Generalizable Neuro-Symbolic Systems for Commonsense Question Answering

Alessandro Oltramari[1,2,3,*], Jonathan Francis[1,2], Filip Ilievski[4], Kaixin Ma[2] and Roshanak Mirzaee[5]

[1] *Bosch Research and Technology Center, Pittsburgh, USA*

[2] *Carnegie Mellon University, Pittsburgh, USA*

[3] *Italian National Research Council, Institute of Cognitive Science and Technology, Trento, Italy*

[4] *Information Sciences Institute, Viterbi School of Engineering, University of Southern California, Marina del Rey, USA*

[5] *Department of Computer Science and Engineering, Michigan State University, East Lansing, USA*

1. Extended abstract

In this contribution [1], we analyze different methods for integrating large language models (LLMs) and knowledge graphs (KGs), presenting a quantitative and qualitative evaluation of these neuro-symbolic strategies on a broad spectrum of question answering benchmark datasets focused on common sense reasoning.

We illustrate how the combination of data-driven algorithms and symbolic knowledge is necessary to enable domain generalizability and robustness in downstream tasks. Our work, despite being prior to the disruptive innovation triggered by ChatGPT [2], remains actual, as it stems from a comprehensive analysis of the intrinsic problems that LLMs have in dealing with high-level forms of inferences requiring spatial, temporal, analogical, causal reasoning. We argue that such problems are dependent on the inner limitations of the neural architectures behind those models and cannot simply be solved by increasing the size of the training data, or by further scaling the number of parameters.

In our original paper we claim that two orthogonal extensions are required for LLMs: first, augmenting them with symbolic knowledge, which is available in large amounts thanks to open-source commonsense KGs; second, extending the resulting neuro-symbolic language models with components that can dynamically perform reasoning. We refer to the former direction as horizontal augmentation, as it focuses on expanding the coverage of the structured knowledge available to LLMs methods, and to the latter as vertical augmentation, as it explores in depth how specific forms of reasoning can be enabled. Our paper mostly focuses on horizontal augmentation, but it also lays the foundations of vertical augmentation. In regards of the former, we study techniques like attention-based knowledge injection, and fine-tuning LLMs with synthetic question-answer pairs generated from KGs: in both techniques, symbolic knowledge is typically translated into latent vectorial representations. However, if sub-symbolic expressions can augment training signals with features transformed from explicit semantic content, such knowledge infusion process doesn't mandate how inferences on the learned knowledge should be performed. Such mechanisms are typically executed using logic-based reasoning, e.g., Region-Connection-Calculus for spatial knowledge [3], Allen's axioms for temporal (or spatial) knowledge [4], etc. (relevant work exists showing how deep learning models can replicate logical reasoning, e.g., [5, 6], but it doesn't follow that any form of reasoning should be reduced to sub-symbolic learning – or at least this is an assumption only for some closely paired neuro-symbolic systems). Implementing inference mechanisms1, and integrating them with large-scale knowledge infusion methods, is what we advocate for in the future of neuro-symbolic AI research and development.

CEUR Workshop Proceedings (CEUR-WS.org)

References

[1] Alessandro Oltramari, Jonathan Francis, Filip Ilievski, Kaixin Ma, Roshanak Mirzaee: Generalizable Neuro-symbolic Systems for Commonsense Question Answering. In Pascal Hitzler, Md Kamruzzaman Sarker (eds.), Neuro-Symbolic Artificial Intelligence: The State of the Art. Frontiers in Artificial Intelligence and Applications Vol. 342, IOS Press, Amsterdam, 2022.

[2] "ChatGPT – Release Notes". Archived from the original on May 4, 2023. Retrieved May 4, 2023.

[3] A. G. Cohn, B. Bennett, J. Gooday, N. M. Gotts, Qualitative spatial representation and reasoning with the region connection calculus, geoinformatica 1 (1997) 275–316.

[4] J. F. Allen, G. Ferguson, Actions and events in interval temporal logic, Journal of logic and computation 4 (1994) 531–579.

[5] M. Ebrahimi, A. Eberhart, P. Hitzler, On the capabilities of pointer networks for deep deductive reasoning, arXiv preprint arXiv:2106.09225 (2021).

[6] A. d. Garcez, S. Bader, H. Bowman, L. C. Lamb, L. de Penning, B. Illuminoo, H. Poon, C. Gerson Zaverucha, Neural-symbolic learning and reasoning: A survey and interpretation, Neuro-Symbolic Artificial Intelligence: The State of the Art 342 (2022) 1.