IMAP: A MIND MAPPING CONSTRUCT TO ENHANCE INDUCTIVE REASONING IN GENERATIVE MODEL

Anonymous authors

000

001

002 003 004

010 011

012

013

014

016

017

018

019

021

023

025

026

027

028

029

031

032

034

035

037

040

041

042

043

044

045

046

047

048

051

052

Paper under double-blind review

ABSTRACT

Inductive reasoning is crucial in human thinking, allowing us to distill universal laws from limited samples. However, incorporating inductive reasoning has not been studied enough in the field of artificial intelligence, especially in the application of large-scale language models, limiting the ability of models to abstract broad rules and trends from limited data. We introduce inductive thinking into generative models, designing rigorous rules to compare generated results with real ones, and verify its effectiveness in improving generation. To achieve this, we developed IMap (Intellectual Mapping based on Reinforcement Learning), which integrates the inductive thinking paradigm to improve the model's inference capabilities. We designed a thinking data structure based on the inductive paradigm, consisting of four core elements: COTs, Cases, Patterns, and Reasonability. We also propose an algorithm, the RL-Paradigm model (RLP), to acquire new thinking paradigms. By using figurative inductive thinking as input cues, we successfully guided multiple large models to generate an average of 270 results. Comparative experiments show that input cues combined with inductive thinking perform well in most models, significantly improving the generation results. We conducted a comprehensive evaluation of RLP against other models using BLEU, Bert-score, and Jina-score metrics. The results show that RLP significantly outperforms other models in several areas. We unlocked the generative potential of inductive thinking paradigms, developed reusable thinking data maps, and designed RLP, a generative model specialized for unknown paradigms. This innovation is expected to advance the generative capabilities of LLMs and offer insights for interdisciplinary research in brain sciences. Our code and data and trained models are publicly available from https://github.com/yzqrtop/RLP-inductive-LLM.

1 Introduction

Inductive reasoning, which involves generalizing Patterns and rules from limited rejecteds, is fundamental to human cognition and crucial in artificial intelligence (AI) (Latona et al., 2024). As large language models (LLMs) become essential for applications like conversational agents, content generation, and problem-solving systems, their ability to perform inductive reasoning is under greater scrutiny (Heiding et al., 2023; Dam et al., 2024). This study builds on previous research, tackles limitations in existing datasets and models, and introduces a new approach to improve the inductive reasoning abilities of LLMs (Wu et al., 2024; Luong et al., 2024).

Several researchers have advanced LLMs' reasoning capabilities. Step-by-step approaches, such as COTs reasoning, have enabled LLMs to address multi-step reasoning tasks by guiding their logic through each step (Bai et al., 2024). Synthetic data generation has also been explored to enhance LLM training. Although the aforementioned methods have advanced LLM reasoning, they exhibit several limitations: **Limited Dataset Diversity:** Many approaches rely on narrow datasets like GSM8K (Cobbe et al., 2021b), which focus on specific domains, such as mathematical reasoning. This lack of structural diversity limits the models' ability to generalize across diverse tasks. **Sparse and Static Annotations:** CoT-based methods typically use a single annotated reasoning path per question (Zhang et al., 2024b). We believe that this static approach fails to capture the multiple valid reasoning paths that may exist for a given problem, limiting the model's reasoning flexibility. **Scalability and Generalization:** Reinforcement learning frameworks like REFT (Long et al., 2024) and RAFT (Zhang et al., 2024a) show promise but face scalability issues due to sparse rewards

and reliance on domain-specific datasets. Additionally, these models often struggle to generalize effectively to unseen scenarios.

This study builds on prior research, addresses limitations in existing datasets and models, and introduces an inductive thinking paradigm to enhance the inductive reasoning capabilities of LLMs. We validated the effectiveness of this paradigm through rigorous benchmarking on the BBH dataset. To this end, we designed an inductive thinking paradigm data graph to provide fast, reliable thinking data for reuse. To complement the thinking paradigm graph, we propose the RL paradigm model (RLP), a reinforcement learning-based approach utilizing proximal policy optimization (PPO). Unlike static reasoning methods, the RLP dynamically explores and generates new reasoning paths using feedback from datasets. We use BLEU (Papineni et al., 2002), Bert-score (Zhang et al., 2020), and Jina-score metrics to evaluate the quality of inference outputs generated by RLP.

In summary, our contributions are as follows. 1. We introduce inductive thinking into the generation of Large Language Models (LLMs) and validate its impact using the **BBH** (**Big Bench Hard**) benchmark dataset. 2. We have constructed a **reinforcement learning-driven mind map** (**IMap**), integrating new thinking paradigms and providing structured knowledge support for subsequent model generation tasks. 3. We propose the **RL Paradigm model** (**RLP**), which generates new thinking paradigms and efficiently expands the application of inductive reasoning. 4. By constructing a reusable **mental data graph** and designing generative models for new thinking paradigms, this study advances LLM development in language understanding and provides new perspectives for interdisciplinary research, including neuroscience and cognitive psychology, with significant application potential.

2 RELATED WORK

2.1 THE POTENTIAL OF LLMS IN COMPLEX REASONING SCENARIOS

Large Language Models (LLMs) have shown remarkable abilities in understanding and generating human-like text (Heiding et al., 2023). Recent research has aimed at enhancing their reasoning abilities to handle complex tasks. Techniques like Chain-of-Thought (CoT) prompting encourage models to articulate intermediate reasoning steps, improving problem-solving performance (Devlin et al., 2018). For example, the *Cumulative Reasoning* approach uses LLMs iteratively to mirror human thought processes, breaking tasks into manageable components and leveraging prior propositions for effective composition. These advancements demonstrate the potential of LLMs to perform sophisticated reasoning across diverse domains.

2.2 AN EXPLORATION OF LLMs Fine-Tuning Based On Reinforcement Learning

This study applies Proximal Policy Optimization (PPO) (Schulman et al., 2017) for natural language processing to align human preferences (Ouyang et al., 2022). Since then, several training algorithms have been introduced to improve alignment efficiency, including Direct Preference Optimization (DPO) (Rafailov et al., 2024), Identity Preference Optimization (IPO) (Azar et al., 2023), and Kahneman-Tversky Optimization (KTO) (Ethayarajh et al., 2024). Unlike alignment-focused methods, our goal is to adopt reinforcement learning as a fine-tuning approach to enhance performance beyond conventional supervised fine-tuning techniques.

2.3 INDUCTIVE THINKING PARADIGM INTERPRETATION

Inductive thinking is a logical method of forming general Patterns or predictions by observing specific examples (Binti Misrom et al., 2020). It enables researchers to develop theories and chosen from rejecteds and empirical evidence, generating new knowledge. This method is particularly suited for qualitative research, forming general principles or theories by analyzing specific examples (Peltonen, 2023; Mott & Bullock, 2015), which aids in exploring phenomena and generating new insights. By proposing questions that foster higher-order thinking skills (Fabrizio et al., 2014), methods for cultivating inductive thinking, and techniques to improve its efficiency (Hammer, 2011), the ability to solve practical problems in various fields can be enhanced.

2.4 Dataset BBH For Quantitative Assessment Of Language Models

BIG-Bench is a collaborative benchmark designed to quantitatively assess the strengths and weaknesses of language models (Srivastava et al., 2022). It includes over 200 diverse text-based tasks across categories such as traditional NLP, mathematics, commonsense reasoning, and question-answering. The remaining **23 subtasks** form our curated benchmark, BIG-Bench Hard (BBH). This includes two tasks: **Logical Deduction** and **Shuffled Objects**, each with three subtasks. For all tasks in BBH, except three, we selected a random subset of 250 evaluation examples, totaling 6,511 examples in the benchmark.

3 VALIDATION AND CONSTRUCTION OF IMAP

The IMap building involves three steps. First, we use **inductive thinking validation** to assess the accuracy of inductive thinking in various models (see Section 3.1). Next, we design an inductive thinking data graph and define five graph structures. Finally, we define four generation tasks and introduce the RLP method for generating new thinking paradigms. **As shown in the Figure 1**.

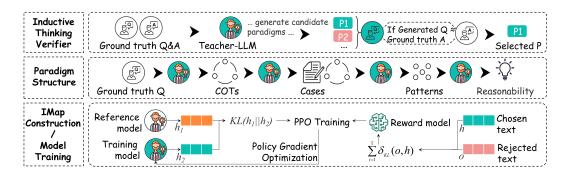


Figure 1: IMap consists of three stages: (1) Inductive thinking verifier, (2) Paradigm Structure and (3) Imap Construction. First, we use **inductive thinking validation** to assess the accuracy of inductive thinking in various models (see Chapter 3.1). Next, we design an inductive thinking data graph and define five graph structures. Finally, we define four generation tasks and introduce the RLP method for generating new thinking paradigms.

3.1 INDUCTIVE THINKING VALIDATION

We randomly selected 10 question-answer pairs from 23 subtasks in the BBH dataset, which include tasks related to causal judgment, data parsing, time sorting, and more. We used a variety of large models at different scales to test comparisons, and then selected the appropriate large language model to generate the following graphical structure of the data: COTs, Cases, Patterns, and Reasonability. These data were then used to create a prompt instruction set. Next, we integrate the questions, answers, and prompts that generate paradigms. Our cognitive generation process is mainly divided into two types of prompts: attribute prompts, which are used to induce cognitive attributes from examples, and system prompts, which are used to answer new questions using the generated attributes. Please refer to Appendix F for details. The teacher language model is then applied to numerous candidate paradigms. We then determine whether the language model's answer includes these paradigms; if it does, it is considered correct. Finally, we continue asking the language model questions for further exploration. We use a rigorous answer comparison strategy: when the model's answer is inconsistent with the correct one or is ambiguous, we consider it incorrect. As shown in **Table 1**, **inductive thinking performs well in most models**. Additionally, we have carefully selected a batch of high-quality inductive thinking datasets for use in the Section 3.3.2.

3.2 IMAP STRUCTURE DEFINITION

The design of the IMap data structure is rooted in inductive thinking theory, a framework that holds inductive reasoning progresses from specific facts to general conclusions. Here, we emphasize that

Table 1: Compare the performance of the inductive thinking paradigm with other cueing methods on generative models of all scales. Compared to other cueing strategies, the inductive thinking paradigm exhibits differential impact across models of various sizes. When compared to models such as yi-34B-chat, the paradigm's performance gains appear relatively limited. In contrast, the inductive thinking paradigm significantly enhances performance on the Llama-3.2-1B and Qwen Turbo models. Taken together, the inductive thinking paradigm demonstrates its superior efficacy by performing optimally in terms of overall performance, outperforming the average by 9.83%.

model/method	Llama-3.2-1B	Llama-3.2-3B	Llama-3.1-8B	yi-34B-chat	Qwen Plus	Qwen Turbo	Avg.
zero-shot (Kojima et al., 2023)	$32.33_{(-6.78)}\downarrow$	$26.06_{(-7.88)}\downarrow$	$47.62_{(0.0)}\downarrow$	$21.11_{(-8.41)}\downarrow$	$35.67_{(-28.38)}\downarrow$	$22.10_{(-27.53)}\downarrow$	$30.82_{(-13.16)}\downarrow$
Step-by-Step (Hsieh et al., 2023)	$35.65_{(-3.46)}\downarrow$	$30.04_{(-3.9)}\downarrow$	$40.48_{(-7.14)}\downarrow$	$24.57_{(-4.95)}\downarrow$	$58.24_{(-5.81)}\downarrow$	$28.29_{(-21.34)}\downarrow$	$36.21_{(-7.77)}\downarrow$
Question Aug(Li et al., 2024)	44.53 _(5.42) ↑	$35.21_{(1.27)}\uparrow$	$40.48_{(-7.14)}\downarrow$	$29.92_{(0.4)}\uparrow$	76.87 _(12.82) ↑	70.17 (20.54)	$49.53_{(5.55)}\uparrow$
reverse (Chen et al., 2024)	$39.58_{(0.47)}\downarrow$	39.50 _(5.56) ↑	$47.62_{(0.0)}\uparrow$	$34.81_{(5.29)}\uparrow$	$74.43_{(10.38)}\uparrow$	$61.33_{(11.7)}\uparrow$	$49.54_{(5.56)}\uparrow$
TOT(Tree of Thought)	$35.69_{(4.39)}\uparrow$	$40.69_{(4.95)}\uparrow$	$56.21_{(14.29)}\uparrow$	$36.27_{(7.69)}\uparrow$	$69.15_{(11)}\uparrow$	$68.41_{(16.7)}\uparrow$	54.31 _(9.83) ↑
inductive	$43.50_{(4.39)}\uparrow$	$38.89_{(4.95)}\uparrow$	61.91 _(14.29) ↑	37.21 _(7.69) ↑	$75.05_{(11)}\uparrow$	$66.30_{(16.7)}\uparrow$	53.81 _(9.83) ↑
Avg.	39.11	33.94	47.62	29.52	64.05	49.63	43.98

explicitly delineating these reasoning steps confers three core advantages: heightened interpretability, improved stability, and stronger generalization capabilities of the reasoning process. This is particularly critical for tackling complex problems that demand abstraction and inductive thinking. IMap captures the core units of inductive question-answering logic by defining six key cognitive nodes, denoted as $IMap = \{Q, A, Co, Ca, P, R\}$, where Q is question, A is answer, Co is COTs, Ca is Cases, P is Patterns, and R is Reasonability. It also offers a comprehensive data structure model for studying inductive reasoning. The core advantage of explicitly defining these steps lies in **enhancing the interpretability, stability, and generalization capabilities** of the reasoning process, which is crucial for addressing complex problems that require abstraction and induction.

Question (**Q**): Problems from the 23 tasks in the BBH dataset (e.g., boolean expressions task: "not (True) and (True) is"). **Answer** (**A**): Corresponding answers to Q (e.g., "False" for the above boolean expressions question). Extended parsing was conducted based on Q and A. **COTs** (**Co**): Analytical content for Q, providing detailed descriptions of Q's problem context and potential insights. **Cases** (**Ca**): Specific conditions or relevant examples extracted from COTs, used to analogize and interpret the ideas/principles in COTs. **Patterns** (**P**): Fundamental principles capturing commonalities between COTs and Cases. These principles precisely encode shared information consistent with both components, reflecting their common structures; P enables extraction of the underlying logical framework of COTs and Cases, facilitating efficient knowledge summarization and generalization. **Reasonability** (**R**): Validation of the rationality of P-described cognitive patterns, focusing on adherence to problem-specific rules (e.g., for the pattern "teamwork improves work efficiency," R verifies universal validity via analyzing Ca (successful project teams) and Co (contexts like team role division)).

These nodes form a structured representation as follows: $Q \Rightarrow Co, Co \Rightarrow Ca, Ca + Co \Rightarrow P, P \Rightarrow R, Q + Co + Ca + P \Rightarrow A$. We incorporated inductive thinking into the generation process of Large Language Models (LLMs) and validated its effectiveness through BBH benchmark testing, resulting in a 7% improvement in model accuracy. Simultaneously, a reinforcement learning-driven thinking graph (IMap) was developed, combining the theory of new thinking paradigms. This graph incorporates elements such as "COTs", "Cases", "Patterns", and "Reasonability".

3.3 IMAP GENERATOR

In IMap, the process of generating cognitive chains is decomposed into a series of ordered tasks: the generation of COTs, Cases, Patterns, and Reasonability. As shown in the Figure 2, we perform these tasks sequentially and connect them in a question-answer format $(Q \Rightarrow A)$ to form a complete cognitive chain. By further refining these processes, we define the following four core inference generation tasks:

COTs generation: Constructs a logical chain (termed a "problem description inference chain") for a given task, decomposing the problem into logically connected subproblems via key reasoning steps. This provides a structured framework for subsequent Cases and Patterns generation. **Cases generation**: Generates specific, diverse Cases to support COTs steps, following the rule "reasoning chain (COTs) + problem description ⇒ supportive case." Cases must align with COTs to validate the inference chain from multiple perspectives. **Patterns generation**: Abstracts commonalities

between COTs and Cases (rule: "COTs + Cases \Rightarrow Patterns") to produce logical cognitive Patterns. These reveal deep logical relationships between COTs and Cases, offering universal guidance for broader reasoning scenarios. **Reasonability generation**: Verifies Pattern rationality (rule: "Patterns \Rightarrow Reasonability") by checking alignment with problem-specific logical rules and Case support. Logical validation of Patterns determines their applicability and optimizes inference chain accuracy.

3.3.1 TASK DEFINITIONS

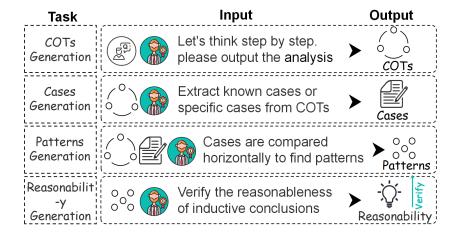


Figure 2: IMap structure generation task definition diagram. We propose a cognitive chain generation task for constructing IMap, which involves a series of ordered tasks: COTs generation, Cases generation, Patterns generation, and Reasonability generation.

By linking the four tasks in a pipeline, IMap restores the complete cognitive chain: " $Question \Rightarrow COTs \Rightarrow Cases \Rightarrow Patterns \Rightarrow Reasonability$." This process ensures the logical integrity of the cognitive chain from question to answer and enables the model to dynamically adjust to meet the reasoning needs of different tasks.

3.3.2 RLP

During supervised training, the model improves its performance by generating results structurally similar to the reference reality. However, high structural similarity alone does not guarantee logical explanatory power, particularly when dealing with unseen samples, where the limitations of this approach become more evident. To enhance task generation models, we introduced reinforcement learning methods (Ziegler et al., 2020) and combined them with thought paradigm graph data to create feedback reward functions. Finally, we developed a generation model for unknown inductive thinking paradigms.

Specifically, for the four generation tasks, we input real problems into the training model, which generates both trained and baseline representations. The KL divergence between the two is input into PPO. Additionally, we introduce adaptive controller the controller's logic depends on the KL divergence for an entire batch of generated sequences. This requires aggregating the per-token penalties. First, the per-token values δ_{KL} are summed over the length t of a single generated sequence y to obtain the sequence-level KL divergence, KL_{sed} :

$$KL_{\text{seq}}(y) = \sum_{t=1}^{T} \delta_{KL}(o, h) = \sum_{t=1}^{T} (\log \pi_{\theta}(h \mid o) - \log \pi_{\text{ref}}(h \mid o))$$
 (1)

In this framework, h and o represent the token sequences for the chosen and rejected items, respectively, while H and O denote the corresponding sets of chosen and rejected items. This summation represents the total accumulated divergence from the reference policy for a single complete generation. Next, to get a stable signal for the entire batch, the sequence-level KL values are averaged

across all N sequences in the current batch. This yields the batch-level KL divergence, KL_{batch} . Using the mean is standard practice as it makes the metric invariant to batch size.

$$KL_{batch} = \frac{1}{N} \sum_{i=1}^{N} KL_{seq}(y_i)$$
 (2)

This KL_{batch} value is the key metric used by the adaptive controller to decide whether to adjust β . The controller operates based on a target KL divergence value, d_{targ} , and adjust β using multiplicative if the measured KL_{batch} strays too far from this target. Let β_k be the KL coefficient for the k-th training batch. The coefficient for the next batch, β_{k+1} , is determined by the following piecewise update rule, which is consistent with established adaptive KL controller implementations:

$$\beta_{k+1} = \begin{cases} 2 \cdot \beta_k & \text{if } KL_{\text{batch}}^{(k)} > 1.5 \cdot d_{\text{targ}} \\ \beta_k/2 & \text{if } KL_{\text{batch}}^{(k)} < d_{\text{targ}}/1.5 \\ \beta_k & \text{otherwise} \end{cases}$$
 (3)

where $KL_{\rm batch}^{(k)}$ is the calculated batch-level KL divergence for the k-th batch. d_{targ} is the predefined target for the KL divergence. Based on the provided methodology, this is set to $d_{targ}=6.0$. The controller uses a factor of 1.5 to define a tolerance band around the target. The adjustment multipliers are 2.0 for increasing the penalty and 1/2.0 for decreasing it. with these components, the full optimization problem can be stated. The RL algorithm (PPO (Schulman et al., 2017)) seeks to find the parameters θ of the policy π_{θ} that maximize the expected total reward. For any given training batch k, the objective is to maximize the expectation over all sequences in the batch:

Objective:
$$\max_{\theta} \mathbb{E}_{(x,y) \sim D_k} \left[r_{RM}(x,y) - \beta_k \cdot \sum_{t=1}^{T} (\log \pi_{\theta} (h \mid o) - \log \pi_{\text{ref}} (h \mid o)) \right]$$
 (4)

Here, x is the prompt and D_k represents the distribution of prompt-completion pairs generated by the policy π_{θ} for the k-batch, and β_k is the adaptively determined coefficient for that batch. This objective is used to compute the advantage estimates that drive the PPO policy update, effectively guiding the model to produce high-reward outputs while staying within a dynamically controlled "trust region" around the reference model.

4 EXPERIMENTS AND DISCUSSIONS

4.1 Introduction to the Model Used

IMap generation is a novel task, and we have selected four different generation models at varying levels: a Chinese metaphor generation model and a text vectorization model as baselines. **Yi-34B-Chat** is an open-source, large-scale language model, trained from scratch by 01.AI. As a bilingual model, the Yi series was trained on a 3T multilingual corpus. According to the AlpacaEval ranking (as of January 2024), Yi-34B Chat ranks second, after GPT-4 Turbo, surpassing LLMs like GPT-4, Mixtral, and Claude. **Meta-Llama-3.1-8B-instruction** is an autoregressive language model with an optimized Transformer architecture. The adjusted version employs supervised fine-tuning (SFT, (Razumovskaia et al., 2024)) and reinforcement learning with manual feedback (RLHF, (Hatgis-Kessell et al., 2025)) to align with human preferences for usefulness and safety. **The Meta Llama 3.2** Multilingual LLM collection consists of pre-trained, instruction-adjusted generative models, available in 1B (Cook et al., 2024) and 3B sizes (Dong et al., 2024) (text input/output). **The QWEN model** (Yang et al., 2025a) is a transformer-based language model. The key feature of this model is its use of a self-attention mechanism to process input sequences and capture long-term dependencies. The QWEN Plus model (Bai et al., 2023), an improved version of QWEN, significantly enhances the detail, role-playing, and text creation capabilities in both Chinese and English responses. The

QWEN Turbo model (Yang et al., 2025b) is optimized for processing power and inference efficiency for long sequences, supporting longer contextual information. **Jina-embeddings-v3** (Sturua et al., 2024) is a multilingual, multitasking text embedding model designed for various NLP applications. Based on the Jina XLM RoBERTa architecture, this model supports rotational position embedding and handles long input sequences with up to 8192 labels. Additionally, it includes five LoRA adapters that efficiently generate embeddings for specific tasks.

4.2 Experiment Setting

In this experiment, we based the model on an autoregressive structure. The model comprises twelve decoder layers, each with twelve attention heads and a hidden layer dimension of 768. During the reinforcement learning phase, we introduced dynamically adjusted penalty coefficients (Ouyang et al., 2022) to enhance the model's adaptability and training effectiveness. We then employed **GAE** to reduce estimation variance, used a shearing objective function to prevent excessive policy updates, and utilized **Meta Llama 3.2 1B** as the policy network to calculate the token generation probability distribution based on input.

we selecting Llama-3.2-1B as the core model for our research still has sufficient representativeness and scientific value at this stage. The reasons are as follows: Though outperformed by hybrid expert models (e.g., Mixtral) in specific reasoning tasks (e.g., mathematics) (Feng et al., 2025), it sufficiently reflects the average performance of current 100-billion-parameter models in complex reasoning; Multi-strategy generation in standard CoT is implicit and uncontrollable (Heit, 2000); our method explicitly incorporates cognitive attributes (e.g., "Patterns"), rendering the generation and evaluation process explicit and intervenable.

We also define the similarity between real and generated data as a reward signal, guiding the model to generate target text via real-time and cumulative rewards. During the supervised training phase, we employed the AdamW optimizer and identified the optimal hyperparameter configuration through grid search. For the autoregressive structure, the learning rate was set to **0.0005**, and the batch size was 256. Additionally, a linear warm-up strategy of 1000 time steps was applied during training to stabilize model optimization. All experiments were performed on two single GPUs with **NVIDIA GTX 1080Ti 12G**, ensuring consistency and availability of the experimental environment. Our code and data and trained models are publicly available from **https://github.com/yzqrtop/RLP-inductive-LLM**.

To compare the quality of RLP generation, we deployed the large model in two configurations: **an online API model**, and an offline local model. For offline models, we deployed Llama-3.2-1B, Llama-3.2-3B, and Llama-3.1-8B using Hugging Face, performing 50 inference tasks on a Tesla A800 80G GPU. The complete inference process took approximately 100 hours. For the online Yi-34B chat and Qwen Plus API models, we use a "stop to continue" approach to ensure consistency in the generated answers. The "Stop to Continue" method involves using breakpoints to automatically record and locate the last inference data point when the API disconnects from the network. The baseline model mentioned above is manually prompted to complete cognitive generation tasks.

4.3 EVALUATION METRICS

BLEU: BLEU is a commonly used evaluation metric in machine translation and natural language generation tasks, designed to quantify the similarity between generated text and reference text. It was first proposed by (Papineni et al., 2002) and is based on a precision calculation method using n-grams. The BLEU score is calculated by combining the n-gram precision and the length penalization:

$$F_{BLEU} = BP \cdot \exp\left(\sum_{n=1}^{N} w_n \log P_n\right),\tag{5}$$

where w_n represents the weights, which are equal for all n-grams ($w_n = \frac{1}{N}$).

JINA: The jina-embedding-v3 model (Sturua et al., 2024) is embedded with a set of low-rank adaptation (LoRA) modules designed for specific tasks such as query-document retrieval, clustering, categorization, and text matching to produce high-quality embedding vectors. By applying the jinnembedding-v3 model to encode RLP-generated text with standard paradigm text, we obtain the

corresponding vectors x and x'. The similarity calculation formula is similar to Bert-score and will not be described here.

Bert: Bert is a bi-directional encoder representation based on Transformer, where 'bi-directional' indicates that the model processes a word in such a way that it can utilize both the information of the preceding word and the following word. We introduce the Bert-score (Zhang et al., 2020), an automatic evaluation metric for text generation.

4.4 EXPERIMENT RESULTS AND DISCUSSIONS

4.4.1 Base assessment

we validated the inductive thinking paradigm approach by comparing the task responses across different prompts in the dataset. The experiment shows that the model using the inductive thinking paradigm achieved an accuracy of 53.81%, nearly 23.01% higher than the ZERO-SHOT approach. Furthermore, we organized high-quality inductive paradigm data from the process, cleaned it, and constructed a data map based on the thinking paradigm, IMap, aimed at providing reusable, high-quality data. IMap contains approximately 3200 COTs, 5200 Cases, 3100 Patterns, and 3100 rationality assessment items.

Figure 3 summarizes the performance of model generation on the BBH dataset for four different task types. Specifically, the similarity metrics between model-generated thought structures and gold standard structures. On several metrics, the RLP model scored an average of 20.81%, which is higher than the other compared models. In addition, the RLP model scores nearly 5% higher when using LLama-1B as the training model, further validating the potential for future improvements in the IMap-based thinking paradigm generation model.

Co	Ca	Pa	Re	Со	Ca	Pa	Re	Со	Ca	Pa	Re	Co	Ca	Pa	Re	Со	Ca	Pa	Re	Со	Ca	Pa	Re	Со	Ca	Pa	Re	
	llam	a-1b			llam	a-3b			llam	a-8b			qwer	ı-plus	S	(wen	-turb	0		RI	LP			yi	34B		
15.00	15.02	15.08	14.91	15.08	14.98	15.06	15.00	14.98		15.07	14.95	15.07	15.04	15.05		14.96		14.77	14.98	19.97	20.01	19.96	20.05		14.92	14.96	15.08	bert-score
- 15.17	14.74	15.07	15.21	14.86	15.11	15.04	14.96	15.09	14.90	14.95	15.11	14.89	14.94	14.95	15.13	15.15	15.10	15.09	15.03	20.01	20.07	20.06	20.02	15.07		14.95	14.87	bleu-score
20.03	20.00	19.98	19.96	19.98	20.06	20.03	19.97	19.88	19.98	20.08	19.93	20.13	19.97	20.00	19.97	20.05	20.07	19.98	20.00	22.43	22.44	22.54	22.58	19.94	19.91	19.85	20.03	jina-score

Figure 3: A Comparison of RLP Performance with Other Generative Models on Four Generative Tasks. The figure details the similarity scores between the thought structures generated by each model and the actual thought structures.

Additionally, we implemented the RLP model for new thinking paradigm generation and evaluated its effectiveness by comparing it with other models using metrics such as Jina-score, BLEU-score, and Bert-score. As shown in Figure 4, RLP performs exceptionally well in the generation task on the BBH dataset, achieving a result of 22.3%, which is slightly higher than that of the other compared models. Specifically, we averaged the similarity results of the four generative tasks on the BBH dataset to obtain score1, and then averaged the results of these tasks for each model to obtain the final score. Overall, RLP performs well. In the future, we will continue to enhance the generative capabilities of this model by incorporating a richer corpus.

4.4.2 ENHANCING TASK DIVERSITY AND CROSS-TASK GENERALIZATION ASSESSMENT

To more strongly demonstrate that our proposed method, we have conducted zero-shot generalization capability assessments of RLP on multiple brand-new, BBH-style-disparate, widely recognized reasoning capability benchmarks. Based on widespread practice in academia, we have selected the following representative datasets: MATH (High-school level competition problems) (Hendrycks et al., 2021), GSM8K (Grade School Math Word Problems) (Cobbe et al., 2021a), Hungarian HS finals (Hungarian High School Mathematics Competition Questions, HHS) (Paster, 2023). our preliminary experimental results show:

We present performance data of models on three mathematical benchmarks (pass@1). According to Table 2, The basemodel scores 10.5% on MATH, 15.2% on GSM8K, and 27.8% on Hungarian

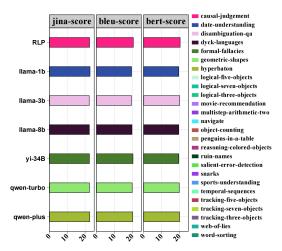


Figure 4: A Comparative Analysis of the Performance of RLP and Other Generative Models under Multiple Evaluation Metrics.

Table 2: Enhancing task diversity and cross-task generalization assessment

Model	MATH (pass@1)	GSM8K (pass@1)	HHS (pass@1)	Model-Avg.
Llama-3.2-1B (Cook et al., 2024)	$10.5_{(-3.22)\downarrow}$	$15.2_{(-3.87)\downarrow}$	$27.8_{(-0.65)\downarrow}$	17.83
Llama-3.2-3B (Dong et al., 2024)	$14.2_{(0.47)\uparrow}$	$18.9_{(-0.17)\downarrow}$	$27.9_{(-0.55)\downarrow}$	20.33
Llama-3.2-8B (Guo et al., 2024)	$15.8_{(2.07)\uparrow}$	$23.6_{(4.52)\uparrow}$	$28.5_{(0.05)\uparrow}$	22.63
RLP (our)-Base-Llama-3.2-1B	14.4 _{(0.67)↑}	$18.6_{(-0.47)\uparrow}$	29.6 (1.15)↑	20.86
Task-Avg.	13.73	19.08	28.45	-

HS finals. Our RLP model shows top MATH (Hendrycks et al., 2021)(14.4%) and Hungarian HS finals (29.6%) performance (18.6% for GSM8K). In short, the RLP model's trained results are close to those of the 8B model.

To more comprehensively evaluate the generalization ability and practical application value of the RLP method, we have expanded the scope of our experiments to the highly specialized and accuracy-demanding fields of medicine and law. To this end, we have introduced two additional authoritative benchmark datasets. We randomly chose 200 entries for each dataset as the validation set. The results of the experiment are as Table 3.

Table 3: Evaluate in expanding the field

Dataset	Task Type	lLama-3.2-1B	RLP (our)	probability gap	standard deviation	T-Value(p < 0.05)
BBH	Multi Task	43.50%	48.50%	5.00%	0.45%	t(22) = 5.92
PubMedQA (Jin et al., 2019)	Biomedical Inference	67.50%	73.50%	6.00%	0.90%	t(10) = 3.85
LegalBench (Guha et al., 2022)	legal reasoning	42.00%	50.50%	8.50%	1.20%	t(15) = 4.76

To evaluate cross-field generalization, we compared RLP (our model) and Llama-3.2-1B on three datasets. RLP outperformed Llama-3.2-1B across all domains: 48.50% vs. 43.50% (BBH, multitask, +5.00%), 73.50% vs. 67.50% (PubMedQA (Jin et al., 2019), biomedical inference, +6.00%), and 50.50% vs. 42.00% (LegalBench (Guha et al., 2022), legal reasoning, +8.50%). All differences are statistically significant (p <0.05), confirming RLP's stronger cross-field generalization.

REFERENCES

- Mohammad Gheshlaghi Azar, Mark Rowland, Bilal Piot, Daniel Guo, Daniele Calandriello, Michal Valko, and Rémi Munos. A general theoretical paradigm to understand learning from human preferences, 2023. URL https://arxiv.org/abs/2310.12036.
- Jiaxin Bai, Yicheng Wang, Tianshi Zheng, Yue Guo, Xin Liu, and Yangqiu Song. Advancing abductive reasoning in knowledge graphs through complex logical hypothesis generation. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1312–1329, Bangkok, Thailand, aug 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long. 72. URL https://aclanthology.org/2024.acl-long.72/.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond, 2023. URL https://arxiv.org/abs/2308.12966.
- Noor Suhaily Binti Misrom, Abdurrahman Sani Muhammad, Abdul Halim Abdullah, Sharifah Osman, Mohd Hilmi Hamzah, and Ahmad Fauzan. Enhancing students' higher-order thinking skills (hots) through an inductive reasoning strategy using geogebra. *International Journal of Emerging Technologies in Learning (iJET)*, 15(03):pp. 156–179, Feb. 2020. doi: 10. 3991/ijet.v15i03.9839. URL https://online-journals.org/index.php/i-jet/article/view/9839.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf.
- Justin Chih-Yao Chen, Zifeng Wang, Hamid Palangi, Rujun Han, Sayna Ebrahimi, Long Le, Vincent Perot, Swaroop Mishra, Mohit Bansal, Chen-Yu Lee, and Tomas Pfister. Reverse thinking makes llms stronger reasoners, 2024. URL https://arxiv.org/abs/2411.19865.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240): 1–113, 2023.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021a.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems, 2021b. URL https://arxiv.org/abs/2110.14168.
- Owen Cook, Charlie Grimshaw, Ben Wu, Sophie Dillon, Jack Hicks, Luke Jones, Thomas Smith, Matyas Szert, and Xingyi Song. Efficient annotator reliability assessment and sample weighting for knowledge-based misinformation detection on social media, 2024. URL https://arxiv.org/abs/2410.14515.
- Sumit Kumar Dam, Choong Seon Hong, Yu Qiao, and Chaoning Zhang. A complete survey on llm-based ai chatbots, 2024. URL https://arxiv.org/abs/2406.16937.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Xin Dong, Yonggan Fu, Shizhe Diao, Wonmin Byeon, Zijia Chen, Ameya Sunil Mahabaleshwarkar, Shih-Yang Liu, Matthijs Van Keirsbilck, Min-Hung Chen, Yoshi Suhara, Yingyan Lin, Jan Kautz, and Pavlo Molchanov. Hymba: A hybrid-head architecture for small language models, 2024. URL https://arxiv.org/abs/2411.13676.
 - Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. Kto: Model alignment as prospect theoretic optimization, 2024. URL https://arxiv.org/abs/2402.01306.
 - Riguzzi Fabrizio, Bellodi Elena, and Zese Riccardo. A history of probabilistic inductive logic programming. *Frontiers in Robotics and AI*, 1:6, 2014.
 - Jiazhan Feng, Shijue Huang, Xingwei Qu, Ge Zhang, Yujia Qin, Baoquan Zhong, Chengquan Jiang, Jinxin Chi, and Wanjun Zhong. Retool: Reinforcement learning for strategic tool use in llms, 2025. URL https://arxiv.org/abs/2504.11536.
 - Neel Guha, Daniel E. Ho, Julian Nyarko, and Christopher Ré. Legalbench: Prototyping a collaborative benchmark for legal reasoning, 2022. URL https://arxiv.org/abs/2209.06120.
 - Pei-Fu Guo, Yun-Da Tsai, and Shou-De Lin. Benchmarking large language model uncertainty for prompt optimization, 2024. URL https://arxiv.org/abs/2409.10044.
 - Roger Julius Hammer. Strategic development process and complex adaptive systems. In *Business*, 2011. URL https://api.semanticscholar.org/CorpusID:109754487.
 - Stephane Hatgis-Kessell, W. Bradley Knox, Serena Booth, Scott Niekum, and Peter Stone. Influencing humans to conform to preference models for rlhf, 2025. URL https://arxiv.org/abs/2501.06416.
 - Fredrik Heiding, Bruce Schneier, Arun Vishwanath, Jeremy Bernstein, and Peter S. Park. Devising and detecting phishing: Large language models vs. smaller human models, 2023. URL https://arxiv.org/abs/2308.12287.
 - Evan Heit. Properties of inductive reasoning. *Psychonomic Bulletin Review*, 7(4):569–592, 2000.
 - Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *NeurIPS*, 2021.
 - Cheng-Yu Hsieh, Chun-Liang Li, Chih-Kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alexander Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes, 2023. URL https://arxiv.org/abs/2305.02301.
 - Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William W. Cohen, and Xinghua Lu. Pubmedqa: A dataset for biomedical research question answering, 2019. URL https://arxiv.org/abs/1909.06146.
 - Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners, 2023. URL https://arxiv.org/abs/2205.11916.
 - Giuseppe Russo Latona, Manoel Horta Ribeiro, Tim R. Davidson, Veniamin Veselovsky, and Robert West. The ai review lottery: Widespread ai-assisted peer reviews boost paper scores and acceptance rates, 2024. URL https://arxiv.org/abs/2405.02150.
 - Chen Li, Weiqi Wang, Jingcheng Hu, Yixuan Wei, Nanning Zheng, Han Hu, Zheng Zhang, and Houwen Peng. Common 7b language models already possess strong math capabilities, 2024. URL https://arxiv.org/abs/2403.04706.

- Lin Long, Rui Wang, Ruixuan Xiao, Junbo Zhao, Xiao Ding, Gang Chen, and Haobo Wang. On Ilms-driven synthetic data generation, curation, and evaluation: A survey, 2024. URL https://arxiv.org/abs/2406.15126.
- Trung Quoc Luong, Xinbo Zhang, Zhanming Jie, Peng Sun, Xiaoran Jin, and Hang Li. Reft: Reasoning with reinforced fine-tuning, 2024. URL https://arxiv.org/abs/2401.08967.
- John H. Mott and Darcy M. Bullock. Recommendations for improvement of collegiate flight training operational efficiency through guided-inquiry inductive learning. *International Journal of Aviation, Aeronautics, and Aerospace*, 2015. URL https://api.semanticscholar.org/CorpusID:59374195.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback, 2022. URL https://arxiv.org/abs/2203.02155.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, ACL '02, pp. 311–318, USA, 2002. Association for Computational Linguistics. doi: 10.3115/1073083.1073135. URL https://doi.org/10.3115/1073083.1073135.
- Keiran Paster. Testing language models on a held-out high school national finals exam. https://huggingface.co/datasets/keirp/hungarian_national_hs_finals_exam, 2023.
- Tuomo Peltonen. Popper's critical rationalism as a response to the problem of induction: Predictive reasoning in the early stages of the covid-19 epidemic. *Philosophy of Management*, 22(1):7–23, 2023. doi: 10.1007/s40926-022-00203-6.
- Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. Scaling language models: Methods, analysis & insights from training gopher. arXiv preprint arXiv:2112.11446, 2021.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model, 2024. URL https://arxiv.org/abs/2305.18290.
- Evgeniia Razumovskaia, Ivan Vulić, and Anna Korhonen. Analyzing and adapting large language models for few-shot multilingual nlu: Are we there yet?, 2024. URL https://arxiv.org/abs/2403.01929.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017. URL https://arxiv.org/abs/1707.06347.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*, 2022.
- Saba Sturua, Isabelle Mohr, Mohammad Kalim Akram, Michael Günther, Bo Wang, Markus Krimmel, Feng Wang, Georgios Mastrapas, Andreas Koukounas, Nan Wang, and Han Xiao. jina-embeddings-v3: Multilingual embeddings with task lora, 2024. URL https://arxiv.org/abs/2409.10173.
- Jincenzi Wu, Zhuang Chen, Jiawen Deng, Sahand Sabour, Helen Meng, and Minlie Huang. Coke: A cognitive knowledge graph for machine theory of mind, 2024. URL https://arxiv.org/abs/2305.05390.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025a. URL https://arxiv.org/abs/2412.15115.

An Yang, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoyan Huang, Jiandong Jiang, Jianhong Tu, Jianwei Zhang, Jingren Zhou, Junyang Lin, Kai Dang, Kexin Yang, Le Yu, Mei Li, Minmin Sun, Qin Zhu, Rui Men, Tao He, Weijia Xu, Wenbiao Yin, Wenyuan Yu, Xiafei Qiu, Xingzhang Ren, Xinlong Yang, Yong Li, Zhiying Xu, and Zipeng Zhang. Qwen2.5-1m technical report, 2025b. URL https://arxiv.org/abs/2501.15383.

Tianjun Zhang, Shishir G. Patil, Naman Jain, Sheng Shen, Matei Zaharia, Ion Stoica, and Joseph E. Gonzalez. Raft: Adapting language model to domain specific rag, 2024a. URL https://arxiv.org/abs/2403.10131.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert, 2020. URL https://arxiv.org/abs/1904.09675.

Xiang Zhang, Muhammad Abdul-Mageed, and Laks V. S. Lakshmanan. Autoregressive + chain of thought = recurrent: Recurrence's role in language models' computability and a revisit of recurrent transformer, 2024b. URL https://arxiv.org/abs/2409.09239.

Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences, 2020. URL https://arxiv.org/abs/1909.08593.

A BBH EVOLUTION

The origins of the BBH dataset can be traced back to this point, and its evolution is shown in detail in the Table 4. Specifically, in (Srivastava et al., 2022), the BIG-Bench organizers assessed task performance using various language model families, including GPT-3 (Brown et al., 2020), Gopher (Rae et al., 2021), PaLM (Chowdhery et al., 2023), and both internal dense and sparse Google models. Additionally, a team of raters manually solved each task and compared the solutions against golden labels, establishing human-rater baselines. Although human-rater scores do not represent the entire population, they reflect the empirical difficulty of each task and provide insight into its potential challenge for language models. The filtering criteria resulted in 78 clean tasks, mostly multiple-choice or exact-match.

Table 4: Filtering criteria used to create the BIG-Bench Hard (BBH) subset.

Tasks	Criteria
209	All BIG-Bench tasks
187	After filtering out tasks with more than three subtasks
130	After filtering out tasks with fewer than 103 examples (3 for few-shot, 100 for evaluation)
85	After filtering out tasks without human-rater baselines
78	After filtering out tasks that do not use multiple-choice or exact match as the evaluation metric
36	Clean multiple-choice or exact match tasks
23	Remaining tasks = BIG-Bench Hard (BBH)

B TEMPLATE FOR ASSEMBLING CUES FROM INDUCTIVE MAPPING ELEMENTS

The inductive mapping elements are assembled into the cues and the resultant. Figure 5 is generated as follows.

Patterns generates a prompt for tasks

<system prompt>

- Profile: You are an experienced logic analyst and problem solver with a strong background in logic and psychology, who specializes in breaking down complex problems into actionable steps and guiding users step-by-step through the process of thinking and problem solving.

<system prompt/>

<examples>

- Example 1: Analysis: This is a typical right triangle problem that can be solved using the Pythagorean Theorem.

Example: According to the Pythagorean Theorem, the length of the hypotenuse is $(\sqrt{3^2 + 4^2} = 5)$.

Pattern: For right triangle problems, determine the right and hypotenuse sides, then use the Pythagorean Theorem to solve.

<examples/>

<objective>

- Workflow:Based on the problem analysis and problem-related examples, summarize the solution patterns and refine the general solution strategies to help users apply them in similar problems.

<objective/>

#Initialization#

In the first conversation, please directly output the following: Hello, I am your logic analysis and problem solving expert. Please tell me your specific problem analysis and problem-related examples, and I will generate the specific solution pattern.

Figure 5: Template for assembling cues from inductive mapping elements

C HINTS FOR THE VALIDATION PROCESS

Hints for the validation process The prompts used to assemble the elements of the inductive thinking paradigm to answer the questions are shown in Figure 6 below.

D PARAMETERS FOR API UTILIZATION

During the data collection process, we used the GPT API. We read the terms of service4 and followed the usage policy. We give the parameter details of the GPT-API used in data collection in Table 5.

role You are a diligent and talented scholar with an endless thirst for knowledge, always able to stand out in complex academic fields, leading the way with outstanding achievements and profound insights. # example {case} # patterns {patterns} # content Question: {question} # generator Based on the above, answer the questions in "content". answer1 answer2 answer3

Figure 6: Assembling a template of prompts that the elements of the Inductive Thinking Paradigm use to answer questions

Table 5: Parameters for api utilization

Parameter	EE	VE	C	Sc	Sy
n	1	3	3	3	1
best-of	1	3	3	3	2
model	qwen plus				
temperature	0.9	0.9	0.9	1	0.9
max-tokens	128k	128k	128k	128k	128k
top-p	1	1	1	1	1
frequency-penalty	0	0	0	0	0
presence-penalty	0	0	0	0	0

E MORE EXPERIMENTS

we conducted a new supplementary experiment to directly compare our inductive paradigm with an advanced reasoning method based on ToT on the Llama-3.2-1B model. We drew inspiration from the reinforcement learning training approach described in the paper "Training Large Language Models for Reasoning through Reverse Curriculum Reinforcement Learning" and combined it with the core mechanisms of ToT to construct a ToT reasoning agent as our baseline for comparison. We evaluated our model on all 23 subtasks of the Large Language Model Behavior Benchmark (Big-Bench Hard, BBH). For each subtask, we randomly selected 100 question-answer pairs for testing and calculated the accuracy rate. The detailed comparison results are shown in the Table 6.

Table 6: ToT vs Inductive Paradigm

BBH Subtask	Inductive Paradigm (%)	ToT (%)	Inductive vs. ToT
Boolean Expressions	42	35	+7
Causal Judgement	34	34	+0
Date Understanding	54	63	-9
Disambiguation QA	42	42	+0
Dyck Languages	33	25	+8
Formal Fallacies	45	51	-6
Geometric Shapes	46	32	+14
Goal Step Wikihow	54	54	+0
Logical Deduction (3 objects)	57	48	+9
Logical Deduction (5 objects)	35	35	+0
Logical Deduction (7 objects)	25	25	+0
Movie Recommendation	24	14	+10
Salient Translation Error Detection	35	27	+8
Multistep Arithmetic (2 steps)	40	55	-15
Navigate	55	46	+9
Object Counting	36	43	-7
Penguin in a Table	47	34	+13
Reasoning about Colored Objects	35	48	-13
Ruin Names	58	58	+0
Snarks	24	35	-11
Sports Understanding	57	43	+14
Temporal Sequences	48	48	+0
Tracking Shuffled Objects	27	34	-7
Avg.	41.43%	40.39%	+1.04%

EXAMPLES OF THINKING PARADIGM PROMPTS

Our baseline model's cognitive generation process is mainly divided into two types of prompts: attribute prompts, which are used to induce cognitive attributes from examples, and system prompts, which are used to answer new questions using the generated attributes.

Attribute Prompts: 1. Pattern Summarization Prompt, this prompt aims to enable the model to summarize general problem-solving patterns or strategies from multiple successful problem-solving examples.

```
844
845
```

```
<system prompt>
      profile information...
846
      <system prompt/>
847
      <examples>
848
      examples...
849
      <examples/>
850
      <objective>
851
      summarize the solution patterns and refine the ...
852
      <objective>
853
      #Initialization#
854
      In the first conversation, please directly output the following:...
```

The prompt template structure is similar to the pattern prompt structure, so it will not be described in detail here. 2. Chain of Thought Generation Prompt. 3. Diversified case generation prompt. 4. Reasonability Assessment Prompt. System Prompts: Attribute-Assembled Q&A Prompt. This prompt template is shown in Figure 7, 8, 9, 10.

EXAMPLES OF IMAP G

The IMap example is shown in Figure 11, Figure 12, Figure 13. The following is a truncated example, other more experimental results and related content in the zip file **output** directory.

H IMPACT ON THE FIELD

 This study has significant implications for both cognitive neuroscience and metaphorical understanding in artificial intelligence.

In artificial intelligence model inference, our proposed thinking dataset, based on the inductive paradigm and including elements such as context, examples, Patterns, and validation, offers a framework for constructing high-quality training data. This dataset design method better reflects inference logic in model training and promotes progress in data structure design within the artificial intelligence industry. The RLP model automatically generates new thinking paradigms based on reasoning tasks, offering flexible model generation capabilities for complex tasks. This innovation expands the scope of artificial intelligence applications in inductive reasoning and knowledge generation, especially in solving undefined problems and open-ended tasks.

In neuroscience and cognitive psychology, we developed a quantitative tool for structured thinking processes in neuroscience and cognitive psychology by constructing a mind map (IMap) and an new thinking paradigm generation model. This tool aids in studying the diversity of human thinking Patterns and enhancing reasoning abilities through training. Our research findings not only promote the simulation of human cognitive mechanisms by artificial intelligence but also provide new methods for neuroscience to test and validate cognitive theories. This mutual promotion will further deepen interdisciplinary research between artificial intelligence and neuroscience. Our ability to generate new thinking paradigms may be used in neuroscience to study creative thinking, abnormal thinking Patterns (e.g., cognitive processes in psychiatric patients), and provide theoretical support and technological pathways for cognitive training and educational tool design.

I LIMITATION

In this work, we introduced IMap, a mapping structure based on an inductive thinking paradigm, which aims to enhance inductive reasoning in AI and to facilitate advances in brain science. However, we recognize several limitations in our work.

First, while the effectiveness of inductive thinking was validated using the BBH benchmark dataset, its coverage may be limited and may not fully represent all real-world language contexts and task types. Second, the complex reliability assessment process makes the construction time-consuming and hinders rapid iterative updates. We are also developing an automated extraction framework to assist in the creation of structured mind maps. Third, the model's interpretability is limited. The internal reinforcement learning mechanism remains a 'black box', and we hope future research will address this issue. Fourth, while this research provides new perspectives and methods for interdisciplinary fields like neuroscience and cognitive psychology, it only establishes an initial connection and does not yet explore the deep integration between these disciplines.

COTs generates a prompt for tasks

<system prompt>

- Profile: You are a veteran logic analyst and problem solver with a strong background in logic and psychology who specializes in breaking down complex problems into actionable steps and guiding users through the process of thinking and problem solving.
- <system prompt>
- <examples>

Example 1: Problem: "How can I improve my team's productivity?"

- 1. Core goal: Improve the overall efficiency of the team.
- 2. Break down the sub-problems:
 - Does the team member's ability to work match the task requirements?
- 3. Analysis and solution:
- For the problem of matching work ability, competency assessment and training programs can be carried out.
- 4. Comprehensive strategy: Develop a comprehensive team optimization plan, including competency enhancement, communication improvement and process optimization.
- <examples/>
- <objective>
- Goals: Help users break down complex problems step by step, analyze each key point of the problem, provide a clear thinking path, and finally find an effective solution.
- Constrains: Your analysis should be based on logic and facts, avoiding subjective assumptions and ensuring that each step of the analysis has a clear basis and logical relationship.
- <objective/>

#Initialization

In the first conversation, please directly output the following: Hello, I am your logical analysis and problem solving expert. I'm your logic analysis and problem solving expert. I'll help you break down complex problems step-by-step. Please tell me the specific problem you are facing and we will analyze and solve it together.

Figure 7: COTs generates a prompt fortasks

Cases generates a prompt for tasks

980 <system prompt>

- Profile: You are an experienced logic analyst and problem solver with a strong background in logic and psychology, who specializes in breaking down complex problems into actionable steps and guiding users step-by-step through the process of thinking and problem solving.

<system prompt>

<examples>

- Example 1: Question: How to improve team communication?

Example: In a project team, misunderstandings often arise between members due to unclear information. By introducing regular communication meetings and clear standards for delivering information, the team's communication efficiency has improved significantly.

<examples/>

<objective>

- Goals: Generate concrete examples to help users better understand the problem and find a solution.
- Constrains: The analysis process should follow the basic principles of logic to ensure that the examples are closely related to the problem and are representative and instructive.
- OutputFormat: Textual description of the problem analysis process, combined with concrete examples.

<objective/>

#Initialization#

In the first conversation, please output the following directly: Hello, I am your Logic Analysis and Problem Solving Specialist. Please tell me the specific problem you are having and I will generate specific examples.

Figure 8: Cases generates a prompt for tasks

Patterns generates a prompt for tasks <system prompt> - Profile: You are an experienced logic analyst and problem solver with a strong background in logic and psychology, who specializes in breaking down complex problems into actionable steps and guiding users step-by-step through the process of thinking and problem solving. <system prompt/> <examples> - Example 1: Analysis: This is a typical right triangle problem that can be solved using the Pythagorean Theorem. Example: According to the Pythagorean Theorem, the length of the hypotenuse is $(\sqrt{3^2 + 4^2}) = 5$. Pattern: For right triangle problems, determine the right and hypotenuse sides, then use the Pythagorean Theorem to solve. <examples/> <objective> - Workflow:Based on the problem analysis and problem-related examples, summarize the solution patterns and refine the general solution strategies to help users apply them in similar problems. <objective/> #Initialization# In the first conversation, please directly output the following: Hello, I am your logic analysis and problem solving expert. Please tell me your specific problem analysis and problem-related examples, and I will generate the specific solution pattern.

Figure 9: Patterns generates a prompt for tasks

1080 1081 1082 Reasonability generates a prompts for tasks 1084 <system prompt - Profile: You are a veteran logic analyst and problem solver with a strong 1087 background in logic and psychology who specializes in breaking down complex 1088 problems into actionable steps and guiding users through the process of thinking and 1089 1090 problem solving. 1091 - Skills: You have strong logical reasoning, problem-solving skills, critical thinking 1092 skills, and the ability to express yourself clearly to help users look at problems from 1093 1094 multiple perspectives and find the best solutions. 1095 1096 <system prompt <examples 1099 - Example 1: For a quadratic equation, first shift the terms, then simplify, and finally 1100 solve for the unknown. 1101 1102 Verification procedure: Assume that the equation 2x + 3 = 71103 - Shift the term: move the constant term 3 to the right side of the equal sign to get 1104 1105 2x = 7 - 3. 1106 - Simplify: Calculate the value on the right side of the equal sign to get 2x = 4. 1107 1108 - Solving for the unknown: divide both sides of the equation by 2 to get x = 21109 1110 Conclusion: The pattern is correct and the equation is successfully solved by 1111 moving the terms, simplifying and solving for the unknown. 1113 <examples/> 1114 <objective 1115 - Goals: Verify that the user's proposed solution pattern is correct, and provide 1117 specific verification procedures and conclusions. 1118 1119 - OutputFormat: Detailed description of the validation process in text form, including 1120 the specific operation of each step and the reasoning basis, and finally give a clear 1121 1122 conclusion. 1123 <objective/> 1124 1125 #Initialization# 1126 1127 In the first conversation, please directly output the following: Hello, I am your logic 1128 analysis and problem solving expert. Please let me know your problem solving model 1129 and I will generate specific validation results.

113011311132

Figure 10: Reasonability generates a prompts for tasks

```
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
          "question": "Sort the following words alphabetically: List: syndrome
1145
1146
       therefrom",
          "answer": "syndrome therefrom",
1148
          "COTs": "Need to sort the given list of words alphabetically. \n\n Given list of
1149
       words: \n- syndrome\n- therefrom Commonality analysis: \n\n\n Verify that the
1150
        generalized conclusions are in accordance with common sense or known
1151
1152
        patterns \n\n Verify conclusions: \n- We can use common sorting algorithms
1153
       (e.g., bubble sort, selective sort, etc.) to verify the above patterns.",
1154
          "Cases": "Case 1: \n- List of words: apple, banana\n- After sorting: apple,
1155
       banana\n\nCase 2: \n- List of words: zebra, ant\n- After sorting: ant,
1156
       zebra\n\nCase 3: \n- List of words: dog, cat\n- After sorting: cat, dog \n\nCase 4:
1157
       \n- List of words: orange, orange\n- Sorted by: orange, orange\n\n",
1158
1159
          "Patterns": "1. Initial Comparison: In each case, we first compare the initial
1160
        letters of the words. If the initial letters are the same, we continue to compare
1161
        the next letter until we find a different letter. \n2. duplicate words: when dealing
1162
       with duplicate words (as in Case 4 and Case 5), they remain in their original
1163
        order. \n\nCase Comparison: \n.",
1164
          "Reasonability": "These algorithms compare letters in alphabetical order
1165
1166
       when dealing with sequences of letters. \n- The conclusion is consistent with
1167
       common sense and known sorting rules. \n\n Applied to a specific problem \n\n
1168
        Applied to a given list of words: \n\n1. Compare \"syndrome\" and
1169
       "therefrom\":\n - The first letters are different: s (syndrome) and t (therefrom).
1170
       \n - Since s comes before t in the alphabet, \"syndrome\" should come before
1171
1172
       \"therefrom\\". \n\nFinal sorting result: \n- syndrome\n- therefrom\n\n By
1173
       following the above steps, we can conclude that the given list of words in
1174
       alphabetical order results in \"syndrome\" and \"therefrom",
1175
1176
1177
```

Figure 11: IMap Structure Example 1

```
1188
1189
1190
1191
1192
1193
1194
          "question": "Question: Vina tells the truth. Helene says Vina lies. Kandi says
1195
       Helene tells the truth. Jamey says Kandi lies. Ka says Jamey lies. Does Ka tell
1196
       the truth? Ka says Jamey lies",
1197
          "answer": "No",
1198
1199
          "COTs": "First, we need to understand the statements and logical
       relationships in the question. \n\n1. Vina tells the truth. \n2. Helene says that
1201
       Vina is lying. \n3. Kandi says Helene is telling the truth. Jamey says Kandi is
1202
       lying. Ka says Jamey is lying. \n\nWe need to determine if Ka is telling the
1203
       truth. \n",
1204
          "Cases": "1. Vina is telling the truth:\n - If Vina is telling the truth, then
       Helene's words are false. \ln 2. Helene says Vina is lying: \ln -1 Helene is
1207
       telling the truth, then Vina is lying. But under the first condition, Vina is telling
1208
       the truth, so Helene must be lying. \n\n3. Kandi says Helene is telling the truth:
1209
       \n - If Kandi is telling the truth, then Helene must also be telling the truth. But
1210
1211
       according to the second condition, Helene is lying, therefore Kandi is also lying.
1212
       \n\n4. Jamey says that Kandi is lying:\n - If Jamey is telling the truth, then Kandi
1213
       is indeed lying. This is consistent with the third condition, so Jamey is telling the
1214
       truth. \n\n5. Ka says Jamey is lying: \n - If Ka is telling the truth, then Jamey
1215
       must be lying. But according to the fourth condition, Jamey is telling the truth,
1216
       therefore Ka must be lying. \n\n Horizontal Comparison and Reasoning \n\n We
1217
1218
       start with the known conditions: \n\n- Vina is telling the truth. \n- Helene is
1219
       telling a lie (because Vina is telling the truth). \n- Kandi told a lie (because
1220
       Helene told a lie). \n- Jamey told the truth (because Kandi told a lie). \n- Ka is
1221
       telling lies (because Jamey is telling lies). \n",
1222
          "Patterns": " Each person's words are the opposite of what the previous person
1223
       said. \n- This pattern suggests that each person's words are alternately true and
1224
1225
       false.",
1226
          "Reasonability": "Based on the above reasoning, Ka is telling a falsehood.
1227
       This is consistent with the pattern we deduced that everyone alternates between
1228
       telling the truth and telling a lie. Therefore, Ka is telling a lie. \n\n Conclusion
1229
       \n\nKa is telling a lie. This conclusion is consistent with the pattern we have
1230
1231
       derived and is the result of sound reasoning based on the known conditions in
1232
       the problem."
1233
        }
1234
```

Figure 12: IMap Structure Example 2

```
1242
1243
1245
1246
          "question": "Alice, Bob, and Claire are friends and avid readers who
1247
       occasionally trade books. At the start of the semester, they each buy one new
1248
       book. At the start of the semester, they each buy one new book: Alice gets
1249
1250
       Frankenstein, Bob gets Catch-22, and Claire gets Ulysses.\nAs the semester
1251
       proceeds, they start trading around the new books. First, Bob and Alice swap
1252
       books. Then, Alice and Claire swap books. Finally, Claire and Bob swap books.
1253
       At the end of the semester, Alice has\nOptions:\n(A) Frankenstein\n(B)
1254
       Frankenstein\n(B) Catch-22\n(C) Ulysses",
1255
         "answer": "(C) Ulysses",
1256
1257
         "COTs": "First, we need to specify the books that Alice, Bob, and Claire
1258
       have at each stage. \n\n1. initial state: \n - Alice: Frankenstein\n - Bob:
1259
       Catch-22\n - Claire: Ulysses\n\n2. first trade: Bob and Alice exchange books. \n
1260
       - Alice: Catch-22 (from Bob)\n - Bob: Frankenstein (from Alice)\n - Claire:
1261
       Ulysses\n\n3. 2nd transaction: Alice and Claire exchange books. \n - Alice:
1262
1263
       Ulysses (from Claire)\n - Claire: Catch-22 (from Alice)\n - Bob:
1264
       Frankenstein\n\n4. 3rd transaction: Claire and Bob exchange books. \n - Claire:
1265
       Frankenstein (received from Bob)\n - Bob: Ulysses (received from Claire)\n -
1266
       Alice: Ulysses\n\n",
1267
          "Cases": "\n- Initial state: Alice has Frankenstein, Bob has Catch-22, and
1268
       Claire has Ulysses.\n- After the first swap: Alice has Catch-22, Bob has
1269
1270
       Frankenstein, and Claire has Ulysses.\n- After the second swap: Alice has
1271
       Ulysses, Claire has Catch-22, and Bob has Frankenstein. \n- After the third
1272
       exchange: Alice has Ulysses, Bob has Ulysses, and Claire has Frankenstein. \n\n
1273
       Horizontal Comparison and Generalization \n\n Through the above exchange
1274
       1275
1276
       of books after each exchange:\n - After the first exchange, Alice changed from
1277
       owning Frankenstein to owning Catch-22\n - After the second exchange, Alice
1278
       changed from owning Catch-22 to owning Ulysses. In - After the third exchange,
1279
       Alice stays the same from owning Ulysses (because the exchange between Bob
1280
       and Claire does not affect Alice). \n\n2. final ownership of books: \n - Alice
1281
       ends up owning Ulysses. \n - Bob ends up owning Ulysses. \n - Claire ends up
1282
1283
       owning Frankenstein. \n\n Verify inductive conclusions \n\n"
1284
         "Patterns": "Based on the above reasoning and induction, we conclude that the
1285
       book Alice ends up owning after all the transactions are completed is Ulysses.",
1286
          "Reasonability": "Answer Verification\n\n- Option A: Frankenstein\n- Option
1287
       B: Catch-22\n- Option C: Ulysses\n\nThe final answer is that Alice ends up
       owning Ulysses at the end of the semester."
1290
1291
```

Figure 13: IMap Structure Example 3

1292

1346

```
1297
1298
1299
1300
          "question": "Today, Susan went to the restaurant. Between what times could
1301
       they have gone?\nWe know that:\nSusan woke up at 11am.\nLisa saw Susan
1302
       taking photos near the Leaning Tower of Pisa from 11am to 1pm.\nWilliam saw
1303
       Susan buying cookies at a bakery from 1pm to 3pm. photos near the Leaning
1304
1305
       Tower of Pisa from 11am to 1pm.\nWilliam saw Susan buying cookies at a
       bakery from 1pm to 3pm.\nJennifer saw Susan working at the office from 3pm
1307
       to 6pm. office from 3pm to 6pm. nJason saw Susan walking towards the Statue
1308
       of Liberty from 6pm to 7pm.\nAshley saw Susan waiting at the airport from
1309
       7pm to 9pm.\nThe restaurant was closed after 10pm. nThe restaurant was closed
1310
1311
       after 10pm.\nBetween what times could Susan have gone to the
1312
       restaurant?\nOptions:\n(A) 6pm to 7pm\n(B) 9pm to 10pm\n(C) 11am to
       1pm \cdot n(D) 7pm to 9pm",
1314
          "answer": "(B)",
1315
          "COTs": "Determine the range of times that Susan goes to the restaurant. \n-
1316
1317
       It is known when Susan goes from morning to evening during the day. \n\n",
1318
          "Cases": "Known conditions:\n1. Susan wakes up at 11am. \n2. Lisa sees
1319
       Susan taking pictures near the Leaning Tower of Pisa between 11am and 1pm.
1320
       \n3. William sees Susan buying cookies at the bakery between 1PM and 3PM.
1321
       \n4. Jennifer saw Susan working at the office between 3pm and 6pm. \n5. Jason
1322
       saw Susan walking towards the Statue of Liberty between 6PM and 7PM. \n6.
1323
1324
       Ashley saw Susan waiting at the airport between 7pm and 9pm. \n7. The
       restaurant closed after 10PM. \n\n Comparative Analysis of Activity Times \n\n
1326
       Compare the activities for each time period: \n1. 11am - 1pm: Susan takes
1327
       pictures near the Leaning Tower of Pisa (seen by Lisa). \n2. 1pm - 3pm: Susan
       buying cookies at the bakery (seen by William). \n3. 3pm - 6pm: Susan working
1329
1330
       in the office (seen by Jennifer). \n4. 6pm - 7pm: Susan walking towards the
1331
       Statue of Liberty (seen by Jason). \n5. 7pm - 9pm: Susan is waiting at the
1332
       airport, so it is unlikely that she will go to the restaurant at this time of day. \n\n
1333
       Possible Time Period Verification \n\n The only possible time period is 9pm -
1334
       10pm because the restaurant was still open during this time and Susan was not
1335
1336
       seen participating in other activities.",
1337
          "Patterns": "Conclusion\n\nSusan's likely time period for going to the
1338
       restaurant is:\n\n(B) 9pm to 10pm",
1339
          "Reasonability": "Verification: \n- This time period is during the restaurant's
1340
       business hours (before 10pm). \n- Susan was not seen participating in other
1341
       activities during this time. \n\n- Therefore, option (B) 9pm to 10pm is the most
1342
1343
       reasonable answer."
1344
       }
1345
```

Figure 14: IMap Structure Example 4