

Neuro-Symbolic Artificial Intelligence

Current Trends

Md Kamruzzaman Saker, Lu Zhou, Aaron Eberhart, and Pascal Hitzler*

Department of Computer Science, Kansas State University, KS, USA

E-mail: hitzler@ksu.edu

Abstract. Neuro-Symbolic Artificial Intelligence – the combination of symbolic methods with methods that are based on artificial neural networks – has a long-standing history. In this article, we provide a structured overview of current trends, by means of categorizing recent publications from key conferences. The article is meant to serve as a convenient starting point for research on the general topic.

Keywords: Artificial Neural Networks, Deep Learning, Knowledge Representation and Reasoning, Neuro-Symbolic Artificial Intelligence

1. What Is Neuro-Symbolic Artificial Intelligence?

We should begin by attempting to define the subject matter. In a general sense, we understand *Neuro-Symbolic Artificial Intelligence* (in short, *NeSy AI*), to be a subfield of the field of Artificial Intelligence (in short, *AI*), which focuses on bringing together, for added value, the *neural* and the *symbolic* traditions in AI. Different spellings are currently in use, that include *neural-symbolic* and *neurosymbolic*, but also *symbolic-subsymbolic* and others – which we consider to be equal. The term *neural* in this case refers to the use of artificial neural networks, or connectionist systems, in the widest sense. The term *symbolic* refers to AI approaches that are based on explicit symbol manipulation. This in general would include things like term rewriting, graph algorithms, and natural language question answering. It is often more narrowly understood, though, as a reference to methods based on formal logic, as utilized, for instance, in the subfield of AI called Knowledge Representation and Reasoning. The lines easily blur, though, and for the purposes of this overview, we will not restrict ourselves to logic-based methods only.

The general promise of NeSy AI lies in the hopes of a best-of-both worlds scenario, where the complementary strengths of neural and symbolic approaches can be combined in a favorable way. On the neural side, the desirable strengths would include trainability from raw data and robustness against faults in the underlying data, while on the symbolic side one would like to retain the inherently high explainability and provable correctness of these systems, as well as the ease of making use of deep human expert knowledge in their design and function. In terms of functional features, utilizing symbolic approaches in conjunction with machine learning – in particular with deep learning, which is currently most busily researched on – one would hope to do better on issues like out of vocabulary handling, training from small data sets, recovery from errors, and in general, explainability, as opposed to systems that rely on deep learning alone.

One of the fundamental differences between neural and symbolic AI approaches, that is relevant for our discussion, is that of representation of information within an AI system. For symbolic systems, representation is explicit and in such terms that are in principle understandable by a human. E.g., a rule such as $\text{square}(x) \rightarrow \text{rectangle}(x)$ is readily understood and manipulated by symbolic means. In neural systems, though, representations are usually by means of weighted connections between (many) neurons and/or simultaneous activations over a (possibly large) number of neurons. In particular, a human observer would not be able to readily recognize what is being represented.

*Corresponding author. E-mail: hitzler@ksu.edu.

In a deep learning context, these distributed representations are called *embeddings*, are learned during training, and are thus an explicit part of the architecture of the deep learning system. On a fundamental level, these two types of representing information are very different indeed. Representations by vectors of real numbers – or even more ephemerally, by connections weighted by real numbers – are in terms of entities within a high-dimensional and continuous, differentiable vector space as required for standard learning algorithms, and are sometimes called sub-symbolic representations. Symbolic representations, on the other hand, are by their nature discrete and their natural relationships (such as implications) do not readily lend themselves to representation into a continuous space.¹ Correspondingly, representational questions are often at the heart of the development of neuro-symbolic approaches.

Neuro-Symbolic AI has a long-standing history, with even deeper roots, such as the McCulloch & Pitts 1943 landmark paper [2], that pre-date the field of AI. A lot of work on the topic is also not explicitly placed as NeSy AI research, as the term itself only begun to be used more systematically since the 1990, and only very recently – as we will see below – has seen an increase in the context of major AI conferences. The general theme is also of importance in the context of the field of Cognitive Science [3].

We know of one series of annual workshops on the topic – the International Workshop on Neural-Symbolic Learning and Reasoning² – that dates back to 2005. But in recent years, several additional venues have been established, including the AAAI Spring Symposium AAAI-MAKE on Combining Machine Learning and Knowledge Engineering,³ the workshop on Combining Symbolic and Sub-Symbolic Methods and Their Applications (CSSA),⁴ the workshop on Cognitive Computation: Integrating neural and symbolic approaches, CoCo,⁵ and the Expo Workshop on Perspectives on Neurosymbolic Artificial Intelligence Research⁶ as well as the workshop on Knowledge Representation & Reasoning Meets Machine Learning, KR2ML,⁷ both at NeurIPS 2020. We will herein focus mostly on recent developments, from the past few years. For the reader interested in getting an overview of previous efforts, we would like to refer to past overview works such as [3–6].

The remainder of this article will focus on providing an overview of recent research contributions to the NeSy AI topic as reflected in the proceedings of leading AI conferences. In order to provide a structured approach to this, we will group these recent papers in terms of a topic categorization proposed by Henry Kautz at an AAAI 2020 address. We will also categorize the same recent papers according to a 2005 categorization proposal [4], and discuss and contrast the two. In Section 2 we will first present the two categorizations. In Section 3 we will then discuss the categorization results, and in Section 4 we will provide a future outlook.

2. Categorizing Neuro-Symbolic Artificial Intelligence

Henry Kautz, in his AAAI 2020 Robert S. Engelmore Memorial Award Lecture, discusses five categories for neuro-symbolic AI systems in his discussion about the "Future of AI":⁸

[symbolic Neuro symbolic] refers to an approach where input and output are presented in symbolic form, however all the actual processing is neural. This, in his words, is the "standard operating procedure" whenever inputs and outputs are symbolic.

[Symbolic[Neuro]] refers to a neural pattern recognition subroutine within a symbolic problem solver, with examples such as AlphaGo, AlphaZero, and current approaches to self-driving cars.

[Neuro \cup compile(Symbolic)] refers to an approach where symbolic rules are "compiled" away during training, e.g. like the 2020 work on Deep Learning For Symbolic Mathematics [7].

¹This gap between the discrete and the continuous can be bridged by mathematical means, e.g. using Cantor Space as in [1]. However the approach did not scale.

²See <http://neural-symbolic.org/>.

³<https://www.aaai-make.info/>

⁴https://alammehwish.github.io/cssa_cikm_2020/

⁵<http://neural-symbolic.org/>

⁶<https://nips.cc/Conferences/2020/Schedule?showEvent=20241>

⁷<https://kr2ml.github.io/2020/>

⁸https://www.youtube.com/watch?v=_cQITY0SPiw, starting at minute 29

[**Neuro** → **Symbolic**] refers to a cascading from a neural system into a symbolic reasoner, such as in the Neuro-Symbolic Concept-Learner [8].

[**Neuro**[**Symbolic**]] refers to the embedding of symbolic reasoning inside a neural engine, where symbolic reasoning is understood as "deliberative, type 2 reasoning" as common, e.g., in business AI, and including an internal model of the system's state of attention: Concepts are decoded to symbolic entities in an attention schema when attention to these concepts is high. A goal in the attention schema then signals that deliberative reasoning be executed.

Kautz presented these as "five ways to bring together the neural and symbolic traditions." In the brief address, only a few examples were given of corresponding work.

In a much older survey from 2005 [4], work on NeSy AI was classified according to eight dimensions, grouped into three aspects:

- Interrelation
 - (1) **Integrated versus hybrid** – this was included mostly for delineation to hybrid systems; indeed the focus was entirely on integrated systems, where symbolic processing functionalities emerge from neural structures and processes. We will have the same emphasis herein.
 - (2) **Neuronal versus connectionist** – this aspect refers to the question whether the artificial neural network architecture is dominantly chosen in order to mimic biological neural networks (neuronal) versus a more technological approach that focused on desirable computational properties as emphasized in the field of AI (connectionist).
 - (3) **Local versus distributed** – this dimension refers to the question how symbolic information is represented within the neural system, namely whether this happens in a very localist way (one or a few dedicated neurons for a symbolic piece of knowledge) or in a highly distributed way, where a representation consists of a larger number of activation values and/or weights.
 - (4) **Standard versus nonstandard** – this refers to the question whether a NeSy AI system draws on artificial neural network architectures that are *standard* in the sense that they are widely used, propagate real numbers and nodes compute only simple functions.
- Language
 - (5) **Symbolic versus logical** – this refers to the symbolic aspect of the NeSy AI system, namely whether it is based on formal logic (logical) or on other symbolic processing (such as structured datatypes like graphs).
 - (6) **Propositional versus first-order** – for the logical aspect only, a delineation whether the underlying logic is propositional in nature, or more expressive like using higher arity predicates and quantifiers.
- Usage
 - (7) **Extraction versus representation** – this aspect refers to the information flow within the system, namely whether it primarily *extracts* symbolic information from a neural substrate, or whether it is primarily based on neural representations of symbolic knowledge.
 - (8) **Learning versus reasoning** – this refers to the core functionality of the system, namely whether its focus is on machine learning or on automated symbolic reasoning.

These eight dimensions presented a view of the existing facets of the field in 2005, and examples were given for each of the dimensions. It is important to note that they were presented as dimensions, and not as binary values, e.g., a system may fall anywhere on a continuum or even fall under both aspects of the dimension (i.e., span the dimension). We gather that the above listed dimensions are mostly self-explanatory as described; further details can be found in [4].

We started out on our investigations with the hypothesis that the NeSy AI field has shifted focus in recent years, e.g. rendering some of the 2005 dimensions obsolete, at least for the time being. We will discuss this in more detail below. At this stage, however, we remark that the Kautz categories listed above seem to address only Interrelation aspects. However they cover these in a rather different way than the 2005 survey, with a focus on more precise architectural description of the system workflows.

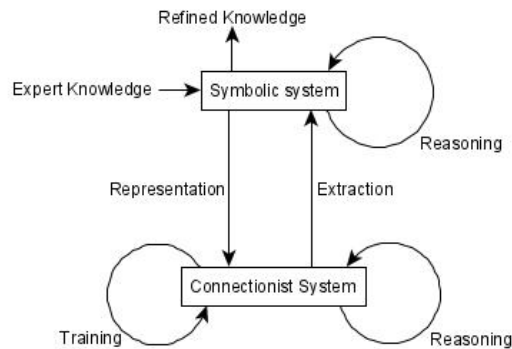


Fig. 1. Neuro-symbolic learning cycle according to [4]

Interestingly, though, a diagram presented in the 2005 survey (see Figure 1) bears striking similarity to a diagram from Kautz' presentation,⁹ and indeed the Kautz categories can be understood as utilizing part of this diagram as architectural substrate, and can serve by itself as a description of the essence of NeSy AI. In particular, in the diagram note that *I/O* is situated on the symbolic side only, while training is situated only on the neural side, whereas reasoning can happen in either part. This is congruent with the best of both worlds perspective discussed earlier, as neural systems excel in training (and in reasoning despite of faults in the data), whereas symbolic systems excel in dealing with explicit knowledge (or serving human-understandable explanations) and sport provably correct reasoning procedures.

3. Neuro-Symbolic Artificial Intelligence: Trends

For the purpose of this overview, and as a guide to recent literature, we had a look at mainstream and top tier AI publications on NeSy AI. We particularly looked at all published papers since 2011 (except ICLR which started in 2013), at top conference venues, and selected all original research papers on NeSy AI; we disregarded overview and survey works. The conferences we looked at, were

- the International Conference on Machine Learning (ICML),
- the Conference and Workshop on Neural Information Processing Systems (NeurIPS formerly NIPS),
- the AAAI Conference on Artificial Intelligence (AAAI)
- the International Joint Conference on Artificial Intelligence (IJCAI), and
- the International Conference on Learning Representations (ICLR).

In order to determine whether a paper falls into the NeSy AI theme, we read the abstract (and sometimes the introduction). As mentioned before, we are aware that not all papers relevant for NeSy AI are phrased in such terms, i.e. we acknowledge that we may have missed a few relevant papers. It still seems a reasonable assumption that the sum of our selected papers represents a valid cross-section. We are also aware that restricting our attention to the above-mentioned five conferences leaves out a lot of relevant work. However our focus was on recent, mainstream AI research, and we believe that our selection is reasonable for this purpose. It is conceivable that an analysis of publications at second-tier conferences and at workshops, or in other fields such as Cognitive Science, may provide a different picture. However this was not the focus of our effort.

The number of papers, per year, for each of the conferences is listed in Table 1. Several things are already remarkable.

- There was only one paper before 2016, and 40 papers since 2017. This alone indicates a shift towards a more mainstream adoption of the NeSy AI research area.

⁹at minute 35

conference	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	total
ICML	0	0	0	0	0	1	3	2	5	6	17
NeurIPS	0	0	0	0	0	0	0	4	2	4	10
AAAI	0	0	0	0	0	1	0	1	1	1	4
IJCAI	1	0	0	0	0	0	2	2	0	1	6
ICLR	N/A	N/A	0	0	0	0	1	1	1	3	6
total	1	0	0	0	0	2	6	10	9	15	43

Table 1

Included paper counts, per year and conference.

category	number of papers	papers
[symbolic Neuro symbolic]	13	[9–17]
[Symbolic[Neuro]]	9	[15, 18–23]
[Neuro \cup compile(Symbolic)]	10	[24–33]
[Neuro \rightarrow Symbolic]	13	[8, 23, 34–44]
[Neuro[Symbolic]]	0	N/A

Table 2

Kautz categories paper count. Two paper fit two categories.

	dimension	(a)	(b)	N/A
Interrelation	integrated (a) vs. hybrid (b)	43	0	0
	neuronal (a) vs. connectionist (b)	0	43	0
	local (a) vs. distributed (b)	2	42	0
	standard (a) vs. nonstandard (b)	43	0	0
Language	symbolic (a) vs. logical (b)	21	24	0
	propositional (a) vs. first-order (b)	3	22	18
Usage	extraction (a) vs. representation (b)	6	37	3
	learning (a) vs. reasoning (b)	19	29	0

Table 3

2005 survey dimensions paper count. Papers may cover both dimensions. N/A was used if the situation seemed unclear or not applicable.

- The total paper count, at these selected conferences, is still rather low compared to the overall number of accepted papers. E.g., IJCAI-PRICAI 2020 had 592 accepted papers, and NeurIPS 2020 had 1900 accepted papers.
- We do not have insights into the topics of rejected papers, but with acceptance rates at these conferences often at 20% or lower, it seems reasonable to expect that the total number of submitted NeSy AI papers, most of which likely published elsewhere in the meantime, was at least about five times higher.

We now manually assessed each of these papers as to their categorization with respect to the two schemes described in Section 2. Doing so was not straightforward at all, as the descriptions of the categories are somewhat ambiguous, and it is not always clear which categorization would be most adequate. Still, from the resulting Table 2 we can see that the first four Kautz categories seem to provide a reasonably balanced perspective on the recent publications at these conferences. The fifth category in fact, as also discussed by Kautz in his address, was meant to be more forward-looking, a goal to aspire to in future research. So we are not surprised to not see any papers in this category.

Categorizations with respect to the 2005 paper are given in Table 3 and a corresponding paper list in Table 4. A number of interesting observations can be made, and we go through each of them in the following.

Regarding the integrated vs. hybrid dimension, we are taking a similar tack as the 2005 survey: this is mostly for delineation, i.e. we focus on integrated approaches.

dimension	papers
integrated	[8–50]
connectionist	[8–50]
local	[17, 38]
distributed	[8–37, 39–50]
standard	[8–50]
symbolic	[9–12, 16–18, 22–25, 31, 32, 34, 36, 38, 39, 41–44]
logical	[8, 13–15, 17, 19–21, 26–30, 33, 35, 37, 40, 41, 45–50]
propositional	[14, 17, 33]
first-order	[8, 13, 15, 19–21, 26–30, 37, 40, 41, 45–50]
extraction	[17, 25, 27, 34, 36, 41]
representation	[8, 12–23, 26–33, 35, 37–50]
learning	[9–13, 17, 18, 22, 24, 25, 29, 31, 32, 34–36, 45, 46, 49]
reasoning	[8, 14–17, 19–21, 23, 25–30, 33, 35, 37–45, 47, 48, 50]

Table 4

2005 survey dimensions paper classification. Dimensions with no entries have been omitted.

Regarding the neuronal vs. connectionist dimension, it is noteworthy but perhaps unsurprising that there were no papers with emphasis on biological plausibility: Much of the current rise of NeSy AI as a topic is fueled by deep learning with a heavy emphasis on desirable computational functionality. The 2005 survey, however, shows a much more balanced picture, and indeed several early works, such as [51], put much more emphasis on biological plausibility. However, it is also conceivable that such papers may not be considered as interesting by typical top tier AI conference reviewers, and thus may seek publication outlets elsewhere.

Similar considerations may apply on both the local vs. distributed and the standard vs. nonstandard network architecture dimension. The recent heavy influence of deep learning, as well as our restriction to top tier AI conferences may explain the lopsided counts: architectures are more-or-less standard deep learning architectures, and in these, knowledge is usually encoded in a distributed fashion using some kind of embeddings in multi-dimensional Euclidean space. In the 2005 survey, nonstandard architectures and localist representations featured much more prominently, and the previously cited [51] is also an example for such work.

The counts on the symbolic vs. logical dimension mark another departure from the situation described in the 2005 paper, which indicated a much stronger emphasis on formal logic in NeSy AI,¹⁰ as opposed to symbolic aspects.

The lopsided count for first-order logic as opposed to propositional logic illustrates another shift in the subfield. It was discussed in the 2005 survey that a majority of NeSy AI work appears to focus on propositional aspects, it was previously understood that developing NeSy AI solutions is much harder when dealing with logics that are more expressive than propositional logics. John McCarthy indeed used the term "propositional fixation" of artificial neural networks to point this out, as early as 1988 [52]. That first-order logics beyond propositional play a more prominent role now is an indicator not only of new and enhanced capabilities, it also means that it is now possible to utilize logics that are of increasing relevance to the Knowledge Representation and Reasoning subfield of AI. This is an excellent and promising development. The N/A papers in this row in Table 3 indicate paper that did not involve logic at all, and instead used symbolic information that is not based on formal logic.

On the extraction vs. representation dimension, we notice that explicit work on extracting symbolic information from trained neural networks is less of an emphasis. This makes sense in the light of Kautz' categories above, where the first three categories would usually not involve an extraction aspect, and the fourth also may not. Knowledge extraction was a more prominent and emphasized topic in the past.

On the learning vs. reasoning dimension, we see a rather balanced count, which indicates that both aspects are not only important, as is to be expected, but can actually be done. In the past, both learning of first-order logic and reasoning with it was particularly difficult for systems based on artificial neural networks, but it appears that large

¹⁰Symbolic but non-logical aspects were of course also relevant, see e.g. [5] where seven out of twelve chapters did not touch on logical aspects.

feature	number of papers	papers
Out of distribution handling	8	[8, 10, 13, 17, 19, 21, 25, 35]
Interpretability	25	[8, 9, 11, 13–15, 17, 19, 20, 25–28, 30, 33, 37–39, 44–50]
Error recovery	2	[39, 45]
Learning from small data	7	[9–11, 25, 26, 28, 45]

Table 5

Some desirable neuro-symbolic features.

strides have been made since the advent of deep learning and the growing understanding that these aspects are of importance for advancing the AI field.

Out of interest, we also looked at the selected 43 papers from a perspective of some possible benefits that may arise from integrating neural and symbolic systems, as indicated in the introduction. The following seemed particularly pertinent to us:

Out of distribution handling: Whether the designed system is able to extrapolate from training input data by means of utilizing symbolic means such as background knowledge, and thus is able to handle scenarios that are significantly different from the training input.

Interpretability: Whether the fact that the system is neuro-symbolic is instrumental to enhanced interpretability of system behavior or outputs, e.g. in terms of making system decisions more transparent and explainable for a human user.

Error recovery: Whether the designed system makes use of its neuro-symbolic design in order to recover more easily from erroneous decisions or outputs.

Learning from small data: Whether the fact that the system is neuro-symbolic is instrumental in making the system trainable on less data than you would expect from a system that is not neuro-symbolic.

As for the previous categorizations, decisions how to classify each paper were often not clear-cut. We usually went with the authors’ own positioning, if there was one. The resulting counts can be found in Table 5. We see that interpretability aspects are most common. Only few papers address out of distribution issues or learning from small data, and there is hardly any work related to error recovery.

4. Paths Forward

Like so many recent topic shifts in AI, it appears that the increased attention to NeSy AI is mostly due to the rise of deep learning, which has shown tremendous success even in areas which previously were mostly in the domain of symbolic approaches, such as Chess or Go game engines. There seems to be a reasonable expectation that deep learning solutions may be much more suitable to address the subsymbolic-symbolic gap than previous connectionist machine learning technology.

However, despite these advances and promises, it seems imperative to explicitly explore and understand the capabilities and limitations of deep learning based symbolic manipulation, as a basis for further progress on combining neural and symbolic aspects in a best of both worlds fashion. Indeed, a systematic exploration of the extent to which deep learning systems can learn straightforward and well-understood symbol manipulation tasks would shed significant light on this question. Possible concrete symbol manipulation tasks for study can be found all over AI and computer science, such as term rewriting, list, tree and graph manipulations, executing formal grammars, elementary algebra, logical deduction. In-depth studies of these from a deep learning perspective would provide systems with elementary capabilities that can then be composed for more complex solutions, or used as modules in larger AI systems [53].

At the same time, it also appears to be important to contrast recent NeSy AI publications with the much larger vision of the subfield. Indeed, the current promise of NeSy AI lies in a favorable combination or integration of deep learning with symbolic AI approaches from the subfield of Knowledge Representation and Reasoning, where complex formal logics dominate. As we have seen in Section 3, current NeSy AI publications at top conferences involve complex logics to a much larger extent than was the case two decades ago, however much of the work is still

mostly on the non-logical symbolic level, and desirable features which conceptually should arise out of a systematic use of neuro-symbolic approaches, such as those listed in Table 5, are often still not met. Whether it is possible, or to what extent, to achieve a much stronger integration of complex formal logics and deep learning, including best-of-both-worlds features, is of course currently not known, and is in fact in itself a fundamental research question that remains to be addressed.

We thus posit that more emphasis is needed, in the immediate future, on deepening the logical aspects in NeSy AI research even further, and to work towards a systematic understanding and toolbox for utilizing complex logics in this context. In particular, it would be helpful for advancing the field to conduct experiments with different levels of sophistication in the logical aspects. For example, systems that utilize "flat" annotations (metadata tags that are simply keywords) are essentially operating on the logical level of "facts" only. We can then ask the question whether such a system can be improved by using, say, a class hierarchy of annotations (such as schema.org), which corresponds to making use of a knowledge base with simple logical implications (i.e., subclass relationships). If the answer to this is affirmative, then more complex logical background knowledge can be attempted to be leveraged. The Semantic Web field of research [54] for example, provides a whole hierarchy of logics on metadata [55] that can be utilized, at the level at which a deep learning system can ingest them for added value, and can thus be an excellent playground for developing fundamental NeSy AI capabilities [56].

An example for a currently emerging line of research compatible with the thoughts just laid out, is that of so-called *deep deductive reasoners* [57], which are deep learning systems trained to perform deductive reasoning over formal logics. The deductive reasoning task can be framed as classification (a statement is or is not a logical consequence of a given theory) or as regression (produce all logical consequences under some finiteness constraints). Different deep learning architectures have already been investigated, on logics that range from very straightforward ones like RDF [55] to first-order predicate logic. While the body of work is still relatively small, it is obvious from the published results [29, 58–61] that deductive reasoning is a very hard task for deep learning, and with increasing hardness and scalability issues as logics become more complex.

5. Conclusions

The recent rise in publications, events, and public addresses on NeSy AI indicates that there is increasing awareness of the importance of the subfield in the mainstream AI community. Arguably, this is also happening at the inflection point in time when we are slowly beginning to explore and understand the inherent limitations of pure deep learning approaches. The use of additional background knowledge is a natural path to attempt to further improve deep learning systems, and much of this line of work falls into the NeSy AI theme.

The broad consensus seems to be that combining neural and symbolic approaches in the sense of NeSy AI is at least a path forward to much stronger AI systems. Arguably, it may also turn out to be a major stepping stone towards human-level artificial intelligence.

References

- [1] P. Hitzler and A.K. Seda, Generalized metrics and uniquely determined logic programs, *Theor. Comput. Sci.* **305**(1–3) (2003), 187–219. doi:10.1016/S0304-3975(02)00709-0.
- [2] W.S. McCulloch and W. Pitts, A Logical Calculus of the Ideas Immanent in Nervous Activity, *Bulletin of Mathematical Biophysics* **5** (1943), 115–133.
- [3] T.R. Besold, A.S. d’Avila Garcez, S. Bader, H. Bowman, P.M. Domingos, P. Hitzler, K. Kühnberger, L.C. Lamb, D. Lowd, P.M.V. Lima, L. de Penning, G. Pinkas, H. Poon and G. Zaverucha, Neural-Symbolic Learning and Reasoning: A Survey and Interpretation, *CoRR abs/1711.03902* (2017). <http://arxiv.org/abs/1711.03902>.
- [4] S. Bader and P. Hitzler, Dimensions of Neural-symbolic Integration – A Structured Survey, in: *We Will Show Them! Essays in Honour of Dov Gabbay, Volume One*, S.N. Artëmov, H. Barringer, A.S. d’Avila Garcez, L.C. Lamb and J. Woods, eds, College Publications, 2005, pp. 167–194.
- [5] B. Hammer and P. Hitzler (eds), *Perspectives of Neural-Symbolic Integration*, Studies in Computational Intelligence, Vol. 77, Springer, 2007. ISBN 978-3-540-73953-1.
- [6] A.S. d’Avila Garcez, L.C. Lamb and D.M. Gabbay, *Neural-Symbolic Cognitive Reasoning*, Cognitive Technologies, Springer, 2009. ISBN 978-3-540-73245-7. doi:10.1007/978-3-540-73246-4.

- [7] G. Lample and F. Charton, Deep Learning For Symbolic Mathematics, in: *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*, OpenReview.net, 2020. <https://openreview.net/forum?id=S1eZYeHFDS>.
- [8] J. Mao, C. Gan, P. Kohli, J.B. Tenenbaum and J. Wu, The Neuro-Symbolic Concept Learner: Interpreting Scenes, Words, and Sentences From Natural Supervision, in: *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, OpenReview.net, 2019. <https://openreview.net/forum?id=rJgMlhRctm>.
- [9] F. Yang, D. Lyu, B. Liu and S. Gustafson, PEORL: Interpreting Symbolic Planning and Hierarchical Reinforcement Learning for Robust Decision-Making, in: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, J. Lang, ed., ijcai.org, 2018, pp. 4860–4866. doi:10.24963/ijcai.2018/675.
- [10] D. Demeter and D. Downey, Just Add Functions: A Neural-Symbolic Language Model, in: *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, AAAI Press, 2020, pp. 7634–7642. <https://aaai.org/ojs/index.php/AAAI/article/view/6264>.
- [11] D. Lyu, F. Yang, B. Liu and S. Gustafson, SDRL: Interpretable and Data-Efficient Deep Reinforcement Learning Leveraging Symbolic Planning, in: *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, AAAI Press, 2019, pp. 2970–2977. doi:10.1609/aaai.v33i01.33012970.
- [12] S. Cao, W. Lu and Q. Xu, Deep Neural Networks for Learning Graph Representations, in: *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, D. Schuurmans and M.P. Wellman, eds, AAAI Press, 2016, pp. 1145–1152. <http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12423>.
- [13] X. Chen, C. Liang, A.W. Yu, D. Song and D. Zhou, Compositional Generalization via Neural-Symbolic Stack Machines, in: *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan and H. Lin, eds, 2020. <https://proceedings.neurips.cc/paper/2020/hash/12b1e42dc0746f22cf361267de07073f-Abstract.html>.
- [14] Y. Xie, Z. Xu, K.S. Meel, M.S. Kankanhalli and H. Soh, Embedding Symbolic Knowledge into Deep Networks, in: *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, H.M. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E.B. Fox and R. Garnett, eds, 2019, pp. 4235–4245. <https://proceedings.neurips.cc/paper/2019/hash/7b66b4fd401a271a1c7224027ce111bc-Abstract.html>.
- [15] R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester and L.D. Raedt, DeepProbLog: Neural Probabilistic Logic Programming, in: *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, S. Bengio, H.M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi and R. Garnett, eds, 2018, pp. 3753–3763. <https://proceedings.neurips.cc/paper/2018/hash/dc5d637ed5e62c36ecb73b654b05ba2a-Abstract.html>.
- [16] F. Arabshahi, S. Singh and A. Anandkumar, Combining Symbolic Expressions and Black-box Function Evaluations in Neural Programs, in: *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, OpenReview.net, 2018. <https://openreview.net/forum?id=Hksj2WWAW>.
- [17] L. de Penning, A.S. d'Avila Garcez, L.C. Lamb and J.C. Meyer, A Neural-Symbolic Cognitive Agent for Online Learning and Reasoning, in: *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, T. Walsh, ed., IJCAI/AAAI, 2011, pp. 1653–1658. doi:10.5591/978-1-57735-516-8/IJCAI11-278.
- [18] J. Jiang and S. Ahn, Generative Neurosymbolic Machines, in: *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan and H. Lin, eds, 2020. <https://proceedings.neurips.cc/paper/2020/hash/94c28dcfc97557df0df6d1f7222fc384-Abstract.html>.
- [19] X. Liang, Z. Hu, H. Zhang, L. Lin and E.P. Xing, Symbolic Graph Reasoning Meets Convolutions, in: *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, S. Bengio, H.M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi and R. Garnett, eds, 2018, pp. 1858–1868. <https://proceedings.neurips.cc/paper/2018/hash/cbb6a3b884f4f88b3a8e3d44c636cbd8-Abstract.html>.
- [20] W.W. Cohen, H. Sun, R.A. Hofer and M. Siegler, Scalable Neural Methods for Reasoning With a Symbolic Knowledge Base, in: *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*, OpenReview.net, 2020. <https://openreview.net/forum?id=BJJguT4YPr>.
- [21] K.K. Teru, E. Denis and W. Hamilton, Inductive Relation Prediction by Subgraph Reasoning, in: *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, Proceedings of Machine Learning Research, Vol. 119, PMLR, 2020, pp. 9448–9457. <http://proceedings.mlr.press/v119/teru20a.html>.
- [22] S. Garg, A. Bajpai and Mausam, Symbolic Network: Generalized Neural Policies for Relational MDPs, in: *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, Proceedings of Machine Learning Research, Vol. 119, PMLR, 2020, pp. 3397–3407. <http://proceedings.mlr.press/v119/garg20a.html>.
- [23] L. Mou, Z. Lu, H. Li and Z. Jin, Coupling Distributed and Symbolic Execution for Natural Language Queries, in: *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, D. Precup and Y.W. Teh, eds, Proceedings of Machine Learning Research, Vol. 70, PMLR, 2017, pp. 2518–2526. <http://proceedings.mlr.press/v70/mou17a.html>.
- [24] C. Xiao, M. Dymetman and C. Gardent, Symbolic Priors for RNN-based Semantic Parsing, in: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, C. Sierra, ed., ijcai.org, 2017, pp. 4186–4192. doi:10.24963/ijcai.2017/585.

- [25] R. Dang-Nhu, PLANS: Neuro-Symbolic Program Learning from Videos, in: *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan and H. Lin, eds, 2020. <https://proceedings.neurips.cc/paper/2020/hash/fe131d7f5a6b38b23cc967316c13dae2-Abstract.html>.
- [26] K. Yi, J. Wu, C. Gan, A. Torralba, P. Kohli and J. Tenenbaum, Neural-Symbolic VQA: Disentangling Reasoning from Vision and Language Understanding, in: *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, S. Bengio, H.M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi and R. Garnett, eds, 2018, pp. 1039–1050. <https://proceedings.neurips.cc/paper/2018/hash/5e388103a391daabe3de1d76a6739ccd-Abstract.html>.
- [27] X. Chen, C. Liang, A.W. Yu, D. Zhou, D. Song and Q.V. Le, Neural Symbolic Reader: Scalable Integration of Distributed and Symbolic Representations for Reading Comprehension, in: *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*, OpenReview.net, 2020. <https://openreview.net/forum?id=ryxjnREFwH>.
- [28] E. Parisotto, A. Mohamed, R. Singh, L. Li, D. Zhou and P. Kohli, Neuro-Symbolic Program Synthesis, in: *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, OpenReview.net, 2017. <https://openreview.net/forum?id=rJ0JwFcex>.
- [29] P. Hohenecker and T. Lukasiewicz, Ontology Reasoning with Deep Neural Networks, *J. Artif. Intell. Res.* **68** (2020), 503–540. doi:10.1613/jair.1.11661.
- [30] J. Xu, Z. Zhang, T. Friedman, Y. Liang and G.V. den Broeck, A Semantic Loss Function for Deep Learning with Symbolic Knowledge, in: *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, J.G. Dy and A. Krause, eds, Proceedings of Machine Learning Research, Vol. 80, PMLR, 2018, pp. 5498–5507. <http://proceedings.mlr.press/v80/xu18h.html>.
- [31] A. Santoro, F. Hill, D.G.T. Barrett, A.S. Morcos and T.P. Lillicrap, Measuring abstract reasoning in neural networks, in: *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, J.G. Dy and A. Krause, eds, Proceedings of Machine Learning Research, Vol. 80, PMLR, 2018, pp. 4477–4486. <http://proceedings.mlr.press/v80/santoro18a.html>.
- [32] R. Trivedi, H. Dai, Y. Wang and L. Song, Know-Evolve: Deep Temporal Reasoning for Dynamic Knowledge Graphs, in: *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, D. Precup and Y.W. Teh, eds, Proceedings of Machine Learning Research, Vol. 70, PMLR, 2017, pp. 3462–3471. <http://proceedings.mlr.press/v70/trivedi17a.html>.
- [33] M. Allamanis, P. Chanthirasegaran, P. Kohli and C. Sutton, Learning Continuous Semantic Representations of Symbolic Expressions, in: *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, D. Precup and Y.W. Teh, eds, Proceedings of Machine Learning Research, Vol. 70, PMLR, 2017, pp. 80–88. <http://proceedings.mlr.press/v70/allamanis17a.html>.
- [34] M. Asai and C. Muise, Learning Neural-Symbolic Descriptive Planning Models via Cube-Space Priors: The Voyage Home (to STRIPS), in: *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, C. Bessiere, ed., ijcai.org, 2020, pp. 2676–2682. doi:10.24963/ijcai.2020/371.
- [35] I. Donadello, L. Serafini and A.S. d’Avila Garcez, Logic Tensor Networks for Semantic Image Interpretation, in: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, C. Sierra, ed., ijcai.org, 2017, pp. 1596–1602. doi:10.24963/ijcai.2017/221.
- [36] M. Asai and A. Fukunaga, Classical Planning in Deep Latent Space: Bridging the Subsymbolic-Symbolic Boundary, in: *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, S.A. McIlraith and K.Q. Weinberger, eds, AAAI Press, 2018, pp. 6094–6101. <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16302>.
- [37] M.D. Cranmer, A. Sanchez-Gonzalez, P.W. Battaglia, R. Xu, K. Cranmer, D.N. Spergel and S. Ho, Discovering Symbolic Models from Deep Learning with Inductive Biases, in: *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan and H. Lin, eds, 2020. <https://proceedings.neurips.cc/paper/2020/hash/c9f2f917078bd2db12f23c3b413d9cba-Abstract.html>.
- [38] A.M. Alaa and M. van der Schaar, Demystifying Black-box Models with Symbolic Metamodels, in: *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, H.M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E.B. Fox and R. Garnett, eds, 2019, pp. 11301–11311. <https://proceedings.neurips.cc/paper/2019/hash/567b8f5f423af15818a068235807edc0-Abstract.html>.
- [39] X. Zhang, A. Solar-Lezama and R. Singh, Interpreting Neural Network Judgments via Minimal, Stable, and Symbolic Corrections, in: *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, S. Bengio, H.M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi and R. Garnett, eds, 2018, pp. 4879–4890. <https://proceedings.neurips.cc/paper/2018/hash/300891a62162b960cf02ce3827bb363c-Abstract.html>.
- [40] B. Dhingra, M. Zaheer, V. Balachandran, G. Neubig, R. Salakhutdinov and W.W. Cohen, Differentiable Reasoning over a Virtual Knowledge Base, in: *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*, OpenReview.net, 2020. <https://openreview.net/forum?id=SJxstlHFPH>.
- [41] S. Amizadeh, H. Palangi, A. Polozov, Y. Huang and K. Koishida, Neuro-Symbolic Visual Reasoning: Disentangling "Visual" from "Reasoning", in: *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, Proceedings of Machine Learning Research, Vol. 119, PMLR, 2020, pp. 279–290. <http://proceedings.mlr.press/v119/amizadeh20a.html>.

- [42] Q. Li, S. Huang, Y. Hong, Y. Chen, Y.N. Wu and S. Zhu, Closed Loop Neural-Symbolic Learning via Integrating Neural Perception, Grammar Parsing, and Symbolic Reasoning, in: *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, Proceedings of Machine Learning Research, Vol. 119, PMLR, 2020, pp. 5884–5894. <http://proceedings.mlr.press/v119/li20f.html>.
- [43] D. Chen, Y. Bai, W. Zhao, S. Ament, J.M. Gregoire and C.P. Gomes, Deep Reasoning Networks for Unsupervised Pattern De-mixing with Constraint Reasoning, in: *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, Proceedings of Machine Learning Research, Vol. 119, PMLR, 2020, pp. 1500–1509. <http://proceedings.mlr.press/v119/chen20a.html>.
- [44] P. Minervini, S. Riedel, P. Stenetorp, E. Grefenstette and T. Rocktäschel, Learning Reasoning Strategies in End-to-End Differentiable Proving, in: *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, Proceedings of Machine Learning Research, Vol. 119, PMLR, 2020, pp. 6938–6949. <http://proceedings.mlr.press/v119/minervini20a.html>.
- [45] R. Vedantam, K. Desai, S. Lee, M. Rohrbach, D. Batra and D. Parikh, Probabilistic Neural Symbolic Models for Interpretable Visual Question Answering, in: *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, K. Chaudhuri and R. Salakhutdinov, eds, Proceedings of Machine Learning Research, Vol. 97, PMLR, 2019, pp. 6428–6437. <http://proceedings.mlr.press/v97/vedantam19a.html>.
- [46] H. Young, O. Bastani and M. Naik, Learning Neurosymbolic Generative Models via Program Synthesis, in: *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, K. Chaudhuri and R. Salakhutdinov, eds, Proceedings of Machine Learning Research, Vol. 97, PMLR, 2019, pp. 7144–7153. <http://proceedings.mlr.press/v97/young19a.html>.
- [47] P. Wang, P.L. Donti, B. Wilder and J.Z. Kolter, SATNet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver, in: *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, K. Chaudhuri and R. Salakhutdinov, eds, Proceedings of Machine Learning Research, Vol. 97, PMLR, 2019, pp. 6545–6554. <http://proceedings.mlr.press/v97/wang19e.html>.
- [48] M. Fischer, M. Balunovic, D. Drachslar-Cohen, T. Gehr, C. Zhang and M.T. Vechev, DL2: Training and Querying Neural Networks with Logic, in: *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, K. Chaudhuri and R. Salakhutdinov, eds, Proceedings of Machine Learning Research, Vol. 97, PMLR, 2019, pp. 1931–1941. <http://proceedings.mlr.press/v97/fischer19a.html>.
- [49] Z. Jiang and S. Luo, Neural Logic Reinforcement Learning, in: *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, K. Chaudhuri and R. Salakhutdinov, eds, Proceedings of Machine Learning Research, Vol. 97, PMLR, 2019, pp. 3110–3119. <http://proceedings.mlr.press/v97/jiang19a.html>.
- [50] C. Xiong, S. Merity and R. Socher, Dynamic Memory Networks for Visual and Textual Question Answering, in: *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, M. Balcan and K.Q. Weinberger, eds, JMLR Workshop and Conference Proceedings, Vol. 48, JMLR.org, 2016, pp. 2397–2406. <http://proceedings.mlr.press/v48/xiong16.html>.
- [51] L. Shastri, Advances in SHRUTI-A Neurally Motivated Model of Relational Knowledge Representation and Rapid Inference Using Temporal Synchrony, *Appl. Intell.* **11**(1) (1999), 79–108. doi:10.1023/A:1008380614985.
- [52] J. McCarthy, Epistemological challenges for connectionism, *Behavioral and Brain Sciences* **11** (1988), 44.
- [53] F. van Harmelen and A. ten Teije, A Boxology of Design Patterns for Hybrid Learning and Reasoning Systems, *J. Web Eng.* **18**(1–3) (2019), 97–124. doi:10.13052/jwe1540-9589.18133.
- [54] P. Hitzler, A review of the semantic web field, *Commun. ACM* **64**(2) (2021), 76–83. doi:10.1145/3397512.
- [55] P. Hitzler, M. Krötzsch and S. Rudolph, *Foundations of Semantic Web Technologies*, Chapman and Hall/CRC Press, 2010. ISBN 9781420090505. <http://www.semantic-web-book.org/>.
- [56] P. Hitzler, F. Bianchi, M. Ebrahimi and M.K. Sarker, Neural-symbolic integration and the Semantic Web, *Semantic Web* **11**(1) (2020), 3–11. doi:10.3233/SW-190368.
- [57] M. Ebrahimi, A. Eberhart, F. Bianchi and P. Hitzler, Towards Bridging the Neuro-Symbolic Gap: Deep Deductive Reasoners, *Applied Intelligence* (2021), to appear.
- [58] B. Makni and J.A. Hendler, Deep learning for noise-tolerant RDFS reasoning, *Semantic Web* **10**(5) (2019), 823–862. doi:10.3233/SW-190363.
- [59] F. Bianchi and P. Hitzler, On the Capabilities of Logic Tensor Networks for Deductive Reasoning, in: *Proceedings of the AAAI 2019 Spring Symposium on Combining Machine Learning with Knowledge Engineering (AAAI-MAKE 2019) Stanford University, Palo Alto, California, USA, March 25-27, 2019., Stanford University, Palo Alto, California, USA, March 25-27, 2019*, A. Martin, K. Hinkelmann, A. Gerber, D. Lenat, F. van Harmelen and P. Clark, eds, CEUR Workshop Proceedings, Vol. 2350, CEUR-WS.org, 2019. <http://ceur-ws.org/Vol-2350/paper22.pdf>.
- [60] M. Ebrahimi, M.K. Sarker, F. Bianchi, N. Xie, A. Eberhart, D. Doran, H. Kim and P. Hitzler, Neuro-Symbolic Deductive Reasoning for Cross-Knowledge Graph Entailment, in: *Proceedings of the AAAI 2021 Spring Symposium on Combining Machine Learning and Knowledge Engineering (AAAI-MAKE 2021), Stanford University, Palo Alto, California, USA, March 22-24, 2021*, A. Martin, K. Hinkelmann, H. Fill, A. Gerber, D. Lenat, R. Stolle and F. van Harmelen, eds, CEUR Workshop Proceedings, Vol. 2846, CEUR-WS.org, 2021. <http://ceur-ws.org/Vol-2846/paper8.pdf>.
- [61] A. Eberhart, M. Ebrahimi, L. Zhou, C. Shimizu and P. Hitzler, Completion Reasoning Emulation for the Description Logic EL+, in: *Proceedings of the AAAI 2020 Spring Symposium on Combining Machine Learning and Knowledge Engineering in Practice, AAAI-MAKE 2020, Palo Alto, CA, USA, March 23-25, 2020, Volume I*, A. Martin, K. Hinkelmann, H. Fill, A. Gerber, D. Lenat, R. Stolle and F. van Harmelen, eds, CEUR Workshop Proceedings, Vol. 2600, CEUR-WS.org, 2020. <http://ceur-ws.org/Vol-2600/paper5.pdf>.