Neuro-Symbolic Methods in Natural Language Processing: A Review

Mst Shapna Akter[1][a], Md Fahim Sultan[1][b] and Alfredo Cuzzocrea[2,3,*][c]
[1]Department of Computer Science and Engineering, Oakland University, Rochester, MI 48309, U.S.A.
[2]iDEA Lab, University of Calabria, Rende, Italy
[3]Dept. of Computer Science, University of Paris City, Paris, France

Keywords: Natural Language Processing, Neuro-Symbolic Techniques, Reasoning, Interpretability.

Abstract: Neuro-Symbolic (NeSy) techniques in Natural Language Processing (NLP) combine the strengths of neural network-based learning with the clear interpretability of symbolic methods. This review paper explores recent advancements in neurosymbolic NLP methods. We carefully highlight the benefits and drawbacks of different approaches in various NLP tasks. Additionally, we support our evaluations with explanations based on theory and real-world evidence. Based on our review, we suggest several potential research directions. Our study contributes in three main ways: (1) We present a detailed, complete taxonomy for the Neuro-Symbolic methods in the NLP field; (2) We provide theoretical insights and comparative analysis of the Neuro-Symbolic methods; (3) We propose future research directions to explore.

1 INTRODUCTION

The recent proliferation of deep learning models in the field of Natural Language Processing (NLP) has resulted in notable advancements, particularly in their performance on benchmark tasks. However, these models are not without limitations (Xu and McAuley, 2023). In particular, they often face tasks that require intricate reasoning or the fusion of diverse fragments of knowledge (Rajani et al., 2020). Further complicating matters is their propensity for data inefficiency and issues pertaining to model generalizability. This is largely attributed to their inherently opaque nature and the absence of a well-defined, structured understanding of the input data they process. In the field of natural language processing (NLP), black box and heuristic methods such as LSTM-DQN (Narasimhan et al., 2015), LSTM-DRQN (Yuan et al., 2018), and CREST (Chaudhury et al., 2020) were used for text-based policy learning. However, these methods showed unsatisfactory results and overfitted the training data. Similarly, the BLINK (Wu et al., 2019) method for short-text and long-text Entity Linking also demonstrated poor performance. To mitigate these issues, the idea of incorporating neuro-symbolic methods in NLP has been proposed. This process involves enhancing a database with new knowledge particles. Early work by (Chaudhury et al., 2021a) has explored using the neuro-symbolic approach to solve text based policy learning. Then, (Jiang et al., 2021a) proposed a neuro-symbolic model for solving entity linking which seems to increase the F1 score by more than 4% over previous state-of-the-art methods on a bechmark dataset. Therefore, more neurosymbolic works have been previously proposed (Gupta et al., 2021; Kimura et al., 2021b; Pacheco et al., 2022b; Zhu et al., 2022; Langone et al., 2020), showing apprealing performance in the benchmark dataset.

Present Work. This manuscript provides a comprehensive overview of recent advancements in neurosymbolic methods applied to NLP.

- **Comprehensive Review With New Taxonomies:** We provide a thorough review of the neuro-symbolic methods used in NLP, accompanied by new taxonomies. We review the research with different NeSy tasks with a comprehensive comparison and summary.
- **Theoretical Insights:** We analyze NeSy methods theoretically, discussing their advantages, disadvantages, and unresolved challenges for future research.

[a] https://orcid.org/0000-0002-9859-6265
[b] https://orcid.org/0009-0009-2550-257X
[c] https://orcid.org/0000-0002-7104-6415

*This research has been made in the context of the Excellence Chair in Big Data Management and Analytics at University of Paris City, Paris, France.

Akter, M. S., Sultan, M. F., Cuzzocrea and A.
Neuro-Symbolic Methods in Natural Language Processing: A Review.
DOI: 10.5220/0013453100003967
In *Proceedings of the 14th International Conference on Data Science, Technology and Applications (DATA 2025)*, pages 274-282
ISBN: 978-989-758-758-0; ISSN: 2184-285X

- **Wide Coverage on Emerging Advances and Outlook on Future Directions:** We examine emerging trends in NeSy methods, including novel models that integrate neural and symbolic approaches. We offer insights into future research directions and areas for improvement.

Related Work. As this field is still in its early stages, there is currently a scarcity of surveys available. Previous work by (Hamilton et al., 2022) acknowledges the significance of reasoning in NeSy, it does not extensively explore the range of reasoning techniques or the challenges associated with their implementation. In contrast, our paper conducts a comprehensive literature review on NeSy methods in natural language processing, providing a systematic understanding of methodologies, comparing different approaches, and offering insights to inspire new ideas in the field.

2 PRELIMINARY ON NEUROSYMBOLIC METHODS

2.1 Neurosymbolic Tasks in NLP Field

Neurosymbolic methods aim to harmonize the learning capabilities of neural networks from data and the reasoning abilities of symbolic systems based on predefined logic. This combination enhances several tasks such as Natural Language Inference (NLI) (Feng et al., 2022a), Linguistic Frameworks (Prange et al., 2022), Sentiment Analysis (Cambria et al., 2022), Question Answering (Gupta et al., 2021; Ma et al., 2019), Entity Learning (Chaudhury et al., 2021a), and Sentence Classification (Sen et al., 2020a).

Natural Language Inference. Given “Bob is a doctor” and “Bob has a medical degree”, a NeSy model would infer that the latter statement is likely true, using neural networks to understand semantics and symbolic logic to make the inference (Feng et al., 2022a).

Linguistic Frameworks. For a sentence like “The ball was thrown by the boy”, neurosymbolic methods use neural networks to process word meanings and a symbolic system to parse grammatical structure (Prange et al., 2022).

Sentiment Analysis. In a sentence such as “I absolutely loved the thrilling plot of the movie!”, neurosymbolic methods would use neural networks to detect positive sentiment and symbolic systems to provide rules-based explanations (e.g., the word “loved” indicates positive sentiment).

Question Answering. If asked, “Where are the internet and 4G services available?”, the task is to extract the relevant answers from the text under concern. The answer is “global world” and incorporate the knowledge with Neural Network (Gupta et al., 2021; Ma et al., 2019).

Entity Linking. From the sentence “Turing was a pioneer in computer science”, neurosymbolic methods use neural networks to identify “Turing” as an entity and symbolic systems to categorize him under “computer science” (Chaudhury et al., 2021a).

Sentence Classification. Sentence Classification focuses on categorizing sentences into predefined classes. Neurosymbolic methods employ neural networks to capture sentence representations and symbolic systems to assign them to appropriate classes based on predefined rules or logic (Sen et al., 2020a).

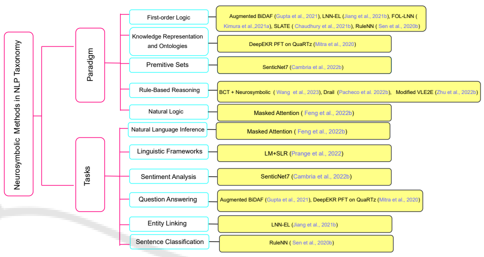

Figure 1: Taxonomy of Neuro-Symbolic Methods in Natural Language Processing.

2.2 Neural Networks in Neurosymbolic Systems

In the field of Natural Language Processing (NLP), neurosymbolic systems combine the advantages of neural networks and symbolic reasoning. While neural networks excel in learning from data and handling noise, symbolic systems bring interpretability and rule-based reasoning. A variety of neural network techniques, including Logical Neural Networks, Neural Natural Logic, reinforcement learning, unsupervised learning, and Neurosymbolic Generative Models, have been effectively applied to facilitate learning in these systems. Besides, Logical Neural Networks (LNNs) combine symbolic logic rules with neural network architectures, forming systems that can be trained using standard deep learning methods while adhering to logical constraints. For example, when classifying a mammal based on “has hair” (A) and “gives live birth” (B), an LNN embeds the rules “If A, then mammal” and “If B, then mammal”, predicting “mammal” if either is true. Network layers mirror these rules, with the final output being their logical (Jiang et al., 2021b; Riegel et al., 2020; Chaudhury et al., 2021a). In addition, Neural Natural Logic integrates neural networks and symbolic logic, transforming logical expressions into vector spaces while retaining logical relationships, thus allowing neural networks to manipulate vector-encoded logic expressions. For example, logical expression "A and B" is represented as vectors **a** and **b**. A neural network learns the "and" operation as function $f(\cdot)$, wherein $f(\mathbf{a},\mathbf{b})$ resembles the vector representation of "A and B". This permits logical reasoning within vector space, harnessing both symbolic logic's power and neural networks' flexibility (Feng et al., 2022b). On the other hand, Reinforcement learning (RL) entails an agent learning decision-making through environment interaction and feedback via rewards or penalties. This can be represented as a Markov decision process (MDP), with components: states (S), actions (A), state transition probability (P), reward function (R), and discount factor (Y). In an NLP context, states may represent stages of text generation, actions can be word choices, and rewards assess sentence fluency and coherence. For example, an RL-based sentence generator is rewarded for grammatically correct sentences, penalized for ungrammatical ones (Wang et al., 2022; Gupta et al., 2021; Kimura et al., 2021b). As for, unsupervised neurosymbolic representation learning merges unsupervised learning, symbolic reasoning, and neural networks. This approach uses domain-specific languages (DSLs) for knowledge representation, which when paired with neural computations, results in clear, well-separated data representations. Examples of this can be seen in Variational Autoencoders (VAEs), which effortlessly combine symbolic programming with deep learning (Zhan et al., 2021). In neurosymbolic generative models infuse high-level structure into the creation of sequence data such as text or music. This approach satisfies relational constraints between example subcomponents, enhancing both high-level and low-level coherence in generated data. This sophisticated method improves the quality of generated data, particularly in low-data environments, by integrating symbolic reasoning into deep generative models (Young et al., 2022).

2.3 Advantages of Neurosymbolic Approaches

Neurosymbolic approaches offer several advantages in addressing the limitations of individual approaches. Using the complementary strengths of both paradigms, neurosymbolic approaches aim to enhance natural language processing tasks. This section discusses the key advantages of neurosymbolic approaches in NLP. Neurosymbolic approaches augment NLP's reasoning capabilities by integrating neural networks' expertise in identifying intricate patterns and learning from large datasets with symbolic reasoning's explicit logical inference. This combination allows neurosymbolic models to unite statistical learning with logical reasoning, fostering a more structured, nuanced understanding of language (Feng et al., 2022a). Besides, neurosymbolic approaches enhance interpretability and explainability over solely neural network models. They offer a clear framework through symbolic reasoning, illuminating the decision-making process, and the explicit representation of knowledge permits identification of reasoning steps and prediction justifications. This clarity is vital in domains that require explanation, such as law, medicine, and critical decision making (Sen et al., 2020a; Verga et al., 2021). For the handling data scarcity, NNNs typically need vastly labeled data for superior performance, a challenging requirement in many NLP tasks due to the effort and cost of obtaining annotated data (Zhao et al., 2020). Neurosymbolic approaches address this by using symbolic reasoning to transfer knowledge across tasks and domains (Deng et al., 2021). This method, through the inclusion of prior knowledge and explicit rules, counters data scarcity and enhances performance even with sparse labeled data. Additionally, Neurosymbolic approaches offer adaptability and flexibility. While neural networks are adept at learning from varied, unstructured data, symbolic reasoning provides a structure for integrating domain-specific rules (**?**). Thus, neurosymbolic models can adapt to different task requirements and include contextual data, all within the bounds of logical constraints.

3 TAXONOMY OF NEUROSYMBOLIC METHODS IN NLP

In this paper, our focus is on examining five commonly employed paradigms in the field of neurosymbolic natural language processing (NLP). First-order Logic, Knowledge Representation and Ontologies 2, Primitive Sets, Rule-Based Reasoning, and Natural Logic. These paradigms have demonstrated remarkable effectiveness in various prominent Neurosymbolic tasks. In the subsequent sections, we provide detailed explanations of each paradigm, as presented in Figure 1.

3.1 First-Order Logic

First-order logic (FOL) serves as a structured formal language that allows for the articulation of relationships and assertions concerning various entities. It is constituted by a variety of logical symbols, including predicates, variables, quantifiers, and connectives (Kimura et al., 2021a). FOL has become a valuable asset for precisely encapsulating knowledge in NLP tasks (Wang et al., 2020). An instance of this can be seen in the representation of the statement, "All cats are mammals," which translates to $\forall x\, \text{Cat}(x) \rightarrow \text{Mammal}(x)$ in FOL. Here, $\forall$ is the universal quantifier, $\text{Cat}(x)$ stands for "x is a cat", and $\text{Mammal}(x)$ signifies "x is a mammal" (Lu et al., 2022). FOL finds utility in the modeling of diverse linguistic phenomena, such as logical inference, semantic relationships, and knowledge representation. (Chaudhury et al., 2021b) presented a symbolic rule learning framework for text-based RL. They employed an MLP with symbolic inputs and a Logical Neural Network (LNN) - a symbolic reasoning-based approach - to learn lifted rules from first-order symbolic abstractions of textual observations. Their results displayed superior generalization to unseen games compared to prior text-based RL methods. Following a similar neuro-symbolic approach, (Jiang et al., 2021b) introduced LNN-EL, an innovation that blends interpretable rules based on FOL with the high performance of neural learning for short text entity linking. On another front, (Kimura et al., 2021b) proposed a technique that involved converting text into FOL and subsequently training the action policy in LNN. Lastly, (Gupta et al., 2021) incorporated the domain knowledge, expressed as FOL predicates, into a deep neural network model named Bidirectional Attention Flow (BiDAF). Equally noteworthy, (Sen et al., 2020b) unveiled a neural network architecture specifically designed to learn transparent models for sentence classification. In this ingenious approach, the models are presented as rules articulated in first-order logic, a variant characterized by well-defined semantics that are readily comprehensible to humans. This approach carries the key advantage of the inherent interpretability of its models, akin to the FOL-based techniques introduced by (Jiang et al., 2021b) and (Kimura et al., 2021b) Each of these works demonstrates the diverse and significant applications of FOL in advancing NLP tasks.

3.2 Knowledge Representation and Ontologies

Symbolic reasoning in NLP often involves the use of knowledge representation formalisms and ontologies. Knowledge representation allows for the explicit representation of knowledge in a structured manner (Mitra et al., 2020). Ontologies provide a formal representation of concepts, relationships, and properties within a specific domain. Common knowledge representation languages in NLP include RDF (Resource Description Framework) and OWL (Web Ontology Language) (Cuzzocrea, 2006). These formalisms facilitate reasoning tasks by defining rules, axioms, and relations between concepts.

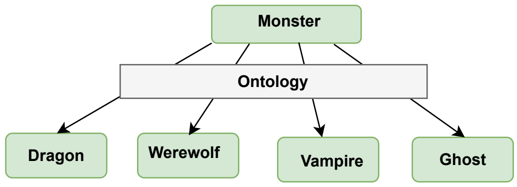

Figure 2: The diagram demonstrates a hierarchical knowledge representation using ontologies, with 'Monster' as the superclass and 'Dragon', 'Werewolf', 'Vampire', and 'Ghost' as subclasses, showcasing the hierarchical structure of knowledge.

3.3 Primitive Sets

In neurosymbolic reasoning, primitive sets are the fundamental operations or predicates from which complex expressions can be constructed. In a symbolic system used for NLP tasks, primitives might include operations for string manipulation, such as concatenation, or predicates to verify certain properties of words or phrases (Cambria et al., 2022). For example, the predicate is_noun(x) might be a primitive that checks whether x is a noun. This can be used to construct more complex expressions, such as 'is _noun(x) AND is _verb(y)', which checks whether x is a noun and y is a verb. These primitives could be used in various NLP tasks like semantic parsing or question answering, where the model needs to understand and manipulate linguistic structures. For instance, in a question-answering task, the system might utilize primitives such as find_entity (x), locate_in_text (x), or extract_answer (x, y), where x and y' represent text or entities in the text. By combining these primitives, the system could parse a question, locate relevant parts of the text, and extract an answer.

3.4 Rule-Based Reasoning

Rule-based systems are designed to generate conclusions or make decisions based on a pre-defined set of rules. Rule-based reasoning systems are characterized by their interpretability and transparency, as the reasoning process follows explicit rules that can be easily understood and audited by humans. For instance, in the study conducted by (Wang et al., 2023), they implemented rule-based symbolic modules for various tasks. Within their arithmetic module, they successfully executed operations such as multiplication where inputting “mul 3 6” produced the result “18”, showcasing rule-based reasoning in solving numerical tasks. Similarly, in the sphere of navigation, their module guided an agent’s movement by generating the next step towards a desired destination. When given the instruction “next step to living room”, the module returned “The next location to move to is: hallway”. This practical application of rule-based spatial reasoning demonstrates the versatile capabilities of such systems. Similarly, (Pacheco et al., 2022a) utilized a rule-based reasoning approach in their DRaiL framework. By defining entities, predicates, and probabilistic rules, they were able to model intricate inter-dependencies among various decisions. These rules, along with a set of constraints, formed the basis of their reasoning process. The integration of these components allowed them to generate complex predictions for given problems, providing a practical demonstration of the effectiveness of rule-based reasoning in natural language understanding tasks. (Zhu et al., 2022) adopted a neuro-symbolic (NS) reasoning approach, a subtype of rule-based reasoning, in their work on vision-language tasks. The query semantics was represented as a functional program, essentially a set of rules derived from the query, which was then executed on the structured representation of the image set to predict an answer. This method showcases how rule-based reasoning can be efficiently implemented even in complex, multimodal domains. (Zhan et al., 2021) employed a rule-based reasoning method in their unsupervised learning framework, using rules to model the relationships and interactions between objects in a scene.

3.5 Comparison and Discussion

The exploration of neurosymbolic methods in Natural Language Processing (NLP) represents a vibrant area of research, which over the years has unfolded a range of methodological paradigms, namely, First-order Logic, Knowledge Representation and Ontologies, Primitive Sets, Rule Based Reasoning, and Natural Logic. The core intention underpinning these paradigms converges towards leveraging the strengths of both neural and symbolic perspectives for enhanced language understanding. Yet, they differ in their theoretical underpinnings and practical applications, each contributing unique strengths and perspectives. In Table 1, we present a comprehensive evaluation of the mentioned paradigms, employing a rating-based approach that encompasses a range of evaluation criteria. For Semantic Understanding, First-order Logic and Natural Logic excel in comprehending and manipulating logical structures in languages, enhancing semantic understanding. In contrast, paradigms like Primitive Sets and Rule Based Reasoning rely on predefined rules or primitives, and their semantic understanding depends on the effectiveness and comprehensiveness of these linguistic encapsulations. Regarding scalability, Knowledge Representation and Ontologies offer a distinct advantage. By employing standardized representation languages like RDF and OWL, these paradigms can cater to large-scale, complex knowledge structures, which is crucial in dealing with extensive or complex language corpora.

4 THEORETICAL INSIGHTS

This section sheds light on the theoretical considerations involved in the application and integration of neuro-symbolic methods in NLP. The core idea behind these methods is to blend the symbolic reasoning capabilities with the learning power of neural networks (Yang et al., 2021). This integration could be treated as a unified system Ψ that takes an input sequence x and produces an output sequence y as $\Psi(x) = y$. For the optimization perspective, the cost function $\mathcal{L}$ in neuro-symbolic methods could be a combination of the loss in the symbolic reasoning component L_{sym} and the loss in the neural learning component L_{nn}, formulated as: $\mathcal{L}(\Theta) = \alpha L_{sym} + \beta L_{nn}$. Here α and β are weights reflecting the significance of each component in a specific task, and Θ represents the model parameters. This raises a key question about how to balance between symbolic reasoning and neural learning, as it largely impacts the model performance. Usually, a neuro-symbolic method will try to learn the best reasoning strategy or symbolic representation by minimizing L_{sym} and enhance the learning capabilities by minimizing L_{nn}. However, it’s critical to note that an overemphasis on symbolic reasoning could lead to a model lacking generalization capabilities, while overfitting on neural learning might cause the model to lose its interpretability and explicit reasoning capability. On the inference side, the outputs from neuro-symbolic methods generally involve both symbolic and neural components. The symbolic part typically includes interpretable rules, logical forms, or other symbolic structures, while the neural component provides the probabilities or confidences over those structures. The ultimate decision would be the one that maximizes the combined confidence score.

Table 1: Comparison of neurosymbolic methods from different evaluation scopes. “SU” indicates semantic understanding, “SC” indicates scalability, “VE” indicates versatility, and “IN” indicates interpretability. We divide the degree into three grades: L (low), M (middle), and H (high), and the ↑ indicates that the higher grade performance is better while the ↓ is the opposite.

Taxonomy	Strategy	Representative Model	Evaluation Score			
			SU↑	SC↑	VE↑	IN↑
First-order (§ 3.1)	Policy Learning	SLATE (Chaudhury et al., 2021b)	H	M	M	H
	Question Answering	Augmented BiDAF (Gupta et al., 2021)	H	M	M	H
	Entity Linking	LNN-EL (Jiang et al., 2021b)	H	M	M	H
	Sentence Classification	RuleNN (Sen et al., 2020b)	H	M	M	H
KR and Ontologies (§ 3.2)	Question Answering	DeepEKR PFT on QuaRTz (Mitra et al., 2020)	H	L	H	H
Primitive Sets (§ 3.3)	Sentiment Analysis	SenticNet7 (Cambria et al., 2022)	M	H	M	M
RB Reasoning (§ 3.4)	Linguistic Framework	Drail (Pacheco et al., 2022b)	H	L	M	H
Natural Logic (§ 3.5)	Natural Language Inference	Masked Attention (Feng et al., 2022b)	H	L	M	H

5 FUTURE DIRECTIONS

While numerous technical strategies have been suggested for Neuro-Symbolic methods as outlined in our survey, several prospective avenues still persist. The creation of reliable and efficient inference strategies presents a significant area for further research. Current techniques such as greedy search (Ma and Hovy, 2015), beam search (Hale et al., 2018), or guided decoding (Chatterjee et al., 2017) have both benefits and drawbacks. Future work should aim to devise strategies that secure top-tier outputs while balancing computational cost. One potential solution could be the development of adaptive multi-modal inference strategies (Bhargava, 2020). These would intelligently switch or combine different strategies based on the nature of the problem and the data at hand. By dynamically choosing or merging the most suitable techniques, this approach could offer the best of all worlds, optimizing output quality and computational efficiency. Neurosymbolic methods excel in various NLP tasks but have potential to expand into areas like Machine Translation (Brants et al., 2007), Text Summarization (Liu and Lapata, 2019), and Dialogue Systems (Wen et al., 2015). This could involve creating a Neuro-Symbolic Multitask Learning Framework (NSMLF) with a core neural network model sharing lower-level representations across tasks and an upper symbolic reasoning layer for task-specific modeling. For instance, shared neural components could learn language patterns from a broad text corpus, while symbolic rules at upper levels provide task-specific precision and interpretability. The NSMLF's design provides task-agnostic flexibility.

6 CONCLUSIONS AND FUTURE WORK

Our paper provides a comprehensive overview of Neurosymbolic NLP and highlights its potential to revolutionize the field. By combining neural networks and symbolic reasoning techniques, neurosymbolic methods offer a unified approach that addresses the limitations of individual paradigms. This integration allows for enhanced semantic understanding, interpretability and scalability in NLP tasks. Moreover, our proposed research directions shed light on the future of this field, offering exciting opportunities for further advancements. As the field of NLP continues to evolve, neurosymbolic methods hold great promise for the development of more advanced and interpretable language models, including emerging machine learning applications (e.g., (Howlader et al., 2018; Camara et al., 2018; Leung et al., 2019)).

ACKNOWLEDGMENTS

This research is supported by the ICSC National Research Centre for High Performance Computing, Big Data and Quantum Computing within the NextGenerationEU program (Project Code: PNRR CN00000013) and by the National Aeronautics and Space Administration (NASA), under award number 80NSSC20M0124, Michigan Space Grant Consortium (MSGC).

REFERENCES

Bhargava, P. (2020). Adaptive transformers for learning multimodal representations. *arXiv preprint arXiv:2005.07486*.

Brants, T., Popat, A. C., Xu, P., Och, F. J., and Dean, J. (2007). Large language models in machine translation. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 858–867, Prague, Czech Republic. Association for Computational Linguistics.

Camara, R. C., Cuzzocrea, A., Grasso, G. M., Leung, C. K., Powell, S. B., Souza, J., and Tang, B. (2018). Fuzzy logic-based data analytics on predicting the effect of hurricanes on the stock market. In *2018 IEEE International Conference on Fuzzy Systems, FUZZ-IEEE 2018, Rio de Janeiro, Brazil, July 8-13, 2018*, pages 1–8. IEEE.

Cambria, E., Liu, Q., Decherchi, S., Xing, F., and Kwok, K. (2022). SenticNet 7: A commonsense-based neurosymbolic AI framework for explainable sentiment analysis. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3829–3839, Marseille, France. European Language Resources Association.

Chatterjee, R., Negri, M., Turchi, M., Federico, M., Specia, L., and Blain, F. (2017). Guiding neural machine translation decoding with external knowledge. In *Proceedings of the Second Conference on Machine Translation*, pages 157–168, Copenhagen, Denmark. Association for Computational Linguistics.

Chaudhury, S., Kimura, D., Talamadupula, K., Tatsubori, M., Munawar, A., and Tachibana, R. (2020). Bootstrapped q-learning with context relevant observation pruning to generalize in text-based games. *arXiv preprint arXiv:2009.11896*.

Chaudhury, S., Sen, P., Ono, M., Kimura, D., Tatsubori, M., and Munawar, A. (2021a). Neuro-symbolic approaches for text-based policy learning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3073–3078.

Chaudhury, S., Sen, P., Ono, M., Kimura, D., Tatsubori, M., and Munawar, A. (2021b). Neuro-symbolic approaches for text-based policy learning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3073–3078, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Cuzzocrea, A. (2006). Combining multidimensional user models and knowledge representation and management techniques for making web services knowledge-aware. *Web Intelligence and Agent Systems: An international journal*, 4(3):289–312.

Deng, S., Zhang, N., Li, L., Chen, H., Tou, H., Chen, M., Huang, F., and Chen, H. (2021). Ontoed: Low-resource event detection with ontology embedding. *arXiv preprint arXiv:2105.10922*.

Feng, Y., Yang, X., Zhu, X., and Greenspan, M. (2022a). Neuro-symbolic natural logic with introspective revision for natural language inference. *Transactions of the Association for Computational Linguistics*, 10:240–256.

Feng, Y., Yang, X., Zhu, X., and Greenspan, M. (2022b). Neuro-symbolic natural logic with introspective revision for natural language inference. *Transactions of the Association for Computational Linguistics*, 10:240–256.

Gupta, K., Ghosal, T., and Ekbal, A. (2021). A neuro-symbolic approach for question answering on research articles. In *Proceedings of the 35th Pacific Asia Conference on Language, Information and Computation*, pages 40–49, Shanghai, China. Association for Computational Lingustics.

Hale, J., Dyer, C., Kuncoro, A., and Brennan, J. (2018). Finding syntax in human encephalography with beam search. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2727–2736, Melbourne, Australia. Association for Computational Linguistics.

Hamilton, K., Nayak, A., Božić, B., and Longo, L. (2022). Is neuro-symbolic ai meeting its promises in natural language processing? a structured review. *Semantic Web*, (Preprint):1–42.

Howlader, P., Pal, K. K., Cuzzocrea, A., and Kumar, S. D. M. (2018). Predicting facebook-users' personality based on status and linguistic features via flexible regression analysis techniques. In *Proceedings of the 33rd Annual ACM Symposium on Applied Computing, SAC 2018, Pau, France, April 09-13, 2018*, pages 339–345. ACM.

Jiang, H., Gurajada, S., Lu, Q., Neelam, S., Popa, L., Sen, P., Li, Y., and Gray, A. (2021a). Lnn-el: A neuro-symbolic approach to short-text entity linking. *arXiv preprint arXiv:2106.09795*.

Jiang, H., Gurajada, S., Lu, Q., Neelam, S., Popa, L., Sen, P., Li, Y., and Gray, A. (2021b). 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)*, pages 775–787, Online. Association for Computational Linguistics.

Kimura, D., Ono, M., Chaudhury, S., Kohita, R., Wachi, A., Agravante, D. J., Tatsubori, M., Munawar, A., and Gray, A. (2021a). Neuro-symbolic reinforcement learning with first-order logic. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3505–3511, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Kimura, D., Ono, M., Chaudhury, S., Kohita, R., Wachi, A., Agravante, D. J., Tatsubori, M., Munawar, A., and Gray, A. (2021b). Neuro-symbolic reinforcement learning with first-order logic. *arXiv preprint arXiv:2110.10963*.

Langone, Rocco and Cuzzocrea, Alfredo and Skantzos, Nikolaos (2020). Interpretable Anomaly Prediction: Predicting anomalous behavior in industry 4.0 settings via regularized logistic regression tools. *Elsevier*, Journal Data & Knowledge Engineering, vol.130 pages 101-850.

Leung, C. K., Braun, P., and Cuzzocrea, A. (2019). Ai-based sensor information fusion for supporting deep supervised learning. *Sensors*, 19(6):1345.

Liu, Y. and Lapata, M. (2019). Text summarization with pretrained encoders. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference* *on Natural Language Processing (EMNLP-IJCNLP)*, pages 3730–3740, Hong Kong, China. Association for Computational Linguistics.

Lu, X., Liu, J., Gu, Z., Tong, H., Xie, C., Huang, J., Xiao, Y., and Wang, W. (2022). Parsing natural language into propositional and first-order logic with dual reinforcement learning. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5419–5431, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Ma, Kaixin and Francis, Jonathan and Lu, Quanyang and Nyberg, Eric and Oltramari, Alessandro(2019). Towards generalizable neuro-symbolic systems for commonsense question answering. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1322–1328, Lisbon, Portugal. Association for Computational Linguistics.

Ma, X. and Hovy, E. (2015). Efficient inner-to-outer greedy algorithm for higher-order labeled dependency parsing. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1322–1328, Lisbon, Portugal. Association for Computational Linguistics.

Mitra, A., Narayana, S., and Baral, C. (2020). Deeply embedded knowledge representation & reasoning for natural language question answering: A practitioner's perspective. In *Proceedings of the Fourth Workshop on Structured Prediction for NLP*, pages 102–111.

Narasimhan, K., Kulkarni, T., and Barzilay, R. (2015). Language understanding for text-based games using deep reinforcement learning. *arXiv preprint arXiv:1506.08941*.

Pacheco, M. L., Roy, S., and Goldwasser, D. (2022a). Hands-on interactive neuro-symbolic NLP with DRaiL. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 371–378, Abu Dhabi, UAE. Association for Computational Linguistics.

Pacheco, M. L., Roy, S., and Goldwasser, D. (2022b). Hands-on interactive neuro-symbolic nlp with drail. In *Proceedings of the The 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 371–378.

Prange, J., Schneider, N., and Kong, L. (2022). Linguistic frameworks go toe-to-toe at neuro-symbolic language modeling. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4375–4391.

Rajani, N. F., Zhang, R., Tan, Y. C., Zheng, S., Weiss, J., Vyas, A., Gupta, A., Xiong, C., Socher, R., and Radev, D. (2020). Esprit: Explaining solutions to physical reasoning tasks. *arXiv preprint arXiv:2005.00730*.

Riegel, Ryan and Gray, Alexander and Luus, Francois and Khan, Naweed and Makondo, Ndivhuwo and Akhalwaya, Ismail Yunus and Qian, Haifeng and Fagin, Ronald and Barahona, Francisco and Sharma, Udit and others (2020). Logical Neural Networks. *arXiv preprint arXiv:2006.13155*.

Sen, P., Danilevsky, M., Li, Y., Brahma, S., Boehm, M., Chiticariu, L., and Krishnamurthy, R. (2020a). Learning explainable linguistic expressions with neural inductive logic programming for sentence classification. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4211–4221.

Sen, P., Danilevsky, M., Li, Y., Brahma, S., Boehm, M., Chiticariu, L., and Krishnamurthy, R. (2020b). Learning explainable linguistic expressions with neural inductive logic programming for sentence classification. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4211–4221, Online. Association for Computational Linguistics.

Wang, R., Jansen, P., Côté, M.-A., and Ammanabrolu, P. (2022). Behavior cloned transformers are neurosymbolic reasoners. *arXiv preprint arXiv:2210.07382*.

Verga, Pat and Sun, Haitian and Soares, Livio Baldini and Cohen, William (2021). Adaptable and Interpretable Neural Memoryover Symbolic Knowledge. In *Proceedings of the 2021 conference of the north american chapter of the association for computational linguistics: human language technologies*, pages 3678–3691.

Wang, R., Jansen, P., Côté, M.-A., and Ammanabrolu, P. (2022). Behavior cloned transformers are neurosymbolic reasoners. *arXiv preprint arXiv:2210.07382*.

Wang, R., Jansen, P., Côté, M.-A., and Ammanabrolu, P. (2023). Behavior cloned transformers are neurosymbolic reasoners. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2777–2788, Dubrovnik, Croatia. Association for Computational Linguistics.

Wang, R., Tang, D., Duan, N., Zhong, W., Wei, Z., Huang, X., Jiang, D., and Zhou, M. (2020). Leveraging declarative knowledge in text and first-order logic for fine-grained propaganda detection. *arXiv preprint arXiv:2004.14201*.

Wen, T.-H., Gašić, M., Mrkšić, N., Su, P.-H., Vandyke, D., and Young, S. (2015). Semantically conditioned LSTM-based natural language generation for spoken dialogue systems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1711–1721, Lisbon, Portugal. Association for Computational Linguistics.

Wu, L., Petroni, F., Josifoski, M., Riedel, S., and Zettlemoyer, L. (2019). Scalable zero-shot entity linking with dense entity retrieval. *arXiv preprint arXiv:1911.03814*.

Xu, C. and McAuley, J. (2023). A survey on dynamic neural networks for natural language processing. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 2370–2381, Dubrovnik, Croatia. Association for Computational Linguistics.

Yang, Y., Zhuang, Y., and Pan, Y. (2021). Multiple knowledge representation for big data artificial intelligence: framework, applications, and case studies. *Frontiers of Information Technology & Electronic Engineering*, 22(12):1551–1558.

Young, H., Du, M., and Bastani, O. (2022). Neurosymbolic deep generative models for sequence data with relational constraints. *Advances in Neural Information Processing Systems*, 35:37254–37266.

Yuan, X., Côté, M.-A., Sordoni, A., Laroche, R., Combes, R. T. d., Hausknecht, M., and Trischler, A. (2018). Counting to explore and generalize in text-based games. *arXiv preprint arXiv:1806.11525*.

Zhan, E., Sun, J. J., Kennedy, A., Yue, Y., and Chaudhuri, S. (2021). Unsupervised learning of neurosymbolic encoders. *arXiv preprint arXiv:2107.13132*.

Zhao, F., Wu, Z., and Dai, X. (2020). Attention transfer network for aspect-level sentiment classification. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 811–821, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Zhu, W., Thomason, J., and Jia, R. (2022). Generalization differences between end-to-end and neuro-symbolic vision-language reasoning systems. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4697–4711, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.