Towards Maintaining and Reusing Complex Event Processing Systems

Marc Roig Vilamala *, Mark Law †, Harrison Taylor *, Tianwei Xing ‡, Luis Garcia ‡,
Dave Braines ¶, Dan Cunnington ¶, Graham White ¶,
Alessandra Russo †, Mani Srivastava ‡, Lance Kaplan §, Alun Preece *, Federico Cerutti * ‖
* Cardiff University (UK), † Imperial College London (UK), ‡ UCLA (USA), § ARL (USA),
¶ IBM-UK (UK), ‖ University of Brescia (Italy)

***Abstract*—Understanding situations formed of patterns of interrelated events is a complex problem: often available training data are sparse and either noisy or potentially manipulated by other members of the coalition, if not by an adversary. In previous research we introduced DeepProbCEP, a hybrid neuro-symbolic architecture that leverages both a neural architecture, to interpret raw data, and logical rules, to express patterns defining complex events, while allowing for end-to-end learning. While DeepProbCEP has many advantages, including its ability of adapting to new contexts by leveraging only sparse data, it suffers from several drawbacks, notably the problem of maintaining logical rules. Inductive Logic Programming (ILP) systems are able to learn logical rules from examples. In this demonstration we use an extension of the recent FastLAS ILP system to learn the definitions of complex events from a small number of examples. This method for automatic derivation of the logical rules using ILP overcomes the previous problem of requiring a manual encoding of the complex event definitions.**

I. Description

This is a collaborative demonstration of tasks 10.2 and 10.3.

Imagine a scenario where we are trying to detect a *shooting* using microphones deployed in a city. In this scenario, *shooting* is a situation of interest that we want to identify from a high-throughput (audio) data stream. Complex Event Processing (CEP) is a type of approach aimed at detecting such situations of interest, called *complex events*, from a data stream using a set of rules. These rules are defined on atomic pieces of information from the data stream, which we call events—or *simple events*, for clarity. Multiple simple events linked with spatio-temporal relations form complex events. For instance, we can define the complex event *shooting* to start when multiple instances of the simple event *gunshot* occur. This would indicate that there is an incidence that needs to be dealt with, which would prompt us to send the relevant authorities. For simplicity, we can assume that when we start to detect *siren* events those authorities have arrived and the situation is being dealt with, which would conclude the complex event. Of course, this is only one of the many situations that we might be interested in detecting on the real world, each of which might require a different response.

In [6] we introduce DeepProbCEP, a hybrid neuro-symbolic architecture that leverages both a neural architecture, to interpret raw data, and logical rules, to express patterns defining complex events, while allowing for end-to-end learning. Compared to simple neural architectures, DeepProbCEP (i) needs fewer labelled data thanks to its end-to-end learning capability, (ii) is robust against noise and adversarial attacks in the form of training data poisoning and (iii) can classify individual events as a by-product of the end-to-end training. Figure 1 illustrates DeepProbCEP [6]: audio signal, processed one second at the time, is classified by a neural network, the results of which are fed into a ProbLog code that allows for event calculus manipulation.

However, DeepProbCEP also suffers from several drawbacks, one of which concerns the elicitation of knowledge to be distilled into logical rules. In [4], [6], [7] we assumed that human experts would be the source of such knowledge and they would thus be responsible to distil it into logical rules. While we strongly believe humans should always have the opportunity to inspect and modify such rules, it is clear that manually writing and maintaining complex and articulated rules that will need adaptation to continuously changing environments is not the way forward. To address the issue, in this demonstration we will show the benefits of using state-of-the-art Inductive Logic Programming (ILP) [3] techniques.

The goal of an ILP system is to learn a set of logical rules that explain some examples in the context of some existing background knowledge (also expressed as logical rules). Compared with other forms of machine learning, ILP has many advantages, including its ability to learn general concepts from relatively few examples. Many ILP systems have severe restrictions limiting the forms of logical rules that can be learned and the types of learning problems that can be solved. For example, many ILP systems require examples of exactly the same concept which is being defined by the learned rules. In this demonstration we use an extension of FastLAS [2], which is capable of *non-observational predicate learning*, meaning that it is possible to encode general *Event Calculus* [1] axioms in the background knowledge and give examples that are only indirectly linked to the learned rules through these axioms. This, together with FastLAS's ability to learn from noisy data, makes FastLAS well suited to learning the complex event definitions needed by DeepProbCEP.

While this does not address all the shortcomings of DeepProbCEP, it is a significant step forward in the direction of providing actionable capabilities that can be deployed and maintained with ease. We discuss such capabilities as part of the DAIS-ITA scenario in the next section.

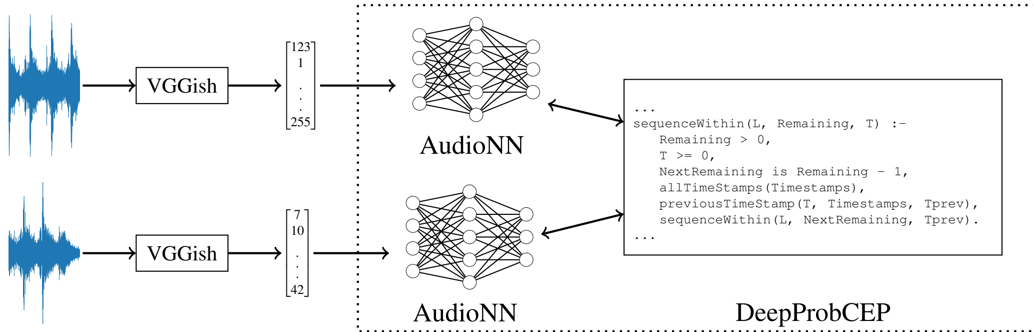

Fig. 1: Overall architecture of DeepProbCEP for the audio setting. Audio files are processed one second at the time using VGGish, a feature extraction system, that returns a vector of 128 values between 1 and 255. This vector is then fed into a feed-forward network that is then connected to a logical layer that uses rules for identifying complex events.

II. Scenario

This demonstration will build on top of the overarching scenario based on NATO Anglova location [5].[1]

As allies of Anglova, our troops are present in particular in the of town of Kristiansand, that recently has become "disquiet." In particular, we have been tasked to monitor and gather intelligence as to the nature and capability of those causing this instability. As part of our activities in the first phase, it is assessed that certain gangster groups seem to be coalescing into pro and anti sides and asserting their dominance, while at the same time extending their network and political influence with the help of a foreign power.

Eventually it is decided that the situation must be contained before the underground force declare a coup d'état. We deploy a battle group to isolate the town and to attempt to arrest/remove the opposition, while causing as little collateral damage as possible. While the manoeuvring is happening and networks are being established we may also add local drone reconnaissance and some covert ground sensors to better understand the patterns of life, and try to identify positions, strengths, persons of interest and their locations. In particular, as part of this demonstration, we assume that we will be having access to audio sensors data, as well as well-labelled datasets for similar situations, albeit not necessarily gathered with the same technology, in the same context, or even in recent time. This hypothesis allows us to demonstrate the need for reusing and maintaining capabilities across multiple campaigns while adapting to new situations minimising efforts.

During the operation we encounter more resistance than initially thought, and the adversary counter by jamming some of our comms capability, destroying some, and isolating parts of the network. In this dynamic, contested environment, we show how our capability for complex event detection provides relevant actionable intelligence that is then fed directly to manned and unmanned headquarters.

[1] https://dais-ita.org/vbc20_wed_1715_exp

III. Demo Requirements

As we received information that AFM2020 will be entirely online, we assume that no requirements are needed to be highlighted at this stage.

Acknowledgement

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

[1] Robert Kowalski and Marek Sergot. A logic-based calculus of events. In *Foundations of knowledge base management*, pages 23–55. Springer, 1989.

[2] Mark Law, Alessandra Russo, Elisa Bertino, Krysia Broda, and Jorge Lobo. Fastlas: Scalable inductive logic programming incorporating domain-specific optimisation criteria. In *AAAI*, pages 2877–2885, 2020.

[3] Stephen Muggleton. Inductive logic programming. *New generation computing*, 8(4):295–318, 1991.

[4] M. Roig Vilamala, L. Hiley, Y. Hicks, A. Preece, and F. Cerutti. A pilot study on detecting violence in videos fusing proxy models. In *2019 22th International Conference on Information Fusion (FUSION)*, pages 1–8, 2019.

[5] N. Suri, K. M. Marcus, C. van den Broek, H. Bastiaansen, P. Lubkowski, and M. Hauge. Extending the anglova scenario for urban operations. In *2019 International Conference on Military Communications and Information Systems (ICMCIS)*, pages 1–7, 2019.

[6] Marc Roig Vilamala, Harrison Taylor, Tianwei Xing, Luis Garcia, Mani Srivastava, Lance Kaplan, Alun Preece, Angelika Kimming, and Federico Cerutti. Deepprobcep: A hybrid neuro-symbolic approach for complex event processing in noisy and adversarial settings. https://dais-ita.org/node/5340, 2020.

[7] Tianwei Xing, Marc Roig Vilamala, Luis Garcia, Federico Cerutti, Lance Kaplan, Alun Preece, and Mani Srivastava. Deepcep: Deep complex event processing using distributed multimodal information. In *2019 IEEE International Conference on Smart Computing (SMARTCOMP)*, pages 87–92. IEEE, 2019.