



Neural-Symbolic Argumentation Mining: An Argument in Favor of Deep Learning and Reasoning

Andrea Galassi^{1*}, Kristian Kersting², Marco Lippi³, Xiaoting Shao² and Paolo Torroni¹

¹ Department of Computer Science and Engineering, University of Bologna, Bologna, Italy, ² Computer Science Department and Centre for Cognitive Science, TU Darmstadt, Darmstadt, Germany, ³ Department of Sciences and Methods for Engineering, University of Modena and Reggio Emilia, Reggio Emilia, Italy

Deep learning is bringing remarkable contributions to the field of argumentation mining, but the existing approaches still need to fill the gap toward performing advanced reasoning tasks. In this position paper, we posit that neural-symbolic and statistical relational learning could play a crucial role in the integration of symbolic and sub-symbolic methods to achieve this goal.

Keywords: neural symbolic learning, argumentation mining, probabilistic logic programming, integrative AI, DeepProbLog, Ground-Specific Markov Logic Networks

OPEN ACCESS

Edited by:

Balaraman Ravindran,
Indian Institute of Technology Madras,
India

Reviewed by:

Elena Bellodi,
University of Ferrara, Italy
Frank Zenker,
Boğaziçi University, Turkey

*Correspondence:

Andrea Galassi
a.galassi@unibo.it

Specialty section:

This article was submitted to
Machine Learning and Artificial
Intelligence,
a section of the journal
Frontiers in Big Data

Received: 21 October 2019

Accepted: 31 December 2019

Published: 22 January 2020

Citation:

Galassi A, Kersting K, Lippi M, Shao X
and Torroni P (2020) Neural-Symbolic
Argumentation Mining: An Argument
in Favor of Deep Learning and
Reasoning. *Front. Big Data* 2:52.
doi: 10.3389/fdata.2019.00052

1. INTRODUCTION

The goal of argumentation mining (AM) is to automatically extract arguments and their relations from a given document (Lippi and Torroni, 2016). The majority of AM systems follows a pipeline scheme, starting with simpler tasks, such as argument component detection, down to more complex tasks, such as argumentation structure prediction. Recent years have seen the development of a large number of techniques in this area, on the wake of the advancements produced by deep learning on the whole research field of natural language processing (NLP). Yet, it is widely recognized that the existing AM systems still have a large margin of improvement, as good results have been obtained with some genres where prior knowledge on the structure of the text eases some AM tasks, but other genres, such as legal cases and social media documents still require more work (Cabrio and Villata, 2018). Performing and understanding argumentation requires advanced reasoning capabilities, which are natural human skills, but are difficult to learn for a machine. Understanding whether a given piece of evidence supports a given claim, or whether two claims attack each other, are complex problems that humans can address thanks to their ability to exploit commonsense knowledge, and to perform reasoning and inference. Despite the remarkable impact of deep neural networks in NLP, we argue that these techniques alone will not suffice to address such complex issues.

We envisage that a significant advancement in AM could come from the combination of symbolic and sub-symbolic approaches, such as those developed in the Neural Symbolic (NeSy) (Garcez et al., 2015) or Statistical Relational Learning (SRL) (Getoor and Taskar, 2007; De Raedt et al., 2016; Kordjamshidi et al., 2018) communities. This issue is also widely recognized as one of the major challenges for the whole field of artificial intelligence in the coming years (LeCun et al., 2015).

In computational argumentation, structured arguments have been studied and formalized for decades using models that can be expressed in a logic framework (Bench-Capon and Dunne, 2007). At the same time, AM has rapidly evolved by exploiting state-of-the-art neural architectures coming from deep learning. So far, these two worlds have progressed largely independently of each other.

Only recently, a few works have taken some steps toward the integration of such methods, by applying techniques combining sub-symbolic classifiers with knowledge expressed in the form of rules and constraints to AM. For instance, Niculae et al. (2017) adopted structured support vector machines and recurrent neural networks to collectively classify argument components and their relations in short documents, by hard-coding contextual dependencies and constraints of the argument model in a factor graph. A joint inference approach for argument component classification and relation identification was instead proposed by Persing and Ng (2016), following a pipeline scheme where integer linear programming is used to enforce mathematical constraints on the outcomes of a first-stage set of classifiers. More recently, Cocarascu and Toni (2018) combined a deep network for relation extraction with an argumentative reasoning system that computes the dialectical strength of arguments, for the task of determining whether a review is truthful or deceptive.

We propose to exploit the potential of both symbolic and sub-symbolic approaches for AM, by combining both results in systems that are capable of modeling knowledge and constraints with a logic formalism, while maintaining the computational power of deep networks. Differently from existing approaches, we advocate the use of a logic-based language for the definition of contextual dependencies and constraints, independently of the structure of the underlying classifiers. Most importantly, the approaches we outline do not exploit a pipeline scheme, but rather perform joint detection of argument components and relations through a single learning process.

2. MODELING ARGUMENTATION WITH PROBABILISTIC LOGIC

Computational argumentation is concerned with modeling and analyzing argumentation in the computational settings of artificial intelligence (Bench-Capon and Dunne, 2007; Simari and Rahwan, 2009). The formalization of arguments is usually addressed at two levels. At the argument level, the definition of formal languages for representing knowledge and specifying how arguments and counterarguments can be constructed from that knowledge is the domain of *structured argumentation* (Besnard et al., 2014). In structured argumentation, the premises and claim of the argument are made explicit, and their relationships are formally defined. However, when the discourse consists of multiple arguments, such arguments may conflict with one another and result in logical inconsistencies. A typical way of dealing with such inconsistencies is to identify sets of arguments that are mutually consistent and that collectively defeat their “attackers.” One way to do that is to abstract away from the internal structure of arguments and focus on the higher-level relations among arguments: a conceptual framework known as *abstract argumentation* (Dung, 1995).

Similarly to structured argumentation, AM too builds on the definition of an argument model, and aims to identify parts of the input text that can be interpreted as argument components (Lippi and Torroni, 2016). For example, if we take

a basic claim/evidence argument model, possible tasks could be claim detection (Aharoni et al., 2014; Lippi and Torroni, 2015), evidence detection (Rinott et al., 2015), and the prediction of links between claim and evidence (Niculae et al., 2017; Galassi et al., 2018). However, in structured argumentation the formalization of the model is the basis for an inferential process, whereby conclusions can be obtained starting from premises. In AM, instead, an argument model is usually defined in order to identify the target classes, and in some isolated cases to express relations, for instance among argument components (Niculae et al., 2017), but not for producing inferences that could help the AM tasks.

The languages of structured argumentation are logic-based. An influential structured argumentation system is *deductive argumentation* (Besnard and Hunter, 2001), where premises are logical formulae, which entail a claim, and entailment may be specified from a range of base logics, such as classical logic or modal logic. In *assumption-based argumentation* (Dung et al., 2009) instead arguments correspond to assumptions, which like in deductive systems prove a claim, and attacks are obtained via a notion of contrary assumptions. Another powerful framework is *defeasible logic programming* (DeLP) (García and Simari, 2004), where claims can be supported using strict or defeasible rules, and an argument supporting a claim is warranted if it defeats all its counterarguments. For example, that a cephalopod is a mollusc could be expressed by a strict rule, such as:

$$\text{mollusc}(X) \leftarrow \text{cephalopod}(X)$$

as these notions belong to an artificially defined, incontrovertible taxonomy. However, since in nature not all molluscs have a shell, and actually cephalopods are molluscs without a shell, rules used to conclude that a given specimen has or does not have a shell are best defined as defeasible. For instance, one could say:

$$\begin{aligned} \text{has_shell}(X) &< \text{mollusc}(X) \\ \sim \text{has_shell}(X) &< \text{cephalopod}(X) \end{aligned}$$

where $<$ denotes defeasible inference.

The choice of logic notwithstanding, rules offer a convenient way to describe argumentative inference. Moreover, depending on the application domain, the document genre, and the employed argument model, different constraints and rules can be enforced on the structure of the underlying network of arguments. For example, if we adopt a DeLP-like approach, strict rules can be used to define the relations among argument components, and defeasible rules to define context knowledge. For example, in a hypothetical claim-premise model, support relations may be defined exclusively between a premise and a claim. Such structural properties could be expressed by the following strict rules:

$$\begin{aligned} \text{claim}(Y) &\leftarrow \text{supports}(X, Y) \\ \text{premise}(X) &\leftarrow \text{supports}(X, Y) \end{aligned}$$

whereby if X supports Y , then X is a claim and Y is a premise. As another abstract example, two claims based on the same premise

may not attack each other:

$$\sim attacks(Y1, Y2) \leftarrow supports(X, Y1) \wedge supports(X, Y2)$$

As an example of defeasible rules, consider instead the context information about a political debate, where a republican candidate, R , faces a democrat candidate, D . Then one may want to use the knowledge that R 's claims and D 's claims are likely to attack each other:

$$attacks(Y1, Y2) \prec auth(Y1, R) \wedge rep(R) \wedge auth(Y2, D) \wedge dem(D)$$

where predicate $auth(A, B)$ denotes that claim A was made by B . There exist many non-monotonic reasoning systems that integrate defeasible and strict inference. However, an alternative approach that may successfully reconcile the computational argumentation view and the AM view is offered by probabilistic logic programming (PLP). PLP combines the capability of logic to represent complex relations among entities with the capability of probability to model uncertainty over attributes and relations (Riguzzi, 2018). In a PLP framework, such as PRISM (Sato and Kameya, 1997), LPAD (Vennekens et al., 2004), or ProbLog (Raedt et al., 2007), defeasible rules may be expressed by rules with a probability label. For instance, in LPAD syntax, one could write:

$$attacks(Y1, Y2) : 0.8 \leftarrow auth(Y1, R) \wedge rep(R) \wedge auth(Y2, D) \wedge dem(D)$$

to express that the above rule holds in 80% of cases. In this example, 0.8 could be interpreted as a weight or score suggesting how likely the given inference rule is to hold. In more recent approaches, such weights could be learned directly from a collection of examples, for example by exploiting likelihood maximization in a *learning from interpretations* setting (Gutmann et al., 2011) or by using a generalization of expectation maximization applied to logic programs (Bellodi and Riguzzi, 2013).

From a higher-level perspective, rules could be exploited also to model more complex relations between arguments or even to encode argument schemes (Walton et al., 2008), for example to assess whether an argument is defeated by another, on the basis of the strength of its premises and claims. This is an even more sophisticated reasoning task, which yet could be addressed within the same kind of framework described so far.

3. COMBINING SYMBOLIC AND SUB-SYMBOLIC APPROACHES

The usefulness of deep networks has been tested and proven in many NLP tasks, such as machine translation (Young et al., 2018), sentiment analysis (Zhang et al., 2018a), text classification (Conneau et al., 2017; Zhang et al., 2018b), relations extraction (Huang and Wang, 2017), as well as in AM (Cocarascu and Toni, 2017, 2018; Daxenberger et al., 2017; Galassi et al., 2018; Lauscher et al., 2018; Lugini and Litman, 2018; Schulz

et al., 2018). While a straightforward approach to exploit domain knowledge in AM is to apply a set of hand-crafted rules on the output of some first stage classifier (such as a neural network), NeSy or SRL approaches can directly enforce (hard or soft) constraints *during training*, so that a solution that does not satisfy them is penalized, or even ruled out. Therefore, if a neural network is trained to classify argument components, and another one¹ is trained to detect links between them, additional global constraints can be enforced to adjust the weights of the networks toward admissible solutions, as the learning process advances. Systems like DeepProbLog (Manhaeve et al., 2018), Logic Tensor Networks (Serafini and Garcez, 2016), or Grounding-Specific Markov Logic Networks (GS-MLN) (Lippi and Frasconi, 2009), to mention a few, enable such a scheme.

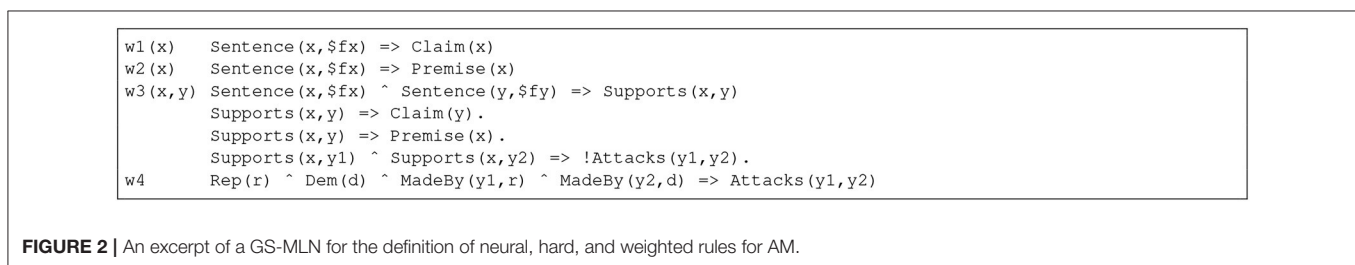
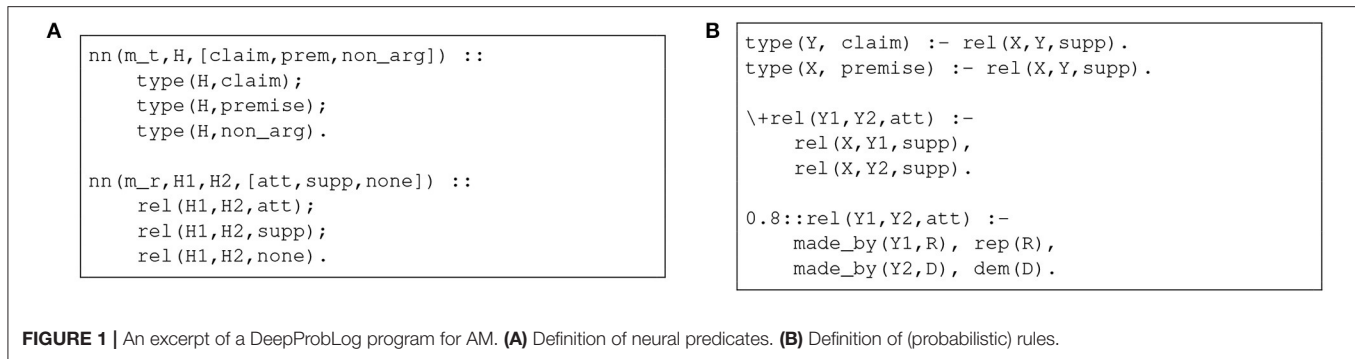
By way of illustration, we report how to implement one of the cases mentioned in section 2 with DeepProbLog and with GS-MLNs. By extending the ProbLog framework, DeepProbLog allows to introduce the following kind of construct:

$$nm(m, \vec{t}, \vec{u}) :: q(\vec{t}, u_1); \dots; q(\vec{t}, u_n).$$

The effect of the construct is the creation of a set of ground probabilistic facts, whose probability is assigned by a neural network. This mechanism allows us to delegate to a neural network m the classification of a set of predicates q defined by some input features \vec{t} . The possible classes are given by \vec{u} . Therefore, in the AM scenario, it would be possible, for example, to exploit two networks m_t and m_r to classify, respectively, the type of a potential argumentative component and the potential relation between two components. **Figure 1A** shows the corresponding DeepProbLog code. These predicates could be easily integrated within a probabilistic logic program designed for argumentation, so as to model (possibly weighted) constraints, rules, and preferences, such as those described in section 2. **Figure 1B** illustrates one such possibility.

The same scenario can be modeled using GS-MLNs. In the Markov logic framework, first-order logic rules are associated with a real-valued weight. The higher the weight, the higher the probability that the clause is satisfied, other things being equal. The weight could possibly be infinite, to model hard constraints. In the GS-MLN extension, different weights can be associated to different groundings of the same formula, and such weights can be computed by neural networks. Joint training and inference can be performed, as a straightforward extension of the classic Markov logic framework (Lippi and Frasconi, 2009). **Figure 2** shows an example of a GS-MLN used to model the AM scenario that we consider. Here, the first three rules model grounding-specific clauses (the dollar symbol indicating a reference to a generic vector of features describing the example) whose weights depend on the specific groundings (variables x and y); the three subsequent rules are hard constraints (ending with a period); the

¹In a multi-task setting, instead of two networks, the same network could be used to perform both component classification and link detection at the same time. In AM, this approach has already been successfully explored (Eger et al., 2017; Galassi et al., 2018).



final rule is a classic Markov logic rule, with a weight attached to the first-order clause.

The kind of approach hereby described strongly differs from the existing approaches in AM. Whereas, Persing and Ng (2016) exploit a pipeline scheme to apply the constraints to the predictions made by deep networks at a first stage of computation, the framework we propose can perform a joint training, which includes the constraints within the learning phase. This can be viewed as an instance of *Constraint Driven Learning* (Chang et al., 2012) and its continuous counterpart, posterior regularization (Ganchev et al., 2010), where multiple signals contribute to a global decision, by being pushed to satisfy expectations on the global decision. Differently from the work by Niculae et al. (2017), who use factor graphs to encode interdependencies between random variables, our approach enables to exploit the interpretable formalism of logic to represent rules. Moreover, the models of NeSy and SRL are typically able to learn the weights or the probabilities of the rules, or even to learn the rules themselves, thus addressing a *structure learning* task.

4. DISCUSSION

After many years of growing interest and remarkable results, time is ripe for AM to move forward in its ability to support complex arguments. To this end, we argue that research in this area should aim at combining sub-symbolic and symbolic approaches, and that several state-of-the-art ML frameworks already provide the necessary ground for such a leap forward.

The combination of such approaches will leverage different forms of abstractions that we consider essential for AM. On the one hand, (probabilistic) logical representations enable to specify AM systems in terms of data, world knowledge and other constraints, and to express uncertainties at a logical and

conceptual level rather than at the level of individual random variables. This would make AM systems easier to interpret—a feature that is now becoming a need for AI in general (Guidotti et al., 2018)—since they could help explaining the logic and the reasons that lead them to produce their arguments, while still dealing with the uncertainties stemming from the data and the (incomplete) background knowledge. On the other hand, AM is too complex to fully specify the distributions of random variables and their global (in)dependency structure a priori. Sub-symbolic models can harness such a complexity by finding the right, general outline, in the form of computational graphs, and processing data.

In order to fully exploit the potential of this joint approach, clearly many challenges have to be faced. First of all, several languages and frameworks for NeSy and SRL exist, each with its own characteristics in terms of both expressive power and efficiency. In this sense, AM would represent an ideal test-bed for such frameworks, by presenting a challenging, large-scale application domain where the exploitation of background knowledge could play a crucial role to boost performance. Inference in this kind of models is clearly an issue, thus AM would provide additional benchmarks for the development of efficient algorithms, both in terms of memory consumption and running time. Finally, although there are already several NeSy and SRL frameworks available, being these research areas still relatively young and in rapid development, their tools are not yet mainstream. Here, an effort is needed in integrating such tools with state-of-the-art neural architectures for NLP.

AUTHOR CONTRIBUTIONS

All authors listed have made a substantial, direct and intellectual contribution to the work, and approved it for publication.

FUNDING

This work has been partially supported by the H2020 Project AI4EU (G.A. 825619), by the DFG project CAML (KE 1686/3-1, SPP 1999), and by the RMU project DeCoDeML.

REFERENCES

- Aharoni, E., Polnarov, A., Lavee, T., Hershovich, D., Levy, R., Rinott, R., et al. (2014). "A benchmark dataset for automatic detection of claims and evidence in the context of controversial topics," in *Proceedings of the First Workshop on Argumentation Mining* (Baltimore, MD: Association for Computational Linguistics), 64–68.
- Bellodi, E., and Riguzzi, F. (2013). Expectation maximization over binary decision diagrams for probabilistic logic programs. *Intell. Data Anal.* 17, 343–363. doi: 10.3233/IDA-130582
- Bench-Capon, T., and Dunne, P. E. (2007). Argumentation in artificial intelligence. *Artif. Intell.* 171, 619–641. doi: 10.1016/j.artint.2007.05.001
- Besnard, P., García, A. J., Hunter, A., Modgil, S., Prakken, H., Simari, G. R., et al. (2014). Introduction to structured argumentation. *Argum. Comput.* 5, 1–4. doi: 10.1080/19462166.2013.869764
- Besnard, P., and Hunter, A. (2001). A logic-based theory of deductive arguments. *Artif. Intell.* 128, 203–235. doi: 10.1016/S0004-3702(01)00071-6
- Cabrio, E., and Villata, S. (2018). "Five years of argument mining: a data-driven analysis," in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18* (Stockholm: International Joint Conferences on Artificial Intelligence Organization), 5427–5433.
- Chang, M.-W., Ratinov, L., and Roth, D. (2012). Structured learning with constrained conditional models. *Mach. Learn.* 88, 399–431. doi: 10.1007/s10994-012-5296-5
- Cocarascu, O., and Toni, F. (2017). "Identifying attack and support argumentative relations using deep learning," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (Copenhagen: Association for Computational Linguistics), 1374–1379.
- Cocarascu, O., and Toni, F. (2018). Combining deep learning and argumentative reasoning for the analysis of social media textual content using small data sets. *Comput. Linguist.* 44, 833–858. doi: 10.1162/coli_a_00338
- Conneau, A., Schwenk, H., Barrault, L., and Lecun, Y. (2017). "Very deep convolutional networks for text classification," in *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers* (Valencia: Association for Computational Linguistics), 1107–1116.
- Daxenberger, J., Eger, S., Habernal, I., Stab, C., and Gurevych, I. (2017). "What is the essence of a claim? cross-domain claim identification," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (Copenhagen: Association for Computational Linguistics), 2055–2066.
- De Raedt, L., Kersting, K., Natarajan, S., and Poole, D. (2016). *Statistical Relational Artificial Intelligence: Logic, Probability, and Computation*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.
- Dung, P. M. (1995). On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.* 77, 321–358. doi: 10.1016/0004-3702(94)00041-X
- Dung, P. M., Kowalski, R. A., and Toni, F. (2009). "Assumption-based argumentation," in *Argumentation in Artificial Intelligence*, eds G. Simari and I. Rahwan (Boston, MA: Springer US), 199–218.
- Eger, S., Daxenberger, J., and Gurevych, I. (2017). "Neural end-to-end learning for computational argumentation mining," in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (Vancouver, BC: Association for Computational Linguistics), 11–22.
- Galassi, A., Kersting, K., Lippi, M., Shao, X., and Torroni, P. (2019). Neural-symbolic argumentation mining: an argument in favour of deep learning and reasoning. arXiv:1905.09103 [Preprint] Available online at: <https://arxiv.org/abs/1905.09103>
- Galassi, A., Lippi, M., and Torroni, P. (2018). "Argumentative link prediction using residual networks and multi-objective learning," in *Proceedings of the 5th Workshop on Argument Mining* (Brussels: Association for Computational Linguistics), 1–10.
- Ganchev, K., Graça, J., Gillenwater, J., and Taskar, B. (2010). Posterior regularization for structured latent variable models. *J. Mach. Learn. Res.* 11, 2001–2049. doi: 10.5555/1756006.1859918
- Garcez, A., Besold, T. R., De Raedt, L., Földiák, P., Hitzler, P., Icard, T., et al. (2015). "Neural-symbolic learning and reasoning: contributions and challenges," in *Proceedings of the AAAI Spring Symposium on Knowledge Representation and Reasoning: Integrating Symbolic and Neural Approaches*, Stanford (Palo Alto, CA: Stanford University), 18–21.
- García, A. J., and Simari, G. R. (2004). Defeasible logic programming: an argumentative approach. *Theory Pract. Log. Program* 4, 95–138. doi: 10.1017/S1471068403001674
- Getoor, L., and Taskar, B. (2007). *Introduction to Statistical Relational Learning*, Vol. 1. Cambridge, MA: MIT Press.
- Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., and Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Comput. Surv.* 51, 93:1–93:42. doi: 10.1145/3236009
- Gutmann, B., Thon, I., and De Raedt, L. (2011). "Learning the parameters of probabilistic logic programs from interpretations," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases* (Athens: Springer), 581–596.
- Huang, Y. Y., and Wang, W. Y. (2017). "Deep residual learning for weakly-supervised relation extraction," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (Copenhagen: Association for Computational Linguistics), 1803–1807.
- Kordjamshidi, P., Roth, D., and Kersting, K. (2018). "Systems AI: a declarative learning based programming perspective," in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18* (Stockholm: International Joint Conferences on Artificial Intelligence Organization), 5464–5471.
- Lauscher, A., Glavaš, G., Ponzetto, S. P., and Eckert, K. (2018). "Investigating the role of argumentation in the rhetorical analysis of scientific publications with neural multi-task learning models," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing* (Brussels: Association for Computational Linguistics), 3326–3338.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* 531, 436–444. doi: 10.1038/nature14539
- Lippi, M., and Frasconi, P. (2009). Prediction of protein β -residue contacts by Markov logic networks with grounding-specific weights. *Bioinformatics* 25, 2326–2333. doi: 10.1093/bioinformatics/btp421
- Lippi, M., and Torroni, P. (2015). "Context-independent claim detection for argument mining," in *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015*, eds Q. Yang and M. Wooldridge (Buenos Aires: AAAI Press), 185–191.
- Lippi, M., and Torroni, P. (2016). Argumentation mining: state of the art and emerging trends. *ACM Trans. Internet Technol.* 16, 10:1–10:25. doi: 10.1145/2850417
- Lugini, L., and Litman, D. (2018). "Argument component classification for classroom discussions," in *Proceedings of the 5th Workshop on Argument Mining* (Brussels: Association for Computational Linguistics), 57–67.
- Manhæve, R., Dumancic, S., Kimmig, A., Demeester, T., and De Raedt, L. (2018). "DeepProbLog: neural probabilistic logic programming," in *Advances in Neural Information Processing Systems*, eds S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Montreal, QC: Curran Associates, Inc.), 3753–3763.

ACKNOWLEDGMENTS

We thank the reviewers for their comments and contributions, which have increased the quality of this work. A preliminary version of this manuscript has been released as a Pre-Print at Galassi et al. (2019).

- Niculae, V., Park, J., and Cardie, C. (2017). "Argument mining with structured SVMs and RNNs," in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (Vancouver, BC: Association for Computational Linguistics), 985–995.
- Persing, I., and Ng, V. (2016). "End-to-end argumentation mining in student essays," in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (San Diego, CA: Association for Computational Linguistics), 1384–1394.
- Raedt, L. D., Kimmig, A., and Toivonen, H. (2007). "ProbLog: a probabilistic prolog and its application in link discovery," in *IJCAI* (Hyderabad), 2462–2467.
- Riguzzi, F. (2018). *Foundations of Probabilistic Logic Programming*. Gistrup: River Publishers.
- Rinott, R., Dankin, L., Alzate Perez, C., Khapra, M. M., Aharoni, E., and Slonim, N. (2015). "Show me your evidence—an automatic method for context dependent evidence detection," in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing* (Lisbon: Association for Computational Linguistics), 440–450.
- Sato, T., and Kameya, Y. (1997). "PRISM: a language for symbolic-statistical modeling," in *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, IJCAI 97* (Nagoya: Morgan Kaufmann), 1330–1339.
- Schulz, C., Eger, S., Daxenberger, J., Kahse, T., and Gurevych, I. (2018). "Multi-task learning for argumentation mining in low-resource settings," in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)* (New Orleans, LA: Association for Computational Linguistics), 35–41.
- Serafini, L., and Garcez, A. D. (2016). Logic tensor networks: deep learning and logical reasoning from data and knowledge. arXiv 1606.04422 [Preprint] Available online at: <https://arxiv.org/abs/1606.04422>
- Simari, G., and Rahwan, I. (2009). *Argumentation in Artificial Intelligence*. Springer US.
- Vennekens, J., Verbaeten, S., and Bruynooghe, M. (2004). "Logic programs with annotated disjunctions," in *Logic Programming*, eds B. Demoen and V. Lifschitz (Berlin; Heidelberg: Springer Berlin Heidelberg), 431–445.
- Walton, D., Reed, C., and Macagno, F. (2008). *Argumentation Schemes*. Cambridge University Press.
- Young, T., Hazarika, D., Poria, S., and Cambria, E. (2018). Recent trends in deep learning based natural language processing. *IEEE Comput. Intell. Mag.* 13, 55–75. doi: 10.1109/MCI.2018.2840738
- Zhang, L., Wang, S., and Liu, B. (2018a). Deep learning for sentiment analysis: a survey. *Wiley Interdiscipl. Rev. Data Mining Knowl. Discov.* 8:e1253. doi: 10.1002/widm.1253
- Zhang, Y., Henao, R., Gan, Z., Li, Y., and Carin, L. (2018b). "Multi-label learning from medical plain text with convolutional residual models," in *Proceedings of the 3rd Machine Learning for Healthcare Conference, Volume 85 of Proceedings of Machine Learning Research*, eds F. Doshi-Velez, J. Fackler, K. Jung, D. Kale, R. Ranganath, B. Wallace, et al. (Palo Alto, CA: PMLR), 280–294.

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 Galassi, Kersting, Lippi, Shao and Torroni. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.