

# A General Neural-symbolic Architecture for Knowledge-intensive Complex Reasoning

Shulin Cao <sup>a,\*</sup>, Zijun Yao <sup>a</sup>, Lei Hou <sup>a</sup> and Juanzi Li <sup>a,\*\*</sup>

<sup>a</sup> *Department of Computer Science and Technology, Tsinghua University, Beijing, China*

*E-mails: caosl19@mails.tsinghua.edu.cn, yaozj20@mails.tsinghua.edu.cn, houlei@tsinghua.edu.cn, lijuanzi@tsinghua.edu.cn*

**Abstract.** The symbolic and connectionist AI are two main routes in the last two generations of AI. Nowadays, researchers have gradually realized that these two routes do not conflict in principle, but can and should integrate and help each other. In this paper, we take knowledge-intensive complex reasoning task as an example to explore the integration of symbolic and connectionist AI. In particular, we design a general neural-symbolic architecture for this task, which is composed of four main components: multiple knowledge sources, knowledge manipulator, reasoning planner, and reasoning conductor. We also introduce existing techniques that can be used to implement each component, and summarize their advantages and potential directions for enhancement in building a neural-symbolic system for knowledge-intensive complex reasoning tasks.

**Keywords:** Knowledge-intensive Complex Reasoning, Neural-symbolic Framework, Large Language Models, Knowledge Bases

## 1. Introduction

Artificial intelligence (AI) aims at creating machine intelligence that can perceive, learn, reason, learn, and interact with their environment in a way that simulates human cognitive abilities. There are two main paradigms adopted in AI in order to achieve this: (1) the symbolic, and (2) the connectionist AI.

The first generation of AI is dominated by the symbolic approach, which believes that a physical symbol system has the necessary and sufficient means for general intelligent action [1]. It inherits the idea of mathematical logic that the physical world could be represented by symbolic systems, and views intelligence as a process of symbolic manipulation and reasoning following logical rules. Its contemporary successor is knowledge engineering, which takes knowledge accumulation, reuse, and knowledge-intensive complex reasoning as its core issues. The second generation of AI is dominated by the connectionist approach. The ideological basis of connectionist AI is that the human brain is physically based on a network of interconnected neurons, and it is possible to create AI by computerizing neural networks that mimic those of living beings. Although the connectionist AI encountered many setbacks in the early stage, in recent years, on the basis of continuously improved deep learning models such as convolutional neural networks (CNN) [2], recurrent neural networks (RNN) [3] and Transformers [4], it has reached a new high level. Recently, large language models (LLMs) such as ChatGPT (GPT 3.5) and GPT-4 [5] have demonstrated astonishing abilities on a variety of domains and tasks, including abstraction, comprehension, vision, coding, mathematics, medicine, law, understanding of human motives and emotions, and more [6]. It is widely recognized by academics and the industry that LLMs marks a new level of intelligence in AI systems, and can be regarded as the

---

\* Shulin Cao and Zijun Yao contribute equally to this paper.

\*\*Corresponding author. E-mail: lijuanzi@tsinghua.edu.cn.

prototype of artificial general intelligence (AGI). LLMs are considered as a landmark success in the connectionist line of AI research.

Traditionally, the symbolic and connectionist approaches have been viewed as competing and mutually reinforcing. In recent years, there has been a widespread rethinking of the relationship between these two classical approaches. Drawing on the analogy from Daniel Kahneman’s “Thinking, Fast and Slow” [7] where System 1 is likened to fast, intuitive perception (akin to neural processes) and System 2 to slow, deliberate cognition (similar to symbolic reasoning), the integration of neural and symbolic approaches, akin to the harmonious cooperation between System 1 and System 2, holds significant promise for advancing AI. Because on one hand, neural approaches, paralleling System 1, are adept at processing vast amounts of sensory data and identifying patterns within this data. However, neural methods have notable limitations. They often operate as “black boxes”, making it challenging to interpret how decisions are made or to trace the reasoning process. This lack of transparency can be problematic in applications where understanding the decision-making process is crucial, such as in healthcare or criminal justice. Additionally, neural networks require large amounts of data to learn effectively, and they can be prone to overfitting, where they perform well on training data but poorly on unseen data. They also struggle with tasks that require complex logical reasoning or handling of abstract concepts, where a clear, structured approach is necessary.

On the other hand, symbolic approaches, mirroring System 2, excel in tasks that demand explicit knowledge representation and rule-based reasoning. The clear, rule-based nature of symbolic AI allows for a level of transparency and explainability that neural methods often lack. However, symbolic systems can be rigid, relying heavily on predefined rules and knowledge bases, which makes them less adaptable to new or unforeseen situations. Furthermore, the development of symbolic systems can be labor-intensive, requiring extensive domain expertise to encode knowledge and rules accurately. Therefore, by combining the intuitive pattern recognition capabilities of neural methods with the structured reasoning abilities of symbolic systems, integrated neural-symbolic models can leverage the strengths of both paradigms. Such models are capable of not only processing sensory data efficiently but also applying logical reasoning and abstract thinking to make sense of the data in a more human-like manner [8]. Zhang Bo also points out that the third-generation AI method should be “double-wheel-driven by data and knowledge”, so as to realize explainable, robust, safe and reliable AI systems [9]. Nowadays, researchers have gradually realized that these two routes do not conflict in principle, but can and should integrate and help each other.




Fig. 1. Knowledge can be stored in different sources, such as structured knowledge bases, unstructured text, and LLMs’ parameters. Color blue denotes concepts, orange denotes entities, green denotes the attributes of entities, and red denotes relations among entities and concepts.

In this paper, we will explore the integration of symbolic and connectionist AI to address knowledge-intensive tasks [10], which are in the form of Question Answering (QA) and demand advanced knowledge-intensive complex reasoning capabilities that extend beyond mere pattern recognition to encompass a deep understanding of entities and their relationships within the world. Specifically, in knowledge-intensive complex reasoning, a complex question is given to the intelligent system, and a concise answer for the question is expected. It requires deep understanding of natural language text (including the question and knowledge corpora), and involves various reasoning capabilities such as multi-hop inference, attribute comparison, and set operation [11, 12]. Here, the knowledge can be stored in different sources, such as structured knowledge bases, unstructured text corpora, the parametric knowledge in LLMs’ parameters, images, etc., as shown in Figure 1. Actually, as John McCarthy states, the 1976 memorandum [13] is of modern interest. In this article, he introduces a story from New York Times, writes 22 questions for it, and points out that an intelligent person or program should be able to answer the questions, which requires deep

understanding and reasoning ability. He also gives guidelines for knowledge-intensive complex reasoning. There are two basic elements: 1) A formalism capable of expressing the assertions of the sentences free from dependence on the grammar of the English language, which can be called artificial natural language—ANL; 2) A data structure for expressing the facts (apart from expressing the sentences). Certain programs can be applied to the data structures.

Inspired by the proposal of John McCarthy, we design our general neural-symbolic architecture for knowledge-intensive complex reasoning, which will be introduced in Section 2. Specifically, the framework is composed of four main construction blocks—knowledge-oriented data manager, knowledge manipulator, reasoning planner, and reasoning conductor. These modules draw advantages from both neural and symbolic methodologies, and take the responsibility of knowledge-intensive complex reasoning in different levels. The overall architecture also integrates human in the loop, to progressively accomplish the targeted tasks and optimize the overall system. Next, we will present the detailed implementations of the four modules in Section 3, by introducing the related works. Finally, we will look into the future and conclude the paper in Section 4.

## 2. Overall Architecture Design

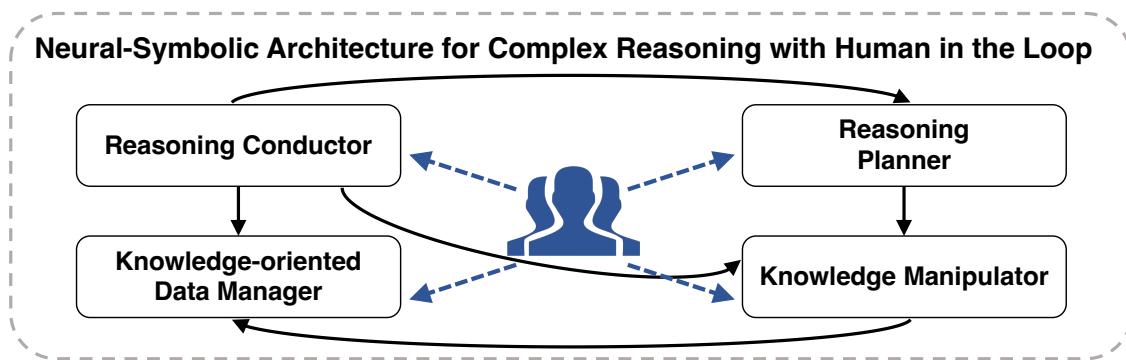


Fig. 2. The proposed neural-symbolic framework to solve knowledge-intensive complex reasoning task. We use solid arrows to indicate calling and access relationship between components and dashed arrows to indicate the interaction between human and each component. The basis of the whole architecture is the knowledge-oriented data manager. Knowledge manipulator performs knowledge operations on the knowledge data. Reasoning planner produces candidate reasoning path by decomposing complex reasoning tasks into assembled knowledge manipulation primitives. Reasoning conductor decides the optimal reasoning path, select knowledge manipulation implementations, and execute on the knowledge manager to obtain the answer. All four components can be implemented with neural methods and symbolic methods or a combination of the two. Our architecture also incorporates human in the loop to update knowledge sources, fix and optimize reasoning plans, and converse with the system to perform multiple-step reasoning tasks.

With the rapid development of large-scale pre-trained models in many different domains [5, 14], neural models show great promise in understanding complex semantics and generalizing to unseen data and tasks. They not only encapsulate tremendous world knowledge, as recognized as parametric knowledge [15], but also show excellent skills in understanding semantics, as demonstrated in solving tasks such as semantic parsing [16] and planning [17]. However, neural methods fail to reliably reason with pre-defined schemata and generate appeared to be correct but actually wrong answers, which is recognized as notorious hallucination issue. Thus, it is necessary to find a way to drive neuron methods to function without hallucination. Luckily, symbols are able to express explicit constraints and knowledge guidance. Although they are less superior than neural methods in generalization, they tend to be more reliable in solving strict knowledge-intensive complex reasoning tasks. Thus, combining the advantages from both the neural and the symbolic perspectives is a promising direction to realize intelligence.

Under such belief, previous researchers developed frameworks that integrate both neuron and symbolic approaches. However, within these existing frameworks, either the neural or the symbolic method predominates, often relegating the other to a secondary role. For example, in framework dominated by neural approaches, symbolic methods are merely utilized as regularization terms for the loss function [18]. Conversely, in frameworks where

symbolic methods take precedence, neural approaches are relegated to serving merely as embedding modules for knowledge representation [19]. In light of this, we express a desire for a framework that unifies neural-symbolic integration, thereby unleashing the potential of both methodologies simultaneously.

We propose a neuro-symbolic framework that mimics the dual-process of human mind [20]. Neural models, as represented by large language models (LLMs), convert rough natural language intentions into well-defined symbols [13, 21]. They resemble system 1 to perform associative thinking. The reasoning is performed with the aid of symbolic methods, which resembles system 2 to perform logical thinking.

We propose to implement such framework as the following architecture, an initiative neural-symbolic realization to solve knowledge-intensive complex reasoning tasks, which is demonstrated in Figure 2. The architecture is composed of four main construction blocks—knowledge-oriented data manager, knowledge manipulator, reasoning planner, and reasoning conductor. These components carry distinct responsibilities, drawing advantages from both neural and symbolic methodologies. Moreover, the proposed framework integrates human involvement in the loop, to progressively accomplish the targeted tasks and optimize the overall system. In the following of this section, we specify the distinct responsibilities of each module and discuss how human contribute to and enhance the overall framework in the loop. We will use question “*Who is taller, China’s most famous basketball player or the basketball player LeBron James?*” as a running example to show how our architecture solves the question.

### 2.1. Knowledge-oriented Data Manager

Knowledge-oriented data manager is the foundation of the overall system. It aggregates knowledge from multiple sources, including not only traditional knowledge bases [22, 23], but also information scattered in raw formats such as plain text [24, 25], model parameters [26], images [27], and tables [28].

To facilitate unified knowledge manipulation, the manager provide knowledge information in a coherent symbolic representation. A typical semantic representation system comprises five core knowledge elements: (1) *Entities* are unique object in the real world; (2) *Concepts* are sets of entities with similar properties; (3) *Relations* describe the connection among entities and concepts; (4) *Attributes* are associated with entities to describe them from certain aspects; (5) *Qualifiers*<sup>1</sup> are constraints on other knowledge items to specify under which conditions that they are true. To answer the example question, we need information about entity “Lebron James” and the most famous basketball player in China. Their heights are naturally attributes. It is worth noting that there are variant definitions for basic knowledge elements. For example, visual knowledge symbols typically defines *Objects*, which shares a similar definition with *Entities*. Thus, the data manager need to provide a compiler to translate different knowledge elements into unified symbols.

Knowledge-oriented data manager brings benefit in twofold. First, the knowledge-oriented data management enables precise knowledge manipulation, in contrary to scattering knowledge in raw data. For example, plain text data only allows for fuzzy similarity matching with query utterances, whereas the knowledge-oriented data manager provides a much more accurate method to access information by posing constraints on the knowledge such as retrieving entities when constraining their concept. This capability has the potential to reduce hallucination when generating answers conditioned on the given data. Second, the knowledge-oriented data manager gathers knowledge complementary with each other from different sources, which broadens the knowledge coverage.

### 2.2. Knowledge Manipulator

Knowledge manipulator provides well-defined knowledge manipulation primitives on the knowledge sources. These primitives have strictly defined behavior, which describes the intended outcome of each individual primitive, without stipulating specific their implementation algorithms and techniques. As a result, each knowledge manipulation primitive can be realized via neural models, symbolic methods, or a combination of both. They can be assembled into programs, the execution of which produce the final result for a complex reasoning task.

Specifically, knowledge manipulation primitives include: (1) *Direct knowledge access*, which retrieves elements from the data structure according to specific requirements, e.g., linking “*Lebron James*” to the entities they refer

---

<sup>1</sup>Also known as *Hyper Relations* or *CVT* in different literature [22, 29].

1 to, a process known as entity disambiguation; (2) *Knowledge processing*, which operates on the outputs of other 1  
2 knowledge operations, *e.g.*, find the most famous one in China from all basketball player entities; (3) *Knowledge* 2  
3 *updating*, which offers the capability to refresh and refine outdated information in the data structure. 3

4 The precise knowledge manipulation brings two key advantages. Firstly, these unified primitives act as an inter- 4  
5 face to incorporate various implementations with complementary capacities for the same objective. This allows for 5  
6 the selection of the optimal implementation to construct the final manipulating operation (how to choose optimal 6  
7 implementation is undertaken by the reasoning conductor in Section 2.4). Secondly, by atomizing basic operations, 7  
8 the overall system imparts explicit meanings to each procedure and preserves the corresponding intermediate results. 8  
9 Consequently, human users can inspect not only the final answer, but also the output of each individual operation. 9  
10 This transparency of execution thereby enhances interpretability. 10

### 11 2.3. Reasoning Planner 11

12 Reasoning planner decomposes knowledge-intensive complex reasoning tasks into assembled knowledge primi- 12  
13 tives. Acting as an interface, it comprehends human requirements and represents them as reasoning processes, *i.e.*, 13  
14 the composition of basic operations, and the execution result of reasoning process is the answer for the knowledge- 14  
15 intensive complex reasoning task. 15

16 Specifically, it takes the language understanding ability of neural models to decompose a complex ques- 16  
17 tion/instruction into basic symbolic units, *i.e.*, an ordered list of steps. In our framework, the format of steps is 17  
18 not constrained, and can be expressed through either natural language or symbolic operations. The reasoning 18  
19 process output by planner is therefore a multi-step program, a chain-of-thought reasoning process, or a sequence of 19  
20 simple questions. For example, for the question “Who is taller, China’s most famous basketball player or the bas- 20  
21 ketball player LeBron James?”, a possible ordered list of steps outputted by the planner is: find all basketball player, 21  
22 then filter by attribute nationality being China, and identify the most famous one. Finally, the identified entity is 22  
23 compared with another found entity, LeBron James, in terms of their heights. 23  
24 24

25 The planning module has multiple different responsibilities, among which the most important is to refine ambigu- 25  
26 ous human intention into well-defined knowledge manipulations. In this process, it allows humans to express their 26  
27 intention in a human-friendly way—natural language. Most recently, large scale pre-trained language models show 27  
28 potential to make plans according to natural language prompts, due to its remarkable language understanding and 28  
29 generation abilities, which sheds light on the feasibility to achieve this goal. Secondly, the planning module extracts 29  
30 higher order intentions from natural language descriptions. For example, answering the question “Did Aristotle Use 30  
31 a Laptop?” is an implicit reasoning task that can be planned with our reasoning planner. 31  
32 32

### 33 2.4. Reasoning Conductor 33

34 Reasoning conductor makes decisions based on the candidate reasoning processes given by the reasoning planner, 34  
35 and executes the knowledge manipulator on the multiple-sourced knowledge to obtain the final answer. Acting as 35  
36 a decision maker, it has the ability to schedule the planner and manipulator flexibly, trying to search an optimal 36  
37 solution for the whole task. 37

38 Specifically, for a natural language question/requirement that describes the complex reasoning task, there may 38  
39 exist multiple ways to plan the reasoning process, and each primitive in the reasoning process can incorporate various 39  
40 implementations. For example, if the user poses a complex question “Who is taller, China’s most famous basketball 40  
41 player or the basketball player LeBron James?”, there are multiple ways to locate the most famous basketball 41  
42 player in China: 1) From the perspective symbolic method, we can identify all the basketball player from China in the 42  
43 knowledge base, and use their node degree in the knowledge graph to measure how famous they are; 2) From the 43  
44 perspective of neuron method, we can retrieve the most famous basketball player from the memory of LLMs with 44  
45 appropriate prompt; 3) As a combination, we can also ask LLMs to read the information of all Chinese basketball 45  
46 players in the knowledge base and identify the most famous one. With so many options, an ideal reasoning conductor 46  
47 need to assess the pros and cons of different plans and choose the optimal one. 47  
48 48

49 The reasoning conductor brings two advantages for our framework. First, it can select an optimal reasoning 49  
50 process from our planning module, and decide the optimal implementation strategies for the knowledge manipulator 50  
51 51

by systematically considering multiple factors comprehensively, including the computational cost, the precision, and recall. Second, it allows users to customize their own priority for decision making, by injecting prior knowledge into the reasoning process, *e.g.*, choosing knowledge base as the primary knowledge sources.

### 2.5. The Overall Architecture with Human in the Loop

Human possess prior-knowledge and evolving posterior knowledge about the world. In the meanwhile, the architecture still suffers out of distribution tasks. To better serve human needs, the proposed architecture also allows human to improve its world understanding and performance with feedbacks. We highlight three most important life cycles that involve both human users and system modules.

The *knowledge updating cycle* allows human users to update the knowledge in storage. This process includes allowing user to inspect knowledge when the answer is less satisfactory and annotate newly involved knowledge. In the *program fixing cycle*, human users are allowed to find error programs given by the reasoning planner and help to select the optimal reasoning path if the solver selects the sub-optimal solution. Humans not only help the system solve current complex reasoning task instantly, but also accumulate training data, which is further used to train better task planner and problem solver. The *human-aided problem solving cycle* allows human user to decompose complex reasoning tasks into simpler ones on the fly in a dialogue. The system coordinate with the human user to solve simple reasoning tasks and assemble their output into the final result We credit knowledge manipulating, planning and conducting plans as fundamental skills of humans. As a such, the human-in-the-loop process does not need to specify the human to be an expert.

## 3. Related Research and Technique towards the Architecture

We introduce existing techniques that can be used to implement the proposed architecture in Section 2. We also summarize their advantages and potential directions for enhancement in building a neural-symbolic system to solve complex reasoning tasks.

### 3.1. Knowledge-oriented Data Manager

We introduce techniques that are used to manage and represent knowledge from different sources, including knowledge base, textual corpora, large language models, and other knowledge sources.

**Knowledge Base.** Knowledge base<sup>2</sup> stores crowd-sourced knowledge in the format  $(h, r, t)$  [22, 23, 30].  $h$  and  $t$  are head node and tail node, while  $r$  is the relation between the head node and the tail node. They are further classified as fact knowledge, hierarchy knowledge, attribute knowledge, and qualifier knowledge. Specifically, for fact knowledge,  $h$  and  $t$  are entities and  $r$  is a relation. For hierarchy knowledge,  $t$  is a concept, and  $h$  is either an entity that falls into the category of  $h$  or a subconcept of  $h$ , where  $r$  denotes `isInstanceOf` and `subClassOf`, respectively. For attribute knowledge,  $h$  is an entity,  $r$  is the attribute name, and  $t$  is the attribute value. Qualifier knowledge recursively specifies the condition under which other knowledge is true, where  $h$  is a triple of the constrained knowledge,  $r$  is an attribute name, and  $t$  is the attribute value.

**Textual Corpora.** Textual corpora are rich in unstructured knowledge. They can be obtained via search engines [31] and aligned with the unified knowledge schemata through a series of information extraction algorithms, such as named entity recognition, entity linking [32], relation extraction [33, 34], and entity typing [35].

In managing textual knowledge, markup languages plays a key role as well-defined symbolic system. GraphQ IR uses sentinel tokens to mark special knowledge elements, *e.g.*, `<C></C>` for concept. Although GraphQ IR [36] is originally proposed to formally represent natural language questions, it shows potential to record textual knowledge in textual data. Other techniques include resource description framework (RDF) and web ontology language (OWL).

---

<sup>2</sup>Also known as knowledge graph.

**Model Parameters.** Large scale pre-trained language models (LLMs) [5, 6] are rich in parametric knowledge. Their knowledge is accessed via knowledge probing [26, 37, 38]. These method constructs template prompts to require LLMs continue to output the desired knowledge following specified knowledge schemata.

**Other Knowledge Sources.** There are also other data structures: scene graph [39] encodes object instances, attributes of objects, and relationships between objects to denote the semantics of an image; tables can be transformed to knowledge graphs with heuristic rules [40].

### 3.2. Knowledge Manipulator

The implementation of knowledge manipulation primitives is twofold: how to design the atomic knowledge operations and how are the atomic knowledge operations executed.

**Knowledge Manipulation Operation Design.** Existing design for knowledge manipulations are classified as three main categories. (1)  *$\lambda$ -calculus Based.*  $\lambda$ -calculus is a general computational model, which is proved to be Turing complete [41].  $\lambda$ -calculus defines atomic operations as  $\lambda$ -functions to represent natural language questions [42]. To facilitate operating on knowledge bases, they evolve into  $\lambda$ -DCS [43–46]. (2) *Database Query Based.* Knowledge bases are usually stored in compatible with database, such as Virtuoso and Neo4j. Thus, it is direct to use graph database query primitives as knowledge manipulation primitives. They include SQL [47] for relational database, SPARQL [48] for graph database, and Cypher for attributed graph database. (3) *Knowledge-oriented Operations.* Most recently, researchers propose knowledge-oriented programming language (KoPL) [12], which designs operates on knowledge elements, such as `Relate` stands for “find entities that have certain relation with the given entity”.

**Knowledge Manipulation Operation Realization.** The realization varies according to the knowledge sources. Here we take `Relate` operation as the representative example to show potential realization strategies.

**Knowledge Graph Based.** The most direct way to realize knowledge manipulation on knowledge base is sub-graph matching. For example, if the triple (Mark Twain, writing language, English) is stored in the knowledge base, then we can directly find the entity that have the relation “writing language” with Mark Twain. To incorporate neural method, we also realize knowledge manipulations in a semantic space using knowledge embedding [49]. There are also attempt to model reasoning on knowledge graph as multi-hop random walk agent [50].

**Textual Corpora Based.** Another method is to employ the knowledge in text corpora with open question answering, e.g., answering the question “What is the written language of Mark Twain?”. Typically, open question answering adopts a two-stage paradigm, i.e., first retrieves relevant documents from the large-scale text corpora, and then performs reading comprehension over the documents to obtain the final answer.

**Language Model Based.** LLMs are capable of answering questions with the rich knowledge stored in models’ parameters. And the question “What is Mark Twain’s writing language?” can be answered by prompting the LLM to perform open-domain question answering.

### 3.3. Reasoning Planner

The planner requires the language understanding ability to translate natural language questions or instructions into explicit reasoning processes. As we introduced in Section 2.3, the translated reasoning process can take the forms of programs, chain-of-thought reasoning steps, or natural language sub-questions.

Here we would like to introduce the reasoning processes in details. (1) For *programs* composed of symbolic knowledge manipulators, Rational Meaning Construction [51] adopts Church programs to offer a structured representation for expressing novel situations and arbitrary problems with respect to a meaningful model over possible world states. W3C officially recommends SPARQL [52] for manipulating and querying RDF stores since 2008 [53]. KoPL [12] is a newly emerged knowledge-oriented programming language designed for querying knowledge base by executing a multi-step program that takes knowledge elements as the arguments. Other programs include Cypher [54],  $\lambda$ -DCS[43],  $\lambda$ -calculus [55], etc. These languages have their own advantages and drawbacks, and can be transpiled to each other with the help of intermediate languages such as GraphQ IR [36]. (2) For *chain-of-thought (CoT) reasoning*, they are composed of step-by-step natural language reasoning steps [56]. With CoT, the ability of LLMs

to perform complex reasoning is improved significantly. (3) For *natural language sub-questions*, a representative is Question Decomposition Meaning Representation (QDMR) in the BREAK dataset [57], which constitutes the ordered list of steps, expressed through natural language, that are necessary for answering a question.

For the planner model, we can take planning as a Seq2Seq translation task and employ an encoder-decoder framework for the model. Such methods usually rely on a large amount of labeled data, *i.e.*, question and reasoning process pairs [12], which are often lacking because such annotation is both expensive and labor-intensive. Another paradigm is to learn from question-answer pairs by taking the answers as weak supervision and searching for gold reasoning processes with reinforcement learning (RL) [58, 59], which is challenging due to combinatorial explosion in searching space along with extremely sparse rewards. Recently, with the remarkable generalizability and language understanding ability of LLMs, a new path is open by inferring the reasoning process with in-context learning [60]. For example, Program-aided Language models (PAL) uses the LLM to read natural language problems and generate programs as the intermediate reasoning steps, and offloads the solution step to a runtime such as a Python interpreter [61].

### 3.4. Reasoning Conductor

The reasoning conductor aims to obtain an optimal solution for the knowledge-intensive complex reasoning task. It needs the flexibility to plan dynamically and consider multiple factors to make a global choice, including whether the chosen knowledge resource is suitable, whether the planning route is the most efficient, and how much computational resource it costs. One possible way to dynamic reasoning conductor is probabilistic programming, *e.g.*, Rational Meaning Construction [51] employs an inference function that computes probabilities over the space of possible worlds consistent with and conditioned on information in the Church program to infer the optimal answer. Another direction is to optimize on tree-like computation graphs. Here we provide some examples.

For the task of answering complex logical queries on incomplete knowledge graphs, Query Computation Tree Optimization (QTO) [62] finds the optimal solution by a forward-backward propagation on a tree-like computation graph, *i.e.*, query computation tree. In particular, QTO utilizes the independence encoded in the query computation tree to reduce the search space, where only local computations are involved during the optimization procedure, *i.e.*, they find a set of entity assignments that maximizes the truth score of a first-order logical query based on the truth value of each one-hop atom provided by a knowledge embedding based link predictor.

For the task of answering complex natural language questions, RoHT [63] first builds the Hierarchical Question Decomposition Tree (HQDT) to understand the semantics of a complex question, and then conducts probabilistic reasoning over HQDT from root to leaves recursively, to aggregate heterogeneous knowledge at different tree levels and search for a best solution considering the decomposing and answering probabilities.

For the task of other complex reasoning tasks that needs search or planning such as Game of 24, ToT [64] utilizes LLMs to perform deliberate decision making. This involves considering multiple reasoning paths, evaluating options internally, and determining the most suitable next steps, as well as looking ahead or backtracking when necessary to make global choices.

### 3.5. Human in the loop

As stated in Section 2.5, it is of great benefit to involving human into the reasoning loop, spanning all four components of the overall architecture.

For the *knowledge-oriented data manager*, it takes the responsibility of providing up-to-date, accurate and comprehensive knowledge, in which humans can play an important role by sharing and editing knowledge. For example, Wikidata employs crowd-sourcing to extend and edit the stored information, even without creating an account, and the form-based interface makes editing easy [65]. The Wikipedia is also committed to “a world in which every single human being can freely share in the sum of all knowledge”.

For the *knowledge manipulator*, it takes the responsibility of accessing, processing, and updating knowledge with atomic actions. And humans can design, update and expand the set of atomic actions with their prior knowledge on the targeted complex reasoning task. For example, for the complex visual tasks such as factual knowledge object



1 tagging and language guided image editing, actions such as image resizing, cropping, filtering, and colorspace  
2 conversions are designed by experts [66].

3 For the *reasoning planner*, it is possible for the planner to improve its performance with the human feedback. For  
4 example, given an interaction interface, the system explains the predicted program step by step in natural language  
5 along with intermediate answers to the user, and then, the user can give feedback for the inspected wrong step, thus  
6 the current planner can accordingly be updated to improve accuracy. In another case, when the planner outputs an  
7 explainable program with transparent reasoning process, e.g., KoPL, the user can even edit with simple graphical  
8 operators, such as “dragging” to add knowledge operators and “slot filling” to designate operator arguments [67].

9 For the *reasoning conductor*, it is also able to improve its performance with human in the loop, if humans have  
10 prior knowledge for problem solving, *i.e.*, humans can select the most optimal solution for the conductor, and such  
11 data can be collected as preference data to train the conductor with reinforcement learning from human feedback  
12 (RLHF) [68]. For example, for WebGPT [31], guidance from humans is central. They collected demonstrations  
13 and comparisons from humans to train neural models to search and navigate the web. Specifically, they collected  
14 examples of humans using the browser to answer questions (demonstrations) to teach neural models to use their  
15 text-based browser. And further, they collected pairs of model-generated answers to the same question and asked  
16 humans which one they preferred (comparisons) to improve the optimize answer quality.

#### 17 18 19 **4. Conclusion and Future Directions**

20  
21 Inspired by the promising capability to incorporate both neural and symbolic models [13, 20], in this paper, we  
22 design an initiative architecture to solve knowledge-intensive complex reasoning tasks. The architecture not only  
23 utilizes the capability to understand rough descriptions and rich knowledge of LLMs, but also enables symbol  
24 system to constraint the behavior of LLMs to reduce hallucination. Although there are many pioneer techniques that  
25 indicates the feasibility of our proposed architecture, as we suggest in Section 3, it still requires efforts to build the  
26 overall architecture. We believe that future works mainly fall primarily into two aspects: the macroscopic and the  
27 microscopic.

28 At the macroscopic level, the challenge lies in selecting and integrating specific techniques to implement the four  
29 key components. This is non-trivial, as we are essentially designing standardised behaviour for different algorithms  
30 to coordinate to achieve the same goal. For example, the problem solver may accept plans given by both LLMs and  
31 traditional syntactic-based semantic parser, which poses a requirement on the consistency between the output of  
32 both sides. This entails not just cherry-picking the best performers within each module but also ensuring that they  
33 can effectively communicate with each other and function as a cohesive system.

34 At the microscopic level, we should not stop the pace to design better implementations for each component and  
35 atomic operations for the knowledge manipulation primitives. We propose two main research directions for future  
36 work. First, we should explore methods to constrain LLMs to perform well-defined symbolic reasonings. Existing  
37 research includes prompt template design and constraint decoding. However, there is still no guarantee to obtain  
38 behaviour following strict schemata. Second, it is also crucial to design coherent and user-friendly interfaces for  
39 each component to incorporate human in the loop. The interface should provide information as close to natural  
40 language as possible to allow human users to inspect the working status of the system. It also needs to provide  
41 guidance and convert human user actions into consistent behaviour of the intervened component in the system. In  
42 addition, we need to explore more interactive systems, such as in the form of dialogue, to carry out the functionality  
43 of the system.

44 In conclusion, we specify the necessity, responsibility, and potential implementation techniques for each com-  
45 ponent to achieve a general architecture for knowledge-intensive complex reasoning tasks. Although there are still  
46 insufficient consideration with regard to the whole design and implementation details, it is intriguing to realize the  
47 whole system to achieve its potential capacity in the near future.

#### 48 49 **References**

- 50  
51 [1] A. Newell, Physical Symbol Systems, *Cogn. Sci.* **4** (1980), 135–183.

- [2] R. Salakhutdinov, Deep learning, in: *The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14, New York, NY, USA - August 24 - 27, 2014*, S.A. Macskassy, C. Perlich, J. Leskovec, W. Wang and R. Ghani, eds, ACM, 2014, p. 1973. doi:10.1145/2623330.2630809.
- [3] D.E. Rumelhart, G.E. Hinton and R.J. Williams, Learning internal representations by error propagation, 1986.
- [4] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L. Kaiser and I. Polosukhin, Attention is All you Need, in: *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, I. Guyon, U. von Luxburg, S. Bengio, H.M. Wallach, R. Fergus, S.V.N. Vishwanathan and R. Garnett, eds, 2017, pp. 5998–6008. <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.
- [5] OpenAI, GPT-4 Technical Report, *ArXiv preprint abs/2303.08774* (2023). <https://arxiv.org/abs/2303.08774>.
- [6] S. Bubeck, V. Chandrasekaran, R. Eldan, J.A. Gehrke, E. Horvitz, E. Kamar, P. Lee, Y.T. Lee, Y.-F. Li, S.M. Lundberg, H. Nori, H. Palangi, M.T. Ribeiro and Y. Zhang, Sparks of Artificial General Intelligence: Early experiments with GPT-4, *ArXiv preprint abs/2303.12712* (2023). <https://arxiv.org/abs/2303.12712>.
- [7] N.F. McGlynn, Thinking fast and slow., *Australian veterinary journal* **92** **12** (2014), N21. <https://api.semanticscholar.org/CorpusID:36031679>.
- [8] A. Goyal and Y. Bengio, Inductive biases for deep learning of higher-level cognition, *Proceedings of the Royal Society A* **478** (2020).
- [9] B. Zhang, Knowledge and Artificial Intelligence, *Sino-German Symposium on "Integrating Symbolic Representation with Numeric Representation for Commonsense Reasoning"* (2019).
- [10] F. Petroni, A. Piktus, A. Fan, P. Lewis, M. Yazdani, N.D. Cao, J. Thorne, Y. Jernite, V. Plachouras, T. Rocktaschel and S. Riedel, KILT: a Benchmark for Knowledge Intensive Language Tasks, in: *North American Chapter of the Association for Computational Linguistics*, 2020. <https://api.semanticscholar.org/CorpusID:221507798>.
- [11] Z. Yang, P. Qi, S. Zhang, Y. Bengio, W. Cohen, R. Salakhutdinov and C.D. Manning, HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering, in: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 2369–2380. doi:10.18653/v1/D18-1259. <https://aclanthology.org/D18-1259>.
- [12] S. Cao, J. Shi, L. Pan, L. Nie, Y. Xiang, L. Hou, J. Li, B. He and H. Zhang, KQA Pro: A Dataset with Explicit Compositional Programs for Complex Question Answering over Knowledge Base, in: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, Dublin, Ireland, 2022, pp. 6101–6119. doi:10.18653/v1/2022.acl-long.422. <https://aclanthology.org/2022.acl-long.422>.
- [13] J. McCarthy, An example for natural language understanding and the AI problems it raises, *Formalizing Common Sense: Papers by John McCarthy* **355** (1990).
- [14] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit and N. Houlsby, An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale, in: *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, OpenReview.net, 2021. <https://openreview.net/forum?id=YicbFdNTTy>.
- [15] A. Roberts, C. Raffel and N. Shazeer, How Much Knowledge Can You Pack Into the Parameters of a Language Model?, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020, pp. 5418–5426.
- [16] A. Drozdov, N. Schärli, E. Akyürek, N. Scales, X. Song, X. Chen, O. Bousquet and D. Zhou, Compositional Semantic Parsing with Large Language Models, in: *The Eleventh International Conference on Learning Representations*, 2022.
- [17] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman et al., Do as i can, not as i say: Grounding language in robotic affordances, *arXiv preprint arXiv:2204.01691* (2022).
- [18] Z. Yao, W. Zhang, M. Chen, Y. Huang, Y. Yang and H. Chen, Analogical inference enhanced knowledge graph embedding, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37, 2023, pp. 4801–4808.
- [19] H. Ren, W. Hu and J. Leskovec, Query2box: Reasoning over knowledge graphs in vector space using box embeddings, *arXiv preprint arXiv:2002.05969* (2020).
- [20] Y. Bengio et al., From system 1 deep learning to system 2 deep learning, in: *Neural Information Processing Systems*, 2019.
- [21] S. Wolfram, Chatgpt gets its' wolfram superpowers', *Recuperado de https://writings.stephenwolfram.com/2023/03/chatgpt-gets-its-wolfram-superpowers* (2023).
- [22] K.D. Bollacker, C. Evans, P.K. Paritosh, T. Sturge and J. Taylor, Freebase: a collaboratively created graph database for structuring human knowledge, in: *SIGMOD Conference*, 2008.
- [23] D. Vrandeic and M. Krotzsch, Wikidata: a free collaborative knowledgebase, *Commun. ACM* **57** (2014), 78–85.
- [24] P.S.H. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W. Yih, T. Rocktäschel, S. Riedel and D. Kiela, Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks, in: *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan and H. Lin, eds, 2020. <https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html>.
- [25] Y. Mao, P. He, X. Liu, Y. Shen, J. Gao, J. Han and W. Chen, Generation-Augmented Retrieval for Open-Domain Question Answering, in: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Association for Computational Linguistics, Online, 2021, pp. 4089–4100. doi:10.18653/v1/2021.acl-long.316. <https://aclanthology.org/2021.acl-long.316>.

- [26] F. Petroni, T. Rocktäschel, S. Riedel, P. Lewis, A. Bakhtin, Y. Wu and A. Miller, Language Models as Knowledge Bases?, in: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Association for Computational Linguistics, Hong Kong, China, 2019, pp. 2463–2473. doi:10.18653/v1/D19-1250. <https://aclanthology.org/D19-1250>.
- [27] K. Tang, Y. Niu, J. Huang, J. Shi and H. Zhang, Unbiased Scene Graph Generation From Biased Training, in: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, IEEE, 2020, pp. 3713–3722. doi:10.1109/CVPR42600.2020.00377.
- [28] X. Dong, E. Gabrilovich, G. Heitz, W. Horn, N. Lao, K. Murphy, T. Strohmman, S. Sun and W. Zhang, Knowledge vault: a web-scale approach to probabilistic knowledge fusion, in: *The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14, New York, NY, USA - August 24 - 27, 2014*, S.A. Macskassy, C. Perlich, J. Leskovec, W. Wang and R. Ghani, eds, ACM, 2014, pp. 601–610. doi:10.1145/2623330.2623623.
- [29] M. Galkin, P. Trivedi, G. Maheshwari, R. Usbeck and J. Lehmann, Message Passing for Hyper-Relational Knowledge Graphs, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Association for Computational Linguistics, Online, 2020, pp. 7346–7359. doi:10.18653/v1/2020.emnlp-main.596. <https://aclanthology.org/2020.emnlp-main.596>.
- [30] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P.N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer and C. Bizer, DBpedia - A large-scale, multilingual knowledge base extracted from Wikipedia, *Semantic Web* **6** (2015), 167–195.
- [31] R. Nakano, J. Hilton, S.A. Balaji, J. Wu, L. Ouyang, C. Kim, C. Hesse, S. Jain, V. Kosaraju, W. Saunders, X. Jiang, K. Cobbe, T. Eloundou, G. Krueger, K. Button, M. Knight, B. Chess and J. Schulman, WebGPT: Browser-assisted question-answering with human feedback, *ArXiv preprint [abs/2112.09332](https://arxiv.org/abs/2112.09332)* (2021). <https://arxiv.org/abs/2112.09332>.
- [32] N.D. Cao, G. Izacard, S. Riedel and F. Petroni, Autoregressive Entity Retrieval, in: *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, OpenReview.net, 2021. <https://openreview.net/forum?id=5k8F6UU39V>.
- [33] X. Han, H. Zhu, P. Yu, Z. Wang, Y. Yao, Z. Liu and M. Sun, FewRel: A Large-Scale Supervised Few-Shot Relation Classification Dataset with State-of-the-Art Evaluation, in: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 4803–4809. doi:10.18653/v1/D18-1514. <https://aclanthology.org/D18-1514>.
- [34] M. Banko, M.J. Cafarella, S. Soderland, M. Broadhead and O. Etzioni, Open Information Extraction from the Web, in: *CACM*, 2007.
- [35] H. Jin, L. Hou, J. Li and T. Dong, Fine-Grained Entity Typing via Hierarchical Multi Graph Convolutional Networks, in: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Association for Computational Linguistics, Hong Kong, China, 2019, pp. 4969–4978. doi:10.18653/v1/D19-1502. <https://aclanthology.org/D19-1502>.
- [36] L. Nie, S. Cao, J. Shi, J. Sun, Q. Tian, L. Hou, J. Li and J. Zhai, GraphQIR: Unifying the Semantic Parsing of Graph Query Languages with One Intermediate Representation, in: *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 2022, pp. 5848–5865. <https://aclanthology.org/2022.emnlp-main.394>.
- [37] Z. Zhong, D. Friedman and D. Chen, Factual Probing Is [MASK]: Learning vs. Learning to Recall, in: *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Association for Computational Linguistics, Online, 2021, pp. 5017–5033. doi:10.18653/v1/2021.naacl-main.398. <https://aclanthology.org/2021.naacl-main.398>.
- [38] H. Peng, X. Wang, S. Hu, H. Jin, L. Hou, J. Li, Z. Liu and Q. Liu, COPEN: Probing Conceptual Knowledge in Pre-trained Language Models, in: *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 2022, pp. 5015–5035. <https://aclanthology.org/2022.emnlp-main.335>.
- [39] J. Johnson, R. Krishna, M. Stark, L. Li, D.A. Shamma, M.S. Bernstein and F. Li, Image retrieval using scene graphs, in: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, IEEE Computer Society, 2015, pp. 3668–3678. doi:10.1109/CVPR.2015.7298990.
- [40] I. Saparina and A. Osokin, SPARQLing Database Queries from Intermediate Question Decompositions, in: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 2021, pp. 8984–8998. doi:10.18653/v1/2021.emnlp-main.708. <https://aclanthology.org/2021.emnlp-main.708>.
- [41] A.M. Turing, Computability and  $\lambda$ -definability, *The Journal of Symbolic Logic* **2**(4) (1937), 153–163.
- [42] Y. Artzi, N. FitzGerald and L. Zettlemoyer, Semantic Parsing with Combinatory Categorical Grammars, in: *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Tutorials)*, Association for Computational Linguistics, Sofia, Bulgaria, 2013, p. 2. <https://aclanthology.org/P13-5002>.
- [43] P. Liang, Lambda Dependency-Based Compositional Semantics, *CoRR [abs/1309.4408](https://arxiv.org/abs/1309.4408)* (2013). <http://arxiv.org/abs/1309.4408>.
- [44] P. Pasupat and P. Liang, Compositional Semantic Parsing on Semi-Structured Tables, in: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Association for Computational Linguistics, Beijing, China, 2015, pp. 1470–1480. doi:10.3115/v1/P15-1142. <https://aclanthology.org/P15-1142>.
- [45] Y. Wang, J. Berant and P. Liang, Building a Semantic Parser Overnight, in: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Association for Computational Linguistics, Beijing, China, 2015, pp. 1332–1342. doi:10.3115/v1/P15-1129. <https://aclanthology.org/P15-1129>.
- [46] P. Pasupat and P. Liang, Inferring Logical Forms From Denotations, in: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, Berlin, Germany, 2016, pp. 23–32. doi:10.18653/v1/P16-1003. <https://aclanthology.org/P16-1003>.

- [47] V. Zhong, C. Xiong and R. Socher, Seq2sql: Generating structured queries from natural language using reinforcement learning, *ArXiv preprint abs/1709.00103* (2017). <https://arxiv.org/abs/1709.00103>.
- [48] Y. Sun, L. Zhang, G. Cheng and Y. Qu, SPARQA: Skeleton-Based Semantic Parsing for Complex Questions over Knowledge Bases, in: *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, AAAI Press, 2020, pp. 8952–8959. <https://aaai.org/ojs/index.php/AAAI/article/view/6426>.
- [49] A. Bordes, N. Usunier, A. García-Durán, J. Weston and O. Yakhnenko, Translating Embeddings for Modeling Multi-relational Data, in: *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, C.J.C. Burges, L. Bottou, Z. Ghahramani and K.Q. Weinberger, eds, 2013, pp. 2787–2795. <https://proceedings.neurips.cc/paper/2013/hash/1cecc7a77928ca8133fa24680a88d2f9-Abstract.html>.
- [50] X.V. Lin, R. Socher and C. Xiong, Multi-Hop Knowledge Graph Reasoning with Reward Shaping, in: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 3243–3253. doi:10.18653/v1/D18-1362. <https://aclanthology.org/D18-1362>.
- [51] L.S. Wong, G. Grand, A.K. Lew, N.D. Goodman, V.K. Mansinghka, J. Andreas and J.B. Tenenbaum, From Word Models to World Models: Translating from Natural Language to the Probabilistic Language of Thought, *ArXiv preprint abs/2306.12672* (2023). <https://arxiv.org/abs/2306.12672>.
- [52] W.W.W. Consortium, SPARQL 1.1 overview (2013).
- [53] E. Prud’hommeaux and A. Seaborne, SPARQL Query Language for RDF, 2008. <http://www.w3.org/TR/rdf-sparql-query/>.
- [54] N. Francis, A. Green, P. Guagliardo, L. Libkin, T. Lindaaker, V. Marsault, S. Plantikow, M. Rydberg, P. Selmer and A. Taylor, Cypher: An Evolving Query Language for Property Graphs, in: *Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018, Houston, TX, USA, June 10-15, 2018*, G. Das, C.M. Jermaine and P.A. Bernstein, eds, ACM, 2018, pp. 1433–1445. doi:10.1145/3183713.3190657.
- [55] H.P. Barendregt et al., *The lambda calculus*, Vol. 3, North-Holland Amsterdam, 1984.
- [56] H. Trivedi, N. Balasubramanian, T. Khot and A. Sabharwal, Interleaving Retrieval with Chain-of-Thought Reasoning for Knowledge-Intensive Multi-Step Questions, in: *Annual Meeting of the Association for Computational Linguistics*, 2022. <https://api.semanticscholar.org/CorpusID:254877499>.
- [57] T. Wolfson, M. Geva, A. Gupta, M. Gardner, Y. Goldberg, D. Deutch and J. Berant, Break It Down: A Question Understanding Benchmark, *Transactions of the Association for Computational Linguistics* **8** (2020), 183–198.
- [58] S. Cao, J. Shi, Z. Yao, X. Lv, J. Yu, L. Hou, J. Li, Z. Liu and J. Xiao, Program Transfer for Answering Complex Questions over Knowledge Bases, in: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, Dublin, Ireland, 2022, pp. 8128–8140. doi:10.18653/v1/2022.acl-long.559. <https://aclanthology.org/2022.acl-long.559>.
- [59] C. Liang, J. Berant, Q. Le, K.D. Forbus and N. Lao, Neural Symbolic Machines: Learning Semantic Parsers on Freebase with Weak Supervision, in: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, Vancouver, Canada, 2017, pp. 23–33. doi:10.18653/v1/P17-1003. <https://aclanthology.org/P17-1003>.
- [60] L. Tianle, X. Ma, A. Zhuang, Y. Gu, Y. Su and W. Chen, Few-shot In-context Learning on Knowledge Base Question Answering, *ArXiv preprint abs/2305.01750* (2023). <https://arxiv.org/abs/2305.01750>.
- [61] L. Gao, A. Madaan, S. Zhou, U. Alon, P. Liu, Y. Yang, J. Callan and G. Neubig, PAL: Program-aided Language Models, *ArXiv preprint abs/2211.10435* (2022). <https://arxiv.org/abs/2211.10435>.
- [62] Y. Bai, X. Lv, J. Li and L. Hou, Answering Complex Logical Queries on Knowledge Graphs via Query Computation Tree Optimization, *ArXiv preprint abs/2212.09567* (2022). <https://arxiv.org/abs/2212.09567>.
- [63] J. Zhang, S. Cao, T. Zhang, X. Lv, J. Shi, Q. Tian, J. Li and L. Hou, Reasoning over Hierarchical Question Decomposition Tree for Explainable Question Answering, *ArXiv preprint abs/2305.15056* (2023). <https://arxiv.org/abs/2305.15056>.
- [64] S. Yao, D. Yu, J. Zhao, I. Shafran, T.L. Griffiths, Y. Cao and K. Narasimhan, Tree of Thoughts: Deliberate Problem Solving with Large Language Models, *ArXiv preprint abs/2305.10601* (2023). <https://arxiv.org/abs/2305.10601>.
- [65] D. Vrandečić and M. Krötzsch, Wikidata: a free collaborative knowledgebase, *Commun. ACM* **57**(10) (2014), 78–85. doi:10.1145/2629489.
- [66] T. Gupta and A. Kembhavi, Visual Programming: Compositional visual reasoning without training, *ArXiv preprint abs/2211.11559* (2022). <https://arxiv.org/abs/2211.11559>.
- [67] Z. Yao, Y.-K. Chen, X. Lv, S. Cao, A. Xin, J. Yu, H. Jin, J. Xu, P. Zhang, L. Hou and J. Li, VisKoP: Visual Knowledge oriented Programming for Interactive Knowledge Base Question Answering, *ArXiv preprint abs/2307.03130* (2023). <https://arxiv.org/abs/2307.03130>.
- [68] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C.L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L.E. Miller, M. Simens, A. Askell, P. Welinder, P.F. Christiano, J. Leike and R.J. Lowe, Training language models to follow instructions with human feedback, *ArXiv preprint abs/2203.02155* (2022). <https://arxiv.org/abs/2203.02155>.