Supplement to Chapter 6: LNN: Logical Neural Networks

Supplement to Chapter 6 from Neuro Symbolic Reasoning and Learning - Current Advances and Future Directions

Additional Discussion

Interpreting Parameters. Note that, unlike PSL/NeurPSL, which assigns weights to rules [13] (which can be thought of as a single weight for an operator), LNN's assign weights to the inputs of an operator. These inputs are associated with the lower-level logical constructs (i.e., atoms for subformulas). As a result, learned weights can be viewed as a level of importance for those components. The weights are not required to sum to 1, so they should be viewed in terms of "absolute" importance, as opposed to relative importance. The bias (β), which is associated with each logical construct, is somewhat less interpretable. While the designers of LNN recommend a common setting for the bias (e.g., $\beta = 1$ as mentioned earlier), they also view it as the "difficulty" involved in satisfying a given logical formula.

In the context of neuro symbolic reasoning, the concept of weighting arguments in a logical construct is singular to LNN's. However, the novelty in the concept with respect to neuro symbolic reasoning is not the weighting itself, but the fact that the weights are learned through gradient descent. The idea had appeared earlier primarily in the context of developing aggregators for use in database queries (e.g., [3]).

Perhaps the greater question to raise is that is such a parameterization scheme viable as we seek to scale neuro symbolic methods? Questions such as should there be shared weights for various constructs, are parameterized operators leading to overfitting, and the relationship between weights for pure neural layers (e.g., lower level perceptual layers) and LNN layers will provide many topics for future research in the area of LNN's.

Interpreting Fuzzy Outputs. With any fuzzy/real-valued logic framework (to include neuro symbolic approaches), interpreting the values associated with the atoms is an important issue. However, in the simplest case, a thresholding of the resulting confidence value can be determined with a typical machine learning validation approach. In other words, the target predicates (and associated ground atoms) can be treated as target classes and identifying a level of confidence (fuzzy value threshold) that provides the desired trade-off between precision and recall is a reasonable approach in many cases. Further, this allows for a direct comparison with standard black-box machine learning approaches.

However, the introduction of the α hyper-parameter in LNN's and the associated interpretation of the real values related to truth or falsehood (e.g., see Figure **??** add an additional layer of complexity. There are three reasons for this. First, the model is returning intervals rather than scalars. Second, there are pre-defined areas of truth and falsehood, and the resulting intervals may be subsumed by one of these levels and/or intersect with one or both levels. Understanding best practices around setting α prior to training as well as establishing proper thresholds for decision making in validation is important in the application of LNN's to real-world systems. This is especially in terms of modularity - as it will be important to understand the semantics of the real-valued intervals provided for target atoms when integrating into other AI platforms.

References

1. Badreddine, S., d'Avila Garcez, A., Serafini, L., Spranger, M.: Logic tensor networks. Artificial Intelligence **303**, 103649 (2022). DOI https://doi.org/10.1016/j.artint.2021.103649. URL https://www.sciencedirect.com/science/article/pii/S0004370221002009
2. Evans, R., Grefenstette, E.: Learning explanatory rules from noisy data. J. Artif. Int. Res. **61**(1), 1–64 (2018)
3. Fagin, R., Wimmers, E.L.: A formula for incorporating weights into scoring rules. Theoretical Computer Science **239**(2), 309–338 (2000). DOI https://doi.org/10.1016/S0304-3975(99)00224-8.
4. Harnad, S.: The symbol grounding problem. Physica D: Nonlinear Phenomena **42**(1-3), 335–346 (1990)
5. Hong, J., Pavlic, T.: Representing prior knowledge using randomly, weighted feature networks for visual relationship detection. In: Combining Learning and Reasoning: Programming Languages, Formalisms, and Representations (2022). URL https://openreview.net/forum?id=iwoNpozn10l
6. Jiang, H., Gurajada, S., Lu, Q., Neelam, S., Popa, L., Sen, P., Li, Y., Gray, A.: LNN-EL: A neuro-symbolic approach to short-text entity linking. In: Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pp. 775–787. Association for Computational Linguistics, Online (2021). DOI 10.18653/v1/2021.acl-long.64. URL https://aclanthology.org/2021.acl-long.64
7. Kifer, M., Subrahmanian, V.: Theory of generalized annotated logic programming and its applications. J. Log. Program. **12**(3&4), 335–367 (1992)

8. Kimura, D., Chaudhury, S., Wachi, A., Kohita, R., Munawar, A., Tatsubori, M., Gray, A.: Reinforcement learning with external knowledge by using logical neural networks. CoRR **abs/2103.02363** (2021). URL https://arxiv.org/abs/2103.02363
9. Kimura, D., Ono, M., Chaudhury, S., Kohita, R., Wachi, A., Agravante, D.J., Tatsubori, M., Munawar, A., Gray, A.: Neuro-symbolic reinforcement learning with first-order logic. In: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, pp. 3505–3511. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic (2021). DOI 10.18653/v1/2021.emnlp-main.283. URL https://aclanthology.org/2021.emnlp-main.283
10. Krizhevsky, A.: Convolutional deep belief networks on CIFAR-10 (2010)
11. Luus, F.P.S., Sen, P., Kapanipathi, P., Riegel, R., Makondo, N., Lebese, T., Gray, A.G.: Logic embeddings for complex query answering. CoRR **abs/2103.00418** (2021). URL https://arxiv.org/abs/2103.00418
12. Ma, M., Gao, J., Feng, L., Stankovic, J.A.: Stlnet: Signal temporal logic enforced multivariate recurrent neural networks. 34th Conference on Neural Information Processing Systems (NeurIPS 2020) URL https://par.nsf.gov/biblio/10231392
13. Pryor, C., Dickens, C., Augustine, E., Albalak, A., Wang, W., Getoor, L.: NeuPSL: Neural probabilistic soft logic. arXiv preprint arXiv:2205.14268 (2022)
14. Riegel, R., Gray, A., Luus, F., Khan, N., Makondo, N., Akhalwaya, I.Y., Qian, H., Fagin, R., Barahona, F., Sharma, U., Ikbal, S., Karanam, H., Neelam, S., Likhyani, A., Srivastava, S.: Logical neural networks (2020). DOI 10.48550/ARXIV.2006.13155. URL https://arxiv.org/abs/2006.13155 block arXiv preprint arXiv:2006.13155 (2020)
15. Sen, P., Carvalho, B.W.S.R.d., Riegel, R., Gray, A.: Neuro-symbolic inductive logic programming with logical neural networks. Proceedings of the AAAI Conference on Artificial Intelligence **36**(8), 8212–8219 (2022). DOI 10.1609/aaai.v36i8.20795. URL https://ojs.aaai.org/index.php/AAAI/article/view/20795
16. Shakarian, P., Simari, G.: Extensions to generalized annotated logic and an equivalent neural architecture. In: IEEE TransAI. IEEE (2022)
17. van Krieken, E., Acar, E., van Harmelen, F.: Analyzing differentiable fuzzy logic operators. Artificial Intelligence **302**, 103602 (2022). DOI https://doi.org/10.1016/j.artint.2021.103602. URL https://www.sciencedirect.com/science/article/pii/S0004370221001533
18. Wang, P., Donti, P.L., Wilder, B., Kolter, J.Z.: Satnet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver. In: K. Chaudhuri, R. Salakhutdinov (eds.) Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA, *Proceedings of Machine Learning Research*, vol. 97, pp. 6545–6554. PMLR (2019). URL http://proceedings.mlr.press/v97/wang19e.html
19. Yang, Z., Ishay, A., Lee, J.: Neurasp: Embracing neural networks into answer set programming. In: C. Bessiere (ed.) Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20, pp. 1755–1762. International Joint Conferences on Artificial Intelligence Organization (2020). Main track