

Neural-Symbolic Integration and the Semantic Web

A Position Paper

Pascal Hitzler ^{a,b}, Federico Bianchi ^c, Monireh Ebrahimi ^a, Md Kamruzzaman Sarker ^a

^a *Wright State University, Dayton, OH, USA*

^b *Kansas State University, Manhattan, KS, USA*

^c *University of Milan – Bicocca, Italy*

Approaches in Artificial Intelligence (AI) based on machine learning, and in particular those employing artificial neural networks, differ fundamentally from approaches that leverage knowledge bases to perform logical deduction and reasoning. The former are *connectionist* or *subsymbolic* AI systems able to solve complex tasks over unstructured data using supervised or unsupervised learning, including problems which cannot reasonably be hand-coded by humans. Subsymbolic methods are generally robust against noise in training or input data and have recently, in the wake of deep learning, been shown to exceed human performance in tasks involving video, audio, and text processing. The latter are *symbolic* systems that thrive under the presence of large amounts of structured data, including for agent planning, constraint solving, data management, integration and querying, and other traditional application areas of knowledge-based systems and formal semantics. Classical rule-based systems, ontologies, and knowledge graphs that power search and information retrieval across the Web are also types of symbolic AI systems.

Symbolic and subsymbolic systems are rather complementary to each other. For example, the key strengths of subsymbolic systems are weaknesses of symbolic ones, and vice versa. Symbolic systems are *brittle*, i.e., susceptible to data noise or minor flaws in the logical encoding of a problem, which stands in contrast to the robustness of connectionist approaches. But subsymbolic systems are generally *black boxes* in the sense that the systems cannot be inspected in ways that pro-

vide insight into their decisions (despite some recent progress on this in the wake of the *explainable AI* effort) while symbolic knowledge bases can in principle be inspected to interpret how a decision follows from input. Most importantly, symbolic and subsymbolic systems contrast in the types of problems and data they excel at. Scene recognition from images appears to be a problem which in general lies outside the capabilities of symbolic systems, for example, while complex planning scenarios appear to be outside the scope of current deep learning approaches.¹

On a more technical level, symbolic and subsymbolic systems differ fundamentally in how they represent data, information, or knowledge. Symbolic systems typically utilize structured representation languages, e.g. stemming from formal logic and the subfield of AI known as *knowledge representation and reasoning*. Trainable artificial neural networks, on the other hand, typically use representations based on high-dimensional Euclidean space, i.e. real-valued vectors, matrices, etc., and it is by no means obvious how reconciliations between these representation forms can be designed.²

The complementary nature of these methods has drawn a divide in the rich field of AI. The di-

¹The topic is being investigated, of course, with some recent progress being made. E.g., [1] report on an application of deep learning to planning, and explicitly frame it as work towards bridging the “subsymbolic-symbolic boundary.”

²It is possible to establish a formal, mathematical bridge in some cases, as e.g. laid out in [22], but so far with limited applicability [3].

vide is technical in nature, as symbol manipulation as captured by logical, deductive reasoning, which lies at the core of symbolic approaches, cannot be sufficiently performed using current subsymbolic systems. Moreover, the training to study subsymbolic systems (involving probability theory, statistics, linear algebra, and optimization) differs from symbolic systems (involving logic and propositional calculus, set and recursion theory, and advanced computability reasoning) so strongly that AI researchers tend to find a side of the divide based on their intellectual interests and background. The divide is also cultural in nature, one of mindsets and prior beliefs, that in the past could sometimes split the academic AI research community by provoking (heated) fundamental discussions. The divide is even geographical, where the European Union holds a much higher prevalence of researchers working on symbolic approaches than in the United States.

Neural-Symbolic Integration [2,4,19],³ as a field of research, addresses fundamental problems related to building a technical bridge between the symbolic and subsymbolic sides of the divide. The promises from a successful bridging of the technological divide are plenty. In the abstract, one could hope for best-of-both-world systems, which combine the transparency and reasoning-ability of symbolic systems with the robustness and learning-capabilities of subsymbolic ones. As such, integrated symbolic-subsymbolic systems may be able to address the knowledge acquisition bottleneck faced by symbolic systems, learn to perform advanced logical or symbolic reasoning tasks even in the presence of noisy or uncertain facts, and even yield self-explanatory subsymbolic models. This is the promise of added value of neural-symbolic integration research for Computer Science. And also more fundamentally, a bridging may shed insights into how natural (human) neural networks can perform symbolic tasks as witnessed by homo sapiens pursuing mathematics, formal logic, and other pursuits that we, introspectively, see as symbolic in nature; this is a basic research problem for Cognitive Science as a discipline.

In the following, we will first lay out, in more detail, promises and possible benefits of neural-

symbolic integration research for the Semantic Web. Then we will look at potential benefits of Semantic Web and neural-symbolic integration research for deep learning. Finally, we will also provide brief pointers to some current research going on in relation to this theme.

1. Benefits of Neural-Symbolic Integration for the Semantic Web

One of the issues that plague the Semantic Web (as well as many other fields in Computer Science and its applications) is the *knowledge acquisition bottleneck*. It refers to the difficult issue of encoding or otherwise storing knowledge, as structured information, for use in Computer Science applications. The manual encoding of such information, e.g. from human experts' knowledge, is a very slow and time-consuming, thus costly, process involving both topic experts and knowledge engineers. At the same time, automated methods are a far cry from producing artifacts (e.g., from textbooks, technical documentations, and other written sources) which would be of sufficient quality for use in intelligent systems applications based on logical inference, such as expert systems, or for data curation and integration.

The underpinnings of key Semantic Web standards, such as RDF [9] and OWL [20], are explicitly logical, which reflects that Semantic Web applications often rely on high data (and schema/ontology) quality, similar to knowledge bases used primarily for deductive reasoning. The knowledge acquisition bottleneck in the Semantic Web field is very noticeable, e.g., given that the creation of ontologies as well as the creation of high-quality knowledge graphs involves high amounts of human expert labour and is correspondingly expensive.

The promise of integrated neural-symbolic systems is that they would be capable of both learning and (deductive) reasoning, and thus that they would be able to acquire, through machine learning, knowledge which is of sufficiently high quality to perform deductive reasoning. This anticipated capability directly addresses the knowledge acquisition bottleneck. There is, thus, a promise in this line of work that integrated neural-symbolic systems will lead to

³See also <http://www.neural-symbolic.org/>

- better methods for automated ontology construction,
- better methods for ontology population (and, thus, knowledge graph construction),
- better methods for ontology alignment,
- better methods for assessing the quality of knowledge graph content,

and similar major lines of research central to the Semantic Web field.

At the same time, integrated neural-symbolic systems carry the promise of being able to perform deductive reasoning – after training – using a (highly parallel) artificial neural network architecture. Consequently, reasoning using such systems can be expected to be extremely fast. This contrasts with traditional deductive reasoning methods, which are usually designed to be provably sound and complete but suffer from long algorithm runtimes. While there has been significant progress on developing highly efficient deductive reasoning engines for Semantic Web content, this remains an issue given ever-increasing availability of data. In fact, the underlying problem is fundamental, as sound and complete reasoning over Semantic Web data necessarily suffers from high computational complexity [21].

Integrated neural-symbolic systems would perform reasoning after training, and presumably this form of reasoning would not be provably sound and complete, but would trade correctness guarantees with higher runtime efficiency, in the spirit of *approximate reasoning* – see e.g. [23] for an exhibition of the underlying rationale. As such, integrated neural-symbolic systems carry a promise to elevate deductive Semantic Web reasoning to much larger amounts of data.

With integrated neural-symbolic systems capable of approximate deductive reasoning, this would furthermore open up possible investigations into combining deductive and inductive reasoning, as well as common-sense reasoning based e.g. on natural language, within a single (artificial neural network based) system.

Side products of such approaches would also be, e.g., entity encodings in formats suitable for artificial neural networks, such as vector or matrix representations. These in turn could be utilized to assess entity similarity with potential applications in data integration. Such encodings could furthermore be used as a sort of compression for data transfer and storage.

2. Benefits of Semantic Web Technologies and Neural-Symbolic Integration for Deep Learning

Semantic Web Technologies are designed for enabling better and more efficient data sharing, discovery, integration and reuse. These data management core capabilities of Semantic Web Technologies are designed to ease the data curation and preparation burden for the training of deep learning systems. Semantic Web data, provided in large amounts and freely available on the Web [38], furthermore provides a rich resource for training data, and deductive reasoning methods over such data can further extend it.

Integrated neural-symbolic systems will furthermore make it possible to utilize background knowledge, given as knowledge graphs or ontologies, as part of deep learning applications. Promises of this include the leveraging of background knowledge and deductive reasoning aspects for improved trainability, but also for interpreting trained deep learning systems by means of background knowledge. The former aspect attempts to reinforce the usefulness of deep learning models through injection of knowledge and has been successfully used in task-oriented conversational AI systems [15] and question answering [32]. The latter aspect touches on the *Explainable AI* theme currently being discussed, which aims at addressing the black-box nature of deep learning systems by making them more transparent, understandable, verifiable, and trustworthy. Most of the current work on this topic attempts to explain system behavior by means of input or output features; however explanations by way of background knowledge carry the promise of being much closer to human conceptualizations, and thus more useful in applications.

Integrated neural-symbolic systems which incorporate deductive reasoning capabilities could furthermore naturally combine these with inferences based on statistics or similarities, including natural-language common-sense reasoning as demonstrated by some deep learning approaches. Such combinations should naturally lead to stronger deep learning systems.

Neural-symbolic systems have already been used on linked datasets like Freebase and DBpedia for different tasks like link prediction [49] and noise tolerant RDFS reasoning [29]. The links between linked datasets could further allow neural-

symbolic systems to both integrate and reason over information coming from different sources. The advantage of this is twofold: firstly, the combined information can be used to extend the amount of training data for neural-symbolic systems; secondly, a neural-symbolic system can be used to learn to reason over a single knowledge graph and then links can be used as entry points to reason over a different one. This could be useful in contexts in which it is costly to learn to reason over a large dataset; one could thus use neural-symbolic methods over a smaller one (or a part of the large one) and then use the learned capabilities over the large one.

3. Selection of Recent Related Work

Deductive Semantic Web Reasoning using Deep Learning

Deductive reasoning over RDF(S) and OWL data has become a part of the standard toolbox for knowledge graphs, and the use of neural-symbolic systems for this purpose has begun to be investigated.

[29] has proposed a noise-tolerant algorithm for deep-learning-based reasoning designed specifically for RDF(S) knowledge graphs. They have introduced a layered graph model representation of RDF graphs based on their predicates, in the form of 3D adjacency matrices where each layer layout forms a graph word. Each input graph and its corresponding entailments then have been represented as a sequence of graph words and have been fed to a neural machine translation model. Their results show noise-tolerant capabilities of their deep model, compared to their symbolic counterpart. However evaluation and training are done on a dataset that uses only one ontology for the inference, i.e., there is no learning of the general logical deduction calculus, and consequently no transfer thereof to new data.

[24] applies Recursive Reasoning Networks (RNN) to OWL RL reasoning where recursive update layers are used to update the individual embeddings using the relations and class memberships in the knowledge base. Their results show the potential of neural-symbolic methods to at-

tain accuracy similar to symbolic methods. However, as for the above mentioned [29], re-training is required for new ontologies to learn the embeddings for the new vocabularies in the ontology, i.e., the approach does not natively support transfer to new data.

[13] addresses the transferability issue by adapting end-to-end memory networks for emulating deductive RDFS reasoning. Transfer was achieved primarily by utilizing a preprocessing step consisting of a normalization. It was demonstrated that the resulting approach can perform reasoning over previously unseen RDFS knowledge graphs.

Knowledge Graph Embeddings

With the recent revival of interest in artificial neural networks, neural link prediction models have been applied extensively for the completion of knowledge graphs, understood in the sense of link-prediction.⁴ These methods [7,53,27,35,37,49,10,56,50,51] heavily rely on the subsymbolic representations, called embeddings, of entities and relations learned through maximization of a scoring objective function over valid factual triples. Thus, the current success of such deep models hinges primarily on the power of those subsymbolic continuous real-valued representations in encoding the similarity/relatedness of entities and relations. For example, TransE [7] for a knowledge base triple s, p, o learns an embedding function by minimizing the distance based on the respective vector representations $d(\mathbf{s} + \mathbf{p}, \mathbf{o})$ where d is a distance function. These methods have been often tested over link prediction tasks.

The use of additional information, such as text, can increase the quality of the representation [52,55,54]. Moreover, embedded representations of knowledge graphs can be extended by considering the logical axioms that appear in a knowledge base, for example, complex logical formulas can be aggregated using fuzzy logic [18].

A recent trend in knowledge graph embedding concerns approaches that use hyperbolic geometry in place of euclidean geometry [34,41]; hyperbolic

⁴Traditionally, "completion" in the context of RDF(S) referred only to materialization of logical consequences; more recently, the term has also been used to refer to the adding of new relationships (graph edges) based on statistical or NLP methods.

geometry generally appears to be more suited to represent hierarchical structures like terminologies and ontologies.

RDF2Vec [39] embeds RDF-based entities and relationships in a vector space by applying word embedding techniques [31] over a virtual document that contains lexicalized graph walks; thus the generated representations are based on token-token co-occurrences.

Explainable Deep Learning

While deep learning is highly successful [25] and even surpasses human capabilities [47,46] in many fields, but it also lacks transparency or interpretability [17,28] of how a decision is being produced from these systems. In safety-critical applications, e.g., in medical, legal or military contexts, this is deemed insufficient. Consequently, researchers are investigating how to produce explanations for the behavior of deep learning systems [16].

Explanations [58,57,42] produced from the deep learning systems are helpful to understand how it is producing the output, and the additional use of domain information helps to enhance [12] the explanation. [36] used an ontology-based deep learning model which predicts human behaviour via Restricted Boltzman Machines [48] and produces explanations of the output using domain ontologies.

[6] shows the use of semantic annotations to label objects in the hidden layers of popular CNN architectures. Labels ranging from colors, materials, textures, parts, objects and scenes help to get a better understanding of hidden parts of the deep network. Although the labels are not semantically structured, this shows that background knowledge can help to improve explainability.

[44] provides a feasibility study on how domain ontologies together with description logic based concept induction [26,43] can be used to explain input output behaviour of trained deep neural networks.

Other Systems for Deductive Reasoning using Deep Learning

The Neural Theorem Prover (NTP) [40] is an extension of the Prolog language in which strict atom unification is replaced with similarity of

atoms in an embedded space; while originally NTP suffered from scalability issues, due to the complexity of the approach, there is evidence that proof-path selection strategies can reduce the complexity impact [33].

DeepProbLog [30] is a programming language that combines a probabilistic logic with neural networks, thus offering a framework that combines the strengths from both approaches.

Logic Tensor Networks [45] (LTNs) combine deep neural networks and first order fuzzy logic. Elements of the logic language are embedded in a vector space (e.g., constants are represented as vectors while predicates are neural tensor networks [49] that have been used on simple reasoning tasks [49,8]). LTNs can be trained over both facts and rules and after training they can be used to make novel logical inferences over data. LTNs have been applied to semantic image interpretation tasks [11] but they have also been shown to have some computational limitations [5].

PossibleWorldNet [14] is a variant of Tree Neural Networks (TreeNN) which has been successfully used for conducting entailment over propositional logic formulas. To evaluate whether A entails B, the PossibleWorldNet generates a set of “possible worlds”, and then evaluates A and B in each of those worlds. Their results show the clear advantage of using this model compared to sequence-to-sequence models which would capture the structure implicitly.

4. Conclusion

In the wake of deep learning, neural-symbolic approaches are receiving renewed attention. We have laid out the promises of neural-symbolic integration research for the Semantic Web field, and vice versa. It appears to be reasonable to expect that the corresponding lines of research will receive growing attention in forthcoming years.

References

- [1] M. Asai and A. Fukunaga. Classical planning in deep latent space: Bridging the subsymbolic-symbolic boundary. In S. A. McIlraith and K. Q. Weinberger, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, New Orleans, Louisiana, USA, February 2-7, 2018*. AAAI Press, 2018.

- [2] S. Bader and P. Hitzler. Dimensions of neural-symbolic integration – A structured survey. In S. N. Artëmov, H. Barringer, A. S. d’Avila Garcez, L. C. Lamb, and J. Woods, editors, *We Will Show Them! Essays in Honour of Dov Gabbay, Volume One*, pages 167–194. College Publications, 2005.
- [3] S. Bader, P. Hitzler, S. Hölldobler, and A. Witzel. A fully connectionist model generator for covered first-order logic programs. In M. M. Veloso, editor, *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007*, pages 666–671, 2007.
- [4] 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.
- [5] F. Bianchi and P. Hitzler. On the capabilities of logic tensor networks for deductive reasoning. In *AAAI Spring Symposium MAKE*, 2019.
- [6] A. O. Bolei Zhou, David Bau and A. Torralba. Interpreting deep visual representations via network dissection. *arXiv preprint arXiv:1711.05611*, 2018.
- [7] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko. Translating embeddings for modeling multi-relational data. In *Advances in neural information processing systems*, pages 2787–2795, 2013.
- [8] S. R. Bowman, C. Potts, and C. D. Manning. Learning distributed word representations for natural logic reasoning. In *2015 AAAI Spring Symposium Series*, 2015.
- [9] D. Brickley and R. Guha, editors. *RDF Schema 1.1*. W3C Recommendation 25 February 2014, 2010. Available from <http://www.w3.org/TR/rdf-schema>.
- [10] K.-W. Chang, W.-t. Yih, B. Yang, and C. Meek. Typed tensor decomposition of knowledge bases for relation extraction. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1568–1579, 2014.
- [11] I. Donadello, L. Serafini, and A. d’Avila Garcez. Logic tensor networks for semantic image interpretation. In *IJCAI International Joint Conference on Artificial Intelligence*, pages 1596–1602, 2017.
- [12] D. Doran, S. Schulz, and T. R. Besold. What does explainable ai really mean? a new conceptualization of perspectives. *arXiv preprint arXiv:1710.00794*, 2017.
- [13] M. Ebrahimi, M. K. Sarker, F. Bianchi, N. Xie, D. Doran, and P. Hitzler. Reasoning over knowledge bases using deep learning. *arXiv preprint arXiv:1811.04132*, 2018.
- [14] R. Evans, D. Saxton, D. Amos, P. Kohli, and E. Grefenstette. Can neural networks understand logical entailment? *arXiv preprint arXiv:1802.08535*, 2018.
- [15] M. Ghazvininejad, C. Brockett, M.-W. Chang, B. Dolan, J. Gao, W.-t. Yih, and M. Galley. A knowledge-grounded neural conversation model. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [16] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi. A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, 51(5):93, 2018.
- [17] D. Gunning. Explainable artificial intelligence (XAI). *Defense Advanced Research Projects Agency (DARPA)*, 2017.
- [18] S. Guo, Q. Wang, L. Wang, B. Wang, and L. Guo. Jointly embedding knowledge graphs and logical rules. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 192–202, 2016.
- [19] B. Hammer and P. Hitzler, editors. *Perspectives of Neural-Symbolic Integration*, volume 77 of *Studies in Computational Intelligence*. Springer, 2007.
- [20] P. Hitzler, M. Krötzsch, B. Parsia, P. F. Patel-Schneider, and S. Rudolph, editors. *OWL 2 Web Ontology Language: Primer (Second Edition)*. W3C Recommendation 11 December 2012, 2012. Available at <http://www.w3.org/TR/owl2-primer>.
- [21] P. Hitzler, M. Krötzsch, and S. Rudolph. *Foundations of Semantic Web Technologies*. Chapman & Hall/CRC, 2010.
- [22] P. Hitzler and A. K. Seda. A note on the relationships between logic programs and neural networks. In D. Sinclair and J. P. Gibson, editors, *4th Irish Workshop on Formal Methods, IWFMM 2000, Maynooth, Ireland, 5th-6th July 2000*, Workshops in Computing. BCS, 2000.
- [23] P. Hitzler and F. van Harmelen. A reasonable semantic web. *Semantic Web*, 1(1):39–44, 2010.
- [24] P. Hohenecker and T. Lukasiewicz. Deep learning for ontology reasoning. *arXiv preprint arXiv:1705.10342*, 2017.
- [25] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436, 2015.
- [26] J. Lehmann and P. Hitzler. Concept learning in description logics using refinement operators. *Machine Learning*, 78(1-2):203–250, 2010.
- [27] Y. Lin, Z. Liu, M. Sun, Y. Liu, and X. Zhu. Learning entity and relation embeddings for knowledge graph completion. In *AAAI*, pages 2181–2187, 2015.
- [28] Z. C. Lipton. The myths of model interpretability. *arXiv preprint arXiv:1606.03490*, 2016.
- [29] B. Makni and J. Hendler. *Deep Learning for Noise-tolerant RDFS Reasoning*. PhD thesis, Rensselaer Polytechnic Institute, 2018.
- [30] R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester, and L. De Raedt. Deepproblog: Neural probabilistic logic programming. In *Advances in Neural Information Processing Systems*, pages 3749–3759, 2018.
- [31] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.
- [32] A. Miller, A. Fisch, J. Dodge, A.-H. Karimi, A. Bordes, and J. Weston. Key-value memory networks for directly reading documents. *arXiv preprint arXiv:1606.03126*, 2016.
- [33] P. Minervini, M. Bosnjak, T. Rocktäschel, E. Grefenstette, and S. Riedel. Scalable neural theorem proving

- on knowledge bases and natural language. 2018.
- [34] M. Nickel and D. Kiela. Poincaré embeddings for learning hierarchical representations. In *Advances in neural information processing systems*, pages 6338–6347, 2017.
- [35] M. Nickel, V. Tresp, and H.-P. Kriegel. Factorizing YAGO: scalable machine learning for linked data. In *Proceedings of the 21st international conference on World Wide Web*, pages 271–280. ACM, 2012.
- [36] N. Phan, D. Dou, H. Wang, D. Kil, and B. Piniowski. Ontology-based deep learning for human behavior prediction with explanations in health social networks. *Information Sciences*, 384:298–313, 2017.
- [37] S. Riedel, L. Yao, A. McCallum, and B. M. Marlin. Relation extraction with matrix factorization and universal schemas. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 74–84, 2013.
- [38] L. Rietveld, W. Beek, R. Hoekstra, and S. Schlobach. Meta-data for a lot of LOD. *Semantic Web*, 8(6):1067–1080, 2017. doi:10.3233/SW-170256.
- [39] P. Ristoski and H. Paulheim. RDF2vec: RDF graph embeddings for data mining. In *International Semantic Web Conference*, pages 498–514. Springer, 2016.
- [40] T. Rocktäschel and S. Riedel. End-to-end differentiable proving. In *Advances in Neural Information Processing Systems*, pages 3788–3800, 2017.
- [41] F. Sala, C. De Sa, A. Gu, and C. Re. Representation tradeoffs for hyperbolic embeddings. In *International Conference on Machine Learning*, pages 4457–4466, 2018.
- [42] W. Samek, T. Wiegand, and K.-R. Müller. Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. *arXiv preprint arXiv:1708.08296*, 2017.
- [43] M. K. Sarker and P. Hitzler. Efficient concept induction for description logics. In *Proceedings of AAAI-19, the Thirty-Third AAAI Conference on Artificial Intelligence, January 27 – February 1, 2019, Honolulu, Hawaii, USA*, 2019.
- [44] M. K. Sarker, N. Xie, D. Doran, M. Raymer, and P. Hitzler. Explaining trained neural networks with semantic web technologies: First steps. In T. R. Besold, A. S. d’Avila Garcez, and I. Noble, editors, *Proceedings of the Twelfth International Workshop on Neural-Symbolic Learning and Reasoning, NeSy 2017, London, UK, July 17-18, 2017.*, volume 2003 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2017.
- [45] L. Serafini and A. S. d. Garcez. Learning and reasoning with logic tensor networks. In *Conference of the Italian Association for Artificial Intelligence*, pages 334–348. Springer, 2016.
- [46] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, et al. Mastering Chess and Shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815*, 2017.
- [47] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354, 2017.
- [48] P. Smolensky. Information processing in dynamical systems: Foundations of harmony theory. Technical report, Colorado Univ at Boulder Dept of Computer Science, 1986.
- [49] R. Socher, D. Chen, C. D. Manning, and A. Ng. Reasoning with neural tensor networks for knowledge base completion. In *Advances in neural information processing systems*, pages 926–934, 2013.
- [50] K. Toutanova, D. Chen, P. Pantel, H. Poon, P. Choudhury, and M. Gamon. Representing text for joint embedding of text and knowledge bases. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1499–1509, 2015.
- [51] T. Trouillon, J. Welbl, S. Riedel, É. Gaussier, and G. Bouchard. Complex embeddings for simple link prediction. In *International Conference on Machine Learning*, pages 2071–2080, 2016.
- [52] Z. Wang and J.-Z. Li. Text-enhanced representation learning for knowledge graph. In *IJCAI*, pages 1293–1299, 2016.
- [53] Z. Wang, J. Zhang, J. Feng, and Z. Chen. Knowledge graph embedding by translating on hyperplanes. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.
- [54] H. Xiao, M. Huang, L. Meng, and X. Zhu. SSP: semantic space projection for knowledge graph embedding with text descriptions. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [55] R. Xie, Z. Liu, J. Jia, H. Luan, and M. Sun. Representation learning of knowledge graphs with entity descriptions. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [56] B. Yang, W.-t. Yih, X. He, J. Gao, and L. Deng. Embedding entities and relations for learning and inference in knowledge bases. *arXiv preprint arXiv:1412.6575*, 2014.
- [57] J. Yosinski, J. Clune, A. Nguyen, T. Fuchs, and H. Lipson. Understanding neural networks through deep visualization. *arXiv preprint arXiv:1506.06579*, 2015.
- [58] Q.-s. Zhang and S.-C. Zhu. Visual interpretability for deep learning: a survey. *Frontiers of Information Technology & Electronic Engineering*, 19(1):27–39, 2018.