Explainability and Performance of Anticipatory Learning Classifier Systems in Non-Deterministic Environments

Romain Orhand
Icube Laboratory - UMR CNRS 7357
Illkirch, France
University of Strasbourg
Strasbourg, France
rorhand@unistra.fr

Anne Jeannin-Girardon
Icube Laboratory - UMR CNRS 7357
Illkirch, France
University of Strasbourg
Strasbourg, France
anne.jeannin@unistra.fr

Pierre Parrend
Icube Laboratory - UMR CNRS 7357
Illkirch, France
ECAM Strasbourg-Europe
Schiltigheim, France
pierre.parrend@ecam-strasbourg.eu

Pierre Collet
Icube Laboratory - UMR CNRS 7357
Illkirch, France
University of Strasbourg
Strasbourg, France
pierre.collet@unistra.fr

ABSTRACT

In the field of Reinforcement Learning, models based on neural networks are highly performing, but explaining their decisions is very challenging. Instead of seeking to open these "black boxes" to meet the increasing demand for explainability, another approach is to used rule-based machine learning models that are explainable by design, such as the Anticipatory Learning Classifier Systems (ALCS). ALCS are able to develop simultaneously a complete representation of their environment and a decision policy based on this representation to solve their learning tasks. This paper focuses on the ability of ALCS to deal with non-deterministic environments used in reinforcement learning problems, while discussing their explainability. Directions for future research are thus highlighted to improve both the performance and the explainability of the ALCS to meet the needs of critical real-world applications.

CCS CONCEPTS

• **Computing methodologies → Rule learning**; *Knowledge representation and reasoning*;

KEYWORDS

Anticipatory Learning Classifier System, Machine Learning, Explainability, Non-determinism, Building Knowledge

ACM Reference Format:
Romain Orhand, Anne Jeannin-Girardon, Pierre Parrend, and Pierre Collet. 2021. Explainability and Performance of Anticipatory Learning Classifier Systems in Non-Deterministic Environments. In *2021 Genetic and Evolutionary Computation Conference Companion (GECCO '21 Companion), July 10–14, 2021, Lille, France.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3449726.3459510

GECCO '21 Companion, July 10–14, 2021, Lille, France

ACM ISBN 978-1-4503-8351-6/21/07…$15.00
https://doi.org/10.1145/3449726.3459510

1 INTRODUCTION

In the field of Reinforcement Learning, the use of deep neural networks (DNN) makes the explanation of the decisions made by these models more complex : indeed, DNN are considered as "black boxes" [1]. Explaining the decisions made by these "black boxes" is increasingly important in many fields such as autonomous driving, medicine, justice, insurance, loan approval, *etc.* : performance alone cannot account for the use of these models for critical applications.

To meet this need for explainability, Anticipatory Learning Classifier Systems (ALCS) can be used to carry out the learning and solving of a task in an explainable way. ALCS are rule-based machine learning algorithms in which rules (called classifiers) are built by articulating "conditions", "actions" and "effects" as described in [5]: ALCS learn to *anticipate* the effects of an action depending on the condition and according to the environment in which they evolve, giving explanatory insights about the built decision policies and environmental representations.

As non-determinism is a feature of many real-world reinforcement learning problems, the work presented in this paper focuses on the ability of ALCS to deal with these non-deterministic properties, while their inherent ability to provide some level of explainability is discussed according to the framework of eXplainable Artificial Intelligence (XAI) [2].

2 ALCS AND NON-DETERMINISM

The Behavioral Sequences [3] and the Probability-Enhanced Predictions (PEP) [4] have been proposed to let the ALCS deal with non-deterministic environments. The Behavioral Sequences solely tackle the Perceptual Aliasing Issue (the perceptual sensors are insufficient to determine the exact state of the environment), while the PEP enable the system to deal with different non-deterministic properties, like noisy perceptual sensors or actions whose results are uncertain.

Behavioral Sequences consists in building sequences of actions to skip the states related to the Perceptual Aliasing Issue (referred to as PAI states). The states encountered by the system between consecutive actions in a sequence are not described by the classifiers: only the states preceding and following the PAI state are represented by the classifiers. As a consequence, there are no reliable classifiers in the PAI state that anticipate the states reachable from the PAI one, meaning the system cannot build a complete and accurate representation of its environment. Moreover, the ALCS do not promote the shortest sequences of actions suited to their environment, meaning that sequences longer than necessary can be chosen, thus favoring sub-optimal decision policies to solve their task. Furthermore, [3] discriminated between classifiers whose Behavioral Sequences make the ALCS loop between identical states, in an attempt to reduce the growth of the population of classifiers, even if this could also favor sub-optimal decision policies: the system may need to follow identical successive states in some environments to develop the optimal decision policy, like in a corridor where the best way to leave it is to always go on.

PEP were introduced in ALCS to enable them to learn a complete and accurate representation of their environment, by permitting the prediction of an ensemble of anticipation depicted by pairs of perceptive attributes and the related probabilities in the classifier anticipation. PEP does not give the ALCS the ability to solve learning tasks related to the Perceptual Aliasing Issue for instance and, the probabilities computed in the PEP, along with the sets of anticipated states they describe, can be inconsistent with the environments used. The updating of probabilities is also sensitive to the perceptions received by the ALCS: a small probability of a PEP attribute may be increased too much, and *vice versa*, thus challenging the convergence to the expected values. A classifier whose effect contains several PEP can depict a set of states that is not representative of the environment, because each pairs of one PEP can be associated with all pairs of the other PEP, thus describing more states than it has truly anticipated.

3 ALCS AND XAI

A key concept in XAI concerns what "explanations" are: this concept should integrate findings of Social and Cognitive Sciences. In this paper, an explanation is defined as an act of transmitting the causes that led to a particular event to someone. In an attempt of stressing such findings within works on XAI, [2] argues that explanations are "contrastive", "social", "selected" and the causes that led to an event are more powerful to provide explanations than probabilistic relationships. In other words, some of the causes that led to a particular event, and not to another one, are more likely to provide an appropriate explanation to the explainee, as long as an interaction between the explainee and the explainer could be set up.

ALCS build their classifiers by differentiating consecutive perceptions of their environment, in order to anticipate the possible changes due to a performed action in a particular situation: the causes of perceptual changes, as well as the changes, are described within the classifiers. Unlike other learning classifier systems, ALCS do not rely on a stochastic process to build their classifiers, therefore correlating the causes with the effects. Iteratively, a population of classifiers is built by fitting its population to the environment, permitting to give the user insights about *why* and *when* a classifier is triggered or not. Classifiers can be then compared with each other, as long as the representations used within the classifiers are meaningful for the users (otherwise, interactions between the ALCS and their users would not be possible).

Nonetheless, the explainability of ALCS can be fostered. For instance, the insights provided by these systems relate to all perceived changes because of the way the ALCS learn, whereas the most appropriate explanations for the user may lie in the unchanged perceptive items or even, in the perceptive items that may be not described by the classifiers.

4 CONCLUSION

Even if the Probability-Enhanced Predictions and the Behavioral Sequences have been successfully used to let the ALCS handle non-deterministic environments, they could be enhanced and used together to improve both the performance and the explainability of the ALCS: among other things, shortest sequences could be promoted through reinforcement or, representations built within the PEP could be more closely-related to the environmental setting of the system.

These two mechanisms have never been used simultaneously in an ALCS, despite their complementary nature. However, coupling both of them should be done carefully, to avoid the building of inadequate classifiers in non-deterministic states. Such inadequate classifiers could prevent the system from achieving its learning task or prevent the user from understanding the environmental representation built by the ALCS. Thus, a coupling of the PEP and the Behavioral Sequences requires the ALCS to be able to detect when a state is a PAI state, and mechanisms should be implemented in order to regulate its population of classifiers.

Additional cognitive mechanism such as selective attention could be implemented to enhance the explanatory drawbacks of the ALCS. For instance, selective attention would permit the system to weight the perceptive items of its classifiers, thus providing more explanatory elements about which items would cause both the triggering of a classifier and a particular change in the environment. Moreover, improving the interaction between ALCS and its users should provide further explanatory items by, for example, allowing users to add their own classifiers to bias the learning or the behavior of the system, or assessing the quality of the provided explanatory items.

Finally, comparing the capabilities of ALCS with other reinforcement learning models through a mutli-objective approach is also a direction for future research, as performance is balanced by explainability. This would allow to better define the possible fields of application of ALCS, that is still an open question.

REFERENCES

[1] Amina Adadi and Mohammed Berrada. 2018. Peeking inside the black-box: A survey on Explainable Artificial Intelligence (XAI). *IEEE Access* 6 (2018), 52138–52160.

[2] Tim Miller. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence* 267 (2019), 1–38.

[3] Romain Orhand, Anne Jeannin-Girardon, Pierre Parrend, and Pierre Collet. 2020. BACS: A Thorough Study of Using Behavioral Sequences in ACS2 (*to appear*). In *International Conference on Parallel Problem Solving from Nature*.

[4] Romain Orhand, Anne Jeannin-Girardon, Pierre Parrend, and Pierre Collet. 2020. PEPACS: Integrating Probability-Enhanced Predictions to ACS2. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference Companion*. 1774–1781.

[5] Wolfgang Stolzmann, Martin Butz, Joachim Hoffmann, and David Goldberg. 2000. First Cognitive Capabilities in the Anticipatory Classifier System. (02 2000).