Characterizing Neuro-Symbolic Reasoning in NLP

Andrew Lyubovsky **Aman Madaan** **Yiming Yang**

Language Technologies Institute, Carnegie Mellon University
Pittsburgh, PA, USA
alyubovs@andrew.cmu.edu
{amadaan, yiming}@cs.cmu.edu

Abstract

Neuro-symbolic reasoning has recently seen a revival as a promising way to bridge the gap between deep learning models that manipulate continuous spaces and the symbolic representations. Despite its promise, the term "neuro-symbolic" remains nebulous and is used in various contexts, from systems that rely on having a rule-based sub-component to systems trained end-to-end using techniques like reinforcement learning. In this work, we conduct a survey of its various interpretations amongst researchers in the natural language processing (NLP) community along with a survey of prior work that focuses on neuro-symbolic methods. We posit that neuro-symbolic models are a host of techniques that allow dense continuous representations to be informed by more abstract discrete representations. Our work serves as a bird's eye view of this promising topic that opens avenues for future work.

1 Introduction

Deep Learning has achieved tremendous success in a wide range of problems including game-playing (Silver et al., 2017; Vinyals et al., 2019), language understanding (Chowdhery et al., 2022; Brown et al., 2020), visual reasoning (Zellers et al., 2022; Yang et al., 2021; Radford et al., 2021), generative modeling (Ramesh et al., 2021; Trinquier et al., 2021), and code generation (Chen et al., 2021; Li et al., 2022).

Despite these successes, several key concerns remain. First, deep learning systems lack explainability by design (Joshi et al., 2021; Danilevsky et al., 2020), making it difficult to contextualize and interpret their decisions. Second, deep learning systems tend to be data and compute hungry (Thompson et al., 2020), often requiring data and compute that is only feasible in large industrial settings. For example, GPT-3 (Brown et al., 2020) was trained on 45TB of data ($\approx$ 500 billion tokens). Third,

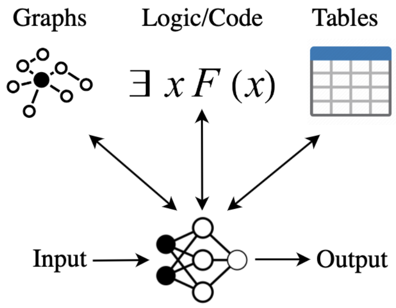

Figure 1: A characterization of neuro-symbolic methods: continuous neural representations are grounded in symbolic structures.

despite impressive performance on tasks that fundamentally require efficient pattern matching, even large models struggle in areas that need a human-like understanding of the world like commonsense reasoning (Marcus, 2021) or detecting toxic content (Hovy and Prabhumoye, 2021).

In the light of these shortcomings, there is now a growing interest in neuro-symbolic reasoning systems, which aim to endow deep learning systems with symbolic reasoning capabilities as a potential way to bridge these gaps (Garcez and Lamb, 2020; Sarker et al., 2021a; Hamilton et al., 2022; d'Avila Garcez et al., 2019; Raedt et al., 2019; Sarker et al., 2021b). A key promise of neuro-symbolic systems is grounding continuous neural representations in discrete symbolic representations, with the hope that neural networks leverage discrete representations for improved performance.

The popularity of neuro-symbolic reasoning has led to a large body of work that seeks to leverage neuro-symbolic techniques (Kramer, 2020; Sarker et al., 2021a). Additionally, as avenues of applying neuro-symbolic reasoning are explored, the exact techniques used to make a system neuro-symbolic are naturally diverging to adapt to new domains and problem settings. Such proliferation has made it challenging to precisely characterize what constitutes neuro-symbolic reasoning. For instance,

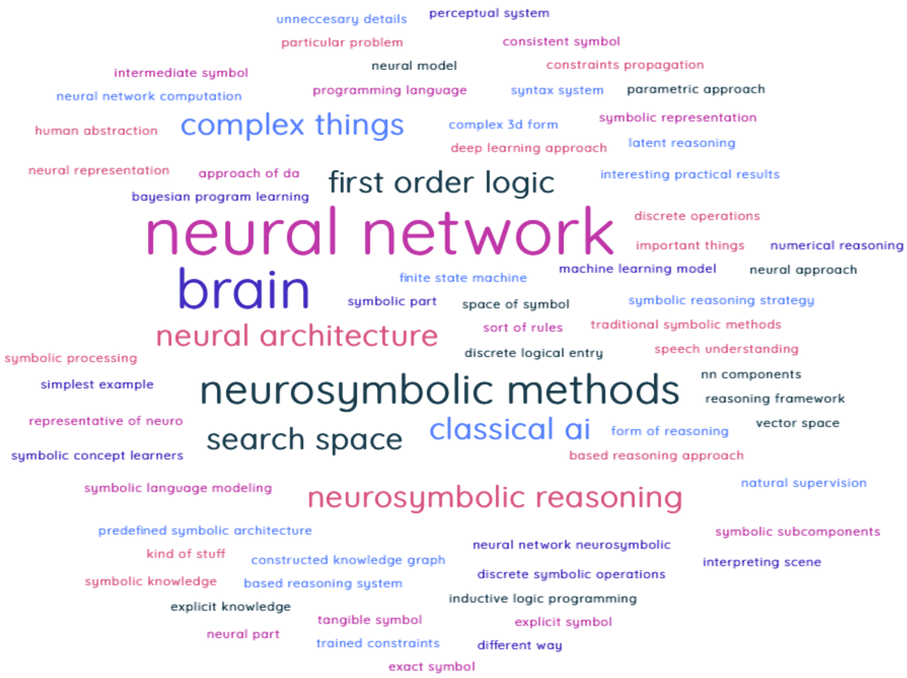

Figure 2: The word cloud is composed of words that were used by respondents when describing neuro-symbolic reasoning. This displays the contexts within which neuro-symbolic AI is understood.

neuro-symbolic techniques have been used in the context of controlling model output (Li et al., 2020; Chen et al., 2020c), leveraging external knowledge (Liang et al., 2018), enforcing certain constraints on the learning method (Demeter and Downey, 2020), and reasoning over graphs (Lamb et al., 2020; Marino et al., 2021). This is in stark contrast with other terms in the field like deep learning, which have a standard usage (i.e., systems that leverage stacked neural networks). The broad interpretation of neuro-symbolic reasoning as a way of combining symbolic representations with neural networks (d'Avila Garcez et al., 2019) is partly the cause, as this leaves room for a lot of interpretation.

In this paper, we take the first steps towards understanding the central theme in the large body of neuro-symbolic reasoning research. The primary goal of this work is not to come up with a strict definition of neuro-symbolic reasoning. Instead, we want to identify conditions that characterize a work as neuro-symbolic. We believe that such an operational definition can help practitioners and researchers better understand existing neuro-symbolic reasoning research, and identify future avenues.

Our approach relies on a survey of researchers and recent work in neuro-symbolic reasoning. Specifically, we conduct a survey of NLP practitioners at academic institutions and industrial labs and ask them to define neuro-symbolic reasoning. Additionally, we curate papers from the last 5 years that mention *neuro-symbolic reasoning* in the title and distill a central theme that unifies these works.

We find that while practitioners diverge in their understanding of neuro-symbolic reasoning, alignment of continuous representations learned by deep neural networks with the discrete world the networks aim to reason about is a key theme. From our survey of papers with neuro-symbolic reasoning in the title, we find neuro-symbolic reasoning approaches view symbolic representations as a way to constrain and guide continuous representations of deep models for solving a task. An overview of our findings from the human and literature survey is shown in Figure 1.

The contributions of this paper are as follows:

- We collect a survey about how NLP researchers interpret the term neuro-symbolic, and we survey the current neuro-symbolic approaches that have been applied to NLP tasks.
- We distill an operational definition of neuro-symbolic reasoning as a way to guide continuous representations in neural networks to leverage relationships that exist in abstract discrete representations.

2 Survey of Interpretations

To gauge the interpretation of the term "Neuro-symbolic", we surveyed 25 graduate students, professors, and research scientists working in the field of NLP, asking them to define the term neuro-symbolic reasoning. We also asked respondents to rate their familiarity with the term on a scale of one ("never heard of the term") and ten ("coined the term"). Figure 3 shows a histogram of the confidence ratings. All participation was voluntary, and participants were fairly compensated.

From the survey, we find multiple common themes amongst the responses that are valuable for characterizing how practitioners understand "Neuro-symbolic" methods. These are presented in Fig. 4, and are summarized next. All the anonymized responses are included in the Appendix.

2.1 Key findings

All respondents note that neuro-symbolic techniques use some combination of neural networks and symbolic representations. However, respondents that are more familiar with the term highlight other aspects. These include having symbolic operations applied over discrete symbols, such as functions. Others note that the symbols are a means of adding information to neural networks, and some note that symbols (such as logical rules) are used to constrain the outputs of the neural networks. Lastly, several responses suggest that neuro-symbolic reasoning is motivated by processes in the brain, specifically referring to the utilization of symbols in human cognitive processes. The existence of multiple hypotheses as to how this occurs contributes to a diverse set of interpretations for "Neuro-symbolic" methods. Specifically, we find the following trends amongst the responses.

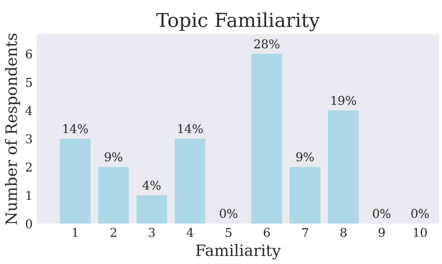

Figure 3: Self reported familiarity of respondents

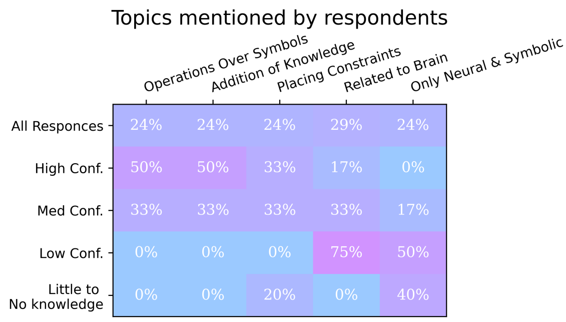

Figure 4: The table displays the percent of responses from each group that refer to different aspects of neuro-symbolic approaches, and responses could contain multiple aspect aspects. These responses suggest a broad understanding of what neuro-symbolic techniques entail.

Key aspects Multiple respondents identify the following as key aspects of neuro-symbolic reasoning:

- **Operations over symbols:** Neurosymbolic approaches often incorporate discrete operations such as mathematical operations or functions, enabling them to execute complicated operations without having to learn them.
- **Addition of knowledge:** Knowledge bases store symbolic information that is integrated with neural networks in some neuro-symbolic techniques.
- **Addition of Constraints:** Symbolic structures such as grammars have constraints that they need to fit. Neural networks can be modified to work in such constrained environments.

Symbolic representations From the survey, we observe that all responses note that neuro-symbolic techniques use some combination of neural networks and symbolic representations. We interpret symbolic representations to be any discrete representation, including natural language (Schuetze, 2016). However, typically symbolic representations refer to the representations that are more structured than text.

Anthropomorphism Several responses suggest that neuro-symbolic reasoning is motivated by processes in the brain or refer to *human thinking* as motivation, specifically to the concept of humans thinking with symbols. This hints that several experts believe symbolic transformation might be important for building a more robust AI.

Familiarity vs. knowledge As expected, the typical response shifts from being vague (*something to do with human-like thinking*) to a more specific (*introducing constraints, adding knowledge*). The exact distribution is displayed in Fig. 4. Note again that the respondents were experts (graduate students and professors) working in NLP.

The diversity of responses shows that there is a growing need to study the exact connotations of the term. We hope that our work takes the first step towards this goal. Towards this end, we surveyed papers that mention neuro-symbolic reasoning in the title.

3 Characterization

The responses from the survey in Section 2 indicate that neuro-symbolic reasoning approaches seek to infuse neural networks with symbolic information. To understand how this fusion is operationalized, we survey research associated with neuro-symbolic reasoning in the last five years and distill a unifying theme.

Sourcing papers Neuro-symbolic reasoning has a long history, with roots in logical programming (Newell and Simon, 1956) and the emergence of logical and algebraic programming languages for symbolic manipulation such as Prolog (Warren et al., 1977), Lisp (Winston and Horn, 1986), and MACSYMA (Martin and Fateman, 1971).

However, as we aim to understand neuro-symbolic reasoning approaches in the context of deep learning, we restrict ourselves to works that integrate neuro-symbolic reasoning with deep learning with a special focus on papers related to NLP tasks. We first identify all papers that include the term "symbolic" in the title and were published in ACL, NAACL, EMNLP, ICML, NeurIPS, ICLR, IJCAI, COLING, EACL, AACL, UAI, AISTATS, or ECML (popular machine learning and natural language processing conferences) between the years of 2016 and 2021. This collection resulted in 126 papers. Next, short papers, papers unrelated to NLP tasks, and papers that did not use neural networks were discarded. In the end, our literature survey focuses on 22 papers in depth and mainly deal with non-artificial data sets.

Summary of findings We find in our survey that while the exact implementation of neuro-symbolic reasoning works varies, a key element in these works is an infusion of more information about the world into the representations learned by deep neural networks. This alignment is achieved by a combination of data structures (e.g., Graph) and rules (e.g., programs). We summarize this in Definition 1.

Definition 1 *Neuro-symbolic methods are a family of approaches that align continuous representations learned by neural networks with the discrete symbols they reason about.*

Further, the distinction between *symbolic* and *neural* AI is not defined by a fixed set of rules. Instead, we find that *all* neural methods use both neural and symbolic representations; they manipulate discrete symbols (characters, words, text, BPE, audio signals, images, video frames, etc.) and are aimed at generating discrete symbols as the output. However, how these systems learn, and the degree to which these systems are aware of the **semantics of the symbols** differentiates them from each other. This difference is precisely the one we leverage in this work to compare and contrast recent works in neuro-symbolic reasoning.

3.1 Promises from Literature

From the literature review, we also observe that neuro-symbolic models elicit positive characteristics that address common limitations of deep learning models. These include the following properties.

- **Robustness** One limitation of deep learning is the difficulty of correcting mistake as these models need to be retrained with new data. By modifying knowledge bases that are integrated with neuro-symbolic models, the models can address unseen domains. (Verga et al., 2021; Yi et al., 2018; Arabshahi et al., 2021).

- **Explainability** Intermediate symbolic results in neuro-symbolic models can be observed by practitioners, and provide an explanation for a model's behavior, making them easier to work with (Bennetot et al., 2019; Stammer et al., 2021; Sarker et al., 2021b).

- **Compositionality** Neuro-symbolic models will often use an iterative approach, where they will make multiple passes, with access to previous outputs. This allows them to combine more basic outputs together to create a more complicated and meaningful solution (Chen et al., 2020c; Liang et al., 2017; Vedantam et al., 2019).

In the next three sections, we explore the different implementations of these approaches. Specifically, we arrange the literature by the type of canonical symbolic structure that is used: Graphs, Tables, and Logical rules.

4 Graphs

Graphs are a ubiquitous data structure that store information about nodes and the connections between them. In general, a graph can be interpreted as a store of (*object 1*, *relation*, *object 2*) triplets, where an edge of type *relation* connects nodes representing *object 1* and *object 2* (Lamb et al., 2020). They are typically used in neuro-symbolic systems to better encode exact relationships between objects or events, and can be leveraged to obtain better representations across all of the nodes.

In the simplest form, neuro-symbolic approaches utilize graphs as a collection of these fact triplets while more advanced applications exploit the rich structure of graphs and can construct chains from such triplets. By directly retrieving triplets from a knowledge base, neuro-symbolic approaches can leverage discrete relationships encoded within the triplets without having to learn them. Verga et al. (2021) use such an approach in order to improve question answering. They use a language model that queries a knowledge base based on a question by identifying a *subject* and *relationship* that the question asks about. For example, if asked "Where was Charles Darwin born?", they would query the table with <Charles Darwin, born_in>, and find <Charles Darwin, born_in, England> stored in the knowledge base. The question is projected into a joint vector space with the <*subject*, *relationship*> pairs from the knowledge base triplets, where the selected pair is used to answer the question. This allows a model to be more robust to new domains as knowledge base triplets from other domains can be added to the knowledge base without retraining the model, and it leverages them to answer questions. A similar approach is used by Moghimifar et al. (2021) to retrieve knowledge base triplets, while Ma et al. (2019) attend to a set of retrieved knowledge triples (and not just a single triple) to improve the task of question answering. Finally, Liang et al. (2017) generate code that processes a set of triplets.

Another set of approaches operates on a chain of such tuples, or equivalently, a path from the graph. This allows for models to combine multiple smaller relationships within a knowledge base together. As a representative approach, Kang et al. (2018) use a knowledge base for the task of scientific entailment (verifying if a fact is supported by a knowledge base). Since multiple scientific facts entail information that is not explicitly stated, they first break down scientific hypothesis into sub-facts that must be true to support the hypothesis. Then, they evaluate each sub-fact corresponding to the knowledge base to verify its validity, retrieving triples based on three similarity measures. They include a neural entailment module that computes whether a sub-fact is entailed in a knowledge base triplet using attention. This allows them to evaluate whether a knowledge base supports a hypothesis by creating a chain of supporting triplets. Arabshahi et al. (2021) uses a similar approach where they store logical rules in a knowledge base and back-chain them using an LSTM to detect what assumptions are made by a statement (natural language inference) (Ebrahimi et al., 2021; Mota and Diniz, 2016).

Complete graph structures are also used in neuro-symbolic approaches, typically using *Graph Neural Networks* or GNNs (Lamb et al., 2020; Alshahrani et al., 2017; Cranmer et al., 2020; Dumancic et al., 2019). GNNs learn vector representations for nodes and links between the nodes. Marino et al. (2021) incorporate graph neural networks for visual question answering tasks. They extract objects in the image and a question and use a knowledge base to form links between the extracted objects, forming a graph. After applying a graph neural network on the graph, they then apply logistic regression to the node representations to find a continuous representation of the graph used to arrive at a final answer to the question (Kirk and Laird, 2019; Yilmaz et al., 2016). Shi et al. (2020) use a similar approach, using Graph attention networks in an intermediate representation for the task of table-based question answering. Similarly, Liang et al. (2018) use image regions to find weights for nodes. Then they form a graph using the nodes and a knowledge base, and improve node representations based on neighboring nodes, leading to better representations of the original features. The symbolic representations allow neighboring nodes to improve the individual node representations, improving the original image's representation.

Finally, some approaches use neural networks to create graphs. Here the neural networks are trained on a knowledge base, and when given a node, they predict the following link and node within a graph. For example, Hwang et al. (2021) address the task of commonsense question answering by generating a graph based on the context given in a question. First, they train a neural network (COMeT) on a knowledge graph (Atomic) that stores triples. To answer questions, they feed the provided situational context as an entity along with the question, which becomes the relation. This process is continued to extend the existing graph. Lastly, they compare the possible answers with the generated nodes to select an answer. Feblowitz et al. (2021) propose a similar approach using a language model that is trained on previous business scenarios to generate a graph of implications from a particular business scenario. This allows neural networks to be grounded in symbolic triplets, thus leveraging the ontology stored within the symbolic representations.

5 Tables

Tables store discrete representations of objects in a structured manner, allowing their use as intermediate symbolic representations. Further, discrete operations can be applied over the tables, allowing for consistent and explainable computations. They are useful when handling multiple objects with similar and task-relevant characteristics. While tables have less defined relationships than graphs, they typically specify object properties more explicitly, allowing for neural components to learn their semantics.

Tables are not only used as information stores but are also used as an abstraction to add structure to input for efficient reasoning. For instance, Yi et al. (2018) extract a table of 3d shapes from an image for visual question answering. An encoded question is used to generate SQL-like code over the extracted table. The SQL code (which might contain operations such as filtering and counting), when executed, derives the answer from the Table. Other neuro-symbolic approaches also ground visual representations in tabular structures (Jiang and Ahn, 2020; Alirezaie et al., 2018). Mou et al. (2017) explore a more conventional use of tables as entity stores and encode the rows and columns to be used in a neural network that retrieves cells (properties of entities) directly from the table. Similarly, Shi et al. (2020) use tables extracted from Wikipedia (TABFACT) for fact verification, and linearize them before feeding to BERT (Chen et al., 2020a).

Zellers et al. (2021) propose a nontraditional application of using tables for the task of commonsense question answering and text generation. They use tables to store a representation of the environment at different time steps, containing objects and their characteristics such as whether a stove is hot. The table also stores labeled actions that were applied to the environment and caused the environment to transition from one state to another. A neural network is then trained to learn how actions cause transitions within the table. This neural network is then embedded into a language model for a downstream task, allowing for the symbolic transitions to be leveraged by the language model.

6 Logic Rules

Functions and logic are symbols that represent discrete operations and discrete relationships. Grounding neural representations in these symbols allow them to organize the symbolic manipulations, allowing them to leverage more abstract concepts (Belle, 2020). For example, first-order logic can restrict the output of a neural network to a specified grammar using Grammar models, which can then be evaluated as a symbolic expression (Li et al., 2020). Programs can also be generated to parse over external symbolic structures such as a knowledge base or text. Lastly, some approaches learn functions in the form of neural networks, breaking down the problem into simpler sub-problems. These approaches are typically used in neuro-symbolic systems when a task relies on a set of known procedures or there exists a set of constraints (e.g., grammars) within a system (Belle, 2020). By combining multiple more basic logic operations, such systems can create compositoinal results.

First order logic expressions combine language expressions with a structured syntax that can be manipulated by a deductive system while retaining their semantics. An example is presented in Table 1. When used in neuro-symbolic systems, they can be add structure to complex data when extracted from images and texts (Dai and Muggleton, 2020; Tsamoura and Michael, 2020). For the task of visual question answering, Amizadeh et al. (2020) extract object visual features from an image, and use them to determine the objects, their attributes, and relationships that are then encoded into a first-order logic statement. After parsing a question into a first-order logic statement, they

FOL Formula	$F(X,Y) = \exists X, \forall Y : Sky(X) \wedge Higher(X,Y)$
Text Explanation	There exists a sky above all objects

Table 1: This is an example of a first-order logic expression. More specifically, the first-order logic formula says that there exists an X, such that for all Y, X is a sky, and X is above Y (Amizadeh et al., 2020).

then evaluate the two statements together to find an answer. Similarly, Stammer et al. (2021) generate first-order logical explanations based on neural activations, improving the explainability of visual question answering models. Alternatively, Elder et al. (2019) improve the table-to-text task by generating an intermediate first-order logic representation, and Arabshahi et al. (2021) create a theorem prover by restricting outputs from a neural network to logic statements stored in a knowledge base and chain them together to create a proof (Rocktäschel and Riedel, 2016). Xiao et al. (2017) address semantic parsing by restricting RNN outputs to those consistent with a context-free grammar. Other similar approaches exist aiming to generate formulas and other high-level representations to fit noisy data (Udrescu et al., 2020; Kramer, 2020; d'Ascoli et al., 2022).

Generated programs are also often used to process symbolic information. To address the task of reading comprehension (answering questions about a given text), Chen et al. (2020c) generate a program that answers a question by processing a text. They use an encoder-decoder model to convert a passage and question into a program that can be applied to a text passage to obtain the answer. This makes the model more compositional, answering more complicated questions about the passage. While this approach generates a program that is applied to a passage, it is more common to parse a table using the program (Parisotto et al., 2017). Shi et al. (2020) generate a lisp-like expression based on a query that gets executed against a table to verify whether the query is correct, and when answering questions about a table Mou et al. (2017) generate a program that selects a cell within a table to answer the question. Similarly, Mao et al. (2019) and Yi et al. (2018) address the visual question answering task by extracting tabular information from images and generating programs from the questions to parse those (Li et al., 2020). Lastly, Liang et al. (2017) generate a program that manipulates knowledge base triplets and stores intermediate results to improve on compositionality of the programs. Other tasks where program generation is used include game playing, instruction generation, planning, and image generation (Asai and Fukunaga, 2018; Asai and Muise, 2020; Chen et al., 2020b; Andersen and Konidaris, 2017; Feinman and Lake, 2021).

Lastly, other approaches aim to simplify the complexity that a neural network needs to learn by breaking down a task into sub-tasks that are solved by separate modules. For example, to improve the task of text generation, Demeter and Downey (2020) create classes of words, where the primary language model generates class tokens in place of cities and dates, and other micro models replace the class tokens with the exact cities and dates that are correct for the context. The language model does not need to learn the relationships associated with specific classes as they can be noisy. Instead, a separate micro model can learn them without learning other aspects of the modeling language. When addressing the task of visual question answering, Vedantam et al. (2019) also create separate neural modules that can solve sub-tasks. Based on a question, these modules are arranged in a manner where outputs from one module are sent into another, as functions within a program. This way, the different modules learn specific processes used when answering questions about images and are not affected by other sub-tasks. Other approaches also integrate preexisting code, simplifying the task that neural components need to learn and often use reinforcement learning (Anderson et al., 2020; Inala et al., 2020; Le-Phuoc et al., 2021).

7 Discussion

Limitations Our approach contains limitations that allow us to deconstruct the different aspects in neuro-symbolic reasoning into a simple characterization. These include passing over technical implementation details, clustering topics together, evading a concrete definition that focuses on a specific interpretation of the term, and excluding the cognitive science motivation for the topic. To retain a simple characterization, we overlook some nontrivial aspects to the approaches' success. For example, some approaches include novel ways of either storing symbolic information or allowing gradients to backpropagate between symbolic and neural components. While it is necessary to understand these details, they are often model-specific and require a comprehensive description of the model to be fully appreciated. Secondly, we do not distinguish between topics when characterizing neuro-symbolic models. While these models differ amongst domains, we do not explore this as we aim to characterize neuro-symbolic models as a whole. Also, while we focus our literature review on models related to more recent NLP tasks, we find that the extracted characterization carries over to all models. Next, some practitioners argue that our characterization does not pinpoint the key promise in neuro-symbolic approaches. However, we intentionally keep the characterization broad as we find that there is a variety of interpretations for neuro-symbolic methods, and we aim to create a description that applies to all of them. Lastly, we find that a key contributor to the popularity of neuro-symbolic methods is their consistency with symbolic manipulation in the brain. While we acknowledge this, we do not focus on this aspect in our characterization as we find that the literature discusses the benefits of these approaches that do not stem from the neuroscience parallel (Srikant and O'Reilly, 2021).

Conclusion In this work, we observe a broad interpretation of what neuro-symbolic approaches entail, aiming to better characterize the term. We first conduct a survey of experts to understand how practitioners interpret the term, and then survey papers to better understand how these approaches are implemented. While all of these approaches aim to embed abstract information from structured symbolic data into neural networks, practitioners interpret neuro-symbolic AI in a more specific scope that varies amongst them. In a research landscape dominated by extremely large language models, neuro-symbolic reasoning approaches might be the key to developing effective and efficient learning machines.

Acknowledgments

This material is partly based on research sponsored in part by the Air Force Research Laboratory under agreement number FA8750-19-2-0200. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory or the U.S. Government. We also thank all of the survey responders for taking the time and effort to complete the survey.

References

Marjan Alirezaie, Martin Längkvist, Michael Sioutis, and Amy Loutfi. 2018. Position paper: Reasoning upon learning: A generic neural-symbolic approach. In Thirteenth International Workshop on Neural-Symbolic Learning and Reasoning (NeSy).

Mona Alshahrani, Mohammad Asif Khan, Omar Maddouri, Akira R Kinjo, Núria Queralt-Rosinach, and Robert Hoehndorf. 2017. Neuro-symbolic representation learning on biological knowledge graphs. Bioinformatics, 33(17):2723–2730.

Saeed Amizadeh, Hamid Palangi, Alex Polozov, Yichen Huang, and Kazuhito Koishida. 2020. Neuro-symbolic visual reasoning: Disentangling. In International Conference on Machine Learning, pages 279–290. PMLR.

Garrett Andersen and George Dimitri Konidaris. 2017. Active exploration for learning symbolic representations. In Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA, pages 5009–5019.

Greg Anderson, Abhinav Verma, Isil Dillig, and Swarat Chaudhuri. 2020. Neurosymbolic reinforcement learning with formally verified exploration. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual.

Forough Arabshahi, Jennifer Lee, Mikayla Gawarecki, Kathryn Mazaitis, Amos Azaria, and Tom. Mitchell. 2021. Conversational neuro-symbolic commonsense reasoning. In AAAI.

Masataro Asai and Alex Fukunaga. 2018. Classical planning in deep latent space: Bridging the subsymbolic-symbolic boundary. In Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018, pages 6094–6101. AAAI Press.

Masataro Asai and Christian Muise. 2020. Learning neural-symbolic descriptive planning models via cube-space priors: The voyage home (to STRIPS). In Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020, pages 2676–2682. ijcai.org.

Vaishak Belle. 2020. Symbolic logic meets machine learning: A brief survey in infinite domains. ArXiv preprint, abs/2006.08480.

Adrien Bennetot, Jean Luc Laurent, Raja Chatila, and Natalia Díaz-Rodríguez. 2019. Towards explainable neural-symbolic visual reasoning. arXiv: Learning.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. Evaluating large language models trained on code.

Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyou Zhou, and William Yang Wang. 2020a. Tabfact : A large-scale dataset for table-based fact verification. In International Conference on Learning Representations (ICLR), Addis Ababa, Ethiopia.

Xinyun Chen, Chen Liang, Adams Wei Yu, Dawn Song, and Denny Zhou. 2020b. Compositional generalization via neural-symbolic stack machines. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual.

Xinyun Chen, Chen Liang, Adams Wei Yu, Denny Zhou, Dawn Song, and Quoc V. Le. 2020c. Neural symbolic reader: Scalable integration of distributed and symbolic representations for reading comprehension. In 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020. OpenReview.net.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. ArXiv preprint, abs/2204.02311.

Miles D. Cranmer, Alvaro Sanchez-Gonzalez, Peter W. Battaglia, Rui Xu, Kyle Cranmer, David N. Spergel, and Shirley Ho. 2020. Discovering symbolic models from deep learning with inductive biases. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual.

Wang-Zhou Dai and Stephen H. Muggleton. 2020. Abductive knowledge induction from raw data. ArXiv preprint, abs/2010.03514.

Marina Danilevsky, Kun Qian, Ranit Aharonov, Yannis Katsis, Ban Kawas, and Prithviraj Sen. 2020. A survey of the state of explainable AI for natural language processing. In Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing, pages 447–459, Suzhou, China. Association for Computational Linguistics.

Stéphane d'Ascoli, Pierre-Alexandre Kamienny, Guillaume Lample, and François Charton. 2022. Deep Symbolic Regression for Recurrent Sequences. ArXiv preprint, abs/2201.04600.

Artur S. d'Avila Garcez, Marco Gori, Luís C. Lamb, Luciano Serafini, Michael Spranger, and Son N. Tran. 2019. Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning. ArXiv preprint, abs/1905.06088.

David Demeter and Doug Downey. 2020. Just add functions: A neural-symbolic language model. In The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020, pages 7634–7642. AAAI Press.

Sebastijan Dumancic, Alberto García-Durán, and Mathias Niepert. 2019. A comparative study of distributional and symbolic paradigms for relational learning. In Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019, pages 6088–6094. ijcai.org.

Monireh Ebrahimi, Md Kamruzzaman Sarker, Federico Bianchi, Ning Xie, Aaron Eberhart, Derek Doran, HyeongSik Kim, and Pascal Hitzler. 2021. Neuro-symbolic deductive reasoning for cross-knowledge graph entailment. In AAAI Spring Symposium: Combining Machine Learning with Knowledge Engineering.

Henry Elder, Jennifer Foster, James Barry, and Alexander O'Connor. 2019. Designing a symbolic intermediate representation for neural surface realization. In Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation, pages 65–73, Minneapolis, Minnesota. Association for Computational Linguistics.

Mark D. Feblowitz, Oktie Hassanzadeh, Michael Katz, Shirin Sohrabi, Kavitha Srinivas, and O. Udrea. 2021. Ibm scenario planning advisor: A neuro-symbolic erm solution. In AAAI.

Reuben Feinman and Brenden M. Lake. 2021. Learning task-general representations with generative neuro-symbolic modeling. In 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net.

Artur d'Avila Garcez and Luis C Lamb. 2020. Neurosymbolic ai: the 3rd wave. ArXiv preprint, abs/2012.05876.

Kyle Hamilton, Aparna Nayak, Bojan Božić, and Luca Longo. 2022. Is neuro-symbolic ai meeting its promise in natural language processing? a structured review. ArXiv preprint, abs/2202.12205.

Dirk Hovy and Shrimai Prabhumoye. 2021. Five sources of bias in natural language processing. Language and Linguistics Compass, 15(8):e12432.

Jena D. Hwang, Chandra Bhagavatula, Ronan Le Bras, Jeff Da, Keisuke Sakaguchi, Antoine Bosselut, and Yejin Choi. 2021. Comet-atomic 2020: On symbolic and neural commonsense knowledge graphs. In AAAI.

Jeevana Priya Inala, Yichen Yang, James Paulos, Yewen Pu, Osbert Bastani, Vijay Kumar, Martin Rinard, and Armando Solar-Lezama. 2020. Neurosymbolic transformers for multi-agent communication. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual.

Jindong Jiang and Sungjin Ahn. 2020. Generative neurosymbolic machines. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual.

Gargi Joshi, Rahee Walambe, and Ketan Kotecha. 2021. A review on explainability in multimodal deep neural nets. IEEE Access.

Dongyeop Kang, Tushar Khot, Ashish Sabharwal, and Peter Clark. 2018. Bridging knowledge gaps in neural entailment via symbolic models. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pages 4940–4945, Brussels, Belgium. Association for Computational Linguistics.

James R. Kirk and John E. Laird. 2019. Learning hierarchical symbolic representations to support interactive task learning and knowledge transfer. In Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019, pages 6095–6102. ijcai.org.

Stefan Kramer. 2020. A brief history of learning symbolic higher-level representations from data (and a curious look forward). In Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020, pages 4868–4876. ijcai.org.

Luís C. Lamb, Artur S. d'Avila Garcez, Marco Gori, Marcelo O. R. Prates, Pedro H. C. Avelar, and Moshe Y. Vardi. 2020. Graph neural networks meet neural-symbolic computing: A survey and perspective. In Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020, pages 4877–4884. ijcai.org.

Danh Le-Phuoc, Thomas Eiter, and Anh Le-Tuan. 2021. A scalable reasoning and learning approach for neural-symbolic stream fusion. AAAI'21.

Qing Li, Siyuan Huang, Yining Hong, Yixin Chen, Ying Nian Wu, and Song-Chun Zhu. 2020. Closed loop neural-symbolic learning via integrating neural perception, grammar parsing, and symbolic reasoning. In Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event, volume 119 of Proceedings of Machine Learning Research, pages 5884–5894. PMLR.

Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, et al. 2022. Competition-level code generation with alphacode. ArXiv preprint, abs/2203.07814.

Chen Liang, Jonathan Berant, Quoc Le, Kenneth D. Forbus, and Ni Lao. 2017. Neural symbolic machines: Learning semantic parsers on Freebase with weak supervision. In Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 23–33, Vancouver, Canada. Association for Computational Linguistics.

Xiaodan Liang, Zhiting Hu, Hao Zhang, Liang Lin, and Eric P. Xing. 2018. Symbolic graph reasoning meets convolutions. In Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada, pages 1858–1868.

Kaixin Ma, Jonathan Francis, Quanyang Lu, Eric Nyberg, and Alessandro Oltramari. 2019. Towards generalizable neuro-symbolic systems for commonsense question answering. In Proceedings of the First Workshop on Commonsense Inference in Natural Language Processing, pages 22–32, Hong Kong, China. Association for Computational Linguistics.

Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B. Tenenbaum, and Jiajun Wu. 2019. The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. In 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net.

Gary Marcus. Experiments testing gpt-3's ability at commonsense reasoning: results. [online]. 2021.

Kenneth Marino, Xinlei Chen, Devi Parikh, Abhinav Kumar Gupta, and Marcus Rohrbach. 2021. Krisp: Integrating implicit and symbolic knowledge for open-domain knowledge-based vqa. 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 14106–14116.

William A Martin and Richard J Fateman. 1971. The macsyma system. In Proceedings of the second ACM symposium on Symbolic and algebraic manipulation, pages 59–75.

Farhad Moghimifar, Lizhen Qu, Terry Yue Zhuo, Gholamreza Haffari, and Mahsa Baktashmotlagh. 2021. Neural-symbolic commonsense reasoner with relation predictors. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers), pages 797–802, Online. Association for Computational Linguistics.

Edjard de S Mota and Yan B Diniz. 2016. Shared multi-space representation for neural-symbolic reasoning. NeSy, 1768:295–318.

Lili Mou, Zhengdong Lu, Hang Li, and Zhi Jin. 2017. Coupling distributed and symbolic execution for natural language queries. In Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017, volume 70 of Proceedings of Machine Learning Research, pages 2518–2526. PMLR.

Allen Newell and Herbert Simon. 1956. The logic theory machine–a complex information processing system. IRE Transactions on information theory, 2(3):61–79.

Emilio Parisotto, Abdel-rahman Mohamed, Rishabh Singh, Lihong Li, Dengyong Zhou, and Pushmeet Kohli. 2017. Neuro-symbolic program synthesis. In 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings. OpenReview.net.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event, volume 139 of Proceedings of Machine Learning Research, pages 8748–8763. PMLR.

Luc De Raedt, Robin Manhaeve, Sebastijan Dumancic, Thomas Demeester, and Angelika Kimmig. 2019. Neuro-symbolic = neural + logical + probabilistic.

Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-shot text-to-image generation. In Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event, volume 139 of Proceedings of Machine Learning Research, pages 8821–8831. PMLR.

Tim Rocktäschel and Sebastian Riedel. 2016. Learning knowledge base inference with neural theorem provers. In Proceedings of the 5th Workshop on Automated Knowledge Base Construction, pages 45–50, San Diego, CA. Association for Computational Linguistics.

Md. Kamruzzaman Sarker, Lu Zhou, Aaron Eberhart, and Pascal Hitzler. 2021a. Neuro-symbolic artificial intelligence: Current trends. ArXiv preprint, abs/2105.05330.

Md. Kamruzzaman Sarker, Lu Zhou, Aaron Eberhart, and Pascal Hitzler. 2021b. Neuro-symbolic artificial intelligence: Current trends. ArXiv preprint, abs/2105.05330.

Hinrich Schuetze. 2016. Nonsymbolic text representation.

Qi Shi, Yu Zhang, Qingyu Yin, and Ting Liu. 2020. Learn to combine linguistic and symbolic information for table-based fact verification. In Proceedings of the 28th International Conference on Computational Linguistics, pages 5335–5346, Barcelona, Spain (Online). International Committee on Computational Linguistics.

David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. 2017. Mastering the game of go without human knowledge. nature, 550(7676):354–359.

Shashank Srikant and Una-May O'Reilly. 2021. Can cognitive neuroscience inform neuro-symbolic inference models? In Is Neuro-Symbolic SOTA still a myth for Natural Language Inference? The first workshop.

Wolfgang Stammer, Patrick Schramowski, and Kristian Kersting. 2021. Right for the right concept: Revising neuro-symbolic concepts by interacting with their explanations. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 3619–3629.

Neil C Thompson, Kristjan Greenewald, Keeheon Lee, and Gabriel F Manso. 2020. The computational limits of deep learning. ArXiv preprint, abs/2007.05558.

Jeanne Trinquier, Guido Uguzzoni, Andrea Pagnani, Francesco Zamponi, and Martin Weigt. 2021. Efficient generative modeling of protein sequences using simple autoregressive models. Nature communications, 12(1):1–11.

Efthymia Tsamoura and Loizos Michael. 2020. Neural-symbolic integration: A compositional perspective. ArXiv preprint, abs/2010.11926.

Silviu-Marian Udrescu, Andrew Tan, Jiahai Feng, Orisvaldo Neto, Tailin Wu, and Max Tegmark. 2020. AI feynman 2.0: Pareto-optimal symbolic regression exploiting graph modularity. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual.

Ramakrishna Vedantam, Karan Desai, Stefan Lee, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. 2019. Probabilistic neural symbolic models for interpretable visual question answering. In Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA, volume 97 of Proceedings of Machine Learning Research, pages 6428–6437. PMLR.

Pat Verga, Haitian Sun, Livio Baldini Soares, and William Cohen. 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, Online. Association for Computational Linguistics.

Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. 2019. Grandmaster level in starcraft ii using multi-agent reinforcement learning. Nature, 575(7782):350–354.

David HD Warren, Luis M Pereira, and Fernando Pereira. 1977. Prolog-the language and its implementation compared with lisp. ACM SIGPLAN Notices, 12(8):109–115.

Patrick Henry Winston and Berthold K Horn. 1986. Lisp.

Chunyang Xiao, Marc Dymetman, and Claire Gardent. 2017. Symbolic priors for rnn-based semantic parsing. In Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017, pages 4186–4192. ijcai.org.

Zhengyuan Yang, Yijuan Lu, Jianfeng Wang, Xi Yin, Dinei Florencio, Lijuan Wang, Cha Zhang, Lei Zhang, and Jiebo Luo. 2021. Tap: Text-aware pre-training for text-vqa and text-caption. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8751–8761.

Kexin Yi, Jiajun Wu, Chuang Gan, Antonio Torralba, Pushmeet Kohli, and Josh Tenenbaum. 2018. Neural-symbolic VQA: disentangling reasoning from vision and language understanding. In Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada, pages 1039–1050.

Özgür Yilmaz, Artur S d'Avila Garcez, and Daniel L Silver. 2016. A proposal for common dataset in neural-symbolic reasoning studies. In NeSy@ HLAI.

Rowan Zellers, Ari Holtzman, Matthew Peters, Roozbeh Mottaghi, Aniruddha Kembhavi, Ali Farhadi, and Yejin Choi. 2021. PIGLeT: Language grounding through neuro-symbolic interaction in a 3D world. 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 2040–2050, Online. Association for Computational Linguistics.

Rowan Zellers, Jiasen Lu, Ximing Lu, Youngjae Yu, Yanpeng Zhao, Mohammadreza Salehi, Aditya Kusupati, Jack Hessel, Ali Farhadi, and Yejin Choi. 2022. Merlot reserve: Neural script knowledge through vision and language and sound. ArXiv preprint, abs/2201.02639.

A Survey Responses

Anonymized survey responses are included in Tables 2 and 3.

Response	Familiarity
In my opinion: *anything* that combines a neural model and a "program" in a synergy that makes the combined model better than any of the two components separately. The "program" can be in any programming language, domain-specific language, or another computational structure such as an automaton.	8
I have heard the term. I interpret it to mean either: 1) neural approaches to creating, interacting with, processing, or reasoning over symbolic structures and representations; or 2) using symbolic reasoning strategies with functions learned via neural networks.	4
Combining knowledge from text (typically via the parameters of a neural network trained on a task on text) and symbolic knowledge (typically a pre-constructed knowledge graph)	6
I understand it as referring to neural architectures that integrate some type of symbolic component (like inductive logic programming or even finite state machines).	6
I interpret it as meaning the imposition of non-trained constraints into a model, whether through the structure of the representation the model is forced to use (e.g. graphs, where nodes represent explicit symbols), or more explicit constraints, like using ILP in models. Another important thing is that the constraints have to map onto the idea of an actual symbol, and in my experience, that tends to be related to the human abstraction that's in place for a particular problem.	6
I interpret it as neural/latent reasoning, grounded with fundamental and intermediate symbols/rules which allow a model/algorithm to systematically reason through these constraints/rules to come to an answer. I think we as humans do possess the ability to reason neurosymbolically, by connecting our latent thought to tangible, or well-defined concepts in the real-world. E.g. we learn a grammar and syntax system in order to communicate. Words can be defined as tangible symbols, which when combined appropriately, can communicate a thought (which is intangible).	6
Neurosymbolic reasoning is essentially a semi-parametric approach, wherein you identify existing discrete logical entries or entities and then perform some form of reasoning on top of that, the reasoning framework being a neural architecture. The symbolic part comes from recognizing the exact symbols or entities and how they are related, or could be related to one another based on a given ontology or ruleset. The neural part helps in generalization.	7
I would like to interpret it as integrating (first-order) logic into the neural network. I think it is logic that makes neural networks Neurosymbolic.	8
To me, a neurosymbolic method means one that combines a deep learning approach with a more traditional symbolic method. There can be different ways to combine them. For example, use a predefined symbolic architecture and then use deep learning to learn representations for the symbols	8
I have never heard the term before. But as I interpret it, it would involve the class of learning architectures and approaches which combines neural subcomponents with symbolic subcomponents.	1
I have heard of it. Don't exactly know what it means but if I had to guess, I would say it refers to a sort of rule-based reasoning system that uses some sort of pre-neural representations (inspired by how brains work). I honestly have no clue.	3
Typically refers to architectures/systems that combine any of the neural-based architectures (NNs, (bi-)LSTMS, autoencoders, etc.) with some source of symbolic processing (e.g. constraint propagation, numerical reasoning) or data (e.g. hand-built ontologies)	8

Table 2: Survey responses

Response	Familiarity
I have heard of it but am not too familiar on what exactly the definition is. If I had to guess, I would assume that it has something to do with the fact that our brains group complex things into abstract "symbols" (simplified representations corresponding with more complex things in the real world). This enables our brains to reason about things without having to pay attention to unneccesary detail, which I don't think current machine learning models are really capable of yet.	4
Haven't heard the term	1
My understanding is that if you combine the neural network computation with discrete symbolic operations together, that would be considered neuro-symbolic. Loosely speaking, if your system has a NN component, and discrete operations like arithmetic, logical and/or, that should also count.	7
If you can do something in under a second, then machines can do it almost as good (this is called perception). If it takes you more than 1 second to something, then machines cannot do it as well because it often involves complex reasoning (this is called cognition or reasoning). It is difficult to reason with interpretable, consistent symbols exclusively in the vector space. On the other hand, the search space of symbols becomes complex and difficult to deal exclusively using symbols with no semantic representation. Neurosymbolic methods combine the best of these worlds by narrowing the search space and learning the mapping between symbols and neural decodings. A primitive and simplest example of a neurosymbolic method could be when beam search is controlled with symbolic rules.	8
Something interpretable and grounded in explicit knowledge	6
My interpretation is that it is a combination of classical AI (like first order logic, knowledge graphs, context free grammar, that kind of stuff) and neural networks. For example, neural networks might be used to model operators and transformations used in classical AI.	2
I have heard it in the setting of AGI for speech understanding on the phonemic level	4
I do not know.	1
My understanding of the term is using first order logic to make neural nets more understandable.	2

Table 3: Survey responses

Tables	Mou et al. (2017),Shi et al. (2020),Yi et al. (2018), Zellers et al. (2021)
Code/Logic	Mou et al. (2017), Shi et al. (2020), Yi et al. (2018),Liang et al. (2017), Demeter and Downey (2020), Arabshahi et al. (2021), Chen et al. (2020c), Amizadeh et al. (2020), Stammer et al. (2021),Mao et al. (2019), Vedantam et al. (2019), Elder et al. (2019), Xiao et al. (2017)
Graphs	Mou et al. (2017),Shi et al. (2020),Verga et al. (2021), Liang et al. (2017), Moghimifar et al. (2021), Kang et al. (2018), Arabshahi et al. (2021), Marino et al. (2021),Ma et al. (2019),Liang et al. (2018), Hwang et al. (2021), Feblowitz et al. (2021)