

# In-Context Learning May Not Elicit Trustworthy Reasoning: A-Not-B Errors in Pretrained Language Models

Pengrui Han<sup>\*†1,3</sup>, Peiyang Song<sup>\*2</sup>, Haofei Yu<sup>3</sup>, Jiaxuan You<sup>3</sup>

<sup>1</sup>Carleton College, <sup>2</sup>California Institute of Technology

<sup>3</sup>University of Illinois Urbana-Champaign

barryhan@carleton.edu, psong@caltech.edu

## Abstract

Recent advancements in artificial intelligence have led to the creation of highly capable large language models (LLMs) that can perform tasks in a human-like manner. However, LLMs exhibit only infant-level cognitive abilities in certain areas. One such area is the A-Not-B error, a phenomenon seen in infants where they repeat a previously rewarded behavior despite well-observed changed conditions. This highlights their lack of inhibitory control – the ability to stop a habitual or impulsive response. In our work, we design a text-based multi-choice QA scenario similar to the A-Not-B experimental settings to systematically test the inhibitory control abilities of LLMs. We found that state-of-the-art LLMs (like Llama3-8b) perform consistently well with in-context learning (ICL) but make errors and show a significant drop of as many as 83.3% in reasoning tasks when the context changes trivially. This suggests that LLMs only have inhibitory control abilities on par with human infants in this regard, often failing to suppress the previously established response pattern during ICL.

## 1 Introduction

In the field of cognitive science, there is a classic phenomenon called the A-Not-B error (Popick et al., 2011; Smith and Gasser, 2005; Vorms, 2012). In a typical A-Not-B task, an experimenter repeatedly hides an attractive toy at one location (e.g., under box A) within the baby’s reach while the baby watches, and has the baby retrieve it several times. Then, during the critical trial, the experimenter moves the toy to a new location (e.g., under box B) that is also within the baby’s reach *while the baby watches*, yet the baby fails to understand the straightforward change of situation, and continues previous actions to search under the original location (box A). This is a key indication of infants’

\*Equal contribution.

†Work done as intern.

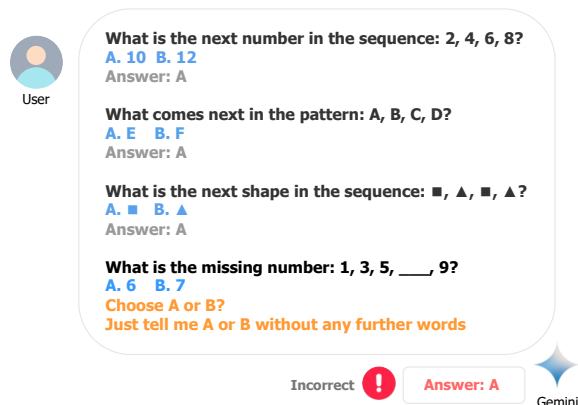


Figure 1: **A-Not-B Style Few-Shot Prompts Mislead Gemini on Elementary School-Level Questions.** This figure presents a simple few-shot prompt that tricks an advanced model Gemini on elementary school-level questions by consistently providing examples with Answer A. This example was collected on Sep 21, 2024, with the possibility that future updates may lead to different results. We provide a screenshot in Appendix §D.

limited *inhibitory control*: the ability to stop or control impulsive responses and instead use deliberate attention and reasoning when the context changes (Geier, 2013; Fiske and Holmboe, 2019). On the other hand, adults have developed strong inhibitory control abilities over the years as they have matured (Geier, 2013). For instance, an adult may initially reach for a frequently used coffee mug on its usual shelf but quickly adapt and reach for it on a new shelf after remembering it was moved the day before. Similarly, evidence has shown that when infants grow to about 12 months (MacNeill et al., 2018), they typically begin to have such inhibitory control abilities, and no longer exhibit A-not-B errors.

Recent advancements in Large Language Models (LLMs) (Saravanan et al., 2023; Zhang et al., 2024a; Chang et al., 2024; Kasneci et al., 2023; Zhao et al., 2023b) have significantly impacted cognitive science-related research (Feng et al., 2024).

Researchers find that LLMs perform well in some cognitive tasks, such as Theory-of-mind (Kosinski, 2023), analogical reasoning (Webb et al., 2023), and moral reasoning (Almeida et al., 2024), which proves that LLMs have some basic cognitive abilities similar with humans. However, they still struggle with other, sometimes embarrassingly simple cognitive settings, including tasks like spatial reasoning (Sharma, 2023; Yang et al., 2019), solving cunningly designed problems (Nezhurina et al., 2024), and reversal learning (Berglund et al., 2024), indicating a lack of robust cognitive abilities. Such failure cases always suggest meaningful directions to improve LLM abilities and alignment with human intelligence. These successes and failures of LLMs in cognitive abilities are better understood through insights from cognition-related theories, including attention (Vaswani et al., 2017), working memory (Li et al., 2022), long-term memory (Zhong et al., 2024), chain-of-thought prompting (Wei et al., 2023) and more. In this work, we draw inspiration from human A-Not-B error experiments to test the inhibitory control ability of LLMs. Our objective is to assess how well LLMs can inhibit established patterns in contexts and to explore strategies for enhancing this ability in practical applications.

To test the inhibitory control ability of LLMs, we adapt the A-Not-B experimental setting to be purely language-based. Like infants developing a habitual response by seeing a toy repeatedly placed in box A, we show LLMs several multiple-choice questions with the same correct option to establish a trivial pattern. Then, during the critical trial, we present a new question where the correct answer is not A (e.g. B). We refer to this prompting construction strategy as *A-Not-B prompting*. Our goal is to see if LLMs can break the established trivial pattern and use deliberate attention and reasoning, similar to humans' inhibitory control.

Surprisingly, as illustrated in Figure 1, even advanced LLMs like Gemini (Team, 2024) fail to answer extremely simple questions correctly under this scenario. In our further experiments, by only using a 3-shot A-Not-B prompt, popular LLMs (eg. Llama3-8b and Qwen1.5-72b) show a significant decrease in accuracy by as much as 83.3% on some reasoning tasks. This indicates that LLMs have weaker cognitive abilities than elementary school children and can cause serious errors when being used in applications.

Furthermore, we perform an in-depth analysis

of why significant errors occur in LLMs with simple A-Not-B prompting. Our findings suggest that during the training stage, model size and training data quality are critical factors—larger models and higher-quality training data can help reduce these errors. In the post-training phase, strategies like self-explanation (Zhao et al., 2023a) have also been effective in mitigating these errors to a certain degree, though not enough to get rid of the errors. Additionally, the type of task and the number of prompt examples significantly impact performance; LLMs are more likely to rely on superficial A-Not-B patterns, particularly in reasoning tasks that involve less straightforward logical patterns and memorization, and when the A-Not-B pattern is more strongly reinforced. Moreover, LLMs are much less resilient to these errors than human adults, indicating that there may be fundamental differences between LLMs' and humans' reasoning processes.

To our knowledge, we are the first to systematically evaluate the inhibitory control abilities of LLMs. We open-source all code and results under a permissive MIT license, to encourage reproduction and further research exploration <sup>1</sup>.

## 2 Related Work

**A-Not-B Error and Inhibitory Control.** The A-Not-B error is a classic cognitive phenomenon observed in infants typically between 8 to 12 months, and resolves as they grow older (Popick et al., 2011; Smith and Gasser, 2005; Vorms, 2012). Researchers note this as a significant milestone in cognitive development in humans that signifies the emergence of inhibitory control (Diamond, 2013; Casey et al., 2000). Inhibitory control is the ability to inhibit a response – such as refraining from seeking an object at a previously habitual location A after explicitly seeing it moved to location B. This ability is among the three core components of executive functions (Diamond, 2013), along with working memory (Baddeley, 1992) and cognitive flexibility (Miyake et al., 2000). These cognitive abilities are crucial for human development and play a vital role in various aspects of cognition (Anderson et al., 2001; Best and Miller, 2010; Blair and Razza, 2007).

Inhibitory control, in particular, emerges early in development and is essential for fundamental human cognitive abilities such as decision-

<sup>1</sup><https://github.com/Peiyang-Song/LLM-A-Not-B-Errors>

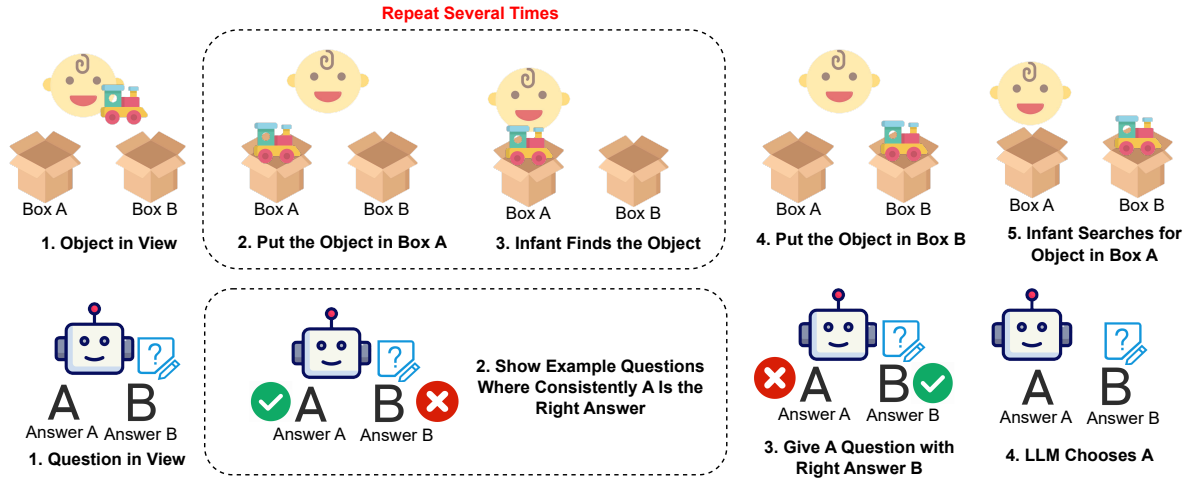


Figure 2: **Illustration of A-not-B Task Performance in Infants and LLMs.** This figure demonstrates the typical A-not-B error using a binary (Boxes A and B) setup. The first sequence showcases an infant’s repeated actions: seeing an object placed in Box A, observing it moved to Box B, yet continuing to search in Box A. This depicts the cognitive phenomenon where prior experience overrides current explicit visual cues. On the bottom sequence, the figure analogously presents a scenario in which LLMs are misled by a consistent answer pattern, illustrating the A-not-B prompting scenarios for LLMs, where they fail to adapt to minimally changed circumstances.

making (Madden and Johnson, 2010; Shoda et al., 1990) and reasoning (Blair, 2002; Hughes, 1998). It is also indispensable for more advanced and complicated human abilities such as social skills (Carlson and Moses, 2001), delay of gratification (Mischel et al., 1989; Duckworth and Seligman, 2005), mental flexibility (Diamond et al., 2002; Cragg and Chevalier, 2012), and more (Anderson, 2002; Best and Miller, 2010). Without sufficient inhibitory control, crucial errors such as social faux pas (Carlson and Moses, 2001; Garon et al., 2008; Kochanska et al., 1997) would occur.

**LLMs and Inhibitory Control.** Humans learn from previous experience (Yip and Wilson, 2010), and inhibitory control allows them to suppress automatic periodic responses and avoid being distracted by past irrelevant patterns, thereby exhibiting more adaptive behaviors (Kolb and Whishaw, 2009; Diamond, 2013). LLMs, like humans, also learn from a vast array of previous experience (Hinton, 2002) including demonstrations (Wang et al., 2024; Shao et al., 2023), examples (Brown et al., 2020; Yıldız et al., 2024), and interactions (Collins et al., 2024; Yang and Deng, 2019). The learning takes place through diverse mechanisms such as ICL (Xie et al., 2021; Coda-Forno et al., 2023), fine-tuning (Han et al., 2024b,a), and pretraining (Yıldız et al., 2024; Raffel et al., 2020).

Previous studies show that LLMs are capable of learning from these experiences (data and con-

texts) (Zhao et al., 2023c; Dong et al., 2022; Tang et al., 2023), extracting meaningful patterns and concepts (Wang et al., 2024; Min et al., 2022), and generalizing effectively (Li et al., 2024b; Song et al., 2024; White et al., 2024). Models trained with these strategies have demonstrated great performance across various benchmarks in language understanding (Wang et al., 2019; Li et al., 2023), question answering (Kwiatkowski et al., 2019; Joshi et al., 2017), reasoning (Clark et al., 2018; Yang et al., 2023), and software engineering (Zhang et al., 2023b; Yan et al., 2024; Poesia et al., 2022).

Given such similarities between LLMs and humans in learning from past experiences, it is important to study whether LLMs suffer from inhibitory control like early-stage human infants. A related field of study is machine unlearning (Zhang et al., 2023a; Tarun et al., 2023; Liu, 2024), which involves forgetting previously learned information under certain scenarios (usually due to privacy or security concerns). Yet as crucial is the ability to suppress previous knowledge without specifically forgetting it, showing adaptability to new situations. To investigate this key cognitive ability in LLMs, we focus on ICL, which does not involve model parameter updates (no modification of what has been learned during past experiences), thus precisely determining if LLMs have sufficient inhibitory control to suppress irrelevant or meaningless information on the fly.

**LLMs and Multiple-Choice QA.** In our experimental setting in Section §3, we test LLMs on multiple-choice questions (MCQs). MCQs are widely used evaluators for LLMs (Zhang et al., 2024b; Li et al., 2024a). Prior research has shown that current LLMs do not perform consistently well on MCQ tasks. For instance, (Zheng et al., 2023) demonstrates that LLMs are sensitive to changes in the positioning of options and tend to prefer certain option IDs potentially due to token bias (Jiang et al., 2024). While such studies explore interesting aspects of LLM performance on MCQs, they primarily focus on the inherent capability of LLMs to perform on MCQs. In contrast, our work investigates whether, and how, external factors, such as minimal changes in context, impact the consistency of LLMs’ performances on MCQs. Such different settings require LLMs to perform inhibitory control and exhibit trustworthy reasoning. We will elaborate more on the specific designs in Section §3.

### 3 Experimental Setup

#### 3.1 Experiment Motivation

Motivated by the original A-not-B experiment, we adapt the scenario to be purely based on natural language, thus creating a parallel experiment for LLMs as shown in Figure 2. In the original experiment, the infant constantly observes the placement of *a certain object* to be in the same *Location A*. In ours, this is implemented by LLMs observing the ground truth answer to an MCQ question to be the same *Option A*. Similar to how the infant then observes the placement of the ball in a different "Location B", LLMs are tasked with a similar-style MCQ from the same domain whose ground truth answer is *Option B*. An A-not-B error is observed if the infant still looks for the object in *Location A*, or if the LLMs choose the wrong *answer A*. We believe that by using such an experiment setup, we are testing similar abilities between the original A-not-B cognitive experiment and the LLM version, because the main difficulty of both lies in inhibiting a previously established trivial pattern of one option being repetitively selected, when moving to a new scenario, and instead eliciting internal knowledge to solve new tasks.

This demonstrates limited inhibitory control because an infant would know well that the ball is in "Location B", if not paying excessive attention to the established pattern from previous demonstra-

tions that "the ball is discovered in Location A". This corresponds to how LLMs *are* able to choose the correct answer (Option B) if it had not seen the previous MCQs with Options A as the correct answer. That is, LLMs *do* have the capability to answer some MCQs correctly yet fail to do so in this mimicked A-not-B scenario, pronounced in an unexpected accuracy drop.

#### 3.2 Models

We select four leading open-source models with different levels of abilities and knowledge. Specifically, we choose two families of LLMs: Llama3 (AI@Meta, 2024) and Qwen-1.5 (Bai et al., 2023). Within each model family, we experiment with both large and small models to investigate the relationship between the impact of A-not-B errors and model sizes. Specifically, we experimented with Llama3-70b, Llama3-8b, Qwen1.5-72b, and Qwen1.5-7b.

#### 3.3 Data

**MCQA Datasets** We choose four representative multiple-choice question-answering (MCQA) datasets, each tasking a particular fundamental reasoning task: the MathQA dataset (Amini et al., 2019) for arithmetic reasoning, the CommonsenseQA dataset (Talmor et al., 2019) for commonsense reasoning, the Winogrande dataset (Sakaguchi et al., 2019) for causal reasoning, and the SciQ dataset (Welbl et al., 2017) for scientific reasoning.

**Preprocessing** We preprocess the datasets to split an MCQA sample into three parts: question statement, choices, and ground truth answer. We modify the MCQA samples so that each has only two choices left, with one of them being the ground truth and the other incorrect, thus loyally resembling the original A-not-B scenario.

#### 3.4 Prompting Settings

We test the models’ performances on the modified datasets in two different prompting settings: (1) the *original* and (2) the *A-not-B* settings.

**Original prompting** We construct a prompt in the standard few-shot paradigm – we first provide  $n$  (typically from 5 to 50) MCQA examples, and then ask one question to the model. We then check if the model’s answer agrees with the ground truth.

**A-Not-B prompting** We reorder the options so that answers for all  $n$  examples we provide are the first one (Option A). Then for the final question we ask, the ground truth is set as the second (Option B). For each model, we test both the *original* and the *A-Not-B* prompting settings with 100 data samples per task and record the success rates. The differences in the success rates then tell the LLMs’ susceptibility to the A-not-B scenario.

## 4 Results

Our results in Table 1 show three key factors that are related to the inhibitory control ability of LLMs – (1) *model size*, (2) *number of few-shot examples*, and (3) *type of reasoning tasks*. We discuss each of them in detail in the rest of this section.

**Impact of Model Size.** In Figure 3, we compare the average rate of change across different numbers of few-shot examples for all four reasoning tasks between large models (Llama3-70b and Qwen1.5-72b) and small models (Llama3-8b and Qwen1.5-7b). The results show that model size significantly impacts performance, with smaller models consistently showing lower resilience (less accuracy) compared to larger models across all shots. Specifically, the larger models exhibit an average drop of 8.7% at 25 shots, while the smaller models show an average drop of 20.8% at 25 shots. This indicates that smaller models with fewer internal parameters and knowledge are more susceptible to the A-Not-B scenario, indicating less inhibitory control abilities.

**Impact of Number of Few-Shot Examples.** In Figure 3, we can also see the average performance of large (Llama3-70b and Qwen1.5-72b) and small (Llama3-8b and Qwen1.5-7b) models with different numbers of few-shot examples. Naturally, as the number of few-shot examples increases, both large and small models are significantly more likely to suffer from an accuracy drop. This indicates that when trivial patterns are reinforced more, it becomes harder for LLMs to inhibit them and instead turn to deliberate attention for reasoning. Such errors can be easily avoided by humans, which we further show in Section 6.

**Impact of Reasoning Tasks.** Experiments demonstrate that models’ vulnerability to the A-Not-B scenario varies across different reasoning tasks, as indicated by the different degrees of performance drop. The most pronounced

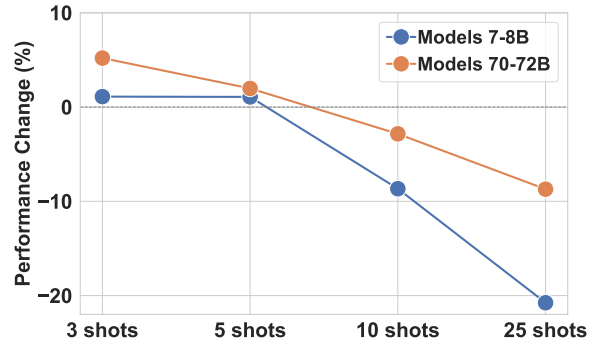


Figure 3: **Increasing Few-Shot Examples and Smaller Model Sizes Lead to Greater Accuracy Decline.** This figure shows the performance impact of increasing few-shot examples on large and small models. Negative values represent a decline in accuracy due to adversarial prompts. (1) More shots generally lead to a greater decline in performance, especially beyond 10 shots. (2) Smaller models (Llama3-8b and Qwen1.5-7b) are consistently more affected than larger models (Llama3-70b and Qwen1.5-72b), indicating higher susceptibility to the A-Not-B scenario.

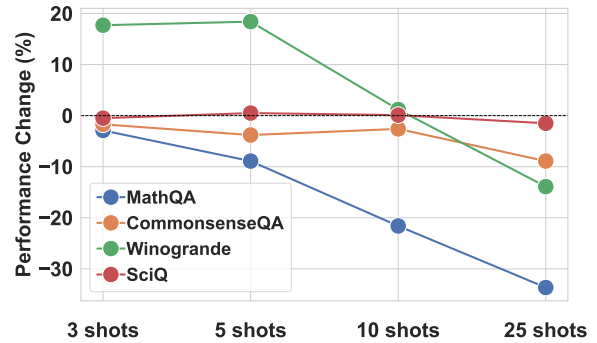


Figure 4: **Different Reasoning Tasks Show Susceptibility to A-Not-B Errors to Varying Degrees.** This figure illustrates how different reasoning tasks experience performance drops under the A-Not-B scenario as the number of shots increases. (1) *Arithmetic reasoning (MathQA)* shows a significant decline in performance across all shots, indicating a high susceptibility. (2) *Commonsense reasoning (CommonsenseQA)* also exhibits A-Not-B errors consistently across all shots but is less impacted than arithmetic reasoning. (3) *Causal reasoning (Winogrande)* begins to show significant performance drops at 25 shots. (4) *Scientific reasoning (SciQ)* displays minimal fluctuations and only shows a slightly more noticeable impact at 25 shots.

drop occurs in the arithmetic reasoning dataset (MathQA (Amini et al., 2019)), which is consistently the most challenging for models. This is likely due to the complexity of arithmetic reasoning, usually demanding multiple steps and requiring advanced skills that lead models to

# of shots	MathQA			CommonsenseQA			Winogrande			SciQ			
	Original	A-not-B	Diff	Original	A-not-B	Diff	Original	A-not-B	Diff	Original	A-not-B	Diff	
Llama3-70b	3	32%	36%	+ 12.5%	84%	88%	+ 4.8%	32%	36%	+ 12.5%	96%	100%	+ 4.2%
	5	36%	42%	+ 16.7%	86%	86%	0.0%	76%	80%	+ 5.3%	98%	100%	+ 2.0%
	10	36%	32%	- 11.1%	86%	88%	+ 2.3%	78%	80%	+ 2.6%	100%	100%	0.0%
	25	28%	24%	- 14.3%	92%	90%	- 2.2%	86%	78%	- 9.3%	100%	100%	0.0%
Llama3-8b	3	62%	22%	- 64.5%	86%	74%	- 14.0%	46%	64%	+ 39.1%	96%	90%	- 6.2%
	5	42%	14%	- 66.7%	92%	82%	- 10.9%	54%	76%	+ 40.7%	94%	94%	0.0%
	10	32%	8%	- 75.0%	94%	82%	- 12.8%	50%	52%	+ 4.0%	98%	96%	- 2.0%
	25	36%	6%	- 83.3%	92%	62%	- 32.6%	50%	28%	- 44.0%	96%	90%	- 6.2%
Qwen1.5-72b	3	66%	68%	+ 3.0%	96%	96%	0.0%	80%	82%	+ 2.5%	96%	98%	+ 2.1%
	5	56%	50%	- 10.7%	94%	92%	- 2.1%	80%	82%	+ 2.5%	94%	96%	+ 2.1%
	10	56%	44%	- 21.4%	94%	92%	- 2.1%	74%	76%	+ 2.7%	92%	96%	+ 4.3%
	25	50%	28%	- 44.0%	94%	92%	- 2.1%	82%	82%	0.0%	92%	94%	+ 2.2%
Qwen1.5-7b	3	64%	88%	+ 37.5%	84%	86%	+ 2.4%	60%	70%	+ 16.7%	96%	94%	- 2.1%
	5	72%	90%	+ 25.0%	88%	90%	- 2.3%	64%	80%	+ 25.0%	94%	92%	- 2.1%
	10	76%	92%	+ 21.1%	92%	94%	+ 2.2%	86%	82%	- 4.7%	94%	92%	- 2.1%
	25	86%	92%	+ 7.0%	98%	96%	- 2.6%	98%	96%	- 2.3%	96%	94%	- 2.1%

Table 1: **LLMs with different sizes and different shots for prompting can be easily misled by A-Not-B style adversarial prompts for arithmetic, commonsense, causal, and scientific reasoning tasks.** Accuracy drops are denoted in green, with accuracy increases shown in red. It includes results with four models, four shot numbers, and four reasoning tasks. Despite the overall trend summarized in Section 4, certain mode-task combinations (e.g. Qwen1.5-7b on mathematical reasoning) show interesting deviations from the overall patterns, which we discuss further in the ablation studies in Section 5.

over-rely on trivial patterns when unsure of how to solve a problem.

In contrast, the commonsense reasoning dataset (CommonsenseQA (Talmor et al., 2019)) shows a more moderate performance drop across all shots, which is consistently less dramatic compared to arithmetic reasoning. This reduced susceptibility can be attributed to the involvement of a certain amount of memorization in commonsense reasoning, which lowers the overall complexity and makes models less susceptible to the A-Not-B scenario, compared to mathematical reasoning.

For causal reasoning (Winogrande (Sakaguchi et al., 2019)), the impact becomes significant at 25 shots. This could be due to causal reasoning tasks typically involving distinct cause-and-effect relationships. When the model encounters a small number of shots, the few-shot examples effectively demonstrate these causal patterns, which can help improve performance if the model learns to follow the demonstrated reasoning. However, as the number of shots increases, the model may start to prioritize the A-Not-B pattern over the causal relationships in the examples. If this happens, the A-Not-B pattern becomes dominant, leading to errors and a noticeable decline in performance at higher shot counts.

For scientific reasoning (SciQ (Welbl et al., 2017)), the accuracy fluctuation is not as evident,

with at most a 1.5% accuracy drop. Since this observation coincides with extremely high accuracy numbers, we suspect an attribution to possible data contamination. Another explanation can be that many specific terms and technologies in these questions were present as-is during the model training, helping LLMs make accurate predictions confidently, thus less likely to be impacted by the A-not-B scenario.

## 5 Ablation Study

In this section, we delve deeper into the implications of interesting findings from Section §4. Specifically, we investigate 3 key research questions that naturally arise from the main results.

**RQ1: Why Qwen1.5-7b has abnormal behavior compared with other models?** *This is because Qwen1.5-7b is biased between A and B tokens during pretraining.*

In our main results, we notice abnormal performance increases in Qwen1.5-7b on mathematical reasoning. Although the increases diverge from the A-not-B scenario we primarily study in this paper, they still suggest untrustworthy reasoning. This is because minor perturbations with no meaningful changes should not lead to any significant accuracy fluctuations if the model is reasoning consistently and reliably.

CHOICE\SHOTS	3	5	10	25
A-NOT-B	+ 37.5	+ 25.0	+ 21.1	+ 7.0
B-NOT-A	- 12.2	- 32.7	0	- 15.3
C-NOT-D	- 11.2	- 7.1	- 28.4	- 34.1
E-NOT-F	- 17.0	- 18.4	- 40.0	- 41.2

Table 2: **Accuracy Drop Across Various Combinations for Different Shot Numbers On Qwen1.5-7b.**

This table illustrates the accuracy changes in Qwen1.5-7b across different variations of the A-not-B scenario (e.g., B-Not-A, C-Not-D, E-Not-F) with a range of shot numbers (3, 5, 10, and 25). Positive values (highlighted in green) indicate an increase in accuracy, while negative values (highlighted in red) represent a decrease. The results show that the A-Not-B configuration leads to an unusual accuracy increase, whereas other combinations consistently result in accuracy drops, suggesting potential internal biases in the model.

To further investigate this anomaly, we conduct an ablation experiment where we minimally alter the A-Not-B scenario to other combinations, such as B-Not-A, C-Not-D, and E-Not-F. We only change the option IDs (e.g., changing A and B to C and D) without altering any other content, keeping the settings symmetric and essentially identical, thus still testing the A-Not-B effect. Given that Qwen1.5-7b displays the most significant accuracy increases in mathematical reasoning, we focus on this task for the ablation study.

As shown in Table 2, we observe A-Not-B errors consistently across all other combinations for Qwen1.5-7b. This consistency suggests that the unusual accuracy increase is indeed an outlier, likely resulting from certain internal biases of the model (Jiang et al., 2024). Observing A-Not-B errors consistently across multiple parallel settings (e.g., C-Not-D) further strengthens the robustness of our experiment design, solidifying our discoveries.

**RQ2: Do A-not-B errors persist in MCQAs with more than two choices?** *Many-shot prompting reveals generalized versions of A-not-B errors.*

We further investigate if LLMs can overcome A-not-B errors in the less challenging case of many-shot MCQA. That is, rather than providing LLMs with two options (A or B) and constantly showing one of them as the correct answer, we provide LLMs with four or five options. The general structure of the prompts is the same – a number of examples followed by a final question. In the original setting, all options (including the answer to the final question) can appear in the demonstrated ex-




Figure 5: **Persistent Vulnerability of LLMs to A-not-B Errors Even in Many-Shot Scenarios.** This figure shows the performance of large (70B) and small (8B) Llama3 models in the many-shot arithmetic reasoning task for A-not-B errors. Despite more options and examples to reduce cognitive biases, both models, especially the small one, still exhibit significant accuracy declines, highlighting LLMs’ internal vulnerability to A-not-B errors.

amples; whereas in the A-not-B setting, all but the correct answer to the final question can appear.

Rather than few-shot experiments, we conduct **many-shot** to offer LLMs sufficient demonstrations. This setting is generalized from the standard A-not-B scenario in Section §3, with more options provided and more examples demonstrated, thus intuitively less challenging. We thus raise the question: Can LLMs overcome A-not-B errors in this generalized (and easier) scenario?

We again choose arithmetic reasoning, the most challenging reasoning task as discussed in Section §4 to investigate this generalized setting. The exact number of many-shot examples is 80. In the original many-shot scenario, 16 examples are provided with each of the five possible options (A, B, C, D, E) as the correct answer. In the A-not-B many-shot scenario, 20 examples are provided for A, B, C, and D, and no examples come with E as its answer. In both settings, the correct answer to the final question is E.

The exact prompts can be found in Figure 9 and Figure 10 in Appendix §C. Experimental results are reported in Figure 5. Significant accuracy drops persist for both large (70B) and small (8B) Llama3 models, with the small one suffering more. The results agree with the main experiments we design in Section §3, indicating that LLMs fail even in this generalized and less challenging version of many-shot A-not-B scenarios. This in turn confirms that the reported failure of LLMs from our main experiments is indeed internal, and cannot be attributed to the limited number of options (A and B).

MODEL \ SHOTS	0	3	5	10	25
Llama3-8b	48%	36%	24%	16%	18%

Table 3: **Self-Explanation Fails to Get Rid of A-not-B Errors in Llama3-8b.** This table illustrates the persistent and significant performance drop of Llama3-8b in the A-not-B scenario, despite the aid of self-explanation. Notably, accuracy still drops as the number of few-shot examples increases, aligning with observations from the main experiment.

**RQ3: Can A-not-B errors be avoided by asking LLMs to self-explain?** *Self-explanation fails to overcome A-not-B errors in challenging arithmetic reasoning tasks.*

Having observed that state-of-the-art open and closed LLMs fail in both embarrassingly simple A-not-B type questions and standard datasets rearranged in the A-not-B fashion, we move a step forward to investigate if some existing methods to improve trustworthy reasoning can help mitigate A-not-B errors. Self-explanation (Zhao et al., 2023a) is a widely used method of this kind that has been reported effective across many reasoning tasks (Zelikman et al., 2022; Zhao et al., 2024a; Zelikman et al., 2024), requiring LLMs to explain themselves while making decisions, thus eliciting its internal reasoning abilities directly. With the aid of self-explanation, we investigate whether LLMs can exhibit better inhibitory control, thus getting rid of A-not-B errors in their reasoning.

Similarly, we choose the most challenging reasoning task as discussed in Section §3 – arithmetic reasoning. Different from previous experiments, we now require LLMs to provide complete reasoning explanations in addition to answers. Results are shown in Table 3. Despite some degree of improvement, significant drops in accuracy are still observed, with an increasing severity as the number of shots increases. Such observations align with our main results in Section §3, indicating that LLMs are not able to easily overcome A-not-B errors even with self-explanation. The results also support the fact that LLMs *consciously* and *persistently* make A-not-B errors in reasoning tasks, hinting at its internal lack of inhibitory control.

## 6 Discussions

### 6.1 Human Performance in A-not-B Scenarios

While humans, particularly adults, generally possess sufficient inhibitory control to overcome the

typical cognitive A-Not-B errors observed in infants, it is uncertain whether they can maintain this resistance in our specific experimental setting. We conduct a human study to verify this.

We select two groups of college math major undergraduate students, ensuring that both groups have similar academic backgrounds for a consistent comparison and to minimize the possibility of random guessing. Each group consists of two participants. The first group is assigned an arithmetic reasoning dataset with 10 MCQs in our original setting of the main experiment, while the second group is given the same arithmetic dataset in the A-not-B setting. To make the experiment more interactive and force humans to digest ("decode") the examples provided, instead of simply showing 10 questions with answers A, we ask the participants to solve all the problems but only count their responses for the final question.

Our findings confirm that humans are highly resilient to A-not-B errors, as the largest performance difference between the two groups is as minimal as 2.0%. This suggests that humans are capable of reasoning trustworthily with sufficient inhibitory control abilities, demonstrating significant resilience to the A-Not-B scenario in our experiments.

### 6.2 Better Pretraining Helps Mitigate A-not-B Errors

To better understand how model pretraining, particularly pretraining data, influences models' resistance to A-not-B errors, we further compare the rates of change in accuracy across different numbers of shots, in Llama3-8b, Llama3-70b, and Llama2-70b. Results are shown in Figure 6. These Llama-family models share the same fundamental architectures, but the Llama3 models have been trained on significantly larger, more novel, and higher-quality datasets compared to Llama2-70b. As illustrated in Figure 6, although Llama2-70b has the same model size as Llama3-70b, it is much more significantly impacted. In fact, Llama2-70b is less resilient to A-not-B errors even compared to the much smaller Llama3-8b, especially as the number of shots increases.

This observation suggests that the quality and quantity of pretraining data can significantly impact LLMs' inhibitory control. When encountering input prompts or contexts, an LLM must apply its internal knowledge, skills, and abilities learned from the pretraining distribution to predict responses. Therefore, better pretraining (both in quality and






Figure 6: **Pretraining Affects Model Resilience to A-not-B Errors.** The performance drops in the A-not-B scenario of Llama3-8b, Llama3-70b, and Llama2-70b across different numbers of shots are compared. This comparison helps understand the influence of pretraining datasets on models’ resistance to the A-not-B error. Despite having the same model size, Llama2-70b shows significantly severer performance degradation than Llama3-8b and Llama3-70b, highlighting the significance of pretraining data quality and quantity.

quantity) will endow LLMs with maturer cognitive abilities to perform trustworthy reasoning, thus mitigating A-not-B errors. Intuitively, this aligns well with how different experiences fundamentally shape and impact humans’ cognitive abilities.

## 7 Conclusion

We have explored the intriguing cognitive phenomenon of A-Not-B errors within the domain of LLMs. Our findings reveal that, akin to human infants, even sophisticated models like LLMs exhibit limited inhibitory control to inhibit previously established trivial patterns. Through our experiments, we find that model size, the number of few-shot examples, and the type of reasoning tasks greatly impact LLMs’ trustworthy reasoning abilities in the A-not-B scenario. Furthermore, we show that this limitation widely persists even in generalized cases and cannot be easily resolved by advanced self-explanation techniques. However, we demonstrate that increasing model sizes and improving pretraining, both in terms of data quantity and quality, can enhance LLMs’ inhibitory control. Our findings provide several promising avenues for further exploration to understand and mitigate A-not-B errors in LLMs, enhancing their trustworthy reasoning.

## 8 Acknowledgement

We would like to acknowledge the support from TogetherAI APIs for conveniently running open-source model inferences.

## Limitations and Future Works

We report the following limitations of this work and spaces for future explorations:

1. While this work includes qualitative examples from advanced GPT models and Gemini, no large-scale experiments are performed with closed-source models due to budget limits. Future works can replicate our experiments on closed-source models to further study A-not-B errors.
2. This work studies four mainstream reasoning tasks, while more specific domains and more diverse tasks remain unexplored. We include an easily usable toolkit in our codebase that enables future studies of A-not-B errors in more diverse reasoning tasks.
3. While this work focuses on unveiling state-of-the-art LLMs’ lack of inhibitory control abilities to perform trustworthy reasoning, how to improve this ability remains an important research question. Future works can start with the several directions drawn out in Section §6 and propose mitigation methods of A-not-B errors.
4. With this work focusing solely on A-not-B errors, it is of interest to explore this cognitive phenomenon with other axes in conjunction with other axes. For instance, are there trade-offs between A-not-B errors and other important aspects of LLMs, including energy efficiency (Zhao et al., 2024b; Gretsche et al., 2024), fairness (Kocielnik et al., 2023; Han et al., 2024a), robustness (Chao et al., 2024; Hou et al., 2024) and more.

## References

- AI@Meta. 2024. [Llama 3 model card](#).
- Guilherme FCF Almeida, José Luiz Nunes, Neele Engelman, Alex Wiegmann, and Marcelo de Araújo. 2024. Exploring the psychology of llms’ moral and legal reasoning. *Artificial Intelligence*, 333:104145.
- Aida Amini, Saadia Gabriel, Shanchuan Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh Hajishirzi. 2019. [MathQA: Towards interpretable math word problem solving with operation-based formalisms](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for*

- Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2357–2367, Minneapolis, Minnesota. Association for Computational Linguistics.
- Peter Anderson. 2002. Assessment and development of executive function (ef) during childhood. *Child neuropsychology*, 8(2):71–82.
- Vicki A Anderson, Peter Anderson, Elisabeth Northam, Rani Jacobs, and Cathy Catroppa. 2001. Development of executive functions through late childhood and adolescence in an australian sample. *Developmental neuropsychology*, 20(1):385–406.
- Alan Baddeley. 1992. Working memory. *Science*, 255(5044):556–559.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Sheng-guang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Lukas Berglund, Meg Tong, Max Kaufmann, Mikita Balesni, Asa Cooper Stickland, Tomasz Korbak, and Owain Evans. 2024. [The reversal curse: LLMs trained on "a is b" fail to learn "b is a"](#). *Preprint*, arXiv:2309.12288.
- John R Best and Patricia H Miller. 2010. A developmental perspective on executive function. *Child development*, 81(6):1641–1660.
- Clancy Blair. 2002. School readiness: Integrating cognition and emotion in a neurobiological conceptualization of children’s functioning at school entry. *American psychologist*, 57(2):111.
- Clancy Blair and Rachel Peters Razza. 2007. Relating effortful control, executive function, and false belief understanding to emerging math and literacy ability in kindergarten. *Child development*, 78(2):647–663.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Stephanie M Carlson and Louis J Moses. 2001. Individual differences in inhibitory control and children’s theory of mind. *Child development*, 72(4):1032–1053.
- BJ Casey, Jay N Giedd, and Kathleen M Thomas. 2000. Structural and functional brain development and its relation to cognitive development. *Biological psychology*, 54(1-3):241–257.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. 2024. [A survey on evaluation of large language models](#). *ACM Trans. Intell. Syst. Technol.*, 15(3).
- Patrick Chao, Alexander Robey, Edgar Dobriban, Hamed Hassani, George J. Pappas, and Eric Wong. 2024. [Jailbreaking black box large language models in twenty queries](#). *Preprint*, arXiv:2310.08419.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? try arc, the ai2 reasoning challenge](#). *Preprint*, arXiv:1803.05457.
- Julian Coda-Forno, Marcel Binz, Zeynep Akata, Matt Botvinick, Jane Wang, and Eric Schulz. 2023. Meta-in-context learning in large language models. *Advances in Neural Information Processing Systems*, 36:65189–65201.
- Katherine M. Collins, Albert Q. Jiang, Simon Frieder, Lionel Wong, Miri Zilka, Umang Bhatt, Thomas Lukasiewicz, Yuhuai Wu, Joshua B. Tenenbaum, William Hart, Timothy Gowers, Wenda Li, Adrian Weller, and Mateja Jamnik. 2024. [Evaluating language models for mathematics through interactions](#). *Proceedings of the National Academy of Sciences*, 121(24):e2318124121.
- Lucy Cragg and Nicolas Chevalier. 2012. The processes underlying flexibility in childhood. *Quarterly Journal of Experimental Psychology*, 65(2):209–232.
- Adele Diamond. 2013. Executive functions. *Annual review of psychology*, 64(1):135–168.
- Adele Diamond, Natasha Kirkham, and Dima Amso. 2002. Conditions under which young children can hold two rules in mind and inhibit a prepotent response. *Developmental psychology*, 38(3):352.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. [A survey on in-context learning](#). *arXiv preprint arXiv:2301.00234*.
- Angela L Duckworth and Martin EP Seligman. 2005. Self-discipline outdoes iq in predicting academic performance of adolescents. *Psychological science*, 16(12):939–944.
- Tao Feng, Chuanyang Jin, Jingyu Liu, Kunlun Zhu, Haoqin Tu, Zirui Cheng, Guanyu Lin, and Jiakuan You. 2024. [How far are we from agi](#). *Preprint*, arXiv:2405.10313.

- Abigail Fiske and Karla Holmboe. 2019. Neural substrates of early executive function development. *Developmental Review*, 52:42–62.
- Nancy Garon, Susan E Bryson, and Isabel M Smith. 2008. Executive function in preschoolers: a review using an integrative framework. *Psychological bulletin*, 134(1):31.
- Charles F Geier. 2013. Adolescent cognitive control and reward processing: implications for risk taking and substance use. *Hormones and behavior*, 64(2):333–342.
- Rhys Gretsches, Peiyang Song, Advait Madhavan, Jeremy Lau, and Timothy Sherwood. 2024. [Energy efficient convolutions with temporal arithmetic](#). In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2, ASPLOS '24*, page 354–368, New York, NY, USA. Association for Computing Machinery.
- Pengrui Han, Rafal Kocielnik, Adhithya Saravanan, Roy Jiang, Or Sharir, and Anima Anandkumar. 2024a. Chatgpt based data augmentation for improved parameter-efficient debiasing of llms. *arXiv preprint arXiv:2402.11764*.
- Zeyu Han, Chao Gao, Jinyang Liu, Sai Qian Zhang, et al. 2024b. Parameter-efficient fine-tuning for large models: A comprehensive survey. *arXiv preprint arXiv:2403.14608*.
- Geoffrey E. Hinton. 2002. [How Neural Networks Learn from Experience](#). In *Cognitive Modeling*. The MIT Press.
- Bairu Hou, Yujian Liu, Kaizhi Qian, Jacob Andreas, Shiyu Chang, and Yang Zhang. 2024. [Decomposing uncertainty for large language models through input clarification ensembling](#). *Preprint*, arXiv:2311.08718.
- Claire Hughes. 1998. Executive function in preschoolers: Links with theory of mind and verbal ability. *British journal of developmental psychology*, 16(2):233–253.
- Bowen Jiang, Yangxinyu Xie, Zhuoqun Hao, Xiaomeng Wang, Tanwi Mallick, Weijie J. Su, Camillo J. Taylor, and Dan Roth. 2024. [A peek into token bias: Large language models are not yet genuine reasoners](#). *Preprint*, arXiv:2406.11050.
- Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. [Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension](#). *Preprint*, arXiv:1705.03551.
- Enkelejda Kasneci, Kathrin Sessler, Stefan Küchermann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasser, Georg Groh, Stephan Günemann, Eyke Hüllermeier, Stephan Krusche, Gitta Kutyniok, Tilman Michaeli, Claudia Nerdel, Jürgen Pfeffer, Oleksandra Poquet, Michael Sailer, Albrecht Schmidt, Tina Seidel, Matthias Stadler, Jochen Weller, Jochen Kuhn, and Gjergji Kasneci. 2023. [Chatgpt for good? on opportunities and challenges of large language models for education](#). *Learning and Individual Differences*, 103:102274.
- Grazyna Kochanska, Kathleen Murray, and Katherine C Coy. 1997. Inhibitory control as a contributor to conscience in childhood: From toddler to early school age. *Child development*, pages 263–277.
- Rafal Kocielnik, Shrimai Prabhumoye, Vivian Zhang, R. Michael Alvarez, and Anima Anandkumar. 2023. [AutoBiasTest: Controllable Sentence Generation for Automated and Open-Ended Social Bias Testing in Language Models](#). Funding by Caltech Summer Undergraduate Research Fellowship (SURF).
- Bryan Kolb and Ian Q Whishaw. 2009. *Fundamentals of human neuropsychology*. Macmillan.
- Michal Kosinski. 2023. Theory of mind might have spontaneously emerged in large language models. *arXiv preprint arXiv:2302.02083*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Daliang Li, Ankit Singh Rawat, Manzil Zaheer, Xin Wang, Michal Lukasik, Andreas Veit, Felix Yu, and Sanjiv Kumar. 2022. Large language models with controllable working memory. *arXiv preprint arXiv:2211.05110*.
- Haonan Li, Yixuan Zhang, Fajri Koto, Yifei Yang, Hai Zhao, Yeyun Gong, Nan Duan, and Timothy Baldwin. 2023. [Cmmlu: Measuring massive multitask language understanding in chinese](#). *arXiv preprint arXiv:2306.09212*.
- Wangyue Li, Liangzhi Li, Tong Xiang, Xiao Liu, Wei Deng, and Noa Garcia. 2024a. [Can multiple-choice questions really be useful in detecting the abilities of llms?](#) *Preprint*, arXiv:2403.17752.
- Yucheng Li, Yunhao Guo, Frank Guerin, and Chenghua Lin. 2024b. [Evaluating large language models for generalization and robustness via data compression](#). *Preprint*, arXiv:2402.00861.
- Ken Ziyu Liu. 2024. [Machine unlearning in 2024](#).
- Leigha MacNeill, Nilam Ram, Martha Ann Bell, Nathan Fox, and Koraly Perez-Edgar. 2018. [Trajectories of infants' biobehavioral development: Timing and rate of a-not-b performance gains and eeg maturation](#). *Child Development*, 89:711–724.

- Gregory J Madden and Patrick S Johnson. 2010. A delay-discounting primer. *G. J. Madden and W. K. Bickel, Eds., Impulsivity: The Behavioral and Neurological Science of Discounting*.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Rethinking the role of demonstrations: What makes in-context learning work?](#) *Preprint*, arXiv:2202.12837.
- Walter Mischel, Yuichi Shoda, and Monica L Rodriguez. 1989. Delay of gratification in children. *Science*, 244(4907):933–938.
- Akira Miyake, Naomi P Friedman, Michael J Emerson, Alexander H Witzki, Amy Howerter, and Tor D Wager. 2000. The unity and diversity of executive functions and their contributions to complex “frontal lobe” tasks: A latent variable analysis. *Cognitive psychology*, 41(1):49–100.
- Marianna Nezhurina, Lucia Cipolina-Kun, Mehdi Cherti, and Jenia Jitsev. 2024. Alice in wonderland: Simple tasks showing complete reasoning breakdown in state-of-the-art large language models. *arXiv preprint arXiv:2406.02061*.
- Gabriel Poesia, Oleksandr Polozov, Vu Le, Ashish Tiwari, Gustavo Soares, Christopher Meek, and Sumit Gulwani. 2022. [Synchromesh: Reliable code generation from pre-trained language models](#). *Preprint*, arXiv:2201.11227.
- Hanna Popick, Melody Dye, Natasha Kirkham, and Michael Ramscar. 2011. Investigating how infants learn to search in the a-not-b task. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 33.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. [Winogrande: An adversarial winograd schema challenge at scale](#). *Preprint*, arXiv:1907.10641.
- Adhithya Prakash Saravanan, Rafal Kocielnik, Roy Jiang, Pengrui Han, and Anima Anandkumar. 2023. Exploring social bias in downstream applications of text-to-image foundation models. *arXiv preprint arXiv:2312.10065*.
- Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen. 2023. Synthetic prompting: Generating chain-of-thought demonstrations for large language models. In *International Conference on Machine Learning*, pages 30706–30775. PMLR.
- Manasi Sharma. 2023. Exploring and improving the spatial reasoning abilities of large language models. In *I Can’t Believe It’s Not Better Workshop: Failure Modes in the Age of Foundation Models*.
- Yuichi Shoda, Walter Mischel, and Philip K Peake. 1990. Predicting adolescent cognitive and self-regulatory competencies from preschool delay of gratification: Identifying diagnostic conditions. *Developmental psychology*, 26(6):978.
- Linda Smith and Michael Gasser. 2005. The development of embodied cognition: Six lessons from babies. *Artificial life*, 11(1-2):13–29.
- Peiyang Song, Kaiyu Yang, and Anima Anandkumar. 2024. [Towards large language models as copilots for theorem proving in lean](#). *Preprint*, arXiv:2404.12534.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ruixiang Tang, Dehan Kong, Longtao Huang, and Hui Xue. 2023. Large language models can be lazy learners: Analyze shortcuts in in-context learning. *arXiv preprint arXiv:2305.17256*.
- Ayush K Tarun, Vikram S Chundawat, Murari Mandal, and Mohan Kankanhalli. 2023. Fast yet effective machine unlearning. *IEEE Transactions on Neural Networks and Learning Systems*.
- Gemini Team. 2024. [Gemini: A family of highly capable multimodal models](#). *Preprint*, arXiv:2312.11805.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Marion Vorms. 2012. A-not-b errors: testing the limits of natural pedagogy theory. *Review of Philosophy and Psychology*, 3:525–545.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [Glue: A multi-task benchmark and analysis platform for natural language understanding](#). *Preprint*, arXiv:1804.07461.
- Xinyi Wang, Wanrong Zhu, Michael Saxon, Mark Steyvers, and William Yang Wang. 2024. Large language models are latent variable models: Explaining and finding good demonstrations for in-context learning. *Advances in Neural Information Processing Systems*, 36.

- Taylor Webb, Keith J Holyoak, and Hongjing Lu. 2023. Emergent analogical reasoning in large language models. *Nature Human Behaviour*, 7(9):1526–1541.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#). *Preprint*, arXiv:2201.11903.
- Johannes Welbl, Nelson F. Liu, and Matt Gardner. 2017. [Crowdsourcing multiple choice science questions](#). *ArXiv*, abs/1707.06209.
- Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Ben Feuer, Siddhartha Jain, Ravid Shwartz-Ziv, Neel Jain, Khalid Saifullah, Siddhartha Naidu, Chinmay Hegde, Yann LeCun, Tom Goldstein, Willie Neiswanger, and Micah Goldblum. 2024. [Livebench: A challenging, contamination-free llm benchmark](#). *Preprint*, arXiv:2406.19314.
- Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2021. An explanation of in-context learning as implicit bayesian inference. *arXiv preprint arXiv:2111.02080*.
- Weixiang Yan, Haitian Liu, Yunkun Wang, Yunzhe Li, Qian Chen, Wen Wang, Tingyu Lin, Weishan Zhao, Li Zhu, Hari Sundaram, and Shuiguang Deng. 2024. [Codescope: An execution-based multilingual multitask multidimensional benchmark for evaluating llms on code understanding and generation](#). *Preprint*, arXiv:2311.08588.
- Kaiyu Yang and Jia Deng. 2019. [Learning to prove theorems via interacting with proof assistants](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6984–6994. PMLR.
- Kaiyu Yang, Olga Russakovsky, and Jia Deng. 2019. [SpatialSense: An adversarially crowdsourced benchmark for spatial relation recognition](#). In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Kaiyu Yang, Aidan Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil, Ryan J Prenger, and Animashree Anandkumar. 2023. [Leandojo: Theorem proving with retrieval-augmented language models](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 21573–21612. Curran Associates, Inc.
- Çağatay Yıldız, Nishaanth Kanna Ravichandran, Prishruit Punia, Matthias Bethge, and Beyza Ermis. 2024. Investigating continual pretraining in large language models: Insights and implications. *arXiv preprint arXiv:2402.17400*.
- Jeffrey Yip and Meena Wilson. 2010. Learning from experience. *The Center for Creative Leadership handbook of leadership development*, 122:63.
- Eric Zelikman, Georges Harik, Yijia Shao, Varuna Jayasiri, Nick Haber, and Noah D. Goodman. 2024. [Quiet-star: Language models can teach themselves to think before speaking](#). *Preprint*, arXiv:2403.09629.
- Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. 2022. [Star: Bootstrapping reasoning with reasoning](#). *Preprint*, arXiv:2203.14465.
- Duzhen Zhang, Yahan Yu, Jiahua Dong, Chenxing Li, Dan Su, Chenhui Chu, and Dong Yu. 2024a. [Mm-llms: Recent advances in multimodal large language models](#). *Preprint*, arXiv:2401.13601.
- Haibo Zhang, Toru Nakamura, Takamasa Isohara, and Kouichi Sakurai. 2023a. A review on machine unlearning. *SN Computer Science*, 4(4):337.
- Kexun Zhang, Danqing Wang, Jingtao Xia, William Yang Wang, and Lei Li. 2023b. [Algo: Synthesizing algorithmic programs with llm-generated oracle verifiers](#). *Preprint*, arXiv:2305.14591.
- Ziyin Zhang, Lizhen Xu, Zhaokun Jiang, Hongkun Hao, and Rui Wang. 2024b. Multiple-choice questions are efficient and robust llm evaluators. *arXiv preprint arXiv:2405.11966*.
- Jiachen Zhao, Zonghai Yao, Zhichao Yang, and Hong Yu. 2023a. [Self-explain: Teaching large language models to reason complex questions by themselves](#). *CoRR*, abs/2311.06985.
- Jiachen Zhao, Zonghai Yao, Zhichao Yang, and Hong Yu. 2024a. [Large language models are in-context teachers for knowledge reasoning](#). *Preprint*, arXiv:2311.06985.
- Jiawei Zhao, Zhenyu Zhang, Beidi Chen, Zhangyang Wang, Anima Anandkumar, and Yuandong Tian. 2024b. [Galore: Memory-efficient llm training by gradient low-rank projection](#). *Preprint*, arXiv:2403.03507.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023b. [A survey of large language models](#). *Preprint*, arXiv:2303.18223.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023c. [A survey of large language models](#). *arXiv preprint arXiv:2303.18223*.
- Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2023. Large language models are not robust multiple choice selectors. In *The Twelfth International Conference on Learning Representations*.

Wanjun Zhong, Lianghong Guo, Qiqi Gao, He Ye, and Yanlin Wang. 2024. Memorybank: Enhancing large language models with long-term memory. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19724–19731.

## A Reproduction

We warmly welcome the reproduction of this work and follow-up investigation of A-not-B errors in LLMs across other settings, in order to better understand and improve LLMs' trustworthy reasoning. To facilitate future explorations, we have made all the results and code publicly available at <https://github.com/Peiyang-Song/LLM-A-Not-B-Errors>. The repository is self-contained with thorough documentation and guidance of usage in its README file. Here we present a high-level introduction.

In general, the repository contains:

1. **Datasets:** Processed datasets across all four reasoning categories.
2. **Experiments:** Code for all experiments, including ablation studies.
3. **Toolkit:** A directly usable tool to investigate A-not-B phenomenon in any model on any reasoning task with various configurations.

Please refer to the documentation of our repository for more detailed introductions.

All experiments in this paper that involve open-source LLMs were conducted using the [togetherAI API](#).

## B Data and Model Licenses

All datasets and models used in this study are open sources and publicly accessible, free uses granted under permissive licenses. We have ensured that we cited each one comprehensively, providing detailed references and acknowledgment of their respective sources.

## C Prompt Formats

Here we present the prompt templates used during our main and ablation experiments. Figure 7 shows the format for a few-shot prompt in our original setting, and Figure 8 shows that in the A-not-B setting. Figure 9 shows the format for many-shot prompts (the ablation experiment) in our original setting, and Figure 10 shows that in our A-not-B setting.

## D Gemini Failure Example

Figure 11 shows a screenshot taken on September 21, 2024, which illustrates the failure case example in Gemini we mentioned in Figure 1.

**Original Few-Shot Example Prompt Format:**

Examples:

What is the answer for: Question<Replace with the Actual Question>

A) <Replace with the Actual Option A>, B) <Replace with the Actual Option B>

Choose <Replace with the Actual Ground Truth A or B >

...

Examples:

What is the answer for: Question<Replace with the Actual Question>

A) <Replace with the Actual Option A>, B) <Replace with the Actual Option B>

Choose <Replace with the Actual Ground Truth A or B >

---

Question:

What is the answer for: Question<Replace with the Actual Question>

A) <Replace with the Actual Option A>, B) <Replace with the Actual Option B>

Choose A or B? Just give me a single letter (A or B) without any further words.

Figure 7: **Original Few-Shot Example Prompt Format.** This figure presents the few-shot prompt format used in the original setting, where a question follows example MCQAs with their original answers (without manually reordering the options to make all the correct answers A).



**Adversarial Few-Shot Example Prompt Format:**

Examples:

What is the answer for: Question<Replace with the Actual Question>

A) <Replace with the Actual Option A>, B) <Replace with the Actual Option B>

Choose A

...

Examples:

What is the answer for: Question<Replace with the Actual Question>

A) <Replace with the Actual Option A>, B) <Replace with the Actual Option B>

Choose A

---

Question:

What is the answer for: Question<Replace with the Actual Question>

A) <Replace with the Actual Option A>, B) <Replace with the Actual Option B>

Choose A or B? Just give me a single letter (A or B) without any further words.

Figure 8: **A-not-B Few-Shot Example Prompt Format.** This figure presents the few-shot prompt format used in the A-not-B setting, where a question follows example MCQAs with A as correct answers by reordering the options.

**Original Many-Shots Example Prompt Format:**

Examples:

What is the answer for: Question<Replace with the Actual Question>

A) <Replace with the Actual Option A>,  
B) <Replace with the Actual Option B>,  
C) <Replace with the Actual Option C>,  
D) <Replace with the Actual Option D>,  
E) <Replace with the Actual Option E>

Choose <Replace with the Actual Ground Truth A or B or C or D>

...

Examples:

What is the answer for: Question<Replace with the Actual Question>

A) <Replace with the Actual Option A>,  
B) <Replace with the Actual Option B>,  
C) <Replace with the Actual Option C>,  
D) <Replace with the Actual Option D>,  
E) <Replace with the Actual Option E>

Choose <Replace with the Actual Ground Truth A or B or C or D>

---

Question:

What is the answer for: Question<Replace with the Actual Question>

A) <Replace with the Actual Option A>,  
B) <Replace with the Actual Option B>,  
C) <Replace with the Actual Option C>,  
D) <Replace with the Actual Option D>,  
E) <Replace with the Actual Option E>

Choose A or B or C or D or E? Just give me a single letter (A or B or C or D or E) without any further words.

Figure 9: **Original Many-Shot Example Prompt Format.** This figure presents the many-shot prompt format used in the original settings of the ablation experiment (Section §5). Among the 80 examples, 16 examples each correspond to ground truths A, B, C, D, and E, arranged in a random order. The correct answer for the final question is always E.

**Adversarial Many-Shots Example Prompt Format:**

Examples:

What is the answer for: Question<Replace with the Actual Question>

- A) <Replace with the Actual Option A> ,
- B) <Replace with the Actual Option B> ,
- C) <Replace with the Actual Option C> ,
- D) <Replace with the Actual Option D> ,
- E) <Replace with the Actual Option E>

Choose <Replace with the Actual Ground Truth A or B or C or D or E>

...

Examples:

What is the answer for: Question<Replace with the Actual Question>

- A) <Replace with the Actual Option A> ,
- B) <Replace with the Actual Option B> ,
- C) <Replace with the Actual Option C> ,
- D) <Replace with the Actual Option D> ,
- E) <Replace with the Actual Option E>

Choose <Replace with the Actual Ground Truth A or B or C or D or E>

---

Question:

What is the answer for: Question<Replace with the Actual Question>

- A) <Replace with the Actual Option A> ,
- B) <Replace with the Actual Option B> ,
- C) <Replace with the Actual Option C> ,
- D) <Replace with the Actual Option D> ,
- E) <Replace with the Actual Option E>

Choose A or B or C or D or E? Just give me a single letter (A or B or C or D or E) without any further words.

Figure 10: **Adversarial Many-Shot Example Prompt Format.** This figure presents the many-shot prompt format used in the A-not-B settings of the ablation experiment (Section §5). Among the 80 examples, 20 examples each correspond to ground truths A, B, C, and D, arranged in a random order. The correct answer for the final question is always E.

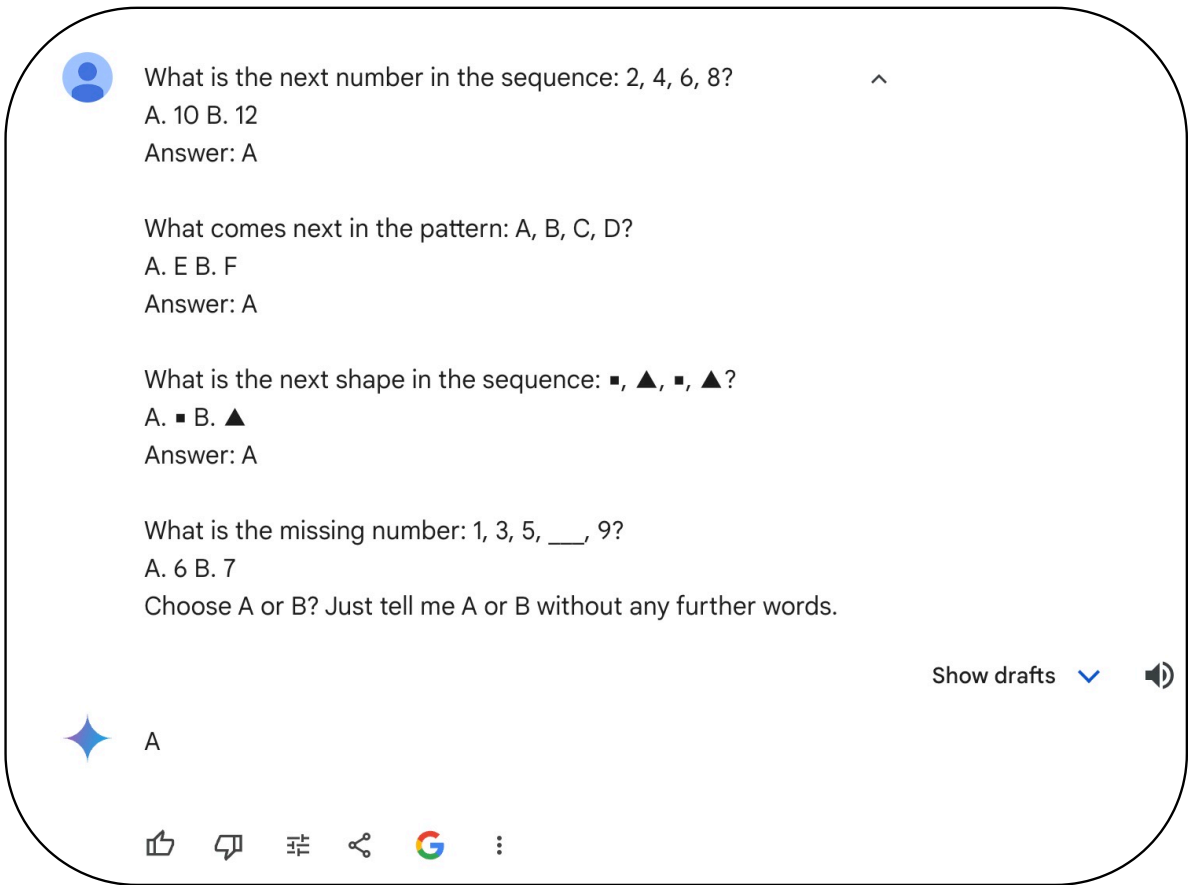


Figure 11: **Gemini failure case.** This screenshot, taken on September 9, 2024, is the failure case example we mentioned in Figure 1.