

Diverge to Induce Prompting: Multi-Rationale Induction for Zero-Shot Reasoning

Po-Chun Chen¹ Hen-Hsen Huang² Hsin-Hsi Chen^{1,3}

¹Department of Computer Science and Information Engineering,
National Taiwan University, Taiwan

²Institute of Information Science, Academia Sinica, Taiwan

³AI Research Center (AINTU), National Taiwan University, Taiwan
pcchen@nlg.csie.ntu.edu.tw,
hhuang@iis.sinica.edu.tw, hhchen@ntu.edu.tw

Abstract

To address the instability of unguided reasoning paths in standard Chain-of-Thought prompting, recent methods guide large language models (LLMs) by first eliciting a single reasoning strategy. However, relying on just one strategy for each question can still limit performance across diverse tasks. We propose Diverge-to-Induce Prompting (DIP), a framework that first prompts an LLM to generate multiple diverse high-level rationales for each question. Each rationale is then elaborated into a detailed, step-by-step draft plan. Finally, these draft plans are induced into a final plan. DIP enhances zero-shot reasoning accuracy without reliance on resource-intensive sampling. Experiments show that DIP outperforms single-strategy prompting, demonstrating the effectiveness of multi-plan induction for prompt-based reasoning.

1 Introduction

Prompt-based reasoning has become a central paradigm for eliciting logical behavior in large language models (LLMs), especially in zero-shot settings (Wei et al., 2022; Kojima et al., 2022). A common approach is zero-shot chain-of-thought (CoT) prompting, which guides models to produce reasoning chains directly but leaves the reasoning process unguided.

To address this limitation, recent work proposes single-strategy prompting methods, such as Plan-and-Solve prompting (Wang et al., 2023a) and Strategic Chain-of-Thought (S-CoT; Wang et al. 2024), where the model first generates a high-level plan or strategy before producing the final answer. While efficient, these approaches typically commit to a single path per question, relying on the model’s first intuition and potentially missing better alternatives. Other methods generate multiple reasoning paths and select among them, for example through external voting, reranking, or sampling-based aggregation (Wang et al., 2022; Zheng et al., 2023;

Suzgun and Kalai, 2024), but these strategies require repeated model calls or additional selection modules, resulting in substantial computational overhead.

We introduce **Diverge-to-Induce Prompting (DIP)**, a prompting framework that first generates multiple high-level rationales, elaborates each into a step-by-step draft plan, and then induces a final plan from these draft plans, which is then used to perform the final inference. Inspired by advances in instruction induction (Honovich et al., 2023; Chen et al., 2024a), DIP enables LLMs to synthesize a high-quality, instance-level draft plan by integrating diverse perspectives. Experiments show that DIP outperforms strong baselines, including state-of-the-art single-strategy prompting methods, in most evaluation settings.

In summary, our main contributions are three-fold: (1) We propose a multi-rationale induction framework that elicits multiple high-level rationales per question, enabling LLMs to integrate diverse perspectives and avoid the blind spots of single-path reasoning; (2) Our method induces a final plan from these draft plans using multiple rationales, without relying on repeated sampling, voting, or external selection modules; and (3) Experiments on BBH and LiveBench Reasoning tasks, spanning a wide range of LLM families, demonstrate that our approach outperforms state-of-the-art single-path reasoning methods in reasoning accuracy.

2 DIP Framework

DIP comprises three main phases: (1) a Divergent Phase, where the model generates multiple high-level rationales and constructs a draft plan for each; (2) an inductive phase, which consolidates these draft plans into a final plan; and (3) an Inference Phase, which produces the final reasoning and answer. Figure 1 illustrates the process, with complete prompt templates provided in Appendix Fig-

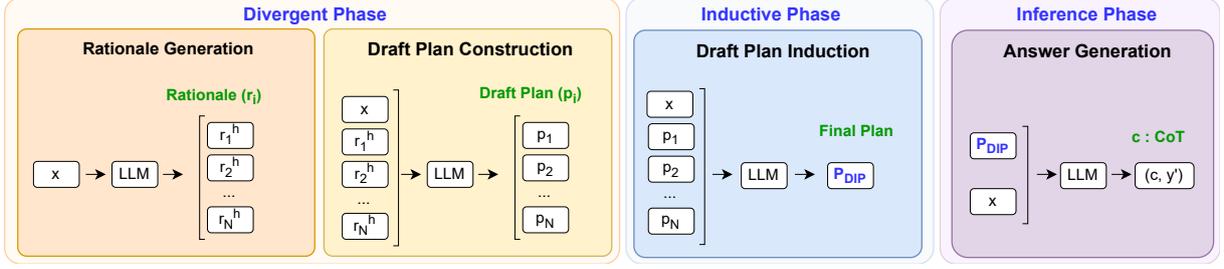


Figure 1: Overview of the DIP framework. The process consists of three main phases: a **Divergent Phase**, where the model generates multiple high-level rationales and constructs a draft plan for each; an **Inductive Phase**, where all plans are induced into a final plan; and an **Inference Phase**, which produces the final reasoning and answer.

ures 3–6. templates for rationale generation, draft plan construction, induction, and answer generation.

2.1 Rationale Generation

Given a question x , the model generates a set of N high-level rationales in a single call:

$$R = \{r_1, r_2, \dots, r_N\} = \text{LLM}(x)$$

Each r_i denotes a distinct high-level rationale for solving the question.

2.2 Draft Plan Construction

Each rationale r_i is expanded by the LLM into a corresponding step-by-step draft plan p_i . All plans are generated in a single model call by providing x and the set R as input, resulting in $P = \{p_1, p_2, \dots, p_N\}$, where each p_i corresponds to r_i .

$$P = \{p_1, p_2, \dots, p_N\} = \text{LLM}(x, R)$$

2.3 Draft Plan Induction

The induction phase induces a single final plan that incorporates the diverse perspectives found in the set of draft plans P for the same question x :

$$P_{\text{DIP}} = \text{LLM}(x, P)$$

2.4 Final Inference

The induced draft plan P_{DIP} is then used to answer the original question:

$$(c, y^*) = \text{LLM}(P_{\text{DIP}}, x)$$

where c denotes the chain-of-thought reasoning and y^* is the final predicted answer.

3 Experiments Setting

3.1 Models

We evaluate DIP across a broad range of LLMs and Large Reasoning Models (LRMs), including both open-source and proprietary systems. Our experiments span six major families: **LLaMA**, **Mistral**, **Gemini**, **GPT**, **Grok**, and the **o-Series**, covering diverse model sizes and capabilities. In particular, our evaluation includes LRMs with extended reasoning capabilities (o4 Mini and Grok 3 Mini), which employ test-time compute for enhanced reasoning. Full model names, versions, and configuration details are provided in Appendix A.

3.2 Datasets

We use the BIG-Bench Hard (BBH) benchmark (Suzgun et al., 2022), a suite of tasks curated from the BIG-Bench collection (Srivastava et al., 2022) that current LLMs struggle to solve at average human performance. We also include the reasoning-type tasks from LiveBench (White et al., 2025), a challenging and contamination-controlled benchmark comprising diverse real-world problems, to evaluate models’ high-precision reasoning abilities.

3.3 Baselines

Z-CoT We apply Zero-shot Chain-of-Thought (Z-CoT) prompting (Kojima et al., 2022), where the model is directly prompted to generate intermediate reasoning steps without demonstrations.

S-CoT Automatic Strategic Chain-of-Thought (S-CoT), proposed by Wang et al. (2024), is a state-of-the-art CoT prompting baseline. It first prompts the model to identify a high-level problem-solving strategy, which is then used to guide step-by-step reasoning.

R-CoT We design a variant of S-CoT called Rationale Chain-of-Thought (R-CoT), in which the

Model	BBH					LiveBench Reasoning				
	Z-CoT	R-CoT	S-CoT	Ours	Δ	Z-CoT	R-CoT	S-CoT	Ours	Δ
Llama 3.3 8B	69.33	61.62	64.35	71.19	+1.86	21.50	16.50	18.50	28.00	+6.50
Llama 3.3 70B	83.54	77.45	80.17	85.51	+1.97	42.00	22.50	44.00	47.00	+5.00
Llama 4 Scout	77.74	79.25	70.32	84.46	+6.72	12.50	8.00	13.50	43.00	+30.50
Llama 4 Maverick	83.59	84.64	78.67	86.20	+2.61	55.00	49.50	50.00	56.50	+1.50
Mistral Small 2	71.07	68.70	69.51	71.65	+0.58	29.50	30.50	28.00	31.00	+1.50
Mistral Small 3	78.78	77.33	77.97	81.04	+2.26	36.50	36.00	35.50	36.50	+0.00
Mistral Small 3.1	77.22	78.43	77.91	78.90	+1.68	39.00	33.50	38.00	39.50	+0.50
Mistral Medium 3	83.48	83.48	84.41	87.01	+3.53	53.00	49.00	53.50	54.50	+1.50
Mistral Large 2	84.06	82.49	82.20	84.87	+0.81	44.50	42.00	43.50	46.50	+2.00
Gemini 2.0 Flash Lite	79.25	81.28	79.71	83.54	+4.29	38.50	43.00	39.00	48.00	+9.50
Gemini 2.0 Flash	80.64	82.03	82.09	85.22	+4.58	54.50	52.00	55.00	58.50	+4.00
GPT 4o mini	79.59	76.12	78.49	82.90	+3.31	29.50	24.50	31.00	32.50	+3.00
GPT 4o	84.87	85.51	85.97	88.06	+3.19	41.50	46.50	42.00	46.00	+4.50
GPT 4.1 Nano	76.17	74.90	76.64	77.91	+1.74	35.00	29.00	29.00	32.00	-3.00
GPT 4.1 Mini	88.93	89.51	88.93	89.86	+0.93	52.00	54.50	63.00	65.00	+13.00
GPT 4.1	89.04	88.87	90.09	92.35	+3.31	65.50	64.50	63.00	70.50	+5.00
o4 Mini	89.68	91.48	90.96	91.59	+1.91	82.50	80.00	84.00	91.00	+8.50
Grok 3	89.33	87.54	89.04	90.61	+1.28	75.00	72.00	77.00	73.00	-2.00
Grok 3 Mini (High)	89.57	88.41	90.78	91.48	+1.91	84.50	83.00	86.50	92.00	+7.50
Grok 3 Mini (Low)	89.04	86.38	89.91	90.26	+1.22	82.00	81.00	84.00	84.50	+2.50

Table 1: Zero-shot performance (%) of different models under Z-CoT, R-CoT, S-CoT, our method, and their difference Δ (Ours – Z-CoT). Blue = improvement, Red = degradation.

model is prompted to generate a rationale instead of a general strategy before reasoning.

3.4 Other Details

In our experiments, we set $N = 5$, meaning that DIP generates five rationales per question. These are elaborated into draft plans and then induced into a final plan for inference.

For model parameters, we set the temperature to 0 and Top-P to 1 whenever these options were supported by the model, ensuring deterministic and reproducible results. Additional implementation and evaluation details are provided in Appendix A.

4 Results and Analysis

We evaluate DIP on BBH and LiveBench Reasoning tasks using both LLMs and LRMs across multiple model families (Table 1). DIP outperforms all baselines, achieving the highest accuracy across all models on BBH, and obtaining the best results in 17 out of 20 settings on LiveBench, demonstrating strong generalization. The performance gains of DIP exhibit different patterns across benchmarks. On BBH, improvements over Z-CoT are consistent, ranging from 0.58 to 6.72 in accuracy. On the more challenging LiveBench tasks, DIP improves performance for most models, with gains ranging from 0.5 to 30.50. In particular, Llama 4 Scout and GPT 4.1 Mini achieve notable improvements of 30.50

Model	DIP	DIP-R	Δ
Llama 4 Scout	84.46	80.23	+4.23
Llama 4 Maverick	86.20	86.03	+0.17
Mistral Small 3	81.04	79.65	+1.39
Mistral Medium 3	87.01	85.39	+1.62
Gemini 2.0 Flash Lite	83.42	84.00	-0.58
Gemini 2.0 Flash	85.22	83.36	+1.86
GPT 4.1 Mini	89.86	89.39	+0.47
O4 Mini	91.59	90.84	+0.75
Grok 3	90.61	90.61	0.00
Grok 3 Mini (Low)	90.26	89.91	+0.35

Table 2: Performance (%) on BBH. $\Delta = \text{DIP} - \text{DIP-R}$. Blue = improvement, Red = degradation.

and 13.00, respectively.

Notably, our approach surpasses both R-CoT and S-CoT in the vast majority of settings. Unlike these baselines, which generate only a single draft plan per question, DIP first explores multiple diverse rationales, constructs corresponding step-by-step plans, and then induces a final plan through draft plan induction. This multi-stage process enables improved accuracy reasoning across model types. Moreover, for complex tasks, DIP helps models such as Llama 4 Scout better follow output format requirements, addressing a common failure case under baseline prompting and further contributing to improved accuracy.

Model	BBH			
	N=1	N=3	N=5	N=7
Llama 4 Scout	83.65	84.12	84.46	82.38
Llama 4 Maverick	86.96	86.20	86.20	87.77
Mistral Small 3	80.06	78.49	81.04	81.97
Mistral Medium 3	86.78	87.88	87.01	88.17
Gemini 2.0 Flash Lite	81.04	83.71	83.42	84.17
Gemini 2.0 Flash	85.10	83.65	85.22	82.55
GPT 4.1 Mini	89.68	89.91	89.86	90.09
O4 Mini	90.49	90.72	91.59	90.03
Grok 3	90.61	90.78	90.61	91.13
Grok 3 Mini (Low)	89.39	91.83	90.26	91.13
Win N=1 Rate	-	7/10	8/10	7/10

Table 3: Performance (%) of different models on BBH as a function of the number of rationales N . For each $N > 1$, the last row reports the number of models that outperform $N = 1$.

5 Discussion

Due to cost constraints, all ablation studies in this section are conducted on two models per series.

Effect of Rationale Generation To assess the impact of rationale generation, we compare DIP with an ablation variant, DIP-R, which omits the Rationale Generation step and directly constructs multiple draft plans for induction. As shown in Table 2, DIP outperforms DIP-R on 9 out of 10 models. This confirms that prompting the model to first generate diverse rationales yields more robust and accurate plans.

Impact of Draft Plan Number We further analyze the effect of varying the number of draft plans (N) used in the induction step. As shown in Table 3, using multiple plans ($N > 1$) outperforms $N = 1$ in most settings. While the optimal N varies across models and datasets, values between 5 and 7 generally yield the best results. This suggests that moderate diversity in draft plans provides sufficient signal for effective induction. However, several models show performance degradation at higher N values (e.g., Llama 4 Scout drops from 84.46% at $N=5$ to 82.38% at $N=7$), suggesting that excessive rationales may introduce noise that interferes with the induction process.

Case Study Figure 2 presents a BBH Salient Translation Error Detection example with the Llama 4 Scout model. In this instance, “Nattern” (colubrid family) is mistranslated as “Lantern family,” constituting a named entity error. Only our

Source
Die Mexikanische Königsnatter oder San-Luis-Potosi-Königsnatter ist eine Schlange aus der Familie der Nattern.
Translation
The Mexican King’s Snake or San Luis Potosi King snake is a snake of the Lantern family.
Options
(A) Modifiers or Adjectives, (B) Numerical Values, (C) Negation or Antonyms, (D) Named Entities, (E) Dropped Content, (F) Facts
Answer
(D) Named Entities

Figure 2: BBH Salient Translation Error Detection example. Identify the error type in a German-to-English translation. Here, “Nattern” (colubrid family) is mistranslated as “Lantern family.”

full DIP setting ($N = 5$) induces a detailed plan that aligns family names and systematically checks all error types, leading to the correct answer. In contrast, both $N = 1$ and DIP-R produce generic or incomplete plans and fail on this case. All other baselines also fail to identify the correct error type. Full intermediate outputs and prompts are provided in the Appendix (see Figures 7 to 14).

Computational Cost Analysis To evaluate whether DIP’s improvements stem from effective plan induction or simply increased computation, we compare against Self-Consistency (SC) (Wang et al., 2022), a widely-used multi-path baseline. Due to cost constraints, we conduct this analysis on Llama 4 Scout and Maverick, the latest open-source Llama models. Table 4 reports token consumption, focusing on output tokens since they typically cost approximately 3× more than input tokens (Together AI, 2025).

DIP achieves substantially better cost-performance trade-offs than sampling baselines. We primarily compare against SC with $k=20$, as prior work shows performance gains typically saturate beyond this point (Wang et al., 2022). On Llama 4 Scout, DIP reaches 84.46% with 1,556 output tokens versus 84.17% with 7,533 tokens for Z-CoT+SC($k=20$) (4.8× fewer tokens). Similar efficiency gains hold on Llama 4 Maverick: DIP uses 7.2× fewer tokens than S-CoT+SC($k=20$) (1,481 vs. 10,612) while achieving higher accuracy (86.20% vs. 85.68%).

Even accounting for total token consumption including DIP’s multi-stage prompting, DIP remains substantially more efficient. For example, on Llama 4 Maverick, DIP uses 2,185 total tokens

Method	Llama 4 Scout (109B)				Llama 4 Maverick (400B)			
	Acc (%)	Tokens			Acc (%)	Tokens		
		Input	Output	Total		Input	Output	Total
Z-CoT	77.74	298.54	376.63	675.17	83.59	298.54	409.84	708.38
S-CoT	70.32	333.54	426.89	760.43	78.67	333.09	530.61	863.70
DIP	84.46	705.08	1,555.51	2,260.59	86.20	704.40	1,480.99	2,185.39
Z-CoT+SC(k=3)	82.55	895.62	1,129.89	2,025.51	85.04	895.62	1,229.52	2,125.14
Z-CoT+SC(k=5)	84.41	1,492.70	1,883.15	3,375.85	85.51	1,492.70	2,049.20	3,541.90
Z-CoT+SC(k=10)	84.17	2,985.40	3,766.30	6,751.70	85.51	2,985.40	4,098.40	7,083.80
Z-CoT+SC(k=15)	84.06	4,478.10	5,649.45	10,127.55	85.22	4,478.10	6,147.60	10,625.70
Z-CoT+SC(k=20)	84.17	5,970.80	7,532.60	13,503.40	86.09	5,970.80	8,196.80	14,167.60
S-CoT+SC(k=3)	76.64	1,000.62	1,280.67	2,281.29	84.46	999.27	1,591.83	2,591.10
S-CoT+SC(k=5)	79.83	1,667.70	2,134.45	3,802.15	85.16	1,665.45	2,653.05	4,318.50
S-CoT+SC(k=10)	82.72	3,335.40	4,268.90	7,604.30	86.09	3,330.90	5,306.10	8,637.00
S-CoT+SC(k=15)	82.90	5,003.10	6,403.35	11,406.45	85.91	4,996.35	7,959.15	12,955.50
S-CoT+SC(k=20)	83.01	6,670.80	8,537.80	15,208.60	85.68	6,661.80	10,612.20	17,274.00
DIP+SC(k=3)	84.87	2,015.02	2,242.73	4,257.75	86.38	1,861.18	2,284.95	4,146.13
DIP+SC(k=5)	84.81	3,515.50	2,929.95	6,445.45	86.38	3,209.18	3,088.91	6,298.09
DIP+SC(k=10)	84.58	7,387.43	4,648.00	12,035.43	86.90	6,698.21	5,098.81	11,797.02
DIP+SC(k=15)	84.99	12,154.98	6,366.05	18,521.03	86.38	11,082.86	7,108.71	18,191.57
DIP+SC(k=20)	84.64	16,922.53	8,084.10	25,006.63	86.55	15,467.51	9,118.61	24,586.12

Table 4: Cost-performance comparison on BBH. Input tokens measure prompt cost, output tokens measure generation cost. For DIP+SC methods, SC is applied only to the final answer generation step.

versus 14,168 for Z-CoT+SC(k=20) (6.5× fewer tokens). Moreover, by applying SC only to the final answer generation step, DIP+SC(k=10) achieves our highest accuracy of 86.90% with 11,797 total tokens, demonstrating that multi-rationale induction and sampling are complementary while maintaining competitive efficiency.

6 Related Work

Inductive Capacities of LLMs Recent studies affirm that LLMs possess strong inductive reasoning abilities, allowing them to abstract rules or instructions from limited examples and generalize to new instances (Wang et al., 2023b; Zhu et al., 2023; Cheng et al., 2024; Sun et al., 2024; He et al., 2024; Yang et al., 2022; Honovich et al., 2022; Chen et al., 2024b, 2025). Most prior work focuses on rule induction at the dataset or task level. By contrast, our work explores **instance-level induction**: for each question, we prompt the model to induce a final solution from multiple draft plans, leveraging LLMs’ inductive capacity at the instance level.

Plan Generation by LLMs LLMs have also been used to generate draft plans before problem solving. Plan-and-Solve and Strategic Chain-of-Thought (S-CoT) prompt the model to produce a high-level plan or strategy to guide step-by-step reasoning (Wang et al., 2023a, 2024). Search-based methods such as Tree-of-Thoughts (ToT), Reason-

ing via Planning, and Graph-of-Thoughts expand multiple reasoning paths (Yao et al., 2023; Hao et al., 2023; Besta et al., 2024).

These search-based approaches, exemplified by ToT, adopt a “learning by doing” paradigm: they immediately execute reasoning steps, explore multiple paths in parallel, and iteratively evaluate partial solutions to guide further exploration. In contrast, DIP follows a “planning before doing” approach. We first generate multiple diverse abstract planning strategies without execution, synthesize them into a final plan through induction, and only then perform a single execution. This approach achieves both the diversity of multi-plan exploration and the efficiency of single-pass inference.

7 Conclusion

We present DIP, a prompting framework that enhances zero-shot reasoning in large language models by first generating multiple diverse high-level rationales, elaborating a draft plan for each, and then inducing a final plan from these. Experiments on BBH and LiveBench Reasoning tasks show that DIP outperforms strong baselines, including state-of-the-art single-strategy prompting methods, across diverse models and tasks. These results demonstrate that multi-rationale induction offers a practical and effective approach for more robust prompt-based reasoning at the instance level.

8 Limitations

While DIP demonstrates improvements on BBH and LiveBench Reasoning tasks, it has several limitations. First, generating and processing multiple rationales and plans per question leads to higher computational cost and more API calls than basic prompting baselines such as Z-CoT. This overhead may limit applicability in cost- or latency-sensitive scenarios, including large-scale or real-time applications. Second, DIP’s performance gains are mainly observed on tasks requiring complex, multi-step reasoning. For tasks that involve only straightforward fact retrieval or commonsense question answering, the framework may offer limited or no benefit, as generating diverse rationales can introduce unnecessary complexity. Finally, our evaluation is restricted to English benchmarks and a subset of LLM architectures. The generalizability of DIP to multilingual tasks, domain-specific reasoning, or other model families remains to be validated.

Acknowledgements

This work was supported by National Science and Technology Council, Taiwan, under grants NSTC 113-2634-F-002-003- and 114-2221-E-002-070-MY3, and Ministry of Education (MOE), Taiwan, under grants NTU-113L9009, NTU-114L9009, and NTU-114L900901.

References

- Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michał Podstawski, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczyk, and Torsten Hoefler. 2024. [Graph of thoughts: solving elaborate problems with large language models](#). In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI’24/IAAI’24/EAAI’24. AAAI Press.
- Po-Chun Chen, Hen-Hsen Huang, and Hsin-Hsi Chen. 2025. [Vqa-induct: Instruction induction for visual question answering](#). In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, CIKM ’25, page 4659–4664, New York, NY, USA. Association for Computing Machinery.
- Po-Chun Chen, Sheng-Lun Wei, Hen-Hsen Huang, and Hsin-Hsi Chen. 2024a. [Induct-learn: Short phrase prompting with instruction induction](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5204–5231, Miami, Florida, USA. Association for Computational Linguistics.
- Po-Chun Chen, Sheng-Lun Wei, Hen-Hsen Huang, and Hsin-Hsi Chen. 2024b. [Induct-learn: Short phrase prompting with instruction induction](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Kewei Cheng, Jingfeng Yang, Haoming Jiang, Zhengyang Wang, Binxuan Huang, Ruirui Li, Shiyang Li, Zheng Li, Yifan Gao, Xian Li, Bing Yin, and Yizhou Sun. 2024. [Inductive or deductive? rethinking the fundamental reasoning abilities of llms](#). *ArXiv*, abs/2408.00114.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony S. Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, and 510 others. 2024. [The llama 3 herd of models](#). *ArXiv*, abs/2407.21783.
- Google. 2025. [Gemini 2.0: Flash, flash-lite and pro](#). FEB. 5, 2025.
- Shibo Hao, Yi Gu, Haodi Ma, Joshua Hong, Zhen Wang, Daisy Wang, and Zhiting Hu. 2023. [Reasoning with language model is planning with world model](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8154–8173, Singapore. Association for Computational Linguistics.
- Kaiyu He, Mian Zhang, Shuo Yan, Peilin Wu, and Zhiyu Zoey Chen. 2024. [Idea: Enhancing the rule learning ability of large language model agent through induction, deduction, and abduction](#).
- Or Honovich, Uri Shaham, Samuel R. Bowman, and Omer Levy. 2022. [Instruction induction: From few examples to natural language task descriptions](#). *ArXiv*, abs/2205.10782.
- Or Honovich, Uri Shaham, Samuel R. Bowman, and Omer Levy. 2023. [Instruction induction: From few examples to natural language task descriptions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1935–1952, Toronto, Canada. Association for Computational Linguistics.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). *ArXiv*, abs/2205.11916.
- MetaAI. 2025a. [Introducing llama api](#). 2025.
- MetaAI. 2025b. [The llama 4 herd: The beginning of a new era of natively multimodal ai innovation](#). April 5, 2025.

- MistralAI. 2024a. [Ai in abundance](#). Sep 17, 2024.
- MistralAI. 2024b. [Pixtral large](#). Nov 18, 2024.
- MistralAI. 2025a. [Medium is the new large](#). May 7, 2025.
- MistralAI. 2025b. [Mistral small 3](#). Jan 30, 2025.
- MistralAI. 2025c. [Mistral small 3.1](#). Mar 17, 2025.
- OpenAI. 2024a. [Gpt-4o mini: advancing cost-efficient intelligence](#). July 18, 2024.
- OpenAI. 2024b. [Hello gpt-4o](#). May 13, 2024.
- OpenAI. 2025a. [Introducing gpt-4.1 in the api](#). April 14, 2025.
- OpenAI. 2025b. [Introducing openai o3 and o4-mini](#). April 16, 2025.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, Agnieszka Kluska, Aitor Lewkowycz, Akshat Agarwal, Alethea Power, Alex Ray, Alex Warstadt, Alexander W. Kocurek, Ali Safaya, Ali Tazarv, and 431 others. 2022. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *ArXiv*, abs/2206.04615.
- Wangtao Sun, Haotian Xu, Xuanqing Yu, Pei Chen, Shizhu He, Jun Zhao, and Kang Liu. 2024. [ItD: Large language models can teach themselves induction through deduction](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Mirac Suzgun and Adam Tauman Kalai. 2024. [Meta-prompting: Enhancing language models with task-agnostic scaffolding](#). *ArXiv*, abs/2401.12954.
- Mirac Suzgun, Nathan Scales, Nathanael Scharli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V. Le, Ed Huai hsin Chi, Denny Zhou, and Jason Wei. 2022. [Challenging big-bench tasks and whether chain-of-thought can solve them](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Together AI. 2025. [Together ai pricing](#). Accessed: May 2025.
- Lei Wang, Wanyu Xu, Yihuai Lan, Zhiqiang Hu, Yunshi Lan, Roy Ka-Wei Lee, and Ee-Peng Lim. 2023a. [Plan-and-solve prompting: Improving zero-shot chain-of-thought reasoning by large language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2609–2634, Toronto, Canada. Association for Computational Linguistics.
- Ruo Cheng Wang, E. Zelikman, Gabriel Poesia, Yewen Pu, Nick Haber, and Noah D. Goodman. 2023b. [Hypothesis search: Inductive reasoning with language models](#). *ArXiv*, abs/2309.05660.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed H. Chi, and Denny Zhou. 2022. [Self-consistency improves chain of thought reasoning in language models](#). *ArXiv*, abs/2203.11171.
- Yu Wang, Shiwan Zhao, Zhihu Wang, Heyuan Huang, Ming Fan, Yubo Zhang, Zhixing Wang, Haijun Wang, and Ting Liu. 2024. [Strategic chain-of-thought: Guiding accurate reasoning in llms through strategy elicitation](#). *ArXiv*, abs/2409.03271.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed H. Chi, F. Xia, Quoc Le, and Denny Zhou. 2022. [Chain of thought prompting elicits reasoning in large language models](#). *ArXiv*, abs/2201.11903.
- Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Benjamin Feuer, Siddhartha Jain, Ravid Shwartz-Ziv, Neel Jain, Khalid Saifullah, Sreemanti Dey, Shubh-Agrawal, Sandeep Singh Sandha, Siddhartha Venkat Naidu, Chinmay Hegde, Yann LeCun, Tom Goldstein, Willie Neiswanger, and Micah Goldblum. 2025. [Livebench: A challenging, contamination-free LLM benchmark](#). In *The Thirteenth International Conference on Learning Representations*.
- xAI. 2025. [Grok 3 beta — the age of reasoning agents](#). February 19, 2025.
- Zonglin Yang, Li Dong, Xinya Du, Hao Cheng, E. Cambria, Xiaodong Liu, Jianfeng Gao, and Furu Wei. 2022. [Language models as inductive reasoners](#). *ArXiv*, abs/2212.10923.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. [Tree of thoughts: deliberate problem solving with large language models](#). In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23*, Red Hook, NY, USA. Curran Associates Inc.
- Huaixiu Steven Zheng, Swaroop Mishra, Xinyun Chen, Heng-Tze Cheng, Ed Huai hsin Chi, Quoc V. Le, and Denny Zhou. 2023. [Take a step back: Evoking reasoning via abstraction in large language models](#). *ArXiv*, abs/2310.06117.
- Zhaocheng Zhu, Yuan Xue, Xinyun Chen, Denny Zhou, Jian Tang, Dale Schuurmans, and Hanjun Dai. 2023. [Large language models can learn rules](#). *ArXiv*, abs/2310.07064.

A Experimental Configurations

Model List and Cost. Table 5 details all model names, exact versions, API providers, and their relative API costs as used in our experiments. The cost column reflects the relative pricing or prompt quota at the time of writing.

Table 7 summarizes the estimated total cost for all experiments, which was approximately \$698, based on the providers’ pricing during the experimental period.

Model Settings. For all experiments, we set the temperature to 0 and Top-P to 1, whenever these options were supported by the model.

Other Settings. To facilitate fair ablation studies on different numbers of rationales and draft plans, we designed our experiments such that all rationales and their corresponding draft plans were generated together in a single batch. Specifically, we prompted the model to produce $N = 9$ rationales and their associated draft plans at once. For each ablation setting, we used only the first k rationales and their corresponding draft plans for induction, ensuring that the initial k inputs were always identical across settings. This setup guarantees that any variation in performance can be attributed solely to the number of rationales and draft plans, rather than differences in their content.

B Use of AI Assistants

We used ChatGPT for grammar refinement and occasional assistance in code development. All language model outputs served solely as references; the final writing and code were thoroughly reviewed and authored by the authors.

C Prompt Templates and Examples

Figures 3 to 6 present the exact prompt templates used at each stage of our DIP framework, covering Rationale Generation, Draft Plan Construction, Draft Plan Induction, and Answer Generation. Each template is provided in full to ensure transparency and reproducibility.

Figures 7 to 14 contain the complete case study examples referenced in the main paper. These examples include all intermediate outputs from the ablation variants and baseline methods, as evaluated using the Llama 4 Scout model. Only the complete DIP method produces the correct answer for these challenging cases, while all other ablations and baselines fail to do so.

Model	Model Version	API Provider	cost
Llama 3.3 8B (MetaAI, 2025a)	Llama-3.3-8B-Instruct	Meta Llama API	-
Llama 3.3 70B (Dubey et al., 2024)	Llama-3.3-70B-Instruct	Meta Llama API	-
Llama 4 Scout (MetaAI, 2025b)	Llama-4-Scout-17B-16E-Instruct-FP8	Meta Llama API	-
Llama 4 Maverick (MetaAI, 2025b)	Llama-4-Maverick-17B-128E-Instruct-FP8	Meta Llama API	-
Mistral Small 2 (MistralAI, 2024a)	mistral-small-2409	Mistral AI	3
Mistral Small 3 (MistralAI, 2025b)	mistral-small-2501	Mistral AI	3
Mistral Small 3.1 (MistralAI, 2025c)	mistral-small-2503	Mistral AI	3
Mistral Medium 3 (MistralAI, 2025a)	mistral-medium-2505	Mistral AI	25
Mistral Large 2 (MistralAI, 2024b)	mistral-large-2411	Mistral AI	75
Gemini 2.0 Flash Lite (Google, 2025)	gemini-2.0-flash-lite-001	Google	8
Gemini 2.0 Flash (Google, 2025)	gemini-2.0-flash-001	Google	10
GPT 4o mini (OpenAI, 2024a)	gpt-4o-mini-2024-07-18	OpenAI	6
GPT 4o (OpenAI, 2024b)	gpt-4o-2024-11-20	OpenAI	100
GPT 4.1 nano (OpenAI, 2025a)	gpt-4.1-nano-2025-04-1	OpenAI	4
GPT 4.1 mini (OpenAI, 2025a)	gpt-4.1-mini-2025-04-14	OpenAI	16
GPT 4.1 (OpenAI, 2025a)	gpt-4.1-2025-04-14	OpenAI	80
o4 mini (OpenAI, 2025b)	o4-mini-2025-04-16	OpenAI	125
Grok 3 Beta (xAI, 2025)	grok-3-beta	xAI	200
Grok 3 Mini Beta (xAI, 2025)	grok-3-mini-beta	xAI	40

Table 5: Mapping of model names to detailed versions, API provider, and relative API cost (in prompt units; lower means cheaper). Note: At the time of writing, the Meta Llama API was in preview and provided free of charge.

Model	BBH					LiveBench Reasoning				
	N=1	N=3	N=5	N=7	N=9	N=1	N=3	N=5	N=7	N=9
Llama 4 Scout	83.65	84.12	84.46	82.38	82.20	43.00	44.00	43.00	46.00	44.00
Llama 4 Maverick	86.96	86.20	86.20	87.77	87.48	55.00	56.50	56.50	56.00	60.50
Mistral Small 3	80.06	78.49	81.04	81.97	80.81	34.00	31.50	36.50	40.50	30.50
Mistral Medium 3	86.78	87.88	87.01	88.17	88.81	53.50	53.50	54.50	59.00	55.00
Gemini 2.0 Flash Lite	81.04	83.71	83.54	84.17	83.42	43.50	44.50	48.00	45.50	43.50
Gemini 2.0 Flash	85.10	83.65	85.22	82.55	82.55	50.00	54.00	58.50	57.50	54.50
GPT 4.1 Mini	89.68	89.91	89.86	90.09	88.81	64.50	67.50	65.00	68.00	65.50
O4 Mini	90.49	90.72	91.59	90.03	90.61	86.50	89.00	91.00	87.50	88.00
Grok 3	90.61	90.78	90.61	91.13	91.48	70.00	69.00	73.00	71.00	71.50
Grok 3 Mini (Low)	89.39	91.83	90.26	91.13	91.30	82.00	82.50	84.50	83.50	85.50
Win N=1 Rate	-	7/10	8/10	7/10	7/10	-	7/10	9/10	10/10	8/10

Table 6: Performance (%) of different models on BBH and LiveBench Reasoning, as a function of the number of rationales N . For each $N > 1$, the last row reports the number of models that outperform $N = 1$.

You are an reasoning expert with outstanding data observation and analysis skills. Your expertise includes designing distinct high-level reasoning frameworks to approach and solve complex problems.

Your current task is to generate **nine** different high-level reasoning frameworks for solving the given question. Each framework should propose a distinct way to guide an LLM's reasoning process, and be sufficiently different in its overall approach or perspective. Avoid low-level operational instructions or implementation steps—focus only on the strategic level.

The input you receive includes the following components:

- <task_information>: Provides specific details about the task.
- <answer_format>: Specifies the required answer format.
- <question>: Contains the question that needs to be addressed.

When formulating each reasoning framework, please consider:

- Clearly outline the overall strategy or approach to solve the problem.
- Focus on intuitive and easily distinguishable reasoning paths.
- Avoid listing procedural steps or implementation details.
- Ensure that each framework presents a unique way of thinking about the problem, not just a minor variation of another.

Your final output **must** strictly follow the exact format below, including the **nine** reasoning frameworks and their structure. Do not omit or modify any part of the format.

<reasoning_frameworks>

<reasoning_framework_1>

[Your brief summary of the unique approach or perspective]

</reasoning_framework_1>

<reasoning_framework_2>

[Your brief summary of the unique approach or perspective]

</reasoning_framework_2>

[...]

<reasoning_framework_9>

[Your brief summary of the unique approach or perspective]

</reasoning_framework_9>

</reasoning_frameworks>

Important Notes:

- You must generate exactly **nine** distinct high-level reasoning frameworks.
- Do not generate an answer to the given question.
- Do not include step-by-step breakdowns or operational instructions.
- Follow the format precisely.

Here is the task information, the answer format, and the question you need to generate reasoning_frameworks:

<task_information>

{task_information}

</task_information>

<answer_format>

{answer_format}

</answer_format>

<question>

{question}

</question>

Figure 3: Prompt template used for Rationale Generation in the DIP framework.

You are an reasoning expert with strong abilities in data observation, analysis, and designing structured reasoning steps. Your expertise includes creating detailed procedural steps based on high-level reasoning frameworks to tackle and solve complex problems.

Your current task is to take the given high-level reasoning frameworks and design step-by-step reasoning procedures that allow an LLM to follow them systematically in order to solve the provided problems.

The input you receive includes the following components:

- ``<task_information>``: Provides specific details about the task.
- ``<answer_format>``: Specifies the required answer format.
- ``<question>``: Contains the question that needs to be addressed.
- ``<reasoning_framework_list>``: A list of high-level reasoning frameworks. Your job is to expand each framework into detailed step-by-step procedures that guide the LLM in solving the task.

When formulating each reasoning framework, please consider:

- Ensure each reasoning framework provides a step-by-step breakdown that progressively refines the solution. Avoid overly broad steps—whenever a step involves multiple decisions, consider breaking it down into smaller steps.
- Clearly outline the overall strategy for solving the problem, based on the corresponding high-level framework summary.
- Break down the solution process into multiple structured steps, ensuring each step is detailed enough to guide a LLM in reasoning through the problem.
- If a step involves multiple logical decisions, consider dividing it into sub-steps for clarity.
- Ensure that each detailed framework remains distinct, reflecting the unique perspective of its high-level summary.
- Take into account the provided task information and answer format to ensure that your frameworks align with the specific requirements of the task.

Your final output **must** strictly follow the exact format below, including the **nine** detailed reasoning frameworks and their structure. Do not omit or modify any part of the format.

`<reasoning_frameworks>`

`<reasoning_framework_1>`

Step 1: [Concise description of the first detailed reasoning step based on high-level framework 1]

Step 2: [Concise description of the second detailed reasoning step]

Step 3: [Concise description of the third detailed reasoning step]

[Add additional steps if needed for clarity and completeness]

`</reasoning_framework_1>`

`<reasoning_framework_2>`

Step 1: [Concise description of the first detailed reasoning step based on high-level framework 2]

Step 2: [Concise description of the second detailed reasoning step]

Step 3: [Concise description of the third detailed reasoning step]

[Add additional steps if needed]

`</reasoning_framework_2>`

[...]

`<reasoning_framework_9>`

Step 1: [Concise description of the first detailed reasoning step based on high-level framework 9]

Step 2: [Concise description of the second detailed reasoning step]

Step 3: [Concise description of the third detailed reasoning step]

[Add additional steps if needed]

`</reasoning_framework_9>`

`</reasoning_frameworks>`

****Important Notes:****

- You must generate exactly **nine** distinct detailed reasoning frameworks.
- Do not generate an answer to the given question. Focus only on constructing the detailed reasoning frameworks based on the provided high-level summaries.
- Ensure the steps are detailed and procedural.
- Follow the format precisely.

Here is the task information, the answer format, the question, and the list of high-level reasoning framework summaries you need to expand into detailed reasoning steps:

`<task_information>`

{task_information}

`</task_information>`

`<answer_format>`

{answer_format}

`</answer_format>`

`<question>`

{question}

`</question>`

`<reasoning_framework_list>`

{reasoning_framework_list}

`</reasoning_framework_list>`

Figure 4: Prompt template used for Draft Plan Generation in the DIP framework.

You are an expert in Natural Language Processing with exceptional data observation and analysis skills. Your expertise includes inductively deriving improved solution rules from provided examples. In this task, you are given task information, an answer format, a specific question, and a list of example reasoning frameworks (reasoning_framework_list). Your objective is to analyze all these inputs and inductively synthesize a detailed and complete reasoning framework that will guide how to solve similar tasks.

Your inductive reasoning should consider:

- The provided task information (e.g., details about the task's subject matter)
- The answer format (e.g., the required structure or type of output)
- The specific question or problem statement
- The reasoning_framework_list, which contains example reasoning frameworks

Your output must include:

1. A clear and detailed <task_instructions> block that contains:
 - **Task Content:** A definition of the task's purpose along with the required activities.
 - **Input Format:** Detailed descriptions of the accepted data types, their formats, and instructions for processing these inputs.
2. A comprehensive <reasoning_framework> block that includes high-level reasoning steps which generalize the approach to solving similar tasks. Ensure that these steps are concise, actionable, and that they support deriving an answer that strictly conforms to the provided answer format.

Present your final output strictly in the following format:

<task_instructions>

Task Content:

[Provide your definition of Task Content here.]

Input Format:

[Provide your description of Input Format here.]

</task_instructions>

<reasoning_framework>

Step 1: [Concise description of the first high-level reasoning step]

Step 2: [Concise description of the second high-level reasoning step]

Step 3: [Concise description of the third high-level reasoning step]

[Add additional steps if needed]

</reasoning_framework>

Note:

Your goal is solely to formulate clear task instructions and a generalizable reasoning framework based on the provided task information, answer format, question, and reasoning_framework_list. Do not attempt to answer the question itself.

Here is the task information, the answer format, the question, and reasoning_framework_list you need to induct reasoning_framework:

<task_information>

{task_information}

</task_information>

<answer_format>

{answer_format}

</answer_format>

<question>

{question}

</question>

<reasoning_framework_list>

{reasoning_framework_list}

</reasoning_framework_list>

Figure 5: Prompt template used for Draft Plan Induction in the DIP framework.

```
<task_instructions>
{task_instructions}
</task_instructions>

<reasoning_framework>
{reasoning_framework}
</reasoning_framework>

<output_format>
{output_format}
</output_format>

Follow these steps carefully:
1. Based on the `reasoning_framework`, provide a step-by-step chain of thought that answers the question. Present your chain of thought in the following format:
  <chain_of_thought>
  [Your step-by-step chain of thought here]
  </chain_of_thought>

2. Based on your `chain_of_thought`, provide the final answer according to the rules specified in the `output_format` section. If unsure and "Output Format" is option, guess the closest option. Present your final answer in the following format:
  <final_answer>
  [Your final answer here]
  </final_answer>

**Note**
- Do not use programming or code to solve this question.
- It's crucial to follow this structure and include both the chain of thought and the final answer in your response.

<question>
{question}
</question>
```

Figure 6: Prompt template used for Answer Generation in the DIP framework.

Method	Rationales or Draft Plan
DIP N=5	<p>Rationales</p> <ol style="list-style-type: none"> (1) Analyze the translation by comparing the source and target sentences to identify changes in named entities, numerical values, modifiers, negation, facts, or dropped content. (2) Focus on identifying changes in the family of the snake in the source and translation sentences to determine the type of error. (3) Examine name changes such as “Mexikanische Königsnatter” → “Mexican King’s Snake” and “San-Luis-Potosi-Königsnatter” → “San Luis Potosi King snake”. (4) Investigate whether the adjective/modifier for the family term shifts from “Nattern” to “Lantern family”. (5) Check for any negation or antonyms introduced in the translation that could change the sentence meaning. <hr/> <p>Draft Plan</p> <p>Plan 1: Step 1: Compare the source sentence “Die Mexikanische Königsnatter . . .” with the translation “The Mexican King’s Snake . . .” to identify changes. Step 2: Analyze the translation for changes in named entities (snake and family names). Step 3: Check for changes in numerical values, modifiers, negation, facts, or dropped content. Step 4: Determine the error type based on the identified changes.</p> <p>Plan 2: Step 1: Identify the family of the snake in the source sentence (“der Nattern”). Step 2: Identify the family of the snake in the translation (“the Lantern family”). Step 3: Compare the two families to determine if there is a change. Step 4: If a change exists, classify it as an error in named entities or modifiers/adjectives.</p> <p>Plan 3: Step 1: Examine the translation of “Mexikanische Königsnatter” → “Mexican King’s Snake” for named-entity changes. Step 2: Examine “San-Luis-Potosi-Königsnatter” → “San Luis Potosi King snake” for named-entity changes. Step 3: Decide if the changes are significant enough to be classified as a named-entity error.</p> <p>Plan 4: Step 1: Identify the modifier/adjective for the family term in the source (“Nattern”). Step 2: Identify the modifier/adjective for the family term in the translation (“Lantern family”). Step 3: Compare the two modifiers/adjectives. Step 4: Classify the change as a modifier/adjective error.</p> <p>Plan 5: Step 1: Check the source sentence for any negations or comparatives. Step 2: Check the translation for any negations or comparatives. Step 3: Compare both sentences to determine if there is a change in negation or antonyms.</p> <hr/> <p>DIP-R N=5</p> <p>Draft Plan</p> <p>Plan 1: Step 1: Identify the key elements in the source sentence and the translation. Step 2: Compare the family names (“Familie der Nattern” vs. “Lantern family”). Step 3: Determine the type of error based on this comparison.</p> <p>Plan 2: Step 1: Analyze the source sentence for specific details. Step 2: Examine the translation for changes in named entities, numerical values, modifiers, negation, or dropped content. Step 3: Classify the error based on the identified changes.</p> <p>Plan 3: Step 1: Check for changes in named entities between the source and translation. Step 2: Verify changes in numerical values, modifiers, or negation. Step 3: Identify if any content has been dropped or if there is a factual error.</p> <p>Plan 4: Step 1: Compare the source sentence and translation for changes in names or locations. Step 2: Examine whether there are changes in numerical values or units. Step 3: Determine if the error involves modifiers, negation, or factual errors.</p> <p>Plan 5: Step 1: Identify the error by comparing the family names in both sentences. Step 2: Classify the type of error (e.g., named entities, numerical values, modifiers). Step 3: Confirm the error type based on the comparison.</p>

Figure 7: BBH Salient Translation Error Detection example: All rationales and plans generated by DIP ($N = 5$) and DIP-R ($N = 5$), where DIP-R omits the Rationale Generation step.

Mthod	Inducted Draft Plan
DIP (N=5)	<p>Step 1: Compare the source sentence “Die Mexikanische Königsnatter oder San-Luis-Potosi-Königsnatter ist eine Schlange aus der Familie der Nattern” with the translation “The Mexican King’s Snake or San Luis Potosi King snake is a snake of the Lantern family” to identify any changes.</p> <p>Step 2: Analyze the translation for changes in named entities, such as the names of the snake and its family.</p> <p>Step 3: Check for changes in numerical values, modifiers, negation, facts, or dropped content.</p> <p>Step 4: Determine the type of error based on the identified changes.</p> <p>More specifically:</p> <ul style="list-style-type: none"> – Identify the family of the snake in the source sentence (“der Nattern”) and in the translation (“the Lantern family”). – Compare the two families to determine if there is a change. – Examine the translation of “Mexikanische Königsnatter” to “Mexican King’s Snake” and “San-Luis-Potosi-Königsnatter” to “San Luis Potosi King snake” for any changes in named entities. – Check if the changes are significant enough to be classified as an error in named entities or modifiers/adjectives. – Verify if there are any changes in negation or antonyms. – Verify if there are any changes in numerical values or facts. – Verify if any significant clause is dropped in the translation. <p>Based on these steps, classify the error into one of the provided options: (A) Modifiers or Adjectives, (B) Numerical Values, (C) Negation or Antonyms, (D) Named Entities, (E) Dropped Content, (F) Facts.</p>
DIP-R (N=5)	<p>Step 1: Identify the key elements in the source sentence and the translation.</p> <p>Step 2: Compare the key elements such as family names, entities, numerical values, modifiers, and content to identify any changes.</p> <p>Step 3: Determine the type of error based on the comparison, considering the categories: Named Entities, Numerical Values, Modifiers or Adjectives, Negation or Antonyms, Facts, and Dropped Content.</p> <p>Step 4: Classify the error into one of the provided options (A, B, C, D, E, F) based on the identified change.</p> <p>Step 5: Confirm the error type based on the comparison and classification.</p>
DIP (N=1)	<p>Step 1: Compare the source sentence with the translation to identify any changes in terms of named entities, numerical values, modifiers or adjectives, negation or antonyms, facts, and dropped content.</p> <p>Step 2: Analyze the translation for changes in named entities, such as the names of the snake and its family, and check if they match the source sentence.</p> <p>Step 3: Check for changes in numerical values, modifiers, negation, facts, or dropped content by carefully examining the source and translated sentences.</p> <p>Step 4: Determine the type of error based on the identified changes by matching them with the given error types: Named Entities, Numerical Values, Modifiers or Adjectives, Negation or Antonyms, Facts, and Dropped Content.</p> <p>Step 5: Select the correct option from the provided list that corresponds to the identified error type.</p>

Figure 8: BBH Salient Translation Error Detection example: Final induced draft plans for each method.

DIP Prompt
<task_instructions>
Task Content:
The task requires detecting the type of error in an English translation of a German source sentence. The error types include Named Entities, Numerical Values, Modifiers or Adjectives, Negation or Antonyms, Facts, and Dropped Content.
Input Format:
- The source sentence in German
- The translated sentence in English
- A list of possible error types
- A set of answer options corresponding to the error types
The task involves comparing the source sentence and the translation to identify any changes and determining the type of error based on the identified changes.
</task_instructions>

Reasoning Framework
<reasoning_framework>
Step 1: Compare the source sentence "Die Mexikanische Königsnatter oder San-Luis-Potosi-Königsnatter ist eine Schlange aus der Familie der Nattern" with the translation "The Mexican King's Snake or San Luis Potosi King snake is a snake of the Lantern family" to identify any changes.
Step 2: Analyze the translation for changes in named entities, such as the names of the snake and its family.
Step 3: Check for changes in numerical values, modifiers, negation, facts, or dropped content.
Step 4: Determine the type of error based on the identified changes.
More specifically:
- Identify the family of the snake in the source sentence ("der Nattern") and in the translation ("the Lantern family").
- Compare the two families to determine if there is a change.
- Examine the translation of "Mexikanische Königsnatter" to "Mexican King's Snake" and "San-Luis-Potosi-Königsnatter" to "San Luis Potosi King snake" for any changes in named entities.
- Check if the changes are significant enough to be classified as an error in named entities or modifiers/adjectives.
- Verify if there are any changes in negation or antonyms.
- Verify if there are any changes in numerical values or facts.
- Verify if any significant clause is dropped in the translation.
Based on these steps, classify the error into one of the provided options:
(A) Modifiers or Adjectives (B) Numerical Values (C) Negation or Antonyms (D) Named Entities (E) Dropped Content (F) Facts
</reasoning_framework>

Prompt Template
<output_format>
The output is a single letter corresponding to the correct option from the list (e.g., A, B, C, etc.).
</output_format>
Follow these steps carefully:
1. Based on the reasoning_framework, provide a step-by-step chain of thought that answers the question. Present your chain of thought in the following format:
<chain_of_thought>
[Your step-by-step chain of thought here]
</chain_of_thought>
2. Based on your chain_of_thought, provide the final answer according to the rules specified in the output_format section.
If unsure and "Output Format" is option, guess the closest option. Present your final answer in the following format:
<final_answer>
[Your final answer here]
</final_answer>

Note:
- Do not use programming or code to solve this question.
- It's crucial to follow this structure and include both the chain of thought and the final answer in your response.

Question
<question>
The following translations from German to English contain a particular error. That error will be one of the following types: Named Entities: An entity (names, places, locations, etc.) is changed to a different entity. Numerical Values: Numerical values (ordinals or cardinals), dates, and/or units are changed. Modifiers or Adjectives: The modifiers and adjectives pertaining to a noun are changed. Negation or Antonyms: Introduce or remove a negation or change comparatives to their antonyms. Facts: Trivial factual errors not pertaining to the above classes are introduced in the translations. Dropped Content: A significant clause in the translation is removed. Please identify that error. Source: Die Mexikanische Königsnatter oder San-Luis-Potosi-Königsnatter ist eine Schlange aus der Familie der Nattern.
Translation: The Mexican King's Snake or San Luis Potosi King snake is a snake of the Lantern family.
The translation contains an error pertaining to
Options:
(A) Modifiers or Adjectives (B) Numerical Values (C) Negation or Antonyms (D) Named Entities (E) Dropped Content (F) Facts
</question>

Model Response
<chain_of_thought>
1. **Identify the source sentence and the translation:**
- Source sentence: "Die Mexikanische Königsnatter oder San-Luis-Potosi-Königsnatter ist eine Schlange aus der Familie der Nattern."
- Translation: "The Mexican King's Snake or San Luis Potosi King snake is a snake of the Lantern family."
2. **Compare named entities:**
- The source sentence mentions "Mexikanische Königsnatter" and "San-Luis-Potosi-Königsnatter" as the names of the snakes.
- The translation mentions "Mexican King's Snake" and "San Luis Potosi King snake," which are similar but not exactly the same due to capitalization and spacing differences, but essentially convey the same information.
3. **Check for changes in family names:**
- The source sentence states the snake is from "der Familie der Nattern."
- The translation states the snake is from "the Lantern family."
- The family name has changed from "Nattern" to "Lantern family."
4. **Determine the type of error:**
- The change from "der Nattern" (which translates to "the colubrids" or a family of snakes) to "the Lantern family" seems to be a change in a named entity or a factual piece of information about the snake's family.
5. **Analyze each option:**
- **Named Entities (D):** The family name "Nattern" (colubrids) has been changed to "Lantern family," which could be considered a change in a named entity.
- **Modifiers or Adjectives (A):** There are no obvious changes in modifiers or adjectives that pertain to a noun in a way that would classify as an error here.
- **Numerical Values (B):** There are no numerical values mentioned in the source or translation.
- **Negation or Antonyms (C):** There is no introduction or removal of a negation.
- **Facts (F):** The change in the family name could also be considered a factual error.
- **Dropped Content (E):** No significant clause appears to be dropped.
6. **Conclusion:**
- The most direct error is the change in the family name from "Nattern" to "Lantern family," which can be classified as a change in a named entity or a factual error. However, given that "Lantern family" is not a recognized family of snakes and "Nattern" refers to a specific family (colubrids), this seems to be a clear case of changing a named entity or introducing a factual error.
7. **Final Classification:**
- Between **Named Entities (D)** and **Facts (F)**, the error seems to lean more towards **Named Entities (D)** because it involves changing the name of a specific family of snakes.
</chain_of_thought>
<final_answer>
D
</final_answer>

Figure 9: Example of the DIP prompt (N=5) and reasoning for translation error detection in the BBH-Induct dataset. The gold (correct) answer for this example is (D) Named Entities.

DIP Prompt

<task_instructions>

Task Content:

The task requires detecting the type of error in an English translation of a German source sentence. The error types to be considered are: Named Entities, Numerical Values, Modifiers or Adjectives, Negation or Antonyms, Facts, and Dropped Content.

Input Format:

- The source sentence in German.
- The translated sentence in English.
- A list of possible error types.

The task involves comparing the source sentence and the translation to identify the error type. The output is a single letter corresponding to the correct option from the list.

</task_instructions>

Reasoning Framework

<reasoning_framework>

Step 1: Identify the key elements in the source sentence and the translation.

Step 2: Compare the key elements such as family names, entities, numerical values, modifiers, and content to identify any changes.

Step 3: Determine the type of error based on the comparison, considering the categories: Named Entities, Numerical Values, Modifiers or Adjectives, Negation or Antonyms, Facts, and Dropped Content.

Step 4: Classify the error into one of the provided options (A, B, C, D, E, F) based on the identified change.

Step 5: Confirm the error type based on the comparison and classification.

</reasoning_framework>

Prompt Template

<output_format>

The output is a single letter corresponding to the correct option from the list (e.g., A, B, C, etc.).

</output_format>

Follow these steps carefully:

1. Based on the reasoning_framework, provide a step-by-step chain of thought that answers the question.

<chain_of_thought>

[Your step-by-step chain of thought here]

</chain_of_thought>

2. Based on your chain_of_thought, provide the final answer according to the rules specified in the output_format section.

If unsure and "Output Format" is option, guess the closest option. Present your final answer in the following format:

<final_answer>

[Your final answer here]

</final_answer>

Note Do not use programming or code to solve this question.

Question

<question>

The following translations from German to English contain a particular error. That error will be one of the following types: Named Entities: An entity (names, places, locations, etc.) is changed to a different entity. Numerical Values: Numerical values (ordinals or cardinals), dates, and/or units are changed. Modifiers or Adjectives: The modifiers and adjectives pertaining to a noun are changed. Negation or Antonyms: Introduce or remove a negation or change comparatives to their antonyms. Facts: Trivial factual errors not pertaining to the above classes are introduced in the translations. Dropped Content: A significant clause in the translation is removed. Please identify that error. Source: Die Mexikanische Königsnatter oder San-Luis-Potosi-Königsnatter ist eine Schlange aus der Familie der Nattern.

Translation: The Mexican King's Snake or San Luis Potosi King snake is a snake of the Lantern family.

The translation contains an error pertaining to

Options:

- (A) Modifiers or Adjectives
- (B) Numerical Values
- (C) Negation or Antonyms
- (D) Named Entities
- (E) Dropped Content
- (F) Facts

</question>

Model Response

<chain_of_thought>

To identify the error type in the translation, let's compare the source sentence and the translation step by step.

1. **Source Sentence:** Die Mexikanische Königsnatter oder San-Luis-Potosi-Königsnatter ist eine Schlange aus der Familie der Nattern.

2. **Translation:** The Mexican King's Snake or San Luis Potosi King snake is a snake of the Lantern family.

Step 1: Identify Key Elements

- Source sentence mentions "Mexikanische Königsnatter" and "San-Luis-Potosi-Königsnatter" as names.
- Translation mentions "Mexican King's Snake" and "San Luis Potosi King snake" as names.
- Source sentence states the family as "der Nattern."
- Translation states the family as "the Lantern family."

Step 2: Compare Key Elements

- The names "Mexikanische Königsnatter" and "San-Luis-Potosi-Königsnatter" are translated to "Mexican King's Snake" and "San Luis Potosi King snake," respectively. This seems to be a change in the naming but let's check if it's an entity change or just a translation.
- The family in the source sentence is "der Nattern" (which translates to "the colubrids" or a general term for snakes, specifically non-venomous snakes), and in the translation, it's "the Lantern family."

Step 3: Determine the Type of Error

- **Named Entities:** The names of the snakes are translated but not significantly altered in terms of entity type (still referring to the same snakes). However, "Lantern family" seems to be a different classification than "der Nattern."
- **Numerical Values:** There's no numerical value mentioned in either sentence.
- **Modifiers or Adjectives:** The modifiers are translated (Mexikanische to Mexican, Königsnatter to King's Snake), which seems accurate.
- **Negation or Antonyms:** There's no introduction or removal of negation.
- **Facts:** The mention of "Lantern family" instead of the actual family "Nattern" (or colubrids) seems to be a factual error regarding the classification of the snakes.
- **Dropped Content:** No significant clause seems to be removed.

Step 4: Classify the Error

Given that "der Nattern" refers to a specific family of snakes (colubrids) and "Lantern family" is not a recognized family of snakes in this context (it seems to be a mistranslation or misclassification), this appears to be a factual error regarding the classification.

Step 5: Confirm the Error Type

The error is best classified as a **factual error** because it misrepresents the family of the snakes. This aligns with option **(F) Facts**.

</chain_of_thought>

<final_answer>

F

</final_answer>

Figure 10: Example of the DIP-R prompt and reasoning for translation error detection in the BBH-Induct dataset, using a step-by-step framework. The gold (correct) answer for this example is (D) Named Entities.

DIP Prompt (N=1)

<task_instructions>

Task Content:

The task requires detecting the type of error in an English translation of a German source sentence. The error types to be considered are: Named Entities, Numerical Values, Modifiers or Adjectives, Negation or Antonyms, Facts, and Dropped Content.

Input Format:

- The source sentence in German.
- The translated sentence in English.
- A list of possible error types.

The task involves comparing the source sentence with the translation to identify any changes and determining the type of error based on the identified changes.

</task_instructions>

Reasoning Framework

<reasoning_framework>

Step 1: Compare the source sentence with the translation to identify any changes in terms of named entities, numerical values, modifiers or adjectives, negation or antonyms, facts, and dropped content.

Step 2: Analyze the translation for changes in named entities, such as the names of the snake and its family, and check if they match the source sentence.

Step 3: Check for changes in numerical values, modifiers, negation, facts, or dropped content by carefully examining the source and translated sentences.

Step 4: Determine the type of error based on the identified changes by matching them with the given error types: Named Entities, Numerical Values, Modifiers or Adjectives, Negation or Antonyms, Facts, and Dropped Content.

Step 5: Select the correct option from the provided list that corresponds to the identified error type.

</reasoning_framework>

Prompt Template

<output_format>

The output is a single letter corresponding to the correct option from the list (e.g., A, B, C, etc.).

</output_format>

Follow these steps carefully:

1. Based on the reasoning_framework, provide a step-by-step chain of thought that answers the question. Present your chain of thought in the following format:

<chain_of_thought>

[Your step-by-step chain of thought here]

</chain_of_thought>

2. Based on your chain_of_thought, provide the final answer according to the rules specified in the output_format section.

If unsure and "Output Format" is option, guess the closest option. Present your final answer in the following format:

<final_answer>

[Your final answer here]

</final_answer>

Note

- Do not use programming or code to solve this question.

- It's crucial to follow this structure and include both the chain of thought and the final answer in your response.

Question

<question>

The following translations from German to English contain a particular error. That error will be one of the following types: Named Entities: An entity (names, places, locations, etc.) is changed to a different entity. Numerical Values: Numerical values (ordinals or cardinals), dates, and/or units are changed. Modifiers or Adjectives: The modifiers and adjectives pertaining to a noun are changed. Negation or Antonyms: Introduce or remove a negation or change comparatives to their antonyms. Facts: Trivial factual errors not pertaining to the above classes are introduced in the translations. Dropped Content: A significant clause in the translation is removed. Please identify that error. Source: Die Mexikanische Königsnatter oder San-Luis-Potosi-Königsnatter ist eine Schlange aus der Familie der Nattern.

Translation: The Mexican King's Snake or San Luis Potosi King snake is a snake of the Lantern family.

The translation contains an error pertaining to

Options:

- (A) Modifiers or Adjectives
- (B) Numerical Values
- (C) Negation or Antonyms
- (D) Named Entities
- (E) Dropped Content
- (F) Facts

</question>

Model Response

<chain_of_thought>

To identify the error in the translation, let's compare the source sentence with the translated sentence step by step.

1. **Named Entities:** The source sentence mentions "Die Mexikanische Königsnatter oder San-Luis-Potosi-Königsnatter" which translates to "The Mexican King's Snake or San Luis Potosi King snake." The names of the snake species seem to be correctly translated, although the formatting and capitalization might slightly differ, they refer to the same entities.

2. **Numerical Values:** There are no numerical values mentioned in either the source or the translated sentence, so there's no error related to numerical values.

3. **Modifiers or Adjectives:** The source sentence describes the snake as "eine Schlange aus der Familie der Nattern," which means "a snake from the family of colubrids" or "a snake of the family of snakes." The translation mentions "a snake of the Lantern family." Here, "Nattern" (which refers to a family of snakes, specifically colubrids) is changed to "Lantern family," which seems to be incorrect as "Lantern" does not correspond to any known family of snakes related to the context.

4. **Negation or Antonyms:** There is no introduction or removal of negation, nor a change to comparatives or their antonyms in the translation.

5. **Facts:** The change from "family of snakes (Nattern)" to "Lantern family" introduces a factual error since "Lantern family" does not correctly represent the family of snakes.

6. **Dropped Content:** There doesn't seem to be a significant clause removed from the translation.

Given these observations, the error seems to pertain to **Named Entities** or **Facts**. However, more precisely, it involves changing the family of the snake from "Nattern" (a correct term for a family of snakes) to "Lantern family," which is not a recognized family of snakes. This change affects the factual accuracy regarding the classification of the snake.

Since "Lantern family" is not a correct family name and represents a change in factual information about the snake's family, and considering that "Named Entities" typically refers to names of specific entities like people, places, or organizations rather than general terms like families of animals, the error is best categorized under **Facts**.

</chain_of_thought>

<final_answer>

F

</final_answer>

Figure 11: Example of DIP Prompt (N=1) and model response for translation error detection in the BBH-Induct dataset. The gold (correct) answer for this example is (D) Named Entities.

Zero-shot CoT Prompt

<task_instruction>
Detect the type of error in an English translation of a German source sentence.
</task_instruction>

Note Do not use programming or code to solve this question.

Prompt Template

<output_format>
The output is a single letter corresponding to the correct option from the list (e.g., A, B, C, etc.).
</output_format>

Follow these steps carefully:

1. Provide a step-by-step chain of thought that answers the question.

<chain_of_thought>
[Your step-by-step chain of thought here]
</chain_of_thought>

2. Based on your chain_of_thought, provide the final answer according to the rules specified in the output_format section.

If unsure and "Output Format" is option, guess the closest option. Present your final answer in the following format:

<final_answer>
[Your final answer here]
</final_answer>

Question

<question>

The following translations from German to English contain a particular error. That error will be one of the following types: Named Entities: An entity (names, places, locations, etc.) is changed to a different entity. Numerical Values: Numerical values (ordinals or cardinals), dates, and/or units are changed. Modifiers or Adjectives: The modifiers and adjectives pertaining to a noun are changed. Negation or Antonyms: Introduce or remove a negation or change comparatives to their antonyms. Facts: Trivial factual errors not pertaining to the above classes are introduced in the translations. Dropped Content: A significant clause in the translation is removed. Please identify that error. Source: Die Mexikanische Königsnatter oder San-Luis-Potosi-Königsnatter ist eine Schlange aus der Familie der Nattern.

Translation: The Mexican King's Snake or San Luis Potosi King snake is a snake of the Lantern family.

The translation contains an error pertaining to

Options:

- (A) Modifiers or Adjectives
- (B) Numerical Values
- (C) Negation or Antonyms
- (D) Named Entities
- (E) Dropped Content
- (F) Facts

</question>

Model Response

<chain_of_thought>

Step 1: Identify the source sentence and its translation.

The source sentence in German is: "Die Mexikanische Königsnatter oder San-Luis-Potosi-Königsnatter ist eine Schlange aus der Familie der Nattern."

The English translation provided is: "The Mexican King's Snake or San Luis Potosi King snake is a snake of the Lantern family."

2: Compare the source sentence and its translation to identify potential errors.

- The source sentence mentions "Mexikanische Königsnatter" and "San-Luis-Potosi-Königsnatter," which are translated to "Mexican King's Snake" and "San Luis Potosi King snake," respectively. This seems accurate in terms of named entities.

- The source sentence states the snake is "eine Schlange aus der Familie der Nattern," which translates to "a snake from the family of Nattern" or more accurately in English, "a snake from the family of colubrids" (since Nattern refers to a family of non-venomous snakes, commonly known as colubrids).

- The translation mentions "the Lantern family," which does not match "der Familie der Nattern" (the family of colubrids or non-venomous snakes).

3: Determine the type of error.

- **Named Entities (D):** The names of the snakes ("Mexican King's Snake" and "San Luis Potosi King snake") seem to be correct translations of "Mexikanische Königsnatter" and "San-Luis-Potosi-Königsnatter."

- **Numerical Values (B):** There are no numerical values mentioned in the source sentence or the translation.

- **Modifiers or Adjectives (A):** The modifiers and adjectives seem to be generally preserved, though the critical error lies elsewhere.

- **Negation or Antonyms (C):** There is no introduction or removal of negation or change of comparatives to their antonyms.

- **Facts (F):** The family mentioned in the translation ("Lantern family") is incorrect. The correct family should be related to "Nattern," which refers to colubrids, not "Lantern," which could imply a different or fictional classification.

- **Dropped Content (E):** The translation does not seem to have dropped a significant clause but has introduced an incorrect fact.

4: Conclusion on the type of error.

The error is in the family of the snake, which is translated as "the Lantern family" instead of something akin to "the family of colubrids" or "the family of non-venomous snakes." This change introduces a **factual error** regarding the classification of the snake.

5: Select the correct option based on the analysis.

Given that the error involves changing a factual piece of information about the snake's family from "Nattern" (colubrids) to "Lantern family," this falls under **Facts (F)**: Trivial factual errors not pertaining to the above classes are introduced in the translations.

</chain_of_thought>

The final answer is:

<final_answer>

F

</final_answer>

Figure 12: Example of Zero-shot CoT Prompt and model response for translation error detection in the BBH-Induct dataset. The gold (correct) answer for this example is (D) Named Entities.

S-CoT Prompt
<task_instruction>
Detect the type of error in an English translation of a German source sentence.
</task_instruction>

Prompt Template
<output_format>
The output is a single letter corresponding to the correct option from the list (e.g., A, B, C, etc.).
</output_format>
Follow these steps carefully:
1. Carefully consider the problem and generate the strategic knowledge that would best guide the problem-solving process.
<strategy>
[Your strategy here]
</strategy>
2. Provide a step-by-step chain of thought that answers the question.
<chain_of_thought>
[Your step-by-step chain of thought here]
</chain_of_thought>
3. Based on your chain_of_thought, provide the final answer according to the rules specified in the output_format section.
If unsure and "Output Format" is option, guess the closest option. Present your final answer in the following format:
<final_answer>
[Your final answer here]
</final_answer>
Note Do not use programming or code to solve this question.

Question
<question>
The following translations from German to English contain a particular error. That error will be one of the following types: Named Entities: An entity (names, places, locations, etc.) is changed to a different entity. Numerical Values: Numerical values (ordinals or cardinals), dates, and/or units are changed. Modifiers or Adjectives: The modifiers and adjectives pertaining to a noun are changed. Negation or Antonyms: Introduce or remove a negation or change comparatives to their antonyms. Facts: Trivial factual errors not pertaining to the above classes are introduced in the translations. Dropped Content: A significant clause in the translation is removed. Please identify that error. Source: Die Mexikanische Königsnatter oder San-Luis-Potosi-Königsnatter ist eine Schlange aus der Familie der Nattern.
Translation: The Mexican King's Snake or San Luis Potosi King snake is a snake of the Lantern family.
The translation contains an error pertaining to
Options:
(A) Modifiers or Adjectives
(B) Numerical Values
(C) Negation or Antonyms
(D) Named Entities
(E) Dropped Content
(F) Facts
</question>

Model Response
<strategy>
To identify the type of error in the translation, let's first analyze the source sentence and the translation given, and then compare them to determine the nature of the discrepancy.
1. **Source Sentence:** Die Mexikanische Königsnatter oder San-Luis-Potosi-Königsnatter ist eine Schlange aus der Familie der Nattern.
2. **Translation:** The Mexican King's Snake or San Luis Potosi King snake is a snake of the Lantern family.
The key components to compare are:
- Names of the species
- Family of the species
</strategy>
<chain_of_thought>
1. **Named Entities and Modifiers:**
- The source sentence mentions "Mexikanische Königsnatter" and "San-Luis-Potosi-Königsnatter" as the names of the snake, which are translated to "The Mexican King's Snake" and "San Luis Potosi King snake". The names are similar but not exactly the same due to the capitalization and slight variation in wording ("King's" vs. "King"), yet they refer to the same entities.
2. **Family of the Species:**
- The source sentence states that the snake is "aus der Familie der Nattern", which translates to "from the family of colubrids" or more generally "from the family of snakes" (Nattern is a German term for a type of snake, specifically the Colubridae family).
- The translation states the snake is "of the Lantern family".
3. **Error Identification:**
- The term "Nattern" refers to a specific family of snakes (Colubridae), which is not the "Lantern family". The "Lantern family" does not correspond to any known classification in biology related to snakes. This indicates a factual error regarding the family classification of the snake.
4. **Type of Error:**
- **Named Entities (D):** The names are slightly adjusted but still refer to the same entities.
- **Numerical Values (B):** There are no numerical values mentioned.
- **Modifiers or Adjectives (A):** The modifiers are adjusted but not in a way that changes the meaning significantly regarding the error type.
- **Negation or Antonyms (C):** There is no introduction or removal of negation.
- **Facts (F):** The family of the snake is incorrectly stated as "Lantern family" instead of a correct biological classification like "Colubridae" or a similar term.
- **Dropped Content (E):** No significant clause is removed.
Given the analysis, the error pertains to **Facts (F)** because a trivial factual error is introduced by stating the snake belongs to the "Lantern family", which is not a correct biological classification.
</chain_of_thought>
<final_answer>
F
</final_answer>

Figure 13: Example of S-CoT Prompt and model response for translation error detection in the BBH-Induct dataset. The gold (correct) answer for this example is (D) Named Entities.

R-CoT Prompt
<task_instruction>
Detect the type of error in an English translation of a German source sentence.
</task_instruction>

Prompt Template
<output_format>
The output is a single letter corresponding to the correct option from the list (e.g., A, B, C, etc.).
</output_format>
Follow these steps carefully:
1. First, carefully consider the question and propose a clear reasoning framework to guide your thinking.
<reasoning_framework>
Step 1: [First high-level reasoning step]
Step 2: [Second high-level reasoning step]
Step 3: [Third high-level reasoning step]
[Add additional steps if necessary]
</reasoning_framework>
2. Next, use the framework you've outlined to explain your reasoning step-by-step.
<chain_of_thought>
1. [Explanation aligned with Step 1]
2. [Explanation aligned with Step 2]
3. [Explanation aligned with Step 3]
[Continue as necessary]
</chain_of_thought>
3. Based on your chain_of_thought, provide the final answer according to the rules specified in the output_format section.
If unsure and "Output Format" is option, guess the closest option. Present your final answer in the following format:
<final_answer>
[Your final answer here]
</final_answer>
Note: Do not use programming or code to solve this question.

Question
<question>
The following translations from German to English contain a particular error. That error will be one of the following types: Named Entities: An entity (names, places, locations, etc.) is changed to a different entity. Numerical Values: Numerical values (ordinals or cardinals), dates, and/or units are changed. Modifiers or Adjectives: The modifiers and adjectives pertaining to a noun are changed. Negation or Antonyms: Introduce or remove a negation or change comparatives to their antonyms. Facts: Trivial factual errors not pertaining to the above classes are introduced in the translations. Dropped Content: A significant clause in the translation is removed. Please identify that error. Source: Die Mexikanische Königsnatter oder San-Luis-Potosi-Königsnatter ist eine Schlange aus der Familie der Nattern.
Translation: The Mexican King's Snake or San Luis Potosi King snake is a snake of the Lantern family.
The translation contains an error pertaining to
Options:
(A) Modifiers or Adjectives
(B) Numerical Values
(C) Negation or Antonyms
(D) Named Entities
(E) Dropped Content
(F) Facts
</question>

Model Response
<reasoning_framework>
Step 1: Identify the key elements in the source sentence and the translation.
Step 2: Compare the key elements to determine the type of error.
Step 3: Analyze the differences and match them to the given error types.
</reasoning_framework>
<chain_of_thought>
1. **Identify key elements in the source sentence and the translation:**
- Source sentence: "Die Mexikanische Königsnatter oder San-Luis-Potosi-Königsnatter ist eine Schlange aus der Familie der Nattern."
- Translation: "The Mexican King's Snake or San Luis Potosi King snake is a snake of the Lantern family."
- Key elements include names (Mexikanische Königsnatter, San-Luis-Potosi-Königsnatter), classification (Familie der Nattern), and the English translation provided.
2. **Compare the key elements:**
- The names "Mexikanische Königsnatter" and "San-Luis-Potosi-Königsnatter" are translated to "Mexican King's Snake" and "San Luis Potosi King snake," respectively. This seems accurate in terms of named entities.
- The family "der Nattern" is translated to "the Lantern family."
3. **Analyze the differences and match them to the given error types:**
- The source sentence mentions "Familie der Nattern," which translates to "family of snakes" or more specifically "Colubridae" in scientific terms.
- The translation mentions "the Lantern family," which is incorrect because "Nattern" refers to a family of non-venomous snakes (Colubridae), not "Lantern."
4. **Determine the error type:**
- The error involves changing "Nattern" (a family of snakes) to "Lantern family," which is a factual error regarding the classification of the snake.
- This does not fit into named entities (D) as the issue is not with the names but with the classification.
- It does not fit into numerical values (B), negation or antonyms (C), or modifiers/adjectives (A) directly.
- It does not fit into dropped content (E) as there is no significant clause removed.
- The closest match is **Facts (F)**, as it involves introducing a trivial factual error regarding the family classification.
</chain_of_thought>
<final_answer>
F
</final_answer>

Figure 14: Example of R-CoT Prompt and model response for translation error detection in the BBH-Induct dataset. The gold (correct) answer for this example is (D) Named Entities.