

# Logical Reasoning over Natural Language as Knowledge Representation: A Survey

Anonymous ACL submission


## Abstract

Logical reasoning is central to human cognition and intelligence. It includes deductive, inductive, and abductive reasoning. Past research of logical reasoning within AI uses formal language as knowledge representation and symbolic reasoners. However, reasoning with formal language has proved challenging (e.g., brittleness and knowledge-acquisition bottleneck). This paper provides a comprehensive overview on a new paradigm of logical reasoning, which uses natural language as knowledge representation and pretrained language models as reasoners, including philosophical definition and categorization of logical reasoning, advantages of the new paradigm, benchmarks and methods, challenges of the new paradigm, possible future directions, and relation to related NLP fields. This new paradigm is promising since it not only alleviates many challenges of formal representation but also has advantages over end-to-end neural methods. This survey focus on transformer-based LLMs explicitly working on deductive, inductive, and abductive reasoning over English representation.


## 1 Introduction

An argument consists of premise(s) and a conclusion. Logical reasoning is a form of thinking in which premises and relations between premises are used in a rigorous manner to infer conclusions that are entailed (or implied) by the premises and the relations (Nunes, 2012). It consists of three reasoning types, namely deductive reasoning, inductive reasoning, and abductive reasoning (Flach and Kakas, 2000) (more illustration on the categorization can be found in §2). It is important since the ability to reach logical conclusions on the basis of prior information is recognized as central to human cognition and intelligence (Goel et al., 2017).

The past research of logical reasoning within AI uses formal language (e.g., first-order logic)



(a) Formal language as knowledge representation and symbolic reasoner



(b) Natural language as knowledge representation and PLM as reasoner

Figure 1: Comparison between the previous paradigm which uses formal representation and symbolic reasoner, and the new paradigm which uses natural language as knowledge representation and LLM as reasoner.

as knowledge representation and symbolic reasoners (Muggleton and Raedt, 1994). This paradigm has resulted in impressive applications such as expert systems (Metaxiotis et al., 2002). However, building and reasoning over formal language have proved challenging (Musen and Van der Lei, 1988), with representative disadvantages of brittleness (an expert system fails as long as its knowledge base does not contain complete knowledge for a problem) and knowledge-acquisition bottleneck (human experts are needed to encode their knowledge with formal representation).

Since the rapid development in language models, natural language has been explored as a new knowledge representation, and large language models (LLMs) have been used as a new reasoner for deductive reasoning (Clark et al., 2020), abductive reasoning (Bhagavatula et al., 2020), and inductive reasoning (Yang et al., 2022b). Therefore, all three reasoning types of logical reasoning have been investigated with natural language as knowledge representation. This research also shows that LLMs can be finetuned or prompted to perform well for each of the reasoning types.

In this paper, we summarize the three previously separately investigated logical reasoning types to-

gether, referred as logical reasoning (from the perspectives of deductive, inductive, and abductive reasoning) over natural language as knowledge representation and LLMs as reasoners (LRNLP), and provide an in-depth survey of LRNLP.

Illustrated in Figure 1, LRNLP means a new paradigm for logical reasoning that uses new knowledge representation (natural language) and new reasoner (LLM). LRNLP can also be seen as a set of tasks on the three reasoning with a constraint of natural language representation and LLM reasoners. The latest methods for the LRNLP tasks are generally modular: multiple LLMs each as one module playing a different function, combined together to perform complex tasks. They make one step of reasoning with one inference of LLM. For complex problems, they usually have access to a knowledge base that stores relevant textual knowledge to be retrieved as premises to support the reasoning process to reach a conclusion, which might be used as a new premise for the next step’s reasoning. By iteratively repeating this process, a final conclusion may be made. Although it looks similar to expert systems, we discuss how LRNLP is possible to overcome many main challenges of the previous paradigm such as brittleness in §3.1.

In addition to the comparison with formal language, in §3.2 we discuss that LRNLP could be viewed as a new type of neural-symbolic method, which has unique advantages over existing neuro-symbolic methods. We also discuss how LRNLP, as a neuro-symbolic method, has advantages over existing end-to-end neural methods (e.g., explainability, controllability, less catastrophic forgetting) in §3.3. These advantages make an LRNLP system possible to deal with many challenging problems.

In the remaining sections of this survey, we review papers on LRNLP (including deductive reasoning §4, inductive reasoning §5, and abductive reasoning §6), and list challenges (§7. Our main focus is to understand the language model’s logical reasoning ability through the three sub-types of logical reasoning to provide finer analysis and avoid ambiguity on which type of reasoning it is conducting. Therefore, we focus on papers that using transformer-based LLMs explicitly working on deductive, inductive, or abductive reasoning tasks. These papers all adopt English as knowledge representation. In §A.1 we discuss the relation of LRNLP to related NLP fields (e.g., commonsense reasoning), which could help to form a clear shape of LRNLP in NLP. For each reasoning sub-type,

we summarize existing task formulations, datasets, and methods under each task.

## 2 Definition and Categorization

There are many subjects related to logical reasoning, including philosophy, logic, and AI. Among them, the definition and categorization aspects of logical reasoning are handled by philosophy research. However, debate exists in philosophy research on the categorization of logical reasoning. We leave a detailed description of the debate in philosophy research in §A.2 and only leave the conclusions here according to philosophy research.

In general, logical reasoning consists of deductive, inductive, and abductive reasoning (Console and Saitta, 2000). Given an argument consisting of premises and a conclusion, we define the sub-type of logical reasoning it involves below:

*Definition for deductive reasoning:* the premises can conclusively provide support for the conclusion, i.e. if the premises are all true, it would be impossible for the conclusion to be false.

*Definition for inductive reasoning:* the premises cannot conclusively provide support for the conclusion, since the conclusion generalizes existing information in premises to new knowledge, which has a wider applicable scope than those in premises.

*Definition for abductive reasoning:* the premises cannot conclusively provide support for the conclusion, since the conclusion contains more specific information over the premises (most commonly used as generating most probable explanations).

Please note that according to Console and Saitta (2000), inductive reasoning and abductive reasoning are not exclusive to each other.

## 3 Advantages of LRNLP

### 3.1 Advantages over Formal Language

Building and reasoning over formal language have proved challenging (Musen and Van der Lei, 1988; Cropper et al., 2022), with disadvantages such as (1) *brittleness* (expert system fails when its knowledge base does not contain complete knowledge for a problem), (2) *knowledge-acquisition bottleneck* (human experts are needed to encode their knowledge with formal representation), (3) *inability to handle raw data such as natural language*, (4) *sensitivity to label errors*, and (5) *failure to recognize different symbols with similar meanings*.

Nevertheless, the new paradigm of logical reasoning, LRNLP, has systematic strengths over these

challenges. Specifically, LLMs contain knowledge themselves (Davison et al., 2019), which makes it possible for them to provide good answers even when some required explicit knowledge is not present in a knowledge base Talmor et al., 2020 (less brittle), and be less affected by input errors (Meng et al., 2021). In addition, with natural language as knowledge representation, such a system can naturally handle raw input, and it is possible to utilize the enormous web corpora to automatically construct rule bases using information extraction Ji, 2018 (less affected by knowledge-acquisition bottleneck); using embeddings for concepts (Mikolov et al., 2013), it semantically “understands” the meaning of symbols and therefore robust for paraphrasing.

### 3.2 Advantages over Neuro-symbolic Systems

LRNLP could be seen as a new type of neuro-symbolic in addition to the existing 6 types(Kautz, 2022), as its goal and design of methodology are typically symbolic (logical reasoning with knowledge bases), while avoiding any symbolic representation, using (currently pure) neural methods. Therefore LRNLP can avoid many bottlenecks of the other neuro-symbolic methods caused by symbolic representation, such as symbolic knowledge acquisition and scalability (Wang and Yang, 2022).

### 3.3 Advantages over E2E Neural Methods

As a neuro-symbolic method, LRNLP systematically has some advantages over end-to-end neural methods, such as interpretability Cambria et al., 2023 (since it is usually stepwise), more controllability (LRNLP reasons following a given knowledge base), and less catastrophic forgetting (LRNLP uses an explicit knowledge base).

## 4 Deductive Reasoning

### 4.1 Existing Task Formulations

Existing tasks for deductive reasoning can be summarized as hypothesis classification, proof generation, proof generation with incomplete information, and implication enumeration. Datasets for tasks are summarized in Table 1. “Proof generation” tab with  $\times$  means it is for hypothesis classification task.

**Hypothesis Classification** Each data example for hypothesis classification task is a tuple  $(theory, hypothesis, correctness)$ , where  $theory$  typically has the form  $(fact^*, rule^*)$ ,  $hypothesis$  is a question, and  $correctness$  can be

Dataset	Human written	Realistic	Multi-step	Theory included	Theory sufficient	Proof generation	Size
D*	$\times$	$\times$	$\checkmark$	$\checkmark$	$\checkmark$	$\times$	500k
ParaRules	$\checkmark$	$\times$	$\checkmark$	$\checkmark$	$\checkmark$	$\times$	40k
Birds-electricity	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\times$	5k
Leap-of-thought	$\times$	$\checkmark$	$\times$	$\checkmark$	$\times$	$\times$	33k
PARARULE-Plus	$\times$	$\times$	$\checkmark$	$\checkmark$	$\checkmark$	$\times$	400k
FOLIO	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\times$	1,435
D*(CWA)	$\times$	$\times$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	500k
D*(OWA)	$\times$	$\times$	$\checkmark$	$\checkmark$	$\times$	$\checkmark$	500k
EntailmentBank	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	1,840
ENWN	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	100

Table 1: Summary of deductive reasoning datasets: D\*, ParaRules & birds-electricity (Clark et al., 2020); leap-of-thought (Talmor et al., 2020); PARARULE-Plus (Bao et al., 2022); FOLIO (Han et al., 2022); D\*(CWA) & D\*(OWA) Tafford et al., 2021; EntailmentBank (Dalvi et al., 2021); ENWN (Sprague et al., 2022).

*True* or *False* (or *Unknown*). This task requires to predict the *correctness* for the *hypothesis* given the *theory*.

**Proof Generation** The proof generation task has the same setting as the hypothesis classification task, except that in addition to predicting a *correctness*, the proof generation task also requires providing a *proof* given *theory* to explain the *correctness*. The *proof* is a directed tree  $(\mathcal{N}, \mathcal{E})$  with nodes  $n \in \mathcal{N}$  and edges  $e \in \mathcal{E}$ . Each node is an item of knowledge in *theory* (usually a *fact* or a *rule*), or a generated intermediate reasoning conclusion, or the *hypothesis* itself; Each edge points from a premise node to a conclusion node to form a deductive argument, which typically needs one-step inference (not multi-step).

### Proof Generation with Incomplete Information

This task is the same as the proof generation task, except that *theory* lacks one *node* to form a complete *proof*. Specifically, given *theory*, it requires to predict the *correctness* of *hypothesis* with a *proof*, as well as recovering the missing *node*.

**Implication Enumeration** Given a *theory*, this task requires to enumerate implications of the *theory*, using deductive reasoning.

## 4.2 Methods

### 4.2.1 Hypothesis Classification

There are mainly three categories of methods for the hypothesis classification task. The first category only conducts the classification task itself; the second category can predict *correctness* as well as generate a *proof*. However, the *correctness* is not necessarily consistent with the predicted *proof*.

Method	Generation based	Inference w/ hypothesis	Stepwise	Proof direction	Heuristic search	Verifier	Human-authored realistic proof	Stage
PProver (Saha et al., 2020)	✗	✓	✗	N/A	N/A	✗	✗	1
multiPProver (Saha et al., 2021)	✗	✓	✗	N/A	N/A	✗	✗	1
EntailmentWriter (Dalvi et al., 2021)	✓	✓	✗	N/A	N/A	✗	✓	1
ProofWriter (Tafjord et al., 2021)	✓	✗	✓	→	✗	✗	✗	2
EVR (Liang et al., 2021)	✓	✗	✓	←	✗	✗	✗	2
IBR (Qu et al., 2022)	✗	✓	✓	←	✓	✗	✗	2
IRGR (Ribeiro et al., 2022)	✓	✓	✓	→	✓	✗	✓	2
SI (Creswell et al., 2022)	✓	✗	✓	→	✓	✗	✗	2
FaiRR (Sanyal et al., 2022b)	✓	✗	✓	→	✓	✗	✗	2
MetGen (Hong et al., 2022)	✓	✗	✓	Both	✓	✗	✓	2
SCSearch (Bostrom et al., 2022)	✓	✗	✓	→	✓	✗	✓	2
ADGV (Sprague et al., 2022)	✓	✗	✓	Both	✓	✓	✓	3
NLProofS (Yang et al., 2022a)	✓	✓	✓	→	✓	✓	✓	3
Entailer (Tafjord et al., 2022)	✓	✓	✓	←	✓	✓	✓	3
Teachme (Dalvi et al., 2022)	✓	✓	✓	←	✓	✓	✗	3

Table 2: Methods for Proof Generation task. “Generation based” means whether *proof* is created by generative inference model, otherwise is by utilizing embeddings to classify nodes and edges of *proof*. “Inference w/ hypothesis” means whether *hypothesis* is provided during inference. → and ← denote forward/backward stepwise proof generation. “Heuristic search” with ✗ means exhaustive search. “Human-authored realistic proof” means whether the dataset adopted uses human-authored *proof*, whose contents are consistent with the real world.

The third category is similar to the second, except that *correctness* always follows *proof*.

Until now, methods from the first category directly use transformer-based LLMs (Vaswani et al., 2017), aiming at analyzing and benchmarking their performance. Specifically, Clark et al. (2020) find that finetuned RoBERTa-large (Liu et al., 2019) can achieve 95%+ accuracy on the test set of D\* and ParaRules datasets; Talmor et al. (2020) further demonstrate that LLMs can be finetuned to reliably perform deductive reasoning using both implicit, pretrained knowledge and explicit natural language statements (*theory*) to make predictions; Han et al. (2022) evaluate finetuned medium-sized language models and few-shot prompting on LLMs on the FOLIO dataset. However, they find that LLM with few-shot prompting only performs slightly better than random results.

The second category methods typically infer LLM only once, and then utilize the final layer embeddings or generations to obtain *correctness* and *proof*. Specifically, PProver (Saha et al., 2020) and multiPProver (Saha et al., 2021) use the [CLS] token to predict *correctness*, and leverage the final layer embeddings of knowledge items in *theory* to generate *proof*; All-At-Once ProofWriter (Tafjord et al., 2021) and EntailmentWriter (Dalvi et al., 2021) generate *correctness* and linearized *proof* at the same time.

The third category methods create a *proof* first, and then predict *correctness* from the *proof*.

§4.2.2 illustrates these methods in detail.

#### 4.2.2 Proof Generation

Methods for this task are summarized in Table 2. Current methods for the proof generation task roughly consist of three stages. In each stage, one key new technique is considered and developed. In stage 1, LLMs are used for forming *proof* in one inference step. In stage 2, modular-based, stepwise frameworks are developed to create *proof* (each module is usually implemented with a single LLM). In stage 3, a verifier is added as a new module to make sure that each reasoning step reflects the belief of LLMs. We summarize the experiment results in §A.8 and the model structures in §A.9.

Methods for stage 1 typically utilize the last layer embeddings (Saha et al., 2020, 2021) or generations (Tafjord et al., 2021; Dalvi et al., 2021) to create *proof*. Methods utilizing embedding typically (1) obtain an averaged embedding for each knowledge item in *theory*, and (2) pass each embedding to a node classifier, and each embedding pairs to an edge classifier to predict nodes and edges for *proof*. Constraints are usually used to enforce the structure of *proof*. Generation methods directly generate linearized *correctness* and full *proof* given linearized *theory* and *hypothesis*.

The motivations of stage 2 methods are generally concerned with end-to-end methods, which is considered to lack interpretability (Liang et al., 2021; Qu et al., 2022; Sanyal et al., 2022b; Bostrom et al., 2022), suffer from compositional generalization



problems (Liang et al., 2021; Creswell et al., 2022), have limited input size (Ribeiro et al., 2022), are not causal (Creswell et al., 2022), and lack constraints on inference validity (Hong et al., 2022).

Methods in stage 2 can be summarized as having two components, an inference module and a reasoning controller. The inference module can be a deduction module (Tafjord et al., 2021; Ribeiro et al., 2022; Creswell et al., 2022; Sanyal et al., 2022b; Bostrom et al., 2022), an abduction module (Liang et al., 2021; Qu et al., 2022), or both (Hong et al., 2022; Sprague et al., 2022). The deduction module performs deductive reasoning, and reasons forwardly from *theory* to *hypothesis* to construct *proof*; the abduction module performs abductive reasoning, and reasons backwardly from *hypothesis* to *theory* to construct *proof*. The reasoning controller in general performs a search process that each step it searches through the *theory* and generated intermediate conclusions space to select (retrieve) premises for the next step inference. The search processes include exhaustive search (Tafjord et al., 2021; Liang et al., 2021) or heuristic search (Qu et al., 2022; Ribeiro et al., 2022; Creswell et al., 2022; Sanyal et al., 2022b; Bostrom et al., 2022; Hong et al., 2022; Sprague et al., 2022). The reasoning controller usually can also stop the search process if it detects the goal.

Motivation of stage 3 methods is similar, basically that stage 2 methods lack explicit verifiers to avoid hallucinating invalid steps (Yang et al., 2022a), and to ensure that the inference processes reflect LLM’s own beliefs (Tafjord et al., 2022).

Methods in stage 3 can be summarized as utilizing explicit verifier(s) (implemented with a LLM) to check the validity of each inference step. One way is to add a new module (additional to the inference module and reasoning controller in stage 2), working as a “fact checker” to verify the generated inference step (Yang et al., 2022a; Tafjord et al., 2022); The other one, called round-trip consistency, is only suitable for methods that use both deduction and abduction modules, where deduction and abduction modules work as the verifier for each other (Sprague et al., 2022).

In addition to the general 3 stages, a new aspect is attended to, which is whether teachable by humans. Build based on Entailer (Tafjord et al., 2022), TeachMe (Dalvi et al., 2022) shows that user corrections can help override erroneous model beliefs, and that a system can gradually improve by accumulating user corrections. Compared to Entailer, it

adds an interaction module and a dynamic memory module to obtain and store human corrections.

#### 4.2.3 Proof with Incomplete Information

ADGV (Sprague et al., 2022) is the only method focusing on this task. It uses both deduction and abduction modules, and the reasoning controller performs heuristic search. The abduction module is used to recover the missing premise.

#### 4.2.4 Implication Enumeration

Tafjord et al. (2021) is the only paper mentioned this task. They compare the performance of “All-At-Once” and “Iterative” ProofWriter on this task. They find that “All-At-Once” performs worse, mainly because it struggles with problems that are more complex than training examples.

### 4.3 Robustness of LLM as Reasoner

The previously introduced methods only focus on solving the deductive reasoning tasks, while it is unclear whether LLMs can be used as robust deductive reasoners. To investigate the problem, Gaskell et al. (2022) create a more challenging synthetic dataset on hypothesis classification task in terms of complexity, and test LLM’s performance on it. They find that with large and complex enough training examples, transformers can perform well on the dataset. In addition, they find that transformers exhibit some degree of generalization and scale-invariance ability; Richardson and Sabharwal (2022) propose an adversarial attack method for synthetic datasets on the hypothesis classification task. They find that transformers are often fooled if the query literally appears within the body of a rule, and transformers struggle to correctly bind variables on either side of a rule; Sanyal et al. (2022a) proposed a synthetic deductive reasoning dataset to evaluate the robustness of language models to minimal logical edits in the inputs and different logical equivalence conditions, and find that LLMs are not robust to their proposed logical perturbations.

## 5 Inductive Reasoning

### 5.1 Existing Task Formulations

Existing tasks for inductive reasoning can be summarized as rule verification and rule generation tasks. Datasets for the tasks are summarized in Table 3. “Generation” tab with  $\times$  means it is for the rule verification task.

Dataset	Human written	Human labeled	Realistic	Rule provided	Not restricted rule types	Generation	Novel scientific hypotheses	Size
property-norm	✗	✗	✓	✗	✗	✗	✗	23k
DEERLET	✗	✓	✓	✓	✓	✗	✗	846
DEER	✓	✓	✓	✓	✓	✓	✓	1.2k
ARC	✗	✗	-	✗	✗	-	✗	1k
OpenD5	✓	✓	✓	✓	✓	✓	-	675
C-LBD	✓	✗	✓	✓	✓	✓	✓	67k
TOMATO	✓	✓	✓	✓	✓	✓	✓	50

Table 3: Summary of inductive reasoning datasets: property-norm (Misra et al., 2022), DEERLET and DEER (Yang et al., 2022b), ARC (Chollet, 2019), OpenD5 (Zhong et al., 2023), C-LBD (Wang et al., 2023a), and TOMATO (Yang et al., 2023b). “Not restricted rule types” means whether the data is not restricted in a specific topic (e.g., taxonomic).

**Rule Verification** Given a generated *rule* and *facts* where the *rule* is generated from, the task is to classify whether the *rule* can be accepted. The current evaluation aspects are from requirements of both inductive reasoning and natural language.

**Rule Generation** Given multiple manually selected *facts* with similar patterns, the task is to induce a *rule* that (1) can entail the *facts*, and (2) is more general than all of the *facts*. Here “more general” means larger information coverage scope. More detailed illustrations can be found in §A.10.

**Scientific Hypotheses Generation** This task is similar to *Rule Generation* task but is much more challenging in that the generated *rule* should not be commonsense knowledge but scientific hypotheses that are even new to humanity.

## 5.2 Methods

*Rule Generation* methods almost always have a *Rule Verification* step after the initial generation of rules. To have a clearer overview, we separately introduce the framing or methods of the two tasks.

### 5.2.1 Rule Verification

Yang et al. (2022b) propose three requirements of rule verification on inductive reasoning from philosophy literature (*rule* and *facts* should not be in conflict; *rule* should reflect reality; *rule* should generalize over *facts*) and one requirement of rule verification from NLP requirement (*rule* should not be trivial or incomplete). They focus on inducing *rule* of many disciplines (e.g., zoology and history) from *facts* as textual observations (e.g. Wikipedia). They implement the verification by LLMs (framing as classification problems).

Another group of works’ (Zhu et al., 2023; Wang et al., 2023b; Qiu and Jiang, 2023) adopted rule verification criteria is compliant with one of the key re-

quirements proposed by Yang et al. (2022b), which is that *rule* and *facts* should not be in conflict. They focus on inducing (executable) *rule* from synthetic *facts* such as a sequence of number (example *rule*: find the smallest number), arithmetic calculation (example *rule*: “6+4=10”), or changes of 2D grid images (example *rule*: executable code for moving the grids). They verify rules by checking the consistency of the labels of annotated examples (*facts*) and the results of *rules*.

### 5.2.2 Rule Generation

Yang et al. (2022b) assume that the inductive reasoning task is so difficult that a proper system should contain a rule populator and (multiple) rule verifiers that filter bad rules from different aspects. Accordingly, they propose a framework named chain-of-language-models (CoLM), where one LLM generates *rules* given *facts*, the other four LLMs filter generated rules mainly based on philosophical requirements of inductive reasoning.

Besides the rule generation and filtering process, Zhu et al. (2023) further propose to generate rules based on chain-of-thought prompting, and verify rules based on whether the rules can be used to deduce the annotated answer correctly; Wang et al. (2023b) further propose that under synthetic datasets, executable code can be generated for the textual rules and verify the rules by executing the code and comparing the results with groundtruth annotation; Qiu and Jiang (2023) further propose a third stage of “rule refinement,” (leveraging feedback and generate again) and that iteratively repeating the three stages can obtain better rules.

### 5.2.3 Scientific Hypotheses Generation

Zhong et al. (2023) focuses on proposing hypotheses (from many disciplines) from a research goal and two comparable corpora. Their method also follows a generate-filter process, where LLMs are used for the filtering stage. Wang et al. (2023a) focus on proposing NLP hypotheses from a seed term and background context. Before the hypotheses generation module, they build knowledge graphs to associate academic terms, and retrieve some of the terms as inspirations. Yang et al. (2023b) focuses on proposing social science and business hypotheses only from a pile of raw web corpora. To utilize raw web corpora, they expand generate-filter modules with a background finder module and an inspiration finder module. They also propose three feedback mechanisms named past feedback, present

Dataset	Human written	Realistic	Multi-step	Theory included	Generation	Size
$\alpha$ NLI	✓	✓	✗	✗	✗	22k
$\alpha$ NLG	✓	✓	✗	✗	✓	76k
AbductionRules	✗	✗	✗	✓	✓	114k
D*-Ab	✗	✗	✓	✓	✓	14k

Table 4: Summary of abductive reasoning datasets:  $\alpha$ NLI and  $\alpha$ NLG (Bhagavatula et al., 2020), AbductionRules (Young et al., 2022), and D\*-Ab (Tafjord et al., 2021). “Realistic” means whether the data is consistent with the real world. “Multi-step” means whether multiple reasoning steps are needed to get the result.

feedback, and future feedback to help the inter-communications between modules to induce more novel, valid, and helpful hypotheses.

## 6 Abductive Reasoning

### 6.1 Existing Task Formulations

Existing tasks for abductive reasoning can be summarized as explanation classification, and explanation generation w/o and w/ theory. Datasets for the tasks are summarized in Table 4. In the table, the “generation” tab and “theory included” tab can be used to determine the task it is used for.

**Explanation Classification** Given observation  $O_1$  at time  $t_1$ , observation  $O_2$  at time  $t_2$  ( $t_2 > t_1$ ), a plausible hypothesis  $h^+$  and a implausible hypothesis  $h^-$  that explain  $O_1$  and  $O_2$ , this task is to select the most plausible hypothesis from  $h^+$  and  $h^-$ .  $O_1$  and  $O_2$  each contains a single sentence.

**Explanation Generation without Theory** Given observation  $O_1$  at time  $t_1$ , observation  $O_2$  at time  $t_2$  ( $t_2 > t_1$ ), this task is to generate a valid hypothesis  $h^+$  given  $O_1$  and  $O_2$ .  $O_1$  and  $O_2$  each is described in a single sentence.

**Explanation Generation with Theory** Given a theory  $C$  and a possible observation  $O$  not provable from  $C$ , the task is to generate a new hypothetical fact  $h$  such that  $C \cup \{h\} \models O$ . Here  $C$  contains multiple facts and rules, where each fact or rule contains a single sentence.  $O$  is in single sentence.

## 6.2 Methods

### 6.2.1 Explanation Classification

Methods for this task generally introduce knowledge in various ways to improve performance. Specifically, Mitra et al. (2019) explore ways to incorporate additional unstructured textual knowledge retrieved from a story corpus through prompt;

Paul and Frank (2020) encode and incorporate knowledge from COMET’s generation (Bosselut et al., 2019) directly into transformer’s internal attention; Lourie et al. (2021) and Paul and Frank (2021) incorporate knowledge by multi-task training; Du et al. (2021) incorporate knowledge with an additional pre-training stage using  $\mathcal{ARL}$  independent story corpora;

In addition to knowledge integration, many different aspects of explanation classification tasks are also investigated. Specifically, Bhagavatula et al. (2020) rewrite the objective using Bayes Rule and formulate a set of probabilistic models that make various independence assumptions on the new objective. They find that the most sophisticated probabilistic model works the best; Zhu et al. (2020) frame this task as a ranking task to also measure the plausibility of hypothesis in addition to discriminating it; Paul and Frank (2021) conduct this task in an unsupervised setting by pretraining on a counterfactual reasoning dataset, which is related to abductive reasoning. Kadikis et al. (2022) propose a method to select suitable LLMs for this task. It is based on the cosine similarity of  $embed(O_1, O_2)$  and  $embed(h_i)$  for each LLM without finetuning. Zhao et al. (2023) assume that different  $h$  are mutually exclusive, and improve performance by incorporating an additional loss item as regularization to enforce an unbalanced probability prediction over different  $h$ . Chan et al. (2023) exploit inter-sentential coherence and the model consistency to develop a prompt tuning model.

### 6.2.2 Explanation Generation without Theory

In general, methods for this task either incorporate knowledge or improve the decoding method to be more suitable for this task.

For knowledge integration, Bhagavatula et al. (2020) utilize textual knowledge generated from COMET and investigate two ways of knowledge integration — via texts or via embeddings, and find that the embedding-based method is more effective; Ji et al. (2020) leverage structural knowledge from ConceptNet (Speer et al., 2017) for this task.

For improving decoding method, Qin et al. (2020) are motivated by the fact that the target  $h^+$  to generate happens before  $O_2$ . They accordingly propose an unsupervised decoding algorithm that can incorporate both past and future contexts.

### 6.2.3 Explanation Generation with Theory

Tafjord et al. (2021) explore the ability of a fine-



tuned T5-11B (Raffel et al., 2020) on  $P(h|C, O)$ . Their results indicate that finetuned T5-11B can reach a high test accuracy of 93% on D\*-Ab.

## 7 Challenges and Opportunities

We list more challenges and opportunities in §A.11.

**Computationally Efficient Reasoner** Many tasks in logical reasoning over formal language have very high algorithmic complexity (Muggleton et al., 2012). Thanks to the low computational cost of each deduction step over formal language, such complex tasks could be possible. However, each deduction step in LRNLP typically costs one inference of an LLM, which makes tasks with high algorithmic complexity nearly prohibitive.

**Robust Deductive Reasoner** Symbolic deductive reasoners are not restricted to training data distributions, while neural deductive reasoners are restricted to their training data (Gontier et al., 2020; Richardson and Sabharwal, 2022). In addition, neural deductive reasoners are also vulnerable to adversarial attacks (Gaskell et al., 2022), while symbolic reasoners are robust to the attacks. The lack of robustness can lead to restricted application domains and incorrect deductive inferences.

**Better Automatic Evaluation Metrics** It is generally difficult to automatically evaluate generative reasoning implications, especially with realistic and not synthetic datasets. The difficulty mainly lies in that the same semantic meaning can be expressed with diversified forms, and that different conclusions might be all acceptable (especially in abductive and inductive reasoning). This may lead to biased evaluation when using automatic metrics.

**More Impacts on (NLP) Applications** As illustrated in §3, overall LRNLP can be seen as a new type of neuro-symbolic method, which takes the advantages from both the symbolic and sub-symbolic aspects. These characteristics make an LRNLP system possible (but might still be challenging) to deal with many (NLP) applications such as medical diagnosis and legal NLP tasks, since many medical and legal problems could be seen as pure logical reasoning problems with very large rule bases (e.g., medical knowledge and laws).

**Probabilistic Inference** In reality, pure deductive reasoning has not always been used. When people include “likely” in their expressions, uncertainty is introduced, which makes the reasoning

process probabilistic; in addition, inductive reasoning and abductive reasoning are by default non-monotonic reasoning. This uncertainty aspect has not been focused in current research. It is probably beneficial to learn from how symbolic reasoning handles uncertainty (Halpern, 2017).

**Reasoning with Incomplete Information** The current proof generation task requires all necessary premises provided to create a proof tree. Only one work (Sprague et al., 2022) focuses on proof generation with the incomplete information task. However, the task they adopt only overlooks one premise, while in reality more might be missing.

**Inductive Reasoning on Web Corpora** Currently, the dataset for rule generation tasks in inductive reasoning provides manually selected facts (Yang et al., 2022b). However, to best leverage a system’s ability to handle natural language, it should be able to work on raw web corpora to induce rules, which leads to a more challenging task of inductive reasoning on web corpora.

**Abductive Reasoning with (Long) Theory** Many tasks such as medical diagnosis conduct abductive reasoning with a long theory (e.g., medical knowledge). However, current abductive reasoning research only covers abductive commonsense reasoning (Bhagavatula et al., 2020) without given theory, or only given short, synthetic, not realistic knowledge as theory (Tafjord et al., 2021).

**Interactions between Reasoning Types** Multiple reasoning types can be used together for complex tasks. Existing works only utilize deductive reasoning with abductive reasoning to create a proof tree (Hong et al., 2022; Sprague et al., 2022). However, many other collaborations are possible, such as using inductive reasoning to collect a (large) rule base, which is to be used as the theory base for deductive reasoning.

## 8 Conclusion

In this survey, we review papers using transformer-based LLMs explicitly working on deductive, inductive, and abductive reasoning over English representation. Specifically, we have introduced the philosophical foundations, advantages of LRNLP, benchmarks and methods, challenges of LRNLP, possible future directions, and the relation of LRNLP to related NLP fields (§A.1).



## Limitations

In consideration of space constraints, this paper focuses more on (1) providing a high-level overview and prospect of the LRNLP field (e.g., advantages and challenges of the field), and (2) delineating the broader evolutionary trajectories of pertinent methodologies. It might not include all the details of the surveyed papers.

## Ethics Statement

This article follows the ACL Code of Ethics. To our knowledge, there are no foreseeable potential risks to use the datasets and methods in this paper.

## References

- A. Aamodt and E. Plaza. 1994. Case-based reasoning:foundational issues,methodological variations,and system approaches. *AI communications*.
- Qiming Bao, Alex Yuxuan Peng, Tim Hartill, Neset Tan, Zhenyun Deng, Michael Witbrock, and Jiamou Liu. 2022. Multi-step deductive reasoning over natural language: An empirical study on out-of-distribution generalisation. The 2nd International Joint Conference on Learning and Reasoning and 16th International Workshop on Neural-Symbolic Learning and Reasoning (IJCLR-NeSy 2022).
- Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Wen-tau Yih, and Yejin Choi. 2020. *Abductive commonsense reasoning*. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. 2020. *PIQA: reasoning about physical commonsense in natural language*. 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*, pages 7432–7439. AAAI Press.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. *COMET: commonsense transformers for automatic knowledge graph construction*. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4762–4779. Association for Computational Linguistics.
- Kaj Bostrom, Zayne Sprague, Swarat Chaudhuri, and Greg Durrett. 2022. *Natural language deduction*

*through search over statement compositions*. *CoRR*, abs/2201.06028.

- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. *A large annotated corpus for learning natural language inference*. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 632–642. The Association for Computational Linguistics.
- Erik Cambria, Qian Liu, Sergio Decherchi, Frank Xing, and Kenneth Kwok. 2022. *Senticnet 7: A commonsense-based neurosymbolic AI framework for explainable sentiment analysis*. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference, LREC 2022, Marseille, France, 20-25 June 2022*, pages 3829–3839. European Language Resources Association.
- Erik Cambria, Lorenzo Malandri, Fabio Mercorio, Mario Mezzanzanica, and Navid Nobani. 2023. *A survey on XAI and natural language explanations*. *Inf. Process. Manag.*, 60(1):103111.
- Chunkit Chan, Xin Liu, Tsz Ho Chan, Jiayang Cheng, Yangqiu Song, Ginny Wong, and Simon See. 2023. Self-consistent narrative prompts on abductive natural language inference. *arXiv preprint arXiv:2309.08303*.
- François Chollet. 2019. *On the measure of intelligence*. *CoRR*, abs/1911.01547.
- Peter Clark, Oyvind Tafjord, and Kyle Richardson. 2020. *Transformers as soft reasoners over language*. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 3882–3890. ijcai.org.
- Luca Console and Lorenza Saitta. 2000. On the relations between abductive and inductive explanation. *Abduction and Induction: Essays on their Relation and Integration*, pages 133–151.
- Antonia Creswell, Murray Shanahan, and Irina Higgins. 2022. *Selection-inference: Exploiting large language models for interpretable logical reasoning*. *CoRR*, abs/2205.09712.
- Andrew Cropper, Sebastijan Dumancic, Richard Evans, and Stephen H. Muggleton. 2022. *Inductive logic programming at 30*. *Mach. Learn.*, 111(1):147–172.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. The pascal recognising textual entailment challenge. In *Machine learning challenges workshop*, pages 177–190. Springer.
- Bhavana Dalvi, Peter Jansen, Oyvind Tafjord, Zhengnan Xie, Hannah Smith, Leighanna Pipatanangkura, and Peter Clark. 2021. *Explaining answers with entailment trees*. In *Proceedings of the 2021 Conference*

- on *Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 7358–7370. Association for Computational Linguistics.
- Bhavana Dalvi, Oyvind Tafjord, and Peter Clark. 2022. [Towards teachable reasoning systems](#). *CoRR*, abs/2204.13074.
- Rajarshi Das, Ameya Godbole, Ankita Naik, Elliot Tower, Manzil Zaheer, Hannaneh Hajishirzi, Robin Jia, and Andrew McCallum. 2022. [Knowledge base question answering by case-based reasoning over subgraphs](#). In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 4777–4793. PMLR.
- Rajarshi Das, Manzil Zaheer, Dung Thai, Ameya Godbole, Ethan Perez, Jay Yoon Lee, Lizhen Tan, Lazaros Polymenakos, and Andrew McCallum. 2021. [Case-based reasoning for natural language queries over knowledge bases](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 9594–9611. Association for Computational Linguistics.
- Joe Davison, Joshua Feldman, and Alexander M. Rush. 2019. [Commonsense knowledge mining from pre-trained 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 2019, Hong Kong, China, November 3-7, 2019*, pages 1173–1178. Association for Computational Linguistics.
- Li Du, Xiao Ding, Ting Liu, and Bing Qin. 2021. [Learning event graph knowledge for abductive reasoning](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 5181–5190. Association for Computational Linguistics.
- Peter A Flach and Antonis C Kakas. 2000. Abductive and inductive reasoning: background and issues. In *Abduction and induction*, pages 1–27. Springer.
- Alexander Gaskell, Yishu Miao, Francesca Toni, and Lucia Specia. 2022. [Logically consistent adversarial attacks for soft theorem provers](#). In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 4129–4135. ijcai.org.
- Vinod Goel, Gorka Navarrete, Ira A Noveck, and Jérôme Prado. 2017. The reasoning brain: The interplay between cognitive neuroscience and theories of reasoning.
- Ben Goertzel, Nil Geisweiller, Lucio Coelho, Predrag Janičić, and Cassio Pennachin. 2011. *Real-World Reasoning: Toward Scalable, Uncertain Spatiotemporal, Contextual and Causal Inference*, volume 2. Springer Science & Business Media.
- Nicolas Gontier, Koustuv Sinha, Siva Reddy, and Christopher Pal. 2020. [Measuring systematic generalization in neural proof generation with transformers](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Joseph Y Halpern. 2017. *Reasoning about uncertainty*. MIT press.
- Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhenting Qi, Martin Riddell, Luke Benson, Lucy Sun, Ekaterina Zubova, Yujie Qiao, Matthew Burtell, David Peng, Jonathan Fan, Yixin Liu, Brian Wong, Malcolm Sailor, Ansong Ni, Linyong Nan, Jungo Kasai, Tao Yu, Rui Zhang, Shafiq R. Joty, Alexander R. Fabbri, Wojciech Kryscinski, Xi Victoria Lin, Caiming Xiong, and Dragomir Radev. 2022. [FOLIO: natural language reasoning with first-order logic](#). *CoRR*, abs/2209.00840.
- Ruixin Hong, Hongming Zhang, Xintong Yu, and Changshui Zhang. 2022. [METGEN: A module-based entailment tree generation framework for answer explanation](#). In *Findings of the Association for Computational Linguistics: NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 1887–1905. Association for Computational Linguistics.
- Jie Huang and Kevin Chen-Chuan Chang. 2022. [Towards reasoning in large language models: A survey](#). *CoRR*, abs/2212.10403.
- Haozhe Ji, Pei Ke, Shaohan Huang, Furu Wei, Xiaoyan Zhu, and Minlie Huang. 2020. [Language generation with multi-hop reasoning on commonsense knowledge graph](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 725–736. Association for Computational Linguistics.
- Heng Ji. 2018. [Information extraction](#). In Ling Liu and M. Tamer Özsu, editors, *Encyclopedia of Database Systems, Second Edition*. Springer.
- Yichen Jiang, Shikha Bordia, Zheng Zhong, Charles Dognin, Maneesh Kumar Singh, and Mohit Bansal. 2020. [Hover: A dataset for many-hop fact extraction and claim verification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 3441–3460. Association for Computational Linguistics.
- Emils Kadikis, Vaibhav Srivastav, and Roman Klinger. 2022. [Embarrassingly simple performance prediction for abductive natural language inference](#). In *Proceedings of the 2022 Conference of the North*

893	<i>American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022</i> , pages 6031–6037. Association for Computational Linguistics.	
898	Henry Kautz. 2022. The third ai summer: Aai robert s. engelmore memorial lecture. <i>AI Magazine</i> , 43(1):93–104.	
901	Janet L Kolodner. 1997. Educational implications of analogy: A view from case-based reasoning. <i>American psychologist</i> , 52(1):57.	
904	Rik Koncel-Kedziorski, Hannaneh Hajishirzi, Ashish Sabharwal, Oren Etzioni, and Siena Dumas Ang. 2015. <a href="#">Parsing algebraic word problems into equations</a> . <i>Trans. Assoc. Comput. Linguistics</i> , 3:585–597.	
908	George Lakoff. 1970. Linguistics and natural logic. <i>Synthese</i> , 22(1-2):151–271.	
910	Zhengzhong Liang, Steven Bethard, and Mihai Surdeanu. 2021. <a href="#">Explainable multi-hop verbal reasoning through internal monologue</a> . In <i>Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021</i> , pages 1225–1250. Association for Computational Linguistics.	
918	Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. <a href="#">Roberta: A robustly optimized BERT pretraining approach</a> . <i>CoRR</i> , abs/1907.11692.	
923	Nicholas Lourie, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. <a href="#">UNICORN on RAINBOW: A universal commonsense reasoning model on a new multitask benchmark</a> . In <i>Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021</i> , pages 13480–13488. AAAI Press.	
933	Bill MacCartney and Christopher D Manning. 2014. Natural logic and natural language inference. In <i>Computing Meaning: Volume 4</i> , pages 129–147. Springer.	
936	John McCarthy. 1990. An example for natural language understanding and the ai problems it raises. <i>Formalizing Common Sense: Papers by John McCarthy</i> , 355.	
940	Yu Meng, Yunyi Zhang, Jiaxin Huang, Xuan Wang, Yu Zhang, Heng Ji, and Jiawei Han. 2021. <a href="#">Distantly-supervised named entity recognition with noise-robust learning and language model augmented self-training</a> . In <i>Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021</i> , pages 10367–10378. Association for Computational Linguistics.	
949	Kostas S. Metaxiotis, Dimitris Askounis, and John E. Psarras. 2002. <a href="#">Expert systems in production planning and scheduling: A state-of-the-art survey</a> . <i>J. Intell. Manuf.</i> , 13(4):253–260.	949
953	Tomás Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. <a href="#">Distributed representations of words and phrases and their compositionality</a> . In <i>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</i> , pages 3111–3119.	953
961	Sewon Min, Victor Zhong, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2019. <a href="#">Multi-hop reading comprehension through question decomposition and rescoring</a> . In <i>Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers</i> , pages 6097–6109. Association for Computational Linguistics.	961
969	Kanishka Misra, Julia Taylor Rayz, and Allyson Ettinger. 2022. <a href="#">A property induction framework for neural language models</a> . <i>CoRR</i> , abs/2205.06910.	969
972	Arindam Mitra, Pratyay Banerjee, Kuntal Kumar Pal, Swaroop Mishra, and Chitta Baral. 2019. <a href="#">Exploring ways to incorporate additional knowledge to improve natural language commonsense question answering</a> . <i>CoRR</i> , abs/1909.08855.	972
977	Stephen H. Muggleton and Luc De Raedt. 1994. <a href="#">Inductive logic programming: Theory and methods</a> . <i>J. Log. Program.</i> , 19/20:629–679.	977
980	Stephen H. Muggleton, Luc De Raedt, David Poole, Ivan Bratko, Peter A. Flach, Katsumi Inoue, and Ashwin Srinivasan. 2012. <a href="#">ILP turns 20 - biography and future challenges</a> . <i>Mach. Learn.</i> , 86(1):3–23.	980
984	Mark A Musen and Johan Van der Lei. 1988. Of brittleness and bottlenecks: Challenges in the creation of pattern-recognition and expert-system models. In <i>Machine Intelligence and Pattern Recognition</i> , volume 7, pages 335–352. Elsevier.	984
989	Terezinha Nunes. 2012. <a href="#">Logical Reasoning and Learning</a> , pages 2066–2069. Springer US, Boston, MA.	989
991	Debjit Paul and Anette Frank. 2020. <a href="#">Social commonsense reasoning with multi-head knowledge attention</a> . In <i>Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020</i> , volume EMNLP 2020 of <i>Findings of ACL</i> , pages 2969–2980. Association for Computational Linguistics.	991
998	Debjit Paul and Anette Frank. 2021. <a href="#">Generating hypothetical events for abductive inference</a> . In <i>Proceedings of *SEM 2021: The Tenth Joint Conference on Lexical and Computational Semantics, *SEM 2021, Online, August 5-6, 2021</i> , pages 67–77. Association for Computational Linguistics.	998



1004	Gabriele Paul. 1993. <a href="#">Approaches to abductive reasoning: an overview</a> . <i>Artif. Intell. Rev.</i> , 7(2):109–152.	1061
1005		1062
1006	Charles Sanders Peirce. 1974. <i>Collected papers of charles sanders peirce</i> , volume 5. Harvard University Press.	1063
1007		1064
1008		1065
1009	Shuofei Qiao, Yixin Ou, Ningyu Zhang, Xiang Chen, Yunzhi Yao, Shumin Deng, Chuanqi Tan, Fei Huang, and Huajun Chen. 2022. <a href="#">Reasoning with language model prompting: A survey</a> . <i>CoRR</i> , abs/2212.09597.	1066
1010		1067
1011		1068
1012		1069
1013	Lianhui Qin, Vered Shwartz, Peter West, Chandra Bhagavatula, Jena D. Hwang, Ronan Le Bras, Antoine Bosselut, and Yejin Choi. 2020. <a href="#">Back to the future: Unsupervised backprop-based decoding for counterfactual and abductive commonsense reasoning</a> . In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020</i> , pages 794–805. Association for Computational Linguistics.	1070
1014		1071
1015		1072
1016		1073
1017		1074
1018		1075
1019		1076
1020		1077
1021		1078
1022	Linlu Qiu and Liwei Jiang. 2023. <a href="#">Phenomenal yet puzzling: Testing inductive reasoning capabilities of language models with hypothesis refinement</a> .	1079
1023		1080
1024		1081
1025	Hanhao Qu, Yu Cao, Jun Gao, Liang Ding, and Ruifeng Xu. 2022. <a href="#">Interpretable proof generation via iterative backward reasoning</a> . In <i>Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022</i> , pages 2968–2981. Association for Computational Linguistics.	1082
1026		1083
1027		1084
1028		1085
1029		1086
1030		1087
1031		1088
1032		1089
1033	Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. <i>J. Mach. Learn. Res.</i> , 21(140):1–67.	1090
1034		1091
1035		1092
1036		1093
1037		1094
1038	Danilo Neves Ribeiro, Shen Wang, Xiaofei Ma, Rui Dong, Xiaokai Wei, Henghui Zhu, Xinchu Chen, Peng Xu, Zhiheng Huang, Andrew O. Arnold, and Dan Roth. 2022. <a href="#">Entailment tree explanations via iterative retrieval-generation reasoner</a> . In <i>Findings of the Association for Computational Linguistics: NAACL 2022, Seattle, WA, United States, July 10-15, 2022</i> , pages 465–475. Association for Computational Linguistics.	1095
1039		1096
1040		1097
1041		1098
1042		1099
1043		1100
1044		1101
1045		1102
1046		1103
1047	Kyle Richardson and Ashish Sabharwal. 2022. <a href="#">Pushing the limits of rule reasoning in transformers through natural language satisfiability</a> . In <i>Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelfth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022</i> , pages 11209–11219. AAAI Press.	1104
1048		1105
1049		1106
1050		1107
1051		1108
1052		1109
1053		1110
1054		1111
1055		1112
1056	Swarnadeep Saha, Sayan Ghosh, Shashank Srivastava, and Mohit Bansal. 2020. <a href="#">Prover: Proof generation for interpretable reasoning over rules</a> . In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020</i> , pages 122–136. Association for Computational Linguistics.	1113
1057		1114
1058		1115
1059		1116
1060		
	Swarnadeep Saha, Prateek Yadav, and Mohit Bansal. 2021. <a href="#">multiprover: Generating multiple proofs for improved interpretability in rule reasoning</a> . In <i>Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021</i> , pages 3662–3677. Association for Computational Linguistics.	
	Merrilee H Salmon. 1989. Introduction to logic and critical thinking.	
	Soumya Sanyal, Zeyi Liao, and Xiang Ren. 2022a. <a href="#">Robustlr: Evaluating robustness to logical perturbation in deductive reasoning</a> . <i>CoRR</i> , abs/2205.12598.	
	Soumya Sanyal, Harman Singh, and Xiang Ren. 2022b. <a href="#">Fairr: Faithful and robust deductive reasoning over natural language</a> . In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022</i> , pages 1075–1093. Association for Computational Linguistics.	
	Min Joon Seo, Hannaneh Hajishirzi, Ali Farhadi, Oren Etzioni, and Clint Malcolm. 2015. <a href="#">Solving geometry problems: Combining text and diagram interpretation</a> . In <i>Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015</i> , pages 1466–1476. The Association for Computational Linguistics.	
	Koustuv Sinha, Shagun Sodhani, Jin Dong, Joelle Pineau, and William L. Hamilton. 2019. <a href="#">CLUTRR: A diagnostic benchmark for inductive reasoning from text</a> . In <i>Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019</i> , pages 4505–4514. Association for Computational Linguistics.	
	Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In <i>Thirty-first AAAI conference on artificial intelligence</i> .	
	Zayne Sprague, Kaj Bostrom, Swarat Chaudhuri, and Greg Durrett. 2022. <a href="#">Natural language deduction with incomplete information</a> . <i>CoRR</i> , abs/2211.00614.	
	Oyvind Tafjord, Bhavana Dalvi, and Peter Clark. 2021. <a href="#">Proofwriter: Generating implications, proofs, and abductive statements over natural language</a> . In <i>Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021</i> , volume ACL/IJCNLP 2021 of <i>Findings of ACL</i> , pages 3621–3634. Association for Computational Linguistics.	
	Oyvind Tafjord, Bhavana Dalvi Mishra, and Peter Clark. 2022. <a href="#">Entailer: Answering questions with</a>	

1117	faithful and truthful chains of reasoning. <i>CoRR</i> , abs/2210.12217.	pages 2369–2380. Association for Computational Linguistics.	1172
1118			1173
1119	Alon Talmor, Oyvind Tafford, Peter Clark, Yoav Goldberg, and Jonathan Berant. 2020. Leap-of-thought: Teaching pre-trained models to systematically reason over implicit knowledge. <i>Advances in Neural Information Processing Systems</i> , 33:20227–20237.	Zonglin Yang, Li Dong, Xinya Du, Hao Cheng, Erik Cambria, Xiaodong Liu, Jianfeng Gao, and Furu Wei. 2022b. <a href="#">Language models as inductive reasoners</a> . <i>CoRR</i> , abs/2212.10923.	1174
1120			1175
1121			1176
1122			1177
1123			
1124	Dung Thai, Dhruv Agarwal, Mudit Chaudhary, Rajarshi Das, Manzil Zaheer, Jay-Yoon Lee, Hannaneh Hajishirzi, and Andrew McCallum. 2023. Machine reading comprehension using case-based reasoning. <i>arXiv preprint arXiv:2305.14815</i> .	Zonglin Yang, Xinya Du, Erik Cambria, and Claire Cardie. 2023a. <a href="#">End-to-end case-based reasoning for commonsense knowledge base completion</a> . In <i>Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics</i> , pages 3509–3522, Dubrovnik, Croatia. Association for Computational Linguistics.	1178
1125			1179
1126			1180
1127			1181
1128			1182
1129	Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. <i>Advances in neural information processing systems</i> , 30.		1183
1130			1184
1131			
1132			1185
1133			1186
1134			1187
1135	Qingyun Wang, Doug Downey, Heng Ji, and Tom Hope. 2023a. <a href="#">Learning to generate novel scientific directions with contextualized literature-based discovery</a> . <i>CoRR</i> , abs/2305.14259.	Zonglin Yang, Xinya Du, Junxian Li, Jie Zheng, Soujanya Poria, and Erik Cambria. 2023b. <a href="#">Large language models for automated open-domain scientific hypotheses discovery</a> . <i>CoRR</i> , abs/2309.02726.	1188
1136			
1137			
1138	Ruocheng Wang, Eric Zelikman, Gabriel Poesia, Yewen Pu, Nick Haber, and Noah D. Goodman. 2023b. <a href="#">Hypothesis search: Inductive reasoning with language models</a> . <i>CoRR</i> , abs/2309.05660.	Zonglin Yang, Xinya Du, Alexander M. Rush, and Claire Cardie. 2020. <a href="#">Improving event duration prediction via time-aware pre-training</a> . In <i>Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020</i> , volume EMNLP 2020 of <i>Findings of ACL</i> , pages 3370–3378. Association for Computational Linguistics.	1189
1139			1190
1140			1191
1141			1192
1142			1193
1143	Wenguan Wang and Yi Yang. 2022. Towards data- and knowledge-driven artificial intelligence: A survey on neuro-symbolic computing. <i>arXiv preprint arXiv:2210.15889</i> .		1194
1144			1195
1145			1196
1146	Wenya Wang and Sinno Jialin Pan. 2022. <a href="#">Deep inductive logic reasoning for multi-hop reading comprehension</a> . In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , ACL 2022, Dublin, Ireland, May 22–27, 2022, pages 4999–5009. Association for Computational Linguistics.	Nathan Young, Qiming Bao, Joshua Bensemann, and Michael Witbrock. 2022. <a href="#">Abductionrules: Training transformers to explain unexpected inputs</a> . In <i>Findings of the Association for Computational Linguistics: ACL 2022, Dublin, Ireland, May 22–27, 2022</i> , pages 218–227. Association for Computational Linguistics.	1197
1147			1198
1148			1199
1149			1200
1150			1201
1151			1202
1152			
1153	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed H. Chi, Quoc Le, and Denny Zhou. 2022. <a href="#">Chain of thought prompting elicits reasoning in large language models</a> . <i>CoRR</i> , abs/2201.11903.	Fei Yu, Hongbo Zhang, and Benyou Wang. 2023. <a href="#">Nature language reasoning, a survey</a> . <i>arXiv preprint arXiv:2303.14725</i> .	1203
1154			1204
1155			1205
1156			
1157	Fangzhi Xu, Qika Lin, Jiawei Han, Tianzhe Zhao, Jun Liu, and Erik Cambria. 2023. Are large language models really good logical reasoners? a comprehensive evaluation from deductive, inductive and abductive views. <i>arXiv preprint arXiv:2306.09841</i> .	Wenting Zhao, Justin T Chiu, Claire Cardie, and Alexander M Rush. 2023. <a href="#">Abductive commonsense reasoning exploiting mutually exclusive explanations</a> . <i>arXiv preprint arXiv:2305.14618</i> .	1206
1158			1207
1159			1208
1160			1209
1161			
1162	Kaiyu Yang, Jia Deng, and Danqi Chen. 2022a. <a href="#">Generating natural language proofs with verifier-guided search</a> . <i>CoRR</i> , abs/2205.12443.	Ruiqi Zhong, Peter Zhang, Steve Li, Jinwoo Ahn, Dan Klein, and Jacob Steinhardt. 2023. <a href="#">Goal driven discovery of distributional differences via language descriptions</a> . <i>CoRR</i> , abs/2302.14233.	1210
1163			1211
1164			1212
1165			1213
1166	Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. <a href="#">Hotpotqa: A dataset for diverse, explainable multi-hop question answering</a> . In <i>Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018</i> , pages 1961–1964. ACM.	Yunchang Zhu, Liang Pang, Yanyan Lan, and Xueqi Cheng. 2020. <a href="#">L2r<sup>2</sup>: Leveraging ranking for abductive reasoning</a> . In <i>Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25–30, 2020</i> , pages 1961–1964. ACM.	1214
1167			1215
1168			1216
1169			1217
1170			1218
1171			1219
			1220
			1221
			1222
			1223
			1224

## A Appendix

### A.1 Relation to Related (NLP) Fields

In this section, we first introduce related NLP fields to general logical reasoning, then introduce fields that are only related to deductive reasoning, inductive reasoning, or abductive reasoning. We hope that this section could be helpful to form a clear shape of LRNLP in NLP.

#### A.1.1 Logical Reasoning

There are some previous works involve the term “logical reasoning”, but do not provide a specification on which sub-type of logical reasoning they involve. In many cases these works are more close to “natural language inference”, which adopts datasets where the data involve a mixture of multiple sub-types of logical reasoning, making it hard to analyze from each sub-type. Therefore we do not include these works in this survey.

**Neuro-Symbolic Computing** Neural-symbolic computing is a hybrid of symbolism and connectionism to exploit advantages from both sides (Wang and Yang, 2022; Cambria et al., 2022). The knowledge representation of its symbolic part basically is a knowledge graph or propositional logic or first-order logic (Wang and Yang, 2022). LRNLP could be seen as a new type of neuro-symbolic in addition to the existing 6 types summarized by Kautz (2022), as its goal and design of methodology are typically symbolic (logical reasoning with knowledge bases), while avoiding any symbolic representation, using (currently pure) neural methods.

**Natural Language Inference** Natural language inference (NLI) is generally considered as the semantic concepts of entailment and contradiction (Bowman et al., 2015). Here logical reasoning tasks can be viewed as special types of NLI focusing on particular reasoning aspects.

**Question Answering** The form of LRNLP looks similar to question answering (QA), however, QA is conducting one-step logical reasoning only when the context provides enough information to answer the question (deductive reasoning), or the answer is a generalization of an argument in context or question (inductive reasoning), or the answer is to provide explanations to the question (abductive reasoning).

**Commonsense Reasoning** Commonsense reasoning (CR) and logical reasoning (LR) are similar in that they both involve “knowledge” and “reasoning”. Compared to LR, CR focuses more on the “knowledge” aspect. Some typical tasks include whether a system has commonsense knowledge (Bosselut et al., 2019; Yang et al., 2020), and whether a system’s answer is commonsense-knowledge-aware (Bisk et al., 2020); LR focuses more on the “reasoning” aspect, e.g., whether a system’s i/o behaviors follow reasoning requirements (Clark et al., 2020).

**Chain of Thoughts** Chain of thoughts (COT) (Wei et al., 2022) is a prompting technique that can elicit the step-by-step reasoning ability of LLMs without finetuning.

COT can potentially be used for each of the three sub-reasoning types of logical reasoning. In fact, for a given (commonsense reasoning) question, some reasoning steps of COT could be deductive, and others can be inductive or abductive. Since the purpose of this paper is to provide a finer analysis on logical reasoning, we do not intentionally cover prompting techniques such as COT.

It is also argued by several modular-based deductive reasoning methods that COT’s reasoning is not causal (Creswell et al., 2022), limited by input size (Ribeiro et al., 2022), and contains unrelated or incorrect steps (Hong et al., 2022; Tafjord et al., 2022).

Overall, it could be interesting to use COT-related methods specifically for deductive, inductive, or abductive reasoning (as opposed to modular-based methods), and it is a less-explored research direction.

#### A.1.2 Deductive Reasoning

**Multi-hop Reasoning** Compared to proof generation, many multi-hop reasoning tasks (Yang et al., 2018; Jiang et al., 2020; Min et al., 2019; Sinha et al., 2019) are much simpler, often being single-branched (Qu et al., 2022), consisting of only 2-3 supporting facts, and are more coarse-grained, involving large chunks of texts such as passages instead of simple, short sentences (Yang et al., 2022a).

Nevertheless, some multi-hop reasoning datasets can be considered as conducting deductive reasoning. For instance, for each data in CLUTRR (Sinha et al., 2019) dataset, a set of facts that can make conclusive support to the target kinship relation is



included in background information as input for each target relation, hence from the philosophical definition (Salmon, 1989), it requires to perform deductive reasoning.

**Mathematical Reasoning** In many mathematical reasoning tasks such as math word problem solving (Koncel-Kedziorski et al., 2015) and geometry problem solving (Seo et al., 2015), the conclusion can be conclusively entailed by the premise. Therefore these tasks belong to deductive reasoning. We do not review math-related papers because we want to focus solely on the challenge of deductive reasoning while mathematical reasoning involves numbers in the text, which introduces additional challenges.

### A.1.3 Inductive Reasoning

**Information Extraction** Information Extraction (IE) is a task of extracting pre-specified types of facts from written texts or speech transcripts, and converting them into structured representations (Ji, 2018). The rule generation task here also extracts rules from facts represented in written texts. The difference is that IE pursues extracting the exact information from existing texts, while inductive reasoning aspires to induce more general rules from existing texts, where the information in rules goes beyond what is exactly stated in the texts.

**Case-based Reasoning** Case-based Reasoning (CBR) is a classic AI subject, whose methods share a general methodology of four steps: retrieve, reuse, revise, and retain (Aamodt and Plaza, 1994). Recently there has been research works devoted to bridging the research of CBR and NLP, by using NLP techniques for CBR challenges (Yang et al., 2023a) and improving NLP tasks with CBR methodologies (Das et al., 2021, 2022; Yang et al., 2023a; Thai et al., 2023). CBR could be seen as a type of analogical reasoning (Kolodner, 1997), and analogical reasoning belongs to inductive reasoning (Salmon, 1989). However, CBR is a different inductive reasoning type than the “generalization” process (from facts to rules) described in Flach and Kakas (2000), but more on the general description on inductive reasoning (Salmon, 1989) that premises cannot conclusively provide support to the conclusion.

### A.1.4 Abductive Reasoning

**Causal Reasoning** In logic research, causal reasoning aims at an epistemological problem of estab-

lishing precise causal relationships between causes and effects. It is generally considered a form of inductive reasoning (Goertzel et al., 2011), since inductive reasoning is to derive rules that lead from one to another. When the focus is to derive possible causes from effects, the problem belongs to abductive reasoning (Goertzel et al., 2011).

## A.2 Full Details About the Definition and Categorization of Logical Reasoning

There are many subjects related to logical reasoning, including philosophy, logic, and AI. Among them, the definition and categorization aspects of logical reasoning are handled by philosophy research. However, debate exists in philosophy research on the categorization of logical reasoning.

One group believes that every argument can be classified as either deduction argument, inductive argument, or fallacy (Salmon, 1989). Without considering fallacy, given that an argument consists of premises and a conclusion, when the premises can conclusively provide support to the conclusion (which means that if the premises of the argument were all true, it would be impossible for the conclusion of the argument to be false), this argument is a deductive argument. Conversely, when the premises can not conclusively provide support to the conclusion, the argument is inductive.

The other group has the same definition of deductive reasoning, but they believe that further categorization of non-deductive reasoning is necessary. Without considering fallacy, they believe in a trichotomy of deductive, inductive, and abductive reasoning (Peirce, 1974). However, even for the second group, the definition and difference between inductive and abductive reasoning are also controversy (Flach and Kakas, 2000).

Nevertheless, Console and Saitta (2000) argue that from the utility perspective of AI, a distinction between inductive and abductive reasoning is possible: both inductive and abductive reasoning provide explanations about the world but their explanations differ in the degree of generality. For instance, an inductive hypothesis allows the validity of properties, observed on a set of individuals, to be generalized to other individuals not in the observations, whereas an abductive one allows unobserved properties to be applied to observed individuals.

Considering that inductive and abductive reasoning can be distinctive enough when formulated in NLP, in this paper, we adopt the second group, particularly Console and Saitta (2000)’s view of

definition and categorization of logical reasoning.

Specifically, the difference between inductive and abductive reasoning is that, both inductive and abductive reasoning provide explanations about the world but their explanations differ in the degree of generality.

For instance, an inductive hypothesis allows the validity of properties, observed on a set of individuals, to be generalized to other individuals not in the observations, whereas an abductive one allows unobserved properties to be applied to observed individuals.

The distinction between inductive and abductive hypotheses strictly parallels the dichotomy *extension* vs. *intension*, or *generality* vs. *informativeness*. In other words, an inductive hypothesis extends or generalizes to unobserved individuals, while an abductive one provides more specific information (e.g., unobserved properties) about existing specific individuals.

For example, if a white ball is found in a bag, inductive reasoning might lead to the conclusion that “all balls in this bag are white”, while abductive reasoning might lead to the conclusion that “someone put the white ball into this bag”.

In this example, the inductive hypothesis generalizes the property of the existing individual (the white ball) to unobserved individuals (other not-seen balls in the bag), while the abductive hypothesis provides more specific information about the current individual (who put this ball to the bag).

To summarize in simple words, in common situations, pure inductive reasoning is to only provide (usually sample to population) generalizations, while pure abductive reasoning is to only provide specific explanations.

Overall, even in the philosophical literature (which takes charge of the research on the definition of logical reasoning), a clear definition for all three types of logical reasoning is rare, but more on the description of the difference between types of logical reasoning (since a clear definition is still under debate). The difference can be illustrated does not mean a precise definition can be given. Nevertheless, considering the above-discussed philosophical literature, we try our best to give a definition below for a more straightforward understanding:

Given an argument consisting of premises and a conclusion, we define the sub-type of logical reasoning it involves below:

*Definition for deductive reasoning:* the premises can conclusively provide support for the conclu-

sion, i.e. if the premises are all true, it would be impossible for the conclusion to be false.

*Definition for inductive reasoning:* the premises cannot conclusively provide support for the conclusion, since the conclusion generalizes existing information in premises to new knowledge, which has a wider applicable scope than those in premises.

*Definition for abductive reasoning:* the premises cannot conclusively provide support for the conclusion, since the conclusion contains more specific information over the premises (most commonly used as generating most probable explanations).

Please note that according to [Console and Saitta \(2000\)](#), inductive reasoning and abductive reasoning are not exclusive to each other, i.e., inductive reasoning and abductive reasoning overlap with each other.

### A.3 Why We Choose Definition in Section 2

Firstly, some other definitions (e.g., from Pieces) are not in conflict with the one we adopted. Secondly, other definitions lack a clear boundary between different types of reasoning, while our adopted definitions clearly delineate such boundaries (e.g., general vs. specific for inductive and abductive reasoning).

To elaborate why there’s no contradiction, specifically, Pierce’s definition is “inference to the most plausible explanation for incomplete observations”. Here “explanation” refers to not guaranteed and specific information. An example about the discussion on “specific” can be found in §A.10. We use definition in §2 because other definitions lack a clear boundary between deductive, inductive, and abductive reasoning, while our adopted definitions clearly delineate such boundaries (e.g., general vs. specific for inductive and abductive reasoning; guaranteed vs. not guaranteed for deductive and the remaining two reasoning).

### A.4 Related Surveys on Reasoning

[Huang and Chang \(2022\)](#); [Qiao et al. \(2022\)](#) mainly reviews the prompting techniques for LLMs, but do not focus on papers that specialized on logical reasoning (the coverage of the two fields are quite different). [Yu et al. \(2023\)](#) also review papers related to reasoning. However, (1) they do not focus on logical reasoning, and do not organize their survey based on the three sub-types of logical reasoning. Particularly, only a small section discusses this topic; (2) their definition on the deductive, inductive, and abductive reasoning lacks

a philosophy foundation (no reference), and is inaccurate. Particularly, from their definition, it is unclear on the difference between inductive and abductive reasoning (e.g., what is the difference between the “more general rule” and “best explanation”? A more general rule, such as Newton’s Law, of course can serve as the best explanation about phenomena related to object movement). In the contrary, this survey’s definition is based on philosophy literature (Console and Saitta, 2000), and our definition can clearly differ between inductive and abductive reasoning. The difference lies in that inductive reasoning is about “general”, and abductive reasoning is about “specific”, while a specific conclusion, such as “it must have rained since the lawn is wet”, is commonly used as “the best explanation”. But the point of abductive reasoning is about “specific”, but not “explanation”, since inductive reasoning can also provide explanation (Flach and Kakas, 2000). We illustrate in §2 that there has been various forms of definition for the three reasoning types in the thousands of years of development of the philosophy research. They are not aware of it and pick pieces of definition without a systematic view, which is the main reason for their confusing definition. We also provide a detailed discussion about the categorization of logical reasoning from a philosophy perspective in § A.2.

Yu et al. (2023) do not stress and organize the survey from the three sub-types of logical reasoning. Xu et al. (2023) provides a comprehensive evaluation of the logical reasoning ability of LLMs. They are not to provide a survey but to use LLMs on the existing logical reasoning datasets.

Lakoff (1970); McCarthy (1990) are the first few works to take a close look at the connection between logical reasoning and natural language. Dagan et al. (2005) proposed evaluating logical reasoning through the comparison of two natural language texts. MacCartney and Manning (2014) apply logical reasoning in natural language inference through iterative editing of natural language.

## A.5 Other Inductive Reasoning Papers

**Implicit Rule Verification** Misra et al. (2022) analyze language model’s ability to generalize novel property knowledge (has sesamoid bones) from concept(s) (robins) to others (sparrows, canaries). As illustrated in §A.10, they analyze the language models’ ability to classify a new fact (but not a rule) as correct or not, given facts. It could be seen

that the correctness of a rule is implicitly predicted by testing multiple facts entailed by the rule.

**Symbolic Rule Generation** Wang and Pan (2022) propose attentive memories with novel differentiable logic operators to induce symbolic rules from texts.

## A.6 Research Trend in the Three Sub-Types of Logical Reasoning

Out of the three reasoning types, deductive reasoning has drawn the most research attention, and has the most abundant of works, especially in 2022. Abductive reasoning has drawn much attention in 2020 and 2021 but has few works in 2022 and 2023. Inductive reasoning is only proposed at the end of 2022, having the least number of works. However, inductive reasoning has attracted much attention since the second half year of 2023.

Two main reasons for the abundance of works in the deductive reasoning domain could be that (1) more challenging benchmarks have been constructed during the last few years, and (2) deductive reasoning could be one of the most commonly used reasoning types in common life. We think the main reason for the little attention drawn to abductive reasoning in recent years is that the benchmarks for abductive reasoning are relatively old and less challenging for LLMs. Inductive reasoning could be a promising research topic since there have been few works in the domain, and it involves very challenging tasks such as proposing new scientific findings.

In general, there has been no framework which is proposed to address all three reasoning domains. However, LLMs generally can exhibit all three reasoning abilities to some extent. It would be interesting for future works to analyze the effect of the pretraining method and scale of LLM on the three reasoning abilities.

## A.7 Relation Between LRNLP and neuro-symbolic

A large proportion of recent papers on deductive reasoning and abductive reasoning leverage a natural language-based knowledge base, and reason over retrieved knowledge from the knowledge base to reach a certain goal (Tafjord et al., 2021; Liang et al., 2021; Qu et al., 2022; Ribeiro et al., 2022; Creswell et al., 2022; Sanyal et al., 2022b; Hong et al., 2022; Bostrom et al., 2022; Yang et al., 2022a; Tafjord et al., 2022; Dalvi et al., 2022). This pattern is very similar to the methodology design



of neuro-symbolic methods, which is to retrieve symbolic knowledge and reason over the retrieved symbolic knowledge. The main difference is that LRNLP adopts natural language as knowledge representation but not symbolic knowledge. Because of the similarity in the methodology design, we consider that LRNLP could be seen as a type of neuro-symbolic methods but without many disadvantages of symbolic representation such as symbolic knowledge acquisition and scalability.

In addition, due to the high similarity in the methodology design to neuro-symbolic, LRNLP also shares some advantages with neuro-symbolic such as explainability. The reason is that the iterative retrieving and reasoning will make the decision-making process more interpretable on the intermediate reasoning steps, and which knowledge is used for each reasoning step.

## A.8 Experiments Summarization

In this section, we summarize the experiment results of an important and literature-abundant task.

Until now there has been only one or two papers working on inductive reasoning. Methods for abductive reasoning generally leverage different resources (such as multi-task, additional knowledge resources, and ancillary loss) and lack a progressive relationship between each other, therefore are less comparable. Currently, the *ProofGeneration* task in deductive reasoning is the most literature-abundant, and methods for this task have progressive relationships with each other. Therefore here we mainly summarize results and analyze for the *ProofGeneration* task.

Table 5 shows the summarized experiment results. We select the most widely used tasks to display their performance. Among the task, the setting of ParaRules is trained on D3 (D\* dataset with depth 3) and tested on the ParaRules test set; the setting of Birds-Electricity is trained on D5 (D\* dataset with depth 5) and tested on bird-electricity set; setting for EntailmentBank is the task 3 which uses full corpus as input (so that many distractors exist in input); setting for OBQA and QuaRTz are zero-shot setting while model pre-trained on another dataset (EntailmentBank).

Among the methods, Creswell et al. (2022) and Bostrom et al. (2022) design unique metrics using EntailmentBank dataset. Sprague et al. (2022) focus on a unique task (proof generation task with incomplete information), therefore we do not list their experiments results in the table. Specifically,

Creswell et al. (2022) work on metric of accuracy. Bostrom et al. (2022) work on metrics of *Goal%*, *#Steps*, where *Goal%* measures the number of valid goals reached by each system, and *#Steps* measures the number of steps expanded before reaching a valid goal. Sprague et al. (2022) work on “proof generation task with incomplete information”, which naturally performs worse than the “proof generation task”.

Overall methods for proof generation tasks tend to use different datasets for evaluation, making them less comparable.

## A.9 Model Structure, Pretraining Data Used, and Experiment Settings

Table 6 shows a collection of model structure, pre-training data usage, and experiment settings for methods in hypothesis classification and proof generation tasks.

In general, RoBERTa-large and T5-11B are the most adopted base models.

## A.10 Meaning of “More General” Required by Inductive Reasoning

This section is collected from Yang et al. (2022b)’s appendix, to help illustrate inductive reasoning.

Given an argument consisting of a premise and a conclusion, if the conclusion involves new information that is not covered by the premise and can not be conclusively entailed by the premise, the argument is an inductive argument (Salmon, 1989).

When the conclusion has a larger scope of information coverage than the premise, and can entail the premise, it can be said that the conclusion is “more general” to the premise (Yang et al., 2022b). In this case, we termed the premise as a “fact”, and the conclusion as a “rule”; When the conclusion contains new pieces of information and cannot entail the premise, as defined by Salmon (1989), the argument is still an inductive argument. But in this case, we termed the premise as a “fact”, and the conclusion as another “fact”.

For instance, if facts that are about cats and dogs are good accompaniment of humans, then some examples of a “more general” rule can be (1) mammals are good accompaniment of humans, or (2) domesticated animals are good accompaniment of humans, or (3) animals with four legs are good accompaniment of human.

In these examples, the rules cover a larger scope than the facts (e.g., mammals compared to cats;

Methods	ParaRules	Birds-Electricity	EntailmentBank (Task 3)							OBQA Accuracy	QuaRTz Accuracy
	Full Accuracy (FA)	Full Accuracy (FA)	Leaves F1	Leaves All-Cor.	Steps F1	Steps All-Cor.	Intermediates F1	Intermediates All-Cor.	Overall All-Correct		
PProver	95.1	80.5	-	-	-	-	-	-	-	-	-
multiPProver	94.5	81.8	-	-	-	-	-	-	-	-	-
EntailmentWriter	-	-	39.7	3.8	7.8	2.9	36.4	13.2	2.9	-	-
ProofWriter	98.5	97.0	-	-	-	-	-	-	-	-	-
EVR	-	63.1	-	-	-	-	-	-	-	-	-
IBR	95.7	93.5	-	-	-	-	-	-	-	-	-
IRGR	-	-	45.6	12.1	16.3	11.8	38.8	36.5	11.8	-	-
Selection-Inference	-	-	-	-	-	-	-	-	-	-	-
FaiRR	98.6	-	-	-	-	-	-	-	-	-	-
MetGen	-	-	34.8	8.7	9.8	8.6	36.7	20.4	8.6	-	-
SCSearch	-	-	-	-	-	-	-	-	-	-	-
ADGV	-	-	-	-	-	-	-	-	-	-	-
NLProofS	-	-	43.2	8.2	11.2	6.9	42.9	17.3	6.9	-	-
Entailer	-	-	-	-	-	-	-	-	-	76.8	74.3
Teachme	-	-	-	-	-	-	-	-	-	77.0	75.9

Table 5: Proof Generation Task Results.

	Model usage	Second-round pretraining data used before finetuning	Experiment setting	Task
RuleTakers	RoBERTa-large (355M)	RACE	finetuning	0
Leap-of-Thought	RoBERTa-large	-	finetuning	0
FOLIO	[BERT and RoBERTa]-[base and large]	-	finetuning & few-shot	0
PProver	RoBERTa-large	-	finetuning	1
multiPProver	RoBERTa-large	-	finetuning	1
EntailmentWriter	T5-11B	-	finetuning	1
ProofWriter	T5-11B	-	finetuning	1
EVR	T5-small (60M)	-	finetuning	1
IBR	RoBERTa-large	-	finetuning	1
IRGR	T5-large (770M) for generation model; MPNet-base-v2 for retriever	-	finetuning	1
Selection-Inference	Chinchilla-7B	-	finetuning & few-shot	1
FaiRR	RoBERTa-large for fact and rule selectors; T5-large for knowledge composer	-	finetuning	1
MetGen	T5-large for deduction and abduction models; ALBERT-xxlarge-v2 for controller	-	finetuning	1
SCSearch	T5-large for deduction model; DeBERTa-v3-large (350M) for goal entailment model	DeBERTa-v3-large on MNLI	finetuning	1
ADGV	T5-large for deduction model; T5-3B for abduction model; DeBERTa-v3-large for goal entailment model	DeBERTa-v3-large on MNLI	finetuning	1
NLProofS	T5-large for deduction model; RoBERTa-large for verifier	-	finetuning	1
Entailer	T5-11B	-	finetuning	1
Teachme	T5-11B	-	human interaction	1

Table 6: A collection of model usage, (second-round) pretraining data usage, and experiment setting for methods in hypothesis classification (denoted as “0” in “Task” column) and proof generation (denoted as “1” in “Task” column) tasks.

domesticated animals compared to cats), and therefore the rules are “more general” than the facts.

“More general” means not only about finding higher taxonomic rank, but can be in unlimited forms. For instance, if the fact is about the Sun rises and falls every day, then some examples of a “more general” rule can be (1) the Earth is the king of the universe or (2) the Earth is rotating itself.

Both rule examples are “more general” than the given fact, since the rule can entail not only the given fact, but also other not mentioned facts such as the observable movements of the other stars in the Milky Way.

## A.11 Other Challenges and Possible Future Directions

**Reliable Rule Generation** Currently, the rule generation method in inductive reasoning relies on out-of-box LLMs, since a finetuned rule generation model could be restricted in a domain. The annotation of an inductive reasoning dataset should only be done by experts and is very time consuming (Yang et al., 2022b). Given the two restrictions, how to improve the quality of generated rules given related facts could be a challenging open problem.

**Reliable Explanation Generation** Abduction is a form of non-monotonic reasoning (Paul, 1993), and potentially has a large search space of conclusions given premises. Therefore, how to generate more (all) reasonable explanations can be challenging (Bhagavatula et al., 2020).

**Building Larger Benchmarks** For complicated reasoning tasks especially in realistic and natural language settings, usually experts are needed for annotation, and the process is very time-consuming (Dalvi et al., 2021; Sprague et al., 2022; Yang et al., 2022b). Therefore it can be challenging to construct significantly larger benchmarks.

**Understanding the Internal Mechanism of LLMs for Reasoning** Until now research works only focused on investigating whether the input/output behaviors of LLMs can be used to simulate a reasoner (Clark et al., 2020) or complete reasoning tasks. However, it is still a challenging open research question to understand the internal mechanism of LLMs for reasoning.

**Reliable Verifier** Most current verifiers (refiners) relies on the internal beliefs of LLMs to select (improve) from generated rules and mitigate halluci-

nation. It’s doubtful whether LLMs have obtained the necessary knowledge.