

Semantic Probabilistic Layers for Neuro-Symbolic Learning

Kareem Ahmed^{1,*}, Stefano Teso², Kai-Wei Chang¹, Guy Van den Broeck¹ and Antonio Vergari³

¹CS Department, UCLA

²CIMeC & DISI, University of Trento

³School of Informatics, University of Edinburgh

Abstract

In this extended abstract, we briefly outline Semantic Probabilistic Layers [1], a new layer that can be plugged into any neural network to guarantee its predictions are consistent with a set of predefined symbolic constraints while being amenable to end-to-end learning via maximum likelihood. SPLs can faithfully, and efficiently, model complex SOP tasks beyond the reach of alternative neuro-symbolic layers. We empirically demonstrate that SPLs outperform these competitors in terms of accuracy on an array of challenging structured-output prediction tasks.

Keywords

Neuro-symbolic Integration, Hard Constraints, Neural Layers, Probabilistic Logics, Trustworthy AI

There is a growing need for trustworthy neural architectures that output predictions consistent with regulation mandated constraints or other requirements. We have recently introduced Semantic Probabilistic Layers [1], a new layer that can be plugged into *any* neural network to guarantee its predictions are consistent with a set of predefined symbolic constraints. To this end, SPLs combine exact probabilistic inference with logical reasoning in a clean and modular way, learning complex distributions and restricting their support to solutions of the constraint. At the same time, SPLs can model intricate correlations, and hard constraints, over the labels, all while being amenable to end-to-end learning via maximum likelihood. As such, they can faithfully, and efficiently, model complex SOP tasks beyond the reach of alternative neuro-symbolic layers. In the published paper [1], we empirically demonstrated that SPLs outperform a number of competitors in terms of accuracy on challenging tasks including hierarchical multi-label classification, pathfinding and preference learning, while retaining perfect constraint satisfaction. The code is available at: <https://github.com/KareemYousrii/SPL>.

References

- [1] K. Ahmed, S. Teso, K.-W. Chang, G. Van den Broeck, A. Vergari, Semantic probabilistic layers for neuro-symbolic learning, *Advances in Neural Information Processing Systems* 35 (2022) 29944–29959.


The 17th International Workshop on Neural-Symbolic Learning and Reasoning, July 3–5, 2023, Siena, Italy

*Corresponding author.

✉ ahmedk@cs.ucla.edu (K. Ahmed)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)