
System-2 Reasoning via Generality and Adaptation

Sejin Kim **Sundong Kim**
Gwangju Institute of Science and Technology
sejinkim@gist.ac.kr sundong@gist.ac.kr

Abstract

While significant progress has been made in task-specific applications, current models struggle with deep reasoning, generality, and adaptation—key components of System-2 reasoning that are crucial for achieving Artificial General Intelligence (AGI). Despite the promise of approaches such as program synthesis, language models, and transformers, these methods often fail to generalize beyond their training data and to adapt to novel tasks, limiting their ability to perform human-like reasoning. This paper explores the limitations of existing approaches in achieving advanced System-2 reasoning and highlights the importance of generality and adaptation for AGI. Moreover, we propose four key research directions to address these gaps: (1) learning human intentions from action sequences, (2) combining symbolic and neural models, (3) meta-learning for unfamiliar environments, and (4) reinforcement learning to reason multi-step. Through these directions, we aim to advance the ability to generalize and adapt, bringing computational models closer to the reasoning capabilities required for AGI.

1 Introduction

Nowadays, Artificial Intelligence (AI) has made significant strides in task-specific applications, mainly through the success of neural networks [15, 28, 41, 61]. Furthermore, AI models based on language models [1, 6] and reinforcement learning models [25, 51] are achieving human-level performance with surprising results on several tasks. However, achieving human-level deep and logical reasoning, referred to as System-2 reasoning [29], a concept central to human cognition and intelligence, remains an open challenge. System-2 reasoning is characterized by abstract thought, logical deduction, and the ability to adapt to novel and complex situations, critical components of Artificial General Intelligence (AGI).

There is growing recognition that for AI to meet the reasoning requirements of AGI, it must excel in generality and adaptation [24], which are two fundamental capabilities that enable models to handle unseen tasks and unpredictable environments. Generality helps AI apply learned knowledge to new contexts, while adaptation enables flexible responses to changing scenarios [14]. These capabilities are also crucial for System-2 reasoning because they are fundamental to human-like intelligence [33].

Despite recent advancements, current AI models are limited by their reliance on training data and task-specific optimization [57]. Neural networks such as Large Language Models (LLMs) [17, 23, 36], and Model-Based RL (MBRL) models [51], while highly capable in constrained settings, often fail to generalize effectively beyond their training distributions. This has led to a critical gap between existing AI models and the level of reasoning required for AGI. In particular, most models demonstrate System-1 reasoning, which refers to fast, pattern-based decision-making, while lacking the deliberative and adaptable qualities inherent in System-2 reasoning [5, 11].

To bridge this gap, it is essential to identify strategies that foster both generality and adaptation in AI models. The Abstraction and Reasoning Corpus (ARC) [14] has been proposed as a benchmark specifically designed to assess an AI model’s ability to generalize and reason abstractly in unfamiliar

scenarios, without relying on extensive training data. Previous approaches to solving ARC, such as program synthesis [2, 3, 4, 8, 9, 20, 26, 37, 38, 50, 65], LLMs [13, 22, 27, 40, 48, 53, 59, 62, 63, 66], and variation of transformers [7, 44, 46] have shown promise but continue to struggle with abstract reasoning and adaptability when applied to novel tasks. For AI to evolve toward AGI-level reasoning, we must advance beyond current methodologies and explore new directions that focus on enhancing System-2 reasoning.

This paper addresses the limitations of existing AI methods in achieving System-2 reasoning and proposes four key research directions: (1) learning human intentions from trajectories, (2) integrating symbolic and neural hybrid models, (3) employing meta-learning for adaptation, and (4) using reinforcement learning to enhance multi-step reasoning. In this paper, we explore each key research direction for improving AI’s System-2 reasoning capabilities and demonstrate how these abilities can be assessed using benchmarks such as the Abstraction and Reasoning Corpus (ARC) [14], which focuses on generality and adaptation. By exploring these directions, we aim to improve AI’s generality and adaptability, bringing AI models closer to the reasoning capabilities required for AGI.

2 Backgrounds

2.1 Two Modes of Reasoning: System-1 and System-2

The concepts of System-1 and System-2 reasoning, introduced in psychology, describe two distinct modes of thinking [29]. System-1 is fast and intuitive, relying on automatic decision-making based on pattern recognition. Humans use System-1 in everyday situations that require little cognitive effort, such as driving or performing simple arithmetic, where decisions are made quickly based on learned experiences. In AI, this can be emulated using neural networks that learn rapidly and, when combined with tree search, replicate fast decision-making [5]. However, System-1 models work well in familiar environments but struggle with new situations and complex tasks.

On the other hand, System-2 reasoning involves slower, more deliberate thought processes that require logical deduction and abstract thinking. In contrast to System-1, System-2 is needed when encountering novel tasks that demand reasoning and planning, rather than relying on pattern recognition. This type of reasoning is essential for solving new challenges and adapting to dynamic environments, as is often the case when humans develop strategies for unfamiliar tasks. Introducing System-2 reasoning into AI would address current AI models’ limitations in handling more abstract, logic-driven tasks and enhance their ability to reason more like humans [11].

System-1 and System-2 differ in their cognitive mechanisms. System-1 is fast, intuitive, and pattern-based, while System-2 is slow, deliberate, and relies on logical reasoning. System-1 operates efficiently in familiar situations but lacks the flexibility to adapt to new environments, whereas System-2 is crucial for handling more complex, unfamiliar tasks. In AGI research, the integration of both systems is essential for achieving human-like reasoning capabilities, enabling AI to not only replicate learned behaviors but also think abstractly and adapt to novel challenges.

2.2 Generality and Adaptation: Two Key Components of AGI Reasoning

While System-2 reasoning enables deep, logical thought processes for handling complex and novel tasks, two critical components are required to fully realize its potential in AI models: generality and adaptation. For AI to achieve human-like intelligence, it must not only excel in task-specific skills but also demonstrate the ability to generalize and adapt across a wide range of environments and tasks. These two capabilities, generality and adaptation, are fundamental components of Artificial General Intelligence (AGI). Generality enables AI models to apply learned knowledge to novel contexts, allowing them to operate in unfamiliar situations without extensive retraining. Adaptation refers to the ability of AI to modify its behavior in response to changing environments, ensuring that it can handle new tasks and challenges as they arise.

Generality in AI is crucial because it allows models to move beyond narrow tasks and extend their capabilities to a broader range of problems. Rather than relying solely on patterns learned from training data, generality enables the system to extract underlying principles and apply them across different domains. For example, a model trained to recognize objects in images can generalize its knowledge to understand new categories with minimal supervision. The generality of an AI model $G(M, \mathcal{T}, \mathcal{K})$ can be mathematically defined as:

$$G(\mathcal{M}, \mathcal{T}, \mathcal{K}) = \frac{1}{N} \sum_{i=1}^N P(T_i | \mathcal{K}) \quad (1)$$

In Eq. 1, $G(\mathcal{M}, \mathcal{T}, \mathcal{K})$ represents the ability of a model \mathcal{M} to generalize across a set of tasks $\mathcal{T} = \{T_1, T_2, \dots, T_N\}$, given the domain knowledge \mathcal{K} . $P(T_i | \mathcal{K})$ refers to the probability that the model can successfully complete a novel task T_i based on the knowledge it has acquired. This formulation highlights how well the model can apply its learned knowledge to new tasks and situations, a key factor in achieving AGI-level reasoning. Achieving this level of generality requires AI to reason abstractly and draw connections between seemingly unrelated tasks, much like how humans operate.

Adaptation, on the other hand, ensures that AI models can adjust to new conditions or tasks without requiring extensive re-engineering or retraining. It reflects the model’s ability to maintain performance despite changes in the environment. For instance, a robotic system deployed in different physical spaces must adapt its behavior to account for various layouts, obstacles, and interactions. Without adaptation, even the most sophisticated AI models would struggle to maintain relevance in constantly evolving real-world environments, thus limiting their long-term effectiveness in AGI applications. The adaptation of an AI model $A(\mathcal{M}, \mathcal{T}, \mathcal{K}, \mathcal{E})$ can be described as:

$$A(\mathcal{M}, \mathcal{T}, \mathcal{K}, \mathcal{E}) = \frac{1}{N} \sum_{i=1}^N P(T_i | \mathcal{K}, E_i) \quad (2)$$

In Eq. 2, $A(\mathcal{M}, \mathcal{T}, \mathcal{K}, \mathcal{E})$ defines the model’s ability to adapt across tasks \mathcal{T} in varying environmental conditions $\mathcal{E} = \{E_1, E_2, \dots, E_N\}$, given the domain knowledge \mathcal{K} . The term $P(T_i | \mathcal{K}, E_i)$ represents the likelihood that the model will successfully complete the task T_i given both global domain knowledge \mathcal{K} and the specific environmental condition E_i . This formulation emphasizes how the ability to adapt depends not only on the model’s knowledge but also on its flexibility to respond to changes in the environment.

The synergy between generality and adaptation is what ultimately defines AGI’s reasoning capabilities. A truly intelligent system must generalize from previous experiences while adapting to new challenges, thus combining both qualities in its reasoning process. Research into enhancing generality and adaptation is therefore essential for developing AI models that can approach the reasoning capabilities of humans and handle tasks that go far beyond simple pattern recognition.

2.3 ARC: A Benchmark for Measuring Generality and Adaptation

The Abstraction and Reasoning Corpus (ARC) is a benchmark to evaluate an AI model’s ability to reason and generalize in ways that mirror human cognitive processes [14]. ARC tasks involve pairs of input and output grids, to predict the correct output for a new, unseen input based solely on a few example pairs. What makes ARC particularly challenging is that it requires systems to infer underlying abstract rules from minimal data, without the support of large datasets or domain-specific information. Each task tests the model’s capacity for abstraction, pattern recognition, and flexible reasoning, and the tasks are structured in such a way that they cannot be solved by brute force or by memorizing patterns.

ARC is specifically designed to assess two core components of AGI, generality and adaptation. For an AI model to succeed in ARC tasks, it must generalize from limited input-output examples and adapt its strategies to solve new, unseen tasks. Thus, ARC challenges AI models to demonstrate both the ability to apply learned knowledge in novel contexts (generality) and the capacity to adjust to varying conditions or rules (adaptation). ARC is not only a benchmark for generality and adaptation but also serves as a comprehensive test for System-2 reasoning in AI. The tasks in ARC require the type of generality and adaptation that are central to System-2 reasoning, making it an ideal tool for evaluating the research directions proposed in this paper. By succeeding in ARC tasks, AI models can demonstrate progress towards achieving System-2 reasoning, which is a crucial step towards AGI.

3 Analysis of Existing Approaches

In addressing ARC tasks, several approaches in AI have been explored, including program synthesis, LLMs, and transformers. In the view of System-2 reasoning, while these approaches have advanced generalization, they still face difficulties when adapting to novel and unfamiliar tasks. By reviewing these methods, we aim to highlight their successes and identify areas where further research is needed to improve both generality and adaptation.

3.1 Program Synthesis Is Bounded in Generality and Adaptation by Initial DSLs

Program synthesis allows AI models to generate programs to solve tasks based on given specifications or data. In the context of ARC, program synthesis facilitates abstract reasoning and interpretability. However, ARC tasks present significant challenges in terms of generality and adaptation, which are both critical for solving novel and dynamic problems. Although research has made significant strides in improving program synthesis through neural methods, object-centric reasoning, and symbolic or logic-based synthesis, these approaches are fundamentally limited by their reliance on an initial domain-specific language (DSL), which constrains both generality and adaptation.

Neural program synthesis has made strides in improving program generation by using neural networks to guide the process [2, 3, 4, 8, 9, 20, 26, 37, 38, 50, 65]. For example, one approach focused on creating interpretable programs [2], while another used a bidirectional search for efficiency [4]. Similarly, object-centric program synthesis leverages object relationships to guide program generation, focusing on enhancing generality by manipulating object characteristics across tasks [20, 37]. On the other hand, symbolic and inductive logic program synthesis applies logical reasoning to decompose ARC tasks into simpler components, improving generality through rule-based reasoning [9, 50]. However, across all these methods, the reliance on predefined DSL limits both generality and adaptation, particularly when faced with tasks that fall outside the scope of the initial DSL. Despite their successes, these approaches face adaptation challenges when problems require flexibility beyond the predefined rules and primitives of the DSL.

These studies highlight a fundamental challenge in program synthesis, where the initial DSL provided for generating programs influences both generality and adaptation. As shown in Eq. 1 and Eq. 2, the domain knowledge (\mathcal{K}) in the form of DSLs directly impacts the model’s ability to generalize to tasks and adapt to unfamiliar environments. Thus, both generality and adaptation of program synthesis are inherently tied to the expressiveness of the DSL, and future research must explore ways to create more dynamic and adaptable DSLs that allow AI to operate across a broader range of domains.

3.2 LLMs Demonstrate Strong Generality but Face Adaptation Constraints

Large Language Models (LLMs) have been successfully applied to a variety of tasks, including abstract reasoning tasks within the Abstraction and Reasoning Corpus (ARC). These models are particularly effective at recognizing patterns and generating solutions based on large-scale pre-training. In the context of ARC, LLMs have been employed to solve tasks by leveraging inductive reasoning, symbolic conversions, and hypothesis refinement. Despite their ability to generalize from large datasets, LLMs still encounter challenges when adapting to novel or unseen problems that require deeper, more abstract reasoning.

Several recent studies have explored different ways to apply LLMs to ARC tasks [13, 22, 27, 40, 48, 53, 59, 62, 63, 66]. For example, one study used LLMs for hypothesis refinement, allowing models to adjust their generated hypotheses based on feedback [48]. Another study focused on converting symbolic problems into natural language explanations, where LLMs could generalize across symbolic reasoning tasks, but faced limitations when dealing with more abstract challenges [53]. In addition, some models have been applied to pattern recognition tasks, utilizing LLMs to identify and generate patterns, though these models were constrained by their reliance on pre-training [40].

While LLMs show significant strengths in generality when solving ARC tasks, this generality ultimately stems from the vast amounts of data they are trained on (\mathcal{K}), making it inevitable that they exhibit weaknesses with untrained data. This inherent limitation of LLMs raises doubts that they are not truly reasoning but rather performing a form of interpolation based on the data they have been trained on [36].

OpenAI’s o1 model attempts to address these limitations through a design that incorporates Chain of Thought (CoT) reasoning for more deliberate problem-solving [43]. This approach could potentially address some challenges faced by traditional LLMs in ARC tasks by allowing the model to refine its thinking process and try different strategies [39]. However, while o1 shows promise, its effectiveness in ARC tasks and potential to achieve AGI remain uncertain [32]. Future research will be crucial in evaluating how such models perform on abstract reasoning tasks and whether they can bridge the gap between pattern recognition and genuine reasoning.

3.3 Transformers Also Enhance Generality but Struggle with Adaptation

Transformers have demonstrated remarkable capabilities in fields such as natural language processing and image recognition. However, when applied to tasks in ARC, transformers often face challenges in terms of generality and adaptation—critical components for solving novel and dynamic problems. Several studies have aimed to adapt transformer models for ARC tasks, leveraging their powerful pattern recognition and attention mechanisms to address complex reasoning tasks. Nevertheless, these approaches remain constrained by their ability to generalize beyond training data and adapt to unfamiliar scenarios.

Various transformer-based models have been proposed to handle ARC tasks more effectively. One approach introduced symmetry priors into the attention mechanisms, allowing the model to handle geometric transformations and generalize across tasks sharing geometric rules [7]. Another model focused on object-centric reasoning, enhancing its generality by leveraging object detection and relationships to solve ARC tasks involving object manipulation [46]. A third study aimed to improve the model’s counting abilities, a crucial aspect of quantitative reasoning in ARC, by refining its recognition of numerical patterns [44]. While these models have demonstrated notable improvements in handling specific ARC tasks, they are still limited by their reliance on predefined structures and priors, which restrict their adaptation to novel or abstract tasks.

Transformer-based approaches in ARC have shown strengths in enhancing generality, particularly in tasks involving geometric transformations, object relationships, or quantitative reasoning. This aligns with the prior knowledge [14], which defines these concepts as foundational components of intelligence. However, these models continue to struggle with adaptation, as they are often constrained by predefined structures like symmetry priors, object detection, or counting mechanisms. Future research must explore ways to make transformers more adaptable to unseen tasks, moving beyond predefined patterns and rules to better handle the dynamic and abstract nature of ARC.

4 Proposed Research Directions Toward System-2 Reasoning

In the previous section, we explored how generality and adaptation are critical components of System-2 reasoning. While existing AI models have shown mixed results in terms of generality, they consistently struggle with adaptation, which limits their ability to handle novel and complex tasks. To address these limitations, this section proposes four research directions aimed at advancing System-2 reasoning. AI must be able to extract abstract knowledge (\mathcal{K}), such as human intentions, from data and combine this knowledge in a generalized way, as seen in neuro-symbolic approaches. Additionally, meta-learning plays a crucial role in enabling AI to adapt efficiently across diverse environments and tasks, building on the foundation of learned models. Finally, reinforcement learning serves as a framework that integrates these research directions, providing the necessary structure to improve both generality and adaptation.

4.1 Learning Human Intentions from Data Supports Logical Deduction

To achieve System-2 reasoning, it is essential to establish a foundation for logical deduction and abstract thinking by capturing human intentions. Humans inherently perform System-2 reasoning through logical and abstract processes, and human intentions contain these crucial elements. Therefore, learning from human intentions, rather than merely imitating human actions, can provide a strong basis for advancing System-2 reasoning. Benchmarks like ARC [14], mini-ARC [30], 1D-ARC [66], ConceptARC [42], LARC [2], and MC-ARC [53] enable the collection of human behavior data to train AI models. Platforms like ARCreate Playground [31], O2ARC [52], ARC-Interactive [56], and

ARC-Game [12] are useful for collecting such data. Additionally, environments like ARCLE [34] provide settings where models can learn from the collected data to improve reasoning and adaptability.

Once data is gathered, methods for extracting human intentions must be applied to enhance AI’s logical deduction capabilities. Several techniques, including Topic Modeling [10], Sequential Pattern Mining (SPM) [55], and Hidden Markov Models (HMM) [49], have been developed to infer intentions from human behavior.

- **Topic Modeling** [10] clusters sub-sequences based on common patterns. It enables AI to infer goals by identifying recurring action sequences. Models must account for temporal dependencies to ensure that relationships between actions are preserved. Frequently co-occurring actions reveal intentions associated with specific patterns, aiding logical deduction and decision-making.
- **Sequential Pattern Mining (SPM)** [55] identifies frequently occurring action sequences and links state transitions to goals. SPM enhances logical deduction by anticipating actions based on observed patterns and tying actions to outcomes.
- **Hidden Markov Models (HMM)** [49] use probabilistic methods to model transitions between actions and hidden states. HMMs are effective for modeling temporal dependencies, inferring goals, and linking short-term actions to long-term objectives.

These approaches are particularly effective for learning intentions from human trajectories in ARC tasks, where understanding and adapting to human strategies are crucial for success. For instance, in a task where the objective is to align a grid diagonally, Topic Modeling could help AI identify repeated rotations and flips as strong indicators of diagonal alignment strategies used by humans. Sequential Pattern Mining can further enhance this understanding by detecting recurring action sequences, such as Rotate(CW) followed by Flip(Horizontal), that consistently lead to correct outcomes. By identifying these patterns, the AI can begin to predict successful strategies and apply them to other, similarly structured tasks. Hidden Markov Models take this reasoning process a step further by modeling the transitions between actions and hidden goals, allowing the AI to generalize learned strategies across a wider variety of novel tasks. Incorporating these methods enables AI models to develop a deeper understanding of human reasoning, enhancing their logical deduction and adaptability, which are fundamental components of achieving System-2 reasoning capabilities. Thus, the combination of these techniques with data collected through platforms like O2ARC [52] not only improves AI’s ability to perform abstract reasoning but also strengthens its potential for handling complex, multi-step tasks in dynamic environments.

4.2 Combining Symbolic and Neural Models Improves Abstraction and Generality

Achieving System-2 reasoning requires models to move beyond simple pattern recognition and incorporate logical, structured reasoning capabilities. One promising approach to achieving this is through hybrid models that combine neural networks with symbolic reasoning. Neural networks excel at learning patterns from data, while symbolic reasoning provides interpretable, rule-based frameworks for logical deduction. By integrating these two methodologies, it is possible to enhance the ability to reason abstractly and generalize to new, unseen tasks, which are essential for System-2 reasoning.

Several methods have been developed to combine symbolic and neural models, enhancing the generality and reasoning abilities.

- **Neuro-Symbolic Programming (NSP)** [45] integrates the pattern recognition strengths of neural networks with the explicit logic of symbolic programming. In this approach, the neural network learns patterns from data, while the symbolic component generates interpretable code that can be applied to reasoning tasks. This hybrid method enables handling of complex tasks requiring both learning from examples and reasoning over abstract concepts, such as program synthesis and abstract reasoning.
- **Differentiable Inductive Logic Programming (∂ ILP)** [19] extends inductive logic programming by integrating it with neural network optimization techniques. This method allows the system to learn logical rules while maintaining the flexibility of neural architectures. ∂ ILP improves the model’s ability to generalize by discovering dynamic rules that can be

applied across various domains, enabling AI to adapt its reasoning to new, unfamiliar tasks with minimal supervision.

- **Symbolic Regression with Neural Networks** [60] combines the power of neural networks with symbolic reasoning to uncover functional relationships between variables. The neural network identifies patterns in the data, and symbolic regression translates these patterns into symbolic expressions. This approach is particularly effective in tasks that require reasoning about mathematical or functional relationships, enhancing the model's generality across different problem spaces.

In ARC tasks, hybrid neuro-symbolic models are well-suited for generalizing from a small number of examples and applying symbolic rules to solve new tasks. For example, NSP can learn how humans apply transformations like rotations and flips to grids, and then generalize these transformations using symbolic logic. ∂ ILP could infer rules based on observed task patterns, while symbolic regression might identify underlying mathematical relationships between grid elements. By combining neural networks and symbolic reasoning, these hybrid models can reason more effectively about the structure of tasks, enhancing their abstract reasoning capabilities and enabling more generalization to various tasks, key aspects of System-2 reasoning.

4.3 Meta-Learning Facilitates Adaptation to Unfamiliar Environments

System-2 reasoning requires AI systems to adapt quickly to new and unfamiliar tasks, a process that can be supported by meta-learning. Unlike traditional learning methods that are focused on task-specific optimization, meta-learning trains models to "learn how to learn," making them more adaptable to various environments. By facilitating faster adaptation and improved generalization, meta-learning methods enable AI systems to respond effectively to novel challenges, which is crucial for achieving System-2 reasoning in tasks that require flexible thinking and logical deduction.

Several meta-learning techniques have been developed to enhance adaptability in AI systems, particularly in environments that demand reasoning about new and evolving tasks.

- **Recurrent Neural Processes (RNPs)** [47, 64] extend neural processes to handle sequential data, making them useful for tasks that require reasoning over time. In RNPs, models learn from time-dependent patterns, allowing them to predict future states based on past experiences. This capability is particularly relevant for multi-step reasoning tasks, enabling AI to adapt dynamically to evolving task requirements.
- **Prototypical Networks** [54] are designed for few-shot learning, where models classify new examples by comparing them to prototype representations. By learning these prototypes, AI systems can generalize across tasks with minimal data, making them efficient in environments that require structural reasoning. Prototypical networks are especially beneficial for System-2 reasoning, where recognizing and adapting to new patterns is key.
- **Meta-Reinforcement Learning (Meta-RL)** [21] combines meta-learning with reinforcement learning, creating a framework where agents not only adapt to new tasks but also optimize their reasoning over time. By incorporating a reward-based system, Meta-RL allows AI to balance reasoning about goals with adaptation in dynamic environments, strengthening its long-term planning and decision-making abilities.

In the context of ARC tasks, where multiple transformations may be required to achieve a goal, meta-learning can significantly enhance adaptability. For instance, in a task where the objective is to align objects diagonally, RNP could track the sequence of transformations and adjust reasoning based on temporal patterns. Prototypical networks could help the system recognize structural patterns involved in diagonal alignment, facilitating generalization across tasks. Meta-RL would provide the framework for balancing long-term planning with real-time reasoning, enabling the system to dynamically adjust its strategy based on feedback. By leveraging meta-learning, AI models can enhance their reasoning and adaptation capabilities, bringing them closer to the level of System-2 reasoning.

4.4 Reinforcement Learning Strengthens Adaptation and Multi-Step Logical Reasoning

System-2 reasoning requires not only logical deduction but also adaptability in dynamic environments and the ability to execute multi-step reasoning. Reinforcement Learning (RL) is particularly suited

for this, as it enables AI systems to learn through trial and error by optimizing actions based on rewards received from the environment. By continuously adjusting actions to maximize long-term rewards, RL enhances AI’s ability to adapt to unfamiliar tasks and perform multi-step reasoning, which are both critical for System-2 reasoning.

Several advanced RL techniques have been developed to strengthen adaptability and multi-step reasoning, providing AI systems with the capacity to plan, adapt, and reason across various contexts.

- **Hierarchical Reinforcement Learning (HRL)** [16] organizes decision-making into layers, where higher-level agents set overarching goals and lower-level agents execute actions to achieve those goals. This hierarchical structure allows for complex reasoning tasks to be broken down into manageable sub-tasks, improving both reasoning and adaptation in multi-step problems.
- **Model-Based Reinforcement Learning (MBRL)** [58] involves creating a model of the environment to predict the outcomes of different actions. By simulating various sequences of actions, MBRL enables AI to plan, reason about potential strategies, and adapt its decisions based on predicted outcomes. This approach is particularly effective for tasks that require long-term planning and reasoning about future states.
- **Relational Reinforcement Learning (RRL)** [18] focuses on learning relationships between objects and actions within a task. By understanding how different entities interact, RRL allows AI systems to reason about object relationships and adapt to new tasks that require spatial reasoning or object manipulation.

In ARC tasks, RL agents can be trained on ARCLE [34] and leveraged to learn optimal sequences of transformations. For instance, HRL can decompose a complex ARC task into sub-tasks, where each layer of the hierarchy handles a different level of abstraction, simplifying the reasoning process. MBRL can simulate the outcomes of different transformations, enabling AI to plan and choose the most effective sequence of actions [35]. RRL, by focusing on object relationships, allows AI to reason about how different transformations affect spatial arrangements, adapting to new tasks that involve object interactions. Through these methods, RL strengthens AI’s ability to perform multi-step logical reasoning and adapt to novel environments, bringing it closer to System-2 reasoning.

5 Conclusion

This paper explored the limitations of existing AI approaches, such as program synthesis, transformers, and large language models (LLMs), in solving ARC tasks, particularly in the context of generality and adaptation. While these methods have demonstrated notable improvements in generality, they consistently struggle with adaptation to novel and unfamiliar environments. Their reliance on predefined structures, patterns, or vast pre-training makes it challenging for current AI systems to achieve the deeper, more flexible reasoning associated with System-2 capabilities.

To address these challenges, we proposed four key research directions aimed at enhancing both generality and adaptation for System-2 reasoning: (1) learning human intentions from action sequences, (2) integrating symbolic and neural hybrid models, (3) employing meta-learning for adaptation, and (4) using reinforcement learning for multi-step logical reasoning. Each of these approaches targets specific limitations of current methods, helping models better handle abstract reasoning, generalize across various domains, and dynamically adapt to novel tasks.

By learning human intentions from action sequences, AI can better capture underlying goals, enabling stronger logical deduction and more human-like reasoning. Symbolic and neural hybrid models enhance abstract reasoning by bridging the gap between data-driven learning and rule-based logic, allowing AI to generalize beyond direct experience. Meta-learning facilitates faster adaptation to new environments by optimizing models to generalize their learning process itself, addressing the flexibility required for unfamiliar tasks. Finally, reinforcement learning strengthens both adaptation and multi-step logical reasoning by providing a framework for sequential decision-making and long-term strategy development.

Advancing towards System-2 reasoning will require these approaches to work in synergy, allowing models to both generalize across diverse tasks and adapt seamlessly to unforeseen scenarios. By pursuing these research directions, AI systems can evolve closer to AGI-level reasoning, demonstrating not only pattern recognition but true human-level reasoning, logical deduction, and abstract thinking.

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. GPT-4 Technical Report. *arXiv:2303.08774*, 2023.
- [2] Samuel Acquaviva, Yewen Pu, Marta Kryven, Theodoros Sechopoulos, Catherine Wong, Gabrielle Ecanow, Maxwell Nye, Michael Tessler, and Joshua B. Tenenbaum. Communicating Natural Programs to Humans and Machines. In *NeurIPS*, 2022.
- [3] James Ainooson, Deepayan Sanyal, Joel P. Michelson, Yuan Yang, and Maithilee Kunda. A Neurodiversity-Inspired Solver for the Abstraction & Reasoning Corpus (ARC) Using Visual Imagery and Program Synthesis. *arXiv:2302.09425*, 2023.
- [4] Simon Alford, Anshula Gandhi, Akshay Rangamani, Andrzej Banburski, Tony Wang, Sylee Dandekar, John Chin, Tomaso Poggio, and Peter Chin. Neural-Guided, Bidirectional Program Search for Abstraction and Reasoning. In *COMPLEX NETWORKS*, 2021.
- [5] Thomas Anthony, Zheng Tian, and David Barber. Thinking Fast and Slow with Deep Learning and Tree Search. In *NeurIPS*, 2017.
- [6] AI Anthropic. The Claude 3 Model Family: Opus, Sonnet, Haiku. *Claude-3 Model Card*, 1, 2024.
- [7] Matti Atzeni, Mrinmaya Sachan, and Andreas Loukas. Infusing Lattice Symmetry Priors in Attention Mechanisms for Sample-Efficient Abstract Geometric Reasoning. In *ICML*, 2023.
- [8] Andrzej Banburski, Anshula Gandhi, Simon Alford, Sylee Dandekar, Sang Chin, and Tomaso Poggio. Dreaming with ARC. In *NeurIPS Workshop on Learning Meets Combinatorial Algorithms*, 2020.
- [9] Shraddha Barke, Emmanuel Anaya Gonzalez, Saketh Ram Kasibatla, Taylor Berg-Kirkpatrick, and Nadia Polikarpova. HYSYNTH: Context-Free LLM Approximation for Guiding Program Synthesis. *arXiv:2405.15880*, 2024.
- [10] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.
- [11] Grady Booch, Francesco Fabiano, Lior Horesh, Kiran Kate, Jonathan Lenchner, Nick Linck, Andreas Loreggia, Keerthiram Murgesan, Nicholas Mattei, Francesca Rossi, et al. Thinking Fast and Slow in AI. In *AAAI*, 2021.
- [12] Alexey Borsky. ARC-Game, 2024.
- [13] Natasha Butt, Blazej Manczak, Auke Wiggers, Corrado Rainone, David Zhang, Michaël Defferrard, and Taco Cohen. CodeIt: Self-Improving Language Models with Prioritized Hindsight Replay. In *ICML*, 2024.
- [14] François Chollet. On the Measure of Intelligence. *arXiv:1911.01547*, 2019.
- [15] David B D’Ambrosio, Saminda Abeyruwan, Laura Graesser, Atil Iscen, Heni Ben Amor, Alex Bewley, Barney J Reed, Krista Reymann, Leila Takayama, Yuval Tassa, et al. Achieving Human Level Competitive Robot Table Tennis. *arXiv:2408.03906*, 2024.
- [16] Peter Dayan and Geoffrey E Hinton. Feudal reinforcement learning. In *NeurIPS*, 1992.
- [17] Gonçalo Hora de Carvalho, Robert Pollice, and Oscar Knap. Show, Don’t Tell: Evaluating Large Language Models Beyond Textual Understanding with ChildPlay. *arXiv:2407.11068*, 2024.
- [18] Sašo Džeroski, Luc De Raedt, and Kurt Driessens. Relational Reinforcement Learning. *Machine Learning*, 43:7–52, 2001.
- [19] Richard Evans and Edward Grefenstette. Learning Explanatory Rules from Noisy Data. *Journal of Artificial Intelligence Research*, 61:1–64, 2018.

- [20] Sébastien Ferré. Tackling the Abstraction and Reasoning Corpus (ARC) with Object-centric Models and the MDL Principle. In *IDA Symposium*, 2024.
- [21] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *ICML*, 2017.
- [22] Liane Galanti and Ethan Baron. Intelligence Analysis of Language Models. *arXiv:2407.18968*, 2024.
- [23] Gael Gendron, Qiming Bao, Michael Witbrock, and Gillian Dobbie. Large Language Models Are Not Strong Abstract Reasoners. In *IJCAI*, 2024.
- [24] Ben Goertzel. Artificial General Intelligence: Concept, State of the Art, and Future Prospects. *Journal of Artificial General Intelligence*, 5(1):1, 2014.
- [25] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering Diverse Domains through World Models. *arXiv:2301.04104*, 2023.
- [26] Céline Hocquette and Andrew Cropper. Relational decomposition for program synthesis. *arXiv:2408.12212*, 2024.
- [27] Shengran Hu, Cong Lu, and Jeff Clune. Automated Design of Agentic Systems. *arXiv:2408.08435*, 2024.
- [28] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly Accurate Protein Structure Prediction with AlphaFold. *nature*, 596(7873):583–589, 2021.
- [29] Daniel Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.
- [30] Subin Kim, Prin Phunyaphibarn, Donghyun Ahn, and Sundong Kim. Playgrounds for Abstraction and Reasoning. In *NeurIPS Workshop on Neuro Causal and Symbolic AI*, 2022.
- [31] Lab42. ARC Playground, 2021.
- [32] Lab42. OpenAI o1 Results on ARC-AGI-Pub, 2024.
- [33] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building Machines That Learn and Think Like People. *Behavioral and Brain Sciences*, 40:e253, 2017.
- [34] Hosung Lee, Sejin Kim, Seungpil Lee, Sanha Hwang, Jihwan Lee, Byung-Jun Lee, and Sundong Kim. ARCLE: The Abstraction and Reasoning Corpus Learning Environment for Reinforcement Learning. In *CoLLAs*, 2024.
- [35] Jihwan Lee, Woochang Sim, Sejin Kim, and Sundong Kim. Enhancing Analogical Reasoning in the Abstraction and Reasoning Corpus via Model-Based RL. In *IJCAI Workshop on Interactions between Analogical Reasoning and Machine Learning*, 2024.
- [36] Seungpil Lee, Woochang Sim, Donghyeon Shin, Sanha Hwang, Wongyu Seo, Jiwon Park, Seokki Lee, Sejin Kim, and Sundong Kim. Reasoning Abilities of Large Language Models: In-Depth Analysis on the Abstraction and Reasoning Corpus. *arXiv:2403.11793*, 2024.
- [37] Chao Lei, Nir Lipovetzky, and Krista A Ehinger. Generalized Planning for the Abstraction and Reasoning Corpus. In *AAAI*, 2024.
- [38] Mintaek Lim, Seokki Lee, Liyew Woletemaryam Abitew, and Sundong Kim. Abductive Symbolic Solver on Abstraction and Reasoning Corpus. In *IJCAI Workshop on Logical Foundation of Neuro-Symbolic AI*, 2024.
- [39] Raffaele Marino. Fast Analysis of the OpenAI O1-Preview Model in Solving Random K-SAT Problem: Does the LLM Solve the Problem Itself or Call an External SAT Solver? *arXiv preprint arXiv:2409.11232*, 2024.

- [40] Suvir Mirchandani, Fei Xia, Pete Florence, Brian Ichter, Danny Driess, Montserrat Gonzalez Arenas, Kanishka Rao, Dorsa Sadigh, and Andy Zeng. Large Language Models as General Pattern Machines. In *CoRL*, 2023.
- [41] Matej Moravčík, Martin Schmid, Neil Burch, Viliam Lisý, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson, and Michael Bowling. DeepStack: Expert-Level Artificial Intelligence in Heads-Up No-Limit Poker. *Science*, 356(6337):508–513, 2017.
- [42] Arseny Moskvihev, Victor Vikram Odouard, and Melanie Mitchell. The ConceptARC Benchmark: Evaluating Understanding and Generalization in the ARC Domain. *Transactions on Machine Learning Research*, 2023.
- [43] OpenAI. Learning to Reason with LLMs, 2024.
- [44] Simon Ouellette, Rolf Pfister, and Hansueli Jud. Counting and Algorithmic Generalization with Transformers. *arXiv:2310.08661*, 2023.
- [45] Emilio Parisotto, Abdel-rahman Mohamed, Rishabh Singh, Lihong Li, Dengyong Zhou, and Pushmeet Kohli. Neuro-Symbolic Program Synthesis. In *ICLR*, 2017.
- [46] Jaehyun Park, Jaegyun Im, Sanha Hwang, Mintaek Lim, Sabina Ualibekova, Sejin Kim, and Sundong Kim. Unraveling the ARC Puzzle: Mimicking Human Solutions with Object-Centric Decision Transformer. In *ICML Workshop on Interactive Learning with Implicit Human Feedback*, 2023.
- [47] Shenghao Qin, Jiacheng Zhu, Jimmy Qin, Wenshuo Wang, and Ding Zhao. Recurrent Attentive Neural Process for Sequential Data. *arXiv:1910.09323*, 2019.
- [48] Linlu Qiu, Liwei Jiang, Ximing Lu, Melanie Sclar, Valentina Pyatkin, Chandra Bhagavatula, Bailin Wang, Yoon Kim, Yejin Choi, Nouha Dziri, et al. Phenomenal Yet Puzzling: Testing Inductive Reasoning Capabilities of Language Models with Hypothesis Refinement. In *ICLR*, 2024.
- [49] Lawrence Rabiner and Biinghwang Juang. An Introduction to Hidden Markov Models. *IEEE ASSP Magazine*, 3(1):4–16, 1986.
- [50] Filipe Marinho Rocha, Inês Dutra, and Vítor Santos Costa. Program Synthesis using Inductive Logic Programming for the Abstraction and Reasoning Corpus. *arXiv:2405.06399*, 2024.
- [51] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering Atari, Go, Chess and Shogi by Planning with A Learned Model. *Nature*, 588(7839):604–609, 2020.
- [52] Suyeon Shim, Dohyun Ko, Hosung Lee, Seokki Lee, Doyoon Song, Sanha Hwang, Sejin Kim, and Sundong Kim. O2ARC 3.0: A Platform for Solving and Creating ARC Tasks. In *IJCAI Demo*, 2024.
- [53] Donghyeon Shin, Seungpil Lee, Klea Lena Kovacec, and Sundong Kim. From Generation to Selection Findings of Converting Analogical Problem-Solving into Multiple-Choice Questions. In *EMNLP*, 2024.
- [54] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical Networks for Few-Shot Learning. In *NeurIPS*, 2017.
- [55] Ramakrishnan Srikant and Rakesh Agrawal. Mining Sequential Patterns: Generalizations and Performance Improvements. In *EDBT*, 1996.
- [56] Simon Strandgaard. ARC-Interactive, 2024.
- [57] Richard Sutton. The Bitter Lesson. *Incomplete Ideas (blog)*, 13(1):38, 2019.
- [58] Richard S Sutton. Integrated Architectures for Learning, Planning, and Reacting Based on Approximating Dynamic Programming. In *ICML*, 1990.

- [59] John Chong Min Tan and Mehul Motani. LLMs as a System of Multiple Expert Agents: An Approach to Solve the Abstraction and Reasoning Corpus (ARC) Challenge. In *IEEE CAI*, 2024.
- [60] Silviu-Marian Udrescu and Max Tegmark. AI Feynman: A Physics-Inspired Method for Symbolic Regression. *Science Advances*, 6(16), 2020.
- [61] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster Level in Starcraft II Using Multi-Agent Reinforcement Learning. *nature*, 575(7782):350–354, 2019.
- [62] Ruo Cheng Wang, Eric Zelikman, Gabriel Poesia, Yewen Pu, Nick Haber, and Noah D Goodman. Hypothesis Search: Inductive Reasoning with Language Models. In *ICLR*, 2024.
- [63] Yile Wang, Sijie Cheng, Zixin Sun, Peng Li, and Yang Liu. Speak It Out: Solving Symbol-Related Problems with Symbol-to-Language Conversion for Language Models. In *ICLR Workshop on AGI*, 2024.
- [64] Timon Willi, Jonathan Masci, Jürgen Schmidhuber, and Christian Osendorfer. Recurrent Neural Processes. *arXiv:1906.05915*, 2019.
- [65] Yudong Xu, Elias B. Khalil, and Scott Sanner. Graphs, Constraints, and Search for the Abstraction and Reasoning Corpus. In *AAAI*, 2023.
- [66] Yudong Xu, Wenhao Li, Pashootan Vaezipoor, Scott Sanner, and Elias B Khalil. LLMs and the Abstraction and Reasoning Corpus: Successes, Failures, and the Importance of Object-Based Representations. *Transactions on Machine Learning Research*, 2024.

NeurIPS Paper Checklist

1. **Claims**

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Abstract, Section 1

2. **Limitations**

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Section 5

3. **Theory Assumptions and Proofs**

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Section 2

4. **Experimental Result Reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification: Since this paper is a positional paper, it does not include any experiments.

5. **Open access to data and code**

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: Since this paper is a positional paper, it does not include any data or code.

6. **Experimental Setting/Details**

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: Since this paper is a positional paper, it does not include any experiments.

7. **Experiment Statistical Significance**

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: Since this paper is a positional paper, it does not include any experiments.

8. **Experiments Compute Resources**

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: Since this paper is a positional paper, it does not include any experiments.

9. **Code Of Ethics**

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification:

10. **Broader Impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: Since this paper is a positional paper, there is no societal impact.

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Since this paper is a positional paper, it does not use any data or models.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: This paper cites every previous work accordingly.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: This paper includes the mathematical definitions of generality and adaptation.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not involve crowdsourcing or human subjects.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This paper does not involve crowdsourcing or human subjects.