Novelty Assessment Report

Paper: Emergent Discrete Controller Modules for Symbolic Planning in Transformers
PDF URL: https://openreview.net/pdf?id=14dlTHVxDX
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2026-01-05

Abstract

Transformers struggle with tasks that require symbolic planning loops, variable updates, and conditional branching, especially under length extrapolation. We introduce discrete controller modules that insert a small set of program primitives (ASSIGN, ADD, COMPARE, BRANCH) into Transformer blocks via a Gumbel–Softmax selector over operations and a compact program state of registers, flags, and optional memory. We prove that the augmented model can simulate any bounded-step program by mapping each primitive step to one controller step, and we bound the deviation of relaxed execution from its discrete trace by $O(\tau+\kappa^{-1})$ (selection temperature τ, comparison sharpness κ). Empirically, the controller-augmented Transformer achieves strong length generalization on algorithmic benchmarks (Sorting, Sum-of-List, BFS), improving longest-length accuracy by up to 20–40 points over strong baselines, and yields consistent gains on symbolic QA (DROP) and program-synthesis-style tasks (RobustFill) with reduced compositionality drop-off. The learned execution is interpretable: operation traces align with ground truth, register roles are linearly decodable, and targeted knockouts cause localized accuracy losses. The approach adds only $\sim$5–7% FLOPs and can be applied sparsely (every P-th layer).

Core Task Landscape

This paper addresses: **Length Generalization in Algorithmic Reasoning with Symbolic Planning**

A total of **50 papers** were analyzed and organized into a taxonomy with **21 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Neural-Symbolic Integration for Planning**
- **Symbolic Planning and Search Methods**
- **Neural Architectures for Algorithmic Reasoning**
- **Symbolic Knowledge Representation and Reasoning**
- **Generalization and Abstraction in Planning**

Complete Taxonomy Tree

- Length Generalization in Algorithmic Reasoning with Symbolic Planning Survey Taxonomy
- Neural-Symbolic Integration for Planning
  - Task and Motion Planning with Neural Components
  - Language-Grounded Task and Motion Planning (3 papers)
    - [1] Text2Motion: from natural language instructions to feasible plans (Kevin Lin, 2023) View paper
    - [3] Inner Monologue: Embodied Reasoning through Planning with Language Models (Huang WenLong, 2022) View paper
    - [8] Compositional Foundation Models for Hierarchical Planning (Ajay, 2023) View paper
  - Hierarchical Scene Graph Planning (2 papers)
    - [4] Hierarchical planning for long-horizon manipulation with geometric and symbolic scene graphs (Yi-Feng Zhu, 2021) View paper
    - [5] Fast Task Planning with Neuro-Symbolic Relaxation (Du Qi-wei, 2025) View paper
  - Learned Operators and Skills for TAMP (4 papers)
    - [7] Learning neuro-symbolic skills for bilevel planning (Athalye, 2022) View paper
    - [9] Learning Type-Generalized Actions for Symbolic Planning (Daniel Tanneberg, 2023) View paper
    - [18] Learning Symbolic Operators for Task and Motion Planning (Tom Silver, 2021) View paper
    - [19] Learning Neuro-Symbolic Abstractions for Robot Planning and Learning (Naman Shah, 2024) View paper
  - Visual Reasoning for Sequential Manipulation (2 papers)
    - [20] Learning to solve sequential physical reasoning problems from a scene image (Driess, 2021) View paper
    - [27] Deep Visual Reasoning: Learning to Predict Action Sequences for Task and Motion Planning from an Initial Scene Image (Driess, 2020) View paper
  - Neuro-Symbolic Reasoning for Long-Horizon Tasks
  - Abductive and Imitation-Based Neuro-Symbolic Learning (2 papers)
    - [11] Learning for Long-Horizon Planning via Neuro-Symbolic Abductive Imitation (Jiejing Shao, 2024) View paper
    - [17] Abductive Learning for Neuro-Symbolic Grounded Imitation (Jie-Jing Shao, 2025) View paper
  - LLM-Based Neuro-Symbolic Planning (4 papers)
    - [15] CoreThink: A Symbolic Reasoning Layer to reason over Long Horizon Tasks with LLMs (Jay Vaghasiya, 2025) View paper
    - [25] Metagent-P: A Neuro-Symbolic Planning Agent with Metacognition for Open Worlds (Li Xiaodong, 2025) View paper
    - [28] Fast and Accurate Task Planning using Neuro-Symbolic Language Models and Multi-Level Goal Decomposition (Minseo Kwon, 2024) View paper

        - [36] JARVIS: A Neuro-Symbolic Commonsense Reasoning Framework for Conversational Embodied Agents (Zheng, 2022) View paper
    - Skill-Centric and Hierarchical Decomposition (3 papers)
        - [14] MOSAIC: A Skill-Centric Algorithmic Framework for Long-Horizon Manipulation Planning (Shaoul, 2025) View paper
        - [31] Rhh-lgp: Receding horizon and heuristics-based logic-geometric programming for task and motion planning (Braun, 2022) View paper
        - [35] Planning with spatial-temporal abstraction from point clouds for deformable object manipulation (Lin, 2022) View paper
    - Embodied Question Answering and Conversational Agents (2 papers)
        - [13] NeuS-QA: Grounding Long-Form Video Understanding in Temporal Logic and Neuro-Symbolic Reasoning (Shah, 2025) View paper
        - [16] Enter the Mind Palace: Reasoning and Planning for Long-term Active Embodied Question Answering (Ginting, 2025) View paper
    - Programmatic and Code-Based Neuro-Symbolic Approaches (3 papers)
    - [26] Codeplan: Unlocking reasoning potential in large language models by scaling code-form planning (J Wen, 2024) View paper
    - [29] Using planning to improve semantic parsing of instructional texts (Vanya Cohen, 2023) View paper
    - [45] Sketch-Plan-Generalize: Learning and Planning with Neuro-Symbolic Programmatic Representations for Inductive Spatial Concepts (Kalithasan, 2024) View paper
- Symbolic Planning and Search Methods
    - Model-Based Symbolic Planning (3 papers)
    - [2] Symbolic Generalization for On-line Planning (Zhengzhu Feng, 2022) View paper
    - [23] Weak, strong, and strong cyclic planning via symbolic model checking (A. Cimatti, 2003) View paper
    - [37] A new representation and associated algorithms for generalized planning (Siddharth Srivastava, 2011) View paper
    - Embodied and Robotic Symbolic Planning (3 papers)
    - [30] Egocentric planning for scalable embodied task achievement (Liu Xiaotian, 2023) View paper
    - [44] Task and Motion Planning for Humanoid Loco-Manipulation (DhÃ©din, 2025) View paper
    - [46] PrimHOI: Compositional Human-Object Interaction via Reusable Primitives (K Jia, 2025) View paper
    - Parallel and Interactive Planning (2 papers)
    - [48] Keep the planner in the loop: parallel planning and execution using Large Language Models (Capitanelli, 2024) View paper
    - [50] Knowledge-based proof planning (Erica Melisâ, 1999) View paper
- Neural Architectures for Algorithmic Reasoning
    - Transformer-Based Algorithmic Learning ★ (2 papers)
    - [0] Emergent Discrete Controller Modules for Symbolic Planning in Transformers (Anon et al., 2026) View paper
    - [12] Beyond a*: Better planning with transformers via search dynamics bootstrapping (Lehnert, 2024) View paper
    - Memory-Augmented Neural Architectures (2 papers)
    - [32] Exploring Patterns of Algorithmic Generalization in Deep Learning (Anil, 2025) View paper
    - [39] Learning Algorithmic Solutions to Symbolic Planning Tasks with a Neural Computer (Tanneberg, 2022) View paper
    - Reinforcement Learning with Symbolic Guidance (3 papers)
    - [33] Learning from Less: Guiding Deep Reinforcement Learning with Differentiable Symbolic Planning (Ye, 2025) View paper
    - [34] Option Discovery for Autonomous Generation of Symbolic Knowledge (Gabriele Sartor, 2021) View paper
    - [38] Enhancing Hierarchical Reinforcement Learning with Symbolic Planning for Long-Horizon Tasks (J Zhang, 2025) View paper
- Symbolic Knowledge Representation and Reasoning
    - Scalable Knowledge Base Reasoning (1 papers)
    - [10] Scalable Neural Methods for Reasoning With a Symbolic Knowledge Base (William W. Cohen, 2022) View paper
    - Symbolic Constraint Analysis and Reasoning (2 papers)
    - [22] Towards Learning to Reason: Comparing LLMs with Neuro-Symbolic on Arithmetic Relations in Abstract Reasoning (Hersche, 2024) View paper
    - [24] Worst-Case Symbolic Constraints Analysis and Generalisation with Large Language Models (Noller, 2025) View paper
    - Multi-Target and Multilingual Symbolic Reasoning (2 papers)
    - [40] Enhancing Large Language Models with Neurosymbolic Reasoning for Multilingual Tasks (Agrawal Ameeta, 2025) View paper
    - [41] Can Large Reasoning Models do Analogical Reasoning under Perceptual Uncertainty? (Camposampiero, 2025) View paper
- Generalization and Abstraction in Planning
    - Scene Abstraction for Generalization (2 papers)
    - [6] Core challenges in embodied vision-language planning (Francis, 2022) View paper
    - [42] Scene Abstraction for Generalizable Long-Horizon Robot Planning (Han, 2024) View paper
    - Goal Inference and Predicate Learning (2 papers)
    - [21] Plan-SOFAI: A neuro-symbolic planning architecture (F Fabiano, 2023) View paper
    - [49] GOALNET: Interleaving Neural Goal Predicate Inference with Classical Planning for Generalization in Robot Instruction Following (Shreya Sharma, 2024) View paper
    - Compositional and Narrative Reasoning (2 papers)
    - [43] Symbolic planning and control of robot motion (C Belta, 2007) View paper
    - [47] Integrating Cognitive, Symbolic, and Neural Approaches to Story Generation: A Review on the METATRON Framework (C Hiram, 2025) View paper

Narrative

Core task: length generalization in algorithmic reasoning with symbolic planning. This field investigates how systems can learn to solve reasoning and planning problems that extend beyond the training distribution in terms of sequence length or problem complexity. The taxonomy reveals several complementary research directions. Neural-Symbolic Integration for Planning explores hybrid architectures that combine learned representations with symbolic reasoning modules, as seen in works like Inner Monologue[3] and Neuro-Symbolic Skills[7]. Symbolic Planning and Search Methods focus on classical planning techniques and their extensions, including approaches like Beyond A-Star[12] that push traditional search boundaries. Neural Architectures for Algorithmic Reasoning examines how neural models—particularly transformers—can be designed or trained to capture algorithmic structure, with studies such as Algorithmic Generalization Patterns[32] analyzing what enables length generalization. Symbolic Knowledge Representation and Reasoning addresses how to encode domain knowledge and constraints, exemplified by Symbolic Constraints LLMs[24]. Finally, Generalization and Abstraction in Planning investigates hierarchical decomposition and abstraction mechanisms that facilitate transfer across problem scales, including works like Sketch-Plan-Generalize[45].

A central tension across these branches concerns the trade-off between end-to-end neural flexibility and the compositional guarantees of symbolic methods. Many recent efforts blend differentiable modules with discrete planning steps to achieve both data efficiency and systematic generalization. Discrete Controller Modules[0] sits within the Neural Architectures for Algorithmic Reasoning branch, specifically under Transformer-Based Algorithmic Learning, and emphasizes modular controller designs that can compose learned primitives for longer horizons. This contrasts with purely symbolic approaches like Symbolic Generalization Planning[2], which rely on explicit operator definitions, and with more integrated neuro-symbolic frameworks such as Neuro-Symbolic Relaxation[5] that soften discrete structures for gradient-based learning. By focusing on discrete modular components, Discrete Controller Modules[0] aims to preserve interpretability and compositionality while leveraging neural pattern recognition, positioning it as a middle ground between fully learned and fully symbolic paradigms.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. Beyond a*: Better planning with transformers via search dynamics bootstrapping

Authors: Lehnert, Lucas, Sukhbaatar, Sainbayar, Lucas Lehnert, et al. (18 authors total) | **Year/Venue**: 2024 | **URL**: View paper

Abstract

While Transformers have enabled tremendous progress in various application settings, such architectures still trail behind traditional symbolic planners for solving complex decision making tasks. In this work, we demonstrate how to train Transformers to solve complex planning tasks. This is accomplished by training an encoder-decoder Transformer model to predict the search dynamics of the A^* search algorithm. We fine tune this model to obtain a Searchformer, a Transformer model that optimally...

Relationship Analysis

Both papers belong to the Transformer-Based Algorithmic Learning category, focusing on training Transformers to solve algorithmic tasks through learning computational patterns. While the original paper introduces discrete controller modules (ASSIGN, ADD, COMPARE, BRANCH) that explicitly encode symbolic planning primitives within Transformer layers to achieve length generalization, the candidate paper (Searchformer) trains Transformers to predict the search dynamics of the A algorithm by generating execution traces as token sequences. The key difference is that the original paper augments Transformers with explicit programmatic control-flow semantics and maintains internal program state (registers, flags, memory), whereas Searchformer relies on next-token prediction of A's search process without explicit symbolic modules, using search dynamics bootstrapping to iteratively improve planning efficiency.

Contributions Analysis

Overall novelty summary. The paper introduces discrete controller modules that augment Transformers with program primitives (ASSIGN, ADD, COMPARE, BRANCH) for algorithmic reasoning tasks requiring symbolic planning loops and length extrapolation. It sits within the Transformer-Based Algorithmic Learning leaf of the taxonomy, which contains only two papers total. This is a relatively sparse research direction compared to more crowded areas like Task and Motion Planning (ten papers across four subcategories) or LLM-Based Neuro-Symbolic Planning (four papers). The sparsity suggests that explicit modular controller designs for Transformers remain underexplored, with most work either pursuing end-to-end neural learning or full neuro-symbolic integration rather than discrete programmatic augmentation.

The taxonomy reveals several neighboring approaches that address length generalization differently. Memory-Augmented Neural Architectures (two papers) explore external memory mechanisms without explicit program primitives, while Reinforcement Learning with Symbolic Guidance (three papers) incorporates symbolic planning through policy learning rather than architectural modules. The Programmatic and Code-Based Neuro-Symbolic Approaches branch (three papers) generates executable programs but typically as outputs rather than integrated execution modules. The paper's approach diverges by embedding discrete controllers directly into Transformer blocks, creating a hybrid that maintains differentiability while enforcing programmatic structure—a design choice that distinguishes it from both pure neural architectures and code-generation frameworks.

Among twenty candidates examined across three contributions, none were found to clearly refute the core claims. The discrete controller module design examined ten candidates with zero refutable matches, suggesting limited prior work on this specific architectural pattern. The formal expressivity result for bounded imperative programs examined three candidates without finding overlapping theoretical analysis. The tight integration mechanism examined seven candidates, again with no clear precedents. This pattern indicates that while the broader problem of length generalization in algorithmic reasoning is well-studied, the specific combination of discrete program primitives, Gumbel-Softmax selection, and register-based state management within Transformer blocks appears novel within the examined literature scope.

The analysis covers top-twenty semantic matches and does not constitute an exhaustive survey of all Transformer augmentation methods or program synthesis techniques. The taxonomy structure suggests the paper occupies a relatively unexplored niche between pure neural algorithmic learning and full neuro-symbolic integration, though the limited search scope means potentially relevant work in adjacent areas (e.g., differentiable programming, neural module networks) may not have been captured. The contribution-level statistics consistently show no clear refutations, but this reflects the examined candidate set rather than definitive novelty across all related literature.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Discrete controller modules with program primitives for Transformers

Description: The authors propose augmenting Transformer blocks with learnable discrete controllers that execute symbolic operations (ASSIGN, ADD, COMPARE, BRANCH) on a latent program state consisting of registers, flags, and optional memory. The controller uses Gumbel-Softmax to select operations differentiably while maintaining interpretable, step-wise execution traces.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Investigating a Novel Transposon Attention Scaffold for Large Scale Transformer Reasoning Patterns

URL: View paper

Brief Assessment

Transposon Attention Scaffold[54] focuses on differentiable functions interacting with attention logits via smooth relaxation mechanisms and relational mapping consistency through symbolic manipulation exercises. This differs fundamentally from the original paper's discrete controller modules that execute symbolic operations (ASSIGN, ADD, COMPARE, BRANCH) using Gumbel-Softmax for differentiable selection while maintaining interpretable, step-wise execution traces with a latent program state.

---

2. Symbolic visual reinforcement learning: A scalable framework with object-level abstraction and differentiable expression search

URL: View paper

Brief Assessment

Symbolic Visual RL[53] focuses on reinforcement learning with symbolic expressions for visual control tasks using genetic programming and object-level abstractions, not on augmenting Transformer architectures with discrete controller modules for symbolic planning.

3. Attention (as Discrete-Time Markov) Chains

URL: View paper

Brief Assessment

Attention Markov Chains[56] interprets attention matrices as discrete-time Markov chains for token importance and segmentation, not discrete symbolic operations (ASSIGN, ADD, COMPARE, BRANCH) with program-like execution traces.

4. Transformer-Assisted Genetic Programming for Symbolic Regression

URL: View paper

Brief Assessment

Transformer Genetic Programming[51] focuses on genetic programming for symbolic regression using transformers as fitness evaluators, not on augmenting transformer architectures with discrete symbolic operations or program primitives like ASSIGN, ADD, COMPARE, BRANCH.

5. UniSymNet: A Unified Symbolic Network Guided by Transformer

URL: View paper

Brief Assessment

UniSymNet[59] focuses on symbolic regression using nested unary operators in symbolic networks, not on augmenting Transformers with discrete control-flow primitives (ASSIGN, ADD, COMPARE, BRANCH) for algorithmic reasoning tasks.

6. AI Reasoning in Deep Learning Era: From Symbolic AI to Neuralâ□□Symbolic AI

URL: View paper

Brief Assessment

Neural-Symbolic AI[57] is a survey paper that reviews the broader paradigm of integrating symbolic logic with neural computation. It does not propose discrete controller modules with specific program primitives (ASSIGN, ADD, COMPARE, BRANCH) for Transformers, nor does it address differentiable operation selection via Gumbel-Softmax within Transformer blocks. The candidate focuses on taxonomizing existing neural-symbolic approaches rather than introducing a novel architectural component for symbolic planning in Transformers.

7. Terminating Differentiable Tree Experts

URL: View paper

Brief Assessment

Differentiable Tree Experts[55] focuses on tree operations using tensor product representations and mixture of experts, not on discrete symbolic operations (ASSIGN, ADD, COMPARE, BRANCH) with program state (registers, flags, memory) as proposed in the original paper.

8. Differentiable tree operations promote compositional generalization

URL: View paper

Brief Assessment

Differentiable Tree Operations[60] focuses on tree-to-tree transformations using tensor product representations and predefined symbolic tree operations (car, cdr, cons), not on augmenting Transformers with discrete controllers for general algorithmic reasoning with registers, flags, and memory.

9. Scallop: From probabilistic deductive databases to scalable differentiable reasoning

URL: View paper

Brief Assessment

Scallop[58] focuses on probabilistic deductive databases and datalog execution for neural-symbolic reasoning, not on augmenting Transformer blocks with discrete controller modules that execute symbolic operations like ASSIGN, ADD, COMPARE, BRANCH.

10. Latent symbol lattices in probabilistic semiosis: An unconventional architectural mechanism for contextual modulation in large language models

URL: View paper

Brief Assessment

Latent Symbol Lattices[52] focuses on probabilistic semiosis and contextual modulation mechanisms in LLMs, not on discrete controller modules with explicit program primitives (ASSIGN, ADD, COMPARE, BRANCH) for algorithmic reasoning tasks.

Contribution 2: Formal expressivity result for bounded imperative programs

Description: The authors provide a theoretical guarantee showing that their controller-augmented Transformer can exactly simulate any program in a bounded imperative class with at most K primitive steps and loop bound B, requiring depth O(K+B). They also bound the approximation error under finite temperature and comparison sharpness parameters.

This contribution was assessed against **3 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. On the expressive power of programming languages

URL: View paper

Brief Assessment

Expressive Power Languages[62] discusses expressiveness of imperative extensions in programming languages generally, while the original paper provides a specific theoretical guarantee for controller-augmented Transformers simulating bounded imperative programs with explicit depth bounds O(K+B).

2. Turing Completeness of Bounded-Precision Recurrent Neural Networks

URL: View paper

Brief Assessment

Turing Completeness RNNs[61] focuses on bounded-precision recurrent neural networks and their Turing completeness properties, not on Transformer architectures augmented with discrete controller modules for simulating bounded imperative programs.

3. Symbolic Feedforward Networks for Probabilistic Finite Automata: Exact Simulation and Learnability

URL: View paper

Brief Assessment

Symbolic Feedforward Networks[63] focuses on simulating probabilistic finite automata using feedforward networks with stochastic matrix operations, not bounded imperative programs with loops and control flow primitives.

Contribution 3: Tight integration of controllers with attention and feed-forward layers

Description: The authors design a method to integrate the discrete controller directly into the Transformer architecture through state-to-token and token-to-state projections, allowing the program state to interact with the token stream via residual connections without requiring external memory modules or post-processing steps.

This contribution was assessed against **7 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. How Do Transformers Learn Variable Binding in Symbolic Programs?

URL: View paper

Brief Assessment

Variable Binding Transformers[65] studies how transformers learn variable binding in symbolic programs through residual stream manipulation and attention mechanisms, but does not propose architectural modifications that integrate discrete controllers with attention/feed-forward layers via state-to-token projections as the original paper does.

2. Deep Learning Models for Estimating Number of Lambda-Term Reduction Steps.

URL: View paper

Brief Assessment

Lambda-Term Reduction Steps[68] focuses on estimating reduction steps in lambda calculus using deep learning models (CNNs, LSTMs, Transformers) with one-hot-encoded term representations. It does not address integration of program state with attention mechanisms or controller modules within Transformer architectures.

3. Momentary Contexts: A Memory and Retrieval Approach for LLM Efficiency

URL: View paper

Brief Assessment

Momentary Contexts[66] focuses on memory retrieval and context reconstruction for LLM efficiency through external database access, not on integrating discrete program-state controllers within Transformer layers. The candidate's memory layer operates between embedding and decoding to retrieve external knowledge, whereas the original contribution embeds discrete control-flow primitives (assign, add, compare, branch) directly into Transformer blocks via state-to-token projections and residual connections.

4. ProTo: Program-Guided Transformer for Program-Guided Tasks

URL: View paper

Brief Assessment

ProTo[67] focuses on program-guided tasks using cross-attention and masked self-attention to integrate program semantics with task specifications, not on integrating discrete controller state with transformer layers via state-to-token projections as in the original paper.

5. Neural Abstract Interpretation

URL: View paper

Brief Assessment

Neural Abstract Interpretation[69] focuses on learning abstract transformers for program analysis domains (intervals, octagons), not on integrating discrete controllers into Transformer architectures for symbolic planning tasks.

6. Learned interpreters: structural and learned systematicity in neural networks for program execution

URL: View paper

Brief Assessment

Learned Interpreters[70] focuses on instruction pointer attention graph neural networks (ipa-gnn) for program execution tasks, not on integrating discrete controllers with Transformer attention/feed-forward layers via state-to-token projections.

7. Thinking Like Transformers

URL: View paper

Brief Assessment

Thinking Like Transformers[64] proposes a programming language (RASP) to model transformer computation abstractly, mapping attention and feed-forward into primitives. This differs from the original paper's tight integration of discrete controllers (assign, add, compare, branch) with transformer layers via state-to-token projections and residual connections for symbolic planning tasks.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Emergent Discrete Controller Modules for Symbolic Planning in Transformers View paper
- [1] Text2Motion: from natural language instructions to feasible plans View paper
- [2] Symbolic Generalization for On-line Planning View paper
- [3] Inner Monologue: Embodied Reasoning through Planning with Language Models View paper
- [4] Hierarchical planning for long-horizon manipulation with geometric and symbolic scene graphs View paper
- [5] Fast Task Planning with Neuro-Symbolic Relaxation View paper
- [6] Core challenges in embodied vision-language planning View paper
- [7] Learning neuro-symbolic skills for bilevel planning View paper
- [8] Compositional Foundation Models for Hierarchical Planning View paper
- [9] Learning Type-Generalized Actions for Symbolic Planning View paper
- [10] Scalable Neural Methods for Reasoning With a Symbolic Knowledge Base View paper

- [11] Learning for Long-Horizon Planning via Neuro-Symbolic Abductive Imitation View paper
- [12] Beyond a*: Better planning with transformers via search dynamics bootstrapping View paper
- [13] NeuS-QA: Grounding Long-Form Video Understanding in Temporal Logic and Neuro-Symbolic Reasoning View paper
- [14] MOSAIC: A Skill-Centric Algorithmic Framework for Long-Horizon Manipulation Planning View paper
- [15] CoreThink: A Symbolic Reasoning Layer to reason over Long Horizon Tasks with LLMs View paper
- [16] Enter the Mind Palace: Reasoning and Planning for Long-term Active Embodied Question Answering View paper
- [17] Abductive Learning for Neuro-Symbolic Grounded Imitation View paper
- [18] Learning Symbolic Operators for Task and Motion Planning View paper
- [19] Learning Neuro-Symbolic Abstractions for Robot Planning and Learning View paper
- [20] Learning to solve sequential physical reasoning problems from a scene image View paper
- [21] Plan-SOFAI: A neuro-symbolic planning architecture View paper
- [22] Towards Learning to Reason: Comparing LLMs with Neuro-Symbolic on Arithmetic Relations in Abstract Reasoning View paper
- [23] Weak, strong, and strong cyclic planning via symbolic model checking View paper
- [24] Worst-Case Symbolic Constraints Analysis and Generalisation with Large Language Models View paper
- [25] Metagent-P: A Neuro-Symbolic Planning Agent with Metacognition for Open Worlds View paper
- [26] Codeplan: Unlocking reasoning potential in large language models by scaling code-form planning View paper
- [27] Deep Visual Reasoning: Learning to Predict Action Sequences for Task and Motion Planning from an Initial Scene Image View paper
- [28] Fast and Accurate Task Planning using Neuro-Symbolic Language Models and Multi-Level Goal Decomposition View paper
- [29] Using planning to improve semantic parsing of instructional texts View paper
- [30] Egocentric planning for scalable embodied task achievement View paper
- [31] Rhh-lgp: Receding horizon and heuristics-based logic-geometric programming for task and motion planning View paper
- [32] Exploring Patterns of Algorithmic Generalization in Deep Learning View paper
- [33] Learning from Less: Guiding Deep Reinforcement Learning with Differentiable Symbolic Planning View paper
- [34] Option Discovery for Autonomous Generation of Symbolic Knowledge View paper
- [35] Planning with spatial-temporal abstraction from point clouds for deformable object manipulation View paper
- [36] JARVIS: A Neuro-Symbolic Commonsense Reasoning Framework for Conversational Embodied Agents View paper
- [37] A new representation and associated algorithms for generalized planning View paper
- [38] Enhancing Hierarchical Reinforcement Learning with Symbolic Planning for Long-Horizon Tasks View paper
- [39] Learning Algorithmic Solutions to Symbolic Planning Tasks with a Neural Computer View paper
- [40] Enhancing Large Language Models with Neurosymbolic Reasoning for Multilingual Tasks View paper
- [41] Can Large Reasoning Models do Analogical Reasoning under Perceptual Uncertainty? View paper
- [42] Scene Abstraction for Generalizable Long-Horizon Robot Planning View paper
- [43] Symbolic planning and control of robot motion View paper
- [44] Task and Motion Planning for Humanoid Loco-Manipulation View paper
- [45] Sketch-Plan-Generalize: Learning and Planning with Neuro-Symbolic Programmatic Representations for Inductive Spatial Concepts View paper
- [46] PrimHOI: Compositional Human-Object Interaction via Reusable Primitives View paper
- [47] Integrating Cognitive, Symbolic, and Neural Approaches to Story Generation: A Review on the METATRON Framework View paper
- [48] Keep the planner in the loop: parallel planning and execution using Large Language Models View paper
- [49] GOALNET: Interleaving Neural Goal Predicate Inference with Classical Planning for Generalization in Robot Instruction Following View paper
- [50] Knowledge-based proof planning View paper
- [51] Transformer-Assisted Genetic Programming for Symbolic Regression View paper
- [52] Latent symbol lattices in probabilistic semiosis: An unconventional architectural mechanism for contextual modulation in large language models View paper
- [53] Symbolic visual reinforcement learning: A scalable framework with object-level abstraction and differentiable expression search View paper
- [54] Investigating a Novel Transposon Attention Scaffold for Large Scale Transformer Reasoning Patterns View paper
- [55] Terminating Differentiable Tree Experts View paper
- [56] Attention (as Discrete-Time Markov) Chains View paper
- [57] AI Reasoning in Deep Learning Era: From Symbolic AI to Neuralâ Symbolic AI View paper
- [58] Scallop: From probabilistic deductive databases to scalable differentiable reasoning View paper
- [59] UniSymNet: A Unified Symbolic Network Guided by Transformer View paper
- [60] Differentiable tree operations promote compositional generalization View paper
- [61] Turing Completeness of Bounded-Precision Recurrent Neural Networks View paper
- [62] On the expressive power of programming languages View paper
- [63] Symbolic Feedforward Networks for Probabilistic Finite Automata: Exact Simulation and Learnability View paper
- [64] Thinking Like Transformers View paper
- [65] How Do Transformers Learn Variable Binding in Symbolic Programs? View paper
- [66] Momentary Contexts: A Memory and Retrieval Approach for LLM Efficiency View paper
- [67] ProTo: Program-Guided Transformer for Program-Guided Tasks View paper
- [68] Deep Learning Models for Estimating Number of Lambda-Term Reduction Steps. View paper
- [69] Neural Abstract Interpretation View paper
- [70] Learned interpreters: structural and learned systematicity in neural networks for program execution View paper