A Roadmap for Neuro-argumentative Learning

Maurizio Proietti[1,*], Francesca Toni[2,*]

[1]IASI-CNR, Rome, Italy
[2]Department of Computing, Imperial College London, UK

Abstract

Computational argumentation (CA) has emerged, in recent decades, as a powerful formalism for knowledge representation and reasoning in the presence of conflicting information, notably when reasoning non-monotonically with rules and exceptions. Much existing work in CA has focused, to date, on reasoning with given argumentation frameworks (AFs) or, more recently, on using AFs, possibly automatically drawn from other systems, for supporting forms of XAI. In this short paper we focus instead on the problem of learning AFs from data, with a focus on neuro-symbolic approaches. Specifically, we overview existing forms of *neuro-argumentative (machine) learning*, resulting from a combination of neural machine learning mechanisms and argumentative (symbolic) reasoning. We include in our overview neuro-symbolic paradigms that integrate reasoners with a natural understanding in argumentative terms, notably those capturing forms of non-monotonic reasoning in logic programming. We also outline avenues and challenges for future work in this spectrum.

Keywords

Computational argumentation, Artificial neural networks, Non-monotonic reasoning

1. Introduction

Computational argumentation (CA) has emerged, since the nineties, as a powerful formalism for knowledge representation and reasoning in the presence of conflicting information (see [1, 2] for recent overviews). It has been shown to capture and generalise, in particular, several forms of non-monotonic reasoning [3, 4], notably required when operating with rules and exceptions naturally giving rise to conflicts (e.g. both the rule “birds fly” and the exception ”penguins do not fly” apply to the same individual *tweety* – a penguin and thus a bird). Also, it is being widely deployed to support forms of explainable AI (XAI) (e.g. see overview in [5]), given the appeal of argumentation in explanations amongst humans, e.g. as in [6], within the broad view that XAI should take findings from the social sciences into account [7].

To date, the bulk of work in CA amounts to defining so-called *argumentation frameworks (AFs)*, which are symbolic representations equipped with semantics/tools for reasoning towards the resolution of conflicts and drawing (argumentatively acceptable) conclusions. In addition, increasingly more attention has been given over the years to combining CA and machine

NeSy'23: 17TH INTERNATIONAL WORKSHOP ON NEURAL-SYMBOLIC LEARNING AND REASONING, Certosa di Pontignano, Siena, Italy, 3–5 July 2023
*Corresponding author.
maurizio.proietti@iasi.cnr.it (M. Proietti); ft@ic.ac.uk (F. Toni)
http://www.iasi.cnr.it/~proietti/ (M. Proietti); https://www.doc.ic.ac.uk/~ft/ (F. Toni)
0000-0003-3835-4931 (M. Proietti); 0000-0001-8194-1459 (F. Toni)

learning (e.g. as overviewed in [8]). Here, we focus specifically on methods combining CA and machine learning with artificial neural networks (NNs)[1], focusing on what we term *neuro-argumentative (machine) learning (NAL)*, resulting from a combination of neural (machine learning) mechanisms and (symbolic) argumentative reasoning using methods in CA. Specifically, in Section 2 we overview existing forms of NAL, including neuro-symbolic paradigms integrating forms of logic programming with a natural understanding in CA terms (see below). Finally, in Section 3 we outline avenues and challenges for future work in the broad spectrum of NAL.

Background on CA. The simplest form of AFs is given by *abstract AFs* (AAFs) [3], which boil down to directed graphs whose nodes are *arguments* (e.g. "*tweety* flies as it is a bird" and "*tweety* does not fly as it is a penguin") and whose edges are *attacks* between them (e.g., the latter argument for *tweety* attacks the former). They are equipped, e.g., with the semantics of *stable extensions*, which are conflict-free sets of arguments attacking every argument they do not contain (leading to accepting the second argument for *tweety* and inferring that it does not fly). Several other forms of CA have been proposed and play a role in existing works in NAL, notably *bipolar AFs (BAFs)* [9], i.e. directed graphs with two types of edges: attacks as in AAFs and *supports* (e.g. between "*tweety* has wings" and the earlier "*tweety* flies as it is a bird"), variations of AAFs and BAFs such as *quantified BAFs* (where arguments are equipped with *base scores* and *dialectical strengths* obtained using *gradual semantics* [10]), *weighted (quantified) BAFs* (where edges are weighted, e.g. see [11]), and forms of *structured* CA [12], where arguments and attacks are drawn from logical formalisms. For example, in *Assumption-Based AFs (ABAFs)* [4, 13, 14, 15], arguments are drawn from "rules" and are supported by "assumptions", and attacks are targeted at the assumptions, by arguments for their "contraries".

CA and non-monotonic reasoning. CA is naturally non-monotonic and indeed there are some close connections between (forms of) CA and several formalisms for non-monotonic reasoning [3, 4], including *normal* logic programs (LPs in short) with negation as failure (NAF) and *answer set programs* (ASPs in short). In particular, ABAFs admit LPs/ASPs as instances. For illustration, consider the LP/ASP $P \cup F$ with rules $P = \{\mathit{flies}(X) \leftarrow \mathit{bird}(X), \mathit{not}\ \neg\mathit{flies}(X),\ \neg\mathit{flies}(X) \leftarrow \mathit{penguin}(X)\}$ (with $\neg$ interpreted syntactically, and not denoting NAF) and facts $F = \{\mathit{penguin}(\mathit{tweety}), \mathit{bird}(\mathit{tweety})\}$. This LP/ASP corresponds to an ABAF with

- rules $P \cup F$,
- all NAF literals in the vocabulary of the LP as assumptions,
- contraries of assumptions $\mathit{not}\ l$ given by l.

Note that, in particular, assumptions include $\mathit{not}\ \neg\mathit{flies}(\mathit{tweety})$, with contrary $\neg\mathit{flies}(\mathit{tweety})$). The semantics of the original LP/ASP is exactly captured by the semantics of the ABAF [4]. For example, stable extensions of ABAFs correspond exactly to stable models of LPs/ASPs [16].

[1] We will refer to several standard NN architectures, notably multi-layer perceptrons (MLPs), convolutional neural networks (CNNs), Long-Short-Term-Memory NNs (LSTMs), and AutoEncoders (AEs).

2. Overview of existing work in NAL

Here, we focus on works at the intersection of CA and NNs positioned as NAL, while ignoring works combining CA and NNs for supporting CA itself, e.g.: NNs performing argumentative reasoning, as in [17, 18, 19]; works on argument mining, using NNs to extract arguments and/or AFs, as in [20]; methods using NNs for computing semantics of AFs, such as [21, 22, 23]; works on the correspondence between AFs of various types and NNs, as in [24, 25].

We summarise existing works on NAL in Table 1 and related works in neuro-symbolic learning with LP/ASP in Table 2, and describe their main characteristics for our purposes below.

Translation of NNs to AFs. A line of work includes methods that take a trained NN as an input and translate it into an AF of some kind. Amongst these methods, *DAX* [26] envisages generic mappings into generalised AFs, with any number of dialectical relations. Several concrete instantiations of this approach focus on mapping NNs (e.g. MLPs or CNNs) into BAFs [27, 28]. Further, *SpArX* [29] relies upon a provable one-to-one-correspondence between MLPs and weighted BAFs [11] to translate MLPs into sparse weighted BAFs, where hidden neurons are clustered, for the purpose of rendering the MLP sparser and more interpretable. These methods aim at using the AFs resulting from the translation for *explaining* the NNs.

Pipeline methods. Some other methods use NNs to extract AFs, and then apply symbolic, argumentative reasoning on the AFs to support some downstream task. For example, *ADA* [30] uses neural classifiers and LSTMs to mine quantified BAFs from textual reviews, and then reasons with them (under gradual semantics) to provide recommendations for movies, evaluated against review aggregations measures. Also, *DEAr* [31] uses AutoEncoders for feature selection from tabular data to extract AAFs, and then reasons with them (under extension-based semantics) to provide explainable classifications from tabular inputs. Further, *Local-HDP-ABL* [32] deploys (possibly neural) feature selection methods with images to extract BAFs for reasoning over the images for explainable classification. In these methods the two components (AF extraction and symbolic, argumentative reasoning) are decoupled within a pipeline, and the argumentative reasoning provides "delayed" feedback to the AF extraction by the NN.

Integrated methods. Finally, argumentative reasoning over AFs has been integrated within NNs as a form of inductive bias, to guide the learning within the NNs. Specifically, *NSAM* [33]

Paper	*Formalism*	*Learning Setting*
DAX [26, 27, 28]	BAFs	Translation of NNs (MLPs, CNNs) to AFs
SpArX [11, 29]	Weighted BAFs	Translation of MLPs to AFs
ADA [30]	Quantified BAFs	AFs learnt by NNs $\rightarrow$ argumentative reasoning
DeAr [31]	AAFs	AFs learnt by AEs $\rightarrow$ argumentative reasoning
Local-HDP-ABL [32]	BAFs	AFs learnt by NNs $\rightarrow$ argumentative reasoning
NSAM [33]	probabilistic semi-AAFs	Reasoning with AFs as inductive bias

Table 1
NAL approaches based on argumentation frameworks (AFs) of various kinds

Paper	*Formalism*	*Learning Setting*
CIL^{2}P[35]	Propositional LP	Translation of NNs to LP
NeSyFOLD[36]	Stratified LP	LP learnt from pre-trained CNN
FFNSL[37]	ASP	LP learnt from pre-trained NNs
NeurASP[38]	Probabilistic ASP	e2e training of NNs via user-defined rules
DeepProbLog[39]	Probabilistic Stratified LP	e2e training of NNs via user-defined rules
NeuroLog[40]	Abductive LP	e2e training of NNs via user-defined rules
embed2sym[41]	ASP	e2e training of NNs via user-defined rules
pix2rule[42]	ASP	e2e learning of ASP

Table 2
Neuro-symbolic approaches based on non-monotonic rules (mappable to AFs)

defines argumentation Boltzmann machines (a form of NNs capturing argumentative knowledge, in the form of *probabilistic semi-abstract AFs*), trained on instances of the argumentative knowledge applicable to given data, to make predictions that can be explained argumentatively.

Related work in logic programming. We include neuro-symbolic paradigms integrating forms of non-monotonic LPs and ASPs, as those have a natural understanding in CA terms. We omit instead neuro-symbolic methods focusing on *positive* logic programs (e.g. see overview in [34]), as they support monotonic reasoning only and are not relevant to an argumentative viewpoint. The LP works are summarised in Table 2. *CIL*^{2}*P* [35] extracts LPs from MLPs, for the purposes of explainability. Some methods follow a pipeline approach, learning LPs from NNs pre-trained for feature extraction on images: *NeSyFOLD* [36] generates stratified LPs from CNNs, and *FFNSL* [37] learns ASPs. Other methods integrate reasoning with LPs/ASPs provided by humans, for the purpose of training NNs *end-to-end* (e2e) to benefit from the knowledge represented in the LPs/ASPs and learn from it as well as from unstructured data such as images: *NeurASP* [38] represents knowledge in the form of probabilistic ASPs, *DeepProbLog* [39] uses probabilistic stratified LPs (i.e. ProbLog programs), *NeuroLog* [40] makes use of abductive LP to supervise the e2e training process, and *embed2sym* [41] integrates clustering for feature extraction with ASPs. Finally, *pix2rule* [42] learns ASPs e2e, through a training regime that processes both images and ASPs, by including a differentiable layer in NNs from which ASP rules can be extracted.

3. Challenges

We conclude by discussing the role that NAL could have in the future development of neuro-symbolic systems in three settings: 1) the NN component is pre-trained, AFs are learnt; 2) AFs are predefined, the NN component is learnt; 3) both components are learnt e2e.

NNs pre-trained, AFs learnt. The simplest way in which CA could be used is within a simple pipeline architecture: an AF is extracted (or learnt) from a pre-trained NN. This approach is very close to standard CA-based forms of XAI as discussed earlier [30, 31, 32]. A more advanced variant of this pipeline architecture could be implemented by representing pre-trained NN modules symbolically as (probabilistic) neural predicates, similarly to [36, 39]. With this representation, together with suitable background knowledge by domain experts, we could learn symbolic concepts as ABAFs (e.g. as suggested in [43]) or as probabilistic ABAFs [44].

Predefined AFs, NNs learnt. In this type of systems NNs are trained e2e: the input to NNs is labeled by symbolic concepts for which we provide a suitable background knowledge defined by means of AFs. In this context, ABA frameworks could be used to guide the training of NNs, again extending the work done in the area of LP and ASP [36, 39]. The ability of ABAFs to formalise various kinds of knowledge/reasoning (e.g. default theories, ontologies, and temporal logics) would be a plus for achieving better accuracy and reliability with smaller data sets.

NNs and AFs both learnt. The most challenging task is to construct systems that are composed of neural modules and symbolic modules where both components are learnt at the same time, while training is done e2e on their composition. A critical sub-task is learning *latent concepts*, that is, the symbolic concepts associated with the output of the neural component. Initial work in this direction, for LP/ASP, could provide a starting point. For example, [45] design an approach based on so-called *policy functions*, similar to reinforcement learning, to learn symbolic knowledge which could be seen as a collection of facts (and thus *monotonic*). While this work makes quite strong hypotheses on the form of the symbolic knowledge to be learnt, it would be interesting to explore whether the approach could be generalised to learn (non-monotonic) AFs. Furthermore, the aforementioned [42] – proposing a specific approach for e2e learning of relations and rules from images – could provide a fruitful starting point to learn AFs.

Although CA-based approaches have not been considered in this setting, they might be advantageous. First, the features of ABAFs useful for learning the two components separately can also be useful for their combined learning. For example, in this more complex scenario we could exploit ABAFs, following an iterative approach, for a variety of tasks such as: represent rich background knowledge, use that knowledge to generate suitable neural model templates, to train them, and extract new knowledge from trained neural models. Further, for this more complex task, we need representation and reasoning mechanisms that cope with the non-monotonicity of knowledge extraction and learning. To this aim, the ability of CA to support various forms of non-monotonic reasoning and to represent the learnt symbolic knowledge in the form of defeasible rules, could play a key role. Notably, we believe that abduction, which can be realised naturally in CA, could be useful for learning latent concepts.

4. Acknowledgments

We would like to thank the anonymous reviewers for constructive criticism. We also thank support from the Royal Society, UK (IEC\R2\222045 - International Exchanges 2022 Cost Share). M. Proietti is a member of the INDAM-GNCS research group, Italy. F. Toni was partially funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 101020934) and by J.P. Morgan and the Royal Academy of Engineering under the Research Chairs and Senior Research Fellowships scheme.

References

[1] K. Atkinson, P. Baroni, M. Giacomin, A. Hunter, H. Prakken, C. Reed, G. R. Simari, M. Thimm, S. Villata, Towards artificial argumentation, AI Magazine 38 (2017) 25–36.

[2] P. Baroni, D. Gabbay, M. Giacomin, L. van der Torre (Eds.), Handbook of Formal Argumentation, College Publications, 2018.

[3] P. Dung, On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games, Artif. Intell. 77 (1995) 321–358. doi:10.1016/0004-3702(94)00041-X.

[4] A. Bondarenko, P. Dung, R. Kowalski, F. Toni, An abstract, argumentation-theoretic approach to default reasoning, Artif. Intell. 93 (1997) 63–101. doi:10.1016/S0004-3702(97)00015-5.

[5] K. Cyras, A. Rago, E. Albini, P. Baroni, F. Toni, Argumentative XAI: A survey, in: Proc. IJCAI, ijcai.org, 2021, pp. 4392–4399. doi:10.24963/ijcai.2021/600.

[6] C. Antaki, I. Leudar, Explaining in conversation: Towards an argument model, Europ. J. of Social Psychology 22 (1992) 181–194.

[7] T. Miller, Explanation in artificial intelligence: Insights from the social sciences, Artif. Intell. 267 (2019) 1–38. doi:10.1016/j.artint.2018.07.007.

[8] O. Cocarascu, F. Toni, Argumentation for machine learning: A survey, in: Proc. COMMA 2016, FAIA 287, IOS Press, 2016, pp. 219–230. doi:10.3233/978-1-61499-686-6-219.

[9] C. Cayrol, M.-C. Lagasquie-Schiex, On the acceptability of arguments in bipolar argumentation frameworks, in: Proc. 8th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, Springer, 2005, pp. 378–389.

[10] P. Baroni, A. Rago, F. Toni, From fine-grained properties to broad principles for gradual argumentation: A principled spectrum, Int. J. Approx. Reason. 105 (2019) 252–286. doi:10.1016/j.ijar.2018.11.019.

[11] N. Potyka, Interpreting neural networks as quantitative argumentation frameworks, in: Proc. AAAI 2021, AAAI Press, 2021, pp. 6463–6470. URL: https://ojs.aaai.org/index.php/AAAI/article/view/16801.

[12] P. Besnard, A. J. García, A. Hunter, S. Modgil, H. Prakken, G. R. Simari, F. Toni, Introduction to structured argumentation, Argument Comput. 5 (2014) 1–4. doi:10.1080/19462166.2013.869764.

[13] P. Dung, R. Kowalski, F. Toni, Assumption-based argumentation, in: Argumentation in Artificial Intelligence, Springer, 2009, pp. 199–218. doi:10.1007/978-0-387-98197-0_10.

[14] F. Toni, A tutorial on assumption-based argumentation, Argument & Computation 5 (2014) 89–117. doi:10.1080/19462166.2013.869878.

[15] K. Cyras, X. Fan, C. Schulz, F. Toni, Assumption-based argumentation: Disputes, explanations, preferences, FLAP 4 (2017).

[16] M. Gelfond, V. Lifschitz, The stable model semantics for logic programming, in: ICLP, MIT Press, 1988, pp. 1070–1080.

[17] N. Slonim, Y. Bilu, C. Alzate, R. Bar-Haim, B. Bogin, F. Bonin, L. Choshen, E. Cohen-Karlik, L. Dankin, L. Edelstein, L. Ein-Dor, R. Friedman-Melamed, A. Gavron, A. Gera, M. Gleize, S. Gretz, D. Gutfreund, A. Halfon, D. Hershcovich, R. Hoory, Y. Hou, S. Hummel, M. Jacovi,

C. Jochim, Y. Kantor, Y. Katz, D. Konopnicki, Z. Kons, L. Kotlerman, D. Krieger, D. Lahav, T. Lavee, R. Levy, N. Liberman, Y. Mass, A. Menczel, S. Mirkin, G. Moshkowich, S. Ofek-Koifman, M. Orbach, E. Rabinovich, R. Rinott, S. Shechtman, D. Sheinwald, E. Shnarch, I. Shnayderman, A. Soffer, A. Spector, B. Sznajder, A. Toledo, O. Toledo-Ronen, E. Venezian, R. Aharonov, An autonomous debating system, Nat. 591 (2021) 379–384. doi:10.1038/s41586-021-03215-w.

[18] P. Goffredo, S. Haddadan, V. Vorakitphan, E. Cabrio, S. Villata, Fallacious argument classification in political debates, in: Proc. IJCAI 2022, ijcai.org, 2022, pp. 4143–4149. doi:10.24963/ijcai.2022/575.

[19] L. Thorburn, A. Kruger, Optimizing language models for argumentative reasoning, in: Proc. ArgML 2022, CEUR Workshop Proceedings 3208, CEUR-WS.org, 2022, pp. 27–44. URL: http://ceur-ws.org/Vol-3208/paper3.pdf.

[20] O. Cocarascu, E. Cabrio, S. Villata, F. Toni, Dataset independent baselines for relation prediction in argument mining, in: Proc. COMMA 2020, FAIA 326, IOS Press, 2020, pp. 45–52. doi:10.3233/FAIA200490.

[21] R. Riveret, D. Korkinof, M. Draief, J. Pitt, Probabilistic abstract argumentation: An investigation with Boltzmann machines, Argument Comput. 8 (2017) 89. doi:10.3233/AAC-170016.

[22] I. Kuhlmann, M. Thimm, Using graph convolutional networks for approximate reasoning with abstract argumentation frameworks: A feasibility study, in: Proc. SUM 2019, LNCS 11940, Springer, 2019, pp. 24–37. doi:10.1007/978-3-030-35514-2_3.

[23] J. Klein, I. Kuhlmann, M. Thimm, Graph neural networks for algorithm selection in abstract argumentation, in: Proc. ArgML 2022, CEUR Workshop Proceedings 3208, CEUR-WS.org, 2022, pp. 81–95. URL: http://ceur-ws.org/Vol-3208/paper6.pdf.

[24] A. S. d'Avila Garcez, D. M. Gabbay, L. C. Lamb, Value-based argumentation frameworks as neural-symbolic learning systems, J. Log. Comput. 15 (2005) 1041–1058. doi:10.1093/logcom/exi057.

[25] A. S. d'Avila Garcez, D. M. Gabbay, L. C. Lamb, A neural cognitive model of argumentation with application to legal inference and decision making, J. Appl. Log. 12 (2014) 109–127. doi:10.1016/j.jal.2013.08.004.

[26] E. Albini, P. Lertvittayakumjorn, A. Rago, F. Toni, DAX: Deep Argumentative eXplanation for neural networks, CoRR abs/2012.05766 (2020). URL: https://arxiv.org/abs/2012.05766.

[27] A. Dejl, P. He, P. Mangal, H. Mohsin, B. Surdu, E. Voinea, E. Albini, P. Lertvittayakumjorn, A. Rago, F. Toni, Argflow: A toolkit for deep argumentative explanations for neural networks, in: Proc. AAMAS 2021, ACM, 2021, pp. 1761–1763. doi:10.5555/3463952.3464229.

[28] P. Sukpanichnant, A. Rago, P. Lertvittayakumjorn, F. Toni, Neural QBAFs: Explaining neural networks under LRP-based argumentation frameworks, in: Proc. AIxIA 2021, LNCS 13196, Springer, 2021, pp. 429–444. doi:10.1007/978-3-031-08421-8_30.

[29] H. Ayoobi, N. Potyka, F. Toni, SpArX: Sparse Argumentative eXplanations for neural networks, CoRR abs/2301.09559 (2023). doi:10.48550/arXiv.2301.09559.

[30] O. Cocarascu, A. Rago, F. Toni, Extracting dialogical explanations for review aggregations with argumentative dialogical agents, in: Proc. AAMAS 2019, IFAAMS, 2019, pp. 1261–1269. URL: http://dl.acm.org/citation.cfm?id=3331830.

[31] O. Cocarascu, A. Stylianou, K. Cyras, F. Toni, Data-empowered argumentation for dialectically explainable predictions, in: Proc. ECAI 2020, FAIA 325, IOS Press, 2020, pp. 2449–2456.
[32] H. Ayoobi, S. H. Kasaei, M. Cao, R. Verbrugge, B. Verheij, Explain what you see: Open-ended segmentation and recognition of occluded 3D objects, CoRR abs/2301.07037 (2023). doi:10.48550/arXiv.2301.07037, to appear at ICRA2023.
[33] R. Riveret, S. N. Tran, A. S. d'Avila Garcez, Neuro-symbolic probabilistic argumentation machines, in: Proc. KR, 2020, pp. 871–881. doi:10.24963/kr.2020/90.
[34] A. Cropper, S. Dumancic, R. Evans, S. H. Muggleton, Inductive logic programming at 30, Mach. Learn. 111 (2022) 147–172. doi:10.1007/s10994-021-06089-1.
[35] A. S. d'Avila Garcez, K. Broda, D. M. Gabbay, Symbolic knowledge extraction from trained neural networks: A sound approach, Artif. Intell. 125 (2001) 155–207. doi:10.1016/S0004-3702(00)00077-1.
[36] P. Padalkar, H. Wang, G. Gupta, NeSyFOLD: A system for generating logic-based explanations from convolutional neural networks, CoRR abs/2301.12667 (2023). doi:10.48550/arXiv.2301.12667.
[37] D. Cunnington, M. Law, J. Lobo, A. Russo, FFNSL: Feed-Forward Neural-Symbolic Learner, Mach. Learn. 112 (2023) 515–569. doi:10.1007/s10994-022-06278-6.
[38] Z. Yang, A. Ishay, J. Lee, NeurASP: Embracing neural networks into answer set programming, in: C. Bessiere (Ed.), Proc. IJCAI 2020, ijcai.org, 2020, pp. 1755–1762. doi:10.24963/ijcai.2020/243.
[39] R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester, L. D. Raedt, Neural probabilistic logic programming in DeepProbLog, Artif. Intell. 298 (2021) 103504. doi:10.1016/j.artint.2021.103504.
[40] E. Tsamoura, T. M. Hospedales, L. Michael, Neural-symbolic integration: A compositional perspective, in: Proc. AAAI 2021, AAAI Press, 2021, pp. 5051–5060. URL: https://ojs.aaai.org/index.php/AAAI/article/view/16639.
[41] Y. Aspis, K. Broda, J. Lobo, A. Russo, Embed2Sym - Scalable neuro-symbolic reasoning via clustered embeddings, in: Proc. KR, 2022. URL: https://proceedings.kr.org/2022/44/.
[42] N. Cingillioglu, A. Russo, pix2rule: End-to-end neuro-symbolic rule learning, in: Proc. IJCLR 2021, CEUR Workshop Proceedings 2986, CEUR-WS.org, 2021, pp. 15–56. URL: http://ceur-ws.org/Vol-2986/paper3.pdf.
[43] M. Proietti, F. Toni, Learning assumption-based argumentation frameworks, in: Proc. ILP, 2022.
[44] P. Dung, P. M. Thang, Towards (probabilistic) argumentation for jury-based dispute resolution, in: Proc. COMMA 2010, FAIA 216, IOS Press, 2010, pp. 171–182. doi:10.3233/978-1-60750-619-5-171.
[45] A. Daniele, T. Campari, S. Malhotra, L. Serafini, Deep symbolic learning: Discovering symbols and rules from perceptions, CoRR abs/2208.11561 (2022). doi:10.48550/arXiv.2208.11561. arXiv:2208.11561.