

A concept-based framework for building agents that can learn and reason across multiple domains.

BY JIAYUAN MAO, JOSHUA B. TENENBAUM, AND JIAJUN WU

Building Intelligent Agents with Neuro-Symbolic Concepts

ONE OF THE long-term goals of artificial intelligence (AI) is to build machines that can continually learn new knowledge from their experiences, ground these experiences in the physical world, and apply the knowledge to their reasoning across different tasks, modalities, and environments. The desired capability

of such agents includes, but is not limited to, describing perceived scenes, answering queries about scenes, making plans to achieve goals, and executing plans in the physical world. We want these machines to be able to learn and solve a wide variety of tasks across different environments, leveraging a feasible amount of data from multiple modalities.

In recent years, we have seen great success in neural-network-based

“end-to-end” learning methods. Still, most of them are tailored to particular tasks and environments, such as categorizing images of objects into a fixed set of labels, translating between languages, and playing video and board games. These systems are usually built on top of relatively simple and monolithic training and inference algorithms (for example, a single neural network trained by stochastic gradient descent). As a re-



» key insights

- We developed a concept-centric paradigm that supports building agents that can learn continually and reason flexibly.
- The agent acquires a vocabulary of neuro-symbolic concepts for objects, relations, and actions, represented through a combination of symbolic programs and neural networks. These concepts are grounded in sensory inputs and actuation outputs and can be composed to solve novel tasks using general-purpose reasoning and planning algorithms.
- The proposed framework offers several advantages, including data efficiency, compositional generalization, continual learning, and zero-shot transfer—key properties for general-purpose AI. These advantages have been demonstrated across applications in vision, language, and robotics.

sult, their success in domain-specific applications relies on the availability of large-scale datasets and computational resources on the particular task of interest. However, annotating high-quality data for reasoning, planning, and control in visual and physical domains is usually labor-intensive

and, in some cases, prohibitively costly. Therefore, there has been limited success in extending these methods to building embodied generalist agents across domains.

In this article, we emphasize the role of *concepts* in learning and reasoning. We draw on rich traditions from philosophy and cognitive science that identify concepts as the basic building blocks of thought (see, for example, readings from Margolis and Laurence²⁴). Humans acquire concepts from the interaction of our evolved cognitive architecture and built-in inductive biases with our experiences in the world, including both our direct percepts and what we learn socially and culturally from interacting and communicating with other humans. Our minds then compose these basic units to form sophisticated compound thoughts: beliefs, desires, and plans. One of the most powerful ways to construct a system of useful concepts for reasoning is to build them out of meanings acquired through language. In particular, we

can consider the granularity of concepts at the level of individual word meanings (of nouns, verbs, and so on) and treat their combinations in a way similar to how words can be combined into phrases and sentences in natural language. Related ideas of composing primitive units of thought have a long tradition in AI, dating back to inductive logic programming,²⁷ statistical relational learning,¹³ and probabilistic-logic programming.²⁰

Our technical proposal is to build a *neuro-symbolic concept* representation. Each concept is a discrete symbol (word or short phrase) that can be grounded onto subsets of the embodied environments. An object concept “orange” grounds to the sets of orange objects; an object relation concept “left” grounds to all object pairs (A,B) such that A is left of B; an action concept “put-left” grounds to all agent action sequences that move the object currently being held to a position on the left of the reference object. This includes pushing the object to the goal, holding it vigorously using tools,

Figure 1. Our framework for building intelligent agents by internalizing a vocabulary of “concepts,” which are represented as compositional programs and neural network embeddings. These concepts can be grounded in various domains: 2D images, videos, 3D scenes, and robotic actions. They can also be recombined based on different types of user queries: visually grounded questions, physical and causal reasoning questions, referring expressions, and manipulation instructions.

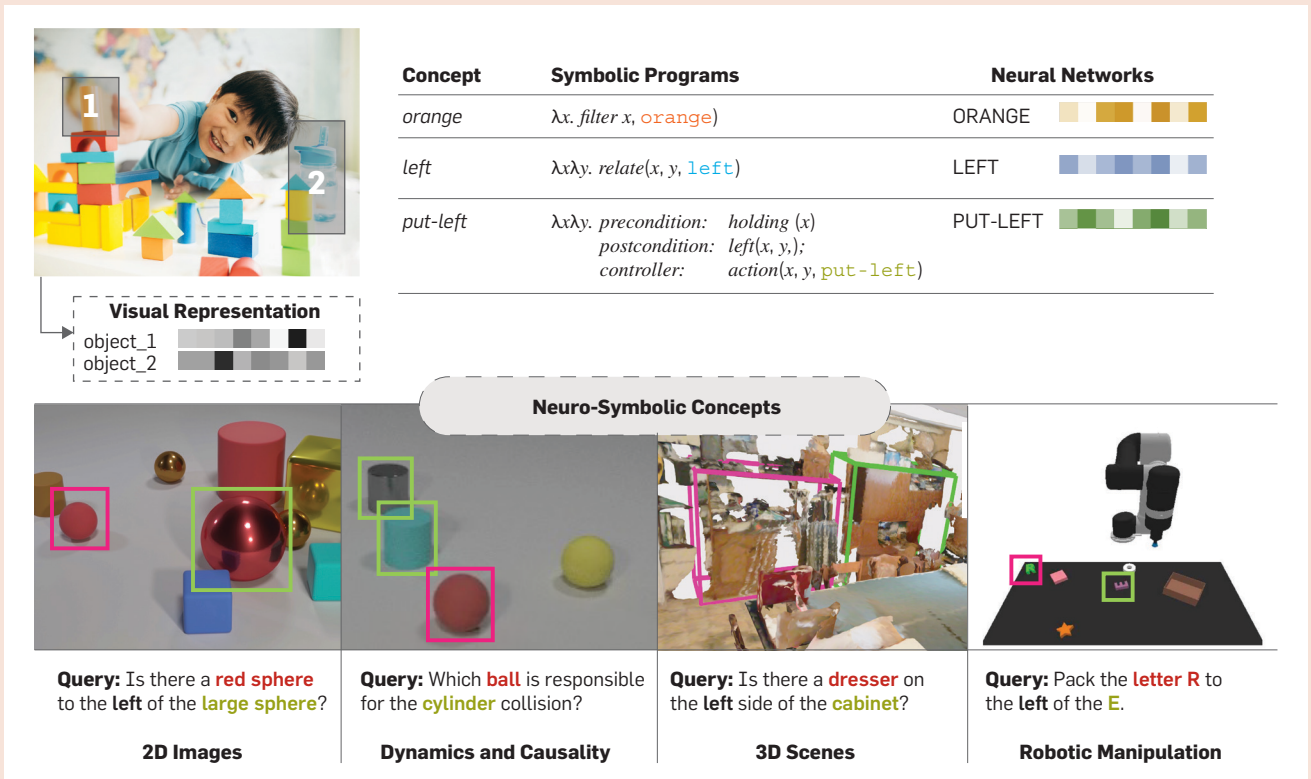


Figure 2. In a neuro-symbolic-concept-centric framework, different concepts such as object categories, properties, relations, and actions are represented as a combination of programmatic and neural representations. In (a), the neural representation connects the concept with sensory and actuation representations. In (b), different concepts can be combined to form new compound concepts.



and so forth. As illustrated in Figure 1, during reasoning, our *concept-centric* framework operates at a higher level of abstraction with this vocabulary of neuro-symbolic concepts.

Neuro-symbolic concepts have two advantages: They can be flexibly grounded on sensory and actuation modalities, and they have strong compositional generalization, since existing concepts can be structurally combined to form new ones. For instance, “orange,” “cylinder,” and “left” can be combined to form a novel concept, such as “Put the orange cylinder to the left of the bottle.” This compositionality naturally suggests a decomposition of the learning problem: We may individually learn how to recognize object shapes (cylinder, bottle), how to recognize object colors (orange), how to reason about object placement (left of the target), and how to move objects from one location to another. During inference time, we can recombine these learned concepts—shapes, colors, relations, and actions—to achieve the specified goals.

Compared with purely end-to-end learning methods, the inherent compositionality of neuro-symbolic concepts makes them much better suited for generalist agent learning. Compositionality simultaneously supports four properties required by generalist agents: data efficiency (because collecting labeled data is especially difficult); compositional generalization (the number of possible scenes and tasks can be exponentially large as the number of objects and their properties increases; we want systems that can generalize to unseen scenes and

unseen goals); continual learning (the system should be able to learn and adapt gradually); and transfer learning (we want to support transfer among different tasks).

In the rest of the article, we first offer a definition of neuro-symbolic concepts. Then we delve into a concrete framework of neuro-symbolic concept learning for visual scene understanding and showcase its ability to learn from little data, continually and generalizably, and be transferable. Finally, we discuss other applications of this approach.

Neuro-Symbolic Concepts

In order to systematically represent the grounding of concepts and how they can be composed, we formally represent each concept c as a tuple of:

$$c = \langle \text{parameter}, \text{program}, \text{neural - nets} \rangle.$$

In essence, our concept representation integrates neural network representations (for instance, vector embeddings), which ground concepts in visual and physical representations, and symbolic representations (particularly, parameterized programs), which characterize how various concepts can be combined. Shown in Figure 2a, the concept “orange” is represented as

$$\text{orange} = \langle \{x\}, \text{filter}(x, \text{ORANGE}), \{\text{ORANGE}\} \rangle,$$

where ORANGE denotes a vector embedding that can be used by the built-in *filter* function to classify orange objects. In practice, this *filter* function can be implemented by comput-

ing the cosine similarity between the neural representation of object x and ORANGE.²¹

Similarly, a relational concept (for example, a prepositional phrase) “left-of” can be defined as:

$$\text{left-of} = \langle \{x, y\}, \text{relate}(x, y, \text{LEFT-OF}), \{\text{LEFT-OF}\} \rangle.$$

The concept relates two objects x and y . The *relate* function will compute the cosine similarity between LEFT-OF and a pairwise representation between (x, y) to determine whether x is left of y in the scene. Actions (for example, verbs) would have three parts in their programs: a controller that can generate sequences of robot-control commands, and the preconditions and postconditions for the action. For example, an action “put-left-of” can be represented as:

$$\begin{aligned} \text{put-left-of} = \langle \{x, y\}, \{ \\ \text{pre} = \text{holding}(x), \\ \text{post} = \text{left-of}(x, y), \\ \text{controller} = \text{PUT-LEFT-OF}, \\ \{\text{PUT-LEFT-OF}\} \rangle. \end{aligned}$$

Here, the preconditions and postconditions of the action can be described with formulas composed from other object-level and relational concepts.

Illustrated in Figure 2b, this representation of concepts enables us to combine existing concepts adhering to symbolic functional composition rules, to form compound concepts such as *orange(x)* and *cylinder(x)* (orange cylinders), or more complex ones such as *orange(x)* and *bottle(y)* and *put-left-of(x, y)* (put the orange

object left of the bottle). To support the formal compositionality of different concepts, all parameters and outputs are typed with primitive types (including objects, events, actions, Booleans, and integers). For instance, object concepts are represented as functions that take the perceptual representation of an object as input and predict classification scores as output, indicating whether the object has the concept. Relational concepts are associated with classifiers that classify object pairs. Meanwhile, action concepts are linked with preconditions (circumstances under which

the action can be executed), post-conditions (the outcomes of executing the action), and controllers that generate agent actions based on the current perceptual state. Depending on the domain and the task, one can choose to implement different primitive operations (for example, *filter* in a visual recognition context).

Leveraging these neuro-symbolic concepts, we can efficiently and effectively learn the grounding of concepts in various domains and recombine them to solve different downstream tasks. For example, as shown in Figure 1, neuro-symbolic concepts can

be grounded in 2D images (such as object properties and object relations), videos (physical events and their relations), 3D scenes (object properties and viewpoint-dependent relations), and finally, robotic-manipulation tasks (object properties, relations, and actions that change them). Therefore, by recombining them through symbolic program structures, we can answer questions, resolve referring expressions, and interpret human instructions. We illustrate this idea in Figure 2. Specifically, we learn the grounding (classifiers and controllers) for individual visual and action

Figure 3. Three challenges in (a) compositional generalization, (b) continual learning of concepts for reasoning, and (c) transfer of learned concepts across domains.

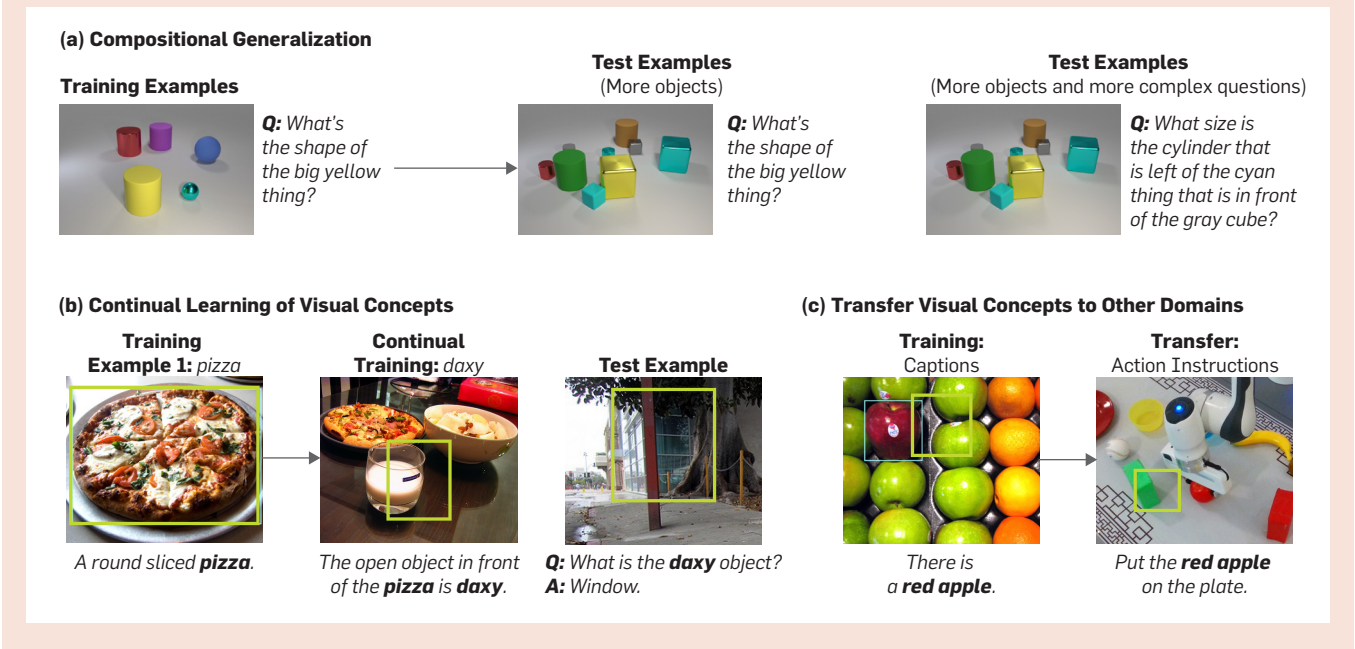
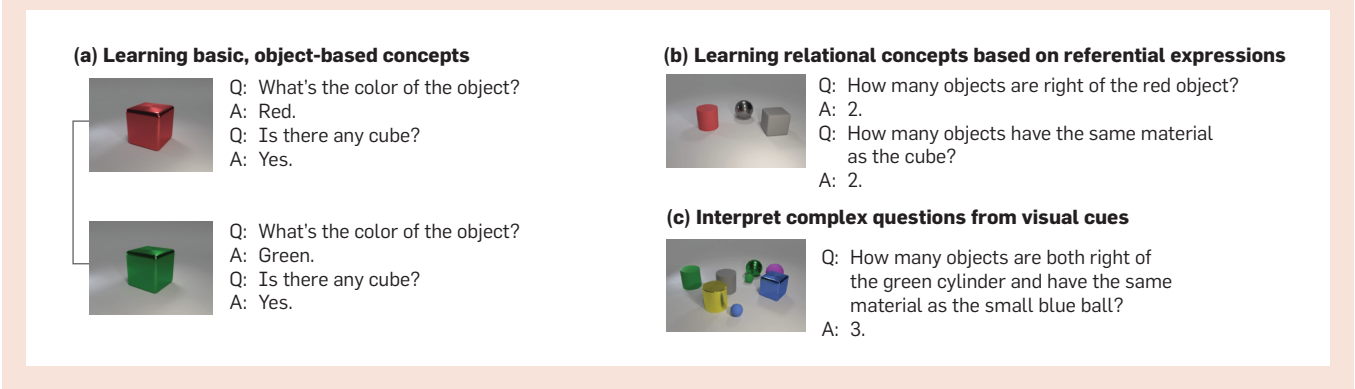


Figure 4. Humans learn visual concepts, words, and semantic parsing jointly and incrementally. (a) Learning visual concepts (red vs. green) starts from looking at simple scenes, reading simple questions, and reasoning over contrastive examples.¹² (b) Afterward, we can interpret referential expressions based on the learned object-based concepts and learn relational concepts (for example, on the right of, the same material as). (c) Finally, we can interpret complex questions from visual cues by exploiting the compositional structure.




concepts and recombine them following a hierarchical program that represents the meaning of the input user query: “Put the orange object left of the bottle.” In the following, we revisit the four important desiderata for generalist agent learning and illustrate how our neuro-symbolic-concept-centric paradigm fulfills all requirements.


Data efficiency. Since learning in embodied environments inevitably involves machine interaction with the physical environment and human annotations, minimizing the amount of data needed to learn a specific concept is crucial. Compared to a monolithic deep neural network, a concept-centric framework gains data efficiency primarily by leveraging modular structures of the learning task. For example, the complex concept “push the orange cylinder” can be decomposed into three individual concepts: two object concepts (“orange” and “cylinder”) and one action concept (“push”). This compositional structure introduces a strong prior into the learning algorithm, allowing it to disentangle the overall learning problem and to perform explicit multi-task learning across heterogeneous data sources (for example, learning the concept of “orange” from images and “push” from robot interaction datasets).

Compositional generalization. Compositionality is often grounded in different aspects across different domains. For example, in the visual concept learning domain (Figure 3a), compositional generalization is expected on at least two levels: the concept composition level (for example, “the big cylinder left of the yellow block”) and the scene composition level (generalization to scenes with a different number of objects compared to training examples). The advantage of concept-centric frameworks is primarily a contribution from the alignment between the composition structure of concepts and the structure of the domain. For example, explicitly reasoning about the set of “orange” and its subset “orange cylinder” in a visual scene makes the reasoning process robust to the total number of objects in a scene.

Continual learning. The demand



Neuro-symbolic concepts have two advantages: They can be flexibly grounded on sensory and actuation modalities, and they have strong compositional generalization, since existing concepts can be structurally combined to form new ones.



for continual concept learning and transfer learning arises due to two challenges. First, it is generally difficult to obtain high-quality “end-to-end” learning examples for embodied agents, that is, from raw perceptual input to robot-control commands. Therefore, a practical system should be able to learn from multiple sources of data, typically of different input-output specifications: unannotated videos, paired images and texts, human demonstrations of skills, and so on. Second, at deployment time, the machine should continually adapt itself to its environment, such as by learning new concepts (for example, an unseen breed of dog or a new type of dish) and new human preferences.

Zero-shot transfer. Unlike monolithic deep neural networks that usually require tuning all parameters while learning new concepts, the concept-centric framework (Figure 3b) naturally allows a flexible introduction of new concepts and adjustment to a single previously learned concept, thanks to the modular structures of concept composition. It also enables a zero-shot transfer of learned concepts across tasks and even domains, such as transferring learned object concepts from the task of image captioning (“The photo shows a dog”) to visual question answering (“How many dogs are there?”), and from the domain of visual concept learning (“apples”) to the domain of robotic manipulation (“Push the apples”; Figure 3c).

Learning Neuro-Symbolic Visual Concepts

Our framework for neuro-symbolic concept learning for visual scene understanding is motivated by how humans learn visual concepts by jointly understanding vision and language.¹² Consider the example shown in Figure 4a. Imagine someone with no prior knowledge of colors is presented with the images of the red and green cubes, paired with the questions and answers. They can easily identify the difference in objects’ visual appearance (in this case, color), and align it to the corresponding words in the questions and answers (*red* and *green*). Other object attributes (for example, shape) can be learned in a similar

way. Starting there, humans are able to inductively learn the correspondence between visual concepts and word semantics (for example, spatial relations and referential expressions, Figure 4b), and unravel compositional logic from complex questions assisted by the learned visual concepts (Figure 4c, also see Abend et al.¹).

This motivated us to build a framework that jointly learns visual perception, words, and semantic language parsing from images and question-answer pairs. Proposed in Mao et al.,²¹ a *neuro-symbolic concept learner (NS-CL)* learns all of these from natural supervision (that is, images and QA pairs), requiring no annotations on images or semantic programs for sentences. Instead, analogous to human concept learning, it learns via curriculum learning. The NS-CL starts by learning representations/concepts of individual objects from short questions (for example, "What's the color of the cylinder?") on simple scenes (less than three objects). By doing so, it learns object-based concepts such as colors and shapes. The NS-CL then learns relational concepts by leveraging these object-based concepts to interpret object referrals (for example, "Is there a box to the right of a cylinder?"). The model iteratively adapts to more complex scenes and highly compositional questions.

Shown in Figure 5, the NS-CL has three modules: a neural network-based perception module, a semantic parser to translate questions into executable programs, and a symbolic program executor. Given an input image, the visual perception module detects objects in the scene and extracts a deep, latent representation for each of them. The semantic parsing module translates an input question in natural language into an executable program, represented in a domain-specific language (DSL) designed for visual question answering (VQA). The DSL covers a set of fundamental operations for visual reasoning, such as filtering out objects with certain concepts or querying the attribute of an object. The generated programs have a hierarchical structure of symbolic, functional modules.

Next, based on the latent program recovered from the question in natu-

ral language, a symbolic program executor executes the program and derives the answer based on the object-based visual representation. Our program executor mainly contains two parts: the concept quantization module and a collection of deterministic functional modules. The concept quantization module classifies object attributes and relations, and the functional modules implement the logic of composing these classification results. To make the execution differentiable with regard to visual representations, we represent the intermediate results in a probabilistic manner: A set of objects is represented by a vector—the attention mask over all objects in the scene. Each element, $\text{Mask}_i \in [0, 1]$, denotes the probability that the i -th object of the scene belongs to the set. Figure 6 shows an illustrative execution trace of a program. The first *filter* operation outputs a mask of length 4 (there are in total four objects in the scene), with each element representing the probability that the corresponding object is selected (that is, the probability that each object is a green cube). The output "mask" on the objects will be fed into the next module (*relate* in this case) as input, and the execution of programs continues. The last module outputs the final answer.

Data efficiency. The NS-CL's modularized design enables interpretable, robust, and accurate visual reasoning. As shown in Figure 7a, it achieves state-of-the-art performance on the CLEVR dataset.¹⁷ More importantly, it enables data-efficient learning of concepts and combinatorial generalization with regard to both visual scenes and semantic programs. Figure 7a highlights comparisons with baselines that do not explicitly learn concepts: When trained on 10% of the CLEVR training data, the NS-CL achieves 98.9% accuracy on the test set, surpassing all baselines by 14%.

Compositional generalization. We also test our model for compositional generalization. As shown in Figure 7b, after being trained on scenes with a small number of objects and simple questions, the NS-CL directly generalizes to more complex scenes and questions, while all baselines show a significant performance drop.

Continual learning. Since our neuro-symbolic learning problem decomposes the problem into learning individual concepts, it naturally supports continual learning of new concepts. As a concrete implementation, in Mei et al.²⁵ we present a meta-learning framework for learning new visual concepts quickly from just one or a few examples, guided by multiple naturally occurring data streams: simultaneously looking at images, reading sentences that describe the objects in the scene, and interpreting supplemental sentences that relate the novel concept to other concepts. The system operates in a class-incremental manner,³⁵ where it continuously receives new examples of an unseen category and builds a new embedding for the novel category. The learned concepts support downstream applications, such as answering questions by reasoning about unseen images. Our model, FALCON, represents individual visual concepts, such as colors and shapes, as embeddings in a high-dimensional space. Given an input image and its paired sentence, our model first resolves the referential expression in the sentence and associates the novel concept with particular objects in the scene. Next, our model interprets supplemental sentences to relate the novel concept with other known concepts, such as "X has property Y" or "X is a kind of Y."¹⁴ Finally, it infers an optimal embedding for the novel concept that jointly 1) maximizes the likelihood of the observed instances in the image and 2) satisfies the relationships between the novel concepts and the known ones. We demonstrate the effectiveness of our model on both synthetic and real-world datasets.

Transfer learning. The learned visual concepts can also be used in other domains, such as image retrieval. With the visual scenes fixed, the learned visual concepts can be transferred directly into the new domain. We only need to learn the semantic parsing of natural language into the new DSL. To do so, we build a synthetic dataset for image retrieval. The dataset contains only simple captions: "There is an <object A> <relation> <object B>." (for example, "There is a box to the right of a cylinder"). The

Figure 5. The NS-CL uses neural symbolic reasoning to bridge the learning of visual concepts, words, and semantic parsing.

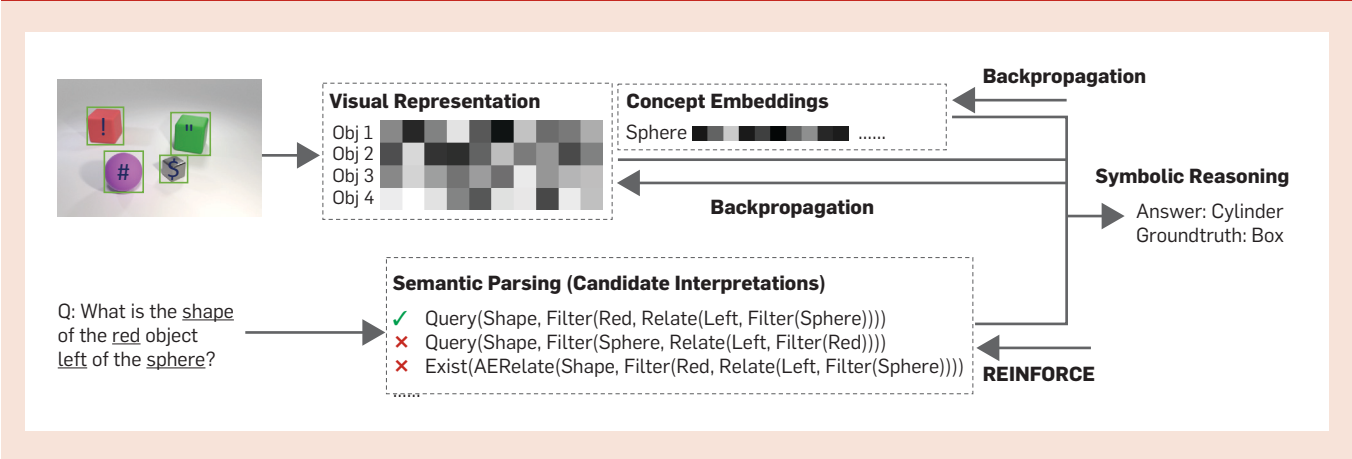


Figure 6. The neuro-symbolic execution procedure of a program based on the visual representation and concept embeddings.

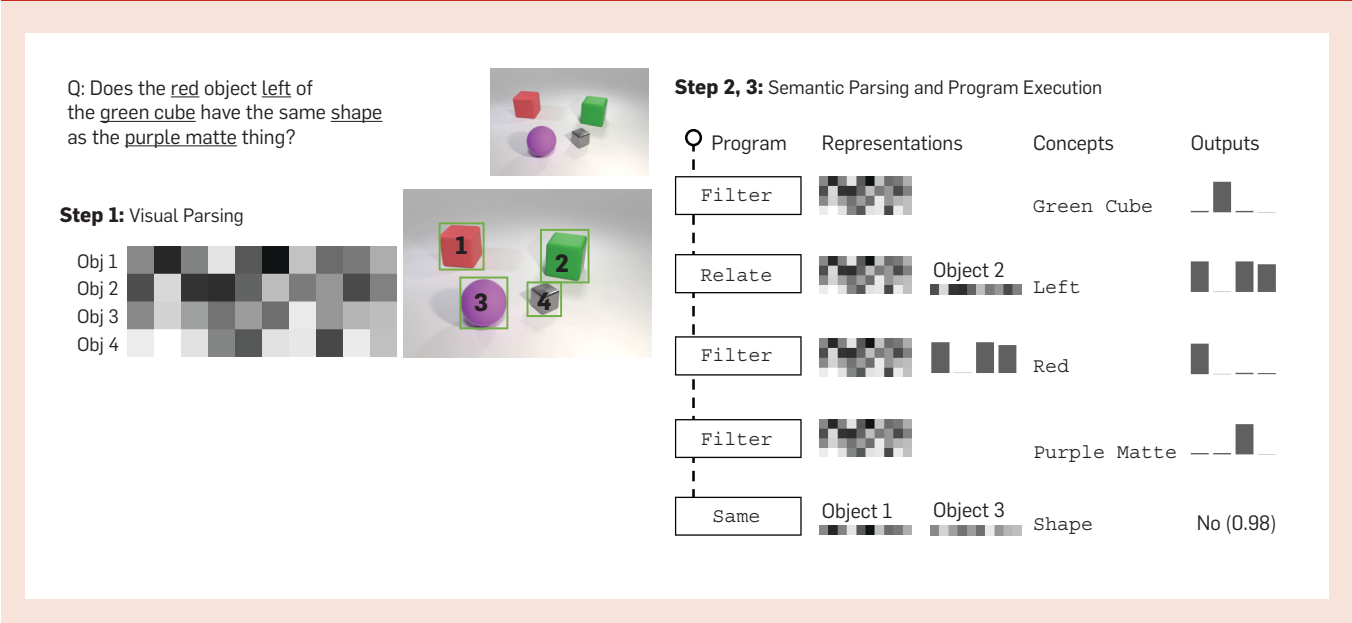
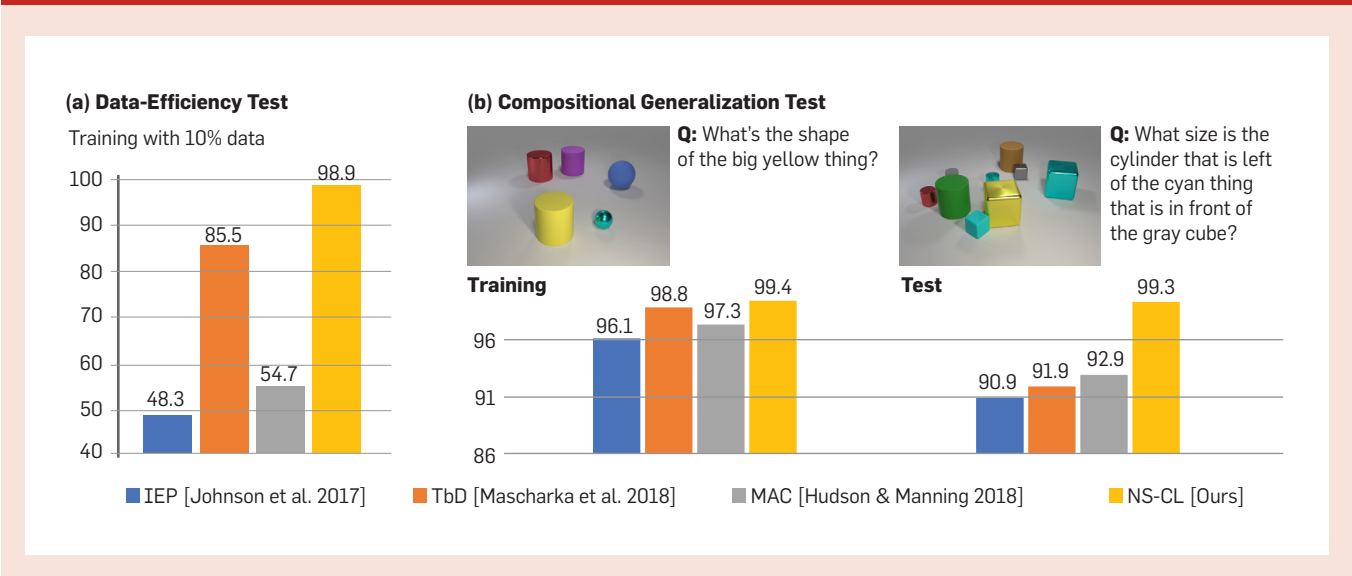


Figure 7. Data-efficiency test and compositional generalization test for the NS-CL.



semantic parser learns to extract corresponding visual concepts (for example, “box,” “right,” and “cylinder”) from the sentence. The program can then be executed on the visual representation to determine if the visual scene contains such relational triples. Note that this functionality cannot be directly implemented on the CLEVR VQA program domain, because questions such as “Is there a box to the right of a cylinder?” can be ambiguous if there exist multiple cylinders in the scene. Due to the entanglement of the visual representation with the specific DSL, baselines trained on CLEVR VQA cannot be applied directly to this task.

Applications

The design principles of the NS-CL—in particular, the visual grounding of concepts through neuro-symbolic reasoning—can naturally generalize to a large body of learning and reasoning tasks.

Accurate and robust image captioning. Wu et al.³⁸ implement a similar idea of neuro-symbolic concept learning to image-caption retrieval tasks. They used pretrained language parsers to translate captions into graphical representations composed of object categories, properties, and relationships. This kind of factorization not only leads to better performance in retrieving accurate descriptions of images, but also improves the robustness of the system with respect to captions that are similar to correct ones (for example, that differ only in one or two words) but inaccurate.

Video and counterfactual reasoning. A line of research^{5,9} has been extending the concept-learning framework to reasoning about physics. The object-centric nature of neuro-symbolic concept learners enables natural integration with learned physics models, which brings the capability to perform predictive and counterfactual reasoning. As an example, in Chen et al.,⁵ the authors ground concepts about physical objects and events from dynamic scenes and language. Building upon a neural object-centric representation, their model is simultaneously also trained to approximate the dynamic interaction among objects with neural

networks. Therefore, after training, it can not only detect and associate objects and events across the frames, but also make future and counterfactual predictions of object interactions (for example, “What will happen if we remove the red block from the scene?”). Later work, such as Ding et al.,⁹ further extends this capability to online inference of physical object properties.

3D concept grounding. The neuro-symbolic framework can also be applied to 3D representations.^{16,29} The variability of the 3D domain induces two fundamental challenges: 1) the expense of labeling and 2) the complexity of 3D grounded language. Hence, essential desiderata for models are to be data-efficient, generalize to different data distributions and tasks with unseen semantic forms, as well as ground complex language semantics (for example, view-point anchoring and multi-object reference: “Facing the chair, point to the lamp on its right”). To address these challenges, Prabhudesai et al.²⁹ propose disentangling different object properties in learning to achieve better performance in few-shot concept learning. Hsu et al.¹⁶ extend the neuro-symbolic concept learner model by introducing functional modules that effectively reason about high-arity relations (that is, relations among more than two objects), key in disambiguating objects in complex 3D scenes. This architecture enables significantly improved performance in terms of data efficiency and generalization, and demonstrates zero-shot transfer to a 3D question-answering task.

Human motion. Endo et al.¹¹ focus on designing models that conduct complex spatiotemporal reasoning over motion sequences. They propose a new framework for learning neural concepts of motion, attribute neural operators, and temporal relations. Unlike the 2D and 3D vision domains, where object segmentations can be readily extracted using preexisting object detectors, motion sequences lack a universal action segmentation methodology. Therefore, their system attempts to temporally localize and ground motion concepts jointly.

Robotic manipulation. Recall

that an important feature of neuro-symbolic concept learning methods is that the learnable modules associated with different concepts are naturally disentangled. Therefore, it directly supports the transfer of learned concepts to other tasks or even other domains (for example, from vision-language domains to robotic-manipulation domains). In Wang et al.³⁷ and Kalithasan et al.,¹⁸ the authors tackle the problem of learning robotic manipulation based on visual input. Both papers exploit the syntactic and semantic structures of language instructions to build robotic-manipulation algorithms composed of object-recognition models and action policies. Kalithasan et al.¹⁸ directly transfer the visual concepts learned by the neuro-symbolic concept learning on images to robotic manipulation, by learning additional object-movement policies with reinforcement learning. Wang et al.,³⁷ by contrast, leverage large-scale pretrained vision-language (VL) models for object property recognition. Compared to a conventional pretraining-finetuning pipeline for leveraging pretrained models for robotics, their method leads to more data-efficient learning and, more importantly, better zero-shot generalization in a variety of unseen objects and tasks.

General Discussion

We have presented a general framework for learning and reasoning that is applicable to various domains and tasks. By leveraging the neuro-symbolic concept representation, our system can continuously learn concepts from data streams in a data-efficient manner and programmatically compose its learned representations to solve new tasks, even previously unseen tasks. This capability allows us to learn and generalize concepts from diverse types of data streams, including image-caption data and robotic demonstrations, in order to solve complex tasks.


Our neuro-symbolic concept learning framework belongs to the broader paradigm of neuro-symbolic AI, a field where researchers explore the synergies between neural networks, symbolic reasoning methods, and probabilistic inference tools. The

idea of connecting neural networks with symbolic entities has its origins in early work on embedding symbolic relationships into vector representations. For example, early research^{7,30} demonstrates how to integrate logical reasoning and neural networks, improving the efficiency, interpretability, and controllability of learning systems.


Building on top of these high-level ideas, and in line with the NS-CL, many neuro-symbolic AI systems have been developed that combine symbolic reasoning mechanisms with neural networks for recognizing object properties and relationships, as well as predicting action commands in interactive environments. By combining perceptual capabilities with tools like forward-chaining theorem provers, answer-set programming solvers, and program-synthesis tools, these frameworks enable reasoning and planning in both visual and physical environments.

In application areas closely related to the NS-CL, Amizadeh et al.² introduce differentiable first-order logic frameworks for reasoning about objects in scenes. Further, Barbiero et al.⁴ combine neural network predictions with fuzzy-logic rule execution, and Shindo et al.³² use differentiable inductive logic programming to recover logic programs representing scene structures for scene reasoning tasks. These methods further formalize the interface between neural and symbolic components using tools such as probabilistic and real-valued logic, supporting formal interpretations of reasoning under uncertainty. These approaches excel not only in data efficiency and compositional generalization but also in offering interpretability of reasoning traces and inferred rules, which is critical in applications that demand transparency.

While much of the prior work has focused on learning simple rules or answering queries about object states and relationships, recent research has extended neuro-symbolic frameworks to more complex reasoning tasks. For example, Wang et al.³⁶ and Yang et al.⁴⁰ tackle visual puzzle-solving, such as Sudoku, using Boolean satisfiability solvers and answer set



By leveraging the neuro-symbolic concept representation, our system can continuously learn concepts from data streams in a data-efficient manner and programmatically compose its learned representations to solve new tasks.



programming tools. End-to-end neural networks for these complex tasks often require significantly more data, which can be difficult to obtain for many practical applications. However, training such complex neuro-symbolic models can be challenging, as backpropagation must occur over long chains of neuro-symbolic computations without intermediate supervision.

Moving toward more abstract, layout-based, and scene-level concepts—an emerging area in machine learning and visual reasoning—Shindo et al.³³ study how patterns of object placements can be learned from just a few examples. Similarly, Hsu et al.¹⁵ investigate reasoning about abstract concepts, such as mazes and treasure maps. These high-level concepts are difficult to interpret by state-of-the-art end-to-end systems like large vision-language models. These studies show that decomposing abstract concepts into smaller, more primitive entities using symbolic structures can significantly improve system performance.

Finally, it is important to highlight other key advantages of neuro-symbolic systems that we have not fully discussed in this article, such as interpretability, controllability, and the ability to integrate with external knowledge bases.^{4,34} These attributes, especially interpretability and safety, are critical in high-stakes decision-making contexts.³⁹


The idea of neuro-symbolic concepts is also closely related to the idea of neural module network compositions,³ since computationally, they are both paradigms for composing neural network modules to solve more complex tasks. However, they differ at both the conceptual and implementation levels. Specifically, in neural module networks, primitive neural networks define functions or transformations that can be applied to inputs, whereas in neuro-symbolic concepts, primitive neural networks handle the grounding of individual concepts, and the operations based on these concepts (for example, *filter* or *count*) are implemented as deterministic functions in DSLs. This disentanglement between grounding and reasoning brings

about significant improvements in data efficiency, compositional generalization, and transferability.


Of course, such improvements come at a cost: Many works on neuro-symbolic concept learning have the limitation of relying on a predefined DSL. This DSL encompasses primitive operators such as *filter* and *relate*, as well as concept symbols such as *orange* and *place*. In the following, we delve into these two parts.

In most domains, by combining object property primitives, relational primitives, and action primitives, along with simple set operations such as intersection, union, and counting, we can construct a highly capable system. Recently, there has been a growing interest in extending the primitive set and number operations to general programming languages, such as Python.⁸ This will significantly improve the expressiveness of the programs by including complex control flows such as loops and recursions. However, in general, this also introduces new challenges in the learning of concepts. Recall that in the NS-CL and many neuro-symbolic concept learning works, the concept representations are learned through backpropagation of the program execution trace. Although there has been much work on backpropagation against program execution traces,²⁸ complex control flows will inevitably introduce intractability of possible execution results, as well as gradient vanishing and explosion problems.

Relying on a predefined set of concepts can be too restrictive in many real-world applications. Three approaches have emerged to address this: grammar-based lexicon learning,²³ concept induction from past experiences, and more recent approaches based on large language models (LLMs). The high-level idea behind grammar-based lexicon learning is that instead of trying to associate each word or phrase with a predefined set of concept names (for example, mapping the word *orange* to a concept ORANGE), we construct the library of concepts by “converting” each word into a concept. This is roughly equivalent to discovering the syntax of each word. For example, if we see the word *orange*



While much of the prior work has focused on learning simple rules or answering queries about object states and relationships, recent research has extended neuro-symbolic frameworks to more complex reasoning tasks.



in a sentence and it is an adjective, then we immediately know the word *orange* should correspond to an object property concept, named ORANGE. This approach is called grammar-based lexicon learning because the system begins with a small set of universal grammar rules and jointly discovers new concepts from the text corpus while learning their grounding.

A second approach involves inducing new concepts from experiences. One method is to represent concepts as “theories composed of other concepts,” in line with the theory-theory of concepts.²⁶ For example, Das et al.⁶ induce new logical concepts by composing previously learned concepts using logic programs. Likewise, Shindo et al.³³ induce transferable scene-level concepts from a few examples, while Ellis et al.¹⁰ learns a “library” of functions built from primitive concepts that can be hierarchically recombined to form complex geometric and scene-level concepts. Another direction explores the invention of new object and relation concepts through contrastive learning.³¹

A third approach is to leverage LLMs such as GPT-4. With the success of LLMs in language-to-code translation, researchers explore the use of these models in extracting concept symbols from natural language queries.^{8,16} In particular, they leverage LLMs that have been trained on Internet-scale text and code corpora to translate natural language queries into programs with concept symbols. These concept symbols are not chosen from a given vocabulary, but are instead automatically generated by LLMs based on the user queries. For each concept symbol that appears in LLM-translated programs, a new concept representation will be initialized and learned.


There are many challenges and future directions for neuro-symbolic concept learning systems. Most approaches focus on relational concepts involving only two objects (or three, in the case of 3D concepts¹⁶), but there are more complex layout concepts and scene-level concepts (for example, mazes¹⁵) that involve many more objects and have variable arities, raising the question of how to handle such complexity. Additionally, while the

NS-CL uses curriculum learning, and some work explores different methods for curriculum construction,¹⁹ the automatic discovery or design of curricula that adapt as humans do when learning new concepts remains unsolved. So far, many concept learning systems can only operate in a pure “class-incremental” learning framework,³⁵ where new concepts build on previous ones; addressing the full “curriculum learning” setting, which requires revision or reversion of previously learned concepts, remains a significant challenge. Another underexplored topic is unsupervised concept learning beyond the current focus on supervised or semi-supervised methods. Finally, another direction for future work is to formalize reasoning under perceptual and other types of uncertainty by incorporating probabilistic inference methods. This includes tools like probabilistic logic programming and probabilistic programming languages.

Up to now, most neuro-symbolic concept learning systems have been developed for particular domains and tasks. For example, we have systems that can learn and reason with concepts for 2D images,²¹ 3D scenes,¹⁶ human motions,¹¹ video events,⁵ and robotic actions.^{22,37} However, they have been built in isolation, with limited shared knowledge among them. An important future direction for neuro-symbolic concept learning is the development of scalable, cross-domain concept libraries. For example, the concept “close to” has grounded meanings across different domains, but the core and abstract notion of distance metrics is really shared across all domains and modalities. From an engineering perspective, similar to recent LLMs such as OpenAI’s GPT and Google’s Gemini, which can solve a wider range of tasks in the language domain based on user instructions, building unified concept representations across domains and modalities would enable us to tackle a broader spectrum of embodied AI problems, spanning from perception to action. From a scientific perspective, connecting and unifying concept representations across domains and modalities could not only bring better data efficiency in learning, but also

enable the grounding of concepts in more abstract scenarios. Ultimately, achieving this level of abstraction and unification would mark a fundamental step toward building truly intelligent systems that can reason, adapt, and act across the vast complexity of the real world.

Acknowledgments

This work is in part supported by AFOSR YIP FA9550-23-1-0127, FA9550-22-1-0387, ONR N00014-23-1-2355, ONR YIP N00014-24-1-2117, ONR MURI N00014-22-1-2740, and NSF RI #2211258, and an AI2050 Senior Fellowship. 

References

- Abend, O. et al. Bootstrapping language acquisition. *Cognition* 164 (2017), 116–143.
- Amizadeh, S. et al. Neuro-symbolic visual reasoning: Disentangling ‘visual’ from ‘reasoning’. In *Intern. Conf. on Machine Learning* (2020).
- Andreas, J., Rohrbach, M., Darrell, T., and Klein, D. Neural module networks. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition* (2016).
- Barbiero, P. et al. Interpretable neural-symbolic concept reasoning. In *Intern. Conf. on Machine Learning* (2023).
- Chen, Z. et al. Grounding physical concepts of objects and events through dynamic visual reasoning. In *Intern. Conf. on Learning Representations* (2021).
- Das, M., Ramanan, N., Doppa, J.R., and Natarajan, S. Few-shot induction of generalized logical concepts via human guidance. *Frontiers in Robotics and AI* 7 (2020), 122.
- Demeester, T., Rocktäschel, T., and Riedel, S. Lifted rule injection for relation embeddings. In *Conf. on Empirical Methods in Natural Language Processing* (2016).
- Dídac, S., Menon, S., and Vondrick, C. ViperGPT: Visual Inference via Python execution for reasoning. In *IEEE/CVF Intern. Conf. on Computer Vision* (2023).
- Ding, M. et al. Dynamic visual reasoning by learning differentiable physics models from video and language. In *Advances in Neural Information Processing Systems* (2021).
- Ellis, K. et al. DreamCoder: Growing generalizable, interpretable knowledge with wake-sleep Bayesian program learning. *Philosophical Transactions of the Royal Society A* 381, 2251 (2023), 20220050.
- Endo, M., Hsu, J., Li, J., and Wu, J. Motion question answering via modular motion programs. In *Intern. Conf. on Machine Learning* (2023).
- Fazly, A., Alshahi, A., and Stevenson, S. A probabilistic computational model of cross-situational word learning. *Cognitive Science* 34, 6 (2010), 1017–1063.
- Friedman, N., Getoor, L., Koller, D., and Pfeffer, A. Learning probabilistic relational models. In *Intern. Joint Conf. on Artificial Intelligence* (1999).
- Han, C. et al. Visual concept metaconcept learning. In *Advances in Neural Information Processing Systems* (2019).
- Hsu, J. et al. What makes a maze look like a maze? In *Intern. Conf. on Learning Representations* (2025).
- Hsu, J., Mao, J., and Wu, J. NS3D: Neuro-symbolic grounding of 3D objects and relations. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition* (2023).
- Johnson, J. et al. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition* (2017).
- Kalithasan, N. et al. Learning neuro-symbolic programs for language guided robot manipulation. In *IEEE Intern. Conf. on Robotics and Automation* (2023).
- Li, Q., Huang, S., Hong, Y., and Zhu, S.C. A competence-aware curriculum for visual concepts learning via question answering. In *European Conf. on Computer Vision* (2020).
- Manhaeve, R. et al. DeepProlog: Neural probabilistic logic programming. In *Advances in Neural Information Processing Systems* (2018).
- Mao, J. et al. The Neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. In *Intern. Conf. on Learning Representations* (2019).
- Mao, J., Lozano-Pérez, T., Tenenbaum, J.B., and Kaelbling, L. PDSketch: Integrated domain programming, learning, and planning. In *Advances in Neural Information Processing Systems* (2022).
- Mao, J. et al. Grammar-based grounded language learning. In *Advances in Neural Information Processing Systems* (2021).
- Margolis, E. and Laurence, S. *Concepts: Core Readings*. The MIT Press (1999).
- Mei, L. et al. FALCON: Fast visual concept learning by integrating images, linguistic descriptions, and conceptual relations. In *Intern. Conf. on Learning Representations* (2022).
- Morton, A. *Frames of Mind: Constraints on the Common-Sense Conception of the Mental*. Oxford University Press (1980).
- Muggleton, S. and Raedt, L.D. Inductive logic programming: Theory and methods. *The J. of Logic Programming* 19 (1994), 629–679.
- Petersen, F., Borgelt, C., Kuehne, H., and Deussen, O. Learning with algorithmic supervision via continuous relaxations. In *Advances in Neural Information Processing Systems* (2021).
- Prabhudesai, M. et al. Disentangling 3D prototypical networks for few-shot concept learning. In *Intern. Conf. on Learning Representations* (2021).
- Rocktäschel, T., Singh, S., and Riedel, S. Injecting logical background knowledge into embeddings for relation extraction. In *Annual Conf. of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (2015).
- Sha, J., Shindo, H., Kersting, K., and Dhami, D.S. Neuro-symbolic predicate invention: Learning relational concepts from visual scenes. *Neurosymbolic Artificial Intelligence* (2024), 1–26.
- Shindo, H., Pfanschilling, V., Dhami, D.S., and Kersting, K. ILP: Thinking visual scenes as differentiable logic programs. *Machine Learning* 112, 5, (2023), 1465–1497.
- Shindo, H., Pfanschilling, V., Dhami, D.S., and Kersting, K. Learning differentiable logic programs for abstract visual reasoning. *Machine Learning* 113, 11 (2024), 8533–8584.
- Skryagin, A., Ochs, D., Dhami, D.S., and Kersting, K. Scalable neural-probabilistic answer set programming. *J. of Artificial Intelligence Research* 78 (2023), 579–617.
- Van de Ven, G.M., Tuytelaars, T., and Tolias, A.S. Three types of incremental learning. *Nature Machine Intelligence* 4, 12 (2022), 1185–1197.
- Wang, P.W., Donti, P., Wilder, B., and Kolter, Z. SATNet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver. In *Intern. Conf. on Machine Learning* (2019).
- Wang, R. et al. Programmatically grounded, compositionally generalizable robotic manipulation. In *Intern. Conf. on Learning Representations* (2023).
- Wu, H. et al. Unified visual-semantic embeddings: Bridging vision and language with structured meaning representations. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition* (2019).
- Yang, W.C., Marra, G., Rens, G., and Raedt, L.D. Safe reinforcement learning via probabilistic logic shields. In *Intern. Joint Conf. on Artificial Intelligence* (2023).
- Yang, Z., Ishay, A., and Lee, J. NeurASP: Embracing neural networks into answer set programming. In *Intern. Joint Conf. on Artificial Intelligence* (2020).

Jiayuan Mao (jiayuanm@seas.upenn.edu) is an assistant professor at the University of Pennsylvania, Philadelphia, PA, USA.

Joshua B. Tenenbaum is a professor at the Massachusetts Institute of Technology, Cambridge, MA, USA.

Jiajun Wu is an assistant professor at Stanford University, Stanford, CA, USA.

© 2026 Copyright held by the owner/author(s).