Computational Models of Argument
F. Toni et al. (Eds.)

doi:10.3233/FAIA220169

EGNN: A Deep Reinforcement Learning Architecture for Enforcement Heuristics

Dennis CRAANDIJK [a,b], Floris BEX [b,c]
[a] *National Police Lab AI, Netherlands Police*
[b] *Department of Information and Computing Sciences, Utrecht University*
[c] *Institute for Law, Technology and Society, Tilburg University*

1. Introduction

An increasing amount of research is being directed towards neuro-symbolic computing, combining learning in neural networks with reasoning and explainability via symbolic representations [4]. One subfield of AI where neuro-symbolic methods are a promising alternative for existing symbolic methods is computational argumentation. Much of the theory of computational argumentation is based on the seminal work by Dung [6], in which he introduces abstract argumentation frameworks (AFs) of arguments and attacks, and several acceptability semantics that define which sets of arguments (*extensions*) can be reasonably accepted. Core computational problems in abstract argumentation are typically solved with handcrafted symbolic methods [1]. However, recently we demonstrated the potential of a deep learning approach by showing that a graph neural network is able to learn to determine almost perfectly which arguments are (part of) an extension [2].

When considering dynamic argumentation - a growing research area where the knowledge about attacks between arguments can be incomplete or evolving - other types of computational problems arise where neuro-sybmolic methods are still unexplored. In [3] we propose our enforcement graph neural network (EGNN), a learning-based approach to the dynamic argumentation problem of enforcement: given sets of arguments that we (do not) want to accept, how to modify the argumentation framework in such a way that these arguments are (not) accepted, while minimizing the number of changes [5]. Here we demonstrate our implementation of an EGNN.

2. Demonstration

When confronted with some problems with a high computational complexity, existing symbolic enforcement solvers exhibit quite a significant drop in runtime performance, limiting their practical applicability. While there is a need for efficient heuristics to address this problem, designing such heuristics takes considerable expert effort and domain knowledge. EGNN is a single architecture that can be trained through deep reinforcement learning to learn enforcement heuristics for all common semantics and enforcement problems, without supervision of an existing solver. EGNN learns a *message passing*

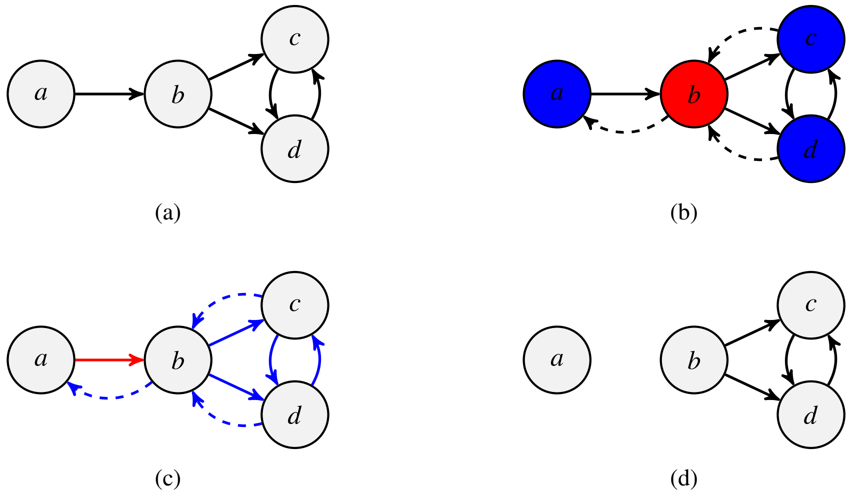

Figure 1. Consider enforcing argument b in the AF $F = (\{a,b,c,d\},\{(a,b),(b,c)(b,d),(c,d),(d,c)\})$. EGNN takes the AF (a), maps it it to a fully connected graph where nodes have a vectorial representation denoting which arguments should be enforced (b). Node vectors are updated through *message passing* and are mapped to an output per edge (c) indicating which edge should be *flipped* (d).

algorithm that predicts which attack relations between arguments should be *flipped* (i.e. added or deleted) in order to enforce the acceptability of (a set of) arguments. Experimental results demonstrate that EGNN can learn near-optimal heuristics for all *extension* and *status* enforcement problems under the most common semantics, and outperforms symbolic solvers with respect to efficiency on enforcement problems that are higher in the complexity hierarchy.

We demonstrate our Python implementation of an EGNN and show: the input, message passing and output steps of the model; the learned heuristics for enforcement problems; how the learned heuristic differs from symbolic algorithms. We do so by graphically demonstrating EGNN's behaviour on an AF (cf. Figure 1).

References

[1] Günther Charwat, Wolfgang Dvorák, Sarah Alice Gaggl, Johannes Peter Wallner, and Stefan Woltran. Methods for solving reasoning problems in abstract argumentation - A survey. *Artificial Intelligence*, 220:28–63, 2015.

[2] Dennis Craandijk and Floris Bex. Deep learning for abstract argumentation semantics. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 1667–1673, 2020.

[3] Dennis Craandijk and Floris Bex. Enforcement heuristics for argumentation with deep reinforcement learning. In *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022*, pages 5573–5581, 2022.

[4] Artur S. d'Avila Garcez, Marco Gori, Luís C. Lamb, Luciano Serafini, Michael Spranger, and Son N. Tran. Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning. *FLAP*, 6(4):611–632, 2019.

[5] Sylvie Doutre and Jean-Guy Mailly. Constraints and changes: A survey of abstract argumentation dynamics. *Argument & Computation*, 9(3):223–248, 2018.

[6] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–358, 1995.