

# Neuro-symbolic Training for Spatial Reasoning over Natural Language

Tanawan Premstri

Department of Computer Science  
Michigan State University  
premsrit@msu.edu

Parisa Kordjamshidi

Department of Computer Science  
Michigan State University  
kordjams@msu.edu

## Abstract

Spatial reasoning based on natural language expressions is essential for everyday human tasks. This reasoning ability is also crucial for machines to interact with their environment in a human-like manner. However, recent research shows that even state-of-the-art language models struggle with spatial reasoning over text, especially when facing nesting spatial expressions. This is attributed to not achieving the right level of abstraction required for generalizability. To alleviate this issue, we propose training language models with neuro-symbolic techniques that exploit the spatial logical rules as constraints, providing additional supervision to improve spatial reasoning and question answering. Training language models to adhere to spatial reasoning rules guides them in making more effective and general abstractions for transferring spatial knowledge to various domains. We evaluate our approach on existing spatial question-answering benchmarks. Our results indicate the effectiveness of our proposed technique in improving language models in complex multi-hop spatial reasoning over text.

## 1 Introduction

Spatial language conveys information about the location and translocation of objects/entities in space. Reasoning over spatial language is required for natural language processing systems to answer "Where" questions and identify the relative spatial relations between objects explained in the text, such as left and right. Figure 1, demonstrates a spatial language comprehension and question-answering problem. Understanding spatial language turns out to be crucial for many AI applications, including question-answering systems, grounding language into visual modality (Zhang et al., 2021; Zhang and Kordjamshidi, 2022; Liu et al., 2023), navigating robots (Sisbot et al., 2007; Yadollahi et al., 2023) and even in more specific

<b>SpartQA-Human</b> <b>Context:</b> There are three blocks called A, B, and C. A is to the left of B and B is to the left of C. In A, there is a small yellow square near and to the left of a large yellow circle. The large circle is near and to the left of a large blue triangle. In B, there is a small yellow square near and below a large blue triangle. There is a medium black triangle near and above the large blue triangle. In C, there is a small black square far to the left and above a small blue circle. The blue circle is near and to the left of a large yellow square. <b>YN Question:</b> Is the small black thing to the right of medium black thing? <b>Answer:</b> Yes <b>FR Question:</b> What is the position of small black thing relative to medium black thing? <b>Answer:</b> Right	<b>ReSQ</b> <b>Context:</b> A tourist group is visiting a school. The school kids are sitting at their desks, singing a song and clapping their hands, with their teacher clapping as well. Six of the tourists are standing behind them and are watching them, one tourist is taking a photo. <b>YN Question:</b> Are the tourists in the school? <b>Answer:</b> Yes <b>StepGame</b> <b>Context:</b> S is positioned below C. The object B is upper and slightly to the left of the object O. U is below S and to the right of S. V and C are next to each other with C at the bottom V on the top. A and U are parallel, and A on the right of U. <b>FR Question:</b> What is the relation of the agent S to the agent C? <b>Answer:</b> below
--	---

Figure 1: Examples of story and questions in Spatial Question Answering (SQA). YN is Yes-No, and FR is Find relation type of questions.

fields such as medical domain (Atif et al., 2007; Datta et al., 2020; Gong et al., 2024).

Despite the high performance of Large Language models (LLMs) in many NLP and vision problems (Brown et al., 2020; OpenAI, 2023), recent studies highlight the shortcomings in complex spatial reasoning over text (Bang et al., 2023; Yang et al., 2023; Mirzaee and Kordjamshidi, 2023; Liu et al., 2023; Chen et al., 2024). This paper focuses on complex, that is, multi-hop spatial reasoning over text. However, our research findings can advance complex reasoning over text in general. Our main idea is to exploit spatial logical rules of reasoning in training language models.

Explicit logic, in combination with LLMs, has been used in some related works. In particular, Yang et al. 2023 integrated explicit symbolic reasoning with LLMs in a pipeline. The LLMs map natural language to symbolic representation, and this formal representation is passed to an Answer Set Programming (ASP) solver to conduct logical reasoning. Similarly, Mirzaee and Kordjamshidi 2023 passed formal representations obtained from LLMs to Prolog to conduct logical reasoning and answer spatial questions. These approaches rely on external reasoning tools and require accurate

formal representations generated at inference time. Obtaining such formal representations for realistic domain natural language is hard and, as shown in (Mirzaee and Kordjamshidi, 2023), limits their effectiveness only to synthetic benchmarks.

Here, we aim to address the limitations of utilizing spatial logic in realistic domains. In our proposed technique, we fine-tune language models using spatial logical rules to guide them in capturing the level of abstraction needed for general spatial reasoning. In particular, we train the models to minimize both cross-entropy loss and violation from logical constraints. We hypothesize that incorporating logical knowledge as a source of supervision, even using synthetic data, enriches the models’ abstract reasoning over spatial concepts, further enhancing their generalizability. A key advantage of our approach is that it does not require full coverage of logical knowledge. Any partially available logical knowledge can be utilized during training. Additionally, formal knowledge representation is not needed at inference time, which is crucial because generating formal representations during inference is a challenging task in itself. Our technique also avoids reliance on external reasoning tools, which can add computational complexity in real-time applications.

We use the available three benchmarks for spatial question answering, SpartQA-Human (Mirzaee et al., 2021), ReSQ (Mirzaee and Kordjamshidi, 2022), and StepGame (Shi et al., 2022) to evaluate our proposed method. Our improved results confirm our hypothesis about the impact of neuro-symbolic training on generalizability. The contribution of this paper can be summarized as follows<sup>1</sup>, (1) We exploit the available logical spatial knowledge in tuning language models to improve their spatial reasoning. (2) We provide extensive experimental results and analysis on the existing benchmarks. Our results consistently confirm the effectiveness of our proposed approach in both encoder-based and generative language models and their generalizability for transferring knowledge to different domains; (3) We show that tuning smaller models with our proposed method generalizes better on out-of-domain and complex reasoning compared to using prompt-engineered, very large language models.

<sup>1</sup>The code, trained models, and data are publicly available at <https://github.com/HLR/SpaRTUNQChain>

## 2 Related Work

Investigating the reasoning capability of NLP models has become an important research topic given the instability in the performance of LLMs (Gendron et al., 2024; Valmeekam et al., 2022; Feng et al., 2024; Chen et al., 2024; Bang et al., 2023). Typically, this ability is assessed through question-answering (QA) benchmarks (Weston et al., 2016; Zhou et al., 2020; Tan et al., 2023). However, fewer studies are focused on spatial reasoning over text. Recent benchmarks, such as SpaRTUN (Mirzaee and Kordjamshidi, 2022) and StepGame (Shi et al., 2022; Li et al., 2024), have been developed to address this issue by providing comprehensive evaluation resources for spatial question answering. The current results using the existing benchmarks indicate that LLMs struggle with spatial reasoning over text, even on synthetic datasets (Mirzaee and Kordjamshidi, 2022; Yang et al., 2023; Bang et al., 2023).

Several works have attempted to improve LLMs’ spatial reasoning capabilities. Mirzaee and Kordjamshidi 2022 fine-tunes LLMs with synthetic data, which improves their multi-hop reasoning abilities, even when applied to realistic domains. Another approach in (Shi et al., 2022) designed a specialized memory network based on a recurrent neural network for deep reasoning tasks, but it was less effective than tuning pre-trained language models (PLMs). Some papers focus on enhancing spatial reasoning by in-context learning. For instance, Sharma 2023 proposed a prompting technique that decomposes the target spatial question into simpler questions and infers the answer to the final query based on the information gained from them. Another approach, proposed by Hu et al. 2024, replaced textual explanations in the Chain-of-Thought (CoT) by symbolic representations. While this method improved spatial understanding in simple environments, it was less applicable to diverse and complex spatial relations settings. Lastly, some studies employ neuro-symbolic modeling that involves LLMs in the pipeline (Yang et al., 2023; Mirzaee and Kordjamshidi, 2023). These approaches extract relations and obtain formal representations using LLMs. Then, the formal representation is passed to an external solver that includes engineered rules of reasoning to conduct logical reasoning and answer spatial questions. While this method significantly improved multi-hop reasoning, it required formal representations

at inference time, which posed a challenge. Furthermore, formalizing commonsense reasoning is difficult, if not impossible, for realistic domains.

Unlike the previous studies, we exploit available spatial logical knowledge as a set of constraints in fine-tuning language models. Compared to hard pipelines, this approach has the advantage that logical knowledge does not need to be complete and cover all possible situations. Our approach utilizes synthetic data and logical representations during training only. Numerous research studies have highlighted the utility of logical knowledge in solving question-answering tasks. For instance, [Prager et al. 2004](#) demonstrates that answering sets of questions jointly using their logical relationships enhances the performance in reasoning tasks. [Amizadeh et al. 2020](#) directly used the logical formalization of the question as input to the model to improve the logical reasoning of vision and language models. [Sun et al. 2022](#) proposed framework with two main components. The first component is for encoding logical knowledge, and the second is for entailment of that knowledge. Lastly, [Kamali et al. 2025](#) proposed a framework that integrates logical reasoning with generalizable grounding interpretation to enhance both grounding and reasoning in vision and language models.

Utilizing logical knowledge has been proven effective in solving many other NLP problems. For example, [Lu et al. 2022](#) uses a heuristic function inspired by the A\* algorithm to restrict the output of generative language models with a set of constraints, while [Qin et al. 2022](#) uses a sampling method looking for the most optimal solution that satisfies the constraints. Both of these techniques apply the constraints during inference, which differs from our proposed method for using them during the training. Lastly, [Rajaby Faghihi et al. 2023](#) provided a benchmark to evaluate the benefits of logical constraints in several NLP problems under the DomiKnowS framework ([Rajaby Faghihi et al., 2021](#)). In this paper, we propose a novel approach for formulating spatial reasoning knowledge as logical constraints and exploit the existing techniques for the actual integration. We utilize constraints solely during training, eliminating the need for access to logical constraints during inference where using them might be problematic for realistic domains and real-time applications.

### 3 Methodology

#### 3.1 Problem Definition

For the spatial QA task, the input includes a textual context  $C$  that describes a scene and a question  $Q$  about the spatial relations between objects in the scene. The output is the answer(s) to the question, denoted as  $A$ . There are two types of questions/answers: Yes/No (YN) and find-relation (FR). The answer domain is restricted to  $A = \{Yes, No\}$  for YN questions. The answers to FR questions depend on the dataset and the types of relations covered. For example, the answer domain for FR can be  $A = \{left, right, above, below, \dots\}$ .

#### 3.2 Backbone Language Models

We use both encoder-based models from the BERT family ([Devlin et al., 2019](#)) as well as generative models like Flan-T5 ([Chung et al., 2024](#)) to assess the effectiveness of our proposed fine-tuning approach. We used other LLMs, such as GPT family and Llama3-8B, as comparison baselines without fine-tuning.

##### 3.2.1 Encoder-based Language Models

We utilize BERT as the backbone architecture to consistently build on top of previous results in [Mirzaee and Kordjamshidi \(2022\)](#). We fine-tune BERT by adding classification layers for QA tasks. We employ a binary classification layer for the YN domain while using multiple binary classification layers for the FR domain. The motivation behind using multiple binary classification layers instead of one multi-label classification layer is that one question can have multiple labels as responses. The number of binary classification layers in the FR domain depends on the number of possible answers,  $|A|$ , in the problem setting. The input to the BERT model is constructed by concatenating the question,  $Q$ , separator token,  $[SEP]$ , and the context  $C$ . After feeding the input to BERT, we use the  $[CLS]$  token from the last layer as the input to the classification layer(s) for final answer prediction.

##### 3.2.2 Generative Language Models

We select Flan-T5 ([Chung et al., 2024](#)) as the open-source generative baseline for fine-tuning. In the YN setting, processing the output is straightforward, as the first token can represent the response directly generated by the model. We compare the probability of the Yes and No tokens and select the one with the highest probability as the answer.

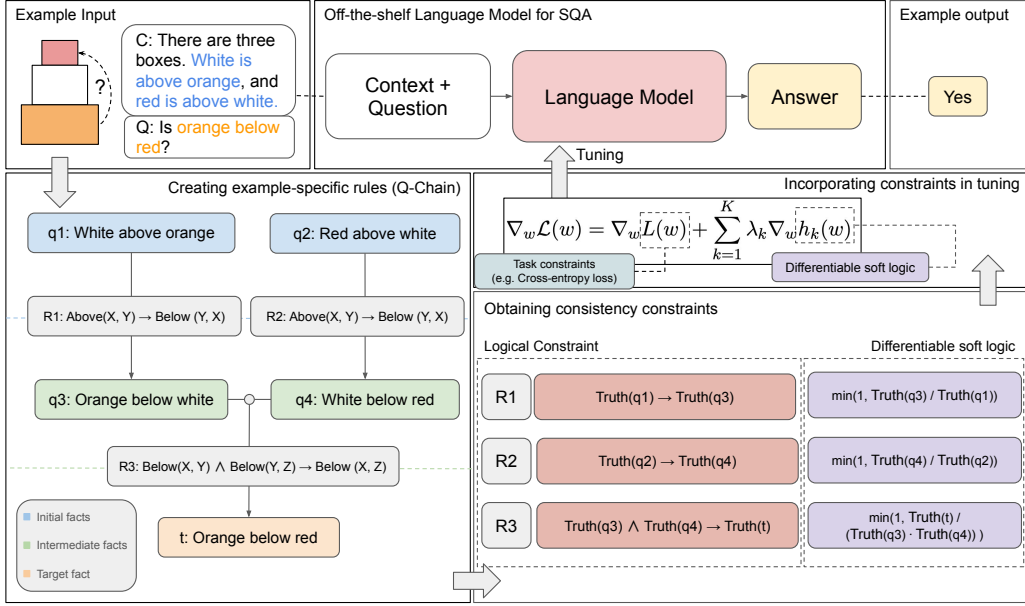


Figure 2: The pipeline for exploiting spatial logic in the training phase involves three steps: 1. Creating the chain of reasoning (Q-Chain) for each training example using annotated facts and rules from the training dataset. 2. Obtaining consistency constraints from spatial logical rules and answers inferred from questions generated using intermediate facts in Q-Chain. 3. Incorporating obtained consistency constraints into the loss function using differentiable soft logic during tuning. For more details, refer to Section 3.3.

However, this approach is not feasible in the FR setting due to the multi-label questions requiring multiple answers. The input prompt to the Flan-T5 model is structured as follows: “Answer based on the context:  $C$   $Q$ .”

For In-context Learning, we select Llama3-8B (AI@Meta, 2024), GPT3.5 (Brown et al., 2020), and GPT-4 (OpenAI, 2023) as the baseline LLMs using *Zero-shot*, and *Few-shot* settings.

**Zero-shot.** We prompt the LLMs to provide answers based on the scene description and the question without including any examples.

**Few-shot.** We randomly sample four pairs of questions and answers from the training dataset and concatenate them to the context and target question to form the few-shot setting.

### 3.3 Training with Spatial Logic

To address the challenge of multi-hop spatial reasoning, we employ symbolic knowledge during training. We assume a formal knowledge base (KB) of rules, referred to as *spatial logic*, is given during training time. This spatial logic does not need to be exhaustive or cover all reasoning aspects; any available knowledge can be used to enhance the reasoning capabilities of the language model. This work uses 79 reasoning rules taken from Mirzaee and Kordjamshidi (2022). These

rules are divided into four categories: converse, inverse, transitive, and transitive + topological. The KB covers rules between 15 spatial relations, including directional and topological relations, such as “If  $x$  is above  $y$ , then  $y$  below  $x$ .” denoted formally as  $Above(x, y) \Rightarrow Below(y, x)$ .

Our main hypothesis is that supervision from high-level logical knowledge allows the model to capture more abstract patterns, thereby improving generalization to other domains. To exploit the spatial logic in the training phase, we follow three steps: (1) Creating the chain of reasoning for each training example. (2) Obtaining consistency constraints from the chain of reasoning. (3) Incorporating constraints in the loss function. Please note that we will need the reasoning chain annotations for fine-tuning with spatial logic. SpaRTUN is the only synthetic dataset that provides such annotated formal representations, making it a suitable source of supervision for our tuning purpose.

**Creating example-specific rules.** We exploit *Spatial Logic* by automatically augmenting each training example with a chain of questions, denoted as *Q-Chain*. This process is illustrated in Figure 2. We infer the answer to the final query from the input context using the *logical resolution tree*, which outlines the logical steps for implication. Our synthetic training data (i.e., SpaRTUN) provides logi-



cal representations of the initial facts expressed in the context. We build the resolution tree using the *forward chaining algorithm* to infer a specific target fact given a set of initial facts. In our example, the initial facts are “The white is above the orange”, denoted as  $q_1$ , and “The red is above the white”, denoted as  $q_2$ . The target fact is “The orange below the red,” denoted as  $t$ . The forward chaining approach is a reasoning method that starts with the initial facts in the knowledge base. It performs an exhaustive search for all applicable inference rules and applies them to the current facts to iteratively infer new ones until the target fact is reached or no new facts are inferred (Nickerson, 1994). In the example of Figure 2, we apply the spatial converse rules to the initial fact  $q_1$  and  $q_2$ . This approach yields the intermediate inferred facts,  $q_3$ , “The orange is below the white,” from  $q_1$ , and  $q_4$ , “The white is below the red,” from  $q_2$ . Finally, we search for the applicable spatial rule to combine  $q_3$  and  $q_4$ , resulting in the target fact,  $t$ , “The orange is below the red.”

**Obtaining consistency constraints.** In this step, we obtain a set of consistency constraints based on the given resolution tree. We denote the truth value of a fact  $q_i$  in the Q-Chain as  $Truth(q_i)$ , which will be True if the fact is True and otherwise is False. We create a consistency constraint for each step of the Q-Chain. For example, in Figure 2, step  $R_1$  is applied to  $q_1$  to infer  $q_3$ . This step implies that the  $Truth(q_1)$  should be consistent with the  $Truth(q_3)$ . In other words,  $q_1$ , “white above orange,” implies  $q_3$ : “orange below white.” The corresponding formal constraint for this example will be  $Truth(q_1) \Rightarrow Truth(q_3)$ .

**Incorporating Constraints.** Our fine-tuning objective is to minimize the violation from the set of logical consistency constraints given for each training example. However, the logical expressions are not differentiable. To form a differentiable loss function, we use a soft logical surrogate of the constraints in the loss. There are various types of conversions for mapping logical constraints into differentiable soft logic, including t-norm Product, t-norm Łukasiewicz, and t-norm Gödel (Li et al., 2019). We use the t-norm Product to convert, which means,  $\neg A$  becomes  $1 - a$ ,  $A \wedge B$  becomes  $ab$ ,  $A \vee B$  becomes  $a + b - ab$ , and  $A \Rightarrow B$  becomes  $\min(1, \frac{b}{a})$ , where  $A$  and  $B$  are concepts with probabilities  $a$  and  $b$  respectively. During the training, the probability of the truth of the facts ( $Truth(q_i)$ ) involved in the constraints is obtained

from the model prediction for computing the loss. We build on previous studies of optimizing an objective that combines the task performance and soft constraints (Li et al., 2019; Asai and Hajishirzi, 2020; Bamas et al., 2020) and form the following objective,

$$\nabla_w \mathcal{L}(w) = \nabla_w L(w) + \sum_{k=1}^K \lambda_k \nabla_w h_k(w), \quad (1)$$

where  $L(w)$  is the task-performance loss function, i.e., Cross-Entropy loss,  $K$  is the number of logical constraints related to the current *Q-Chain*,  $h_k$  measures the violation from a given logical constraint in soft logic differentiable form, and  $\lambda_k$  is a learning parameter indicating the importance of obeying  $k^{th}$  rule. We set the  $\lambda_k$  to 1 for all constraints to ensure equal importance.

In fact, in this work, we utilize the implementation in the DomiKnowS framework (Rajaby Faghihi et al., 2021) to integrate the logical constraints. This framework provides a declarative language for incorporating symbolic knowledge as logical constraints. We present the original logical forms in the program, and DomiKnowS automatically converts them into a differentiable form for inclusion in the loss function. We employ an optimization technique referred to as the Primal-Dual (PD) program that uses a dual formulation of the original objective (Bamas et al., 2020) for fine-tuning. Programming for building a model in DomiKnowS requires defining a graph of concepts and relationships, along with logical rules between concepts. The program for our problem setting is provided in Appendix E.

### 3.4 In-context Learning with Spatial Logic

We design a set of experiments to analyze and compare the LLMs’ performance against our proposed fine-tuning method. We utilize spatial logical reasoning to manually create the rationale of few-shot examples for our in-context learning settings, including Chain-of-Thoughts (Wei et al., 2022), Logical Representation, and Chain-of-Symbols (Hu et al., 2024).

**Chain-of-Thought (CoT).** To enable CoT, we provide instructions that require an explanation before answering the question. For each example response, we include a manually crafted rationale based on spatial logical rules. We provide these examples with instructions and target questions as input to the model to generate the answer. Multiple

examples of this CoT setting are demonstrated in Appendix C.2.

**Logical Representation (LR).** In this setting, we follow the same prompt structure as CoT but use logical representations instead of natural language. We obtain the first order logical representations (predicate-argument structures) by manually converting the CoT explanations. The predicate-argument structure  $R(obj_i, obj_j)$  represents the relation  $R$  between  $obj_i$  and  $obj_j$ , where  $obj_i$  and  $obj_j$  denote object  $i$  and object  $j$ , respectively. An example of the LR counterpart of the CoT phrase is in Table 1.

**Chain of Symbol (CoS).** This symbolic representation was introduced by Hu et al. (2024). We explicitly follow the proposed process in Hu et al. (2024) for converting spatial expressions from natural language into symbolic representation. An example of the CoS counterpart of CoT expression is shown in Table 1.

Domain	Spatial context
CoT	large red square is to the left of a small green square
LR	Left(large red square, small green square)
CoS	(large, red, square) < (small, green, square)

Table 1: An example of a spatial relation used in Chain of Thought (CoT), Logical Representation (LR), and Chain of Symbol (CoS).

**Step-by-Step Prompting.** From the SpartaQA-Human dataset, we observed that scene descriptions are often complex and lengthy. To deal with this issue, we introduce an additional prompting step using LLMs to break down the context before answering the question. Each generated line of context contains a simple spatial description. For example, given the original context, “*There exists a big red square and a medium red square in a block called A. The big one is touching the right edge of the block.*” the expected context generated by the LLM would be “*The big red square in block A.*  
<newline> *The big red square is touching the right edge of block A...*”. More examples are provided in Appendix C. We refer to this prompting technique as *Step-by-Step*, which can be applied before any ICL method.

## 4 Experiments

### 4.1 Datasets

Our experiments are conducted on all existing datasets, including SpartaTUN, SpartaQA-Human, ReSQ, and StepGame.

**SpartaTUN** (Mirzaee and Kordjamshidi, 2022) is a synthetic SQA dataset. This dataset contains both YN and FR types of questions that require multi-hop reasoning to answer. It covers a wide range of spatial relations.<sup>2</sup> The answer for YN is  $A = \{Yes, No\}$ , while the answer for the FR is a subset of 15 spatial relations, e.g.,  $\{left, behind\}$ , and  $\{above, inside\}$ . The dataset provides the chain of reasoning annotations that we use to create the  $Q$ -Chains for our proposed training method. This dataset is used only for training because it is synthetic and the baselines performance reported in (Mirzaee and Kordjamshidi, 2022) on its test split is already high.

**SpartQA-Human (Ver.1)** (Mirzaee et al., 2021) is a small human-annotated SQA dataset. The dataset contains both YN and FR types of questions. We only use the YN question types for fine-tuning and evaluation. The YN questions require more reasoning steps and are shown to be more challenging based on previous research results in Mirzaee and Kordjamshidi (2022).

**SpartQA-Human (Ver.2)** We follow the methodologies from (Mirzaee et al., 2021) to increase the size of the human annotated dataset. The objective is to enrich the SpartaQA-Human with a more extensive test using diverse patterns and questions. Annotators included authors and an undergraduate student who was paid as a research assistant. We extended both types of YN and FR questions.

**ReSQ** (Kordjamshidi et al., 2017; Mirzaee and Kordjamshidi, 2022) is a small realistic SQA dataset with multi-hop reasoning. The depth of reasoning is smaller than other synthetic datasets, but it usually requires commonsense knowledge to infer the answer. The annotations include the depth of reasoning denoted as  $k$ , which can be 1 or 2. The questions that required complex reasoning and commonsense knowledge are annotated as *unclassified* depth. 12.30% of the test examples fall under the  $k=1$  split, 23.93% fall under the  $k=2$  split, and the rest are unclassified. This dataset only has YN questions. We use this dataset to evaluate the performance of our models in realistic domains.

**StepGame** (Shi et al., 2022) is a synthetic SQA dataset that contains a diverse set of multi-hop reasoning questions involving spatial relations, with annotated reasoning depth denoted as  $k$ . The questions in this dataset include up to 10 hops

<sup>2</sup>The dataset also includes Don’t Know questions, but for simplicity, we apply the closed-world assumption and label those questions as No.

of reasoning, allowing for a comprehensive evaluation of models’ reasoning capability at varying depths. This dataset contains only FR questions. The domain of answer is  $a = \{left, right, above, below, overlap, lower-left, lower-right, upper-left, upper-right\}$ .

## 4.2 Experimental Models for Tuning

**BERT-family.** We use *bert-based-uncased* as the initial checkpoint, referred to as BERT. This model is only fine-tuned with the target dataset in each experiment. In some of the initial experiments we used RoBERTa as well. However, due to insignificance of the differences, we continued other model variations with BERT. The results of RoBERTa can be found in Appendix B.

**BERT-T.** We fine-tune the model with SpaRTUN first. Then, we further fine-tune with the target dataset of each experiment.

**BERT-T + Q-Chain.** We adhere to the methodology explained in Section 3.3 to integrate logical knowledge into the BERT-T model via fine-tuning with constraints using the DomiKnowS framework (Rajaby Faghihi et al., 2021).

**Flan-T5.** We follow the methodology described in Section 3.2.2. We selected the *flan-t5-base* in this experiment as the starting checkpoint. To reduce computational costs by decreasing the number of training parameters, we employ the LoRA adapter (Hu et al., 2022).

**Flan-T5-T.** We utilize the same setting as the Flan-T5 model. We fine-tune the model first with SpaRTUN and then with the target dataset.

**Flan-T5-T + Q-Chain.** We integrate the logical knowledge into Flan-T5-T model using proposed method in section 3.3. The backbone is the same as the baseline model without any changes.

**Experimental Setup.** We used the same hyperparameters across all experiments. The learning rate was selected from  $\{1e^{-5}, 8e^{-6}, 1e^{-6}\}$ . For SpartaQA-Human and ReSQ, the number of epochs was 100, while for StepGame, it was 30. For fine-tuning with SpaRTUN, we used 12 epochs for BERT and 8 epochs for Flan-T5. The final hyperparameters were chosen using the development set of the target dataset. In all experiments, we used Cross-Entropy Loss as the loss function and Adam as the optimizer, with the *amsgrad* parameter set to *True*. The experiments were run on 8 A6000 GPUs, requiring approximately 100 GPU hours.

### 4.2.1 Experimental Models for In-context

We evaluate the performance of three selected LLMs for in-context learning setting including GPT3.5, GPT-4, and Llama3-8B, using ReSQ and SpartaQA-Human to compare with our fine-tuning approach. The selected examples and prompts for each in-context learning method can be found in the Appendix C.

**Zero-shot.** We directly ask LLMs to answer the question given the corresponding context.

**Few-shot.** We provide LLMs with four randomly selected examples from the training data.

**CoT.** We augment the *few-shot* setting with reasoning explanations along with the *CoT* examples.

**LR.** We provide the few-shot examples with the reasoning explanation in the logical format described in Section 3.4.

**CoS.** We further alter the *few-shot* examples with the CoS symbolic forms explained in Section 3.4.

## 4.3 Results and Discussion

**Realistic Domain.** ReSQ serves as a realistic SQA domain. Table 2 shows that using the *Q-chain* is effective for both BERT and Flan-T5, with notable improvement on Flan-T5. To better understand these results, we analyzed the performance across different splits of ReSQ based on the manually annotated depth of reasoning required to answer the questions. The first two splits include questions requiring one or two hops of reasoning, denoted as  $k = 1$  and  $k = 2$ . The *unclassified* split covers questions where reasoning depth is hard to determine and rely more on commonsense knowledge. An example of each case is provided in Figure 3.

Our observation from Table 2 reveals that our model consistently improves performance on  $k = 2$  but negatively impacts BERT on  $k = 1$  and the *unclassified* categories. This suggests that our fine-tuning approach is more effective when deeper reasoning is needed but is less effective for questions requiring commonsense knowledge. On the other hand, LLMs outperform fine-tuning approaches on ReSQ, specifically in the *unclassified* subset. They consistently show 2-13% higher performance than Flan-T5+T+*Q-chain* (lines 7-13). These results indicate that one of LLMs’ superior performances is due to their commonsense knowledge while our approach performs better in multi-hop reasoning. Nevertheless, we observe that logic-based fine-tuning enhances Flan-T5 more significantly than BERT on the *unclassified* subset. This

Model	SpartQA-Human		ReSQ				#
	Ver.1	Ver.2	k=1	k=2	unclassified	All	
BERT	54.54	53.57	70.67	56.85	60.66	60.98	1
BERT-T	55.94	58.03	<b>76.00</b>	54.79	<b>61.18</b>	61.15	2
BERT-T+ <i>Q-Chain</i> (Ours)	<b>59.44</b>	<b>58.92</b>	72.00	<b>58.90</b>	59.90	<b>61.31</b>	3
Flan-T5	54.54	60.71	74.67	56.16	61.44	61.80	4
Flan-T5-T	49.65	57.14	81.33	54.79	61.44	62.30	5
Flan-T5-T+ <i>Q-Chain</i> (Ours)	<b>55.94</b>	<b>61.61</b>	<b>81.33</b>	<b>57.53</b>	<b>63.75</b>	<b>64.43</b>	6
GPT3.5 (zero-shot)	58.04	58.03	74.67	60.95	66.58	66.22	7
GPT3.5 (few-shot)	62.23	58.92	84.00	68.49	68.12	70.16	8
GPT3.5 (LR)	64.33	59.82	-	-	-	-	9
GPT3.5 (CoS)	60.14	64.28	-	-	-	-	10
GPT3.5 (CoT)	65.73	71.43	<b>86.67</b>	67.12	68.64	70.49	11
GPT-4 (zero-shot)	<b>77.62</b>	68.75	84.00	73.97	<b>76.86</b>	<b>77.05</b>	12
Llama-3 (zero-shot)	61.54	50.89	80.00	64.38	67.35	68.20	13
Llama-3 (few-shot)	62.94	60.71	82.67	69.86	71.46	72.46	14
Llama-3 (LR)	60.14	<b>72.32</b>	-	-	-	-	15
Llama-3 (CoS)	67.13	61.60	-	-	-	-	16
Llama-3 (CoT)	67.83	70.54	82.76	<b>76.03</b>	67.10	71.15	17
PistaQ*	75.52	-	-	-	-	47.21	18

Table 2: Accuracy of SpartQA-Human and ReSQ with various models. For ReSQ,  $k$  is the number of the reasoning steps required to answer the questions. *Unclassified* indicates the cases in which  $k$  was challenging for human annotators to determine. The results from PistaQ\* reported from [Mirzaee and Kordjamshidi \(2023\)](#).

Model	k=1	k=2	k=3	k=4	k=5	k=6	k=7	k=8	k=9	k=10
BERT	98.51	95.53	91.68	66.71	49.11	41.47	41.47	32.09	28.94	28.16
BERT-T	98.50	95.32	<b>93.26</b>	<b>76.78</b>	<b>66.36</b>	58.76	53.70	46.27	42.71	40.12
BERT-T+ <i>Q-Chain</i> (Ours)	<b>98.70</b>	<b>96.45</b>	93.03	74.58	64.95	<b>59.04</b>	<b>54.38</b>	<b>49.23</b>	<b>45.36</b>	<b>44.05</b>
Llama-3 (few-shot)	38.01	27.87	24.15	21.27	19.75	18.03	16.88	15.52	15.17	14.70
GPT3 (few-shot)	55.00	37.00	25.00	30.00	32.00	29.00	21.00	22.00	34.00	31.00
GPT3 (CoT)	61.00	45.00	30.00	35.00	35.00	27.00	22.00	24.00	23.00	25.00
GPT3(d2) + ASP	92.60	89.90	89.10	93.80	92.90	91.60	91.20	90.40	89.00	88.30

Table 3: Accuracy of StepGame on several models including results of GPT3 reported from [Yang et al. 2023](#).

Model	Raw Context	Step by step
GPT3.5 (zero-shot)	58.04	63.64
GPT3.5 (few-shot)	62.23	64.33
GPT3.5 (CoT)	65.73	67.83
GPT3.5 (LR)	64.33	59.44
GPT3.5 (CoS)	60.14	58.74
GPT-4 (zero-shot)	<b>77.62</b>	<b>78.32</b>

Table 4: The accuracy of LLMs on SpartQA-Human.

suggests that the *Q-chain* can facilitate complex reasoning, particularly when applied to a model with stronger commonsense capabilities. Note that applying LR and CoS is not straightforward on ReSQ due to the challenges associated with extracting logical representations from realistic domains text. This challenge is evident in the results reported for PistaQ\*, which employs a similar approach. PistaQ\* achieved only 47.21% accuracy on ReSQ, despite significant improvements observed in synthetic domains such as SpartQA-Human.

**Synthetic Domain with More Complex Logical Reasoning.** SpartQA-Human and StepGame are synthetic domains used in our experiments. According to Table 2, we consistently observe improvements with our proposed *Q-chain* in the synthetic domain which typically requires considerable reasoning hops. Moreover, the gap between small

<b>ReSQ: k = 1</b> <b>Context:</b> A brown house with large windows and flags above the entrance, and the white writing Hotel Peregrino <b>Question:</b> Are the flags above the entrance? <b>Label:</b> Yes <b>Llama3 (few-shot):</b> Yes (✓) <b>Flan-T5-T+Q-Chain (Ours):</b> Yes (✓)	
<b>ReSQ: k = 2</b> <b>Context:</b> Room with two beds on the right, a table with two chairs on the left and a large window in the back, there is also a mirror on the left. <b>Question:</b> Are the chairs to the left of the beds? <b>Label:</b> Yes <b>Llama3 (few-shot):</b> No (✗) <b>Flan-T5-T+Q-Chain (Ours):</b> Yes (✓)	
<b>ReSQ: k = unclassified</b> <b>Context:</b> A room with grey walls, two single beds with grey, green and yellow striped bedcover and a night table with a bedside lamp in between. <b>Question:</b> What is the relation of the agent S to the agent C? <b>Label:</b> Yes <b>Llama3 (few-shot):</b> Yes (✓) <b>Flan-T5-T+Q-Chain (Ours):</b> No (✗)	

Figure 3: Three examples of each category in the ReSQ dataset with responses from Llama3 (Few-shot) and Flan-T5 + Q-Chain (Ours).

PLMs and LLMs is much smaller in this domain compared to the realistic domain (ReSQ). This is expected, as we concluded in the results of the realistic domain in Table 2, that LLMs are better at commonsense reasoning than complex multi-hop spatial reasoning. The hypothesis is further supported when assessing the proposed method on StepGame. As can be observed in Table 3, the fine-tuning method consistently demonstrates significant positive differences in all reasoning steps compared to LLMs. The struggle of GPT3 on reasoning over StepGame is also investigated in ([Yang et al.](#),



<b>StepGame: k = 1</b> <b>Context:</b> T is below C. H and M are vertical and H is above M. J is below and to the left of H. <b>Question:</b> What is the relation of the agent M to the agent H? <b>Label:</b> Below <b>Llama3 (few-shot):</b> Below(✓) <b>BERT-T+Q-Chain (Our):</b> Below(✓)	<b>StepGame: k = 10</b> <b>Context:</b> K is at the bottom of T. V is to the right of G horizontally. W and L are next to each other with W on the left and L on the right. Z is positioned below and to the right of X. T is positioned in the lower right corner of O. X and N are parallel, and X is under N. G is diagonally to the bottom left of M. V is below and to the left of O. N is on the lower right of M. T is diagonally to the bottom left of I. H and I are parallel, and H on the left of I. L is at the bottom of I. K is on the top of Z and is on the same vertical plane. U and B are vertical and U is below B. <b>Question:</b> What is the relation of the agent N to the agent X? <b>Label:</b> Above <b>Llama3 (few-shot):</b> Below(✗) <b>BERT-T+Q-Chain (Our):</b> Above(✓)
<b>StepGame: k = 4</b> <b>Context:</b> R is at the bottom and A is on the top. R is below E at 4 o'clock. W is positioned right to M. M and R are parallel, and M is under R. S is on the left side and above W. Z is above J. <b>Question:</b> What is the relation of the agent M to the agent W? <b>Label:</b> Left <b>Llama3 (few-shot):</b> Right(✗) <b>BERT-T+Q-Chain (Our):</b> Left(✓)	

Figure 4: Three examples of each category in the StepGame dataset with responses from Llama3 (Few-shot) and BERT-T + Q-Chain (Our).

2023). Our proposed method consistently improves by 1-4% on a higher number of reasoning hops ( $k=6$  to  $k=10$ ), similar to the observation results of ReSQ. These results confirm our primary hypothesis that our proposed method equips the models with a higher level of logical abstraction to conduct more complex logical reasoning. We provide three examples of responses from this dataset to qualitatively demonstrate the improvements in Figure 4. We also included results from GPT-3(d2) + ASP and PistaQ\*, both of which demonstrate significant effectiveness in multi-hop reasoning. These models perform relation extraction to obtain a formal logical representation using LLMs, while reasoning is handled separately using external tools with engineered rules of reasoning. However, these pipelines face two key limitations. First, they require formal representations at test time, effective only on synthetic data. Second, they do not fine-tune the language models for better reasoning, relying instead on external tools. This approach contrasts with ours, which focuses on directly improving the language models’ reasoning capabilities.

**In-context Learning.** This section focuses only on comparing the variations of in-context learning for spatial reasoning. As shown in Table 2 and Table 4, CoT proved to be the most effective approach for spatial reasoning in both GPT and Llama. We also evaluated the impact of integrating formal representations in in-context learning, referred to as the LR approach. As evident in Table 4 (line 4), using formal representations marginally degraded performance. In most cases, utilizing natural language in the CoT proved to be more effective. Additionally, we evaluate CoS symbolic representations, finding that CoT and LR outperformed CoS using the same examples. This highlights the advantage of natural

language reasoning over symbolic reasoning for in-context learning. We further applied step-by-step prompting explained in Section 3.4 to evaluate more sophisticated prompting techniques on LLMs. Table 4 shows that this technique achieves 2-5% better results than zero-shot, few-shot, and CoT using raw context. The results demonstrate that more straightforward and shorter contexts improve LLMs’ spatial understanding. This is consistent with previous results on other domains and tasks (Li et al., 2023; Liu et al., 2024). Lastly, we compare open-source Llama and closed-source GPT models. Our results indicate that Llama-3 is competitive with GPT-3.5, showing only minor differences in the YN domain as shown in Table 2. Thus, we plan to prioritize open-source models in future research.

## 5 Conclusion

Given the importance of spatial reasoning in many real-world applications, we focus on improving this reasoning skill in language models. We equip LLMs with neuro-symbolic supervision to leverage logical knowledge during fine-tuning. This practice boosts their capability of capturing the necessary level of abstraction for spatial reasoning and enhances the generalizability for transferring knowledge across domains. Our constraint-based training technique achieves higher accuracy than other competitive Spatial Question-answering baselines across various benchmarks. Furthermore, the results indicate that our model performs better in scenarios requiring more reasoning steps. Lastly, we compare our models with state-of-the-art LLMs. Experimental comparisons show that while large LLMs like GPT3.5 perform better in commonsense reasoning, we achieve higher performance in multi-hop spatial question-answering with tuning comparatively smaller language models like BERT with spatial logic.

## Limitations

While we improve the reasoning capability of the models, our approach by no means solves the general reasoning problem. The trustworthiness and reliability of the LLMs are still a research challenge. Our models will need GPU resources to run, which might be expensive. Our research is based on open-source data and code, and the results will be reproducible. We do not see any ethical concerns in our research approach or the impact of our

results. Our data, being limited to the spatial reasoning domain, does not include any specific type of bias that can harm minority people belonging to any particular gender or race.

## Acknowledgment

This project is partially supported by the Office of Naval Research (ONR) grant N00014-23-1-2417. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of Office of Naval Research. We thank anonymous reviewers for their constructive feedback, which greatly helped us improve this manuscript.

## References

- AI@Meta. 2024. [Llama 3 model card](#).
- Saeed Amizadeh, Hamid Palangi, Alex Polozov, Yichen Huang, and Kazuhito Koishida. 2020. [Neuro-symbolic visual reasoning: Disentangling "Visual" from "Reasoning"](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 279–290. PMLR.
- Akari Asai and Hannaneh Hajishirzi. 2020. [Logic-guided data augmentation and regularization for consistent question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5642–5650, Online. Association for Computational Linguistics.
- Jamal Atif, Céline Hudelot, Geoffroy Fouquier, Isabelle Bloch, and Elsa D Angelini. 2007. From generic knowledge to specific reasoning for medical image interpretation using graph based representations. In *IJCAI*, pages 224–229.
- Etienne Bamas, Andreas Maggiori, and Ola Svensson. 2020. [The primal-dual method for learning augmented algorithms](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 20083–20094. Curran Associates, Inc.
- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet V. Do, Yan Xu, and Pascale Fung. 2023. [A multitask, multilingual, multimodal evaluation of ChatGPT on reasoning, hallucination, and interactivity](#). In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–718, Nusa Dua, Bali. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Boyuan Chen, Zhuo Xu, Sean Kirmani, Brain Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. 2024. SpatialVLM: Endowing vision-language models with spatial reasoning capabilities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14455–14465.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2024. [Scaling instruction-finetuned language models](#). *Journal of Machine Learning Research*, 25(70):1–53.
- Surabhi Datta, Yuqi Si, Laritza Rodriguez, Sonya E Shooshan, Dina Demner-Fushman, and Kirk Roberts. 2020. [Understanding spatial language in radiology: Representation framework, annotation, and spatial relation extraction from chest x-ray reports using deep learning](#). *Journal of Biomedical Informatics*, 108:103473.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Chun Feng, Joy Hsu, Weiyu Liu, and Jiajun Wu. 2024. Naturally supervised 3d visual grounding with language-regularized concept learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13269–13278.
- Gaël Gendron, Qiming Bao, Michael Witbrock, and Gillian Dobbie. 2024. [Large language models are not strong abstract reasoners](#). In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, pages 6270–6278. International Joint Conferences on Artificial Intelligence Organization. Main Track.

- Shizhan Gong, Yuan Zhong, Wenao Ma, Jinpeng Li, Zhao Wang, Jingyang Zhang, Pheng-Ann Heng, and Qi Dou. 2024. [3dsam-adapter: Holistic adaptation of sam from 2d to 3d for promptable tumor segmentation](#). *Medical Image Analysis*, 98:103324.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Hanxu Hu, Hongyuan Lu, Huajian Zhang, Yun-Ze Song, Wai Lam, and Yue Zhang. 2024. [Chain-of-symbol prompting for spatial reasoning in large language models](#). In *First Conference on Language Modeling*.
- Danial Kamali, Elham J. Barezi, and Parisa Kordjamshidi. 2025. [Nesycoco: A neuro-symbolic concept composer for compositional generalization](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Parisa Kordjamshidi, Taher Rahgooy, Marie-Francine Moens, James Pustejovsky, Umar Manzoor, and Kirk Roberts. 2017. Clef 2017: Multimodal spatial role labeling (msprl) task overview. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, pages 367–376, Cham. Springer International Publishing.
- Fangjun Li, David C. Hogg, and Anthony G. Cohn. 2024. [Advancing spatial reasoning in large language models: An in-depth evaluation and enhancement using the stepgame benchmark](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(17):18500–18507.
- Tao Li, Vivek Gupta, Maitrey Mehta, and Vivek Srikrumar. 2019. [A logic-driven framework for consistency of neural models](#). 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)*, pages 3924–3935, Hong Kong, China. Association for Computational Linguistics.
- Yucheng Li, Bo Dong, Frank Guerin, and Chenghua Lin. 2023. [Compressing context to enhance inference efficiency of large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6342–6353, Singapore. Association for Computational Linguistics.
- Fangyu Liu, Guy Emerson, and Nigel Collier. 2023. Visual spatial reasoning. *Transactions of the Association for Computational Linguistics*, 11:635–651.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. [Lost in the middle: How language models use long contexts](#). *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Ximing Lu, Sean Welleck, Peter West, Liwei Jiang, Jungo Kasai, Daniel Khashabi, Ronan Le Bras, Lianhui Qin, Youngjae Yu, Rowan Zellers, Noah A. Smith, and Yejin Choi. 2022. [NeuroLogic a\\*esque decoding: Constrained text generation with lookahead heuristics](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 780–799, Seattle, United States. Association for Computational Linguistics.
- Roshanak Mirzaee and Parisa Kordjamshidi. 2022. [Transfer learning with synthetic corpora for spatial role labeling and reasoning](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6148–6165, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Roshanak Mirzaee and Parisa Kordjamshidi. 2023. [Disentangling extraction and reasoning in multi-hop spatial reasoning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3379–3397, Singapore. Association for Computational Linguistics.
- Roshanak Mirzaee, Hossein Rajaby Faghihi, Qiang Ning, and Parisa Kordjamshidi. 2021. [SPARTQA: A textual question answering benchmark for spatial reasoning](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4582–4598, Online. Association for Computational Linguistics.
- Raymond S. Nickerson. 1994. [Chapter 13 - the teaching of thinking and problem solving](#). In Robert J. Sternberg, editor, *Thinking and Problem Solving*, volume 2 of *Handbook of Perception and Cognition*, pages 409–449. Academic Press, San Diego.
- OpenAI. 2023. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- John Prager, Jennifer Chu-Carroll, and Krzysztof Czuba. 2004. [Question answering using constraint satisfaction: QA-by-dossier-with-constraints](#). In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL-04)*, pages 574–581, Barcelona, Spain.
- Lianhui Qin, Sean Welleck, Daniel Khashabi, and Yejin Choi. 2022. [Cold decoding: Energy-based constrained text generation with langevin dynamics](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 9538–9551. Curran Associates, Inc.
- Hossein Rajaby Faghihi, Quan Guo, Andrzej Uszok, Aliakbar Nafar, and Parisa Kordjamshidi. 2021. [DomiKnowS: A library for integration of symbolic domain knowledge in deep learning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 231–241, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.



- Hossein Rajaby Faghihi, Aliakbar Nafar, Chen Zheng, Roshanak Mirzaee, Yue Zhang, Andrzej Uszok, Alexander Wan, Tanawan Premisri, Dan Roth, and Parisa Kordjamshidi. 2023. [Gluecons: A generic benchmark for learning under constraints](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(8):9552–9561.
- Manasi Sharma. 2023. [Exploring and improving the spatial reasoning abilities of large language models](#).
- Zhengxiang Shi, Qiang Zhang, and Aldo Lipani. 2022. [Stepgame: A new benchmark for robust multi-hop spatial reasoning in texts](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11321–11329.
- Emrah Akin Sisbot, Luis F Marin, and Rachid Alami. 2007. Spatial reasoning for human robot interaction. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2281–2287. IEEE.
- Haitian Sun, William W. Cohen, and Ruslan Salakhutdinov. 2022. [Reasoning over logically interacted conditions for question answering](#). *Preprint*, arXiv:2205.12898.
- Qingyu Tan, Hwee Tou Ng, and Lidong Bing. 2023. [Towards benchmarking and improving the temporal reasoning capability of large language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14820–14835, Toronto, Canada. Association for Computational Linguistics.
- Karthik Valmeekam, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. 2022. [Large language models still can’t plan \(a benchmark for LLMs on planning and reasoning about change\)](#). In *NeurIPS 2022 Foundation Models for Decision Making Workshop*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.
- Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M. Rush, Bart van Merriënboer, Armand Joulin, and Tomas Mikolov. 2016. [Towards ai-complete question answering: A set of prerequisite toy tasks](#). In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Elmira Yadollahi, Miguel Alexandre Monteiro, and Ana Paiva. 2023. [Learning spatial reasoning in virtual vs. physical games with robots](#). In *Proceedings of the 11th International Conference on Human-Agent Interaction*, HAI ’23, page 162–170, New York, NY, USA. Association for Computing Machinery.
- Zhun Yang, Adam Ishay, and Joohyung Lee. 2023. [Coupling large language models with logic programming for robust and general reasoning from text](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 5186–5219, Toronto, Canada. Association for Computational Linguistics.
- Yue Zhang, Quan Guo, and Parisa Kordjamshidi. 2021. [Towards navigation by reasoning over spatial configurations](#). In *Proceedings of Second International Combined Workshop on Spatial Language Understanding and Grounded Communication for Robotics*, pages 42–52, Online. Association for Computational Linguistics.
- Yue Zhang and Parisa Kordjamshidi. 2022. [LOViS: Learning orientation and visual signals for vision and language navigation](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5745–5754, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Ming Zhou, Nan Duan, Shujie Liu, and Heung-Yeung Shum. 2020. [Progress in neural nlp: Modeling, learning, and reasoning](#). *Engineering*, 6(3):275–290.

## A Data Statistic

The data statistic of each dataset is shown in Table 5.

Dataset	Train	Test	Dev
SpaRTUN (YN)	20633	3232	3152
SpaRTUN (FR)	18400	2818	2830
Q-Chain (YN)	56987	-	-
Q-Chain (FR)	46750	-	-
SpartQA-Human(Ver.1)	161	143	51
SpartQA-Human(Ver.2)	200	112	60
ReSQ	1008	610	333
StepGame	50000	5000	100000

Table 5: Size of each SQA benchmark used in experiments

## B RoBERTa results

We provide results from the same setting as the BERT using RoBERTa in Table 6 and found that the overall results are slightly lower than BERT for typical transfer learning. So, we decided to use only BERT for the remaining experiments.

## C In-context Learning Examples

The example of each in-context learning prompt and example is indicated below.

### C.1 Few-Shot

**system** : You will be given a story and question. Then, you will answer either only Yes or No based on given story. Candidate answer: [Yes, No]



Model	SpartQA-Human		ReSQ				#
	Ver.1	Ver.2	$k=1$	$k=2$	unclassified	All	
BERT	54.54	53.57	70.67	56.85	60.66	60.98	1
RoBERTa	54.54	-	76.00	56.16	58.87	60.33	2
BERT-T	55.94	58.03	76.00	54.79	61.18	61.15	2
RoBERTa-T	49.65	-	64.67	57.87	55.78	56.72	4
BERT-T+ <i>Q-Chain</i> (Our)	59.44	58.92	72.00	58.90	59.90	61.31	3

Table 6: Accuracy of SpartQA-Human and ReSQ with BERT and RoBERTa for comparison between these models.

**user** : There exist a big red square, a big red triangle, a medium red square, and a small green circle in a block called A. The triangle and medium square are touching the bottom edge of the block. The big and medium square are touching the right edge of the block. And the circle is above the big square which is above the medium square. There is another block called B to the left of block A. A medium green square is touching block B's left edge and below a medium red square. The medium red square is above and to the left of a small green square. Also a medium red triangle is below and to the left of the small square. Is the medium red triangle below the red square in block B?

**assistant** : Yes

**user** : There are three blocks A, B and C. Block A has a medium white triangle and a small red rectangle. There exists a medium white rectangle to the left of the triangle and to the right of the red rectangle. Above block A there is block B. It has a small white oval touching the bottom edge of the block and a medium white triangle to the right of it. Below block A there is block C. It has a small green triangle touching the top edge of the block. A medium white triangle is below a small green rectangle. Another small green triangle is below and to the right of the white triangle. Are all rectangles above a white triangle?

**assistant** : Yes

**user** : There is a block called A. Inside A, a black triangle above and to the left of a black circle. There is another block called B which is to the right of A. B has two yellow circles, a blue circle, and a blue square. The blue circle is touching the left edge of B. One of the yellow circles is to the right of the blue circle and to the left of the blue square. The blue square is touching the right edge of B and above the other yellow circle, Which also touches the right edge of B. Does block B have all of the circles inside it?

**assistant** : No

**user** : There is a block called A. A has a large blue square near and below a medium blue triangle.

Far to the left of the medium blue triangle is a small yellow square. Far above the medium blue triangle is a small blue triangle. There is another block called B, which is to the right of A. In B, There is a small yellow circle touching the left edge of B. Far to the right of the circle is a large yellow square. Near and above the large square is a small yellow square. There is another block called C, which is to the right of B. In C, there is a small blue circle near and to the left of a large yellow triangle. The large yellow triangle is near and to the left of a large blue triangle. The large blue triangle is far above a large blue circle. The large blue circle is touching the bottom edge of C. Are the large blue things in C to the left of the large blue thing in A?

**assistant** : No

**user** : There is a block called A. In A there is a red triangle Which is above a black circle. The black circle is touching the bottom of A. A yellow square is touching the top edge of A, and below that, a red oval is touching the right edge of A. There is a second block called B, to the left of A. Inside B are 3 squares: a brown, a green, and a red. The green square is above the red square and touching the left edge of B. The brown square is touching both the bottom edge and the right edge of B. The red square is touching the left edge of the brown square. Is the oval to the right of the brown thing?

## C.2 Chain of Thought

**system** : You will be given story and question. Then, you will answer either Yes or No with explanation based on given story step by step. Candidate answer: [Yes, No]

**user** : There exist a big red square, a big red triangle, a medium red square, and a small green circle in a block called A. The triangle and medium square are touching the bottom edge of the block. The big and medium square are touching the right edge of the block. And the circle is above the big square which is above the medium square. There is another block called B to the left of block A.

A medium green square is touching block B's left edge and below a medium red square. The medium red square is above and to the left of a small green square. Also a medium red triangle is below and to the left of the small square. Is the medium red triangle below the red square in block B?

**assistant** : From information provided: The medium red square is above a small green square in B. Therefore, the small green square is below the medium red square in B. The medium red triangle is below the small square in B, and the small green square is below the medium red square in B. Therefore, the medium red triangle is below the medium red square in B. Answer: Yes

**user** : There are three blocks A, B and C. Block A has a medium white triangle and a small red rectangle. There exists a medium white rectangle to the left of the triangle and to the right of the red rectangle. Above block A there is block B. It has a small white oval touching the bottom edge of the block and a medium white triangle to the right of it. Below block A there is block C. It has a small green triangle touching the top edge of the block. A medium white triangle is below a small green rectangle. Another small green triangle is below and to the right of the white triangle. Are all rectangles above a white triangle?

**assistant** : From given information: A medium white triangle is below a small green rectangle in C. Then, small green rectangle is above medium white triangle in C. A small red rectangle in block A, a medium white triangle in C, and block C is below block A. Then, a medium white triangle is below a small red rectangle. A medium white triangle is below a small red rectangle. Therefore, a small red rectangle is above a medium white triangle. A medium white rectangle is in A, a medium white triangle in C, and block C is below block A. Then, a medium white triangle is below a medium white rectangle. A medium white triangle is below a medium white rectangle. Therefore, a medium white rectangle is above a medium white triangle. A medium white rectangle is above a medium white triangle. Thus, all rectangles from given story are above a white triangle. Answer: Yes

**user** : There is a block called A. Inside A, a black triangle above and to the left of a black circle. There is another block called B which is to the right of A. B has two yellow circles, a blue circle, and a blue square. The blue circle is touching the left edge of B. One of the yellow circles is to the right of the blue circle and to the left of the blue square.

The blue square is touching the right edge of B and is above the other yellow circle, which is also touching the right edge of B. Does block B have all of the circles inside it?

**assistant** : From given relation, A black triangle is above a black circle in A. Then, there is a black circle in A. Therefore, not all circles are in B. Answer: No

**user** : There is a block called A. A has a large blue square near and below a medium blue triangle. Far to the left of the medium blue triangle is a small yellow square. Far above the medium blue triangle is a small blue triangle. There is another block called B, which is to the right of A. In B, there is a small yellow circle touching the left edge of B. Far to the right of the circle is a large yellow square. Near and above the large square is a small yellow square. There is another block called C, which is to the right of B. In C, there is a small blue circle near and to the left of a large yellow triangle. The large yellow triangle is near and to the left of a large blue triangle. The large blue triangle is far above a large blue circle. The large blue circle is touching the bottom edge of C. Are the large blue things in C to the left of the large blue thing in A?

**assistant** : Based on the provided content. Block B is to the right of block A. Block C is to the right of block B. Therefore, block C is to the right of block A. In block C, the yellow triangle is near and to the left of large blue triangle. Therefore, there is large blue triangle in C. Block A has a large blue square near and below a medium blue triangle. Therefore, there is large blue square in A. Then, large blue triangle in C is to the right of large blue square in A. So, large blue triangle in C is not to the left of large blue square in A. Answer: No

**user** : There is a block called A. In A there is a red triangle which is above a black circle. The black circle is touching the bottom of A. A yellow square is touching the top edge of A, and below that, a red oval is touching the right edge of A. There is a second block called B which is to the left of A. Inside B there are 3 squares, a brown one, a green one, and a red one. The green square is above the red square and touching the left edge of B. The brown square is touching both the bottom edge and the right edge of B. The red square is touching the left edge of the brown square. Is the oval to the right of the brown thing?

### C.3 Logical Representation (LR)

**system** : You will be given story and question. Then, you will answer either Yes or No with explanation based on given story step by step. Candidate answer: [Yes, No]

**user** : There exist a big red square, a big red triangle, a medium red square, and a small green circle in a block called A. The triangle and medium square are touching the bottom edge of the block. The big and medium square are touching the right edge of the block. And the circle is above the big square which is above the medium square. There is another block called B to the left of block A. A medium green square is touching the left edge of block B and is below a medium red square. The medium red square is above and to the left of a small green square. Also a medium red triangle is below and to the left of the small square. Is the medium red triangle below the red square in block B?

**assistant** : Below(Medium red triangle in B, red square in block B) Answer: Yes

**user** : There are three blocks A, B and C. Block A has a medium white triangle and a small red rectangle. There exists a medium white rectangle to the left of the triangle and to the right of the red rectangle. Above block A there is block B. It has a small white oval touching the bottom edge of the block and a medium white triangle to the right of it. Below block A there is block C. Its has a small green triangle touching the top edge of the block. A medium white triangle is below a small green rectangle. Another small green triangle is below and to the right of the white triangle. Are all rectangles above a white triangle?

**assistant** : Below(C, A) => Above(A, C), Contain(A, Small red rectangle) + Above(A, C) + Contain(C, Medium white triangle) => Above(Small red rectangle, Medium White triangle), Above(A, C) & Contain(A, Medium white rectangle) & Contain(C, Medium white triangle) => Above(Medium white rectangle, Medium White triangle), Below(Medium white triangle in C, small green rectangle in C) => Above(small green rectangle, Medium white triangle) Answer: Yes

**user** : There is a block called A. Inside A, there is a black triangle that is above and to the left of a black circle. There is another block called B which is to the right of A. In B, there are two yellow circles, a blue circle, and a blue square. The blue circle is touching the left edge of B. One of the

yellow circles is to the right of the blue circle and to the left of the blue square. The blue square is touching the right edge of B and is above the other yellow circle, Which is also touching the right edge of B. Does block B have all of the circles inside it?

**assistant** : Contain(A, Black Circle) =>  $\neg$ Contain(B, Black Circle),  $\neg$ Contain(B, Black Circle) => Some x,  $\neg$ Contain(B, x) where x is Circle. Answer: No

**user** : There is a block called A. In A, there is a large blue square near and below a medium blue triangle. Far to the left of the medium blue triangle is a small yellow square. Far above the medium blue triangle is a small blue triangle. There is another block called B, which is to the right of A. In B, There is a small yellow circle touching the left edge of B. Far to the right of the circle is a large yellow square. Near and above the large square is a small yellow square. There is another block called C, which is to the right of B. In C, there is a small blue circle near and to the left of a large yellow triangle. The large yellow triangle is near and to the left of a large blue triangle. The large blue triangle is far above a large blue circle. The large blue circle is touching the bottom edge of C. Are the large blue things in C to the left of the large blue thing in A?

**assistant** : Question: Left(large blue thing in C, large blue thing in A). Related relation from story: Right(B, A), Right(C, B), Contain(C, large blue triangle), Contain(A, large blue square) Steps: Right(C, B) + Right(B, A) => Right(C, A) Contain(C, large blue triangle) + Right(C, A) + Contain(A, large blue square) => Right(large blue triangle in C, large blue square in A) Right(large blue triangle in C, large blue square in A) => !Left(large blue thing in C, large blue thing in A) Answer: No

### C.4 Chain of Symbol

**system** : You will be given story and question. Then, you will answer either Yes or No with explanation based on given story step by step. Candidate answer: [Yes, No]

**user** : There exist a big red square, a big red triangle, a medium red square, and a small green circle in a block called A. The triangle and medium square are touching the bottom edge of the block. The big and medium square are touching the right edge of the block. And the circle is above the big square which is above the medium square. There is another block called B to the left of block A. A medium green square is touching the left edge of

block B and is below a medium red square. The medium red square is above and to the left of a small green square. Also a medium red triangle is below and to the left of the small square. Is the medium red triangle below the red square in block B?

**assistant** : Let analyze the relations given from context to answer question.B: 1. (medium, red, triangle)  $\downarrow$ ,  $<$  (small, red, square)Based on the information, B: 1. (medium, red, triangle)  $\downarrow$ ,  $<$  (small, red, square)2. (medium, red, triangle)  $\downarrow$  (small, red, square)Answer: Yes

**user** : There are three blocks A, B and C. Block A has a medium white triangle and a small red rectangle. There exists a medium white rectangle to the left of the triangle and to the right of the red rectangle. Above block A there is block B. It has a small white oval touching the bottom edge of the block and a medium white triangle to the right of it. Below block A there is block C. Its has a small green triangle touching the top edge of the block. A medium white triangle is below a small green rectangle. Another small green triangle is below and to the right of the white triangle. Are all rectangles above a white triangle?

**assistant** : Let consider the related relations to answer the questions, A: 1. (small, red, rectangle) 2. (medium, white, rectangle) C: 1. (small, green, rectangle) = edge 2. (medium, white, triangle)  $\downarrow$  (small, green, rectangle) C  $\downarrow$  A Based on the information, let determined the relation of all rectangle and white triangles, 1. C  $\downarrow$  A. Therefore, A  $\uparrow$  C 2. A  $\uparrow$  C, A: (small, red, rectangle), and C: (medium, white, triangle). Therefore, A: (small, red, rectangle)  $\uparrow$  C: (medium, white, triangle) 3. A  $\uparrow$  C, A: (medium, white, rectangle), and C: (medium, white, triangle). Therefore, A: (medium, white, rectangle)  $\uparrow$  C: (medium, white, triangle). 4. C: (medium, white, triangle)  $\downarrow$  (small, green, rectangle). Therefore, (small, green, rectangle)  $\uparrow$  (medium, white, triangle).5. A: (small, red, rectangle)  $\uparrow$  C: (medium, white, triangle), A: (medium, white, rectangle)  $\uparrow$  C: (medium, white, triangle), and C: (small, green, rectangle)  $\uparrow$  (medium, white, triangle), therefore, all rectangle above a white triangle. Answer: Yes

**user** : There is a block called A. Inside A, there is a black triangle that is above and to the left of a black circle. There is another block called B which is to the right of A. In B, there are two yellow circles, a blue circle, and a blue square. The blue circle is touching the left edge of B. One of the yellow circles is to the right of the blue circle and

to the left of the blue square. The blue square is touching the right edge of B and is above the other yellow circle, Which is also touching the right edge of B. Does block B have all of the circles inside it?

**assistant** : Let consider the related relations to answer the questions,A: 1. (black, triangle)  $\uparrow <$  (black, circle)B: 1. (yellow, circle1) right (blue, circle)2. (blue, square)  $\uparrow$  (yellow, circle2) 3.(blue, circle) =left edge of B Based on the information, let determined the relation of all rectangle and white triangles, 1. A: 1. (black, triangle)  $\uparrow <$  (black, circle). Therefore, A: (black, circle)2. A: (black, circle). Therefore, B doesn't have all circle. Answer: No

**user** : There is a block called A. In A, there is a large blue square near and below a medium blue triangle. Far to the left of the medium blue triangle is a small yellow square. Far above the medium blue triangle is a small blue triangle. There is another block called B, which is to the right of A. In B, There is a small yellow circle touching the left edge of B. Far to the right of the circle is a large yellow square. Near and above the large square is a small yellow square. There is another block called C, which is to the right of B. In C, there is a small blue circle near and to the left of a large yellow triangle. The large yellow triangle is near and to the left of a large blue triangle. The large blue triangle is far above a large blue circle. The large blue circle is touching the bottom edge of C. Are the large blue things in C to the left of the large blue thing in A?

**assistant** : Let consider the related relations to answer the questions,B  $>$  A,C  $>$  BA: 1. (large, blue square)  $\sim \downarrow$  (medium, blue square)C: 1. (large, yellow, triangle)Based on the information, let determined the relation large blue things in C and the large blue thing in A,1. C  $>$  B and B  $>$  A. Therefore, C  $>$  A2. A: 1. (large, blue, square)  $\sim \downarrow$  (medium, blue square). Therefore, A: (large, blue, square)3. C  $>$  A, A: (large, blue, square), and C: (large, yellow, triangle).Therefore, C: (large, yellow, triangle)  $>$  A: (large, blue, square).4. C: (large, yellow, triangle)  $>$  A: (large, blue, square).Therefore, C: (large, yellow, triangle)  $!<$  A: (large, blue, square)large blue things in C is not to the left of the large blue thing in A. Answer: No

**user** : There is a block called A. In A there is a red triangle Which is above a black circle. The black circle is touching the bottom of A. A yellow square is touching the top edge of A and below that there is a red oval is touching the right edge of A.



There is a second block call B Which is to the left of A. Inside B there are 3 squares, a brown one, a green one, and a red one. The green square is above the red square and touching the left edge of B. The brown square is touching both the bottom edge and the right edge of B. The red square is touching the left edge of the brown square. Is the oval to the right of the brown thing?

### C.5 Step-by-Step Prompting

**Raw context :** There exist a big red square, a big red triangle, a medium red square, and a small green circle in a block called A. The triangle and medium square are touching the bottom edge of the block. The big and medium square are touching the right edge of the block. And the circle is above the big square which is above the medium square. There is another block called B to the left of block A. A medium green square is touching the left edge of block B and is below a medium red square. The medium red square is above and to the left of a small green square. Also a medium red triangle is below and to the left of the small square.

**Splitting context :** A big red square in block A.

A big red triangle in block A.

A medium red square in block A.

A small green circle in block A.

The triangle is touching the bottom edge of block A.

The medium square is touching the bottom edge of block A.

The big square is touching the right edge of block A.

The medium square is touching the right edge of block A.

The circle is above the big square in A.

The big square is above the medium square in A.

Block B is to the left of block A.

A medium green square is touching the left edge of block B.

A medium green square is below a medium red square in B.

The medium red square is above a small green square in B.

The medium red square is left of a small green square in B.

The medium red triangle is below the small square in B.

The medium red triangle is left the small square in B.

### D Spatial Logical Rules

The conversion from spatial logical rules proposed in (Mirzaee and Kordjamshidi, 2022) to logical constraints used in our experiment is shown in Table 7.

Rule Type	Rule	Logical Constraints
Converse	$above(x, y) \Rightarrow below(y, x)$ $below(x, y) \Rightarrow above(y, x)$ $left(x, y) \Rightarrow right(y, x)$ $right(x, y) \Rightarrow left(y, x)$ $front(x, y) \Rightarrow behind(y, x)$ $behind(x, y) \Rightarrow front(y, x)$ $coveredby(x, y) \Rightarrow cover(y, x)$ $cover(x, y) \Rightarrow coveredby(y, x)$ $inside(x, y) \Rightarrow contain(y, x)$ $contain(x, y) \Rightarrow inside(y, x)$	$Truth(q_i) \Rightarrow Truth(q_j)$
Symmetric	$near(x, y) \Rightarrow near(y, x)$ $far(x, y) \Rightarrow far(y, x)$ $touch(x, y) \Rightarrow touch(y, x)$ $disconnected(x, y) \Rightarrow disconnected(y, x)$ $overlap(x, y) \Rightarrow overlap(y, x)$	$Truth(q_i) \Rightarrow Truth(q_j)$
Transitivity	$left(x, y) \wedge left(y, z) \Rightarrow left(x, z)$ $right(x, y) \wedge right(y, z) \Rightarrow right(x, z)$ $above(x, y) \wedge above(y, z) \Rightarrow above(x, z)$ $below(x, y) \wedge below(y, z) \Rightarrow below(x, z)$ $behind(x, y) \wedge behind(y, z) \Rightarrow behind(x, z)$ $front(x, y) \wedge front(y, z) \Rightarrow front(x, z)$ $inside(x, y) \wedge inside(y, z) \Rightarrow inside(x, z)$ $contain(x, y) \wedge contain(y, z) \Rightarrow contain(x, z)$ $inside(x, y) \wedge coveredby(y, z) \Rightarrow inside(x, z)$ $contain(x, y) \wedge cover(y, z) \Rightarrow contain(x, z)$ $inside(x, y) \wedge left(y, z) \Rightarrow left(x, z)$ $inside(x, y) \wedge right(y, z) \Rightarrow right(x, z)$ $inside(x, y) \wedge above(y, z) \Rightarrow above(x, z)$ $inside(x, y) \wedge below(y, z) \Rightarrow below(x, z)$ $inside(x, y) \wedge behind(y, z) \Rightarrow behind(x, z)$ $inside(x, y) \wedge front(y, z) \Rightarrow front(x, z)$ $inside(x, y) \wedge near(y, z) \Rightarrow near(x, z)$ $inside(x, y) \wedge far(y, z) \Rightarrow far(x, z)$ $inside(x, y) \wedge disconnected(y, z) \Rightarrow disconnected(x, z)$ $coveredby(x, y) \wedge left(y, z) \Rightarrow left(x, z)$ $coveredby(x, y) \wedge right(y, z) \Rightarrow right(x, z)$ $coveredby(x, y) \wedge above(y, z) \Rightarrow above(x, z)$ $coveredby(x, y) \wedge below(y, z) \Rightarrow below(x, z)$ $coveredby(x, y) \wedge behind(y, z) \Rightarrow behind(x, z)$ $coveredby(x, y) \wedge front(y, z) \Rightarrow front(x, z)$ $coveredby(x, y) \wedge near(y, z) \Rightarrow near(x, z)$ $coveredby(x, y) \wedge far(y, z) \Rightarrow far(x, z)$ $coveredby(x, y) \wedge disconnected(y, z) \Rightarrow disconnected(x, z)$	$Truth(q_i) \wedge Truth(q_j) \Rightarrow Truth(q_k)$

Table 7: The conversion from spatial logical rules proposed in (Mirzaee and Kordjamshidi, 2022) to logical constraints used in our experiment.

## E DomiKnowS Concepts and Graphs

Graph and concepts declaration of our work inside the DomiKnowS framework.

### E.1 Graph for YN

```

story = Concept(name="story")
question = Concept(name="question")
story_contain, = story.contains(question)
answer_class = question(name="
    answer_class", ConceptClass=

```

```

EnumConcept,
    values=["yes", "
        no"])
symmetric = Concept(name="symmetric")
s_quest1, s_quest2 = symmetric.has_a(
    arg1=question, arg2=question)
ifL(andL(answer_class.yes('x'), existsL(
    symmetric('s', path=('x', symmetric)
))),
    answer_class.yes(path=('s', s_quest2
        )))
reverse = Concept(name="reverse")

```

```

r_quest1, r_quest2 = reverse.has_a(arg10
=question, arg20=question)
ifL(andL(answer_class.yes('x'), existsL(
reverse('r', path=('x', reverse)))),
answer_class.no(path=('r', r_quest2)
))
ifL(andL(answer_class.no('x'), existsL(
reverse('r', path=('x', reverse)))),
answer_class.yes(path=('r', r_quest2
)))

transitive = Concept(name="transitive")
t_quest1, t_quest2, t_quest3 =
transitive.has_a(arg11=question,
arg22=question, arg33=question)
ifL(andL(answer_class.yes('x'), existsL(
transitive("t", path=('x',
transitive))), answer_class.yes(path
=('t', t_quest2))),
answer_class.yes(path=('t', t_quest3
)))

transitive_topo = Concept(name="
transitive_topo")
tt_quest1, tt_quest2, tt_quest3,
tt_quest4 = transitive_topo.has_a(
arg111=question, arg222=question,
arg333=question, arg444=question)
ifL(andL(answer_class.yes('x'), existsL(
transitive("t", path=('x',
transitive_topo))),
answer_class.yes(path=('t',
tt_quest2)), answer_class.
yes(path=('t', tt_quest3)))
,
answer_class.yes(path=('t',
tt_quest4)))

```

## E.2 Graph for FR on SpARTUN

```

story = Concept(name="story")
question = Concept(name="question")
story_contain, = story.contains(question
)

left = question(name="left")
right = question(name="right")
above = question(name="above")
below = question(name="below")
behind = question(name="behind")
front = question(name="front")
near = question(name="near")
far = question(name="far")
disconnected = question(name="
disconnected")
touch = question(name="touch")
overlap = question(name="overlap")
coveredby = question(name="coveredby")
inside = question(name="inside")
cover = question(name="cover")
contain = question(name="contain")
output_for_loss = question(name="
output_for_loss")

# Only one label of opposite concepts
exactL(left, right)
exactL(above, below)
exactL(behind, front)
exactL(near, far)

```

```

exactL(disconnected, touch)

# Inverse Constrains
inverse = Concept(name="inverse")
inv_question1, inv_question2 = inverse.
has_a(arg1=question, arg2=question)

# First inverse relation, allow inverse
back and forth
inverse_list1 = [(above, below), (left,
right), (front, behind), (coveredby,
cover),
(inside, contain)]

for ans1, ans2 in inverse_list1:
ifL(andL(ans1('x'), existsL(inverse(
's', path=('x', inverse))),
ans2(path=('s', inv_question2)))

ifL(andL(ans2('x'), existsL(inverse(
's', path=('x', inverse))),
ans1(path=('s', inv_question2)))

# 2 PMD := entropy + beta *
constraint_loss (Train with no-
constraint first then working on)
# symmetric
inverse_list2 = [(near, near), (far, far
), (touch, touch), (disconnected,
disconnected), (overlap, overlap)]
for ans1, ans2 in inverse_list2:
ifL(andL(ans1('x'), existsL(inverse(
's', path=('x', inverse))),
ans2(path=('s', inv_question2)))

# Transitive constrains
transitive = Concept(name="transitive")
tran_quest1, tran_quest2, tran_quest3 =
transitive.has_a(arg11=question,
arg22=question, arg33=question)

transitive_1 = [left, right, above,
below, behind, front, inside,
contain]

for rel in transitive_1:
ifL(andL(rel('x'),
existsL(transitive("t",
path=('x', transitive))
),
rel(path=('t', tran_quest2)
)),
rel(path=('t', tran_quest3)))
# Transitive of cover and contain
transitive_2 = [(coveredby, inside), (
cover, contain)]
for rel1, rel2 in transitive_2:
ifL(andL(rel2('x'),
existsL(transitive("t",
path=('x', transitive))
),
rel1(path=('t', tran_quest2
))),
rel2(path=('t', tran_quest3)))

# Transitive of inside/cover with
position
transitive_3_1 = [inside, coveredby]
transitive_3_2 = [left, right, above,
below, behind, front, near, far,

```

```

disconnected]
for rel1 in transitive_3_1:
    for rel2 in transitive_3_2:
        ifL(andL(rel1('x'),
                    existsL(transitive("t",
                                      path=('x',
                                             transitive))),
                    rel2(path=('t',
                               tran_quest2))),
            rel2(path=('t', tran_quest3)
                ))

# Transitive + topo constrains
tran_topo = Concept(name="
    transitive_topo")
tran_topo_quest1, tran_topo_quest2,
    tran_topo_quest3, tran_topo_quest4 =
    tran_topo.has_a(arg111=question,
        arg222=question, arg333=question,
        arg444=question)
# (x inside y) + (h inside z) + (y
    direction z) => (x direction h)
tran_topo_2_1 = [inside, coveredby]
tran_topo_2_2 = [left, right, above,
    below, behind, front, near, far,
    disconnected]
for rel1 in tran_topo_2_1:
    for rel2 in tran_topo_2_2:
        ifL(andL(rel1('x'),
                    existsL(tran_topo('to',
                                      path=('x',
                                             tran_topo))),
                    rel1(path=('to',
                               tran_topo_quest2))),
            rel2(path=('to',
                       tran_topo_quest3)
                ),
            rel2(path=('to',
                       tran_topo_quest4)))

tran_topo_3_1 = [left, right, above,
    below, behind, front, near, far,
    disconnected]
tran_topo_3_2 = [contain, cover]
for rel1 in tran_topo_3_1:
    for rel2 in tran_topo_3_2:
        ifL(andL(rel1('x'),
                    existsL(tran_topo('to',
                                      path=('x',
                                             tran_topo))),
                    rel1(path=('to',
                               tran_topo_quest2))),
            rel2(path=('to',
                       tran_topo_quest3)
                ),
            rel1(path=('to',
                       tran_topo_quest4)))

```

```

lower_left = question(name="lower-left")
lower_right = question(name="lower-right")
upper_left = question(name="upper-left")
upper_right = question(name="upper-right")
overlap = question(name="overlap")

```

### E.3 Graph for FR on StepGame

```

story = Concept(name="story")
question = Concept(name="question")
story_contain, = story.contains(question
)

left = question(name="left")
right = question(name="right")
above = question(name="above")
below = question(name="below")

```