

Neuro-Symbolic Frameworks: Conceptual Characterization and Empirical Comparative Analysis

Sania Sinha

SINHASA3@MSU.EDU

Tanawan Premsri

PREMSRIT@MSU.EDU

Danial Kamali

KAMALIDA@MSU.EDU

Parisa Kordjamshidi

KORDJAMS@MSU.EDU

Department of Computer Science and Engineering, Michigan State University

Editors: Leilani H. Gilpin, Eleonora Giunchiglia, Pascal Hitzler, and Emile van Krieken

Abstract

Neurosymbolic (NeSy) frameworks combine neural representations and learning with symbolic representations and reasoning. Combining the reasoning capacities, explainability, and interpretability of symbolic processing with the flexibility and power of neural computing allows us to solve complex problems with more reliability while being data-efficient. However, this recently growing topic poses a challenge to developers with its learning curve, lack of user-friendly tools, libraries, and unifying frameworks. In this paper, we characterize the technical facets of existing NeSy frameworks, such as the symbolic representation language, integration with neural models, and the underlying algorithms. A majority of the NeSy research focuses on algorithms instead of providing generic frameworks for declarative problem specification to leverage problem solving. To highlight the key aspects of Neurosymbolic modeling, we showcase three generic NeSy frameworks - *DeepProbLog*, *Scallop*, and *DomiKnowS*. We identify the challenges within each facet that lay the foundation for identifying the expressivity of each framework in solving a variety of problems. Building on this foundation, we aim to drive transformative action and encourage the community to rethink this problem in novel ways.

Keywords: Neurosymbolic, Comparing NeSy frameworks, DomiKnowS, DeepProbLog, Scallop, Combining learning and reasoning

1. Introduction

Symbolic or *good old-fashioned AI* focused on creating rule-based reasoning systems (Hayes-Roth, 1985) exemplified by early works such as the Physical Symbol System (Augusto, 2021; Newell, 1980) and ELIZA (Weizenbaum, 1966). However, drawbacks such as limited scalability due to the need to explicitly define domain predicates and rules for each task, lack of robustness in handling messy real-world data, and low computational efficiency led to a decline in the popularity of this paradigm, shifting the focus toward neural computing and deep learning. **Deep Learning** (LeCun et al., 2015; Ahmad et al., 2019) revolutionized AI, as nuanced relationships in data could be learned by backpropagation through multiple layers of processing and creating abstract representations of data. However, it led to a loss of explainability (Li et al., 2023a), dependence on large amounts of data, lack of generalizability for unobserved situations, and rising concerns about its environmental sustainability (Bender et al., 2021). **Neurosymbolic AI** (Hitzler and Sarker, 2022; Bhuyan et al., 2024), a combination of symbolic AI and reasoning with neural networks,

attempts to incorporate the capabilities of both worlds and create systems that are data and time efficient, generalizable, and explainable. Neurosymbolic models have been applied to several real-world applications (Bouneffouf and Aggarwal, 2022) in safety-critical areas (Lu et al., 2024) such as healthcare (Hossain and Chen, 2025) and autonomous driving (Sun et al., 2021). Several techniques have been proposed for this integration (Kautz, 2022; Jayasingha et al., 2025), trying to combine the advantages and mitigate the disadvantages from both symbolic and neural methods. However, due to the lack of unified libraries to facilitate this research and the focus on specific algorithms rather than generic frameworks, this research becomes less impactful. Moreover, the few generic frameworks tend to vary in problem formulation, implementation, algorithms, and flexibility of use. This poses a challenge in being able to compare their performance uniformly or identify a research direction that improves on previous work. To alleviate this issue, we provide a comparative study with the following key contributions, a) Identifying the main components of existing NeSy frameworks, b) Comparison of frameworks across the identified facets, c) Highlighting the requirements for the next generation of NeSy frameworks, building upon the drawbacks of the current systems and the possible interplays between the neural and symbolic components. We plan to expand this study to cover more frameworks, while the three selected ones are used to explain the aspects of our characterization. These frameworks are demonstrated with four example tasks detailed in this paper tying the comparative facets concretely with a technical implementation.¹

2. Neurosymbolic Frameworks

A NeSy framework should provide flexibility for modeling both neural and symbolic components and their interplay in a unified declarative framework, going beyond specific underlying algorithms and techniques (Kordjamshidi et al., 2016, 2015). On the symbolic side, a generic framework should support a symbolic representation language that can be seamlessly connected to neural components and cover various underlying symbolic reasoning mechanisms. On the neural side, we need to have the flexibility of connecting to various architectures, including various loss functions, sources of supervision, and training paradigms. More importantly, a NeSy framework should provide a modeling language for seamless integration of these two paradigms in building pipelines, joint decision-making models and, in general, for arbitrary configurations of multiple models. Such a NeSy framework should support neuro-symbolic training and inference beyond specific integration algorithms. We distinguish between *NeSy techniques* and *NeSy frameworks*. By techniques, we mean when problem-specific algorithms are provided (Lample and Charton, 2020; Burattini et al., 2002). For example, **AlphaGo** (Silver et al., 2016) introduced a reinforcement learning solution to Go, using Monte Carlo Tree Search as a symbolic component inside a neural network. Another example is **NS-CL** (Mao et al., 2019) (Neuro-Symbolic Concept Learner) that integrates neural perception with symbolic reasoning to learn visual concepts and compositional language grounding for VQA tasks. Many other techniques and algorithms are proposed for the interplay between the two paradigms (Badreddine et al., 2022; Cohen et al., 2017; Smolensky et al., 2016; Lima et al., 2005; Sathasivam, 2011; Serafini and d’Avila Garcez, 2016; Lamb et al., 2021; Zhang et al., 2024) some of which are employed in the generic

1. <https://github.com/HLR/nesy-examples>

frameworks, such as Inference Masked Loss (Guo et al., 2020), Semantic Loss (Xu et al., 2018), Primal-Dual (Nandwani et al., 2019). NeSy techniques often lack the generality of frameworks, which are designed as broader tools intended for practical use and extensibility with new integration algorithms and with the capability of programming and configuring the two parts and their interplay.




Figure 1: An overview of various components of a NeuroSymbolic AI framework. NeuroSymbolic AI is at the intersection of symbolic reasoning and neural learning, depicted in green and purple respectively. The interplay techniques can be classified into 6 types, shown in the middle. Arrows indicate the connections between components (not the computation flow).

In this work, we illustrate these components by focusing on a selection of representative, generic NeSy frameworks. These frameworks are chosen for their contribution to advancing the development of general-purpose approaches. **DeepProbLog** (Manhaeve et al., 2021) is a probabilistic logic programming language that incorporates neural predicates in logic programming with an underlying differentiable translation of logical reasoning. The probabilistic logic programming component is built on top of ProbLog (De Raedt et al., 2007). **DomiKnowS** (Rajaby Faghihi et al., 2021; Faghihi et al., 2023, 2024) is a declarative learning-based programming framework (Kordjamshidi et al., 2019) that integrates symbolic domain knowledge into deep learning. It is a Python framework that facilitates the incorporation of logical constraints that represent domain knowledge with neural learning in PyTorch. **Scallop** (Huang et al., 2021; Li et al., 2023b, 2024) is a framework that includes a flexible symbolic representation based on relational data modeling, using declarative logic programming similar to DeepProbLog while it is based on Datalog (Abiteboul et al., 1995) instead of Prolog. We also point to less general frameworks such as **LEFT** (Hsu et al., 2023) that is designed for grounding language in visual modality and compositional reasoning over concepts. The framework consists of an LLM interpreter that converts natural language to logical programs. The generated programs are directed to a differentiable, domain-independent, and soft first-order logic-based executor. LEFT is limited to tasks requiring grounding language in vision such as visual question answering (Johnson et al., 2017; Yi et al., 2018; Liu et al., 2019). Building on this foundation, NeSyCoCo, (Kamali et al., 2025) has been introduced to address the limitations of LEFT, particularly the

challenge of lexical variety of natural language and alleviating the issue of non-canonical predicate logic representations for handling unseen concepts. NeSyCoCo extends LEFT’s approach by using distributed word representations to connect a wide variety of linguistically motivated predicates to neural modules, thus alleviating the reliance on a predefined predicate vocabulary. **PyReason** (Aditya et al., 2023) is a library built to support reasoning on top of outputs from neural networks. The neural component produces outputs such as labels or concept scores, while the symbolic component does graph-based reasoning with logical rules declared in a graph structure. It can produce an explanation trace for inference and has a memory-efficient implementation. **PLoT** (Wong et al., 2023) (Probabilistic Language of Thought) is a *proposed* framework leveraging neural and probabilistic modeling for generative world modeling. It models thinking with probabilistic programs and meaning construction with neural programs. The goal is to provide a language-driven unified thinking interface. **CCN+** (Giunchiglia et al., 2024) is a framework that modifies the output layer of a neural network to make results compliant with requirements that can be expressed in propositional logic. A requirement neural layer, ReqL, is built on top of the neural network. The standard cross-entropy loss is adapted into a ReqLoss to learn from the constraints in the ReqL layer. While CCN+ is a general framework, the underlying neurosymbolic technique follows similar ideas that integrate logic in the architecture in (Li and Srikumar, 2019). **DeepLog** (Derkinderen et al., 2025) is another *proposed* neurosymbolic AI framework that unifies logic and neural computation using a declarative paradigm. It introduces the DeepLog language, an annotated neural extension of grounded first-order logic that can abstract over various logics and apply them either within the model architecture or in the loss function. It employs computational algebraic circuits implemented on GPUs, forming a neurosymbolic abstract machine. DeepLog allows efficient specification and execution of diverse neurosymbolic models and inference tasks in a declarative fashion.

Framework	Symbolic		Model Dec	Interplay		LLM
	Lang	Knowledge Rep		Algorithm	Eff	
CCN+	None	Propositional Logic	✗	Learning from constraints	✗	✗
DomiKnowS	None	Concepts & Relations graph, FOL	✓	Learning from constraints	✗	Faghihi et al. (2024)
DeepProbLog	ProbLog	FOL	✗	Deductive Reasoning	✗	✗
LEFT	None	FOL	✗	Differentiable soft logic	✗	Hsu et al. (2023)
PyReason	None	FOL	✗	Graph Reasoning	✓	✗
Scallop	DataLog	FOL	✗	Deductive Reasoning	✓	Li et al. (2024)

Table 1: Comparison of frameworks and key factors. Lang: External language required, Knowledge Rep: Knowledge Representation, Model Dec: Model Declaration flexibility, Algorithm: Supported algorithm(s) for learning and inference, Eff: Computational efficiency considerations, LLM: Use of Large Language Models, FOL: First Order Logic.

In this paper, we characterize the more generic NeSy frameworks based on: a) Symbolic knowledge representation language, b) Representation and flexibility of Neural Modeling,

c) Model Declaration, d) Interplay between symbolic and sub-symbolic systems, and e) The usage of LLMs. Figure 1 shows the connection between these aspects. The neural representations and the symbolic representations are the two main components of a neurosymbolic framework. The neural representation guides learning and obtains supervision from the data, while the symbolic representations leverage symbolic reasoning and exploit symbolic knowledge during training or inference. Table 1 shows an overview of the frameworks across selected factors. In the next sections, we focus on **DomiKnowS**, **DeepProbLog**, and **Scallop** to provide a deeper investigation of the challenges in each component. Depending on the framework, certain tasks are easier to implement, while others require more effort. The chosen frameworks enable us to solve the same task in multiple frameworks.

3. Symbolic Knowledge Representation

Generic Neuro-Symbolic (NeSy) systems and frameworks use symbolic knowledge representation languages to encode constraints, facts, probabilities, and rules. Frameworks vary in how they represent and integrate this symbolic knowledge. Many employ classical formal logic, grounded in well-defined syntax and semantics, and adapt these representations and reasoning mechanisms within a unified integration framework. Some frameworks build on established formalisms such as logic programming or constraint satisfaction. In contrast, others take an entirely new hybrid semantics, while preserving conventional symbolic syntax. Table 2 compares the implementation of symbolic knowledge (concepts or facts) for the MNIST Sum task. In general, the domain knowledge consists of the two concepts of *digits* and the *sum*.

As can be seen, DomiKnowS represents a part of symbolic domain knowledge as a graph $G(V, E)$, where the nodes are the concepts in the domain and the edges are the relationships between them. Each node can have properties. More complex knowledge beyond entities and relations is expressed with a pseudo first-order logical language with quantifiers designed in Python. DomiKnowS mostly interprets the symbolic knowledge as logical constraints, such as the implementation of `sum_combinations` in the given example. Unlike the other frameworks, DomiKnowS does not build on predefined formal semantics. It follows a FOL-like syntax for symbolic logical representations, making it independent of the formal semantics of an underlying formal language and allows more flexibility of representations and adaptation to underlying algorithms in the framework. DeepProbLog, on the other hand, utilizes logical predicates that are originally a part of the probabilistic logic programs (Ng and Subrahmanian, 1992) of ProbLog (De Raedt et al., 2007), for its symbolic representation. These neural predicates obtain their probability distributions from the underlying neural models. Probabilistic facts, neural facts, and neural annotated disjunctions (nAD) whose probabilities are supplied by the neural component of the program can be added. Here, `digit` is a neural predicate as indicated by the use of `nn(...)`. DeepProbLog follows the formal semantics of Prolog (Clocksin and Mellish, 2003), followed by ProbLog, its probabilistic extension. Finally, Scallop adopts a relational data model for symbolic knowledge representation (Kolaitis and Vardi, 1990). Scallop is built on top of the syntax and formal semantics of Datalog and its probabilistic extensions, relaxing the exact semantics of ProbLog. It allows for the expression of common reasoning, such as aggregation, negation, and recursion. Similar to DeepProbLog, some of these predicates in the symbolic part obtain their probability distribution from neural models, such as `digit_1` and

DomiKnowS	
1	image = Concept(name='image')
2	digit = image(name='digits', ConceptClass=EnumConcept, values=digits)
3	image_pair = Concept(name='pair')
4	pair_d0, pair_d1 = image_pair.has_a(digit0=image, digit1=image)
5	s = image_pair(name='summations', ConceptClass=EnumConcept, values=summations)
6	# ...
7	sum_combinations.append(andLgetattr(digit, d0_nm)(path=('x', pair_d0)), getattr(digit, d1_nm)(path=('x', pair_d1))))
DeepProbLog	
1	nn(mnist_net, [X], Y, [0,1,2,3,4,5,6,7,8,9]) :: digit(X,Y).
2	addition(X,Y,Z) :- digit(X,X2), digit(Y,Y2), Z is X2+Y2.
Scallop	
1	self.scl_ctx = scallopy.ScallopContext(provenance=provenance, k=k)
2	self.scl_ctx.add_relation("digit_1", int, input_mapping=list(range(10)))
3	self.scl_ctx.add_relation("digit_2", int, input_mapping=list(range(10)))
4	self.scl_ctx.add_rule("sum_2(a + b) :- digit_1(a), digit_2(b)")

Table 2: Comparison of Symbolic Representation across frameworks.

`digit_2`. Additionally, while ProbLog requires exhaustive search for computations, Datalog can use top-k results and exploit database optimizations, making Scallop algorithmically more time-efficient than DeepProbLog.

4. Neural Models Representations

The other core component of a NeSy system is the neural modeling that is integrated with the symbolic knowledge discussed above. The neural models are mostly wrapped up under the logical predicate names in most of the frameworks that have an explicit logical knowledge representation language. To best leverage the reasoning capabilities of the symbolic system available and the ability of neural models to learn abstract representations from data, the neural models are used as abstract concept learners for the concepts defined as logical predicates in the symbolic representation. The neural model representation is often used to predict probability distributions for the symbolic concepts based on raw sensory inputs. Given that neurosymbolic frameworks are explicit about the symbolic concepts and relations, reading data into the symbolic representation requires special considerations for processing raw data, putting data into structure and **Data Modeling**. The aspect of data modeling is often ignored in current NeSy frameworks and it is done by the programmers in an ad-hoc pre-processing step. The neural modeling is often written using standard deep learning libraries, such as PyTorch (Paszke et al., 2019).

DomiKnowS
<pre> 1 class Net(): 2 # Neural Network 3 image_batch['pixels', image_contains.reversed] = JointSensor(image['pixels'], 4 forward=make_batch) 5 image['logits'] = ModuleLearner('pixels', module=Net()) # ... </pre>
DeepProbLog
<pre> 1 class MNIST_Net(): 2 # Neural Network 3 network = MNIST_Net() 4 net = Network(network, "mnist_net", batching=True) 5 net.optimizer = torch.optim.Adam(network.parameters(), lr=1e-3) 6 model = Model("models/addition.pl", [net]) 7 # </pre>
Scallop
<pre> 1 class MNISTSum2Net(nn.Module): 2 def __init__(self, provenance, k): 3 # Neural Network 4 self.mnist_net = MNISTNet() 5 # Scallop Context 6 self.scl_ctx = scallopy.ScallopContext(provenance=provenance, k=k) 7 # 8 class Trainer(): 9 def __init__(self, train_loader, test_loader, model_dir, learning_rate, 10 loss, k, provenance): 11 self.model_dir = model_dir 12 self.network = MNISTSum2Net(provenance, k) 13 self.optimizer = optim.Adam(self.network.parameters(), lr=learning_rate) # ... </pre>

Table 3: For neural integration, DomiKnowS utilizes sensors and readers for reading in data, while a learner connects to a network. DeepProbLog connects the neural network to the ProbLog file, requiring data handling to construct the terms and queries from the raw data. Scallop has an additional layer on top of the standard network that adds the symbolic context.

Table 3 shows snippets of neural modeling expressions across frameworks, highlighting differences in implementation. Scallop utilizes relatively standard neural modeling using PyTorch, while needing an added context of symbolic rules. Although integrated into Python, the context relation and rule setup are verbatim from DataLog and only passed as a parameter to a function, which requires familiarity with DataLog and its semantics.

DeepProbLog, on the other hand, needs manual configuration of the raw data and processing into queries built for ProbLog, on top of other standard neural components. This processed data is passed into the neural network which is then connected to a ProbLog program, such as `addition.pl` in the figure. DomiKnowS’s neural component is built in PyTorch. Unlike other frameworks, DomiKnowS has built-in components called *Readers*, *Sensors* and *Module learners* for a more *transparent data modeling*, making the connection to neural components and feeding data into them explicitly in the program. This provides more flexibility in connecting the concepts to deterministic or probabilistic functions that can interact with other symbolic concepts. The module learner can also use custom models. This makes the interaction with raw data structured, transparent, and controllable.

5. Model Declaration

Most frameworks utilize neural components as abstract concept learners and use a symbolic component to reason over the learned concepts. Each learner is a model and *model declaration* refers to the flexibility of modularizing and connecting different learners. Each learner can receive supervision independently. In most neurosymbolic frameworks, the supervision from data is usually provided based on the final output of the end-to-end model. For example, in an MNIST Sum task used throughout and detailed in Section 8, the neural and symbolic components are trained based on the final output of the sum, without access to individual digit labels in a semi-supervised setting. The task loss, e.g., a Cross-Entropy Loss, is computed, and errors are backpropagated through the differentiable operations that led to the output generation. For example, in DeepProbLog, we can declare a single loss function associated with the entire neural component. Gradient computations differ across frameworks depending on whether losses are defined individually for each neural output or specified as a single global loss function. However, there remains a need for models capable of incorporating supervision at multiple levels of their symbolic representations. In DomiKnowS, loss computation can be defined for each symbol. Since each concept is linked to both learning modules and ground-truth labels, their losses can be integrated seamlessly. This enables joint training of all concepts alongside the target task, allowing each concept to be optimized more effectively—leveraging available data without relying solely on the target task’s output. In other words, it provides the flexibility of building pipelines of decision making, obtaining distant supervision in addition to joint training and inference. Model declaration is very well manifested in modern tools and libraries for building agentic frameworks, where multiple agents are structured into pipelines or graphs to make communicative decisions by passing the output of one agent as input to another (Plaat et al., 2025).

6. Interplay between Symbolic and Sub-symbolic

Kautz (2022) provides a characterization of the possible interplays between symbolic and sub-symbolic components. This interplay of neurosymbolic can be explained by the concept of System 1 and System 2 thinking described in Kahneman (2011). Research in this field aims to create an ideal integration that seamlessly supports ”thinking fast and slow” (Booch et al., 2021; Fabiano et al., 2023). Here, System 1 refers to the fast neural processing, while

System 2 corresponds to the slower, more deliberate symbolic reasoning. Different methods for the integration of symbolic reasoning and neural programming have been explored such as employing *logical constraint satisfaction*, *integer linear programming*, *differentiable reasoning*, *probabilistic logic programming*. In this section, we will discuss a system-level algorithmic comparison of the different frameworks.

DomiKnowS models the inference as an integer linear programming problem to enforce the model to follow constraints expressed in first-order logical form (Van Hentenryck et al., 1992). The objective of the program is guided by the neural components, and the framework supports multiple training algorithms for learning from constraints. The Primal-Dual formulation (Nandwani et al., 2019) converts the constrained optimization problem into a min-max optimization with Lagrangian multipliers for each constraint, augmenting the original loss with a soft logic surrogate to minimize constraint violations. Sampling-Loss (Ahmed et al., 2022), inspired by semantic loss (Xu et al., 2018) and samples a set of assignments for each variable based on the probability distribution of the neural modules' output. Integer Linear Programming (ILP) (Cropper and Dumančić, 2022) formulates an optimization objective based on Inference-Masked Loss (Guo et al., 2020) to constrain the model during training. The training goal is to adjust the neural models to produce legitimate outputs that adhere to the given constraints. At prediction time, ILP can also be applied to enforce final predictions that comply with given constraints. DomiKnowS relies on the off-the-shelf optimization solver Gurobi (Gurobi Optimization, LLC, 2024). **DeepProbLog** models each problem as a probabilistic logical program that consists of neural facts, probabilistic facts, neural predicates, and a set of logical rules. A joint optimization of the parameters of the logic program is done alongside the parameters of the neural component. Neural network training is done using learning from entailment (Frazier and Pitt, 1993) while in ProbLog, gradient-based optimization is performed on the underlying generated Arithmetic Circuits (Shpilka et al., 2010), which is a differentiable structure. The Arithmetic Circuits are transformed from a Sentential Decision Diagram (Darwiche, 2011) generated by ProbLog. Algebraic ProbLog (Kimmig et al., 2011) is used to compute the gradient alongside probabilities using semirings (Eisner, 2002). **Scallop** is similar in its setup to DeepProbLog where it creates an end-to-end differentiable framework combining a symbolic reasoning component with a neural modeling component. They aim to relax the formal semantics required by the use of ProbLog in DeepProbLog and instead rely on a symbolic reasoning language extending DataLog, built into their framework. They have a customizable provenance semiring framework (Green et al., 2007), where different provenance semirings, such as extended max-min semiring and top-k proofs semiring, allow learning using different types of heuristics for gradient calculations. Table 4 compares the computational efficiency of these models at training and inference time on a single training/testing example. As theoretically suggested, Scallop is expected to outperform other frameworks in inference and training speed, owing to its memory and time-efficient implementation in Rust. The results in Table 4 support this expectation, with Scallop achieving the fastest inference time, on par with DomiKnowS. In practice, DeepProbLog achieves slightly faster training performance than Scallop. This discrepancy may be due to overhead unrelated to the core algorithmic complexity. DomiKnowS exhibits slower training due to the overhead of uploading the entire graph of data into memory.

7. Role of Large Language Models

Large foundation models hold significant promise for overcoming the bottleneck of acquiring symbolic representations, which are essential for symbolic reasoning and consequently in neurosymbolic frameworks.

Source of Symbolic Knowledge: The symbolic knowledge in neuro-symbolic systems, which is integrated with the neural component, can originate from several distinct sources. While most systems require explicit, hand-crafted symbolic knowledge, earlier classical logic-based learning research can be used for automatically learning rules from data by using inductive logic programming (Nienhuys-Cheng and de Wolf, 1997; Bratko and Muggleton, 1995) or mining constraints. Nowadays, even LLMs can be utilized to generate symbolic knowledge (Pan et al., 2023a; Mirzaee and Kordjamshidi, 2023; Acharya et al., 2024; Xu et al., 2024a). Several neurosymbolic frameworks and systems have tried utilizing large foundation models to generate the symbolic knowledge, based on the task or query, to overcome the labor-intensive nature of hand-crafting rules for every single task and the time required in the automatic learning of symbolic knowledge from data (Ishay et al., 2023; Xu et al., 2024a; Yang et al., 2024). Extraction of symbolic representations from Foundation Models has become possible given the vast implicit knowledge stored within these models, such as LLMs and multimodal models, which are trained on massive and diverse corpora (Li et al., 2024; Petroni et al., 2019). These models can generate symbolic content (e.g., candidate rules, knowledge graph triples, or logic statements), perform reasoning that mimics symbolic inference, or act as components alongside symbolic modules (Fang and Yu, 2024). For example, LLMs can be prompted to extract facts from unstructured text, effectively populating a symbolic knowledge graph (Yao et al., 2025). Techniques like Symbolic Chain-of-Thought inject formal logic into the LLM’s reasoning process, improving accuracy and explainability on logical reasoning tasks (Xu et al., 2024b). However, foundation models are prone to hallucinations and lack the strict logical guarantees of traditional symbolic systems (Zheng et al., 2024). Therefore, integrating foundation models often requires careful prompting, verification steps to ensure reliability (Xu et al., 2024b). **Generation of inputs to symbolic engines:** LLMs have also been used to generate translations from raw inputs, specially natural language, to symbolic language that is then fed into a symbolic reasoner. In examples such as Logic-LM (Pan et al., 2023b), LLMs are leveraged to convert a natural language query into symbolic language that is then solved by a symbolic reasoner. This method improves the performance of unfinetuned LLMs on logical reasoning-based tasks. DomiKnowS (Faghihi et al., 2024) takes this a step further by enabling users to describe problems in natural language which LLMs then use to generate relevant concepts and relationships. Through a user-interactive process, these concepts and relationships are refined iteratively. Finally, the LLM translates the user-defined constraints from natural language into first-order logic representations before converting them into DomiKnowS syntax. Some systems use LLMs in multiple capacities. In VIERA (Li et al., 2024), which is built on top of Scallop, 12 foundation models can be used as plugins. These models are treated as stateless functions with relational inputs and outputs. These foundation models can be either language models like GPT (OpenAI et al., 2024) and LLaMA (Touvron et al., 2023), vision models such as OWL-ViT (Minderer et al., 2022) and SAM (Kirillov et al., 2023), or multimodal models such as CLIP (Radford et al., 2021). These models can be

used to extract facts, assign probabilities, or for classification, and are treated as "foreign predicates" in their interface. An older version, DSR-LM (Zhang et al., 2023) of this utilized BERT-based language models for perception and relation extraction, combined with a symbolic reasoner for question answering. LEFT, on the other hand, uses LLMs both for the generation of the concepts that are used for grounding and as an interpreter to generate the first-order logic program corresponding to a natural language query, that is solved by the symbolic executor.

8. Example Tasks

NeSy frameworks formulate problems in various ways based on their implementation and symbolic interpretation. In **DomiKnowS**, the symbolic reasoning part is formulated as a logical constraint solving problem. The domain is represented as a graph $G(V, E)$, where the nodes are the concepts in the domain and the edges are the relationships between them. Each node can have properties. The final logical constraints apply to the graph concepts. In **DeepProbLog**, the symbolic reasoning problem is interpreted as probabilistic logic programs in ProbLog. In **Scallop**, similarly, the problem is viewed as a combination of the neural and the symbolic components where the symbolic part is a probabilistic logical program similar to DeepProbLog with further optimized inference. In **LEFT**, the problem is limited to the application of concept learning and grounding language into visual modality. Here, the neural model is composed of feature extractors, object and relation classifiers (concept learners), and a first-order logic program generator for a given question. In this section, we will compare the problem formulations in each of these frameworks for a set of tasks. **Note** that we only include LEFT for the visual question answering task due to the domain-specific nature of the framework. All code associated with these tasks and referenced in this section is maintained publicly on GitHub.²

8.1. MNIST Sum

The MNIST Sum task is an extension of the classic MNIST handwritten digit recognition task (Lecun et al., 1998) where given two images of digits, the task is to output their sum that is a whole number. The training examples consist of the two images of the digits and the ground-truth label of their sum. The individual labels of the digits are not available for training.

8.1.1. DOMIKnowS

Problem Specification. DomiKnowS formulates the problem using graph representations of concepts, relations, and logic. For performing the MNIST Sum task in DomiKnowS, the first concept defined is *image* concept representing visual information. The *digit* concept, a subclass of image, is introduced to represent the output class, ranging from 0 to 9. To establish relationships between digit images, the *image pair* concept is defined as an edge connecting two digit concepts. The sum concept is then introduced under image pair to represent the summation of the two digit concepts and the ground-truth output of the program. For this task, three constraints are defined. The first two constraints utilize

2. <https://github.com/HLR/nesy-examples>

Framework	MNIST Sum			Toy-NER		
	Training Time (ms)	Testing Time (ms)	Memory (MB)	Training Time (ms)	Testing Time (ms)	Memory (MB)
DomiKnowS	37.72	2.34	573.80	14.86	14.29	576.82
DeepProbLog	5.84	3.24	739.72	20.74	20.47	3767.08
Scallop	6.50	2.35	581.36	1.50	1.05	297.1

Framework	Math-Inference			Simple VQA		
	Training Time (ms)	Testing Time (ms)	Memory (MB)	Training Time (ms)	Testing Time (ms)	Memory (MB)
DomiKnowS	77.46	69.58	1039.30	81.02	60.52	1413.44
DeepProbLog	5.54	5.07	1588.30	755.13	363.81	1346.23
Scallop	0.948	0.223	345.48	13.17	22.66	768.4
LEFT	N/A	N/A	N/A	6.19	3.44	755.1

Table 4: Time and Space efficiency for each framework across 4 tasks. Time and memory records are averaged over 5 runs. Times are in milliseconds per sample. Memory utilization is in megabytes and indicates the overall memory for training.

exactL to ensure that the predicted digit and sum values belong to only one valid class. Another constraint enforces that the expected sum value matches the sum of the two digit predictions. This is implemented using *ifL* constraints, which verify whether the predicted digits form one of the possible solutions for a valid sum. If multiple solutions exist, the *orL* constraint ensures that at least one of the answers corresponds to the predicted digits.

Neural Modeling. The model declaration comprises standard neural modeling components, including data loading, pre-processing, neural network definition, and loss function specification. The process begins with the *ReaderSensor*, which reads the input image. Next, a relation concept is defined using another sensor, *JointSensor*, to establish connections between images. The module learner is then employed to generate an initial prediction for the digit concept, which is subsequently passed to another sensor, *FunctionalSensor*, to compute the sum of two images.

8.1.2. DEEPPROBLOG

Problem Specification. DeepProbLog formulates a problem regarding probabilistic facts, neural facts, and neural annotated disjunctions (nAD). In the MNIST Sum task, the fact X is defined to represent the input image. A neural network function is then introduced to map X to its corresponding digit, denoted as $\text{digit}(X, Y)$. To enforce constraints about the summation and the ground-truth sum, a function is defined to compute the sum of two digits.

Neural Modeling. The neural modeling follows a standard neural network setup, such as a CNN-based classifier. It is preceded by data loading and pre-processing, which are performed separately from the ProbLog program. Thus, the neural model used in DeepProbLog can be initialized independently of the DeepProbLog model. Once the neural model is initialized, the framework passes it along with a probabilistic program as input.

8.1.3. SCALLOP


Problem Specification. Scallop formulates the problem in terms of relations, values, and (Horn) rules derived from Datalog. As discussed earlier, the concepts and constraints defined in this framework are similar to those in DeepProbLog. However, these rules can be directly embedded into a Scallop program through its API. The process begins by establishing the concepts *digit1* and *digit2* to represent the digit values of two given images. Based on the summation of these two values, it must be equal to the *sum_2* logical reasoning module, which serves as the ground truth for this task.

Neural Modeling. Unlike DeepProbLog, the neural modeling is integrated with Scallop’s relation and rule declaration. The neural modeling remains a standard neural network.


8.2. Shapes

The **Shapes** dataset is a synthetic VQA benchmark designed to evaluate elementary spatial reasoning. Each sample consists of a 128×128 pixel image where the task is to answer the fixed question: “Is there a red shape above a blue circle?”. The primary rationale for using a synthetic dataset is to create a controlled experimental environment. This approach allows us to isolate specific reasoning skills—such as attribute binding and relational understanding—from the perceptual complexities and spurious correlations.

Positive examples are generated by programmatically placing a red object (circle, square, or triangle) above a blue circle. In negative examples, this specific spatial configuration is absent. To increase task complexity, every image also contains one to three randomly placed, non-overlapping distractor objects with varying shapes and colors. The dataset comprises 2,000 images, divided into perfectly balanced training and testing sets of 1,000 samples each. Examples of this benchmark are shown in Figure 2.



True example: A red triangle is above a blue circle.



False example: No red shape is above a blue circle.

Figure 2: Examples from the Shapes dataset for the question “Is there a red shape above a blue circle?”

8.2.1. COMMON PERCEPTION-REASONING INTERFACE

We decompose the system into: (i) a neural perception unit producing distributions over object attributes and pairwise relations; and (ii) a symbolic reasoning unit that consumes these as soft facts over the object domain and evaluates a single existential rule encoding the query. Supervision is a binary yes/no label optimized via a cross-entropy loss on the final answer distribution.

8.2.2. DomiKnowS

Problem Specification. DomiKnowS formulates the problem as a graph representation. It first declares the *image* concept to represent the visual input. The *objects* concept is then defined to represent the individual objects contained within the image. An explicit connection declaration is established in the framework between images and objects, indicating that each image may contain multiple objects. Under the image concept, two additional concepts—*color* and *shape*—are introduced to represent the attributes of each object for each shape and color considered. These two concepts serve as the outputs of the neural modeling component. Next, DomiKnowS defines a *relation* concept between pairs of objects to capture their relational structure within the same image. Finally, symbolic reasoning is expressed using *existsL*, which denotes the existence of a particular combination of queries. The inference takes the following form:

```
existsL(is_red(X), is_blue(Y), is_circle(Y), relation(Rel))
```

where the X and Y represent the first object and second object, Rel represent the relation of X and Y .

Neural Modeling. DomiKnowS uses a class of functions called *ReaderSensors* to process the raw input data. A *ModuleLearner* is then employed to query the object-centric encoder, producing representations of the objects within the image. These representations are subsequently passed to three *ModuleLearner* to predict the color (red and blue) and shape (circle) attributes of the objects. A *CompositionCandidateSensor* is used to connect pairs of objects within the image, providing the information needed to compute their relations. Finally, the framework aggregates all information through predefined logical expressions, which are used to infer the final output.

8.2.3. DEEPPROBLOG

Problem Specification. We use a ProbLog program with neural annotated disjunctions for `color(Image, 0, C)`, `shape(Image, 0, S)`, and `relation(Image, 01, 02, R)`. The decision rule is:

```
answer(Image, yes) :- obj(01), obj(02),
    01 /= 02,
    color(Image, 01, red),
    color(Image, 02, blue),
    shape(Image, 02, circle),
    relation(Image, 01, 02, R).
```

A complementary `answer(Image, no)` ensures a normalized binary outcome. The object domain `obj/1` is declared up to n_{\max} derived from annotations. Batched queries are issued as `answer(tensor(batch(I)), Y)` and solved with an exact engine.

Neural Modeling. An object-centric encoder crops and embeds objects; heads output categorical distributions for color and shape per object and for relation per ordered object pair. These heads are wrapped as DeepProbLog networks and bound to the program’s neural predicates, enabling end-to-end training from the `answer` predicate.

8.2.4. SCALLOP

Problem Specification. We declare unary predicates over object indices for attributes (`red(i)`, `blue(j)`, `circle(j)`) and a binary predicate for the chosen spatial relation `R(i, j)`. The query is encoded as a Horn rule:

```
ans() :- red(i), blue(j), circle(j), above(i, j).
```

Scallop maps soft facts to weighted relations under a differentiable provenance and produces a normalized {yes,no} answer.

Neural Modeling. Per-object (color, shape) and per-pair (relation) distributions from the encoder are converted to Scallop facts, respecting masks for padded objects and pairs. The Scallop context composes these facts with the rule to yield the answer distribution, providing gradients to the neural unit.

8.2.5. LEFT

Problem Specification. We express the query in first-order logic executed by LEFT’s generalized FOL executor:

```
exists(Object, lambda x: exists(Object, lambda y:
    above(y, x) and red(x) and circle(y) and blue(y)))
```

The executor grounds over the object domain and evaluates the formula using attribute predicates for unary properties and a binary predicate for the relation, returning a scalar decision consistent with the query.

Neural Modeling. The object-centric encoder yields per-object attribute scores (for color/shape) and per-pair relation scores. These tensors parameterize the executor’s predicates, forming a differentiable perception–reasoning pipeline trained on the binary supervision.

8.3. Toy NER

For a simplified, toy version of the Named-Entity Recognition task ([Tjong Kim Sang and De Meulder, 2003](#)), we create a dataset of randomly generated embeddings representing persons and locations. The objective is to learn the concepts of ”works_in”, i.e., whether a person works in a location, and ”is_real_person”, i.e., whether a given embedding is a person or not. These concepts are learned using indirect supervision on two queries, that compose the atomic concepts.

```
constraint1 = is_real_person(P1) AND works_in(P1,L1) AND is_real_person(P2)
            AND works_in(P2,L2)
constraint2 = is_real_person(P2) AND works_in(P2,L2) OR is_real_person(P3)
            AND works_in(P3,L3)
```

Here, $P_1, P_2, P_3, L_1, L_2, L_3$ are input embeddings with P s representing persons and L s representing locations. Thus, the final task is to predict the output of these two constraints (true or false), given 6 embeddings corresponding to 3 persons and 3 locations.

8.3.1. DOMIKNOWS

Problem Specification. To perform the Toy NER task, DomiKnowS begins by constructing a graph representation of the problem. It first defines the *person* and *location* concepts to represent the entity types of interest. Under the *person* concept, three sub-concepts are declared to represent three individual persons, and under the location concept, three sub-concepts are introduced to represent distinct locations. A *pair* relation is then defined to connect a person and a location. This relation is used to create three separate *work_in[i]* concepts, each representing the output relation between person i and location i . Finally, the inference queries are expressed in a form-based manner, using *andL* to represent logical conjunctions (AND) between concepts and *orL* to represent logical disjunctions (OR) in the desired queries above.

Neural Modeling. The model declaration follows the standard pipeline of neural components. It begins with a *ReaderSensor* to process the raw input data. A *ModuleLearner* is then invoked, using the embeddings of either person or location to generate predictions based on the given representations. A *JointSensor* is subsequently employed to connect the person and location concepts in order to form the *work_in[i]* relation specified in the problem definition. Finally, model inference is carried out based on predefined queries, yielding the final output.

8.3.2. DEEPPROBLOG

Problem Specification. DeepProbLog declares 6 concepts, *is_real_person[i]* for each person i and *works_in[i]* for each location i . These concepts are used to formulate the two constraints with the final query, *check*, being the conjunction of the two. Note that, DeepProbLog does not support multiple simultaneous queries and hence, the two constraints cannot be trained with individual supervision for labels of each.

Neural Modeling. DeepProbLog uses the two neural networks for classification of the two concepts. The main challenge is the creation of the dataloader that reads from the raw data and converts it into the format necessary to send queries with the appropriate value substitutions for variables. This integration with ProbLog needs to be done by the user from scratch.

8.3.3. SCALLOP

Problem Specification. Scallop utilizes 6 concepts similar to DeepProbLog. The value of the two constraints is merged using a final module *check(P1 * P2 * W1 * W2 + P3 * W3)*. Here, the results of concepts *is_real_person[i]* ($P1, P2, P3$) and *works_in[i]* ($W1, W2, W3$) are passed instead of the embeddings themselves. This context is passed through python itself, instead of a separate Datalog file.

Neural Modeling. The neural modeling component is standard with flexible loss declarations. To adapt from the raw data, an embedding generator *NERDataset* is created which is the same as the *JSONDataset* utilized in the DeepProbLog solution. However, there is no need to manually adapt the data to generate queries to the datalog component, as the forward function of the neural network is the final query and only requires the outputs of the neural networks passed to it.

8.4. Math Equation Inference

The Math Equation Inference task is designed to evaluate whether a neural network can learn local mathematical concepts from global supervision. The input to this task consists of two lists, each containing six real numbers sampled uniformly from the range $[-1, 1]$. The objective is to determine whether the model can learn from a global condition that encompasses the properties of the first list, the properties of the second list, and the relationship between them. We consider two properties, that is, $\sum_{i=0}^8 x_i > 0$ and $\sum_{i=0}^8 |x_i| > 0.5$. For relations between the two lists, we examine two cases: (i) whether the first elements of the lists have the same sign, and (ii) whether their last elements have opposite signs. These concepts are learned through indirect supervision using global queries that compose the atomic concepts, expressed as follows:

property1(L1) AND property2(L2) AND relation(L1, L2)

where $L1, L2$ are randomly generated lists of six real numbers, property1 and property2 are drawn from the set of considered properties, and relation is drawn from the set of considered relations.

8.4.1. DOMIKnowS

Problem Specification. For implementing this task in DomiKnowS, the first concept defined is the *problem* concept, which represents the overall mathematical inference problem. Next, the *lst* concept is introduced to represent a list of considered numbers. Under this *lst* concept, all possible condition concepts—*is_cond1* and *is_cond2*—are defined to represent the output class of each list. Subsequently, two relation concepts, *is_relation1* and *is_relation2*, are introduced between pairs of *lst* instances to capture the two relational conditions described in the problem statement. Importantly, all possible conditions and relations must be defined, even if some are not ultimately used in the final inference. To define the inference query, the logical operator *andL* is employed to combine the three relevant concepts. For example: *andL(is_cond1(L1), is_relation1(L1, L2), is_cond2(L2))*. This query corresponds to the case where the sum of $L1$ is greater than 0 (*is_cond1*), the first elements of $L1$ and $L2$ share the same sign (*is_relation1*), and the sum of the absolute values of $L2$ is greater than 0.5 (*is_cond2*).

Neural Modeling. DomiKnowS begins with a *ReaderSensor* to process the two input lists of numbers. Then, *CompositionCandidateSensor* is employed to connect the two lists, forming the intermediate representation required for relation prediction in later stages. Four neural networks are instantiated using *ModuleLearner*—two dedicated to property concepts and two to relation concepts. These networks are trained to predict the respective properties and relations. The predicted concepts are subsequently integrated into the final query, which specifies the target relation for the given problem and serves as the prediction label during training and evaluation. It is important to note that DomiKnowS requires all possible inference queries to be explicitly represented in the graph. Each query is set as either active or inactive to obtain the correct prediction during inference.

8.4.2. DEEPPROBLOG

Problem Specification. In this task, DeepProbLog defines four neural predicates: two for the considered properties (*is_property[i]_nn* for object property i), and two for the considered

relations ($is_relation[i]_nn$ for relation i). These four neural predicates are then combined to define the final query concept, *inference*. One *inference* concept is created for each possible combination of the condition of $L1$ (2 possibilities), the condition of $L2$ (2 possibilities), and the relation between $L1$ and $L2$ (2 possibilities). In total, eight inference concepts are formulated. During evaluation, only the relevant concept corresponding to the input configuration is activated and used to produce the final prediction for the problem.

Neural Modeling. DeepProbLog employs four neural networks, corresponding to the defined concepts: two for property concepts and two for relation concepts. The pipeline then proceeds with data loading, where all possible candidate lists of numbers are read separately and pre-processed to obtain the required model inputs. This also includes manually connecting the list of numbers within the same problem. Then, all defined neural networks are called to obtain all possible relations and properties output. Lastly, the query must be constructed in the exact predefined pattern to ensure that the system produces the correct output aligned with the desired inference condition.

8.4.3. SCALLOP

Problem Specification. Similar to DomiKnowS and DeepProbLog on this task, Scallop begins with the declaration of four concepts, that is, two property concepts and two relation concepts. However, Scallop only defines the one final module, $inference(P1 * P2 * R)$, that accumulates the probability based on the output of the properties of $L1$ and $L2$, and the relation between $L1$ and $L2$. The output is defined in the neural modeling part to get the connection between the final module and the defined concept.

Neural Modeling. The neural modeling component in Scallop follows the standard pipeline. However, Scallop requires a manually defined connection between the model output and the declared final module. This requirement arises because the final module is specified in a general form, rather than being tied to a particular combination of inference concepts. In contrast, other frameworks explicitly define every possible query corresponding to a combination of properties and relations and refer only to those queries during inference.

9. Discussion and Future Direction

Table 1 summarizes the comparative aspects of existing frameworks and outlines future directions for optimizing, as we observe many columns marked with ' \mathbf{X} ', implying most frameworks present challenges that hinder the application and flexibility of the frameworks. While current frameworks are functional, future developments should take a more holistic approach that considers all aspects from an end-user perspective, aiming to improve usability as general-purpose libraries and foster wider adoption of neurosymbolic methods.

Symbolic Representation. The generic neurosymbolic frameworks provide a formal knowledge representation language of their choice. The selected languages often are based on pure logical formalisms with established formal semantics, for example, Datalog or Prolog. However, we argue that knowledge representation for neurosymbolic frameworks needs to be an innovative language designed for this integration purpose with adaptable semantics with learning as the pivotal concept (Kordjamshidi et al., 2019). Restricting these frameworks to classical AI formalisms and formal semantics limits the level of extension that can be made and restricts the support of various algorithms and types of integration.

Neural Modeling. Most of the examined frameworks leave neural modeling and the task of connecting the symbolic and sub-symbolic components, up to the user. This connection usually requires low-level data pre-processing, which is time consuming to implement. A lack of user-friendly libraries discourages developers from using neurosymbolic methods to solve downstream tasks. There is a need for abstractions in these frameworks (Kordjamshidi et al., 2022) that improve user experience and remove the need for users to implement such data processing from scratch.

Model Declaration. There is a need to be explicit about the low-level components of the neural architecture, enabling us to design interactions between neural and symbolic components and connect them as intended. The goal is to provide flexibility in designing arbitrary loss functions and connecting them to data for supervising concepts at various neural layers, which will allow any symbol to be learnable.

Types of Interplay. Considering Kautz (2022)’s classification, current frameworks are limited in supporting one or two ways of interactions. DeepProbLog and Scallop utilize one form of implementation, while DomiKnowS has multiple settings as detailed in Section 6. One of the key challenges is determining the appropriate level of abstraction in a neural model after which reasoning should occur. The classification types demonstrate how a neural model can identify the relevant symbolic representations and suggest that neurosymbolic frameworks could leverage these models to learn and route inputs to the corresponding symbolic reasoning system. However, it remains unclear what level of abstraction is most effective for solving the end task in practice.

LLM. Drawbacks often associated with employing symbolic AI into neural computing, such as creation of the symbolic knowledge for integration, can be mediated with the use of LLMs and foundation models. LLMs have the potential to alleviate the classical issues in symbolic processing. Their vast knowledge can also be utilized to reduce the need for rebuilding neural components, allowing for flexible connections with different symbolic components.

10. Conclusion

Neurosymbolic AI presents a promising path forward in addressing the limitations of purely symbolic or neural approaches to AI. By integrating symbolic reasoning with neural learning, NeSy frameworks offer a balance between interpretability, data usage, time efficiency, and generalization. In this paper, we characterize the core components of NeSy frameworks and provide an analysis of some existing ones - DeepProbLog, Scallop, and DomiKnowS, illustrating their comparative facets. We identified a number of facets such as symbolic knowledge and data representation, neural modeling, model declaration, their method of integrating the symbolic and sub-symbolic systems, and LLMs integration. We identify key challenges in each facet that can guide us toward building the next generation of neurosymbolic frameworks. Unifying ideas in the field and building flexible frameworks by incorporating strengths in every facet will ease the learning curve associated with NeSy systems and improves their application. Future NeSy frameworks should aim to provide user-friendly interfaces, scalability, algorithm coverage and seamless integrations with foundation models. Recent advances in large language models (LLMs) and vision–language models (VLMs) provides promising solutions to longstanding knowledge engineering challenges, fostering more effective and scalable integration of symbolic representations, and advancing research in neurosymbolic AI.

Acknowledgments

This project is partially supported by the Office of Naval Research (ONR) grant N00014-23-1-2417. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Office of Naval Research. We thank Uzair Mohammad for his editing and suggestions for Figure 1.

References

Serge Abiteboul, Richard Hull, and Victor Vianu. *Foundations of Databases: The Logical Level*. Addison-Wesley Longman Publishing Co., Inc., USA, 1st edition, 1995. ISBN 0201537710.

Kamal Acharya, Alvaro Velasquez, and Houbing Herbert Song. A survey on symbolic knowledge distillation of large language models. *IEEE Transactions on Artificial Intelligence*, 5(12):5928–5948, 2024. doi: 10.1109/TAI.2024.3428519.

Dyuman Aditya, Kaustuv Mukherji, Srikanth Balasubramanian, Abhiraj Chaudhary, and Paulo Shakarian. Pyreason: Software for open world temporal logic, 2023. URL <https://arxiv.org/abs/2302.13482>.

Jamil Ahmad, Haleem Farman, and Zahoor Jan. *Deep Learning Methods and Applications*, pages 31–42. Springer Singapore, Singapore, 2019. ISBN 978-981-13-3459-7. doi: 10.1007/978-981-13-3459-7_3. URL https://doi.org/10.1007/978-981-13-3459-7_3.

Kareem Ahmed, Tao Li, Thy Ton, Quan Guo, Kai-Wei Chang, Parisa Kordjamshidi, Vivek Srikumar, Guy Van den Broeck, and Sameer Singh. Pylon: A pytorch framework for learning with constraints. In Douwe Kiela, Marco Ciccone, and Barbara Caputo, editors, *Proceedings of the NeurIPS 2021 Competitions and Demonstrations Track*, volume 176 of *Proceedings of Machine Learning Research*, pages 319–324. PMLR, 06–14 Dec 2022. URL <https://proceedings.mlr.press/v176/ahmed22a.html>.

Luis M. Augusto. From symbols to knowledge systems: A. newell and h. a. Simon’s contribution to symbolic AI. *Journal of Knowledge Structures and Systems*, 2(1):29–62, 2021.

Samy Badreddine, Artur d’Avila Garcez, Luciano Serafini, and Michael Spranger. Logic tensor networks. *Artificial Intelligence*, 303:103649, 2022. ISSN 0004-3702. doi: <https://doi.org/10.1016/j.artint.2021.103649>. URL <https://www.sciencedirect.com/science/article/pii/S0004370221002009>.

Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’21, page 610–623, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383097. doi: 10.1145/3442188.3445922. URL <https://doi.org/10.1145/3442188.3445922>.

Bikram Pratim Bhuyan, Amar Ramdane-Cherif, Ravi Tomar, and T. P. Singh. Neuro-symbolic artificial intelligence: a survey. *Neural Computing and Applications*, 36(21):12809–12844, July 2024. ISSN 1433-3058. doi: 10.1007/s00521-024-09960-z. URL <https://doi.org/10.1007/s00521-024-09960-z>.

Grady Booch, Francesco Fabiano, Lior Horesh, Kiran Kate, Jonathan Lenchner, Nick Linck, Andreas Loreggia, Keerthiram Murgesan, Nicholas Mattei, Francesca Rossi, and Biplav Srivastava. Thinking fast and slow in ai. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17):15042–15046, May 2021. doi: 10.1609/aaai.v35i17.17765. URL <https://ojs.aaai.org/index.php/AAAI/article/view/17765>.

Djallel Bouneffouf and Charu C. Aggarwal. Survey on applications of neurosymbolic artificial intelligence, 2022. URL <https://arxiv.org/abs/2209.12618>.

Ivan Bratko and Stephen Muggleton. Applications of inductive logic programming. *Commun. ACM*, 38(11):65–70, November 1995. ISSN 0001-0782. doi: 10.1145/219717.219771. URL <https://doi.org/10.1145/219717.219771>.

E. Burattini, A. de Francesco, and M. De Gregorio. Nsl: a neuro-symbolic language for monotonic and non-monotonic logical inferences. In *VII Brazilian Symposium on Neural Networks, 2002. SBRN 2002. Proceedings.*, pages 256–261, 2002. doi: 10.1109/SBRN.2002.1181487.

William F Clocksin and Christopher S Mellish. *Programming in PROLOG*. Springer Science & Business Media, 2003.

William W Cohen, Fan Yang, and Kathryn Rivard Mazaitis. Tensorlog: Deep learning meets probabilistic dbs. *arXiv preprint arXiv:1707.05390*, 2017.

Andrew Cropper and Sebastijan Dumančić. Inductive logic programming at 30: A new introduction. *J. Artif. Int. Res.*, 74, September 2022. ISSN 1076-9757. doi: 10.1613/jair.1.13507. URL <https://doi.org/10.1613/jair.1.13507>.

Adnan Darwiche. Sdd: A new canonical representation of propositional knowledge bases. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, volume 22, page 819, 2011.

Luc De Raedt, Angelika Kimmig, and Hannu Toivonen. Problog: a probabilistic prolog and its application in link discovery. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, IJCAI'07, page 2468–2473, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc.

Vincent Derkinderen, Robin Manhaeve, Rik Adriaensen, Lucas Van Praet, Lennert De Smet, Giuseppe Marra, and Luc De Raedt. The deeplog neurosymbolic machine, 2025. URL <https://arxiv.org/abs/2508.13697>.

Jason Eisner. Parameter estimation for probabilistic finite-state transducers. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 1–8, 2002.

Francesco Fabiano, Vishal Pallagani, Marianna Bergamaschi Ganapini, Lior Horesh, Andrea Loreggia, Keerthiram Murugesan, Francesca Rossi, and Biplav Srivastava. Plan-SOFAI: A neuro-symbolic planning architecture. In *Neuro-Symbolic Learning and Reasoning in the era of Large Language Models*, 2023. URL <https://openreview.net/forum?id=ORAhay0H4x>.

Hossein Rajaby Faghihi, Aliakbar Nafar, Chen Zheng, Roshanak Mirzaee, Yue Zhang, Andrzej Uszok, Alexander Wan, Tanawan Premsri, Dan Roth, and Parisa Kordjamshidi. Gluecons: A generic benchmark for learning under constraints, 2023. URL <https://arxiv.org/abs/2302.10914>.

Hossein Rajaby Faghihi, Aliakbar Nafar, Andrzej Uszok, Hamid Karimian, and Parisa Kordjamshidi. Prompt2demodel: Declarative neuro-symbolic modeling with natural language, 2024. URL <https://arxiv.org/abs/2407.20513>.

Chuyu Fang and Song-Chun Yu. Large language models are neurosymbolic reasoners. *arXiv preprint arXiv:2401.09334*, 2024. URL <https://arxiv.org/html/2401.09334v1>.

Michael Frazier and Leonard Pitt. Learning from entailment: An application to propositional horn sentences. In *Proceedings of the Tenth International Conference on International Conference on Machine Learning*, pages 120–127, 1993.

Eleonora Giunchiglia, Alex Tatomir, Mihaela Cătălina Stoian, and Thomas Lukasiewicz. Ccn+: A neuro-symbolic framework for deep learning with requirements. *International Journal of Approximate Reasoning*, 171:109124, 2024. ISSN 0888-613X. doi: <https://doi.org/10.1016/j.ijar.2024.109124>. URL <https://www.sciencedirect.com/science/article/pii/S0888613X24000112>. Synergies between Machine Learning and Reasoning.

Todd J. Green, Grigoris Karvounarakis, and Val Tannen. Provenance semirings. In *Proceedings of the Twenty-Sixth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, PODS '07, page 31–40, New York, NY, USA, 2007. Association for Computing Machinery. ISBN 9781595936851. doi: 10.1145/1265530.1265535. URL <https://doi.org/10.1145/1265530.1265535>.

Quan Guo, Hossein Rajaby Faghihi, Yue Zhang, Andrzej Uszok, and Parisa Kordjamshidi. Inference-masked loss for deep structured output learning. In Christian Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 2754–2761. International Joint Conferences on Artificial Intelligence Organization, 7 2020. doi: 10.24963/ijcai.2020/382. URL <https://doi.org/10.24963/ijcai.2020/382>. Main track.

Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2024. URL <https://www.gurobi.com>.

Frederick Hayes-Roth. Rule-based systems. *Communications of the ACM*, 28(9):921–932, 1985.

Pascal Hitzler and Md Kamruzzaman Sarker. Neuro-symbolic artificial intelligence: The state of the art. 2022.

Delower Hossain and Jake Y Chen. A study on neuro-symbolic artificial intelligence: Health-care perspectives, 2025. URL <https://arxiv.org/abs/2503.18213>.

Joy Hsu, Jiayuan Mao, Joshua B. Tenenbaum, and Jiajun Wu. What's left? concept grounding with logic-enhanced foundation models, 2023. URL <https://arxiv.org/abs/2310.16035>.

Jiani Huang, Ziyang Li, Bingshong Chen, Karan Samel, Mayur Naik, Le Song, and Xujie Si. Scallop: From probabilistic deductive databases to scalable differentiable reasoning. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 25134–25145. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/d367eef13f90793bd8121e2f675f0dc2-Paper.pdf.

Adam Ishay, Zhun Yang, and Joohyung Lee. Leveraging large language models to generate answer set programs. In Pierre Marquis, Tran Cao Son, and Gabriele Kern-Isbner, editors, *Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning, KR 2023*, Proceedings of the International Conference on Knowledge Representation and Reasoning, pages 374–383. Association for the Advancement of Artificial Intelligence, 2023. doi: 10.24963/kr.2023/37.

Prashani Jayasingha, Bogdan Iancu, and Johan Lilius. Neurosymbolic approaches in ai design – an overview. In *2025 IEEE Symposium on Trustworthy, Explainable and Responsible Computational Intelligence (CITREx Companion)*, pages 1–5, 2025. doi: 10.1109/CITRExCompanion65208.2025.10981497.

Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.

Daniel Kahneman. *Thinking, fast and slow*. macmillan, 2011.

Danial Kamali, Elham J. Barezi, and Parisa Kordjamshidi. Nesycoco: A neuro-symbolic concept composer for compositional generalization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(4):4184–4193, Apr. 2025. doi: 10.1609/aaai.v39i4.32439. URL <https://ojs.aaai.org/index.php/AAAI/article/view/32439>.

Henry Kautz. The third ai summer: Aaai robert s. engelmore memorial lecture. *AI Magazine*, 43(1):105–125, Mar. 2022. doi: 10.1002/aaai.12036. URL <https://ojs.aaai.org/aimagazine/index.php/aimagazine/article/view/19122>.

Angelika Kimmig, Guy Van den Broeck, and Luc De Raedt. An algebraic prolog for reasoning about possible worlds. *Proceedings of the AAAI Conference on Artificial Intelligence*, 25(1):209–214, Aug. 2011. doi: 10.1609/aaai.v25i1.7852. URL <https://ojs.aaai.org/index.php/AAAI/article/view/7852>.

Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything, 2023. URL <https://arxiv.org/abs/2304.02643>.

Phokion G. Kolaitis and Moshe Y. Vardi. On the expressive power of datalog: tools and a case study. In *Proceedings of the Ninth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, PODS '90, page 61–71, New York, NY, USA, 1990. Association for Computing Machinery. ISBN 0897913523. doi: 10.1145/298514.298542. URL <https://doi.org/10.1145/298514.298542>.

Parisa Kordjamshidi, Dan Roth, and Hao Wu. Saul: Towards declarative learning based programming. volume 2015, 12 2015.

Parisa Kordjamshidi, Daniel Khashabi, Christos Christodoulopoulos, Bhargav Mangipudi, Sameer Singh, and Dan Roth. Better call Saul: Flexible programming for learning and inference in NLP. In Yuji Matsumoto and Rashmi Prasad, editors, *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3030–3040, Osaka, Japan, December 2016. The COLING 2016 Organizing Committee. URL <https://aclanthology.org/C16-1285/>.

Parisa Kordjamshidi, Dan Roth, and Kristian Kersting. Declarative learning-based programming as an interface to ai systems. *Frontiers in Artificial Intelligence*, 5, 2019. URL <https://api.semanticscholar.org/CorpusID:195069123>.

Parisa Kordjamshidi, Dan Roth, and Kristian Kersting. Declarative learning-based programming as an interface to ai systems. *Frontiers in Artificial Intelligence*, Volume 5 - 2022, 2022. ISSN 2624-8212. doi: 10.3389/frai.2022.755361. URL <https://www.frontiersin.org/journals/artificial-intelligence/articles/10.3389/frai.2022.755361>.

Luis C. Lamb, Artur Garcez, Marco Gori, Marcelo Prates, Pedro Avelar, and Moshe Vardi. Graph neural networks meet neural-symbolic computing: A survey and perspective, 2021. URL <https://arxiv.org/abs/2003.00330>.

Guillaume Lample and François Charton. Deep learning for symbolic mathematics. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=S1eZYeHFDS>.

Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.

Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015. doi: 10.1038/nature14539. URL <https://doi.org/10.1038/nature14539>.

Bo Li, Peng Qi, Bo Liu, Shuai Di, Jingen Liu, Jiquan Pei, Jinfeng Yi, and Bowen Zhou. Trustworthy ai: From principles to practices. *ACM Comput. Surv.*, 55(9), January 2023a. ISSN 0360-0300. doi: 10.1145/3555803. URL <https://doi.org/10.1145/3555803>.

Tao Li and Vivek Srikumar. Augmenting neural networks with first-order logic. In Anna Korhonen, David Traum, and Lluís Màrquez, editors, *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 292–302, Florence, Italy,

July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1028. URL <https://aclanthology.org/P19-1028/>.

Ziyang Li, Jiani Huang, and Mayur Naik. Scallop: A language for neurosymbolic programming. *Proc. ACM Program. Lang.*, 7(PLDI), June 2023b. doi: 10.1145/3591280. URL <https://doi.org/10.1145/3591280>.

Ziyang Li, Jiani Huang, Jason Liu, Felix Zhu, Eric Zhao, William Dodds, Neelay Velingker, Rajeev Alur, and Mayur Naik. Relational programming with foundational models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(9):10635–10644, March 2024. ISSN 2159-5399. doi: 10.1609/aaai.v38i9.28934. URL <http://dx.doi.org/10.1609/aaai.v38i9.28934>.

Priscila Lima, Mariela Morveli-Espinoza, Glauzia K E Pereira, and Felipe França. Satyrus: A sat-based neuro-symbolic architecture for constraint processing. pages 137–142, 01 2005. doi: 10.1109/ICHIS.2005.97.

Runtao Liu, Chenxi Liu, Yutong Bai, and Alan L. Yuille. Clevr-ref+: Diagnosing visual reasoning with referring expressions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

Zhen Lu, Imran Afzidi, Hong Jin Kang, Ivan Ruchkin, and Xi Zheng. Surveying neuro-symbolic approaches for reliable artificial intelligence of things. *Journal of Reliable Intelligent Environments*, 10(3):257–279, 2024. doi: 10.1007/s40860-024-00231-1. URL <https://doi.org/10.1007/s40860-024-00231-1>.

Robin Manhaeve, Sebastijan Dumančić, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. Neural probabilistic logic programming in deepproblog. *Artificial Intelligence*, 298:103504, 2021. ISSN 0004-3702. doi: <https://doi.org/10.1016/j.artint.2021.103504>. URL <https://www.sciencedirect.com/science/article/pii/S0004370221000552>.

Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B. Tenenbaum, and Jiajun Wu. The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision, 2019. URL <https://arxiv.org/abs/1904.12584>.

Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, Xiao Wang, Xiaohua Zhai, Thomas Kipf, and Neil Houlsby. Simple open-vocabulary object detection with vision transformers, 2022. URL <https://arxiv.org/abs/2205.06230>.

Roshanak Mirzaee and Parisa Kordjamshidi. Disentangling extraction and reasoning in multi-hop spatial reasoning. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3379–3397, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.221. URL [https://aclanthology.org/2023.findings-emnlp.221/](https://aclanthology.org/2023.findings-emnlp.221).

Yatin Nandwani, Abhishek Pathak, Mausam, and Parag Singla. A primal dual formulation for deep learning with constraints. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/cf708fc1decf0337aded484f8f4519ae-Paper.pdf.

Allen Newell. Physical symbol systems. *Cognitive Science*, 4(2):135–183, 1980. ISSN 0364-0213. doi: [https://doi.org/10.1016/S0364-0213\(80\)80015-2](https://doi.org/10.1016/S0364-0213(80)80015-2). URL <https://www.sciencedirect.com/science/article/pii/S0364021380800152>.

Raymond Ng and V.S. Subrahmanian. Probabilistic logic programming. *Information and Computation*, 101(2):150–201, 1992. ISSN 0890-5401. doi: [https://doi.org/10.1016/0890-5401\(92\)90061-J](https://doi.org/10.1016/0890-5401(92)90061-J). URL <https://www.sciencedirect.com/science/article/pii/089054019290061J>.

Shan-Hwei Nienhuys-Cheng and Roland de Wolf. *What is inductive logic programming?* Springer, 1997.

OpenAI, Josh Achiam, and Steven Adler et al. Gpt-4 technical report, 2024. URL <https://arxiv.org/abs/2303.08774>.

Liangming Pan, Alon Albalak, Xinyi Wang, and William Wang. Logic-LM: Empowering large language models with symbolic solvers for faithful logical reasoning. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3806–3824, Singapore, December 2023a. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.248. URL <https://aclanthology.org/2023.findings-emnlp.248/>.

Liangming Pan, Alon Albalak, Xinyi Wang, and William Yang Wang. Logic-lm: Empowering large language models with symbolic solvers for faithful logical reasoning, 2023b. URL <https://arxiv.org/abs/2305.12295>.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library, 2019. URL <https://arxiv.org/abs/1912.01703>.

Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H. Miller, and Sebastian Riedel. Language models as knowledge bases?, 2019. URL <https://arxiv.org/abs/1909.01066>.

Aske Plaat, Max Van Duijn, Niki Van Stein, Mike Preuss, Peter Van der Putten, and Kees Joost Batenburg. Agentic large language models, a survey. *Journal of Artificial Intelligence Research*, 84, December 2025. ISSN 1076-9757. doi: 10.1613/jair.1.18675. URL <http://dx.doi.org/10.1613/jair.1.18675>.

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. Learning transferable visual models from natural language supervision, 2021. URL <https://arxiv.org/abs/2103.00020>.

Hossein Rajaby Faghihi, Quan Guo, Andrzej Uszok, Aliakbar Nafar, and Parisa Kordjamshidi. DomiKnowS: A library for integration of symbolic domain knowledge in deep learning. In Heike Adel and Shuming Shi, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 231–241, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-demo.27. URL <https://aclanthology.org/2021.emnlp-demo.27/>.

Saratha Sathasivam. Learning rules comparison in neuro-symbolic integration. *International Journal of Applied Physics and Mathematics*, 1(2):129, 2011.

Luciano Serafini and Artur S d’Avila Garcez. Learning and reasoning with logic tensor networks. In *Conference of the italian association for artificial intelligence*, pages 334–348. Springer, 2016.

Amir Shpilka, Amir Yehudayoff, et al. Arithmetic circuits: A survey of recent results and open questions. *Foundations and Trends® in Theoretical Computer Science*, 5(3–4):207–388, 2010.

David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016. doi: 10.1038/nature16961. URL <https://doi.org/10.1038/nature16961>.

Paul Smolensky, Moontae Lee, Xiaodong He, Wen-tau Yih, Jianfeng Gao, and li Deng. Basic reasoning with tensor product representations. 01 2016. doi: 10.48550/arXiv.1601.02745.

Jiankai Sun, Hao Sun, Tian Han, and Bolei Zhou. Neuro-symbolic program search for autonomous driving decision module design. In Jens Kober, Fabio Ramos, and Claire Tomlin, editors, *Proceedings of the 2020 Conference on Robot Learning*, volume 155 of *Proceedings of Machine Learning Research*, pages 21–30. PMLR, 16–18 Nov 2021. URL <https://proceedings.mlr.press/v155/sun21a.html>.

Erik F. Tjong Kim Sang and Fien De Meulder. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147, 2003. URL <https://aclanthology.org/W03-0419/>.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan

Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Bin Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023. URL <https://arxiv.org/abs/2307.09288>.

Pascal Van Hentenryck, Helmut Simonis, and Mehmet Dincbas. Constraint satisfaction using constraint logic programming. *Artificial Intelligence*, 58(1):113–159, 1992. ISSN 0004-3702. doi: [https://doi.org/10.1016/0004-3702\(92\)90006-J](https://doi.org/10.1016/0004-3702(92)90006-J). URL <https://www.sciencedirect.com/science/article/pii/000437029290006J>.

Joseph Weizenbaum. Eliza—a computer program for the study of natural language communication between man and machine. *Commun. ACM*, 9(1):36–45, January 1966. ISSN 0001-0782. doi: 10.1145/365153.365168. URL <https://doi.org/10.1145/365153.365168>.

Lionel Wong, Gabriel Grand, Alexander K. Lew, Noah D. Goodman, Vikash K. Mansinghka, Jacob Andreas, and Joshua B. Tenenbaum. From word models to world models: Translating from natural language to the probabilistic language of thought, 2023. URL <https://arxiv.org/abs/2306.12672>.

Fangzhi Xu, Zhiyong Wu, Qiushi Sun, Siyu Ren, Fei Yuan, Shuai Yuan, Qika Lin, Yu Qiao, and Jun Liu. Symbol-LLM: Towards foundational symbol-centric interface for large language models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13091–13116, Bangkok, Thailand, August 2024a. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.707. URL <https://aclanthology.org/2024.acl-long.707/>.

Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang, and Guy Van den Broeck. A semantic loss function for deep learning with symbolic knowledge. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5502–5511. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/xu18h.html>.

Xinyu Xu, Pan Wang, Shusen Dong, Ningyu Wu, Yue Zhang, and Baobao Chang. Faithful logical reasoning via symbolic chain-of-thought. *arXiv preprint arXiv:2405.18357*, 2024b. URL <https://arxiv.org/abs/2405.18357>.

Xiaocheng Yang, Bingsen Chen, and Yik-Cheung Tam. Arithmetic reasoning with LLM: Prolog generation & permutation. In Kevin Duh, Helena Gomez, and Steven Bethard,

editors, *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 699–710, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.nacl-short.61. URL <https://aclanthology.org/2024.nacl-short.61/>.

Liang Yao, Jiazen Peng, Chengsheng Mao, and Yuan Luo. Exploring large language models for knowledge graph completion, 2025. URL <https://arxiv.org/abs/2308.13916>.

Kexin Yi, Jiajun Wu, Chuang Gan, Antonio Torralba, Pushmeet Kohli, and Joshua B. Tenenbaum. Neural-symbolic vqa: Disentangling reasoning from vision and language understanding. In *Advances in Neural Information Processing Systems*, pages 1039–1050, 2018.

Hanlin Zhang, Jiani Huang, Ziyang Li, Mayur Naik, and Eric Xing. Improved logical reasoning of language models via differentiable symbolic programming, 2023. URL <https://arxiv.org/abs/2305.03742>.

Honghua Zhang, Po-Nien Kung, Masahiro Yoshida, Guy Van den Broeck, and Nanyun Peng. Adaptable logical control for large language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=58X9v92zRd>.

Danna Zheng, Mirella Lapata, and Jeff Z. Pan. How reliable are llms as knowledge bases? re-thinking facultality and consistency, 2024. URL <https://arxiv.org/abs/2407.13578>.