logic rules used to define the complex events. This code defines that the complex events happen if a specific pattern of simple events happens within a given window of time. For this purpose, we have created the clause *sequence*, which allows us to easily define which patterns of simple events generate the complex events. For more details on the code, see <https://github.com/dais-ita/DeepProbCEP>.

For the experimentation in this paper, we set a maximum number of epochs of 100. However, in order to avoid overfitting we also make use of early stopping with a patience of 10 epochs. This means that if the performance of the system on the validation dataset does not improve for 10 epochs, we end the training early. We will then use the weights that

²In order to make it compatible with DeepProbLog, we use a PyTorch implementation of VGGish, available at <https://github.com/harritaylor/torchvggish>

performed the best in the validation dataset for testing.

5 Datasets generation

In this section, we describe how we have generated the datasets used to evaluate our approach. All the datasets are generated using Urban Sounds 8K (Salamon, Jacoby, and Bello 2014), a dataset containing over eight thousand short audio files (4 seconds or less) that contain sounds from 10 different classes: *air_conditioner*, *car_horn*, *children_playing*, *dog_bark*, *drilling*, *engine_idling*, *gun_shot*, *jackhammer*, *siren*, and *street_music*.

The first dataset used in our experiments are the *base datasets*. These datasets allow us to evaluate how the size of the sliding window affects the performance of the system. In our approach, we use this sliding window to define the maximum amount of time that can pass between the first and last simple events that will be aggregated into a complex event. As such, if a set of simple events follow the pattern we have defined but they are too far apart in a temporal sense, no complex event will be generated. This allows us to define that simple events that are separated by large amounts of time have no relation to each other.

We also want to evaluate how robust our approach is, as defined in the fourth objective from Section 1. For this purpose, we have generated datasets that perform an adversarial attack in the form of data poisoning. For this attack a percentage of the labels in the training dataset have been randomly changed to an incorrect value. Different percentage values are used to evaluate how this affects our approach. We call this dataset type *random noise dataset*.

In the following sections we will give more details on how both types of datasets have been generated.

For both types of datasets, we are using the same definitions for the complex events. Specifically, we are looking for patterns in which the sound that occurs in the last position of our sliding window also appears in another position within the window size. Each of the 10 sound classes in Urban Sounds 8K generates a different class of complex event.

5.1 Base dataset

In this section, we describe how we generated the base dataset. The process used to generate the base dataset allows us to change the window size by changing the value of *Window*. *Window* is a positive integer that indicates the number of timestamps between the first and last simple events that form a complex event. For this paper, we have generated datasets with window sizes of 2, 3, 4 and 5.

In order to generate the base dataset, we use the different folds from Urban Sounds 8K. Out of the 10 folds provided by the original dataset, 8 are used to generate our training dataset, 1 is used to generate our validation dataset and the last fold is used to generate our testing dataset. The steps to generate the base datasets are shown in Figure 2, which illustrates the following steps:

1. We take all the audio files from the original dataset and randomly shuffle them into a sequence S of simple events, where each audio file represents one simple event. Therefore, the length of S is the number of audio files in

the original dataset. Simple events can be accessed by their index like so $S[I]$. For each of them we can access the file itself and the class it contains using $S[I].audio$ and $S[I].class$, respectively. In order to have a consistent length for all simple events, only the first second of each audio file is used.

2. We create a list C that will indicate for each timestamp whether a complex event happens. We initialize this list with *null*, which hereinafter represents that no complex event happens at the specified timestamp.
3. For each timestamp T where $0 < T < len(S)$:
 - (a) If the pattern for one of the complex events occurs, mark $C[T]$ as the corresponding complex event. More specifically, if there exists P such that $T - Window < P \leq T$ and $S[P].class = S[T].class$, mark $C[T]$ as ceN , where N is the value of $S[T].class$. This means that if a sound occurs at the last position in the window (T) and somewhere else within the window (P), a complex event is generated.
 - (b) Otherwise, leave $C[T]$ marked as the null class.
4. Finally, if this is the training dataset, generate the training sequence of simple events TS , which will only contain the audio files, but not the ground truth of which class they represent, as these should not be available when performing end-to-end training. Therefore, $TS[I] = S[I].audio$ for $0 < I < len(S)$.

Before using these datasets for training, they are also balanced in order to avoid overfitting for a specific class. This results in a training dataset with 1000 training points for each window size.

5.2 Random noise datasets

Given the complexity of the definition of some of the real world complex events, it can sometimes be hard to correctly label when a certain complex event is happening. This can lead to errors on the training dataset, which might affect the accuracy of the system after training. In Section 1, we defined that one of our objectives was to be robust against adversarial conditions. As such, we want to evaluate how much of an effect training using a noisy dataset has on the performance of our approach. Of course, this is not an issue with our synthetically generated dataset. However, using a synthetically generated dataset offers us the opportunity of artificially introducing noise in a controlled manner. This allows us to evaluate how well our approach might perform when used on a real dataset, which might contain an unknown level of noise.

We use a poisoning attack where we randomly change a percentage of the training labels for another random label to simulate this noise. For instance, assume that, based on the complex event definitions, a certain timestamp is marked as *ceSirens*. However, if this case is affected by the noise it will be labelled as another complex event, which is selected randomly every time. We will call the datasets generated using this process *random noise datasets*.

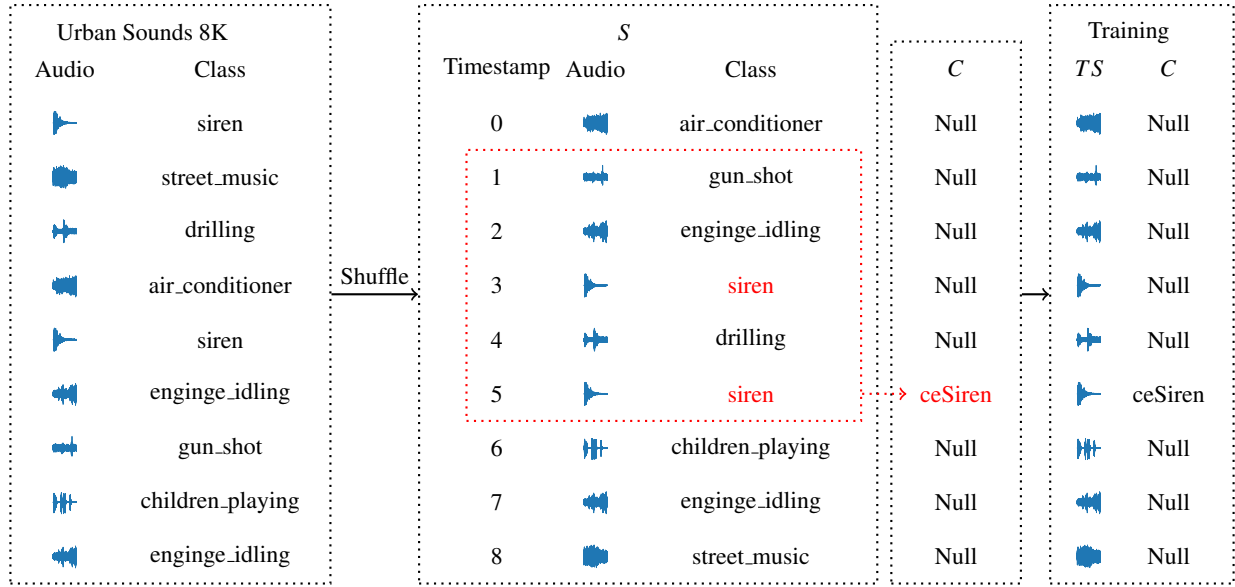


Figure 2: Diagram representing how the datasets used in this paper are generated. A window of 5 has been used. After randomly shuffling the audio-class pairs from the Urban Sounds 8K dataset, we detect that at timestamp 5 we have two instances of the class *siren* within the given window. Therefore, we mark timestamp 5 in *C* as *ceSiren*. We can also observe that there is a pair of *engine_idling* on timestamps 2 and 7. However, the distance between them is bigger than the given window, and therefore that does not result in a complex event. Finally, for the training dataset we remove the ground truth for the sound class, as we are doing end-to-end training.

In order to generate the random noise datasets, we use the same steps described to generate the base dataset, explained above in Section 5.1. However, when labelling a certain timestamp as a complex event (Step 3a), there is probability of changing that label to a random complex event. This probability is determined by the noise percentage parameter. We have generated datasets ranging from 0.0 to 0.6 noise, with a step of 0.2. For this, 0.0 means that no data has been poisoned, while 1.0 would mean that all the data has been poisoned. Finally, the datasets are balanced. These datasets also have a size of 1000 training points.

It is important to note that this attack is only performed on the training dataset. This means that the testing dataset will maintain the ground truth, which will allow us to see how the system would perform in a real life scenario.

This kind of noise might appear both due to a malicious agent, and to the difficulty of labelling the sophisticated scenarios where a CEP system would be useful: for instance, different annotators might be having different consideration on what constitutes the a complex event.

6 Experimental analysis

In this section, we explore the accuracy results for our approach after training with the synthetic datasets explained in Section 5.

All the values displayed on the graphs and tables in the following sections are the result of averaging the accuracies of 3 different executions.

Window size	Accuracy	STD
2	0.8657	0.0041
3	0.7645	0.0109
4	0.7069	0.0191
5	0.6401	0.0225

Table 1: Average accuracy results and standard deviation for complex events classification by window size.

6.1 Performance with base dataset

In Table 1 we can see the results of training our approach on a balanced dataset with 1000 training data points. As shown in the table, the performance of the approach is fairly good with a window size of 2. However, the performance does decrease as the window size increases. This could have been expected, as the problem gets more complex as the window size increases. This is because a bigger window size contains more simple events, which makes it more likely that the system will incorrectly classify one of them. For example, this can cause the system to predict that a complex event is happening when it is not, thus reducing the performance of the system.

6.2 Robustness against poisoning attack with random noise

As explained above, we also want to know how robust our approach is against adversarial conditions. For this purpose, we have trained the system with the random noise datasets




Figure 3: Evaluation of system’s performance with different levels of noise in the training data. The horizontal axis indicates the percentage of training data points where the class has been randomly selected. The vertical axis indicates the accuracy of the system for either individual classifications or for complex event detection. The error bars represent the standard deviation for each case. The error bars have been slightly offset in their horizontal axis to avoid overlapping, which would make them hard to read.

(explained above in Section 5.2) and evaluated how the performance is affected. The results are shown in Figure 3. As seen in the graph, while there is a slight decrease in performance when training with the 20% noise dataset, it does not seem to significantly impact the system. By contrast, with a percentage of noise of 40% or higher, the performance seems much less consistent. While, in some cases, the system is still able to train correctly, in others it performs very significantly worse. This is what causes the high standard deviation that can be seen in the graph. However, it does not seem likely that the user would not realize that almost half of the training points in the dataset are incorrectly labelled. As such, we would argue that our approach is robust against reasonable amounts of noise.

7 Conclusion and future work

In this paper we have presented a hybrid neuro-symbolic approach capable of performing CEP on subsymbolic data. More specifically, we have demonstrated that our approach is capable of detecting complex events from an audio stream after training end-to-end. We have also shown that our approach is quite robust against a random noise poisoning attack, thus fulfilling all four objectives defined in Section 1.

As part of future work, we are considering on evaluating the performance on other types of subsymbolic data. This should be possible by changing the architecture of the neural network based on the type of subsymbolic data we want to use.

An other area on which future research could be applied is on the time efficiency of the approach. Due to the use of a logic layer, our approach is slower in both training time and

inference time when compared to approaches that are implemented using a neural network, such as Neuroplex (Xing et al. 2020). This is mostly due to the cost of generating the arithmetic circuit used to calculate the output for the logic layer. DeepProbLog offers a cache functionality to reduce the amount of times this arithmetic circuit has to be generated. However, some further research will be needed to make the most out of this functionality for problems that deal with a temporal aspect.

8 Acknowledgements

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

- Anicic, D.; Rudolph, S.; Fodor, P.; and Stojanovic, N. 2012a. Real-time complex event recognition and reasoning—a logic programming approach. *Applied Artificial Intelligence - AAI* 26:6–57.
- Anicic, D.; Rudolph, S.; Fodor, P.; and Stojanovic, N. 2012b. Stream reasoning and complex event processing in ETALIS. *Semantic Web* 3:397–407.
- Board, D. I. 2019. Ai principles: Recommendations on the ethical use of artificial intelligence by the department of defense. *Supporting document, Defense Innovation Board*.
- De Raedt, L., and Kimmig, A. 2015. Probabilistic (logic) programming concepts. *Machine Learning* 100(1):5–47.
- De Raedt, L.; Kimmig, A.; and Toivonen, H. 2007. ProbLog: A probabilistic prolog and its application in link discovery. *IJCAI International Joint Conference on Artificial Intelligence* 2468–2473.
- Hershey, S.; Chaudhuri, S.; Ellis, D. P.; Gemmeke, J. F.; Jansen, A.; Moore, R. C.; Plakal, M.; Platt, D.; Saurous, R. A.; Seybold, B.; et al. 2017. Cnn architectures for large-scale audio classification. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 131–135. IEEE.
- Liu, K.; Liu, W.; Gan, C.; Tan, M.; and Ma, H. 2018. T-c3d: Temporal convolutional 3d network for real-time action recognition.
- Luckham, D. C. 2002. *The Power of Events: An Introduction to Complex Event Processing in Distributed Enterprise Systems*. USA: Addison-Wesley Longman Publishing Co., Inc.
- Manhaeve, R.; Dumancic, S.; Kimmig, A.; Demeester, T.; and De Raedt, L. 2018. Deepproblog: Neural probabilistic logic programming. In *NIPS2018*. 3749–3759.

- Manhaeve, R.; Dumančić, S.; Kimmig, A.; Demeester, T.; and De Raedt, L. 2021. Neural probabilistic logic programming in deepproblog. *Artificial Intelligence* 298:103504.
- Mishra, S.; Jain, M.; Siva Naga Sasank, B.; and Hota, C. 2018. An ingestion based analytics framework for complex event processing engine in internet of things. In Mondal, A.; Gupta, H.; Srivastava, J.; Reddy, P. K.; and Somayajulu, D., eds., *Big Data Analytics*, 266–281. Cham: Springer International Publishing.
- Pearl, J. 2009. *Causality: Models, Reasoning, and Inference*. Cambridge University Press.
- Roig Vilamala, M.; Hiley, L.; Hicks, Y.; Preece, A.; and Cerutti, F. 2019. A pilot study on detecting violence in videos fusing proxy models. In *2019 22th International Conference on Information Fusion (FUSION)*, 1–8.
- Roldán, J.; Boubeta-Puig, J.; Luis Martínez, J.; and Ortiz, G. 2020. Integrating complex event processing and machine learning: An intelligent architecture for detecting iot security attacks. *Expert Systems with Applications* 149:113251.
- Salamon, J.; Jacoby, C.; and Bello, J. P. 2014. A dataset and taxonomy for urban sound research. In *22nd ACM International Conference on Multimedia (ACM-MM'14)*, 1041–1044.
- Teymourian, K.; Rohde, M.; and Paschke, A. 2012. Knowledge-based processing of complex stock market events. In *Proceedings of the 15th International Conference on Extending Database Technology, EDBT '12*, 594–597. New York, NY, USA: Association for Computing Machinery.
- Xing, T.; Garcia, L.; Vilamala, M. R.; Cerutti, F.; Kaplan, L.; Preece, A.; and Srivastava, M. 2020. *Neuroplex: Learning to Detect Complex Events in Sensor Networks through Knowledge Injection*. New York, NY, USA: Association for Computing Machinery. 489–502.