

Neuromyotonic Fusion: Integrating Symbolic Reasoning with Deep Learning Architectures

Pradipta Gorai¹, Sangita Bose²

¹Department of MCA, Swami Vivekananda University-
700121, West Bengal, India
pradiptagorai1@gmail.com

²Department of Computer Science & Engineering, School of
Engineering, Swami Vivekananda University,
Barrackpore, Kolkata- 700121, West Bengal, India
sangitabose21@gmail.com

Abstract

Because computational context comprehension refers to an agent's capacity to combine different knowledge sources for decision-making, it is typically recognized as a

necessary precondition for advanced machine reasoning skills, such those found in artificial intelligence (AI). Two traditional approaches in the search for such machine sense-making capabilities are data-driven and knowledge-driven methodologies.

While data-driven approaches aim to simulate the statistical regularities of occurrences through real-world observations, they lack mechanisms for organically adding outside information and are nonetheless challenging to comprehend. Conversely, knowledge-driven methods, combine structured knowledge bases, perform symbolic reasoning based on axiomatic principles, and are more interpretable in their inferential processing; however, they often lack the ability to estimate the statistical salience of an inference. To combat these issues, we propose the use of hybrid AI methodology as a general framework for combining the strengths of both approaches. Specifically, we inherit the concept of neuro-symbolism as a way of using knowledge-bases to guide the

learning progress of deep neural networks. We further ground our discussion in two applications of neuro-symbolism and, in both cases, show that our systems maintain interpretability while achieving comparable performance, relative to the state-of-the-art.

1. Introduction

Human brain naturally possesses the ability to perceive context, which helps us make decisions in challenging sensory circumstances. For example, humans can combine information from auditory and visual modalities to do a number of activities, from operating a motor vehicle to creating reasonable conclusions drawn from common sense. This sense-making ability was described by Allen Newell and Herbert Simon in their theory of cognition [30, 31]. Through sensory stimuli, humans accumulate experiences, generalize, and reason over them, storing the resulting knowledge in memory; during task-execution, the dynamic combination of live experience and distilled knowledge enables humans to make decisions that are efficient in terms of time and assess the quality of those decisions by taking into account external feedback.

One of the long-standing objectives of artificial intelligence (AI) practice and research, both in business and academia, has been to give robots this sense-making capacity.

Knowledge-driven and data-driven approaches are two traditional strategies in the quest for this kind of machine intelligence. In addition to being essential for enhancing machine autonomy, sense-making is also necessary to enable smooth communication with people.

Because they have similar mental representations of the physical world and social

environment, humans are able to interact successfully with one another [16]. These models are essential for understanding how decision-making processes work and for fostering

reciprocal trust because they make contextual knowledge apparent. Similarly, we might claim that “explainable AI” is dependent on human ability to observe the decision-

making processes that underpin machine sense-making and is a consequence or affordance of computational context comprehension [20].

2. Literature Review

A new field of study called “neuromyotonic fusion” aims to combine deep learning systems with symbolic thinking, which has historically relied on logic and symbols. The goal of this fusion is to produce more potent and understandable AI systems by utilizing the advantages of both symbolic and connectionist methodologies. A survey of the literature on neuromyotonic fusion identifies a number of important themes and areas for further study.

2.1. Inspiration and Context:

Scholars frequently start off by outlining the drawbacks of exclusively connectionist or symbolist methods. While deep learning systems might not be transparent or interpretable, symbolic reasoning might have trouble managing uncertainty and learning from data.

The goal of neuromyotonic fusion is usually to get beyond these restrictions by fusing the data-driven and pattern recognition skills of neural networks with the organized and symbolic knowledge representation of neural networks.

2.2. Methods of Integration:

Symbolic reasoning and deep learning have been combined in a number of ways. These include creating hybrid designs that incorporate both symbolic and connectionist elements, and integrating symbolic representations into neural networks for the purpose of learning symbolic rules.

In order to successfully combine symbolic and neural representations, methods including logic-guided learning, neuromyotonic knowledge graphs, and attention processes have been investigated.

2.3. Uses:

Applications of neuromyotonic fusion in several disciplines are covered in the literature. Natural language comprehension,

robotics, knowledge graph reasoning, and scientific discovery are a few examples. Researchers investigate how, in certain domains, neuromyotonic systems can perform better than conventional methods.

2.4. Explainability and Interpretability:

One major challenge is how to handle the black-box nature of deep learning models. The need of developing accurate and interpretable neuromyotonic models is frequently emphasized in the literature. Creating techniques to draw significant conclusions from the integration of neuronal and symbolic representations is part of this.

2.5. Obstacles and Unanswered Questions:

Typically, the review lists the difficulties in implementing neuromyotonic fusion, including the requirement for efficient learning algorithms that can handle material that is both symbolic and non-symbolic, guaranteeing scalability, and addressing the semantic difference between neural and symbolic representations.

It is common to identify open problems and future research areas, such as investigating innovative neuromyotonic architectures, creating benchmark datasets, and figuring out the theoretical foundations of such fusion.

2.6. Experiments and Case Studies:

Important case studies and tests that have been carried out to confirm the efficacy of neuromyotonic techniques are summarized in the literature review. These investigations might demonstrate how, for some tasks, neuromyotonic fusion works better than purely symbolic or purely neural techniques.

2.7. Comparing with Similar Methods:

It's common to examine comparisons with alternative strategies like reinforcement learning, meta-learning, and transfer learning. The goal of research is to demonstrate the special

benefits of neuromyotonic fusion in the resolution of challenging issues.

2.8. Moral Aspects to Take into Account:

The literature review may discuss neuromyotonic fusion-related ethical problems, particularly those pertaining to responsibility, justice, and prejudice. It talks about how important it is to create systems that follow moral guidelines.

3. Methodology

An interdisciplinary method called neurosymbolic fusion combines deep learning architectures—which are excellent at extracting patterns and representations from data—with symbolic reasoning, which entails working with symbols and abstract representations. By combining the best features of both methods, this integration seeks to solve complicated issues that could call for both data-driven learning and symbolic reasoning. The following is a general approach to neurosymbolic fusion:

3.1. Describe the domain of the problem:

Clearly state the issue you wish to solve and indicate if data-driven learning, symbolic reasoning, or both are required.

3.2. Significant Illustration:

Determine the problem's symbolic elements. Describe the entities, connections, and regulations that control the issue domain. Logic-based representations are commonly employed in symbolic thinking.

3.3. Information Display:

Identify the problem's data-related components. Determine what kinds of data are available and how to encode them so that a neural network may use them as input. This might contain text, pictures, structured data, etc.

3.4. Select the Architecture of Neurosymbolism:

Choose an architecture that combines deep learning and symbolic reasoning. Symbolic modules added to neural networks, neurosymbolic systems, or hybrid models combining the two are common techniques.

3.5. Data-to-Symbolic Mapping:

Provide a method for converting symbolic representations into a neural network-input format. Encoding symbols into vectors or other numerical representations might be required for this.

3.6. Maps Data to Symbolic Domains:

Create a system that maps the neural network's output back into symbolic representations in a similar manner. Decoding numerical outputs into symbolic representations could be required for this.

3.7. Generation of Training Data:

Create or compile a dataset with relevant data and symbolic representations. The neural network will be trained using this dataset. Make sure the dataset is reflective of the issued domain and diverse.

3.8. Training Models:

Utilizing the provided dataset, train the integrated neurosymbolic model. Adjust the model such that the data-driven learning and the symbolic reasoning components are balanced. Try different architectures and hyperparameters to maximize performance.

3.9. Assessment:

Utilizing the proper measures, assess the neurosymbolic model's performance. Evaluate the model's ability to represent both data-driven patterns and symbolic reasoning. Examine its effectiveness in comparison to methods that are solely neurological or symbolic.

3.10. Repeated Fine Tuning:

Refine the model iteratively based on the evaluation outcomes. To enhance overall performance, modify neural network design, training methods, or symbolic rules.

3.11. Explainability and Interpretability:

Make sure the outcomes of the neurosymbolic model are comprehensible and interpretable. Understanding how the

model blends learnt patterns with symbolic thinking depends on this.

3.12. Practical Implementation:

If the neurosymbolic model performs satisfactorily, apply it to real-world situations. Keep an eye on how it performs in production and update as needed.

Recall that the precise features of the issue domain, as well as the meticulous construction and training of the integrated model, determine how successful neurosymbolic fusion will be.

Conclusion

To sum up, neurosymbolic fusion is a promising area of artificial intelligence that combines the strengths of deep neural networks and logical reasoning in data-driven learning with the smoothness of symbolic systems. This combined technique provides a potent answer for large issues requiring both symbolic comprehension and pattern recognition, beyond the limitations of separate approaches. Neurosymbolic fusion, which bridges the gap between implicit feature acquisition and explicit knowledge representation, has the potential to improve decision-making in a variety of fields.

The technique described focuses on the definition of problem domains, the creation of suitable mappings, and the iterative refinement of models through a systematic integration of symbolic and neural components. By carefully coordinating these stages, neurosymbolic models are able to accurately represent the complexities of real-world situations and provide outcomes that are understandable and comprehensible. Therefore, the use of neurosymbolic fusion in real-world scenarios provides a comprehensive and cooperative method of problem-solving, proving the ability to address issues that are beyond the scope of simply symbolic or data-driven techniques. Neurosymbolic fusion is a prime example of the possibilities of multidisciplinary collaboration as this field of study develops. It shows how to combine the advantages of many AI paradigms to create more resilient and intelligent systems.

References

‘inadequate safety culture’ contributed to uber automated test vehicle crash - ntsb calls for federal review process for automated vehicle testing on public roads. <https://ntsb.gov/news/press-releases/Pages/NR20191119c.aspx>. (Accessed on 11/26/2019).

ALSHARGI, F., SHEKARPOOR, S., SORU, T., AND SHETH, A. Metrics for evaluating quality of embeddings for ontological concepts. In AAAI 2019 Spring Symposium on Combining Machine Learning with Knowledge Engineering (AAAI-MAKE) (2019).

ANDERSON, J. R. How can the human mind occur in the physical universe?, vol. 3. Oxford University Press, 2009.

BAKER, C. F., FILLMORE, C. J., AND LOWE, J. B. The berkeley framenet project. In Proceedings of the 17th international conference on Computational linguistics-Volume 1 (1998), Association for Computational Linguistics, pp. 86–90.

BAUER, L., WANG, Y., AND BANSAL, M. Commonsense for generative multi-hop question answering tasks. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (Brussels, Belgium, Oct.-Nov. 2018), Association for Computational Linguistics, pp. 4220–4230.

BORDES, A., USUNIER, N., GARCIA-DURAN, A., WESTON, J., AND YAKHNENKO, O. Translating embeddings for modeling multi-relational data. In Advances in Neural Information Processing Systems(2013), pp. 2787–2795.

BOSELUT, A., RASHKIN, H., SAP, M., MALAVIYA, C., CELIKYILMAZ, A., AND CHOI, Y. COMET:

Commonsense transformers for automatic knowledge graph construction. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (Florence, Italy, July 2019), Association for Computational Linguistics, pp. 4762–4779.

CAESAR, H., BANKITI, V., LANG, A. H., VORA, S., LIONG, V. E., XU, Q., KRISHNAN, A., PAN, Y., BALDAN,

G., AND BEIJBOM, O. nuscenes: A multimodal dataset for autonomous driving. arXivpreprint arXiv:1903.11027 (2019).

CHEN, M., TIAN, Y., YANG, M., AND ZANIOLO, C. Multilingual knowledge graph embeddings

for cross-lingual knowledge alignment. In Proceedings of the 26th International Joint Conference on Artificial Intelligence (2017), AAAI Press, pp. 1511–1517.

DAVIS, E. Representations of commonsense knowledge. Morgan Kaufmann, 2014